A habituation account of change detection in same/different judgments

Eddy J. Davelaar · Xing Tian ·
Christoph T. Weidemann · David E. Huber

Abstract We investigated the basis of change detection in a short-term priming task. In two experiments, participants were asked to indicate whether or not a target word was the same as a previously presented cue. Data from an experiment measuring magnetoencephalography failed to find different patterns for "same" and "different" responses, consistent with the claim that both arise from a common neural source, with response magnitude defining the difference between immediate novelty versus familiarity. In a behavioral experiment, we tested and confirmed the predictions of a habituation account of these judgments by comparing conditions in which the target, the cue, or neither was primed by its presentation in the previous trial. As predicted, cue-primed trials had faster response times, and target-primed trials had slower response times relative to the neither-primed baseline. These results were obtained irrespective of response repetition and stimulus–response contingencies. The behavioral and brain activity data support the view that detection of change drives performance in these tasks and that the underlying mechanism is neuronal habituation.

E. J. Davelaar (✉)
Department of Psychological Sciences, Birkbeck,
University of London,
Malet Street,
London WC1E 7HX, UK
e-mail: eddy.davelaar@gmail.com

X. Tian
New York University,
New York, NY, USA

C. T. Weidemann
Swansea University,
Swansea, Wales, UK

D. E. Huber
University of California,
San Diego, La Jolla, CA, USA

A question of fundamental interest in the study of perception, memory, and other cognitive processes is whether the processes of the brain are generally tuned toward the detection of familiar items and events, or whether the brain detects that which is new in the environment. The study of episodic memory has focused on the process of familiarity, demonstrating that behavioral responses are sensitive to relative degrees of memory strength (e.g., Hintzman & Curran, 1995). Furthermore, the timing and cortical location of these familiarity responses have been evidenced with neural recordings (Curran, 2000; Henson, Cansino, Herron, Robb, & Rugg, 2003). For long-term memory, it is sensible to focus on the detection of old rather than new items. More specifically, there is evidence that familiarity captures long-term statistical regularities that are useful for a wide variety of tasks and judgments (e.g., Hertwig, Herzog, Schooler, & Reimer, 2008). Whereas it is nearly uniformly accepted that long-term memory is based on familiarity detection rather than novelty detection, this distinction is less well studied for shorter time scales, such as in perceptual tasks or in short-term working memory tasks. To investigate the roles of familiarity detection and novelty detection (i.e., change detection) in perceptual tasks, we examined situations that

involved the detection of change with respect to stimuli that were presented close to each other in time and space.

Closely related paradigms have demonstrated that even with full attention and awareness, participants may be unable to detect change. For example, in studies of "change blindness," participants are shown alternations of photos or abstract objects (e.g., Rensink, O'Regan, & Clark, 1997; Simons, 1996), and they experience difficulty identifying the change across alternations. In other work, changes across scenes are frequently missed, irrespective of whether these changes are abrupt or gradual (Rensink et al., 1997; Simons, Franconeri, & Reimer, 2000). Although these change detection paradigms suffice for investigating the limits of the cognitive system, they are not ideal for studying the mechanisms by which change detection occurs. This is because these detection paradigms involve visual displays with a high density of objects. Research has shown that an increase in perceptual load narrows the attentional focus (e.g., Williams, 1989), which draws attention away from the feature dimension along which a change might occur, thus producing inattentional blindness (Most, Scholl, Clifford, & Simons, 2005; Simons, 2000). Therefore, in tasks with a high density of objects, change blindness is attributable to load-induced narrowing of spatial attention, because attention is more likely to be allocated to an area where no change occurs. This sensitivity to the focus of attention renders these paradigms unsuitable for investigating the *mechanisms* underlying change detection.

To isolate the mechanisms of change detection, we used a paradigm with low attentional load. In this paradigm, a first stimulus, S1, is followed by a second stimulus, S2, to which a response is required indicating how it compares to S1. Such tasks featured in a debate on the architecture underlying same/different judgments (see Farell, 1985, for a review). This debate contrasted single- versus dual-process models (a dual-process model involves the separate processes of sameness detection and change detection), with the evidence favoring single-process models. Single-process models assume that a match value is obtained by contrasting S2 with S1, and that this value, which is along a match/nonmatch dimension, is compared to a criterion in order to determine the appropriate response (Ratcliff, 1985; Van Zandt, Colonius, & Proctor, 2000). However, these models do not specify whether the criterion is placed on the magnitude of a match signal (i.e., making a decision based on familiarity) or on the magnitude of a mismatch signal (i.e., making a decision based on novelty). In this article, we advocate a single-process model of change detection, and we test the claim that the nature of this process is to calculate the degree of mismatch (novelty).

Building on a theory of short-term priming (Huber & O'Reilly, 2003; see also Huber, 2008b), we favor the view

that the perceptual system primarily employs mechanisms for detecting differences between stimuli rather than similarities. This theory explained a series of experiments (Huber, Shiffrin, Lyle & Quach, 2002a; Huber, Shiffrin, Lyle, & Ruys, 2001; Huber, Shiffrin, Quach, & Lyle, 2002) that used a two-alternative forced choice (2AFC) paradigm in which a briefly flashed (and masked) target stimulus needed to be identified. On each trial, prime stimuli presented prior to the target matched either the target or the incorrect choice alternative in the 2AFC (i.e., the foil). When comparing target-primed trials to foil-primed trials, the results were equal and opposite, suggesting that priming primarily induced a preference effect. Furthermore, with short prime durations (50 ms), there was a preference to choose the primed stimulus, but with longer prime durations (2,000 ms), there was a preference against choosing the primed stimulus (Huber, 2008b). In other words, if the target choice was primed for a long duration, participants tended to choose the foil instead of the target, reducing accuracy. Thus, following a sufficiently long exposure to a preceding stimulus, performance was better when the target represented a change.

As reviewed in Huber (2008b), this pattern of results is readily explained by proposing that the perceptual system habituates to any stimulus attended (e.g., a prime) and that the amount of habituation becomes greater with longer prime durations (Huber, 2008b; Huber & O'Reilly, 2003). According to this theory, when a prime is very salient, the primed choice word (target or foil) is still in a habituated state when the choice words appear, slowing down perceptual processing for that choice. Assuming that the decision is guided by a sense of perceptual fluency, the participant then chooses the alternative that is less habituated, or conversely, more novel. This theory was first instantiated as a formal model called "Responding Optimally with Unknown Sources of Evidence" (ROUSE; Huber et al., 2001). We focus on the neurocomputational version of ROUSE (nROUSE; Huber & O'Reilly, 2003), which is a neural model for novelty detection. This neural model produces habituation through "synaptic depression," as we explain next.

In most neurocomputational models of cognitive processes, the activation of internal representations increases with increasing stimulus duration (but see Gotts & Plaut, 2002, and Huber & O'Reilly, 2003). However, these models are at odds with the common observation that neuronal firing rates change nonmonotonically with changes in the duration of relevant stimuli (Tsodyks & Markram, 1997). When excitatory pyramidal cells receive constant input, their firing rates typically increase sharply and reach a peak within 150 ms, after which they decrease toward a lower level that is above the baseline firing rate. This neural fatigue has been termed "synaptic depression," because a

variety of cellular mechanisms serve to produce a temporary reduction in the efficacy of the synapses, with a recovery period needed before synaptic resources are replenished. One such mechanism is the amount of neurotransmitter available within a spiking neuron. If a neuron has been inactive, it will have sufficient neurotransmitter available to effectively transmit its activation to other neurons. However, if that neuron has been active recently, its neurotransmitter supplies may be temporarily depleted, in which case the firing of the neuron is ineffective in terms of transmitting activation to other neurons. To illustrate the patterns of synaptic depression, consider Fig. 1, in which the output activation of a simulated neuron is shown as a function of stimulus duration (see Huber & O'Reilly, 2003, for details). As shown in the top panel, without synaptic depression, output activation increases monotonically with stimulus duration, and this pattern of response is identical when the cell is driven by a second input after a short delay. However, as shown in the bottom panel, which includes synaptic depression, the output activation is nonmonotonically related to stimulus duration. Furthermore, repeating the stimulus after a short delay leads to a decreased neuronal response as compared to the first presentation (e. g., compare the heights of the two peaks).

Using a model of habituation through synaptic depression, Huber and O'Reilly (2003) showed that for the short-term priming paradigm, the time needed to achieve peak activation for each choice (i.e., perceptual fluency) captured both
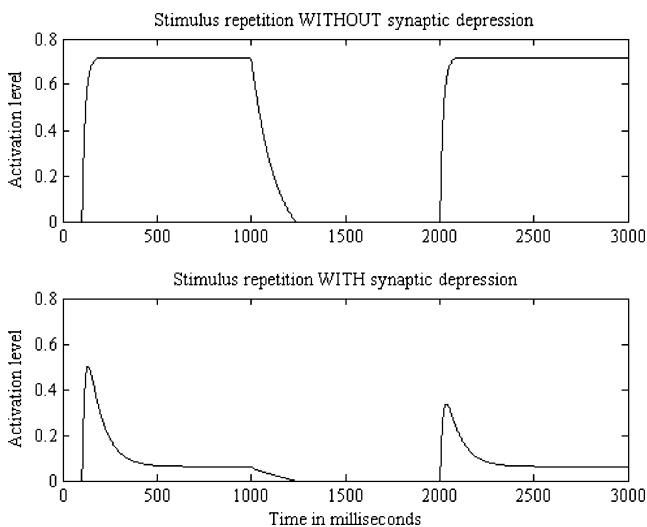

Fig. 1 Illustration of the effect of synaptic depression on the output activation of units in a neural network and of the relation with repetition-induced suppression. In both simulations, the stimulus is presented in the interval from 100 to 1,000 ms, and then re-presented from time = 2,000 ms. The top panel shows the typical monotonic activation profile, without synaptic depression or repetition-induced suppression. The bottom panel shows the nonmonotonic activation profile due to synaptic depression and the presence of repetition-induced suppression

positive priming following short prime durations and negative priming following long prime durations. Similar to the interactive-activation model of McClelland and Rumelhart (1981), this was done with a three-layer (visual, orthographic, and lexical–semantic) network, although the representations at each layer were localist, and synaptic depression was used for all simulated neurons. Thus, habituation was imposed at all levels of processing. For longer prime durations, the peak of the primed response was lower and later. Beyond modeling behavioral results, the same model was used to make detailed predictions regarding the time course of neural responses, which were confirmed in a combined electroencephalogram (EEG)/magnetoencephalogram (MEG) study (Huber, Tian, Curran, O'Reilly, & Woroch, 2008). Using parameters previously constrained by behavioral results, this model produced an accurate account of the ERP waveforms, which provided additional support to the claim that representations beyond the low-level visual stage are subject to habituation. Similar to this approach, in the present study we used MEG to address the earliest signal that reflects novelty detection.

If neural habituation is the mechanism behind behavioral negative priming with immediate repetitions, this may also explain the ubiquitous finding that neural responses experience a "repetition suppression," which is used as an important marker in neuroimaging studies (Grill-Spector & Malach, 2001). Indeed, the mechanism of synaptic depression has been demonstrated to capture repetition-induced suppression in neural responses (Miller & Desimone, 1994). The neurophysiological mechanisms underlying repetition-induced suppression are under intense investigation (Fuhrmann, Cowan, Segev, Tsodyks, & Stricker, 2004), as is the precise mapping of habituation at the neural level, as compared to the repetition suppression seen in neuroimaging studies (Chouinard, Morrissey, Kohler, & Goodale, 2008; Guo, Lawson, & Jiang, 2007; Henson, Mouchlianitis, Matthews, & Kouider, 2008; Horner & Henson, 2008; James & Gauthier, 2006; Sawamura, Orban, & Vogels, 2006; for a review, see Grill-Spector, Henson, & Martin, 2006). The relation between neural repetition suppression and behavioral priming is still poorly understood, and dissociations have been observed between these phenomena (Dobbins, Schnyer, Verfaellie, & Schacter, 2004; Ganel et al., 2006; Sayres & Grill-Spector, 2006).

It is important to note that a close relation exists between neural habituation and information processing models, in which a prediction is made about what to expect on the basis of the recent past. For instance, Chance, Nelson, and Abbott (1998) implemented synaptic depression in a model of V1 cells, demonstrating increased sensitivity to transient changes. In other words, by resetting baseline activation by adapting to the current stimulus, perception is sensitive to change relative to the current stimulus. Using a hidden

Markov model (HMM), Huber (2008a) formalized this relation between perceptual predictions and synaptic depression: A cascaded HMM was found to produce the same dynamics as synaptic depression. The HMM model produced less of a response to the repetition of an input because recent events provided a strong prediction for the perceptual evidence of the repeated item. Thus, repeated items were discounted such that any new perceptual evidence that might suggest a second occurrence of an item was instead attributed to this prediction based on the recent past. From this perspective, the reason for repetition suppression is to provide perceptual prediction, whereas synaptic depression is the mechanism by which repetition suppression is neurally implemented.

In applying this priming model to the present study, we do not claim that any particular level of processing is more important in terms of neural habituation. Because the present study contrasted repetitions (which are the same at all levels) with unrelated words (which are different at all levels), we cannot specify whether the observed effects were due to orthographic/phonological or lexical–semantic habituation (or some other higher-level representation). Furthermore, we do not claim that synaptic depression is the precise mechanism underlying these effects, considering that other neural mechanisms can produce similar habituation dynamics. Instead, our claim is that the dynamics of habituation, as they exist at some or all levels of processing, enable change detection between what just occurred and what appears subsequently.

By testing a habituation account both with behavioral priming and with MEG, we approached this issue with a single comprehensive model that is based on neural dynamics. There are two components to our account: (1) that same/different judgments for an immediate comparison between cue and target words are based on a novelty signal rather than a familiarity signal or different signals for *same* and *different* trials, and (2) that habituation is the mechanism that underlies this novelty signal. We tested the first component in Experiment 1 by comparing the topographic similarity and response magnitudes of neural responses for *same* versus *different* trials. This was done with words that never repeated across trials, so as to eliminate any priming/habituation effects from previous trials. We tested the second component in Experiment 2 by examining repetition priming between one trial and the next; if habituation underlies novelty detection and if habituation can linger from recent presentations of a word, this account predicts that last-trial priming of a cue word (i.e., even greater habituation to the cue) should facilitate *different* trials, whereas last-trial priming of a target word should produce the opposite effect (i.e., target habituation, as if the target was just seen as the cue, rather than during the last trial). Critically, this was done with a small number of words that

were reused many times across trials, so as to equate any effects of long-term repetition priming. In this manner, we isolated the short-term habituation effect from the last trial. Furthermore, we contrasted this account with an account of repetition-induced suppression that is based on associative learning between a repeated stimulus and its behavioral response (e.g., Logan, 1990). In Experiment 2, we differentiated between these accounts using a paradigm in which the same repeated target word from one trial to the next might involve different behavioral responses on different trials (e.g., a response of "same" on trial *n*–1, but a response of "different" on trial *n*).

Experiment 1

We propose that change detection when a cue word is immediately followed by a target word is based on habituation and that the detection of sameness between cue and target is the absence of change detection. In other words, the brain does not employ a separate mechanism for detection of immediate repetitions for visually presented words. On this account, habituation to the cue word serves to produce a large response when the target is different from rather than repeats the cue (because the neural representation for a repeated target is already habituated). We do not claim that sameness is the absence of any signal (cf. Johnson, Spencer, Luck & Schöner, 2009). Instead, we suggest that a sameness signal is *created* by a competitive comparison between a single source of novelty and a referent (or bias). Such a view is compatible with the literature showing that some populations of cells respond more strongly for repeated stimuli, while others respond less strongly (Miller & Desimone, 1994).

In this same/different judgment paradigm, the obvious alternative account of performance is that the mechanism that triggers change detection may be the absence of familiarity rather than the presence of novelty. For this alternative to explain performance, it needs to be assumed that presentation of an item temporarily boosts the familiarity response for that item. This boost serves as the signal behind accurate performance, because a target that is the same as the cue is potentiated by its presentation as the cue, which ensures that the familiarity response to that target is above the familiarity detection criterion and thus no change is detected. These accounts are not mutually exclusive, and it is possible that both potentiation of familiarity and depression of novelty play a hand in change detection.

To differentiate between these accounts, Experiment 1 used MEG to determine whether the neural response pattern differs between *same* and *different* trials. MEG is a noninvasive measure of neural activity that has the same millisecond temporal resolution as EEG, but better spatial resolution. This

increased spatial resolution makes MEG ideally suited to measure the pattern of neural activity in response to the target word. Additionally, we investigated whether the magnitude of the neural signal differs between same and different targets. A relative increase for *same* trials would suggest that repetitions produce potentiation, whereas a decrease would support the habituation account. A number of MEG studies have shown that habituation occurs in the time window of 100–350 ms after onset, with nonlexical and lexical habituation effects appearing in the early versus late parts of this window, respectively (see Monahan, Fiorentino, & Poeppel, 2008, for a review). Although these studies evidenced neural habituation with word stimuli, their focus was mainly on repetition priming effects instead of on determining mechanisms of change detection. To investigate this mechanism in Experiment 1, we used a multivariate sensor analysis to determine whether different brain networks are used for change detection versus detection of sameness. For each individual, we compared the patterns across all 157 sensors for *same* and *different* trials to ascertain whether patterns differed for the two types of trials, which would suggest the involvement of different processes. However, if the pattern remained the same but varied in response magnitude, this would suggest a single underlying process.

Because we were interested in the mechanism underlying change detection rather than the task-related decision process, we focused our analyses on the earliest MEG component that shows a difference for a target word that differs from a word presented immediately prior to the target. Previous work with repetition priming in a perceptual identification task revealed that the M170 is affected by immediate repetitions of visually presented words. Furthermore, this component is modulated by prime duration (Huber et al., 2008), as would be expected if this component were affected by habituation. Therefore, we focused on the M170 as the critical signal in the present experiment. The presentation sequence used in the present experiment follows closely from the MEG experiment of Huber et al. (2008). However, the task in that experiment was identification of a briefly flashed and then masked target word presented immediately after a prime word. In contrast, we now call the first word a "cue" because the task in the present experiment was a speeded same/different judgment of a target word that remained visible until a response was given. Similar to the prime durations used in the Huber et al. experiment, the present experiment used cue durations of 150 versus 2,000 ms to induce different levels of habituation.

Method

Participants A group of 17 paid participants were recruited from the University of Maryland for Experiment 1. All participants were right-handed native English speakers with no history of neurological disorder.

Materials We constructed a pool of 1,500 five-letter words. These words had a minimum written language frequency of 2 per million (Kučera & Francis, 1967). All words were displayed in uppercase Times New Roman font in yellow font color against a black background, and subtended less than 3° of visual angle. Stimulus materials were projected on a screen inside the MEG chamber with presentation times synchronized to the vertical refresh. The refresh rate was 75 Hz, providing display increments of 13.3 ms.

Design Experiment 1 conformed to a 2 × 2 within-subjects design, crossing the variables cue type (same as vs. different from target) and cue duration (short vs. long: 150 vs. 2,000 ms).

Behavioral procedure Experimental sessions lasted for about half an hour and included 48 blocks of eight experimental trials. Participants were instructed to refrain from blinking or moving during each block. Between each block of trials, participants received feedback about the total number of correct answers in the previous block. There were two repetitions of each of the four unique combinations of experimental conditions (short–same, short–different, long–same, long–different), which were randomly intermixed in each block. Word selection occurred randomly without replacement, such that each word was only seen on one trial. This was necessary to eliminate any priming effects from previous trials. Participants were fully informed as to the nature of the conditions and were warned that there were equal numbers of *same* and *different* trials. Two blocks of eight practice trials were presented prior to the 48 blocks of experimental trials.

At the beginning of each trial, a fixation line consisting of five short, dashed lines appeared for 1,000 ms in the center of the screen, followed by two instances of a cue word, one presented immediately above and one immediately below the fixation line. After 2,000 ms (long) or 150 ms (short), the cue words were followed by the target, presented in the center of the screen. Participants were asked to judge whether the target word was the same as or different than the cue word by pressing the corresponding button. They were encouraged to respond as quickly and accurately as possible. The target word remained on the screen until participants responded.

MEG procedure Magnetic signals were recorded using a 160-channel (157 data channels and 3 reference channels), whole-head axial gradiometer system (KIT, Kanazawa, Japan). In order to check head position within the MEG scanner, five electromagnetic coils were attached to the head of a participant 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 3-

D digitizer software (Source Signal Imaging, Inc.) and digitizing hardware (Polhemus, Inc.). The coils were localized with respect to the MEG sensors, both at the beginning and at the end of the experiment.

During the experiment, participants were supine inside a magnetically shielded room. Before the visual word experiment, they listened to 200 combined repetitions of 250-Hz and 1-kHz versions of a 50-ms sinusoidal tone (ISI randomized between 750 and 1,550 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 a review, see Roberts, Ferrari, Stufflebeam, & Poeppel, 2000). A dipole-like pattern (i.e., a source–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 scanner.

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. Raw data were noise-reduced using CALM (Adachi, Shimogawara, Higuchi, Haruta, & Ochiai, 2001). A 1,000-ms time period that was time-locked to the long-duration cue word was extracted and averaged over the 192 long-cue-duration trials. These were used as standard responses of the M170 response to a singly presented word (see the Analysis of MEG Data section below). There was no baseline correction (i.e., subtraction of a prestimulus interval) of the data, because statistical comparisons were only between *same* and *different* trials rather than between short- versus long-duration cues. In other words, the prestimulus interval was functionally identical for the conditions that were compared, so baseline correction would only serve to reduce statistical power by introducing an additional source of variance. A second 1,000-ms time period that was time-locked to the target onset was extracted and averaged over the 96 trials for each of four conditions. Trials with amplitudes >3pT (~5%) were considered artifacts and were discarded. The averages were low-pass filtered with a cutoff frequency of 20 Hz.

Analysis of MEG data Unlike EEG, analysis of MEG is typically done separately for each individual, because MEG provides a much more spatially localized response, and different individuals position their heads differently in relation to the sensors with MEG. Therefore, the same sensor may produce very different, or even opposite, responses (i.e., opposite sides of a local electromagnetic dipole) for different individuals. To address this issue, Tian and Huber (2008) developed multivariate measurement techniques that use the

pattern of responses across all sensors by normalizing these patterns in relation to a "standard response," which is the pattern observed in some other condition for that individual. We briefly describe these techniques here and point the interested reader to the discussion by Tian, Poeppel, and Huber (2011) that accompanies the release of the TopoTool-box suite of MATLAB routines that implement these techniques. These files, as well as a tutorial, can be downloaded from Dr. Tian's website (https://files.nyu.edu/xt235/public/). In these techniques, rather than examining separate sink (magnetic influx)/source (magnetic outflux) patterns for the M170 of each individual, all 157 sensors are used to ascertain whether a difference in the M170 between conditions is due to a different pattern across the sensors or whether the M170 in different conditions corresponds to the same pattern, but with different response magnitudes of that pattern. These measurements are achieved through a vector dot product over sensors that is either normalized (similarity between patterns) or not normalized (response magnitude). This is the same mathematical distinction as the difference between a correlation (which is normalized for variance) versus a regression slope. However, in the present case, these numbers can be geometrically interpreted in terms of the 157-dimensional sensor space. Thus, these measures are interpreted as similarity or magnitude rather than correlation or regression. The similarity measure produces the cosine of the angle between patterns in sensor space ("angle test of response similarity"), whereas the response magnitude measure produces the projection of one pattern onto the standard response ("projection test of response magnitude"). The magnitude measure was obtained by determining the peak latency using the root-mean square and taking the average magnitude of the 40-ms time window around the peak latency.

It is important to note that the angle test seeks the same measurement goal as the vector normalization technique of McCarthy and Wood (1985), although the two techniques are mathematically different. The angle test first takes the dot product of the two sensor vectors and then normalizes against vector length, whereas the technique of McCarthy and Wood first normalizes against vector length and then uses ANOVA to compare the two normalized vectors. This distinction is important, and unlike McCarthy and Wood's technique, which can produce an apparent topographic difference as an artifact of using vectors of sensor values that are not zero-centered (Urbach & Kutas, 2002), the angle test is similar to a correlation analysis and does not suffer from this problem.

Results

Behavioral results Collection of behavioral responses from individuals in the MEG chamber was constrained by the

recording setup. More specifically, although participants were told to respond quickly and accurately, they were also strongly urged not to move (particularly their head and shoulders) and not to blink, as this would cause artifacts. They responded by pressing a mouse button, which was out of sight down by their side (they were lying supine). As a result, responses times were slow and variable. However, accuracy was nearly perfect, indicating that they were indeed performing the requested task. Given these limitations, not much can be gained from their behavioral responses. Nevertheless, we report them for completeness. Accuracy and mean response times are shown in Table 1. Participants responded more accurately with long cue durations than with short cue durations [$F(1, 16) = 6.31$, $MSE = 0.001$, $p < .05$]. We found no significant effects for the response times. Given the relatively long response times (> 600 ms), it was unlikely that the MEG response of interest, the M170, was contaminated by motor processes.

Angle test of response similarity We first compared the target M170 responses in the four conditions (same/different targets following short/long cue durations) using the angle test to determine whether reliable differences existed between these conditions in the topographic pattern across the 157 sensors. Split-half reliability was determined by comparing trials from the first half of the experiment with trials from the second half (first/second). This comparison was done separately for first- versus second-half topographic patterns from the same condition as well as first- versus second-half patterns from different conditions. Thus, the within-condition comparisons provided a measure of the null hypothesis variance between the first versus second halves of the experiment, against which differences between conditions were compared for hypothesis testing.

These between/within condition comparisons were calculated for each subject. The two within-condition angles (same-first vs. same-second as well as different-first vs. different-second) were averaged for each participant to yield an overall within-condition value (for short duration, $M = 0.527$, $SD = 0.254$; for long duration, $M = 0.476$, $SD = 0.199$). The two between-condition angles (same-first vs. different-second as well as different-first vs. same-second) were also

averaged for each participant, to yield an overall between-condition value (for short duration, $M = 0.535$, $SD = 0.222$; for long duration, $M = 0.477$, $SD = 0.201$). In a dependent-samples test across the 17 participants, there was no significant difference in the similarity (angle test) of the between-condition comparisons versus the within-condition comparisons for both cue durations [short duration, $t(16) = 0.193$, $p = .849$; long duration, $t(16) = 0.017$, $p = .987$]. Thus, the topographic patterns across the target responses were not found to be different. This suggests that the same distribution of cortical responses was involved in the *same* and *different* conditions. Therefore, any magnitude differences were taken to indicate change in the magnitude of the underlying cortical responses rather than differential recruitment of cortical areas in some conditions as compared to others.

Projection test of response magnitude Because there did not appear to be differences in the M170 response patterns between same versus different targets, we next tested whether there was a difference in the M170 response magnitude. This was achieved through the geometric projection of the target word response onto a "standard response" pattern that is specific to each individual. The standard response should be a situation that involves the same kind of presentation and attentional state as the experimental condition. Therefore, for the present experiment, we used the M170 response to the cue in the long-duration conditions. Because half of all trials used a long-duration cue, this provided a fairly reliable M170 response pattern for each individual that was evoked by the onset of a single word in this same/different paradigm. The M170 response to the target for each of the four experimental conditions was then projected onto this standard response to yield four M170 response magnitude values for each individual.

Figure 2 highlights the need for a measure of response magnitude that disentangles the neural response to the target word from the neural response to the cue word. This figure shows global field power at each millisecond from the onset of the target word. Global field power is the square root of the average squared value across the 157 sensors (root-mean square). For comparison, this figure also shows the response to the cue word in the long conditions, which gives an indication of the response profile to a word presented in isolation (i.e., the response to a word in isolation without a prior cue word). As can be seen in the figure, the grand average waveforms for the two short-duration conditions overlap with the MEG response to the short-duration cue, whereas the response to the long cue is similar to the response to a target following a long cue (i.e., a relative lack of overlap for long conditions). The projection test estimates the magnitude of the M170 signal that is present within this overlapping signal, which should help isolate the component of the short conditions that is

Table 1 Results of Experiment 1, showing means of behavioral (accuracy, correct response time) measures (with standard deviations)

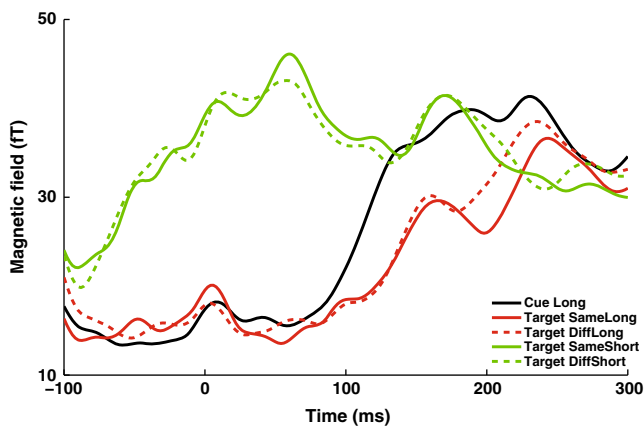| | Accuracy | | Response Time (ms) | |
| | Cue Duration | | | |
Trial Type	150 ms	2,000 ms	150 ms	2,000 ms
Same	.95 (.05)	.96 (.05)	615 (176)	623 (175)
Different	.96 (.03)	.98 (.02)	613 (183)	639 (216)

Fig. 2 Grand average root-mean-square waveforms (of all 157 channels across 17 participants) in response to the cue in the long-cue-duration condition (black) and to the target in the short- (green) and long- (red) cue-duration conditions in Experiment 1. The *different* trials are the dashed lines, and the *same* trials are the solid lines. The high level of activity in the responses to the target in the short conditions was caused by overlap with ongoing responses to the preceding cue. This overlap does not influence the comparison within the short-duration conditions

due to the presentation of the target. Furthermore, because the level of overlap is equal between *same* and *different* trials that test the same cue duration, *same* and *different* trial response magnitudes can be meaningfully compared.

Table 2 presents the M170 response magnitudes across trial types and cue durations. The magnitude of the MEG signal was higher following a long-duration cue [$F(1, 16) = 6.11$, $MSE = 0.45$, $p < .05$], but this was qualified by an interaction between trial type and cue duration [$F(1, 16) = 3.12$, $MSE = 0.01$, $p < .1$] due to an effect of trial type in the long-cue-duration condition [$t(16) = 2.463$, $p < .05$], but not in the short-cue-duration condition ($t < 1$). No significant effects were obtained for peak latency (all $ps > .3$). Response times did not correlate (within each condition) with MEG magnitude ($ps > .5$) or peak latency ($ps > .6$).

Discussion

In Experiment 1, we tested the hypothesis that there is a single change detection process underlying same/different

Table 2 Results of Experiment 1, showing means of MEG (magnitude, latency) measures (with standard deviations)

Trial Type	Projected Magnitude		Peak Latency (ms)	
	Cue Duration			
	150 ms	2,000 ms	150 ms	2,000 ms
Same	.098 (.26)	.239 (.15)	176 (13)	177 (16)
Different	.111 (.28)	.297 (.16)	176 (13)	178 (16)

judgments to target words presented immediately after cue words. We found that in the temporal window around the M170 in response to the target word, the response patterns across the 157 sensors were similar, regardless of whether the target was the same as or different than the cue word, which suggests that the same identification processes were involved for both trial types. Critically, as predicted, there was a main effect such that the magnitude of response was larger for *different* trials (novelty detection) than for *same* trials (lack of novelty detection). Additionally, there was an interaction with cue duration such that the effect of trial type was larger following a long-duration cue, which demonstrates that the signal underlying change detection with words is sensitive to the timing of the comparison word. This difference in the magnitudes of the target words' M170s as a function of cue duration replicates previous findings with a perceptual identification task (Huber et al., 2008), except that the present experiment used target words that were clearly visible. Furthermore, this result was predicted by a model of change detection based on habituation. Specifically, because habituation is greater following a long-duration cue, a target that differs from the cue was expected to produce a larger boost of additional activation (i.e., novelty). This result favors an account based on habituation and novelty detection rather than familiarity.

The habituation model supposes that a novelty response in the *different* condition is the signal that underlies accurate performance. Therefore, the model predicts that there should be a magnitude difference between the *same* and *different* trials for both the short- and long-cue-duration conditions, considering that even though accuracy was lower in the short cue duration, it was still very good. Thus, even though the novelty response was predicted to be weaker in the short-cue condition, it should still be present. However, the MEG projection test failed to find a reliable difference in the short-cue condition (although the trend was in the expected direction). Nevertheless, the failure to detect such a difference depends on statistical power. More specifically, even if a small but highly reliable novelty response underlies performance, MEG measurements may fail to detect this response with sufficient reliability.[1]

From our proposal, one might expect that the amplitude or latency of the M170 would predict response time. In other words, a stronger or a quicker novelty response should serve to promote faster accurate behavioral

---

[1] In other words, the neural signal that drives the decision is of high reliability (even though it is small), whereas the signal picked up at the MEG sensors is hugely variable due to electronic noise, noise from other brain regions, inadequate spatial sampling, and so on. On top of all that, the signal underlying MEG requires that neurons be spatially aligned (if they are not aligned, their signals cancel out). Thus, for many reasons, there could be too much noise with MEG to pick up on a predictably small signal.

responses. For example, Noguchi, Inui, and Kakigi (2004) used a repetition priming paradigm with letters while recording MEG, and they found that the peak latency, but not the peak amplitude, correlated with RT. In our experiment, we found no such correlation (indeed, response times were relatively uniform across conditions). However, Noguchi et al. examined a completely different neural signal, which occurred at 252 ms, almost 100 ms later than the M170 we focused on. A recent study has shown that event-related components that occur with longer delays after stimulus onset are more strongly associated with task difficulty and task-relevant perceptual processes, whereas earlier components are more associated with low-level perceptual categorization processes (Philiastides, Ratcliff, & Sajda, 2006; Philiastides & Sajda, 2006; Ratcliff, Philiastides, & Sajda, 2009). Our results suggest that the M170 is an early neural signal that discriminates between *same* and *different* trials. Furthermore, the consistency between the present results and the results of Huber et al. (2008) with a different task demonstrates that the M170 to words is not closely involved in decision processes.

Outline of experiment 2

Experiment 1 revealed that the first neural response to a target word that distinguished *same* and *different* trials was (a) similar in its topographic distribution between the two trial types, thus implicating the same neural signal for both, and (b) larger in magnitude for *different* trials, suggesting that the signal driving performance is detection of novelty rather than familiarity. These results fit well with the neurocomputational model of short-term priming (Huber & O'Reilly, 2003), for which the dynamics are shown in Fig. 1. We now turn to a counterintuitive, but necessary, prediction from this model, and test this prediction by examining response times in a same/different judgment task.

We used the proposal that habituation underlies change detection to make testable predictions for performance across subsequent *pairs* of cue/target trials (see Fig. 3). The task was the same as in Experiment 1, but now with a structure to the trials so that we could examine the effect of repetitions from one trial to the next, and thereby the effect of habituation on change detection. Because we were interested in change detection, we focused on trials in which the cue and target differed, and looked at the effects of priming the cue versus priming the target, with priming provided by the preceding cue–target trial. As a result of this focus, Experiment 2A used a go/no-go task, with go responses when the target differed from the cue, so for this experiment the only data available were for *different* trials. Next, we considered the predictions of a habituation account in this paradigm for cue and target priming.
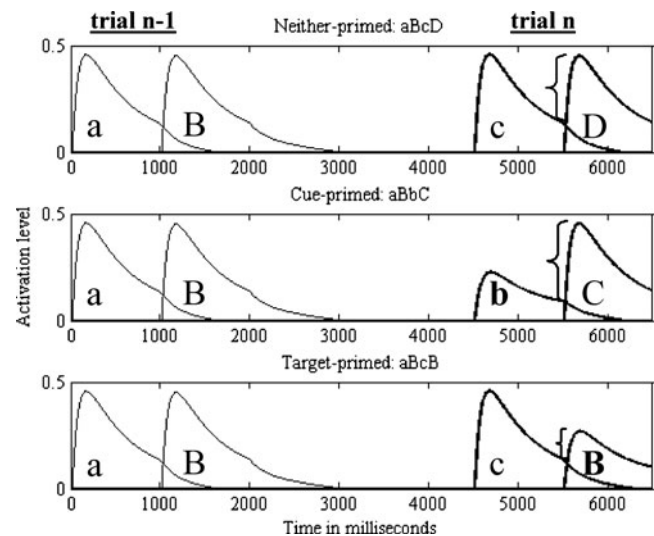

Fig. 3 Predictions from a neural network model with synaptic depression, showing activation levels as a function of time. For each panel, the activation profiles are for the cue and target of the previous trial (n − 1) and then the cue and target for the current trial (n), in that order. Lowercase letters indicate the cue word, which was presented in lowercase, whereas uppercase letters indicate the target word, which was presented in uppercase. Thus, the sequence of letters for each panel indicates different conditions in accord with repetitions (i.e., priming) from the target of the previous trial to the next trial. Shown are conditions in which both the previous trial and the subsequent trial involve cues that are different from the targets. For Experiment 2A, this would involve go responses on both trials, whereas in Experiment 2B, this would involve "different" responses on both trials. The panels represent scenarios in which the previous target does not prime either the cue or the target (top), primes the cue (middle), or primes the target (bottom) of the current trial. The critical comparison between the three conditions is the difference between the height of the final peak and the height of the activation profile at the onset of the target stimulus. This is highlighted by the brackets. As compared to the neither-primed condition, this difference is larger for the cue-primed and smaller for the target-primed condition. See the text for details

We used the following four-letter notation to indicate cue and target identity for a pair of trials in the critical conditions. Each letter symbolizes a word, and the order of the letters reflects the order of the presentation of the corresponding words (i.e., cue and target for the previous trial and cue and target for the present trial, in that order). Repeated letters indicate words that repeated within or across the two trials (the case of the letters reflects the case in which the words were presented: lower for cues vs. upper for targets). For trials that were preceded by a *same* or a *different* trial, the first two letters were aA or aB, respectively. To refer to the neither-primed conditions, the first two letters would differ from the last two letters (i.e., aAbC or aBcD). For cue-primed conditions, the second and third letters are the same (i.e., aAaB or aBbC). For target-primed conditions, the second and fourth letters are identical (i.e., aAbA and aBcB).

Figure 3 provides an illustration of the predictions of the habituation account, showing neural output activation

levels as a function of time during the sequential presentation of a cue–target pair on trial $n - 1$ followed by trial n. On each trial, a cue is presented for 1 s, followed by a target for 1 s, and then a break of 2.5 s between trials. On presentation of any stimulus, the neural activation in the model follows a nonmonotonic function due to habituation. When neither the cue nor the target of trial n is primed (upper panel of Fig. 3), the neural responses are similar for both cue and target. However, if the cue appeared on the preceding trial (middle panel), the neural response for the cue word is habituated on trial n. Based on the results from Experiment 1, we assume that the novelty response is calculated by monitoring the degree of additional activation with the onset of the target word. Thus, performance is related to the difference between the degree of activation at the offset of the cue and the peak activation value in response to the target, as shown by the brackets on each panel. For any given trial, this difference is affected by priming from the previous trial, because habituation tends to linger for seconds or tens of seconds. In the case of cue priming (middle panel of Fig. 3), habituation increases the difference between cue and target as compared to the neither-primed condition (top panel). However, in the case of target priming (bottom panel), this difference is smaller, because the habituated target representation does not achieve as high a peak value. If this difference in activation is the signal supporting change detection, the increased signal in the cue-primed condition should lead to better performance (i.e., faster responses) relative to the neither-primed condition. Likewise, the decreased signal in the target-primed condition should reduce performance (i.e., slower responses). Thus, an account based on change detection and habituation predicts a pattern of costs and benefits, with last-trial priming depending on which word is primed.

Although Fig. 3 shows specific predictions, these predictions reflect the neural response rather than behavior. Therefore, our predictions regarding response time behavior are qualitative and only address the ordinal pattern in the data. The amount of novelty is assumed to feed into a decision process, which is needed to make quantitative predictions about accuracy and response time. Therefore, our account is focused on the process that takes place before the actual decision stage. In order to contrast this account with several other alternative explanations that have appealed to mechanisms within the decision stage, we refer to this model as a *nondecision model* (e.g., something prior to the decision stage, such as the speed or the quality of the perceptual response).

We can contrast these predictions to those formulated based on simple response repetition (i.e., a bias in favor of the response given on the last trial). Unlike the habituation model, response repetition predicts no difference between

priming the cue versus priming the target, each as compared to the baseline neither-primed condition (see Table 3). Next, consider predictions from theories that explain priming effects through associative learning of stimulus–response contingencies (e.g., Hommel, 2004; Logan, 1988, 1990; Rothermund, Wentura, & De Houwer, 2005). These theories predict that if a word is repeated as a target on two consecutive trials and the same response is required (either "same" or "different") on both trials, then performance should be better than in the neither-primed baseline and better than when different responses are required on the two trials (see Table 3). Note that this is exactly the opposite prediction from the habituation account (cf. Fig. 3). In addition, the habituation account predicts that "different" responses to primed targets will be slowed, regardless of the response given on the previous trial to that same target word. Whereas the habituation model is focused on processes that occur before the decision stage, the response repetition account and the stimulus–response account are examples of decision-stage models. As such, these types of models compete directly with the theoretical accounts of same/different judgments reviewed in the introduction. A third decision-stage model, an interference model, is one in which the target might be mistakenly compared with the preceding cue and the target of the previous trial. When these comparisons are congruent, the corresponding response time is faster than when the comparison is incongruent. As can be seen in Table 3, this alternative model makes the same prediction as the habituation model when the targets on successive trials repeat, but not when the cue was a target in the previous trial.[2] More specifically, the interference account does not predict an effect for the cue-primed condition, because in both the cue-primed and neither-primed conditions the target was not seen on the previous trial, and thus mistakenly comparing the target to the words of the previous trial will suggest the correct response of "different" for both conditions. This same logic explains why the stimulus–response theory also predicts no effect in the cue-primed condition—because the target was not seen on the previous trial for both the cue-primed and neither-primed conditions, it has no associated response.

We tested the predictions of the habituation account in an experiment in two parts that involved priming from previous trials by repeating the target word across trials. In Experiment 2A, we investigated change detection by asking participants to make a response whenever a change occurred. This was achieved through a go/no-go task in which a response was required when the target word was *different* than the preceding cue word. To directly contrast the habituation account with the stimulus–response learning

---

[2] We thank an anonymous reviewer for suggesting this type of decision-stage account.

Table 3 Predictions of three decision-based models and the habituation model for the change in response times between a primed trial (cue, target, or both) versus the corresponding neither-primed trial in Experiment 2

Neither-Primed Condition	Primed Trial	Decision-Stage Models			Non-Decision-Stage Model
		Response Repetition	Stimulus–Response	Interference	Habituation
aAbC	aAaB	Slower	No difference	No difference	**Faster**
	aAbA	**Slower**	**Slower**	**Slower**	**Slower**
aBcD	aBbC	**Faster**	No difference	No difference	**Faster**
	aBcB	Faster	Faster	**Slower**	**Slower**
aBcC	aBbB	Slower	Slower	**Faster**	**Faster**

Observed results are shown in boldface

account, we conducted Experiment 2B, which was identical to Experiment 2A in terms of the actual sequence of presentations, but which measured change detection with same/different rather than go/no-go responses. Note that both of these tasks concern change detection within a trial rather than between trials. Thus, stimuli presented in the previous trial provided no information about the required response on the current trial.

Beyond priming of the cue versus priming of the target, which contrasts the predictions of the habituation account with the three decision-stage accounts, we also tested the specific predictions of the habituation model regarding the duration and timing of events. Specifically, we investigated the role of cue duration. On trials in which the cue duration is brief, the activation in response to the cue will be high at the onset of the target, so a target that is different from the cue will not provide much *additional* activation (i.e., change detection will be difficult). However, if the cue duration is long, the response to the cue will be strongly habituated, and it will be easy to detect a target that is different. Therefore, we manipulated the duration of the cue word (200 vs. 1,000 ms). This was done in a blocked design to allow anticipation of the target onset, which would otherwise provide a source of noise in the response times. Provided that last-trial priming was equivalent in all conditions, the habituation model predicted that longer cue durations should make it easier to detect changes, regardless of whether the cue or target was primed. Thus, the difference between the cue-primed and target-primed conditions should diminish with increasing cue duration. However, because the experiment was run with cue duration in a blocked fashion (i.e., if the cue duration on trial *n* was long, the cue duration on trial *n* − 1 was also long), and because a previous *same* trial introduces greater priming than a previous *different* trial, this introduced some complexities into the predictions that are considered in greater detail in the Discussion section.

The durations differed somewhat from those used in Experiment 1. Experiment 1 needed to include a very long

cue duration to eliminate overlap in the neural response between cue and target, so 2,000 ms was used for the long-cue condition. In contrast, in Experiment 2 we used only 1,000 ms for the long-cue condition to reduce the total time of the experiment. A more important difference between Experiment 1 versus Experiment 2 was how often words repeated across trials. In Experiment 1, words never appeared on more than one trial. However, in Experiment 2, the same words not only repeated across contiguous trials, but also on many other trials throughout the experiment. This was a crucial element of the design. As discussed in the introduction, habituation is a short-term effect that biases against attended stimuli, leading to slow responses. In addition to the deleterious effect of habituation, there are also positive long-term priming effects (e.g., Hertwig et al., 2008). Thus, each word on trial *n* − 1 would produce a beneficial long-term priming effect together with a short-lived habituation effect, both of which might affect performance on trial *n*. In order to separate these two opposing forces, we used a small pool of words and drew with replacement from this pool in order to equate the words in terms of the beneficial long-term priming effect. In other words, regardless of last-trial priming, all of the words on every trial (after practice) were subject to the benefits of long-term priming. This allowed us to isolate the short-term habituation effect by comparing cue- or target-primed trials against the baseline neither-primed trials.

Experiment 2

Method

Participants Groups of 56 and 53 college-aged students from Indiana University participated in exchange for course credit in Experiments 2A and 2B, respectively.

Design Both Experiments 2A and 2B conformed to a 3 × 2 × 2 within-subjects factorial design, crossing the factors

prime type (neither-primed, cue-primed, target-primed), previous response (Exp. 2A, no go vs. go; Exp. 2B, same vs. different), and cue duration (200 vs. 1,000 ms).

Materials The materials used in both experiments consisted of the following 12 four-letter words: *back*, *give*, *most*, *turn*, *hold*, *main*, *such*, *type*, *blue*, *kind*, *stop*, *farm*, *best*, *down*, *late*, and *size*. All words could appear as either cues or targets, with cue words presented in lowercase and target words in uppercase black letters on a white background in font size 18 (see Fig. 4). Each unique condition was repeated 12 times, which were defined by the relation between the target of trial *n* − 1 and the cue and target of trial *n*. Thus, a cue-primed trial was a trial on which the target of the previous trial was the same word as the cue of the current trial (e.g., Trial 3 in Fig. 4). A target-primed trial was a trial on which the target of the previous trial was the same as the target of the current trial (e.g., Trial 4 in Fig. 4). A particular cue–target pair on a given trial could provide a measurement of priming from the previous trial, or it might serve as the prime for the subsequent trial. There were equal numbers of *same* and *different* trials included in the analyses. The full list of trials presented to the participant used more *different* than *same* trials, because unanalyzed filler trials (e.g., Trial 1 in Fig. 4) were used at the beginning of each subblock of seven trials. A subblock of seven trials was a sequence of trials that started with a *different* trial, followed by six experimental trials. Several seven-trial sequences were created that together covered all combinations with 12 observations per condition. This was necessary to balance the design and to make sure that participants were engaging in the task at the beginning of each subblock (as measured by a keypress). Overall, there were 192 *different* trials and 144 *same* trials.

Procedure Participants were seated in front of a computer screen. Each trial started with a black fixation stimulus, "++," for 1,000 ms, followed by a cue word in lowercase that

was presented for either 200 or 1,000 ms. Immediately after the cue word, the target word appeared and remained on the screen for 2,000 ms or until the participant pressed the space bar. In Experiment 2A, the task of the participants was to press the space bar as quickly as possible only when the target word was *different* from the cue word. No response was to be given when the two words were the same. Figure 4 shows an example of a sequence. In Experiment 2B, participants responded with the "z" key or the "/" key to indicate whether the target word was the same as or different than the cue word. The mapping of the keys to the "same" and "different" responses was counterbalanced across participants. To train participants on this key–response mapping, we split the 56 practice trials into two blocks of 28 trials. In the first practice block, the labels "same" and "different" appeared in the left and right lower corners of the screen, in correspondence to the key mapping for that participant. In the second practice block, these labels were absent. If no response was made within 2,000 ms, "Too slow" was presented in red font color.

Feedback was given after every trial for 1,500 ms. A correct go trial was followed by the response time in blue font color, and a correct no-go trial was followed by "0.000 s" in blue font color. All incorrect trials were followed by "XXXX" in red font color. Before the experimental trials, there were 56 practice trials, in the following order: 7 long–7 short–7 long–7 short–14 long–14 short. After the practice block, participants received feedback about their accuracy in percentage correct. The main experiment had four blocks of 84 trials. Each block consisted of 12 subblocks, which were grouped into sets of 3 that used the same cue duration. These groupings were such that participants experienced trials of different cue durations in the following order: 21 long–21 short–21 long–21 short. After each block, participants received feedback about their percentage correct (across all correct trials). All stimuli (fixation, cue, target, and feedback) were presented in the center of the computer screen.

Fig. 4 Example sequence with the three conditions used in Experiment 2. Although the sequence has four *different* trials, the first is considered a filler trial. Trials 2, 3, and 4 are a neither-primed, a cue-primed, and a target-primed trial, respectively
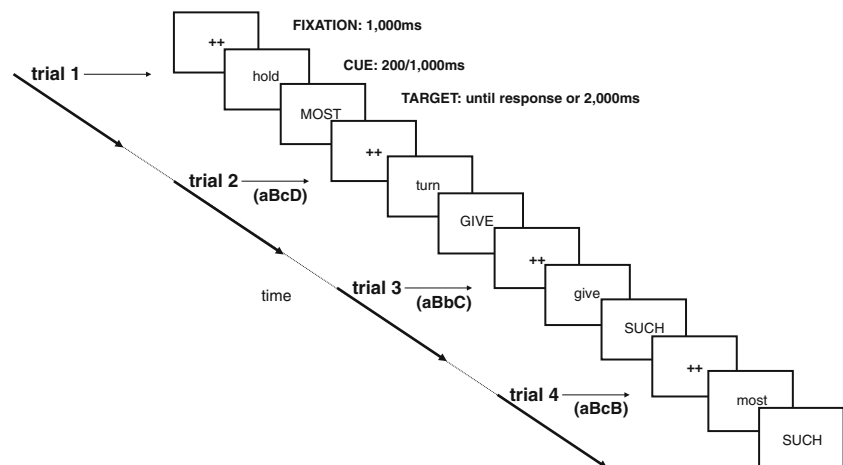
Results

We analyzed the data with experiment (2A vs. 2B) as a between-subjects factor. The detailed results and summaries of the data of each experiment are presented in Online supplemental materials. Figures 5 and 6 show the results of Experiments 2A and 2B, respectively. Only RT data from trials for which the preceding and current trials were correct were included in the analyses. Table 4 presents results from the omnibus $3 \times 2 \times 2 \times 2$ (prime type, cue duration, previous response, experiment) mixed factorial ANOVA conducted on response times and accuracy for the conditions in which the current trial involved a cue and target that were different.

Accuracy There were main effects of cue duration ($p < .01$), prime type ($p < .01$), and experiment ($p < .001$). The experiments differed in the effects of the previous trial and the effect of prime type, as indicated by significant interactions between experiment and previous trial type ($p < .05$) and between experiment and prime type ($p < .001$). These differences are fully unpacked in the Supplemental materials.

Response times For response times to *different* trials, the ANOVA revealed a main effect of prime type ($p < .001$), reflecting slower RTs for target-primed trials and faster RTs for cue-primed trials. There was also a main effect of experiment ($p < .001$), reflecting a general slowing with two versus one response alternative.

Of the two-way interactions that did not involve the experiment factor, all were significant ($ps < .05$). The only two-way interaction with the experiment factor that was significant was previous response ($p < .001$), which again reflected the difference in number of response alternatives between the two experiments.

Notably, the analysis revealed that none of the three-way interactions involving the experiment factor were significant, demonstrating that the predicted interaction between cue duration and prime type existed for both experiments. This interaction was due to a reduced prime type effect for long-duration-cue trials. As predicted, the advantage for priming the cue and the disadvantage for priming the target were reduced as cue duration increased.

The three-way interaction between cue duration, previous response, and prime type was significant ($p < .001$), which was due to a smaller effect of prime type with a long cue duration only after *different* trials (see Figs. 5 and 6). This three-way interaction was stronger in Experiment 2B than in Experiment 2A (see Figs. 5 and 6), as indicated by the four-way interaction. A closer look at the raw data of Experiment 2A revealed that the four-way interaction was driven by unusually fast RTs for the baseline neither-primed condition, aBcD, only for the short-cue-duration go trials of Experiment 2A. This eliminated the cue-priming effect only for this condition.

Discussion

In Experiments 2A and 2B, we tested the hypothesis that the mechanism of change detection is habituation, which diminishes the response to repeated stimuli. Using a neural model of synaptic depression, we made several specific predictions regarding the pattern of RTs for different priming conditions. Synaptic depression produces a stronger response to a target when it is preceded by a different cue stimulus as compared to a repetition. By manipulating the amount of depression of the cue or target stimulus through priming from the previous trial, we tested the predictions that RTs in the cue-primed condition should be faster than those in the neither-primed condition and that RTs in the target-primed condition should be slower than those in the neither-primed condition. This pattern of results was confirmed in Experiment 2 and constitutes the main test of the habituation account.

Two results from these experiments are inconsistent with an account based on stimulus–response learning (see Table 3). First, both versions of the experiment revealed slower responses in the target-primed (aBcB) conditions

Fig. 5 Results of Experiment 2A for go trials only: priming effects (RT in baseline minus RT in priming condition) of correct response times as a function of prime type, previous trial type, and cue duration. The error bars represent standard errors of the means
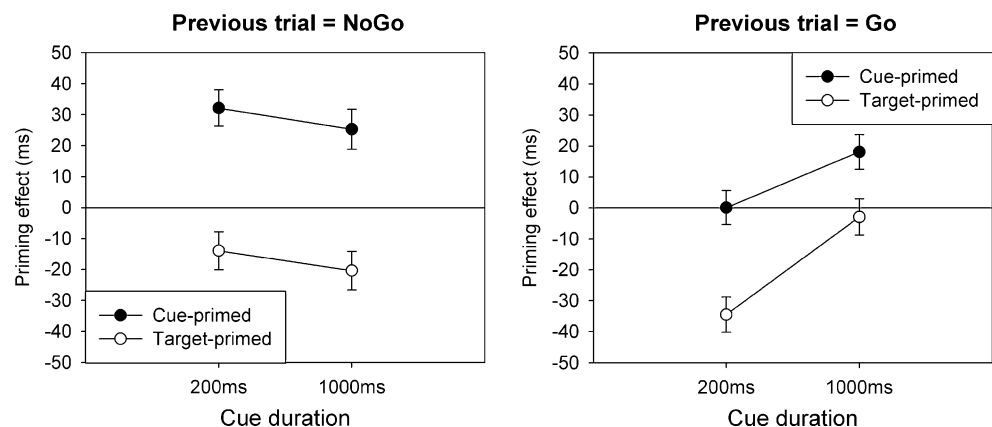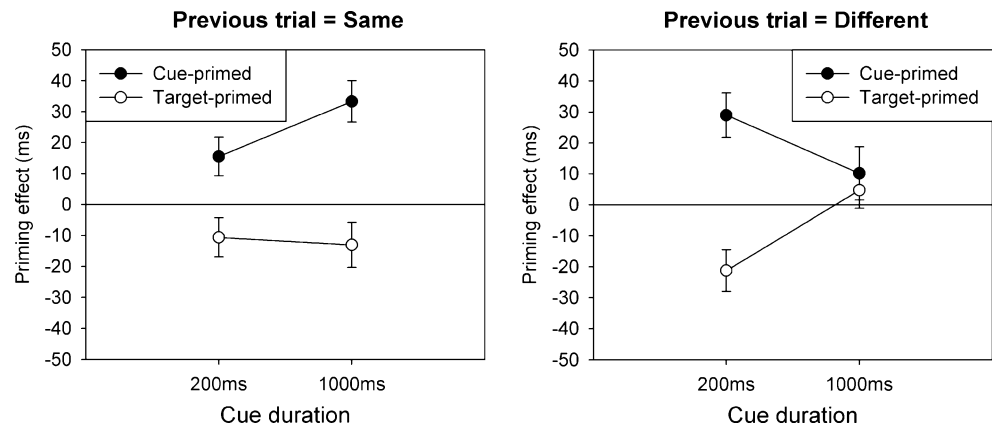
Fig. 6 Results of Experiment 2B for *different* trials only: priming effects (RT in baseline minus RT in priming condition) of correct response times as a function of prime type, previous trial type, and cue duration. The error bars represent standard errors of the means

than in the baseline neither-primed conditions (aBcD), whereas stimulus–response theories (as well as the interference account) predicted the opposite. Second, for the same/different task of Experiment 2B, for which stimulus–response theories predict slower RTs for both-primed (aBbB) than for neither-primed conditions (due to inconsistent stimulus–response mappings), we observed the opposite pattern. These results suggest that processes other than stimulus–response learning have a profound influence on the response time.

We predicted that longer cue durations would make it easier to detect changes, which should tend to dilute the effect of priming the cue or target. Although the duration× prime type interaction did not reach significance in either experiment, this was apparently due to a lack of power, because pooling the data revealed this interaction, which was not further modulated by the experiment factor. The lack of this two-way interaction was in part due to its being

masked by a three-way interaction, which included the type of response on the preceding trial. In particular, the expected two-way interaction was only found for trials that followed a *different* trial. The habituation model predicts that increasing the cue duration should diminish the priming difference between the cue-primed and target-primed conditions, but this would only be true if priming from the last trial were equivalent regardless of cue duration. However, the experimental design manipulated cue duration in a blocked fashion, and this introduced an unanticipated confounding factor—as cue duration increased, this increased the amount of priming when the last trial was a *same* trial (i.e., aA) but not when the last trial was a *different* trial (i.e., aB). Thus, when the last trial was a *same* trial, increasing the cue duration produced two counteracting effects: (1) An increased cue duration on the current trial served to enhance novelty detection, and thus reduce priming effects, and (2) an increased cue duration on

Table 4 Results of an analysis of variance on mean correct response times and accuracy for go trials and *different* trials across Experiments 2A and 2B

Effect	df	Accuracy		Response Time	
		MSE	F	MSE	F
Duration (Dur)	1,106	0.002	8.331[**]	4,528.403	0.057
Previous (Prev)	1,106	0.002	0.330	2,314.273	0.670
Prime Type (Prime)	2,212	0.002	5.899[**]	1,246.080	103.842[***]
Experiment (Exp)	1,106	0.007	21.058[***]	58,631.012	16.219[***]
Dur×Prev	1,106	0.002	0.007	1,712.231	9.300[**]
Dur×Prime	2,212	0.001	1.441	1,102.257	4.020[*]
Prev×Prime	2,106	0.002	0.822	1,261.257	4.650[*]
Dur×Exp	1,106	0.002	0.211	4,528.403	0.007
Prev×Exp	1,106	0.002	4.808[*]	2,314.273	69.182[***]
Prime×Exp	2,212	0.002	8.560[***]	1,246.080	1.377
Dur×Prev×Prime	2,212	0.001	0.449	1,308.523	9.168[***]
Dur×Prev×Exp	1,106	0.002	0.012	1,712.231	1.420
Dur×Prime×Exp	2,212	0.001	1.527	1,102.257	0.267
Prev×Prime×Exp	2,212	0.002	1.428	1,261.257	1.206
Dur×Prev×Prime×Exp	2,212	0.001	0.126	1,308.523	5.655[**]

[*] $p < .05$, [**] $p \leq .01$, [***] $p \leq .001$

the previous trial served to enhance the degree of habituation for the primed word, and thus increase priming effects. As reported in the Online supplementary materials, we modeled the exact sequence of events across the two trials with different parameter values and found that the observed three-way interaction between response type of the last trial, prime type, and cue duration was predicted by the model, depending on the balance between the rate of habituation and the rate of recovery from habituation.

Whereas Experiments 2A and 2B confirmed our central hypotheses, they produced different response priming effects. In particular, in Experiment 2A, RTs were *faster* after go trials than after no-go trials, whereas in Experiment 2B, RTs were *slower* after *different* trials than after *same* trials. This main effect of the previous trial in Experiment 2A was expected when considering the literature examining response suppression. The slowdown in RTs after successful stopping (Rieger & Gauggel, 1999; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008) has been explained as an increase in the response threshold through online monitoring of decision conflict (Davelaar, 2009; Emeric et al., 2007; Verbruggen & Logan, 2008). It has been hypothesized that this conflict is used to control performance through a neural network that underlies behavioral inhibition (Aron & Poldrack, 2006). The influence of this network in the context of a go/no-go paradigm is a transient increase in the effective response threshold (cf. Frank, 2006), which is diminished, or absent, when a *same* trial requires an overt response (such as in Exp. 2B). The increase in RTs after no-go trials can be accounted for as the aftereffect of a transient increase in response threshold.

A second difference between the two experiments was the overall faster RTs in Experiment 2A as compared to Experiment 2B, irrespective of cross-trial aftereffects. This difference in RTs when comparing the number of response alternatives has been documented and interpreted in a study comparing a go/no-go task to two-choice lexical decision (Gomez, Ratcliff, & Perea, 2007).

General discussion

We investigated a theoretical account of change detection that posits a crucial role for novelty detection, as implemented through synaptic depression (Huber, 2008b; Huber & O'Reilly, 2003). In Experiment 1, we obtained support for a critical assumption of the habituation account. Using MEG methodology, we failed to find evidence of different cortical sources for *same* and *different* trials, so the data were consistent with the claim that both responses result from the same novelty detection process. Crucially, the magnitude of the neural response was larger for *different*

than for *same* trials, which makes an alternative interpretation based on familiarity detection rather than change detection less plausible. This result can be related to the findings seen at the level of single cells in terms of frequency of spiking (Tsodyks & Markram, 1997). Furthermore, even though we observed this result by examining activity over all 157 sensors, it replicated previous MEG studies that had found reduced signal for repeated stimuli using a sensor-of-interest approach (Noguchi et al., 2004).

In Experiment 2, we confirmed the predictions from a neural habituation account of change detection. Assuming that presentation of a word on the previous trial produces lingering habituation, this account explains the finding of faster responding in the cue-primed conditions and slower responding in the target-primed conditions. Habituation predicted enhanced novelty detection for priming of the cue, due to a bigger change in activation when a non-habituated target followed a habituated cue. This account predicted decreased novelty detection for priming of the target, due to a smaller change in activation when a habituated target followed a nonhabituated cue. Indeed, both experiments revealed the predicted pattern, regardless of the response compatibility between subsequent trials. This suggests that there was little or no stimulus–response learning, which was appropriate, considering that only 12 words were used, such that the same words appeared on many trials, with an equal numbers of "same" and "different" responses. By repeating words in this fashion, we reduced the effects of stimulus–response learning and equated the effects of long-term repetition priming, thus isolating the short-lived habituation effect from the last trial.

The predictions of the habituation account differ from those of decision-stage theories, such as those based on stimulus–response learning. In Experiment 2, we found that nonmatch responses to a target stimulus were *faster* when the cue stimulus had been presented on the preceding trial, and that nonmatch responses to a target stimulus were *slower* when the target stimulus had been presented on the preceding trial. Importantly, this pattern occurred regardless of response repetition and of the number of overt responses (i.e., it occurred in both experiments). To highlight this issue, consider the situation in which the same target word appears on consecutive trials and the same response was given on both trials (e.g., a response that the cue and target were "different," such as with aBcB). According to theories of the repetition of specific stimulus–response mappings (Logan, 1990), the time to respond should be faster on the second trial as compared to a neutral condition, because the repeated presentation of the target activates the associated response. However, according to the habituation view, repeated presentation of the target leads to less activation of the target representation and thus to lower novelty,

resulting in *slower* response times. Slower responses were indeed observed. This result does not rule out the existence of stimulus–response associations, but, at a minimum, it suggests that the role of stimulus–response associations was small relative to the effect of habituation.

The MEG data from Experiment 1 are consistent with the view that same/different judgments are based on a single signal (rather than separate "same" and "different" signals). However, these data do not place strong constraints on the architecture of the decision process—which, in the spirit of sequential sampling models, for example, could involve separate accumulators for the different responses that are driven by this signal. Such models have featured in the debate regarding single- and dual-process models of same/different judgments (Farell, 1985; Ratcliff, 1985; Van Zandt et al., 2000). The present work extends those models by addressing whether the signal that is used is one of sameness (i.e., match, familiarity) or of different-ness (i.e., mismatch, novelty). We found evidence favoring the view that the judgment is based on a single signal that conveys the amount of novelty (instead of familiarity).

Habituation likely works on multiple levels, ranging from low-level visual features to orthography, phonology, and semantics (see Huber, 2008b). Our data do not allow us to discriminate between the different contributions of different types of representations. We note, however, that cue-primed conditions always included a case change between the prime and the cue (aAaB or aBbC), whereas both prime and target were presented in uppercase for the target-primed conditions (aAbA or aBcB). The fact that the predictions of the habituation account were verified independently of case changes suggests that habituation occurs at representational levels beyond pure perceptual levels. Taken together, these behavioral and neural results add support to the cumulative body of evidence that repetition-induced suppression is the result of neural fatigue, and that this fatigue is present throughout a number of neural processing levels, producing concomitant effects on behavior.

Grill-Spector et al. (2006) reviewed a number of explanations of repetition-induced suppression of neuronal responses, all of which assumed that a repeated stimulus is processed by the same neural substrate. A habituation model posits that repetition-induced suppression is due to reduced neural responses, whereas a sharpening model posits that fewer neurons participate in the processing of a repeated stimulus. A third alternative, the facilitation model, posits that the same neurons are faster in processing the repeated stimulus. These explanations can be contrasted with associative learning accounts (e.g., Hommel, 2004; Logan, 1988, 1990; Rothermund et al., 2005), which explain repetition-induced suppression as a consequence of simply bypassing the initial processing levels. Our

results strongly support the habituation model, because it is the only account consistent with the behavioral and MEG data presented here. This habituation account has been used to explain negative repetition priming in a face identification task (Rieth & Huber, 2010) and in a word identification task, as evidenced behaviorally (e.g., Huber, 2008b) and by means of ERP/MEG (Huber et al., 2008). The present work extended this account to repetition priming of same/different judgments with words. This suggests that the model can be applied quite generally, although it remains to be seen whether this account can explain the entirety of the repetition suppression phenomena.

Unlike the present study of change detection, which involves a comparison of two words, change detection has been widely studied with more complex tasks in which several objects are presented simultaneously (Awh, Barton, & Vogel, 2007; Luck & Vogel, 1997; Rouder et al., 2008). After a blank interval, equal numbers of objects are presented at the same locations, and participants have to indicate whether the second display contains an object substitution. The results typically show high accuracy with less than four objects, and set-size-dependent decreases in accuracy with more than four objects. This has been interpreted as reflecting the capacity limitation of visual working memory (Cowan, 2001; Luck & Vogel, 1997; but see Wilken & Ma, 2004, for an alternative view).

The view that neural habituation is the mechanism underlying change detection is appealing in light of these limitations of visual working memory capacity. In contrast to accounts that involve comparisons of actively maintained objects, change detection through habituation is computationally efficient, in that it subtracts from rather than adds to the maintenance load. Although a number of models have been proposed to account for performance in tasks of visual working memory (Palmer, 1990; Pashler, 1988; Wilken & Ma, 2004), these do not usually constitute a full process model of the task used by Luck and Vogel (1997; but see Johnson, Spencer, & Schöner, 2009). Some models that have posited active maintenance of multiple objects (e.g., Macoveanu, Klingberg, & Tegnér, 2006) have not (yet) been extended to account for change detection tasks in visual working memory. We propose that an extension involving habituation of neural representations could provide a sensible account of change detection in these tasks.

In a recent model based on dynamic field theory, it was proposed that change detection in visual working memory tasks is mediated by an interactive network of three neural layers (Johnson, Spencer, & Schöner, 2009; Johnson, Spencer, et al. 2009). The model is able to generate repetition suppression through inhibition of the input layer that comes from an intermediate layer, which in turn is maintained by a memory layer. Although the model is much

more complex than the single-layer model we have used, it does show the profile of habituation. Therefore, this model is also well-suited to capture the qualitative profile of Experiment 2. However, it is unclear whether this three-layer model, which spans from early visual areas to prefrontal areas, is able to capture the pattern of MEG topography seen in Experiment 1. In any case, the fact that this model and the habituation model both include neurally plausible mechanisms for repetition suppression and are both consistent with our behavioral results lends further credence to our claim that change detection arises through a single nondecision novelty detection process.

Conclusion

Our investigation was driven by the question of whether detection of change on brief time scales is based on both familiarity and novelty, or mainly on one of these signals. Combining evidence from MEG recordings with theory-driven analyses of choice response data allowed us to identify that novelty detection likely underlies both detection of change and detection of sameness in a word change detection task, as measured with same/different responses and with go/no-go responses. Furthermore, we found support for the claim that this novelty detection is enabled by transient habituation of neural responses. This interpretation contrasts with the central role of familiarity in theories of episodic recognition memory, which operate at much longer time scales. Such apparent discrepancies across tasks suggest that cognitive processing flexibly adapts to task demands, exploiting the dynamics of different signals (e.g., short-lived novelty vs. long-lived familiarity) that could aid performance.

References

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

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *Journal of Neuroscience, 26,* 2424–2433. doi:10.1523/JNEUROSCI.4682-05.2006

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science, 18,* 622–628. doi:10.1111/j.1467-9280.2007.01949.x

Chance, F. S., Nelson, S. B., & Abbott, L. F. (1998). Synaptic depression and the temporal response characteristics of V1 cells. *Journal of Neuroscience, 18,* 4785–4799.

Chouinard, P. A., Morrissey, B. F., Kohler, S., & Goodale, M. (2008). Repetition suppression in occipital-temporal visual areas is modulated by physical rather than semantic features of objects. *NeuroImage, 41,* 130–144.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *The Behavioral and Brain Sciences, 24,* 87–185. doi:10.1017/S0140525X01003922

Curran, T. (2000). Brain potentials of recollection and familiarity. *Memory & Cognition, 28,* 923–938.

Davelaar, E. J. (2009). Conflict-monitoring and (meta)cognitive control. In J. Mayor, N. Ruh, & K. Plunkett (Eds.), *Connectionist models of behaviour and cognition II* (pp. 241–252). Singapore: World Scientific.

Dobbins, I. G., Schnyer, D. M., Verfaellie, M., & Schacter, D. L. (2004). Cortical activity reductions during repetition priming can result from rapid response learning. *Nature, 428,* 317–319. doi:10.1038/nature02400

Emeric, E. E., Brown, J. W., Boucher, L., Carpenter, R. H. S., Hanes, D.P., Harris, R., … Schall, J. D. (2007). Influence of history on saccade countermanding performance in humans and macaque monkeys. *Vision Research, 47,* 35–49.

Farell, B. (1985). "Same"–"different" judgments: A review of current controversies in perceptual comparisons. *Psychological Bulletin, 98,* 419–456. doi:10.1037/0033-2909.98.3.419

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks, 19,* 1120–1136.

Fuhrmann, G., Cowan, A., Segev, I., Tsodyks, M., & Stricker, C. (2004). Multiple mechanisms govern the dynamics of depression at neocortical synapses of young rats. *The Journal of Physiology, 557,* 415–438.

Ganel, T., Gonzalez, C. L. R., Valyear, K. F., Culham, J. C., Goodale, M. A., & Kohler, S. (2006). The relationship between fMRI adaptation and repetition priming. *NeuroImage, 32,* 1432–1440.

Gomez, P., Ratcliff, R., & Perea, M. (2007). A model of the go/no-go task. *Journal of Experimental Psychology. General, 136,* 389–413. doi:10.1037/0096-3445.136.3.389

Gotts, S. J., & Plaut, D. C. (2002). The impact of synaptic depression following brain damage: A connectionist account of "access" and "degraded-store" semantic impairments. *Cognitive, Affective, & Behavioral Neuroscience, 2,* 187–213.

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: Neural models of stimulus-specific effects. *Trends in Cognitive Sciences, 10,* 14–23. doi:10.1016/j.tics.2005.11.006

Grill-Spector, K., & Malach, R. (2001). fMR-adaptation: A tool for studying the functional properties of human cortical neurons. *Acta Psychologica, 107,* 293–321.

Guo, C., Lawson, A. L., & Jiang, Y. (2007). Distinct neural mechanisms for repetition effects of visual objects. *Neuroscience, 149,* 747–759.

Henson, R. N., Cansino, S., Herron, J. E., Robb, W. G. K., & Rugg, M. D. (2003). A familiarity signal in human anterior medial temporal cortex. *Hippocampus, 13,* 259–262.

Henson, R. N., Mouchlianitis, E., Matthews, W. J., & Kouider, S. (2008). Electrophysiological correlates of masked face priming. *NeuroImage, 40,* 884–895.

Hertwig, R., Herzog, S. M., Schooler, L. J., & Reimer, T. (2008). Fluency heuristic: A model of how the mind exploits a by-product of information retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34,* 1191–1206.

Hintzman, D. L., & Curran, T. (1995). When encoding fails: Instructions, feedback, and registration without learning. *Memory & Cognition, 23,* 213–226. doi:10.3758/BF03197223

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences, 8,* 494–500. doi:10.1016/j.tics.2004.08.007

Horner, A. J., & Henson, R. N. (2008). Priming, response learning and repetition suppression. *Neuropsychologia, 46,* 1979–1991.

Huber, D. E. (2008a). Causality in time: Explaining away the future and the past. In M. Oaksford & N. Chater (Eds.), *The probabilistic mind: Prospects for rational models of cognition* (pp. 351–376). Oxford: Oxford University Press.

Huber, D. E. (2008b). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology. General, 137,* 324–347. doi:10.1037/0096-3445.137.2.324

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,* 403–430. doi:10.1016/S0364-0213(03)00012-0

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002a). Mechanisms of source confusion and discounting in short-term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 28,* 1120–1136. doi:10.1037/0278-7393.28.6.1120

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review, 108,* 149–182. doi:10.1037/0033-295X.108.1.149

Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002b). Mechanisms of source confusion and discounting in short-term priming: 1. Effects of prime duration and prime recognition. *Memory & Cognition, 30,* 745–757. doi:10.3758/BF03196430

Huber, D. E., Tian, X., Curran, T., O'Reilly, R. C., & Woroch, B. (2008). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology. Human Perception and Performance, 34,* 1389–1416. doi:10.1037/a0013625

James, T. W., & Gauthier, I. (2006). Repetition-induced changes in BOLD response reflect accumulation of neural activity. *Human Brain Mapping, 27,* 37–46. doi:10.1002/hbm.20165

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009a). A dynamic neural field model of visual working memory and change detection. *Psychological Science, 20,* 568–577.

Johnson, J. S., Spencer, J. P., & Schöner, G. (2009b). A layered neural architecture for the consolidation, maintenance, and updating of representations in visual working memory. *Brain Research, 1299,* 17–32.

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

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review, 95,* 492–527. doi:10.1037/0033-295X.95.4.492

Logan, G. D. (1990). Repetition priming and automaticity: Common underlying mechanisms? *Cognitive Psychology, 22,* 1–35. doi:10.1016/0010-0285(90)90002-L

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281. doi:10.1038/36846

Macoveanu, J., Klingberg, T., & Tegnér, J. (2006). A biophysical model of multiple-item working memory: A computational and neuroimaging study. *Neuroscience, 141,* 1611–1618. doi:10.1016/j.neuroscience.2006.04.080

McCarthy, G., & Wood, C. C. (1985). Scalp distributions of event-related potentials: An ambiguity associated with analysis of variance models. *Electroencephalography and Clinical Neurophysiology, 62,* 203–208.

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,* 375–407. doi:10.1037/0033-295X.88.5.375

Miller, E. K., & Desimone, R. (1994). Parallel neuronal mechanisms for short-term memory. *Science, 263,* 520–522. doi:10.1126/science.8290960

Monahan, P. J., Fiorentino, R., & Poeppel, D. (2008). Masked repetition priming using magnetoencephalography. *Brain and Language, 106,* 65–71.

Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What you see is what you set: Sustained inattentional blindness and the capture of awareness. *Psychological Review, 112,* 217–242. doi:10.1037/0033-295X.112.1.217

Noguchi, Y., Inui, K., & Kakigi, R. (2004). Temporal dynamics of neural adaptation effect in the human visual ventral stream. *Journal of Neuroscience, 24,* 6283–6290.

Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology. Human Perception and Performance, 16,* 332–350.

Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics, 44,* 369–378. doi:10.3758/BF03210419

Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural representation of task difficulty and decision making during perceptual categorization: A timing diagram. *Journal of Neuroscience, 26,* 8965–8975.

Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex, 16,* 509–518.

Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. *Psychological Review, 92,* 212–225. doi:10.1037/0033-295X.92.2.212

Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Sciences, 106,* 6539–6544.

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368–373. doi:10.1111/j.1467-9280.1997.tb00427.x

Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology, 90,* 509–518.

Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. *Journal of Experimental Psychology. Human Perception and Performance, 36,* 596–618.

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,* 114–129.

Rothermund, K., Wentura, D., & De Houwer, J. (2005). Retrieval of incidental stimulus–response associations as a source of negative priming. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 31,* 482–495. doi:10.1037/0278-7393.31.3.482

Rouder, J. N., Morey, R. D., Cowan, N., Zwilling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of discrete-capacity models of visual working memory. *Proceedings of the National Academy of Sciences, 105,* 5975–5979.

Sawamura, H., Orban, G. A., & Vogels, R. (2006). Selectivity of neuronal adaptation does not match response selectivity: A single-cell study of the fMRI adaptation paradigm. *Neuron, 49,* 307–318.

Sayres, R., & Grill-Spector, K. (2006). Object-selective cortex exhibits performance-independent repetition suppression. *Journal of Neurophysiology, 95,* 995–1007.

Simons, D. J. (1996). In sight, out of mind: When object representations fail. *Psychological Science, 7,* 301–305.

Simons, D. J. (2000). Attentional capture and inattentional blindness. *Trends in Cognitive Sciences, 4,* 147–155. doi:10.1016/S1364-6613(00)01455-8

Simons, D. J., Franconeri, S. L., & Reimer, R. L. (2000). Change blindness in the absence of a visual disruption. *Perception, 29,* 1143–1154. doi:10.1068/p3104

Tian, X., & Huber, D. E. (2008). Measures of spatial similarity and response magnitude in MEG and scalp EEG. *Brain Topography, 20,* 131–141. doi:10.1007/s10548-007-0040-3

Tian, X., Poeppel, D., & Huber, D. E. (2011). TopoToolbox: Using sensor topography to calculate psychologically meaningful measures from event-related EEG/MEG. *Computational Intelligence and Neuroscience*. doi:10.1155/2011/674605

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, 94,* 719–723.

Urbach, T. P., & Kutas, M. (2002). The intractability of scaling scalp distributions to infer neuroelectric sources. *Psychophysiology, 39,* 791–808.

Van Zandt, T., Colonius, H., & Proctor, R. W. (2000). A comparison of two response time models applied to perceptual matching. *Psychonomic Bulletin & Review, 7,* 208–256. doi:10.3758/BF03212980

Verbruggen, F., & Logan, G. D. (2008). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology. General, 137,* 649–672. doi:10.1037/a0013170

Verbruggen, F., Logan, G. D., Liefooghe, B., & Vandierendonck, A. (2008). Short-term aftereffects of response inhibition: Repetition priming or between-trial control adjustments? *Journal of Experimental Psychology. Human Perception and Performance, 34,* 413–426. doi:10.1037/0096-1523.34.2.413

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision, 4*(12), 11:1120–1135. doi:10.1167/4.12.11

Williams, L. J. (1989). Foveal load affects the functional field of view. *Human Performance, 2,* 1–28. doi:10.1207/s15327043hup0201_1