Repetition and the brain: neural models of stimulus-specific effects

Kalanit Grill-Spector[1,2], Richard Henson[3] and Alex Martin[4]

[1]Department of Psychology, Stanford University, Stanford, CA 94305, USA
[2]Neurosciences Program, Stanford University, Stanford, CA 94305, USA
[3]MRC Cognition and Brain Sciences Unit, Cambridge, UK
[4]Laboratory of Brain and Cognition; National Institute of Mental Health, Bethesda, Maryland 20892-1366, USA

One of the most robust experience-related cortical dynamics is reduced neural activity when stimuli are repeated. This reduction has been linked to performance improvements due to repetition and also used to probe functional characteristics of neural populations. However, the underlying neural mechanisms are as yet unknown. Here, we consider three models that have been proposed to account for repetition-related reductions in neural activity, and evaluate them in terms of their ability to account for the main properties of this phenomenon as measured with single-cell recordings and neuroimaging techniques. We also discuss future directions for distinguishing between these models, which will be important for understanding the neural consequences of repetition and for interpreting repetition-related effects in neuroimaging data.

Introduction

When stimuli are repeated, neural activity is usually reduced. This neural repetition effect has been reported at multiple spatial scales, from the level of individual cortical neurons in monkeys [1–3] to the level of hemodynamic changes (measuring the pooled activation of millions of neurons) in humans using functional magnetic resonance imaging (e.g. fMRI [4–10]). Repetition-related reductions also occur at multiple temporal scales, both in their longevity – from milliseconds [3] to minutes [9] and days [11] – and in the latency of their expression [12,13]. The phenomenon also occurs in multiple brain regions, and across an impressively large number of experimental conditions.

This stimulus-specific reduction in neural activity has been referred to as adaptation [14–16], mnemonic filtering [17], repetition suppression [18], decremental responses [19] and neural priming [20]. We will use 'repetition suppression' (RS) to refer to decreased neural responses following stimulus repetition. As will be apparent below, it is important to distinguish the scale at which RS arises, because the correspondence across scales (e.g. neural firing and hemodynamic responses) might not be simple.

Interest in repetition effects has recently intensified, for two main reasons. First, repetition effects have proved useful for inferring the nature of representations across different stages of a processing stream. This approach has been used behaviorally (e.g. using visual aftereffects to infer the nature of orientation tuning [21] or face representation [22,23]), with single-cell recording [24], and more recently has become popular with fMRI, particularly given claims that it enables improved spatial resolution [7] (Box 1). The second reason for heightened interest is the possibility that RS might be the neural correlate of priming [25–27]. Priming refers to improved processing of a repeated stimulus according to some behavioral measure (e.g. greater accuracy in identifying the stimulus, or faster response times to make a decision about it), and often occurs under the same experimental conditions as RS. Nevertheless, it is important to note that, under certain conditions, priming can be associated with increased activity, rather than reduction (for discussion of repetition enhancement effects and changes in frontal cortices that might contribute to priming, see [27,28]).

The purpose of this review is to consider several kinds of neural models that have been proposed to account for repetition suppression (RS). We focus primarily on studies using visually presented objects and their effects on the ventral object processing stream, to maximize overlap between monkey and human studies. We evaluate the neural models in terms of their ability to account for the main properties of RS as measured with single-cell recordings, fMRI and electroencephalogram/magnetoencephalogram (EEG/MEG), discuss implications of these models for interpreting experimental results, and propose directions for distinguishing between the models.

Repetition suppression as measured with single-cell recording

Stimulus-specific repetition-related reductions in firing rates have been found in physiological recordings of neurons in macaque inferior temporal (IT) cortex [14,15, 17–19] (Figure 1). These repetition effects have been reported for awake behaving animals performing various visual tasks (e.g. match to sample [17], recognition memory [19,29]), as well as in anesthetized animals [30], and occur for both behaviorally relevant and irrelevant

Box 1. fMRI-adaptation: a method to characterize the nature of neural representations

The method of fMRI-adaptation is hypothesized to be able to characterize the functional properties of neural populations at subvoxel resolutions [16]. First, one measures the basic repetition suppression (RS) effect induced by repetitions of identical stimuli. This is done by adapting cortex with a repetitive presentation of specific stimulus or by measuring RS on subsequent repetitions after a single presentation of a stimulus. In general, repetition of an identical stimulus produces an immediate decrease in the fMRI signal (Figure 2). After adaptation, subjects are presented with a stimulus that is varied along one dimension. If the underlying neural representation is insensitive to the change in the stimulus then the fMR signal will be reduced similar to the reduction produced by repetitions of identical stimuli. Alternatively, if the neurons are sensitive to the transformation, the fMR signal will return to the original (non-adapted) level.

fMRI-adaptation has been widely used by researchers to examine selectivity in object-selective cortex to object transformations [7,37, 46], object format [36], perceived shape [79], contour completion [48] and face representation [7,80,81]. It has also been used to probe higher-level conceptual representations using object pictures [82] and words [83]. These studies have documented that fMRI-adaptation in occipitotemporal cortex is not limited to the identical image, but also occurs (albeit to a lesser extent) to transformed versions of the same object, to different exemplars that share the same name (e.g. two different umbrellas), and even to different words that are conceptually related [83]. Finally, fMRI-adaptation has also been used to examine sensitivity to specific types visual information, such as orientation tuning [84,85], color [85,86] and sensitivity to motion information [87–89].

stimuli [2]. RS effects are stimulus specific in the sense that they do not appear to reflect a global reduction in the firing of a population of neurons to all subsequent stimuli [17]. Nonetheless, the precise definition of a 'stimulus' is important, because neural RS can exhibit invariance to some changes in stimulus dimensions (such as the size or position of an object [24]).

The degree of RS depends on several factors, although the precise effect of each factor can vary with brain region. RS persists despite many intervening stimuli, particularly in more anterior regions of IT cortex [2,31]. RS increases with more repetitions of the same stimulus, such that firing rates resemble an exponentially decreasing function of presentation number [1]. This reduction in firing rates occurs primarily for visually excited neurons, [14] and the greatest reduction tends to occur for neurons that were most active on the first presentation [1]. Importantly, several studies indicate that RS onsets rapidly, as fast as 70–80 ms in some perirhinal neurons [31] and with a mean population latency of around 150 ms in IT [15]. Indeed, Xiang and Brown [31] suggested that these effects are too fast for 'top-down' influences.

An important concern for the present discussion is the proportion of neurons showing RS. The best available estimates suggest that significant RS occurs for ∼50–67% of visually responsive neurons in IT cortex [2,15]. There are two possible interpretations of these data. One is that they reflect two different neural populations: One showing RS, and the other showing no reduction in response [19]. Alternatively, there might be a single population, in which all neurons show some degree of RS (with the failure to observe RS in some neurons simply the result of limited

power to detect a significant reduction). However, the exact relationship between RS and the degree of stimulus specificity has not been determined. This is a crucial issue that we will return to later (see section below: 'Distinguishing the neural models').
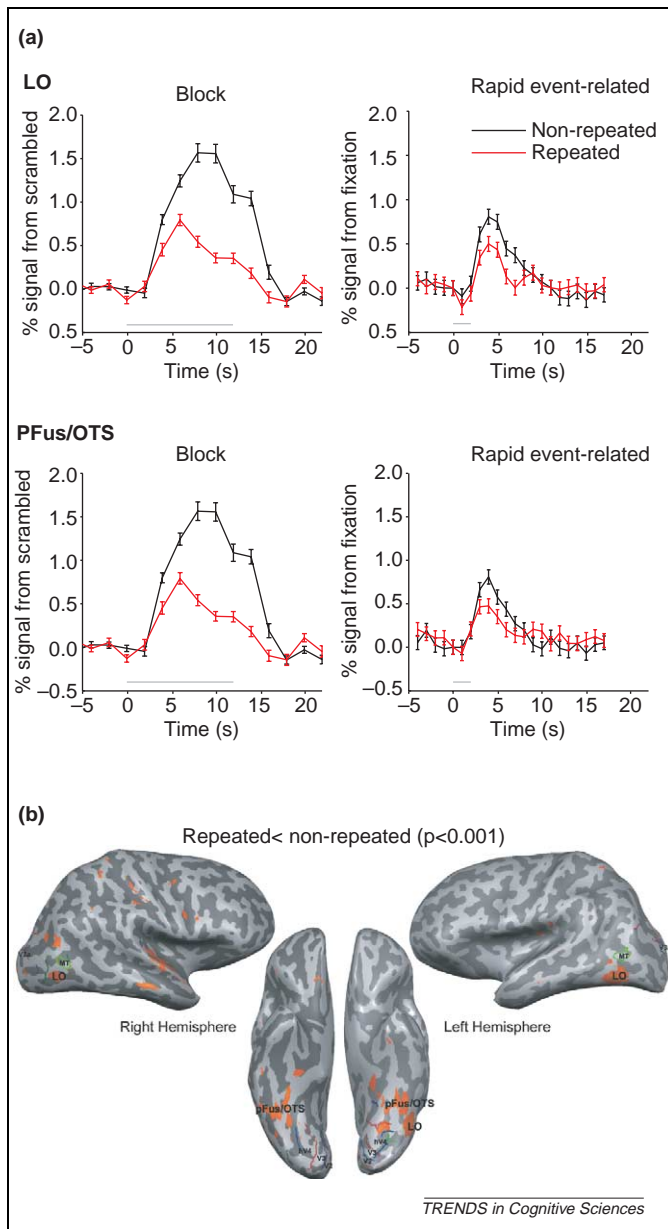
Repetition suppression as measured with fMRI

The basic phenomenon of RS measured with fMRI, also referred to as fMRI-adaptation, has been replicated many times in ventral temporal cortex (see review [32]) as well other areas such as medial temporal [6] and frontal cortex [33–35]. Many adaptation paradigms have been used to measure RS, including multiple repetitions of the same stimulus without intervening items [7] (Figure 2a: block design), or after a single presentation with either no [36] or many [37] intervening items (Figure 2a: rapid event-related). Thus, the properties of RS listed below might vary not only across different brain regions, but also as a function of the paradigm and task [38] (Box 2).

As with single-cell data, fMRI responses tend to decrease monotonically with the number of repetitions [9,16,39], often reaching a plateau after six to eight repetitions [8,39]. RS is maximal when there are no intervening stimuli between repeats [13,39], but has also been observed with tens of intervening stimuli [9] and even after multiple days between presentations [11,35,40]. Other factors that affect the level of RS are stimulus contrast [41], attention [42–44] and the duration of the initial stimulus presentation [45]. RS has also been reported for subliminal stimuli [10]. Some studies have argued that RS occurs across a wide expanse of object-responsive cortex, and not only in regions within that cortex that show maximal responses to particular stimuli [40,41]. However, the relationship between RS and stimulus selectivity has not been tested directly.

The first study to use fMRI-adaptation to probe the sensitivity of object representations to object transformations (size, position, viewpoint, and illumination) was conducted by Grill-Spector and colleagues [7] (Box 1). They found that the posterior section of lateral occipital cortex (LO), which is known to be object selective, showed sensitivity to many object transformations including size and position; however, more anterior object-selective regions were to some extent invariant to size and position, but sensitive to object rotation or illumination. A later
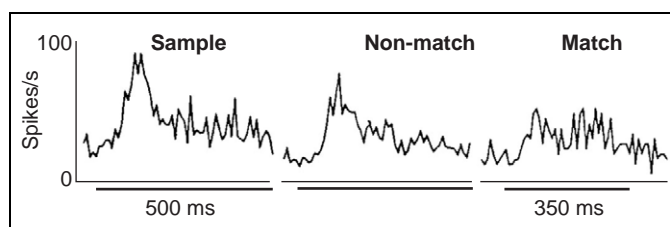
Figure 1. Object repetition effects in a neuron in macaque IT cortex. Averaged responses to stimuli appearing as samples (first presentation), non-matches (first presentation of a new stimulus after a sample) or matches (repeated presentation of the sample stimulus). Responses are averaged across six stimuli, in 10-ms bins. The bar beneath each graph shows stimulus presentation. Y-axis denotes firing rate in spikes per second. (Reproduced with permission from [1].)

Figure 2. Object-repetition effects measured with fMRI: single subject data. **(a)** Time course of fMRI activation in object-selective cortex. Object-selective cortical regions were defined as regions that showed higher activation for animal stimuli than for scrambled animals, with p < 0.001 at the voxel level [94]. Black: first presentation of the stimulus; Red: repeated versions of the same stimulus. Horizontal gray bar indicates stimulus presentation. Error bars indicate the standard error of the mean across trials for this subject. Block design (left graphs): images were presented in blocks of 12 animals (either different animals or all the same). Blanks of 250 ms appeared between consecutive images. 0–2 catch trials of scrambled images were included in each block to control for attention. Rapid event-related design (right graphs): images of animals were presented in rapid sequence, all images repeated 8 times, with the number of intervening stimuli between repeats varied from zero to >16 [39]. **(b)** Activation maps. Colored regions showed significantly lower activation for repeated than for non-repeated images, with p < 0.005 at the voxel level (block design experiment). (Adapted from [7,39].)

study by Vuilleumier and colleagues [37] used an event-related paradigm and reported similar results regarding size invariance, but found differential sensitivity to object rotation across hemispheres: the left hemisphere displayed invariance to object rotation, the right did not. A related study by James *et al.* [46] found higher invariance to object rotation in the ventral than the dorsal stream.

The different levels of rotation invariance found across studies could depend on the objects used, the degree of rotation and the type of paradigm used.

More recently, researchers have used fMRI-adaptation paradigms with monkeys [47–49]. This provides an important link between methodology (single-cell vs. fMRI) and species (monkeys vs. humans). For example, Sawamura *et al.* [49] found comparable size invariance, as inferred via fMRI-adaptation, for objects presented to both monkeys and humans.

Repetition effects as measured with EEG/MEG

Repetition effects have also been studied by measuring changes in the electrical (EEG) or magnetic (MEG) field, usually recorded above the scalp. These effects reflect changes in the amplitude and/or synchrony of local field potentials (LFPs) caused by transmembrane currents in large numbers of neurons. Most EEG studies examine event-related potentials (ERPs), which reflect changes in electrical potential during the few hundred milliseconds following stimulus onset, averaged across trials. The earliest object repetition effects are typically observed ∼200 ms after stimulus onset [13,50–52]. Henson *et al.* [13], found effects of repeating the same view of an object as early as 160–190 ms when there were no intervening objects; with one or more intervening objects, repetition effects only emerged from ∼200 ms. Schendan and Kutas [50] showed that repetition produced an ERP difference between 140 and 250 ms that was greater for objects repeated with the same view than with a different view, whereas later time windows (400–700 ms) showed repetition effects with greater generalization across views. ERP studies using long-lag repetition of faces, including studies that record directly from the cortical surface (in patients undergoing surgery [53]), tend to find that repetition effects emerge only after the first face-specific ERP component, the 'N170' [54,55]. Together, these studies suggest that when there are intervening stimuli between object repetitions, repetition effects on ERPs tend to occur soon after initial stimulus-specific components (>200 ms; Figure 3a), unless there is a change in view, in which case the effects emerge later (>400 ms).

Other EEG studies concentrate on changes in the power of electrical or magnetic oscillations that are induced by stimulus repetition (high-frequency oscillations are not observed in ERPs if they are not phase-locked across trials). Some studies report decreased high frequency (>40 Hz) power around 220–350 ms for repetition of familiar objects across lags of one or two intervening objects [52] (Figure 3b). Such changes in power in certain frequency bands have been shown to correlate with the blood oxygenation-level-dependent (BOLD) changes measured by fMRI [56].

Recent methods for distributed source localization of high-density EEG/MEG data provide estimates of spatio-temporal patterns of activity over the cortex and suggest that prefrontal activity might precede repetition effects in ventral temporal cortex. For example, Dale and colleagues [12] estimated that the onset for repetition effects (for words) recorded by MEG was ∼250 ms in ventral cortex, which was after activity had reached frontal cortex

---

Box 2. Experimental assay of repetition suppression (RS) as measured by BOLD

Several procedural issues need to be kept in mind when comparing RS effects across different studies: the number of repetitions; time between repetitions; the paradigm used (prolonged repetition, i.e. adaptation, versus brief presentations); whether repeated stimuli are compared with initial presentations of the same stimuli or with initial presentations of different stimuli; whether the repetitions are task relevant or irrelevant [27]. RS is a relative measure, so it is important to remember that even the baseline condition (e.g. a series of different, non-repeated objects) might have some degree of RS owing to shared stimulus properties.
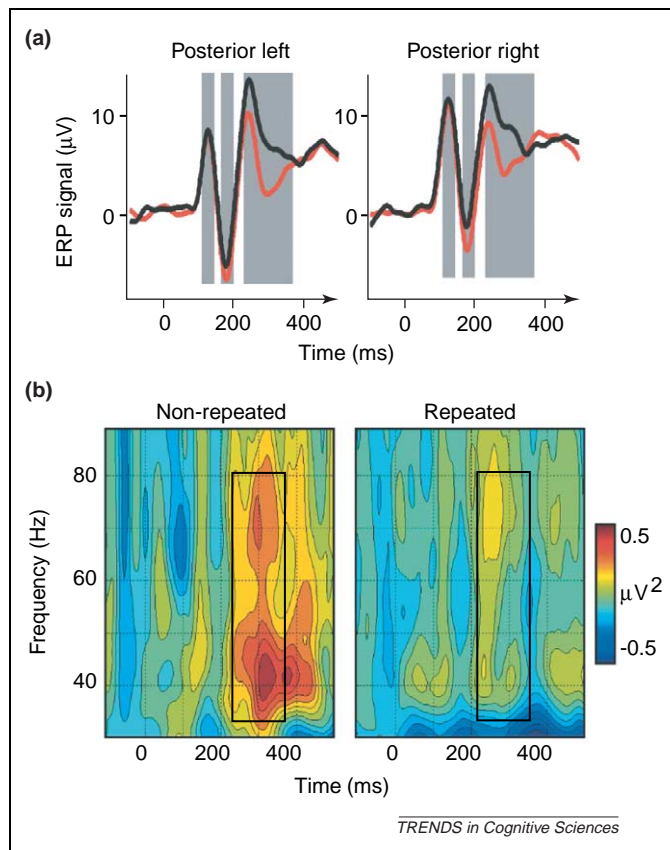
It is also important to distinguish studies that examine RS during intermixed initial and repeated presentations (event-related design) from those that compare the mean activity during a block of stimuli (repeat vs. non-repeat blocks). Block designs have the advantage of prolonged repetition that increases the magnitude of RS. However, additional factors, such as attention and expectation, might also contribute to differences in the BOLD signal. Thus it is important to control these factors across conditions. Subliminal priming paradigms overcome many of these problems [10], although subliminal effects tend to be smaller (and might be qualitatively different [60]).

Other designs compare the responses of two stimuli presented sequentially, in which the pair consists of either two presentations of the same or different stimuli [36]. This approach combines the advantage that short-lag repetition effects tend to be larger with the advantage of a reduced predictability of repetition (if the order of 'same' and 'different' pairs is randomized). However, it has a potential disadvantage that the RS effect might be a small proportion of the total BOLD response to both stimuli within a pair. Further, depending on the interval between stimuli, one might have to consider additional factors such as potential non-linearities in the relationship between neural activity and BOLD response at short interstimulus intervals.

(~185 ms). A recent study that localized broadband (4–50 Hz) MEG power reported increased power for repeated vs. first presentations of objects from 120–230 ms in frontal cortex, followed by decreases from 350–440 ms in

fusiform cortex [57]. These studies indicate multiple repetition effects across both space and time. They also suggest that some (long-lag) repetition effects in ventral temporal cortex might arise relatively late (contrary to the single-cell data), perhaps after interactions with frontal cortex.
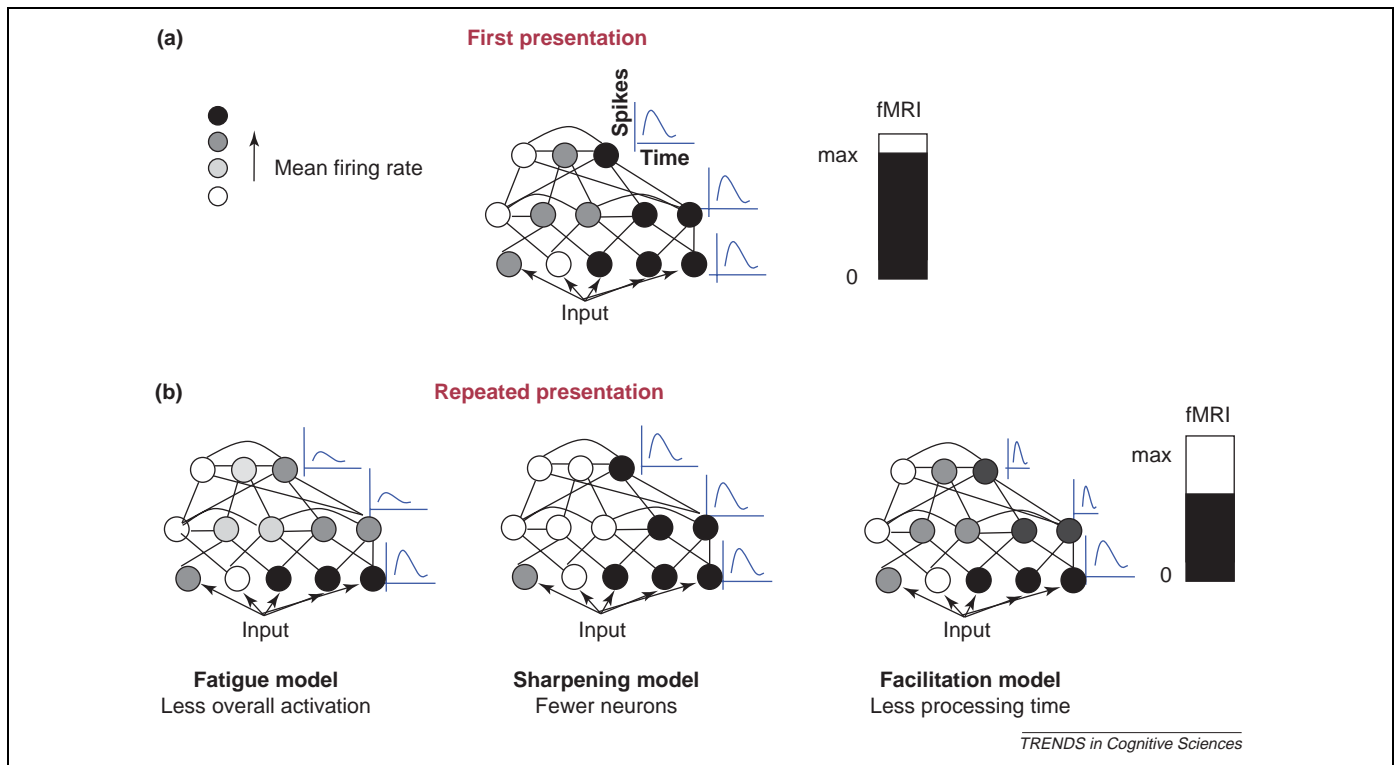
Relating different types of data

Relating the data recorded from single cells by electrodes (firing rates, multi-unit activity: MUA, or LFP) to the data recorded by fMRI (BOLD) and by EEG/MEG (synchronous changes in LFPs over many neurons) is non-trivial. In general terms, increases in MUA correlate with increases in LFP, and increases in LFP correlate with increases in BOLD [58] and presumably EEG/MEG power. However, the special constraints of each technique give rise to situations where their measures might dissociate. For example, single-cell recording is biased towards large excitatory neurons and might not capture changes in smaller cells. Moreover, changes in LFPs from presynaptic activity might not always produce changes in postsynaptic firing (e.g. if excitatory and inhibitory presynaptic activity cancel out). Thus, in principle, presynaptic activity in neurons could increase, causing increases in BOLD and MEG/EEG power, even if the firing rates of neurons in that region did not. Furthermore, changes in the synchrony but not amplitude of LFPs could, in principle, be reflected in EEG/MEG signals, but not BOLD. Conversely, changes in LFPs in neurons that are not aligned might have detectable BOLD signals without detectable EEG/MEG signals. In addition, because the BOLD signal integrates over several seconds, transient effects or small shifts in the latency of neural activity might be difficult to detect with fMRI. Finally, issues concerning experimental design need to be taken into account when comparing across methodologies or even within a methodology (Box 2).

A related issue is that repetition effects in one region might change the balance of inputs this region sends to a second region, thus producing effects beyond the region that showed RS. For example, Tolias and colleagues [59] found that neurons in macaque visual area V4, which are not selective for motion direction, become selective for motion direction after adaptation to motion [47]. They

Figure 3. Object-repetition effects measured with EEG. **(a)** ERPs from posterior left and right electrodes, in response to line drawings of familiar objects, repeated with 1–2 intervening objects. Data averaged across 11 participants. Black: first presentation of the stimulus; Red: repeated versions of the same stimulus. The difference between repeated and non-repeated curves is only reliable by the third time-window (230 ms onwards); the early deflections within the first two time-windows did not show significant repetition effects in this study. **(b)** Amplitude of EEG as a function of time and frequency for first (left plot) and repeated (right plot) presentations, using same paradigm as in (a). Power spectrograms were calculated on the basis of single trials, after subtracting the mean evoked response, and therefore reflect components that are not phase-locked to stimulus onset (see [52] for more details). Note the decrease in amplitude (color change towards green) from 220–350 ms (black box) for repeated relative to first presentations. This decrease occurs in the frequency band above 30 Hz (gamma band). (Data reproduced with permission from [52].)

Figure 4. Models for RS. To distinguish between our models we will use a schematic neural network consisting of three putative cortical processing stages. **(a)** The visual stimulus is assumed to cause activity in the input layer (corresponding to early visual cortex) before being processed in a hierarchical sequence of stages. The blue graphs indicate spiking (as a function of time) of the neurons with highest response at each stage (indicated by black circles). **(b)** Because the BOLD signal integrates neuronal activity over time, all three of these models predict reduced BOLD for repeated stimuli, but for different reasons: Fatigue model (left, lower firing rates); Sharpening model (centre, fewer neurons responding); Facilitation model (right, shorter duration of neural processing).

suggest that V4 neurons normally receive balanced inputs from motion-selective neurons in the middle temporal area, MT, but that after adaptation the input from MT becomes unbalanced, resulting in motion-selective responses in V4. This finding is relevant for the use of RS to infer the nature of representations in different brain regions (Box 1).

Models for repetition suppression

There are multiple potential neural causes of the RS measured with single-cell recording, fMRI or EEG/MEG. We now describe three models that have been sketched previously in the literature (Figure 4): (i) the *Fatigue model*, whereby the amplitude of firing of stimulus-

responsive neurons decreases [2,16]; (ii) the *Sharpening model*, whereby fewer neurons respond [1,18,26]; and (iii) the *Facilitation model,* whereby the latency [60] and/or duration of neural activity is shortened [3,27]. An important consideration to keep in mind is how each of these models increases processing efficiency.

Fatigue model
According to this model, all neurons initially responsive to a stimulus show a proportionally equivalent reduction in their response to repeated presentations of the same stimulus. As a consequence the mean population firing rate declines but there are no changes in the pattern of relative responses across neurons, or in the temporal

---

Box 3. Potential neural mechanisms underlying repetition suppression (RS)

At least four mechanisms might play a role in RS (i.e. the reduction in spiking rates following repetition) at the neural level:

(1) Firing-rate adaptation. This reflects reduced excitability of neurons owing to an increase in potassium ion currents that hyperpolarize the membrane potential and increase its conductance [61]. This in turn reduces the impact of synaptic inputs and hence lowers the probability of neural spiking. This effect tends to be short lived, decaying over hundreds of milliseconds, although possibly lasting tens of seconds (when potassium currents are activated by sodium ions [90]).

(2) Synaptic depression. This is a temporary reduction in synaptic efficacy that is believed to reflect reductions in presynaptic neuro-transmitter release [91]. The effect depends on pre- (rather than post-) synaptic activity and can last tens of seconds for excitatory synapses (it tends to be shorter-lived for inhibitory synapses).

(3) Long-term depression (LTD). This longer-lasting mechanism reflects plasticity changes owing to correlated pre- and postsynaptic activity and normally entails multiple stages, including decreases in calcium concentration, perhaps even gene expression and ultimately protein synthesis. LTD reflects a reduction in synaptic efficacy [92], normally follows low-frequency stimulation and can last at least hours.

(4) Long-term potentiation (LTP). Like LTD, this long-lasting mechanism entails multiple stages and might come in several forms (depending on the precise receptor type). LTP reflects an increase in synaptic efficacy. Both LTP and LTD have been best studied for excitatory (glutamatergic) synapses, for which they tend to increase or decrease, respectively, the probability of spiking following presynaptic input. Although controversial, similar mechanisms might also operate on inhibitory synapses (from interneurons [93]), in turn causing either decreases or increases in spiking probability respectively.
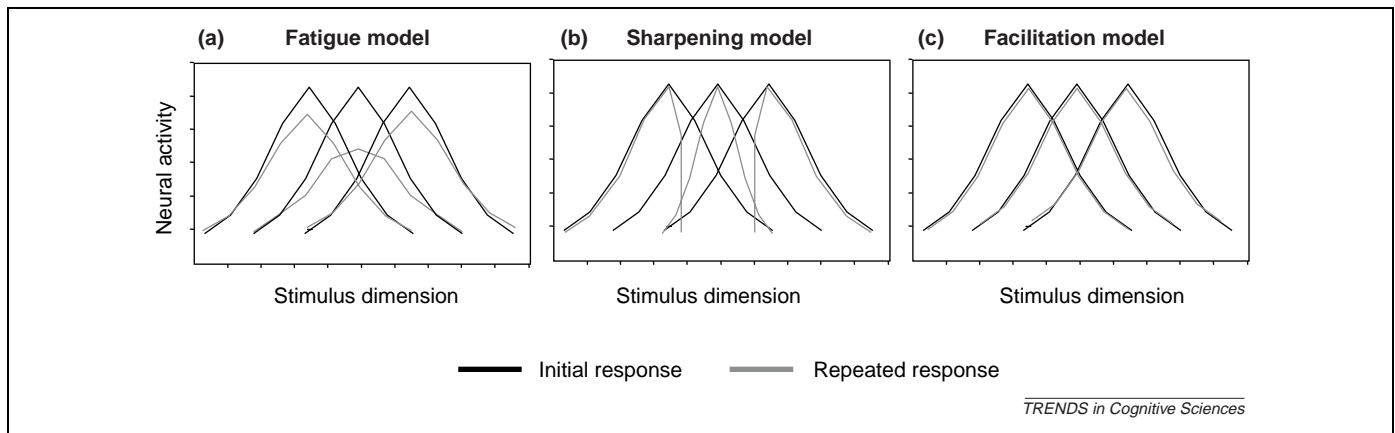
Figure 5. Implications of models for neural tuning following repetition. Tuning curves before (black) and after (gray) repetition. The central stimulus in each case was repeated. **(a)** Fatigue model suggests that repetition will reduce overall response in proportion to initial response, but the preferred stimulus and tuning width will not change. **(b)** Sharpening model predicts that repetition will narrow the tuning bandwidth, that is, sharpen the response around the preferred stimulus. **(c)** Facilitation model does not make strong predictions for any particular change in tuning curves (although in some kinds of attractor networks, the tuning curves might even widen).

window in which neurons are responding. One mechanism for fatigue could be firing-rate adaptation (Box 3), in which the reduction in a neuron's firing rate is proportional to its initial response [1,41] (similar to a gain mechanism [61]). However, this mechanism does not explain the specificity of RS, that is, why the neuron's response is reduced to some stimuli, yet resumes high firing rates to other stimuli. An alternative mechanism is reduced synaptic efficacy of specific synapses from connected neurons (synaptic depression; Box 3). In this manner, only specific patterns of presynaptic input to the neuron (which are stimulus dependent) reduce its firing rate. This type of mechanism has been implicated in early visual cortex and usually occurs with prolonged repetitive stimulation.

One prediction from this model is that the amount of RS will be greater in neurons that respond optimally to a stimulus than in other neurons (Figure 5a). As a result, the sensitivity of the system to stimuli that are different from the repeating stimulus is increased, thereby providing a mechanism for 'novelty detection'. Reducing the firing rate might also help prevent saturation of the neural response function by increasing its dynamic range. Another advantage hypothesized for such a mechanism is that it reduces redundancies in the neural code, which increases the efficiency of information encoding [62].

However, it is not immediately clear how reduced firing rates can account for increased speed and accuracy of processing repeated stimuli (in addition to increased sensitivity to novel stimuli), as arises in priming. One explanation is provided in a computational model by Gotts (S. Gotts, PhD thesis, Carnegie Mellon University, USA, 2003), in which a reduction in the mean and variance of firing rates allows for greater synchrony of neural responses. Greater synchrony of presynaptic input is believed to be more effective in triggering a postsynaptic response [63], which would imply more rapid transmission of information throughout the network, resulting in faster responses (priming). A key prediction of this model is that synchrony can increase while stimulus-specific firing rates decrease. Although there is evidence in support of this possibility [64], others have argued for the opposite effect [65]. An increase in synchrony might also be difficult

to reconcile with observations of reduced oscillatory power following stimulus repetition [52], although it is possible that decreased amplitude of local field potentials out-weighs the increased synchrony of those potentials.

Sharpening model

Desimone and colleagues [18] and Wiggs and Martin [26] have suggested that repetition results in sparser representation of stimuli. According to this model, some – but not all – neurons that initially responded to a stimulus will show RS to subsequent presentation of that stimulus. Importantly, it is the neurons that code features irrelevant to identification of a stimulus that exhibit RS (Figures 4,5b). Thus, repetition-related changes are viewed primarily as a learning process, in which representations (e.g. tuning curves, Figure 5b) are 'sharpened' and, as a consequence, the distributed representation becomes sparser, resulting in fewer responsive neurons in total. An important difference between the Sharpening and Fatigue models is that, for Sharpening, many of the neurons that are optimally tuned to the repeating stimulus should show little or no response reductions, rather than exhibit the greatest response reduction, as in the Fatigue model.

Sparser representations clearly have adaptive value in terms of a reduced metabolic cost. Also, because the representation becomes sharper (tuning curves become narrower, Figure 5b), the neurons become more sensitive to change. Sparser representations might also allow for more efficient or faster processing, although this depends on the manner in which their information is read-out by downstream neurons. Because the Sharpening model suggests a changed and improved representation for repeated stimuli, this model has been widely used to explain priming [26,27,45]. However, a recent study suggests that some RS in object-selective cortex might reflect response learning and implies that object representations do not necessarily reorganize as a consequence of repetition [66].

The mechanism underlying the formation of sparser representations is unknown but could reflect inhibition from lateral connections between neurons within

a population. For example, Norman and O'Reilly [67] used a competitive Hebbian learning rule to simulate the sharpening of representations with repetition (within medial temporal cortex). Initially, many neurons respond weakly to a distributed input pattern representing the stimulus. Through competitive interactions, the neurons with the strongest initial response become 'stronger', and inhibit the 'weaker' neurons. Thus some neurons show increased firing rates following repetition, whereas others show decreased firing rates. By assuming that the number of 'strong' units is less than the number of 'weak' units, the population response decreases with repetition because there are more neurons showing reduced activity than showing increased activity. If information from only those neurons with high firing rates is 'read out' by downstream neurons, their increased firing rate following repetition (despite the global decrease) could explain the faster processing of repeated stimuli.

Facilitation model
At its simplest, this model predicts that repetition causes faster processing of stimuli, that is, shorter latencies or shorter durations of neural firing. One example is the 'accumulation' model of James *et al.* [60], in which stimulus information is accrued faster following repetition. Given that the hemodynamic signal measured by fMRI integrates over several seconds of neural activity, a shorter duration of activity results in a decreased amplitude of the fMRI signal. A shorter duration of neural activity is also consistent with earlier peaking of the fMRI response [68], and might explain why decreases in firing rate can appear to arise after the initial visual response [15]: the neurons initially fire robustly to both first and repeated presentations, but this firing stops sooner for repeated presentations.

An extension of the Facilitation model assumes that the cause of this faster processing is synaptic potentiation between neurons following an initial stimulus presentation, and that this potentiation can occur at many levels in the processing stream. As a consequence, information flows through the stream more rapidly, and hence identification of a repeated stimulus occurs faster. In terms of attractor neural network models, synaptic potentiation might be viewed as deepening the basin of attraction, resulting in a shorter time for the network to settle on a stable pattern corresponding to identification of the stimulus. An example of such a dynamical, network model is sketched by Friston [69]. The key idea behind this model is that the firing rate of the long-range excitatory (output) neurons in a population codes 'prediction error' [70], which is the difference between bottom-up input ('evidence') and top-down input ('prediction'). The dynamics of the network are such that prediction error decreases over time after stimulus onset, and synaptic changes serve to accelerate this decrease when the stimulus is repeated (i.e. repetition improves prediction).

This emphasis on recurrent activity between many levels of the processing stream is consistent with the spatiotemporal pattern of repetition effects emerging from MEG/EEG data. Moreover, if interregional interactions require an initial volley of activity through the network

[71], this model could further explain the relatively late onset of long-lag repetition effects recorded with EEG/MEG. However, such a model would not explain why much earlier repetition effects have been observed in some neurons (e.g. 75–100 ms, which is thought to be too early for feedback [31]); further, this model does not necessarily predict decreases in the peak firing rate of individual neurons.

Each of the above models would clearly benefit from further elaboration, including instantiation as detailed computational models. It is possible that different models might apply in different brain regions and under different experimental conditions (e.g. different paradigms/tasks). We next consider some experiments that might distinguish between the models.

Distinguishing the neural models
There are three main directions in which these models can be distinguished: (i) examining the relationship between RS and stimulus selectivity; (ii) examining the effect of repetition on the tuning of cortical responses along a stimulus dimension; and (iii) examining the temporal window in which processing occurs for new and repeated stimuli.

Examining the relationship between RS and stimulus selectivity
The models differ in their predictions on whether RS is strongest for the preferred stimulus, or for non-preferred stimuli. The Sharpening model predicts that neurons showing little or no RS to a repeated stimulus are highly selective for that stimulus. By contrast, both the Fatigue and Facilitation models predict that RS is proportional to the initial response. Thus, neurons that respond optimally for a stimulus should show the largest suppression. These hypotheses can be tested with single-cell recording.

Examining the effect of repetition on neural tuning
Another way to distinguish the models would be to find a single dimension (e.g. motion, orientation), along which stimuli differ, and examine the effect of repetition on the tuning curves of different neurons along that dimension. The models differ in their prediction of how repetition will change the shape of neuronal tuning. According to the Fatigue model, repetition reduces the response in proportion to the initial response, but the tuning width does not change (Figure 5a). Most probably, the reduction will be maximal for tuning curves centered on the location of the repeating stimulus along the stimulus dimension, and lesser for tuning curves centered further away. This is consistent with physiological data in V1 for orientation [72], spatial frequency [73] and motion direction [74]. By contrast, according to the Sharpening model, repetition 'sharpens' tuning curves (Figure 5b). This is consistent with studies of learning-related changes in IT cortex and V4 [75–77]. Finally, the Facilitation model does not suggest any particular effect on tuning curves (Figure 5c). Indeed, even a widening of the curves might be possible if repetition enlarged the attractor basin in an attractor network model.

An elegant study by Kohn and Movshon [74] demonstrates how measuring tuning to a stimulus dimension before and after repetition can be used to distinguish between models. Kohn and Movshon examined the effect of repetition (using prolonged adaptation) to the preferred, non-preferred and null motion direction on tuning curves of MT neurons. Their data reveal that the tuning of MT neurons to motion direction narrows (sharpens) around the repeating stimulus when adapting a cell with its preferred motion direction and, further, that tuning curves shift towards the repeating stimulus when adapting the neurons with a non-preferred stimulus. However, sharpening of tuning curves did not occur in V1; rather V1 data were consistent with the Fatigue model. These findings provide evidence that stimulus repetition can have different neural consequences depending on brain region. Thus, different models of RS could apply for different cortical sites.

Examining the temporal window of processing for new and repeated stimuli

The models can also be distinguished in the temporal domain. In particular, the Facilitation model suggests that the latency and/or duration of the response to repeated items will be shorter than to first presentations. The Fatigue and Sharpening models do not suggest a difference in the temporal processing window for repeated stimuli. The latency and duration of processing might be examined via single-cell recordings and/or EEG/MEG techniques.

Another prediction from one version of the Facilitation model is that facilitation will be more pronounced after top-down information is processed. As a result, the largest difference between repeated and non-repeated stimuli should occur after initial sensory processing (e.g. after the first 200 ms). This is consistent with the observation that long-lag repetition effects typically emerge after initial stimulus-specific ERP/ERF components. Unlike the Facilitation model, the Fatigue model suggests that RS occurs during the entire processing window. This predicts that the earliest selective responses are modulated by repetition. For the Sharpening model, the onset of repetition effects might depend on whether sharpening is implemented in terms of changes in feedforward (input) synapses, or via changes from lateral connections within a region. In the latter case, RS might be delayed until selective neurons have inhibited non-selective neurons, which would be consistent with RS occurring subsequent to the initial stimulus-specific response.

We would like to emphasize that any empirical data relevant to the above predictions are likely to depend on other factors, such as the lag between repetitions. One of the central outstanding questions is whether different models apply at different time scales. One possibility is that the mechanisms related to the Fatigue model operate during immediate repetitions of a stimulus within a few hundred milliseconds, and reflect transient stimulus-specific effects that onset rapidly. By contrast, the effects of repetition across many intervening stimuli might be more consistent with the Sharpening or Facilitation models, and reflect long-term learning that leads to

changes in the spatial pattern of stimulus-selective responses and/or dynamics of those responses.

Conclusions

Adaptation paradigms have become increasingly popular for examining the processing characteristics of different cortical regions of the brain. Clearly, progress has been made using RS to infer the nature of representations in different cortical regions, or as a marker for increased processing efficiency, without a complete understanding of its neural basis. Nonetheless, it is our belief that the specific neural mechanisms matter, because interpretation and design of experiments depend on the nature of the underlying neural model. For example, models differ as to whether RS reflects quantitative or qualitative changes in representations. One important possibility we raised is that of multiple models – which vary in their relevance across space, time and task – that might parallel the multiplicity of potential neural/synaptic mechanisms (Box 3). Although it is possible that a single model could apply under all conditions, we think it is likely that different models are needed to explain RS in different conditions: Some effects might be short-lived, others might be permanent. Further, the magnitude of RS might vary as a function of task and stimuli. Finally, it is not yet known whether the same or different mechanisms operate in different brain regions.

Clearly, many questions remain regarding the neural basis of RS (Box 4). Progress will be aided primarily by a better understanding of the effects of stimulus repetition on single neurons. This will include concurrent recording from multiple electrodes to measure firing rates and their correlations between neurons. Progress can also be made using more global measures in humans, such as fMRI and EEG/MEG, provided important differences between these measurements are kept in mind. Improvements in the spatial resolution of fMRI [78] are likely to be important for distinguishing between models. The use of similar stimuli and paradigms in physiological experiments and fMRI/EEG/MEG experiments will be key for comparisons across different types of data. Most informative will be concurrent recording of repetition effects using all these methods. These future experiments will yield important empirical data that will validate (or refute) the theoretical predictions laid out in this review.

Box 4. Questions for future research

• What is the relationship between repetition suppression (RS) and stimulus selectivity?
• How does repetition affect the tuning curves of neurons?
• How do the Fatigue and Sharpening models account for improved behavioral performance (priming)?
• In what time scale does facilitation occur?
• Does facilitation require interactions between different brain regions, in particular feedback?
• Does RS induced by prolonged adaptation and RS induced by one (or several) presentations with intervening stimuli involve similar or distinct neural mechanisms?
• How does RS in lower level regions (e.g. V1) differ from RS in high-level regions (e.g. MT or LO)?

Acknowledgements

References

1 Li, L. *et al*. (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J. Neurophysiol.* 69, 1918–1929

2 Miller, E.K. and Desimone, R. (1994) Parallel neuronal mechanisms for short-term memory. *Science* 263, 520–522

3 Sobotka, S. and Ringo, J.L. (1996) Mnemonic responses of single units recorded from monkey inferotemporal cortex, accessed via transcommissural versus direct pathways: a dissociation between unit activity and behavior. *J. Neurosci.* 16, 4222–4230

4 Demb, J.B. *et al*. (1995) Semantic encoding and retrieval in the left inferior prefrontal cortex: a functional MRI study of task difficulty and process specificity. *J. Neurosci.* 15, 5870–5878

5 Buckner, R.L. *et al*. (1995) Functional anatomical studies of explicit and implicit memory retrieval tasks. *J. Neurosci.* 15, 12–29

6 Stern, C.E. *et al*. (1996) The hippocampal formation participates in novel picture encoding: evidence from functional magnetic resonance imaging. *Proc. Natl. Acad. Sci. U. S. A.* 93, 8660–8665

7 Grill-Spector, K. *et al*. (1999) Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* 24, 187–203

8 Jiang, Y. *et al*. (2000) Complementary neural mechanisms for tracking items in human working memory. *Science* 287, 643–646

9 Henson, R. *et al*. (2000) Neuroimaging evidence for dissociable forms of repetition priming. *Science* 287, 1269–1272

10 Naccache, L. and Dehaene, S. (2001) The priming method: imaging unconscious repetition priming reveals an abstract representation of number in the parietal lobes. *Cereb. Cortex* 11, 966–974

11 van Turennout, M. *et al*. (2000) Long-lasting cortical plasticity in the object naming system. *Nat. Neurosci.* 3, 1329–1334

12 Dale, A.M. *et al*. (2000) Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron* 26, 55–67

13 Henson, R.N. *et al*. (2004) The effect of repetition lag on electrophysiological and haemodynamic correlates of visual object priming. *Neuroimage* 21, 1674–1689

14 Sobotka, S. and Ringo, J.L. (1994) Stimulus specific adaptation in excited but not in inhibited cells in inferotemporal cortex of macaque. *Brain Res.* 646, 95–99

15 Ringo, J.L. (1996) Stimulus specific adaptation in inferior temporal and medial temporal cortex of the monkey. *Behav. Brain Res.* 76, 191–197

16 Grill-Spector, K. and Malach, R. (2001) fMR-adaptation: a tool for studying the functional properties of human cortical neurons. *Acta Psychol. (Amst.)* 107, 293–321

17 Miller, E.K. *et al*. (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J. Neurosci.* 13, 1460–1478

18 Desimone, R. (1996) Neural mechanisms for visual memory and their role in attention. *Proc. Natl. Acad. Sci. U. S. A.* 93, 13494–13499

19 Brown, M.W. and Xiang, J.Z. (1998) Recognition memory: neuronal substrates of the judgement of prior occurrence. *Prog. Neurobiol.* 55, 149–189

20 Maccotta, L. and Buckner, R.L. (2004) Evidence for neural effects of repetition that directly correlate with behavioral priming. *J. Cogn. Neurosci.* 16, 1625–1632

21 Graham, N. (1989) *Visual Pattern Analyzers*, Oxford University Press

22 Webster, M.A. *et al*. (2004) Adaptation to natural facial categories. *Nature* 428, 557–561

23 Leopold, D.A. *et al*. (2001) Prototype-referenced shape encoding revealed by high-level aftereffects. *Nat. Neurosci.* 4, 89–94

24 Lueschow, A. *et al*. (1994) Inferior temporal mechanisms for invariant object recognition. *Cereb. Cortex* 4, 523–531

25 Schacter, D.L. and Buckner, R.L. (1998) Priming and the brain. *Neuron* 20, 185–195

26 Wiggs, C.L. and Martin, A. (1998) Properties and mechanisms of perceptual priming. *Curr. Opin. Neurobiol.* 8, 227–233

27 Henson, R.N. and Rugg, M.D. (2003) Neural response suppression, haemodynamic repetition effects, and behavioural priming. *Neuropsychologia* 41, 263–270

28 Schacter, D.L. *et al*. (2004) Specificity of priming: a cognitive neuroscience perspective. *Nat. Rev. Neurosci.* 5, 853–862

29 Sobotka, S. and Ringo, J.L. (1993) Investigation of long-term recognition and association memory in unit responses from inferotemporal cortex. *Exp. Brain Res.* 96, 28–38

30 Miller, E.K. and Desimone, R. (1993) Scopolamine affects short-term memory but not inferior temporal neurons. *Neuroreport* 4, 81–84

31 Xiang, J.Z. and Brown, M.W. (1998) Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* 37, 657–676

32 Kourtzi, Z. and Grill-Spector, K. (2005) fMRI adaptation: a tool for studying visual representations in the primate brain. In *Fitting the Mind to the World: Adaptation and After-Effects in High-Level Vision* (Rhodes, G. and Clifford, C.W.G., eds), pp. 173–188, Oxford University Press

33 Wagner, A.D. *et al*. (1997) Dissociations between familiarity processes in explicit recognition and implicit perceptual memory. *J. Exp. Psychol. Learn. Mem. Cogn.* 23, 305–323

34 Buckner, R.L. and Koutstaal, W. (1998) Functional neuroimaging studies of encoding, priming, and explicit memory retrieval. *Proc. Natl. Acad. Sci. U. S. A.* 95, 891–898

35 van Turennout, M. *et al*. (2003) Modulation of neural activity during object naming: effects of time and practice. *Cereb. Cortex* 13, 381–391

36 Kourtzi, Z. and Kanwisher, N. (2000) Cortical regions involved in perceiving object shape. *J. Neurosci.* 20, 3310–3318

37 Vuilleumier, P. *et al*. (2002) Multiple levels of visual object constancy revealed by event-related fMRI of repetition priming. *Nat. Neurosci.* 5, 491–499

38 Henson, R.N. *et al*. (2002) Face repetition effects in implicit and explicit memory tests as measured by fMRI. *Cereb. Cortex* 12, 178–186

39 Sayres, R.A., and Grill-Spector, K. (2003) Object-selective cortex exhibits performance independent repetition-suppression. *J. Neurophysiol.* (Epub ahead of print, 19 Oct)

40 Chao, L.L. *et al*. (2002) Experience-dependent modulation of category-related cortical activity. *Cereb. Cortex* 12, 545–551

41 Avidan, G. *et al*. (2002) Analysis of the neuronal selectivity underlying low fMRI signals. *Curr. Biol.* 12, 964–972

42 Murray, S.O. and Wojciulik, E. (2004) Attention increases neural selectivity in the human lateral occipital complex. *Nat. Neurosci.* 7, 70–74

43 Eger, E. *et al*. (2004) BOLD repetition decreases in object-responsive ventral visual areas depend on spatial attention. *J. Neurophysiol.* 92, 1241–1247

44 Yi, D.J. and Chun, M.M. (2005) Attentional modulation of learning-related repetition attenuation effects in human parahippocampal cortex. *J. Neurosci.* 25, 3593–3600

45 Zago, L. *et al*. (2005) The rise and fall of priming: how visual exposure shapes cortical representations of objects. *Cereb. Cortex* 15, 1655–1666

46 James, T.W. *et al*. (2002) Differential effects of viewpoint on object-driven activation in dorsal and ventral streams. *Neuron* 35, 793–801

47 Tolias, A.S. *et al*. (2001) Motion processing in the macaque: revisited with functional magnetic resonance imaging. *J. Neurosci.* 21, 8594–8601

48 Kourtzi, Z. *et al*. (2003) Integration of local features into global shapes: monkey and human fMRI studies. *Neuron* 37, 333–346

49 Sawamura, H. *et al*. (2005) Using functional magnetic resonance imaging to assess adaptation and size invariance of shape processing by humans and monkeys. *J. Neurosci.* 25, 4294–4306

50 Schendan, H.E. and Kutas, M. (2003) Time course of processes and representations supporting visual object identification and memory. *J. Cogn. Neurosci.* 15, 111–135

51 Doniger, G.M. *et al*. (2001) Visual perceptual learning in human object recognition areas: a repetition priming study using high-density electrical mapping. *Neuroimage* 13, 305–313

52 Gruber, T. and Muller, M.M. (2005) Oscillatory brain activity dissociates between associative stimulus content in a repetition priming task in the human EEG. *Cereb. Cortex* 15, 109–116

53 Puce, A. *et al*. (1999) Electrophysiological studies of human face perception: III. Effects of top-down processing on face-specific potentials. *Cereb. Cortex* 9, 445–458

54 Schweinberger, S.R. *et al*. (2004) N250r: a face-selective brain response to stimulus repetitions. *Neuroreport* 15, 1501–1505

55 Eimer, M. (2000) Event-related brain potentials distinguish processing stages involved in face perception and recognition. *Clin. Neurophysiol.* 111, 694–705

56 Brookes, M.J. *et al*. (2005) GLM-beamformer method demonstrates stationary field, alpha ERD and gamma. *Neuroimage* 26, 302–308

57 Gilbert, J. *et al*. (2005) Top-down modulation during object priming: evidence from MEG. *J Cogn Neurosci* (Suppl.) p. 132

58 Logothetis, N.K. *et al*. (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412, 150–157

59 Tolias, A.S. *et al*. (2005) Neurons in macaque area V4 acquire directional tuning after adaptation to motion stimuli. *Nat. Neurosci.* 8, 591–593

60 James, T.W. and Gauthier, I. Repetition-induced changes in BOLD response reflect accumulation of neural activity. *Hum. Brain Mapp.* (in press)

61 Carandini, M. and Ferster, D. (1997) A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science* 276, 949–952

62 Muller, J.R. *et al*. (1999) Rapid adaptation in visual cortex to the structure of images. *Science* 285, 1405–1408

63 Fries, P. *et al*. (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563

64 Stopfer, M. and Laurent, G. (1999) Short-term memory in olfactory network dynamics. *Nature* 402, 664–668

65 Chawla, D. *et al*. (1999) The relationship between synchronization among neuronal populations and their mean activity levels. *Neural Comput.* 11, 1389–1411

66 Dobbins, I.G. *et al*. (2004) Cortical activity reductions during repetition priming can result from rapid response learning. *Nature* 428, 316–319

67 Norman, K.A. and O'Reilly, R.C. (2003) Modeling hippocampal and neocortical contributions to recognition memory: a complementary-learning-systems approach. *Psychol. Rev.* 110, 611–646

68 Henson, R.N. *et al*. (2002) Detecting latency differences in event-related BOLD responses: application to words versus nonwords and initial versus repeated face presentations. *Neuroimage* 15, 83–97

69 Friston, K. (2005) A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 360, 815–836

70 Rao, R.P. and Ballard, D.H. (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87

71 Sugase, Y. *et al*. (1999) Global and fine information coded by single neurons in the temporal visual cortex. *Nature* 400, 869–873

72 Dragoi, V. *et al*. (2002) Dynamics of neuronal sensitivity in visual cortex and local feature discrimination. *Nat. Neurosci.* 5, 883–891

73 Movshon, J.A. and Lennie, P. (1979) Pattern-selective adaptation in visual cortical neurones. *Nature* 278, 850–852

74 Kohn, A. and Movshon, J.A. (2004) Adaptation changes the direction tuning of macaque MT neurons. *Nat. Neurosci.* 7, 764–772

75 Baker, C.I. *et al*. (2002) Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. *Nat. Neurosci.* 5, 1210–1216

76 Sigala, N. and Logothetis, N.K. (2002) Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 415, 318–320

77 Rainer, G. *et al*. (2004) The effect of learning on the function of monkey extrastriate visual cortex. *PLoS Biol.* 2, E44

78 Beauchamp, M.S. *et al*. (2004) Unraveling multisensory integration: patchy organization within human STS multisensory cortex. *Nat. Neurosci.* 7, 1190–1192

79 Kourtzi, Z. and Kanwisher, N. (2001) Representation of perceived object shape by the human lateral occipital complex. *Science* 293, 1506–1509

80 Winston, J.S. *et al*. (2004) fMRI-adaptation reveals dissociable neural representations of identity and expression in face perception. *J. Neurophysiol.* 92, 1830–1839

81 Rotshtein, P. *et al*. (2005) Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nat. Neurosci.* 8, 107–113

82 Koutstaal, W. *et al*. (2001) Perceptual specificity in visual object priming: functional magnetic resonance imaging evidence for a laterality difference in fusiform cortex. *Neuropsychologia* 39, 184–199

83 Wheatley, T. *et al*. Automatic priming of semantically related words reduces activity in the fusiform gyrus. *J. Cogn. Neurosci.* (in press)

84 Boynton, G.M. and Finney, E.M. (2003) Orientation-specific adaptation in human visual cortex. *J. Neurosci.* 23, 8781–8787

85 Engel, S.A. (2005) Adaptation of oriented and unoriented color-selective neurons in human visual areas. *Neuron* 45, 613–623

86 Wade, A.R. and Wandell, B.A. (2002) Chromatic light adaptation measured using functional magnetic resonance imaging. *J. Neurosci.* 22, 8148–8157

87 Huk, A.C. and Heeger, D.J. (2002) Pattern-motion responses in human visual cortex. *Nat. Neurosci.* 5, 72–75

88 Culham, J.C. *et al*. (2000) Independent aftereffects of attention and motion. *Neuron* 28, 607–615

89 Tootell, R.B. *et al*. (1995) Visual motion aftereffect in human cortical area MT revealed by functional magnetic resonance imaging. *Nature* 375, 139–141

90 Schwindt, P.C. *et al*. (1989) Long-lasting reduction of excitability by a sodium-dependent potassium current in cat neocortical neurons. *J. Neurophysiol.* 61, 233–244

91 Markram, H. and Tsodyks, M. (1996) Redistribution of synaptic efficacy between neocortical pyramidal neurons. *Nature* 382, 807–810

92 Ito, M. (1989) Long-term depression. *Annu. Rev. Neurosci.* 12, 85–102

93 Laezza, F. *et al*. (1999) Long-term depression in hippocampal interneurons: joint requirement for pre- and postsynaptic events. *Science* 285, 1411–1414

94 Grill-Spector, K. (2003) The neural basis of object perception. *Curr. Opin. Neurobiol.* 13, 159–166