Neuron
Article

Effects of Long-Term Visual Experience on Responses of Distinct Classes of Single Units in Inferior Temporal Cortex

Luke Woloszyn[1,*] and David L. Sheinberg[1,*]
[1]Department of Neuroscience, Brown University, Providence, RI 02912, USA
*Correspondence: luke.woloszyn@gmail.com (L.W.), david_sheinberg@brown.edu (D.L.S.)
DOI 10.1016/j.neuron.2012.01.032

SUMMARY

Primates can learn to recognize a virtually limitless number of visual objects. A candidate neural substrate for this adult plasticity is the inferior temporal cortex (ITC). Using a large stimulus set, we explored the impact that long-term experience has on the response properties of two classes of neurons in ITC: broad-spiking (putative excitatory) cells and narrow-spiking (putative inhibitory) cells. We found that experience increased maximum responses of putative excitatory neurons but had the opposite effect on maximum responses of putative inhibitory neurons, an observation that helps to reconcile contradictory reports regarding the presence and direction of this effect. In addition, we found that experience reduced the average stimulus-evoked response in both cell classes, but this decrease was much more pronounced in putative inhibitory units. This latter finding supports a potentially critical role of inhibitory neurons in detecting and initiating the cascade of events underlying adult neural plasticity in ITC.

INTRODUCTION

Visual perception is a consequence of the concerted activity of neurons throughout the visual system. At the same time, the response properties of single neurons in the visual system depend on visual experience for their proper development (Hubel and Wiesel, 1965). Therefore, to understand visual perception, one must understand the effects of visual experience. Although receptive field properties of cortical neurons in early visual areas become less plastic with age (Hubel and Wiesel, 1970), neurons later in the visual hierarchy exhibit plasticity well into adulthood. In particular, neurons in the functionally mature inferior temporal cortex (ITC)—a collection of areas in the primate brain hypothesized to underlie visual object recognition (DiCarlo and Cox, 2007; Logothetis and Sheinberg, 1996; Tanaka, 1996)—can adapt their responses to the statistics of visual input (Erickson and Desimone, 1999; Li and DiCarlo, 2008, 2010; Miyashita, 1988) and to a behavioral task's perceptual demands (Baker et al., 2002; Freedman et al., 2006; Kobatake et al., 1998; Logothetis et al., 1995; Op de Beeck et al., 2006). Neuronal activity in ITC is thus a joint product of accrued past experience and current input, and its investigation can shed light on the question of how memory and perception interact continuously at the level of single neurons.

Visual experience with a set of objects can be induced experimentally by mere exposure (Anderson et al., 2008; Freedman et al., 2006), by discrimination training (Baker et al., 2002; Freedman et al., 2006; Kobatake et al., 1998; Logothetis et al., 1995; Sigala and Logothetis, 2002), or by explicit memorization (Sakai and Miyashita, 1991). To infer the impact of visual experience on ITC, neuronal responses to familiar or learned stimuli are compared to a pre-exposure baseline (De Baene et al., 2008), to responses in untrained subjects (Kobatake et al., 1998), or most commonly, to responses to novel or unlearned stimuli (Anderson et al., 2008; Baker et al., 2002; Freedman et al., 2006; Logothetis et al., 1995; Miyashita et al., 1993). The resulting neuronal changes remain a matter of debate. Early studies reported that single neurons in ITC, on average, developed strong responses to a small (and different) subset of learned stimuli, which were larger than the maximal responses across the unlearned set (Kobatake et al., 1998; Logothetis et al., 1995; Miyashita, 1993; Sakai and Miyashita, 1994). Such strengthening of specific responses could amplify the neurons' impact on downstream areas, which would, in theory, facilitate behavior driven by recognition of well-known objects. However, recent studies have reported no change or even decreased maximal responses to familiar as compared to novel stimuli as well as a concomitant experience-dependent decrease in the overall population response (Anderson et al., 2008; Baker et al., 2002; Freedman et al., 2006; Op de Beeck et al., 2007, 2008). These divergent findings have been attributed to more unbiased single-unit selection procedures, to comparisons within rather than across animals, and to more finely controlled stimulus exposure protocols. Interestingly, while both firing rate increases and decreases can increase single-cell selectivity (i.e., narrow the tuning bandwidth), recently reported modulations have been on the order of a few spikes per second (Baker et al., 2002; Cox and DiCarlo, 2008; De Baene et al., 2008; Freedman et al., 2006), leading some to propose that visual experience results only in subtle neuronal plasticity in ITC (Op de Beeck and Baker, 2010). Behavioral data, on the other hand, indicate that the impact of visual experience on recognition behavior can be large (Gauthier and Tarr, 1997; Logothetis et al., 1995; Mruczek and Sheinberg, 2007).

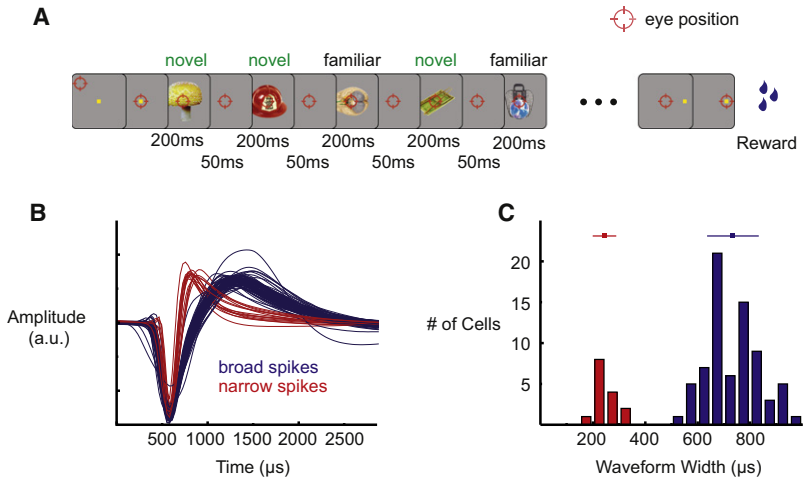

Figure 1. Experimental Paradigm and Spike Waveform Clustering

(A) Passive fixation task during which ten stimuli were presented for 200 ms each with a 50 ms interstimulus interval. Familiar and novel stimuli were interleaved.

(B) All recorded spike waveforms, aligned by their troughs and labeled according to their cluster membership. Waveform amplitudes have been normalized by their heights. a.u., arbitrary units.

(C) Distribution of spike widths (trough-to-peak durations) and the two clusters that emerged from the k-means algorithm. The bars above the distributions show mean ± SD of the respective distributions.

Two factors have impeded progress in our understanding of the effects of visual experience on single-unit responses in ITC. First, it is unclear with which stimuli to sample the tuning functions of individual ITC neurons. Advances have been made on this issue (Brincat and Connor, 2004, 2006; Rust and Dicarlo, 2010; Sáry et al., 1993; Tanaka, 1996; Yamane et al., 2008), but we are far from predicting responses to arbitrary visual patterns. The lack of increased responses and small selectivity increases to learned stimuli could thus be a result of not selecting the appropriate images to drive individual neurons; using large stimulus sets can partially ameliorate this issue. The second problem has been the averaging of responses over several distinct cell classes. We know that cortex comprises many different cell types (Connors and Gutnick, 1990; Markram et al., 2004; Peters and Jones, 1984), which mediate different functions within circuits. One means of distinguishing cell classes is by the shapes of their extracellularly recorded spikes (Barthó et al., 2004; Mitchell et al., 2007; Niell and Stryker, 2008). Data indicate that neurons that generate narrow spikes correspond primarily to fast-spiking inhibitory cells, whereas broad-spiking neurons correspond primarily to excitatory pyramidal cells (Barthó et al., 2004; Henze et al., 2000; Kawaguchi and Kubota, 1997; McCormick et al., 1985; Nowak et al., 2003). No studies to date, however, have probed the potential differential effect of visual experience on distinct cell classes in ITC.

Here, we show that experience caused putative excitatory neurons to respond much more robustly to their best familiar compared to their best novel stimuli. In contrast, familiarity caused a dramatic decrease in the maximum and average rates of putative inhibitory neurons. Together, the results suggest that visual experience can profoundly alter visual object representations in ITC.

RESULTS

To understand how long-term sensory input sculpts the responses of individual ITC neurons, we first familiarized each of two monkeys with 125 color images of real-world objects (Hemera Photo-Objects: Vol. 1, 2, and 3) (see Figure S1A available online). The monkeys were trained to both passively fixate the stimuli and to perform a short-term memory task with them. This exposure phase lasted between 3 months (monkey I) and 12 months (monkey D), resulting in an estimated number of exposures equal to 1,000 (monkey I) and 3,000 (monkey D) repetitions per image, split roughly evenly between the two tasks. Once familiarization was completed, we recorded the activity of well-isolated single units in ITC (n = 50 from monkey D; n = 38 from monkey I) in a passive fixation task (Figure 1A). Each neuron was screened with 125 familiar and 125 novel stimuli. The 125 novel stimuli were picked randomly on a daily basis from the same database as the familiar set (for examples, see Figures S1B–S1D). We recorded all units deemed visual by inspection of online stimulus-locked rastergrams. Both monkeys provided qualitatively similar data, so the results have been combined across subjects. Any notable differences are acknowledged (see Figure S3 for the main results split by monkey).

As a means of correlating visual response properties with specific cell classes, we characterized the recorded sample of single units by the trough-to-peak widths of their extracellular spike waveforms (Figures 1B and 1C). Consistent with previous studies (Diester and Nieder, 2008; Hussar and Pasternak, 2009; Mitchell et al., 2007), we observed that the distribution of these widths was bimodal, and we thus divided the neurons via a k-means algorithm into two categories: broad spiking and narrow spiking (Figure 1C). Previous results have suggested that narrow spikes correspond primarily to inhibitory, fast-spiking interneurons, whereas broad spikes correspond primarily to excitatory pyramidal neurons (Barthó et al., 2004; Connors and Gutnick, 1990; McCormick et al., 1985). For clarity, we thus refer to the narrow-spiking neurons as putative inhibitory and to the broad-spiking ones as putative excitatory.

Example Cells

Figures 2A–2G show the activity of seven representative single units. Each unit was stimulated with the same set of 125 familiar stimuli but with a different set of 125 novel stimuli. The top five rows (Figures 2A–2E) correspond to putative excitatory cells. In general, these units exhibited an enhanced response to the best familiar compared to the best novel stimulus. This advantage, however, was restricted to the highest ranked stimuli (with the notable exception of the unit shown in Figure 2C).

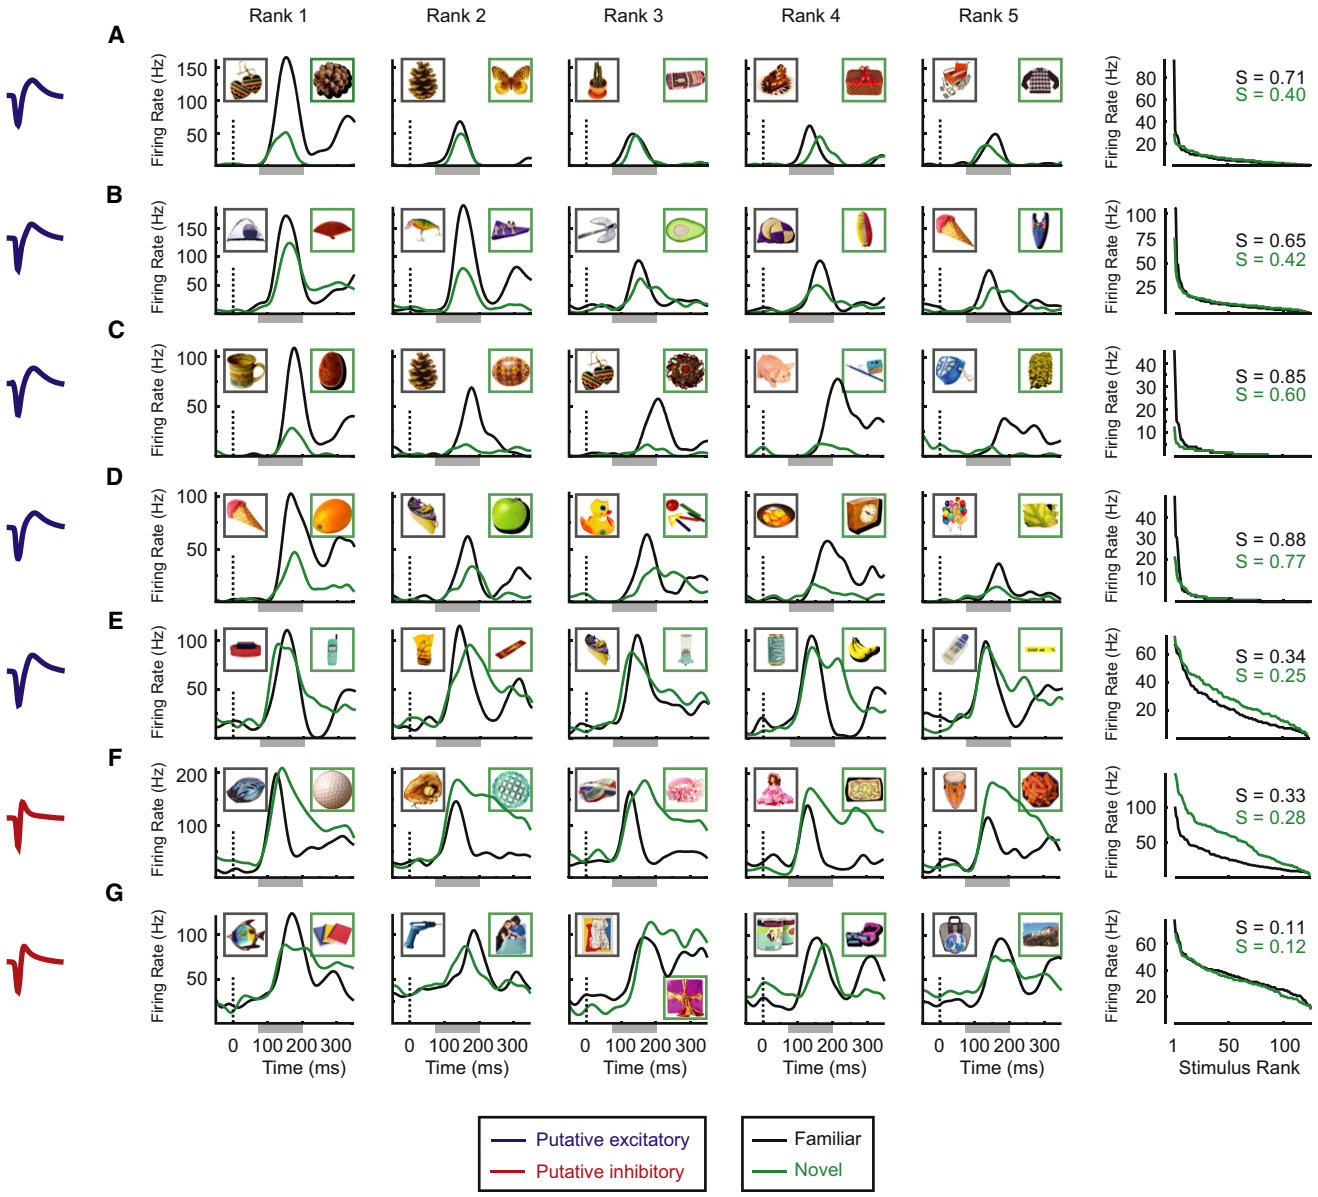

Figure 2. Example Neuronal Responses to Familiar and Novel Stimuli

(A–E) Five representative putative excitatory cells. (F and G) Two representative putative inhibitory cells. In all rows the column on the far left shows both the mean spike waveform of each cell and the cluster to which the waveform was assigned (blue, broad spike; red, narrow spike). In the middle-five columns are plotted the spike density functions (SDFs, spike times convolved with a Gaussian kernel with $\sigma = 20$ ms) for the top five stimuli from the familiar set (black) and the top five stimuli from the novel set (green). These rankings were determined not on the basis of the peak value of the SDF but rather from the spike counts in the interval 75–200 ms after stimulus onset, which is shown as a light-gray bar abutting the time axis. The insets in these graphs show the actual familiar and novel images eliciting the response. The column on the far right shows each neuron's entire distribution of mean firing rates, sorted according to rank. Again, the mean firing rates were computed from the spike counts in the interval 75–200 ms after stimulus onset, and the rankings were done independently for the familiar and novel sets. The numbers in the top right of the rank plots show the magnitude of the sparseness metric that was used to quantify single-cell selectivity.

Furthermore, note that the best familiar stimulus elicited a robust firing rate that reached a peak level of around 100 Hz in every neuron, suggesting that we were able to find highly effective stimuli for activating these neurons. The increased firing rates of putative excitatory cells to top-ranked familiar stimuli compared to top-ranked novel stimuli translated directly into increased selectivity (sparseness) for the familiar stimulus set (Figures 2A–2E, right column).

The bottom two rows (Figures 2F and 2G) correspond to putative inhibitory cells. Putative inhibitory cells nearly always showed a greater response to the best novel compared to the best familiar stimulus, an effect that appeared after the initial

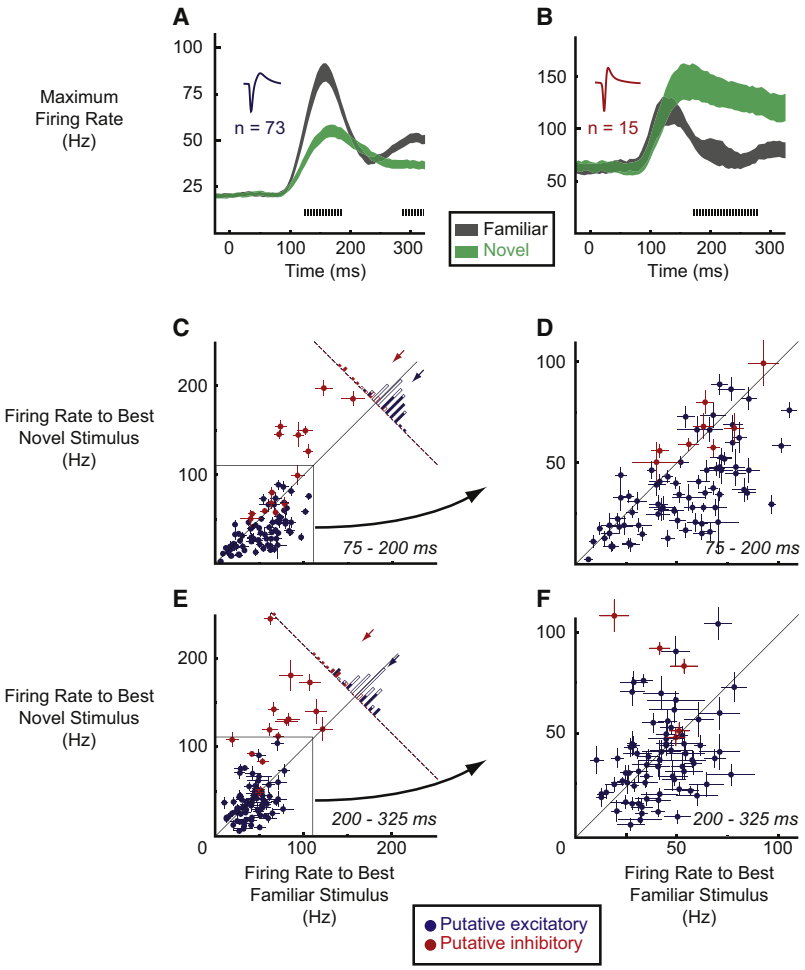

Figure 3. Visual Experience Increases Maximum Responses of Putative Excitatory Cells but Decreases Maximum Responses of Putative Inhibitory Cells

(A and B) Sliding window analyses (step size = 5 ms, window size = 50 ms) of maximum firing rates to familiar (black) and novel (green) stimulus sets, averaged separately for the putative excitatory (A) and putative inhibitory (B) cells. Shaded regions indicate ± SEM. Tick marks denote the time points at which the differences between the maximum familiar responses and maximum novel responses achieved statistical significance according to a permutation test ($p < 0.05$).

(C) Distribution of individual cells' responses to the best familiar (x axis) and best novel (y axis) stimulus during the early epoch (75–200 ms). Each data point represents the activity of a single unit. Cells are color labeled according to cluster membership (blue, putative excitatory; red, putative inhibitory). Error bars represent mean ± SEM across individual repetitions of the best familiar or best novel stimulus. Histogram in the top right shows the distribution of differences for both subpopulations. Shaded bars show individually significant cases ($p < 0.05$, Mann-Whitney U test). Arrows denote mean maximum response differences across either the putative excitatory (blue) or inhibitory (red) cells.

(D) A magnified view of the plot in (C), emphasizing the distribution of effects in the putative excitatory cells.

(E and F) Same as in (C) and (D) but for the late epoch (200–325 ms).

visual transient. These units also responded with an elevated rate to a much larger portion of stimuli than putative excitatory cells, regardless of stimulus set (Figures 2F and 2G, right column), and their firing rates could reach high peak values (~200 Hz; see Figure 2F). In addition, note that the reduced firing rates of putative inhibitory cells to familiar stimuli could span the entire range of ranks (Figure 2F, right column). While these experience-dependent firing rate changes could also result in selectivity increases, these were less reliable than those observed in putative excitatory cells (Figures 2F and 2G, right column).

Visual Experience Increases Maximum Responses of Putative Excitatory Cells but Decreases Maximum Responses of Putative Inhibitory Cells

We began with a simple question: Did experience with a set of stimuli result in the emergence of stronger ITC responses, and if so, did this effect depend on cell class? Because neurons in ITC can exhibit marked selectivity, and thus fail to be activated by many stimuli independent of experience, we narrowed the focus of this query to just the maximum responses. In particular, for every neuron we extracted a pair of mean firing rates: one elicited by the single most effective familiar stimulus, and one by the single most effective novel stimulus.

To gain insight into the time course of experience-dependent maximum firing rate differences, we first computed this statistic with a sliding window (step size = 5 ms; window size = 50 ms). In Figure 3A we see that, averaged across the population of putative excitatory cells, the maximum responses to the familiar set were much greater than to the novel set, and this difference emerged at about the same time as the onset of the visual response (earliest significant difference = 120 ms; $p < 0.05$, permutation test, corrected for multiple comparisons; see Supplemental Experimental Procedures). In contrast, averaged across the population of putative inhibitory cells (Figure 3B), the maximum responses to the familiar set were much smaller than to the novel set, and this difference did not emerge until after the initial visual transient (earliest significant difference = 170 ms).

We next examined experience-dependent maximum firing rate differences in individual units. We divided the data into two time epochs: an early epoch of 75–200 ms, and a late epoch of 200–325 ms. In Figures 3C–3F, we plot for each epoch, and at two different scales to emphasize the distribution of putative excitatory units, the magnitude of each cell's response to its single best familiar and to its single best novel stimulus. In the early epoch (Figures 3C and 3D), the majority of putative excitatory cells (blue points) lie below the diagonal line, indicating that for these neurons the best familiar stimulus elicited a stronger response than the best novel stimulus. Averaged across the population of putative excitatory cells, the firing rate to the best familiar stimulus was 16.55 ± 2.22 Hz

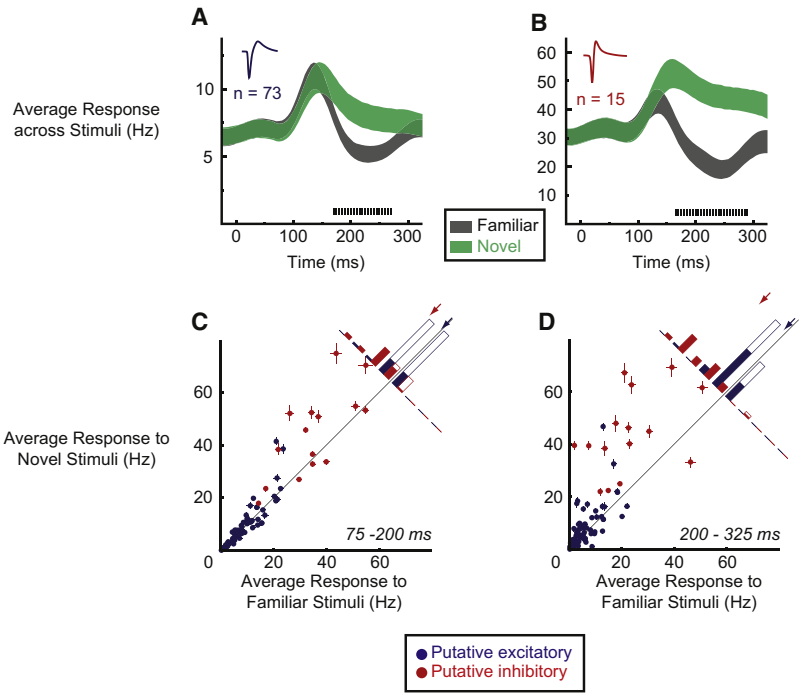

Figure 4. Visual Experience Decreases Average Stimulus-Evoked Responses of Putative Excitatory and Inhibitory Cells

Conventions same as in Figure 3 with the notable exception that the metric of interest is the average, not maximum, response across the 125 familiar or 125 novel stimuli. Error bars in (C) and (D) represent mean ± SEM across the 125 familiar or 125 novel (mean) firing rates. Individually significant cases in histograms of (C) and (D) were determined with a t test ($p < 0.05$).

(mean ± SEM) greater than the firing rate to the best novel stimulus (blue arrow in Figure 3C; $p < 0.001$, paired t test), an increase of nearly 50% (52.69 Hz compared to 36.14 Hz). In the late epoch (Figures 3E and 3F), this difference diminished (blue arrow in Figure 3E, familiar − novel, 4.40 ± 2.41 Hz; $p = 0.07$).

Putative inhibitory cells led to a different distribution of maximum firing rate differences (Figures 3C and 3E, red points). In both the early (Figure 3C) and late (Figure 3E) epochs, most putative inhibitory cells were driven to a much higher firing rate by their best novel than by their best familiar stimulus (red points above unity diagonal). In the early epoch the population-averaged difference in maximum firing rate was 27.63 ± 7.97 Hz in favor of the novel set (red arrow in Figure 3C; $p = 0.004$, paired t test) but significant only in one monkey (compare Figures S3C and S3D), whereas in the late epoch it rose to 53.65 ± 12.11 Hz (red arrow in Figure 3E, novel − familiar; $p < 0.001$) and became significant in each monkey.

Visual Experience Decreases the Average Stimulus-Evoked Firing Rate in Putative Excitatory and Inhibitory Cells

We next asked how neuronal responses to familiar and novel stimuli differ when averaged across the entire ensemble of stimuli. Such an analysis offers a glimpse into ITC neurons' more typical firing rate modulations, that is, their stimulus-evoked firing rates to a randomly chosen, as opposed to their most effective, stimulus. We computed for each cell its average stimulus-evoked response, which we defined as the average over the mean firing rates to each of the 125 stimuli within either the familiar or novel set (Figures 4A–4D). Paralleling previous reports that have grouped neurons into two distinct classes based on extracellular spike waveform (Diester and Nieder, 2008; Mitchell et al., 2007), we first note that putative inhibitory units had much larger stimulus-driven activity than putative excitatory units. This can be appreciated by comparing the axes in Figure 4A (putative excitatory) and Figure 4B (putative inhibitory) and by comparing the blue (putative excitatory) and red (putative inhibitory) points in Figures 4C and 4D. To quantify this difference, we compared the average stimulus-evoked firing rates of putative excitatory cells to those of putative inhibitory cells within each unique combination of stimulus set (familiar/novel) and time epoch (early/late). All comparisons were highly significant (mean ± SEM Hz for putative excitatory versus putative inhibitory: familiar early, 8.62 ± 0.70 versus 35.12 ± 3.24; familiar late, 5.90 ± 0.60 versus 22.96 ± 3.54; novel early, 9.20 ± 0.92 versus 44.26 ± 4.21; novel late, 7.79 ± 0.91 versus 44.00 ± 4.01; $p < 0.001$ for every comparison, uncorrected, two-sample t tests). Because it has been shown that current injections can drive fast-spiking inhibitory units to very high firing rates (McCormick et al., 1985), the higher average responses of narrow-spiking units further support the labeling of this cell class as putative inhibitory. We observed a similar difference in firing rates when we looked at spontaneous activity, which we took as the last 500 ms of the fixation epoch (putative excitatory, 5.20 ± 0.68 Hz; putative inhibitory, 15.01 ± 2.87 Hz; $p = 0.004$, two-sample t test).

Notably, we found that in both cell classes the novel set elicited higher average responses than the familiar set (Figures 4A–4D). Like the maximum response effect in putative inhibitory units, these experience-dependent differences in average firing rate emerged, in both cell classes, after the initial visual transient (Figures 4A and 4B). In particular, in the early epoch (Figure 4C), the population-averaged difference for the putative excitatory cells was small and not significant (familiar − novel, mean ± SEM, −0.59 ± 0.42 Hz; $p = 0.17$, paired t test), and whereas the difference was larger and significant in the putative inhibitory subset (familiar − novel, −9.14 ± 2.85 Hz; $p = 0.006$), it was only observed in one monkey (compare Figures S3C and S3D). It was in the late epoch (Figure 4D) that population-averaged differences in average firing rate for both classes of cells became significantly different from zero (familiar − novel; putative excitatory, −1.90 ± 0.67 Hz, $p = 0.006$; putative inhibitory, −21.04 ± 4.01 Hz, $p < 0.001$; in one monkey the putative excitatory effect was marginally significant, $p = 0.09$). Consistent with these

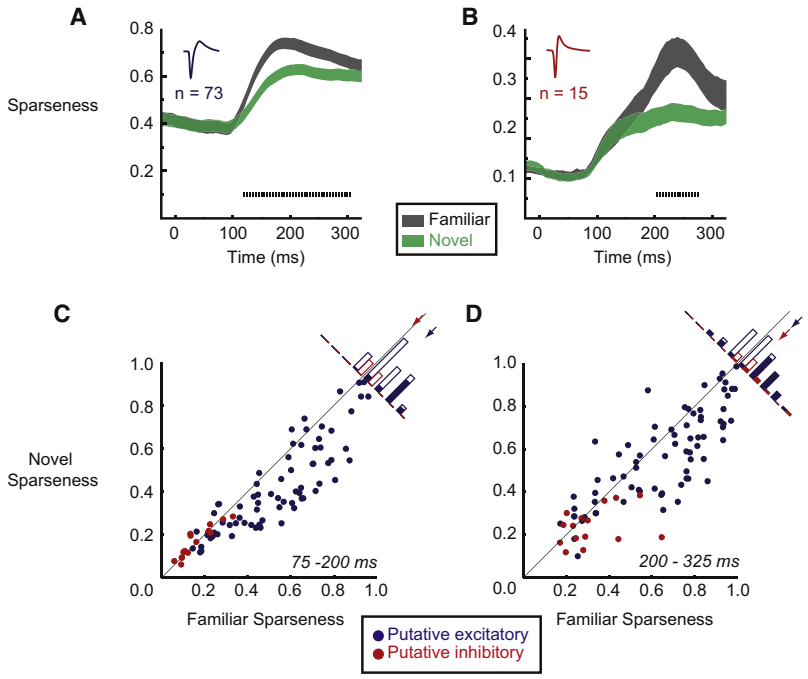

Figure 5. Visual Experience Increases Selectivity of Putative Excitatory and Inhibitory Cells

Same conventions as in Figures 3 and 4, except that the metric investigated is sparseness across the 125 familiar or 125 novel stimuli. Individually significant cases in histograms of (C) and (D) were determined with a permutation test ($p < 0.05$).

observations, we also observed that experience led to decreases in the proportion of stimuli eliciting a significant elevation in firing rate and to increases in the proportion of stimuli eliciting a significant reduction in firing rate (Figure S4). Furthermore, although both cell classes showed reduced average responses to familiar stimuli, this decrease was much larger in putative inhibitory than excitatory cells (early epoch, $p = 0.001$; late epoch, $p < 0.001$; two-sample t tests; early epoch effect not significant in the same monkey whose effects tended to arise later), which can be seen by comparing the red and blue arrows in the histograms of Figures 4C and 4D.

Visual Experience Increases Selectivity of Putative Excitatory Cells

To convey information, neurons modulate their firing rates. The greater and/or more reliable this modulation, the more informative the neuron's firing rate becomes about the presence (or absence) of some stimulus. Because we have shown that visual experience not only led to an increase in maximum response (in putative excitatory cells) but also to a decrease in average response, we have already implicated visual experience in sharper stimulus selectivity. Here, we make this idea explicit.

To capture increases in selectivity with a single metric, we computed the value of (lifetime) sparseness (Olshausen and Field, 2004; Rolls and Tovee, 1995; Vinje and Gallant, 2000; Zoccolan et al., 2007) (see Experimental Procedures). Sparseness quantifies how much of a single neuron's total firing rate, across a stimulus set, is concentrated within a few stimuli. A neuron with high sparseness will be quiet most of the time, but there will be a few stimuli that elicit robust firing rates. By definition, this is a selective neuron. An unselective neuron, one with low sparseness, will respond with an elevated firing rate to many stimuli. We calculated the sparseness of cells' responses across the familiar and novel stimulus sets, first with a sliding window (Figures 5A and 5B) and then in the previously defined early and late epochs (Figures 5C and 5D).

As with the average response analyses, one of the more conspicuous features of the data was that putative inhibitory units had much lower sparseness than putative excitatory units for every combination of stimulus set and epoch (mean ± SEM putative excitatory versus putative inhibitory; familiar early, 0.53 ± 0.03 versus 0.16 ± 0.02; familiar late, 0.65 ± 0.03 versus 0.32 ± 0.04; novel early, 0.42 ± 0.02 versus 0.17 ± 0.02; novel late, 0.57 ± 0.02 versus 0.24 ± 0.02; $p < 0.001$ for every comparison, uncorrected, two-sample t tests). The broad tuning of putative inhibitory units is consistent with recent functional data (Kerlin et al., 2010; Liu et al., 2009; Sohya et al., 2007) as well as neuroanatomical data showing that these units can receive highly convergent and heterogeneous input from the surrounding excitatory population (Bock et al., 2011).

Importantly, we found that the sparseness of putative excitatory cells was significantly greater for familiar than novel stimuli, in both the early and late epochs (compare black and green curves in Figure 5A; see blue points and arrows in Figures 5C and 5D; mean ± SEM familiar − novel; early epoch, 0.11 ± 0.01; late epoch, 0.08 ± 0.02; $p < 0.001$ in both instances, paired t tests).

In the putative inhibitory population, we observed a somewhat different and less conclusive set of results. First, note that the familiar sparseness for this population of cells did not reach its peak value until late in the visual response (black curve in Figure 5B). Averaged across the population of narrow-spiking neurons, sparseness for familiar stimuli was significantly greater than for novel stimuli only in the late epoch (compare black and green curves in Figure 5B, see red points and arrows in Figures 5C and 5D; mean ± SEM familiar − novel; early epoch, −0.01 ± 0.01, $p = 0.43$; late epoch, 0.08 ± 0.04, $p = 0.04$; paired t tests) and only in one monkey (late epoch, monkey D, $p = 0.19$; monkey I, $p = 0.01$).

The selectivity analyses argue that the sparseness of putative excitatory, and possibly putative inhibitory cells, in ITC is not a static property but rather one that visual experience can increase. In general, sparseness can be increased either by increasing the proportion of near-zero responses (Tolhurst et al., 2009) or by increasing the response magnitude to a subset of the most effective stimuli. We have already shown that in the early epoch, putative excitatory cells had higher maximum responses to familiar than novel stimuli. Could this difference

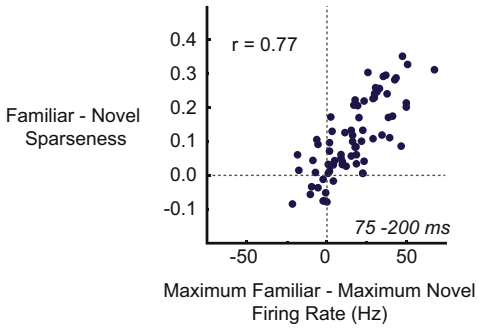

Figure 6. For Putative Excitatory Cells the Experience-Dependent Increase in Maximum Response Predicts the Experience-Dependent Increase in Selectivity

The difference between familiar and novel sparseness is plotted as a function of the difference between maximum familiar and maximum novel responses. Each point represents a single putative excitatory unit. Maximum responses and sparseness values were taken from the early epoch (75–200 ms).

account for these cells' increased sparseness? We addressed this question by subtracting for each putative excitatory cell its maximum response across the novel set from its maximum response across the familiar set and then by correlating these differences with the differences between familiar and novel sparseness (Figure 6). Indeed, the experience-dependent increase in maximum response of putative excitatory cells was a good predictor of how much more selective individual cells were to stimuli within familiar compared to novel sets (Pearson's $r = 0.77$, $p < 0.001$; $r = 0.80$ in monkey D, $r = 0.75$ in monkey I). No such relationship was observed in the late epoch ($r = 0.00$; $p = 0.998$) or in the early or late epochs of putative inhibitory cells (early, $r = 0.27$, $p = 0.33$; late, $r = -0.06$, $p = 0.82$) (data not shown). We further confirmed the robust contribution of the differences in maximum firing rates to selectivity changes with a randomization procedure (Figure S6). We conclude that, in the early epoch, experience-dependent increases in the putative excitatory cells' maximum responses contributed to a sparser (more selective) representation of familiar compared to novel stimuli. It is important to note that this conclusion is different from the more traditional concept of a sparse neuron as an infrequently active neuron (Haider et al., 2010; Rolls and Tovee, 1995; Tolhurst et al., 2009; Vinje and Gallant, 2000). Here, the analyses suggest that increased sparseness resulted in a neuron that fired more spikes to its preferred stimulus. Only in the late epoch (of both putative excitatory and inhibitory cells) did we find that the experience-dependent increases in sparseness could be better accounted for by decreases in the proportion of familiar stimuli eliciting a significantly elevated response (data not shown).

Visual Experience Does Not Impede the Ability of Putative Excitatory and Inhibitory Cells to Discriminate between Novel Stimuli

In our experiments, visual experience caused marked differences in neuronal responses to familiar versus novel stimuli. Nonetheless, novel stimuli elicited robust activity from the population of recorded ITC neurons, indicating that neuronal activity in ITC can contribute to the recognition of both stimulus sets. Could ITC neurons discriminate as well among members of the novel set as of the familiar set? We probed this question with a receiver operating characteristic (ROC) analysis. In particular, we performed ROC analyses on all possible pairwise combinations of stimuli (within a set), each time summarizing the discriminability of the two firing rate distributions with the area under ROC curve (AUC) (Rust and Dicarlo, 2010). We took the average of the AUC values as a metric of overall discriminability, which captured how well, on average, a single neuron's spike counts could discriminate between the identities of any two arbitrarily chosen stimuli.

We first note that putative inhibitory cells conveyed more information about stimuli, familiar and novel, than did putative excitatory cells (Figures 7A and 7B, compare blue to red points) (mean ± SEM putative excitatory versus putative inhibitory; familiar early, 0.673 ± 0.008 versus 0.702 ± 0.011; familiar late, 0.648 ± 0.007 versus 0.698 ± 0.011; novel early, 0.665 ± 0.008 versus 0.729 ± 0.013; novel late, 0.682 ± 0.009 versus 0.778 ± 0.009; $p = 0.04$ for familiar early comparison, where the difference was not significant in one monkey; $p < 0.001$ for all other comparisons, familiar late comparison was not significant in same monkey, uncorrected, two-sample t tests). This finding is consistent with the broader tuning of putative inhibitory cells, which allowed them to respond in a stimulus-selective manner to more than just the top few stimuli.

Notably, we found that spike counts of both putative excitatory and inhibitory cells could be used to discriminate between novel stimuli as well as, or even better than, familiar stimuli. The only case in which the familiar set fared better was the early epoch of putative excitatory cells, but this difference was small and not significant in either monkey separately (Figure 7A, blue points and arrow; mean familiar AUC = 0.673, mean novel AUC = 0.665; $p = 0.046$, paired t test). Furthermore, note that the late epoch of putative excitatory cells more than compensated for this initial difference (Figure 7B, blue points and arrow; mean familiar AUC = 0.648, mean novel AUC = 0.682; $p < 0.001$). For the putative inhibitory cells, the novel stimuli could be better discriminated in both epochs (Figures 7A and 7B, red points; early epoch, mean familiar AUC = 0.702, mean novel AUC = 0.729 $p = 0.004$; late epoch, mean familiar AUC = 0.698, mean novel AUC = 0.778, $p < 0.001$), with one monkey showing much stronger and reliable differences than the other. Visual experience, therefore, did not prevent neurons in ITC from contributing reliably to the encoding of both familiar and novel stimuli.

Given that putative inhibitory cells had lower sparseness than putative excitatory cells but were better able to discriminate between any two arbitrarily chosen images, we wondered whether there was a relationship between sparseness and mean pairwise AUC values. In Figures 7C and 7D, we have plotted individual cells' sparseness and mean pairwise AUC values for the early and late epochs (putative inhibitory units are indicated by open symbols). For both familiar (Figures 7C and 7D, black points and lines) and novel (green points and lines) stimuli, we observed a strong linear correlation between the two metrics. The correlation held even when we restricted the analysis to just the putative excitatory cells (Figures 7C and 7D, filled

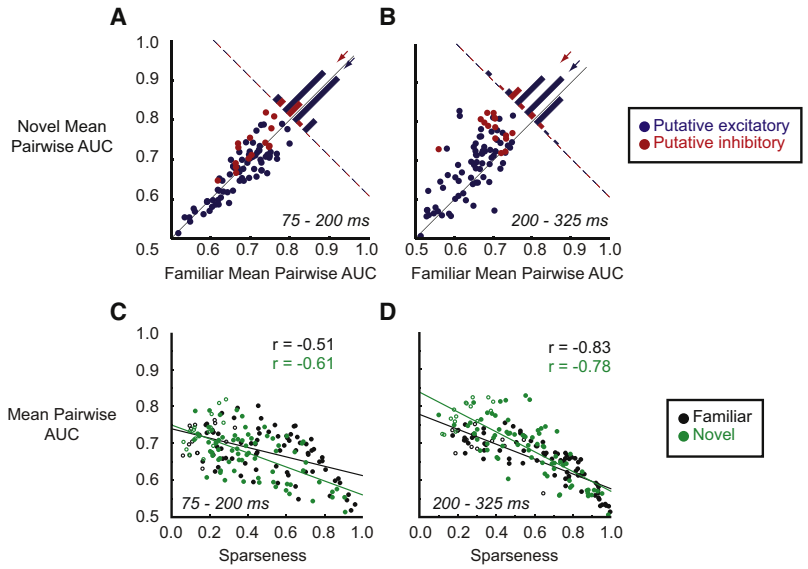

Figure 7. Putative Excitatory and Inhibitory Units Can Discriminate between Stimuli within the Familiar and Novel Sets

(A) Distribution of individual cells' familiar (x axis) and novel (y axis) mean pairwise AUC values during the early epoch (75–200 ms). Each data point represents the mean pairwise AUC value of a single unit. Cells are color labeled according to cluster membership (blue, putative excitatory; red, putative inhibitory). Histogram in the top right shows the distribution of AUC differences for both subpopulations. For clarity all bars are shaded, but this does not indicate significance. Arrows denote mean differences across either the putative excitatory (blue) or inhibitory (red) cell class.

(B) Same as in (A) but for the late epoch (200–325 ms).

(C) Relationship between sparseness and mean pairwise AUC value for familiar (black) and novel (green) stimuli during the early epoch (75–200 ms). Putative inhibitory units are indicated by open circles.

(D) Same as in (C) but for the late epoch (200–325 ms).

circles). This suggests that an increase in sparseness precluded a neuron from discriminating stimuli at the lower end of its firing rate distribution. Because visual experience led to a considerable increase in sparseness, we conclude that individual ITC neurons contributed to the encoding of a smaller number of familiar compared to novel stimuli.

DISCUSSION

Here, we asked whether visual long-term experience's effects on single-neuron responses in ITC vary with cell type. We first showed that the best stimulus from the familiar set drove putative excitatory cells much more robustly than the best stimulus from the novel set. This effect was reversed for putative inhibitory cells. We further showed that, on average, both putative excitatory and putative inhibitory neurons responded with a smaller response to a randomly chosen familiar compared to novel stimulus, but this difference was much larger in the putative inhibitory population. We then went on to show that experience increased sparseness in putative excitatory neurons and, to a lesser degree, in putative inhibitory neurons. For the putative excitatory neurons, the experience-dependent increase in sparseness could be well accounted for by an increased firing rate to the top familiar stimulus. Finally, we demonstrated that the experience-dependent modifications have a minimal impact on the ability of ITC neurons to discriminate between the stimuli in the novel set. In Figure 8, we provide a schematic summarizing the observed firing rate changes in both classes of neurons.

Methodological Approach

Neurons in neocortex can be classified on the basis of morphology, physiology, connectivity, laminar distribution, neurotransmitter content, and/or expression of calcium-binding proteins, to name the most common schemes (Markram et al., 2004). In extracellular recording studies, most of these characteristics remain unknown, leading many to simply average results over all recorded cells, potentially obscuring important cell class-dependent differences. However, a growing body of evidence supports the utility of dividing extracellularly recorded spikes into putative excitatory and inhibitory classes based on spike shape (Barthó et al., 2004; Johnston et al., 2009; Tamura et al., 2004). The technique's foundation rests on results suggesting that fast-spiking, parvalbumin-positive inhibitory interneurons express an abundance of Kv3 voltage-gated potassium channels, which endow them with their unique narrow action potentials (Kawaguchi and Kubota, 1997; McCormick et al., 1985; Rudy and McBain, 2001).

As with any classification scheme, caution should be exercised with this method's application. Indeed, a recent electrophysiological study from the primary motor cortex of the monkey showed that pyramidal tract neurons can also emit narrow spikes (Vigneswaran et al., 2011). Whether such results will be extended to cortical areas with a less-specialized corticospinal projection, a more representative distribution of cell types, and a more typical laminar profile remains an open question, but it is unlikely neuronal classification based on spike waveform alone can represent a one-to-one mapping (Nowak et al., 2003). Nonetheless, the method offers an important first step for dividing a sample of neurons into putatively different cell classes, i.e., it is better than no division at all if functional differences between the two classes can be shown to exist (Diester and Nieder, 2008; Hussar and Pasternak, 2009; Mitchell et al., 2007). For ease of exposition we thus assume this division in the following discussion.

Putative Excitatory Cells

Several studies have explored the impact of visual experience on the maximum response magnitude of single ITC neurons. Early work showed that the best familiar stimulus elicits a higher firing rate than the best novel stimulus (Kobatake et al., 1998; Miyashita, 1993; Sakai and Miyashita, 1994). More recent work, however, has revealed that the best familiar and best novel

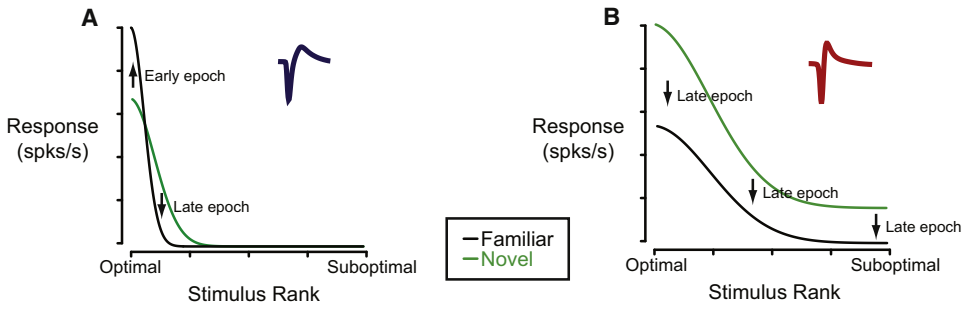

Figure 8. Schematic Representation of Experience-Dependent Firing Rate Changes in Putative Excitatory and Putative Inhibitory Units

(A) Firing rates of putative excitatory neurons are arranged in descending order of effectiveness. In this cell class, visual experience increased responses to the most effective stimuli, particularly in the early epoch, and decreased responses to moderately effective stimuli, especially in the late epoch. spks, spikes.

(B) Same as in (A) but for putative inhibitory units. Visual experience caused a much more widespread and noticeable decline in firing rates of these neurons. This change was most prominent in the late epoch.

stimuli, on average, evoke equivalent firing rates (Baker et al., 2002; Freedman et al., 2006; Op de Beeck et al., 2007). Here, we have provided data reconciling these disparate results by showing that whether experience increases or decreases the maximum response depends on both cell class and over what time epoch firing rates are computed. In particular, putative excitatory cells responded more strongly to the best familiar stimulus, but only in the early epoch, whereas putative inhibitory cells responded more strongly to the best novel stimulus, particularly in the late epoch. Given that excitatory cells are estimated to outnumber inhibitory cells by a ratio of about 4:1 (Markram et al., 2004), can the averaging across the two cell classes account for the recent absence of maximum response differences? In principle, this is possible because the absolute magnitude of the experience-dependent maximum response modulation was much larger for the putative inhibitory than putative excitatory cells (absolute difference, putative excitatory early phase = 16.55 Hz, putative inhibitory late phase = 53.65 Hz). Indeed, calculating firing rates over the window 75–325 ms post-stimulus onset and collapsing across the two cell classes leads to a much reduced and in one monkey a nonsignificant maximum response difference between the familiar and novel stimulus sets (two monkeys combined, best familiar − best novel = 2.64 Hz, paired t test, $p = 0.40$; monkey D, −0.21 Hz, $p = 0.97$; monkey I, 6.40 Hz, $p = 0.02$; in the monkey in which the difference remained significant, the difference decreased from 11.93 Hz when computing it from early epoch spike counts of putative excitatory cells alone, nearly a 50% decrease).

Another potential explanation as to why some reports have failed to observe an enhanced response to the best familiar stimulus concerns the size of the stimulus sets. In the studies where the best familiar stimulus failed to elicit a stronger response, the familiar and novel sets each consisted of no more than 20 stimuli (Baker et al., 2002; Freedman et al., 2006; Op de Beeck et al., 2007). Conversely, each of the studies that have reported stronger familiar responses used stimulus sets with at least that many stimuli (Kobatake et al., 1998; Logothetis et al., 1995; Miyashita, 1993; Sakai and Miyashita, 1994). With a small and/or relatively homogeneous stimulus set, it is plausible that the lack of enhanced familiar responses is a consequence of exploring only the low-response regions of the high-dimensional image space in which ITC responses lie, regions in which responses to familiar and novel stimuli are similar. Consistent with this proposal, when we randomly selected smaller subsets of familiar and novel responses (from our own data set), and thus were more likely to exclude the response from the best familiar stimulus, we observed that the population level difference in maximum firing rates decreased (Figure S5). Further supporting the suggestion that the differences in maximum firing rate depend on finding the appropriate stimuli, two of the studies that failed to observe an enhanced familiar response reported the firing rates to the best familiar stimuli to be <25 Hz (Baker et al., 2002; Freedman et al., 2006). Because this value presumably included both excitatory and inhibitory neurons, it is likely to be even lower for just excitatory neurons. In the present study we recorded from putative excitatory cells that had an average maximum response to the familiar set of 52.69 Hz (taken over the epoch 75–200 ms) and a peak maximum response, depending on the monkey, of around 70–110 Hz.

What could the increased response magnitude of the putative excitatory cells to the best familiar stimulus reflect in terms of the underlying neuronal circuitry? Because the experience-dependent enhancement was present at the time of visual response onset, the most parsimonious explanation is to posit a potentiated excitatory input from areas upstream of ITC, such as V4 (Seltzer and Pandya, 1978). This hypothesis is consistent with the present conception of ventral visual stream function. In particular, the ventral visual stream is thought to elaborate on the shape, color, and texture attributes of visual input (Anzai et al., 2007; Brincat and Connor, 2004, 2006; Gallant et al., 1993; Hubel and Wiesel, 1959; Kobatake and Tanaka, 1994; Logothetis et al., 1996; Pasupathy and Connor, 1999; Rust and Dicarlo, 2010; Tanaka, 1996; Tanaka et al., 1991; Yamane et al., 2008). The gradual increase in optimal stimulus complexity as one traverses the ventral pathway has been interpreted as an increase in sensitivity for particular combinations of local features. This sort of image transformation makes explicit, and thus easier to readout, the higher-order correlations present in the visual input. This process is thought to culminate in ITC. Because the local feature responses of neurons at early stages

in the visual system can be recombined in a virtually infinite number of ways, there is no need for their experience-dependent modification beyond that observed in the critical period. Indeed, modification of these building blocks of stimulus encoding could dramatically disrupt responses of downstream neurons dependent on a stable foundation of local responses. The particular combinations of local features that the organism learns to recognize, however, will depend on its recent perceptual history. We propose that one of ITC's computational roles is to learn and encode with a higher maximum response those conjunctions that occur frequently and reliably. To do so, neurons in ITC strengthen the influence of those synaptic inputs that have a tendency to frequently and reliably excite them. Such learning can be implemented through classical Hebbian plasticity mechanisms, and in particular, NMDA receptor (NMDAR)-mediated long-term potentiation (LTP) (Feldman, 2009). Supporting this hypothesis, stimulus-specific, NMDAR-mediated response potentiation has previously been reported in mouse visual cortex (Frenkel et al., 2006). It will be important for future studies to determine whether the neuronal changes to the stimuli we used can or cannot be detected earlier in the visual system (Rainer et al., 2004; Yang and Maunsell, 2004). Under our proposed scheme, such changes should be minimal.

We showed that a direct result of experience-dependent maximum response increases in putative excitatory cells is increased sparseness (selectivity) for stimuli within the familiar set. This is consistent with earlier work (Kobatake et al., 1998; Logothetis et al., 1995) but stands in contrast to recent data showing that selectivity increases in ITC are a consequence only of decreased responses to stimuli at the lower end of the firing rate distribution (Baker et al., 2002; Freedman et al., 2006). While we were able to replicate the decrease in average stimulus-evoked responses, this effect's presence (Freedman et al., 2006), as well as its relationship to increased selectivity, held only in the late phase of the visual response. The late emergence of this suppression suggests that experience not only strengthens feed-forward input but also likely prunes and/or weakens synaptic connections within ITC (Feldman, 2009). Taken together, these results argue that experience steers putative excitatory neurons to contribute to the encoding of only their most effective stimuli at the expense of less-effective stimuli. Supporting this assertion, we showed that there is an inverse relationship between the selectivity of neurons and their ability to discriminate arbitrarily chosen pairs of stimuli. We speculate that a smaller population of projection neurons each firing many, very informative spikes may be better at driving downstream neurons and thus have more impact on perceptually guided behavior compared to a large population of neurons each firing a few, less-informative spikes.

Putative Inhibitory Cells

Putative inhibitory cells also showed average response decreases to familiar stimuli. The magnitude of this effect, however, was much larger in the inhibitory population. This observation adds to recent reports showing that behavioral factors can affect putative inhibitory cells to a much greater degree (Mitchell et al., 2007; Niell and Stryker, 2010). One intriguing possible role for increased inhibitory output is that it serves to detect novelty and initiate the cascade of events that underlie the subsequent plasticity. Research over the past decade has revealed that critical period plasticity within primary visual cortex is closely linked with the maturation of GABAergic transmission, with anecdotal reports implicating, in particular, inhibition mediated by parvalbumin-positive interneurons (Hensch, 2005). Indeed, a recent report indicates that interneurons of this class broaden their orientation tuning in parallel with the onset of the critical period (Kuhlman et al., 2011). We thus propose that the increased activity of our putative inhibitory cells is the neurochemical trigger for the robust selectivity changes within the putative excitatory population. If this hypothesis is true, the challenge will be to elucidate what allows the inhibitory cells within ITC to mediate plasticity into adulthood. That is, even though in primary visual cortex critical period plasticity can be prematurely triggered by enhancing GABAergic transmission, the plastic window still has a finite duration, and importantly, once it ends, it cannot be reinitiated (Fagiolini and Hensch, 2000). Further work suggests that there is a developmental trajectory intrinsic to inhibitory cells, which allows them to control the temporal specificity of plasticity (Southwell et al., 2010). Whether this maturational program is in some important ways different in inhibitory cells further along the visual hierarchy, where plasticity can extend into adulthood, is a question for future research.

Our observation that putative inhibitory cells were much less selective than putative excitatory cells, regardless of stimulus set and time epoch analyzed, is consistent with a previous result (Zoccolan et al., 2007). In areas where columnar structure with regard to some feature dimension is well defined (e.g., orientation columns in cat and primate primary visual cortex), inhibitory neurons have narrow tuning. In areas lacking such an organization (e.g., primary visual cortex of mice and rabbits), inhibitory neurons have broader tuning. Thus, an emerging view is that the amount of selectivity within the inhibitory population reflects the degree to which excitatory neurons with similar receptive field properties are in spatial proximity to one another (Bock et al., 2011; Cardin et al., 2007; Kerlin et al., 2010; Liu et al., 2009; Sohya et al., 2007). To the extent that this hypothesis is true, our results indicate that columnar organization within ITC, with respect to the stimulus set employed, is moderate at best (Fujita et al., 1992; Tsunoda et al., 2001). Otherwise, we should have seen selectivity values within the putative inhibitory population mirror the selectivity values within the putative excitatory population. Importantly, we can extend this line of reasoning and propose that inhibitory activity serves as a proxy for the amount of surrounding excitatory activity. Viewed in this light, the massive increase in the average response of our putative inhibitory population to the novel stimuli further speaks to the robust effects that experience exerts on neuronal circuitry in ITC. In other words the increased inhibitory activity is consistent with the hypothesis that novel compared to familiar stimuli activate a much larger number of excitatory cells and/or drive them, on average, to fire many more spikes. It is worth noting that perhaps the reason why putative inhibitory cells are better at detecting the novelty of stimuli is because they "listen" to the summed excitatory output of a fairly large collection of surrounding neurons. In this manner, the massive increase in inhibitory output would serve to not only signal novelty but also

to maintain an appropriate level of excitatory to inhibitory balance. In fact, maintenance of this balance could be crucial to the normal operation of this sensory circuit while it undergoes robust remodeling. Alternatively, another nonmutually exclusive hypothesis is that this balance is important for putting the brakes on too much plasticity occurring too rapidly. Answers to these questions await further experimental exploration.

EXPERIMENTAL PROCEDURES

All experimental procedures were in accordance with the guidelines published in the National Institutes of Health Guide for the Care and Use of Laboratory Animals and were approved by the Brown University Institutional Animal Care and Use Committee. Two male adult macaque monkeys were used in this study. Standard operant conditioning techniques were used to train the subjects to fixate and to press buttons for a small liquid reward. Eye movements were recorded using the EyeLink II video tracking system (SR Research, Osgoode, Ontario, Canada) running at 500 Hz. When the monkeys were ready for recordings, we implanted custom chambers that allowed for a dorsal access to ITC (Horsley-Clark coordinates, +15 anterior, +20 lateral). Based on reconstructed electrode trajectories, we believe most of our recordings took place from the lateral convexity of ITC, ventral to the lower bank of the superior temporal sulcus (STS) and lateral to the perirhinal cortex (Figure S2). Recordings were obtained with fine tungsten microelectrodes (Alpha Omega Engineering, Alpharetta, GA, USA, or Frederic Haer Company, Bowdoinham, ME, USA). Single units were isolated online using a threshold and dual-amplitude windows, while analog signals were streamed to disk for offline analysis.

All stimuli used were taken from Hemera Photo-Objects Vols. 1, 2, and 3 (Hemera Technologies), subtended about 2° × 2° of visual angle at a viewing distance of 90 cm, and were presented centrally on top of a uniform gray background. Both monkeys were familiarized with the same set of 125 stimuli (Figure S1A). During the familiarization phase the monkeys saw the images in either a passive fixation task or in a delayed match-to-sample task. When the familiarization phase was completed, we began the recordings. All recordings were obtained during a passive fixation task in which eye position was constrained to be within 1° of the center of the screen, as ten stimuli (no repeats) were presented. At the end of the stimulus presentation epoch, an extrafoveal square target was presented (eccentricity = 6°) to which the monkey had to saccade to obtain its juice reward. Because the goal of this experiment was to compare neuronal responses to familiar and novel stimuli, for every recording session we selected a new set of 125 never before seen stimuli. Although the selection process was random, we used the scale invariant feature transform and the dot product of normalized color histograms to eliminate from this novel set stimuli which looked either too similar to the familiar ones or to one another (see Supplemental Experimental Procedures).

We attempted to record from every well-isolated and visually responsive unit in ITC. To avoid a neuronal selection bias, the vast majority of visually responsive units (n = 40/50, 80% for monkey D; n = 35/38, 92% for monkey I) were found and isolated with an independent set of 50 initially novel stimuli that gradually became familiar as the recording sessions accumulated. Thus, the results presented here are not a consequence of selecting units that we knew ahead of time would be responsive to familiar items. All neurons reported were held for at least five repetitions of each unique stimulus, but most were held for ten (n = 46/50, 92% for monkey D; n = 35/38, 92% for monkey I).

We divided the sample of neurons into two classes based on the widths (trough-to-peak durations) of their extracellularly recorded spike waveforms. Clustering was performed with a k-means algorithm. We labeled the broad-spiking class as putative excitatory and the narrow spiking as putative inhibitory.

Although we recorded the neuronal activity in a rapid serial visual presentation paradigm to allow each one of the large number of unique stimuli to be presented many times while simultaneously maintaining single-unit isolation, the stimulus presentation durations (200 ms) and interstimulus durations (50 ms) were long enough to allow for a separate analysis of the early and late components of the neuronal response. The early phase was defined as the epoch 75–200 ms, and the late phase was defined as the epoch 200–325 ms, both relative to stimulus onset. The main firing rate metrics used throughout this study were the maximum response and the average response. The maximum response was defined as the maximum across the mean firing rates to the 125 stimuli in either the familiar or novel set. The average response was defined as the average over the mean firing rates.

To determine, for a single cell, whether the maximum response across the familiar set was significantly different from the maximum response across the novel set, we used the Mann-Whitney U test (histograms in Figures 3C and 3E). To compare statistically the average stimulus-evoked response across the 125 familiar stimuli to that across the 125 novel stimuli, we used a t test (histograms in Figures 4C and 4D). To assess whether population-averaged data were different from a null hypothesis, we applied the appropriate (paired or unpaired) t tests, always two-tailed. As a measure of selectivity, we used the sparseness metric (Olshausen and Field, 2004; Rolls and Tovee, 1995; Vinje and Gallant, 2000; Zoccolan et al., 2007). This metric takes the form $S = (1 - A)/(1 - 1/n)$, where $A = (\sum_i^n r_i/n)^2 / \sum_i^n (r_i^2/n)$, n is the number of stimuli, and r_i are the mean firing rates to a set of stimuli. S takes values between 0 and 1. We evaluated the significance of sparseness differences between the familiar and novel sets with a randomization test (histograms in Figures 5C and 5D). We also used randomization test (corrected for multiple comparisons) to determine the time points at which the sliding window firing rates from two conditions, averaged across the population of neurons, were different from one another (see Supplemental Experimental Procedures for more details on the randomization tests). To establish how well a single neuron's spike counts could discriminate between any two randomly chosen stimuli within either the familiar or novel sets, we used the AUC, which measures the discriminability of two spike count distributions (Green and Swets, 1966). In particular, we computed all pairwise AUC values in the set of 125 familiar or 125 novel stimuli, reflected about 0.5 values below 0.5 (e.g., 0.35 became 0.65), and took their average (Figure 7).

SUPPLEMENTAL INFORMATION

Supplemental Information includes six figures and Supplemental Experimental Procedures and can be found with this article online at doi:10.1016/j.neuron.2012.01.032.

ACKNOWLEDGMENTS

We wish to thank all members of the D.L.S. lab for their helpful comments and suggestions throughout the course of this experiment. We also acknowledge John Ghenne's expert animal care. This research was supported in part by NIH Grant #EY14681 (to D.L.S.), NSF Grant #SBE-0542013 (to D.L.S.), and NIH Grant #T32 EY018080-04 (to L.W.).

Accepted: January 26, 2012
Published: April 11, 2012

REFERENCES

Anderson, B., Mruczek, R.E., Kawasaki, K., and Sheinberg, D. (2008). Effects of familiarity on neural activity in monkey inferior temporal lobe. Cereb. Cortex *18*, 2540–2552.

Anzai, A., Peng, X., and Van Essen, D.C. (2007). Neurons in monkey visual area V2 encode combinations of orientations. Nat. Neurosci. *10*, 1313–1321.

Baker, C.I., Behrmann, M., and Olson, C.R. (2002). Impact of learning on representation of parts and wholes in monkey inferotemporal cortex. Nat. Neurosci. *5*, 1210–1216.

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. J. Neurophysiol. *92*, 600–608.

Bock, D.D., Lee, W.C., Kerlin, A.M., Andermann, M.L., Hood, G., Wetzel, A.W., Yurgenson, S., Soucy, E.R., Kim, H.S., and Reid, R.C. (2011). Network

anatomy and in vivo physiology of visual cortical neurons. Nature *471*, 177–182.

Brincat, S.L., and Connor, C.E. (2004). Underlying principles of visual shape selectivity in posterior inferotemporal cortex. Nat. Neurosci. *7*, 880–886.

Brincat, S.L., and Connor, C.E. (2006). Dynamic shape synthesis in posterior inferotemporal cortex. Neuron *49*, 17–24.

Cardin, J.A., Palmer, L.A., and Contreras, D. (2007). Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. J. Neurosci. *27*, 10333–10344.

Connors, B.W., and Gutnick, M.J. (1990). Intrinsic firing patterns of diverse neocortical neurons. Trends Neurosci. *13*, 99–104.

Cox, D.D., and DiCarlo, J.J. (2008). Does learned shape selectivity in inferior temporal cortex automatically generalize across retinal position? J. Neurosci. *28*, 10045–10055.

De Baene, W., Ons, B., Wagemans, J., and Vogels, R. (2008). Effects of category learning on the stimulus selectivity of macaque inferior temporal neurons. Learn. Mem. *15*, 717–727.

DiCarlo, J.J., and Cox, D.D. (2007). Untangling invariant object recognition. Trends Cogn. Sci. *11*, 333–341.

Diester, I., and Nieder, A. (2008). Complementary contributions of prefrontal neuron classes in abstract numerical categorization. J. Neurosci. *28*, 7737–7747.

Erickson, C.A., and Desimone, R. (1999). Responses of macaque perirhinal neurons during and after visual stimulus association learning. J. Neurosci. *19*, 10404–10416.

Fagiolini, M., and Hensch, T.K. (2000). Inhibitory threshold for critical-period activation in primary visual cortex. Nature *404*, 183–186.

Feldman, D.E. (2009). Synaptic mechanisms for plasticity in neocortex. Annu. Rev. Neurosci. *32*, 33–55.

Freedman, D.J., Riesenhuber, M., Poggio, T., and Miller, E.K. (2006). Experience-dependent sharpening of visual shape selectivity in inferior temporal cortex. Cereb. Cortex *16*, 1631–1644.

Frenkel, M.Y., Sawtell, N.B., Diogo, A.C., Yoon, B., Neve, R.L., and Bear, M.F. (2006). Instructive effect of visual experience in mouse visual cortex. Neuron *51*, 339–349.

Fujita, I., Tanaka, K., Ito, M., and Cheng, K. (1992). Columns for visual features of objects in monkey inferotemporal cortex. Nature *360*, 343–346.

Gallant, J.L., Braun, J., and Van Essen, D.C. (1993). Selectivity for polar, hyperbolic, and Cartesian gratings in macaque visual cortex. Science *259*, 100–103.

Gauthier, I., and Tarr, M.J. (1997). Becoming a "Greeble" expert: exploring mechanisms for face recognition. Vision Res. *37*, 1673–1682.

Green, D.M., and Swets, J.A. (1966). Signal Detection Theory and Psychophysics (New York: John Wiley and Sons, Inc).

Haider, B., Krause, M.R., Duque, A., Yu, Y., Touryan, J., Mazer, J.A., and McCormick, D.A. (2010). Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation. Neuron *65*, 107–121.

Hensch, T.K. (2005). Critical period mechanisms in developing visual cortex. Curr. Top. Dev. Biol. *69*, 215–237.

Henze, D.A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K.D., and Buzsáki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. J. Neurophysiol. *84*, 390–400.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. J. Physiol. *148*, 574–591.

Hubel, D.H., and Wiesel, T.N. (1965). Binocular interaction in striate cortex of kittens reared with artificial squint. J. Neurophysiol. *28*, 1041–1059.

Hubel, D.H., and Wiesel, T.N. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. J. Physiol. *206*, 419–436.

Hussar, C.R., and Pasternak, T. (2009). Flexibility of sensory representations in prefrontal cortex depends on cell type. Neuron *64*, 730–743.

Johnston, K., DeSouza, J.F., and Everling, S. (2009). Monkey prefrontal cortical pyramidal and putative interneurons exhibit differential patterns of activity between prosaccade and antisaccade tasks. J. Neurosci. *29*, 5516–5524.

Kawaguchi, Y., and Kubota, Y. (1997). GABAergic cell subtypes and their synaptic connections in rat frontal cortex. Cereb. Cortex *7*, 476–486.

Kerlin, A.M., Andermann, M.L., Berezovskii, V.K., and Reid, R.C. (2010). Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. Neuron *67*, 858–871.

Kobatake, E., and Tanaka, K. (1994). Neuronal selectivities to complex object features in the ventral visual pathway of the macaque cerebral cortex. J. Neurophysiol. *71*, 856–867.

Kobatake, E., Wang, G., and Tanaka, K. (1998). Effects of shape-discrimination training on the selectivity of inferotemporal cells in adult monkeys. J. Neurophysiol. *80*, 324–330.

Kuhlman, S.J., Tring, E., and Trachtenberg, J.T. (2011). Fast-spiking interneurons have an initial orientation bias that is lost with vision. Nat. Neurosci. *14*, 1121–1123.

Li, N., and DiCarlo, J.J. (2008). Unsupervised natural experience rapidly alters invariant object representation in visual cortex. Science *321*, 1502–1507.

Li, N., and DiCarlo, J.J. (2010). Unsupervised natural visual experience rapidly reshapes size-invariant object representation in inferior temporal cortex. Neuron *67*, 1062–1075.

Liu, B.H., Li, P., Li, Y.T., Sun, Y.J., Yanagawa, Y., Obata, K., Zhang, L.I., and Tao, H.W. (2009). Visual receptive field structure of cortical inhibitory neurons revealed by two-photon imaging guided recording. J. Neurosci. *29*, 10520–10532.

Logothetis, N.K., and Sheinberg, D.L. (1996). Visual object recognition. Annu. Rev. Neurosci. *19*, 577–621.

Logothetis, N.K., Pauls, J., and Poggio, T. (1995). Shape representation in the inferior temporal cortex of monkeys. Curr. Biol. *5*, 552–563.

Logothetis, N.K., Leopold, D.A., and Sheinberg, D.L. (1996). What is rivalling during binocular rivalry? Nature *380*, 621–624.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. Nat. Rev. Neurosci. *5*, 793–807.

McCormick, D.A., Connors, B.W., Lighthall, J.W., and Prince, D.A. (1985). Comparative electrophysiology of pyramidal and sparsely spiny stellate neurons of the neocortex. J. Neurophysiol. *54*, 782–806.

Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron *55*, 131–141.

Miyashita, Y. (1988). Neuronal correlate of visual associative long-term memory in the primate temporal cortex. Nature *335*, 817–820.

Miyashita, Y. (1993). Inferior temporal cortex: where visual perception meets memory. Annu. Rev. Neurosci. *16*, 245–263.

Miyashita, Y., Date, A., and Okuno, H. (1993). Configurational encoding of complex visual forms by single neurons of monkey temporal cortex. Neuropsychologia *31*, 1119–1131.

Mruczek, R.E., and Sheinberg, D.L. (2007). Context familiarity enhances target processing by inferior temporal cortex neurons. J. Neurosci. *27*, 8533–8545.

Niell, C.M., and Stryker, M.P. (2008). Highly selective receptive fields in mouse visual cortex. J. Neurosci. *28*, 7520–7536.

Niell, C.M., and Stryker, M.P. (2010). Modulation of visual responses by behavioral state in mouse visual cortex. Neuron *65*, 472–479.

Nowak, L.G., Azouz, R., Sanchez-Vives, M.V., Gray, C.M., and McCormick, D.A. (2003). Electrophysiological classes of cat primary visual cortical neurons in vivo as revealed by quantitative analyses. J. Neurophysiol. *89*, 1541–1566.

Olshausen, B.A., and Field, D.J. (2004). Sparse coding of sensory inputs. Curr. Opin. Neurobiol. *14*, 481–487.

Op de Beeck, H.P., and Baker, C.I. (2010). The neural basis of visual object learning. Trends Cogn. Sci. *14*, 22–30.

Op de Beeck, H.P., Baker, C.I., DiCarlo, J.J., and Kanwisher, N.G. (2006). Discrimination training alters object representations in human extrastriate cortex. J. Neurosci. *26*, 13025–13036.

Op de Beeck, H.P., Wagemans, J., and Vogels, R. (2007). Effects of perceptual learning in visual backward masking on the responses of macaque inferior temporal neurons. Neuroscience *145*, 775–789.

Op de Beeck, H.P., Deutsch, J.A., Vanduffel, W., Kanwisher, N.G., and DiCarlo, J.J. (2008). A stable topography of selectivity for unfamiliar shape classes in monkey inferior temporal cortex. Cereb. Cortex *18*, 1676–1694.

Pasupathy, A., and Connor, C.E. (1999). Responses to contour features in macaque area V4. J. Neurophysiol. *82*, 2490–2502.

Peters, A., and Jones, E.G. (1984). Classification of cortical neurons. In Cellular Components of the Cerebral Cortex, A. Peters and E.G. Jones, eds. (New York: Plenum), pp. 107–121.

Rainer, G., Lee, H., and Logothetis, N.K. (2004). The effect of learning on the function of monkey extrastriate visual cortex. PLoS Biol. *2*, E44.

Rolls, E.T., and Tovee, M.J. (1995). Sparseness of the neuronal representation of stimuli in the primate temporal visual cortex. J. Neurophysiol. *73*, 713–726.

Rudy, B., and McBain, C.J. (2001). Kv3 channels: voltage-gated K+ channels designed for high-frequency repetitive firing. Trends Neurosci. *24*, 517–526.

Rust, N.C., and Dicarlo, J.J. (2010). Selectivity and tolerance ("invariance") both increase as visual information propagates from cortical area V4 to IT. J. Neurosci. *30*, 12978–12995.

Sakai, K., and Miyashita, Y. (1991). Neural organization for the long-term memory of paired associates. Nature *354*, 152–155.

Sakai, K., and Miyashita, Y. (1994). Neuronal tuning to learned complex forms in vision. Neuroreport *5*, 829–832.

Sáry, G., Vogels, R., and Orban, G.A. (1993). Cue-invariant shape selectivity of macaque inferior temporal neurons. Science *260*, 995–997.

Seltzer, B., and Pandya, D.N. (1978). Afferent cortical connections and architectonics of the superior temporal sulcus and surrounding cortex in the rhesus monkey. Brain Res. *149*, 1–24.

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

Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K., and Tsumoto, T. (2007). GABAergic neurons are less selective to stimulus orientation than excitatory neurons in layer II/III of visual cortex, as revealed by in vivo functional Ca2+ imaging in transgenic mice. J. Neurosci. *27*, 2145–2149.

Southwell, D.G., Froemke, R.C., Alvarez-Buylla, A., Stryker, M.P., and Gandhi, S.P. (2010). Cortical plasticity induced by inhibitory neuron transplantation. Science *327*, 1145–1148.

Tamura, H., Kaneko, H., Kawasaki, K., and Fujita, I. (2004). Presumed inhibitory neurons in the macaque inferior temporal cortex: visual response properties and functional interactions with adjacent neurons. J. Neurophysiol. *91*, 2782–2796.

Tanaka, K. (1996). Inferotemporal cortex and object vision. Annu. Rev. Neurosci. *19*, 109–139.

Tanaka, K., Saito, H., Fukada, Y., and Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. J. Neurophysiol. *66*, 170–189.

Tolhurst, D.J., Smyth, D., and Thompson, I.D. (2009). The sparseness of neuronal responses in ferret primary visual cortex. J. Neurosci. *29*, 2355–2370.

Tsunoda, K., Yamane, Y., Nishizaki, M., and Tanifuji, M. (2001). Complex objects are represented in macaque inferotemporal cortex by the combination of feature columns. Nat. Neurosci. *4*, 832–838.

Vigneswaran, G., Kraskov, A., and Lemon, R.N. (2011). Large identified pyramidal cells in macaque motor and premotor cortex exhibit "thin spikes": implications for cell type classification. J. Neurosci. *31*, 14235–14242.

Vinje, W.E., and Gallant, J.L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision. Science *287*, 1273–1276.

Yamane, Y., Carlson, E.T., Bowman, K.C., Wang, Z., and Connor, C.E. (2008). A neural code for three-dimensional object shape in macaque inferotemporal cortex. Nat. Neurosci. *11*, 1352–1360.

Yang, T., and Maunsell, J.H. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. J. Neurosci. *24*, 1617–1626.

Zoccolan, D., Kouh, M., Poggio, T., and DiCarlo, J.J. (2007). Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. J. Neurosci. *27*, 12292–12307.