

Encoding of Partially Occluded and Occluding Objects in Primate Inferior Temporal Cortex

Tomoyuki Namima^{1,2} and  Anitha Pasupathy^{1,2}

¹Department of Biological Structure, University of Washington, Seattle, Washington 98195, and ²Washington National Primate Research Center, University of Washington, Seattle, Washington 98195

Object segmentation—the process of parsing visual scenes—is essential for object recognition and scene understanding. We investigated how responses of neurons in macaque inferior temporal (IT) cortex contribute to object segmentation under partial occlusion. Specifically, we asked whether IT responses to occluding and occluded objects are bound together as in the visual image or linearly separable reflecting their segmentation. We recorded the activity of 121 IT neurons while two male animals performed a shape discrimination task under partial occlusion. We found that for a majority (60%) of neurons, responses were enhanced by partial occlusion, but they were only weakly shape selective for the discriminanda at all levels of occlusion. Enhancement of IT responses in these neurons depended largely on the area of occlusion but only minimally on the color and shape of the occluding dots. In contrast to the above group of neurons, a sizable minority responded best to the unoccluded stimulus and showed strong selectivity for the shape of the discriminanda. In these neurons, response magnitude and shape selectivity declined with increasing levels of occlusion. Simulations revealed that the response characteristics of both classes of neurons were consistent with a model in which the responses to the occluded shape and the occluders are weighted separately and linearly combined. Overall, our results support the hypothesis that information about occluded and occluding stimuli are linearly separable and easily decodable from IT responses and that IT neurons encode a segmented representation of the visual scene.

Key words: macaque monkey; object recognition; shape discrimination; shape representation; ventral visual pathway; visual cortex

Significance Statement

Recognizing partially occluded objects can be challenging, yet the primate visual system achieves it rapidly and effortlessly. For successful recognition in the face of occlusion, segmentation of the occluded and occluding objects is a critical first step. Using a combination of experimental data and simulations, here we demonstrate that responses of neurons in macaque IT cortex, the highest stage of the form processing pathway, reflect occluded and occluding stimuli as segmented components and are not bound together as they appear in the visual image. These results support the idea that segmentation and perception of occluded and occluding stimuli are directly mirrored in the responses of neurons in the highest form processing stages.

Received Nov. 25, 2020; revised Apr. 27, 2021; accepted May 5, 2021.

Author contributions: T.N. and A.P. designed research; T.N. performed research; T.N. analyzed data; and T.N. and A.P. wrote the paper.

This work was supported by the National Institutes of Health (National Eye Institute Grant R01EY018839 to A.P., Vision Core Grant P30EY01730 to the University of Washington, and Office of Research Infrastructure Programs Grant OD010425 to the Washington National Primate Research Center). We thank Taekjun Kim, Dina V. Popovkina, and Dean Pospisil for discussions and comments, and Amber Fyall, Zachary Lindbloom-Brown, and the Instrumentation Services Core at the Washington National Primate Research Center for technical support.

The authors declare no competing financial interests.

Correspondence should be addressed to Anitha Pasupathy at pasupat@u.washington.edu.

<https://doi.org/10.1523/JNEUROSCI.2992-20.2021>

Copyright © 2021 the authors

Introduction

In natural vision, partial occlusions are pervasive, and they affect the visual image in two fundamental ways. First, some features of the occluded object may be hidden behind occluders, making object recognition difficult. Second, the bounding contours of occluded and occluding objects meet, making segmentation and boundary assignment difficult. Despite these challenges, the primate brain appears to effortlessly navigate object recognition under occlusion. To gain further insights into the neural basis of this capacity, here we focus on the inferior temporal (IT) cortex, the final stage of object processing along the ventral visual pathway in the nonhuman primate (Holmes and Gross, 1984; Horel et al., 1987) and investigate how occluded and occluding objects are encoded when animals perform a shape discrimination task.

When objects are partially occluded, information about the occluded and occluding objects are spatially bound together in the visual image. Perceptually, however, the component objects become segregated. In this study, we want to determine how the information about occluding and occluded objects are encoded in the responses of IT neurons, that is, whether the information is linearly separable and easily decodable (thereby consistent with perception), or whether the information is bound together as in the visual image.

Much of what we know about encoding in IT cortex has come from experiments in which isolated shapes or objects are presented as visual stimuli. A few studies have investigated how multiple simultaneous objects may be encoded (Miller et al., 1993; Zoccolan et al., 2005, 2007; McMahan and Olson, 2009; Bao and Tsao, 2018), but these studies have used nonoverlapping stimuli where object segmentation can be trivial, and recognition can proceed unhindered by occlusions. In one previous study, Kovacs et al. (1995) explored IT representations under occlusion in passively fixating animals and demonstrated that the responses of many IT neurons declined in magnitude, whereas a few showed enhanced responses when simple geometric shapes were subjected to partial occlusion. Missal et al. (1997, 1999) demonstrated that responses to a foreground shape were also suppressed in the presence of a larger background stimulus. To build on these previous studies and develop a quantitative model for how signals from multiple overlapping stimuli are multiplexed in IT neurons, we targeted brain regions in anterior IT cortex. We recorded the activity of single neurons to unoccluded and partially occluded stimuli as macaque monkeys were engaged in a shape discrimination task. As in previous studies in V4 and prefrontal cortex (Kosai et al., 2014; Fyall et al., 2017), the level of occlusion was systematically titrated by varying the diameter of randomly positioned occluders. Across neurons, we studied how responses to preferred, less preferred, and nonpreferred stimuli were modulated by different levels of occlusion, occluder colors, and occluder shapes. Overall, our results suggest that responses in anterior IT cortex to partially occluded stimuli are influenced by two factors—one that reflects the shape of the occluded stimulus and another that reflects the area of the occluding dots but not its shape or color. Model simulations suggest that the observed IT responses can be simply described in terms of an additive model.

Materials and Methods

Experimental subjects

Two adult male rhesus macaques (*Macaca mulatta*), weighing 8.2 kg (Monkey M) and 13.9 kg (Monkey O) participated in these experiments. Animals were surgically implanted with custom-built head posts attached to the skull with orthopedic screws. After behavioral training, a metal recording chamber was implanted on the parietal skull surface of each monkey, followed by a craniotomy in a subsequent surgery, to allow vertical access to anterior IT cortex. All animal procedures conformed to National Institutes of Health guidelines and were approved by the Institutional Animal Care and Use Committee at the University of Washington.

Recording site and electrophysiology

Recording sites were determined based on stereotaxic coordinates and structural magnetic resonance images (MRI) collected before hardware implantation. The sites were located on the posterior bank of anterior superior temporal sulcus or the inferior temporal gyrus around anterior middle temporal sulcus (AMTS; monkey M, A13-16 and L16-19 on the left hemisphere; monkey O, A8-14 and L16-31 on the right hemisphere). Custom-built plastic grids with 1 mm spacing were placed in the metal recording chamber before every recording session. We used two grids

with 0.5 mm offset along X and Y dimensions to facilitate denser sampling of 0.7 mm. A metal guide tube was advanced within the grid hole to ~5–10 mm below superficial dura. An epoxy-coated tungsten electrode (200 μ m; FHC) or U-probe (210 μ m, 24 contacts with 100 μ m spacing; Plexon) was inserted through a metal guide tube using a hydraulic microdrive (MO-91AOil Hydraulic Micromanipulator, Narishige). The depth for recording location was identified based on the visual response and the depth profile of electrode penetrations. The depth profile was assessed with the comparison of the cortical transitions of gray and white matter with the structural MRI. Recording procedures used were adapted from Namima et al. (2014).

For data acquisition, a MAP Data Acquisition System (Plexon) or a Cerebus System (Blackrock Microsystems) was used. Voltage signals from the electrode contacts were amplified and filtered (100 Hz–8 kHz with Plexon MAP system; 250 Hz–8 kHz with Cerebus System) and stored for off-line analysis. Waveform signals from single neurons were manually isolated using an Offline Sorter (Plexon) after each recording session. After completion of recording in monkeys M and O, the electrode penetration tracks around the AMTS were visually confirmed after autopsy.

Visual stimulation

Animals were seated in front of a visual display at a distance of 57 cm. Visual stimuli were presented on a CRT monitor (1600 \times 1200 pixels, 97 Hz frame rate) calibrated with a spectrophotograph radiometer (PR-650, PhotoResearch). Stimulus onset and offset times were monitored based on photodiode detection of synchronized pulses in the lower left corner of the monitor. Animals' eye position was monitored using a 1 kHz infrared eye tracking system (Eyelink 1000, SR Research). Stimulus presentation and behavioral events were controlled by custom software written in Python (<https://github.com/mazerj/pype3>).

Experimental Design

Behavioral task. During each recording session, monkeys performed a sequential shape discrimination task (Fig. 1A) previously used in our investigations of V4 and ventrolateral prefrontal cortex (vLPFC; Kosai et al., 2014; Fyall et al., 2017). Each trial began when the animal acquired and maintained fixation on a white dot (0.1° diameter) within 1.1° of visual angle for a 200 ms duration. Then, a simple two-dimensional (2D) shape was presented at central fixation (reference) for 600 ms, followed by a 200 ms interstimulus interval (ISI). This was followed by the presentation of an unoccluded or partially occluded 2D shape stimulus (test) parafoveally at an eccentricity of ~2–9.5° for 600 ms. The animal was required to maintain fixation during the presentation of reference and test stimuli. Although the specific positioning of the test stimulus on any given day was random, the range of positions sampled matched the range in our V4 and vLPFC studies. Next, the test stimulus was extinguished, and after a brief 50 ms delay, the central fixation was replaced by two target dots on either side of the fixation spot at a distance of 6° along the horizontal meridian. Monkeys were then required to saccade to the right or left target dot to report whether the reference and test shapes were the same or different, respectively. A saccade to the correct target within 500 ms of test stimulus offset was rewarded with drops of juice. Error trials were terminated without reward.

Visual stimuli

For each recording session, a pair of discrimination stimuli was chosen from a set of simple geometric shapes previously used to study V4 (Pasupathy and Connor, 2001; Fig. 1B). For 11 recording sessions, we first screened responses to many shapes and chose two stimuli—one that evoked strong responses and a second that evoked weaker responses—based on the shape preferences of one of the recorded neurons. For the remaining 110 sessions, we chose the stimulus pairs randomly as we often recorded responses from multiple neurons simultaneously. However, this mixed strategy (some customized, some random) allowed us to gain a comprehensive understanding of how occlusion modulates IT responses for a range of stimuli—preferred, less preferred, and nonpreferred—unlike our V4 experiments where we customized stimuli for each neuron. In each session, the chosen stimuli served as reference

and test, resulting in four experimental conditions. For each recording session we also chose two colors, one for the discrimination stimuli and another for the occluding dots. For 34 of 121 recording sessions, we customized the colors to match the preferences of one of the neurons under study (Kosai et al., 2014); the discrimination stimuli were presented in a preferred color, and the occluders were presented in a nonpreferred color. For the majority of sessions (87 of 121), however, we chose two contrasting colors at random. Visual stimuli were presented against an achromatic gray background (mean luminance 8 cd/m^2). The shapes and occluders were presented at one of four mean luminances (4 cd/m^2 , 8 cd/m^2 , 12 cd/m^2 , or 18 cd/m^2) and were thus either darker, equiluminant, or brighter than the background. The reference and test stimuli were identical in size, but the size varied across days in accordance with the eccentricity of the test stimulus position consistent with our previous studies in V4 and vlPFC (Kosai et al., 2014).

To characterize how occlusion modulates IT responses, we systematically titrated the levels of occlusion by varying the diameter of the occluding dots (Fig. 1C). On each trial, 36 dots were randomly positioned within a square region that fully encompassed the discrimination stimulus. The level of occlusion was quantified as the percentage of shape stimulus area that was visible. For each recording session we included unoccluded test stimuli (100% visible area) and four or five other levels of occlusion. We typically included a higher proportion of unoccluded stimuli to keep animals motivated (median number of repetitions for unoccluded stimuli was 15). Occluder position was randomized on every trial, even on different repetitions of the same occlusion level. Kosai et al. (2014) has details on how occluders were positioned.

Control experiments

Stimulus color. To test whether the effect of partial occlusion on the IT responses depended on the specific color preferences of the neuron, we performed control experiments in which we swapped the colors for the discrimination stimuli and the occluding dots. The color-reversed control experiment was conducted on 28 neurons after the completion of the main experiment, using the exact same shapes and occlusion levels as in the main experiment.

Occluder shape. To determine how the shape of the occluding dots influenced responses, we presented elliptical occluding dots instead of the circular dots used in our main experiment above. Ellipses were created by deforming the orthogonal axes of the circular dot by a 2:0.5 ratio. The orientation of each ellipse was randomly chosen in 45° increments. Position and orientation of individual dots were shuffled across trials. For this control experiment, we tested two shapes for the occluders (circle or ellipse) and two occlusion levels (small or large; percentage of visible area = 90 or 72%). We collected data from nine neurons where we completed at least seven repetitions of each condition.

Temporal asynchrony. To determine how the responses to the occluding dots alone compared with the responses to the occluding and occluded stimulus combination, on a subset of trials for nine neurons, we presented the occluding dots and the occluded shape with a temporal asynchrony. Specifically, after the presentation of the reference stimulus and the ISI, the occluding dots were presented alone for 400 ms, and

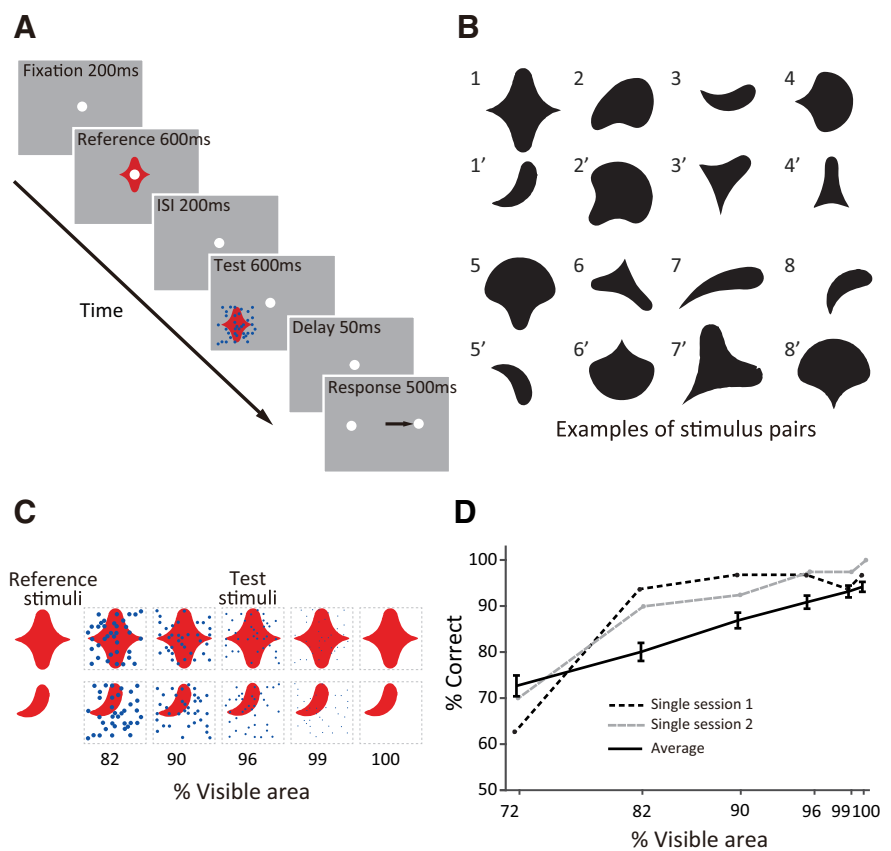


Figure 1. Behavioral task and example stimuli. **A**, Sequence of events on a single trial of the shape discrimination task. After the animal acquires fixation, unoccluded stimuli were presented during the reference epoch, and either partially occluded or unoccluded stimuli were presented during the test epoch. The animal was required to report whether the sequentially presented shapes were the same or different with a rightward or leftward saccade, respectively. **B**, Example shape pairs used as discriminanda during our recording sessions. Shapes were chosen from a simple 2D shape set (Pasupathy and Conner, 2001). Shape pairs with the same identifying integer were used in the same session. **C**, Schematic illustration of partially occluded stimuli. The visible area, that is, the unoccluded area of the occluded shapes, was titrated by varying the diameter of randomly positioned dots. **D**, Animal behavioral performance as a function of occlusion. Example single sessions (dashed lines) and average across sessions (solid line) are shown. Error bars indicate SEM.

then the shape to be discriminated appeared behind the occluders. After a further 600 ms duration, the occluders and the occluded shape were both extinguished and replaced by the target dots following a 50 ms delay. The control trials with temporal asynchrony were randomly interleaved in the main experiment but were tested only for two occlusion levels (percentage of visible area = 90 or 72%) of the five that were tested for the main experiment.

Data Analysis

Data inclusion. Spike signals and event time stamps were analyzed off-line with custom-built MATLAB code (MathWorks). Spike signals from 186 single IT neurons were recorded. Here we analyze responses of 121 neurons (M, 28; O, 93) that passed the following three criteria: behavioral performance ($>75\%$ correct on unoccluded stimuli), visual responsiveness (see below) and number of repeats (≥ 7 repetitions across all task conditions). Median number of repetitions across the 121 included neurons was 10.

Responsiveness criteria. For all analyses, we included only those neurons that showed a statistically significant response above baseline during the test epoch. To identify responsive neurons, we asked whether responses to any of the unoccluded/occluded stimuli in each recording session presented during the test epoch was significantly above baseline (two-sample *t* test, $p < 0.01$; Bonferroni corrected). Responses during the test epoch were assessed by computing the average firing rate during two windows of 50–300 ms and 50–150 ms after the test stimulus onset,

to ensure that neurons with narrow transient responses were not excluded. Baseline firing rate was assessed during the 200 ms window just before the onset of the reference stimulus. We also excluded neurons if the responses during the test epoch were not significantly greater than those during the preceding ISI (200 ms window). This criterion facilitated the rejection of neurons with lingering responses to the reference stimuli but no test stimulus-driven responses.

For all the analyses described below we used the longer test epoch (50–300 ms) for 85/121 neurons that showed statistically significant responses above baseline during this time epoch. For the remaining 36 neurons, with statistically significant responses during the 50–150 ms epoch, we used the shorter (50–150 ms) time window. Results were similar when we used the shorter or longer time window for all neurons (data not shown).

Shape selectivity. Shape selectivity was quantified by the area under the curve (AUC) of the receiver operating characteristics (ROC) curve derived from two spike count distributions for the two discrimination shapes during the test epoch. To normalize across occlusion levels, the spike count distributions were first z-scored within each occlusion level and then combined across levels (Britten et al., 1992). The shape selectivity index (AUC_{shape}) ranged from 0.5 (unselective) to 1.0 (selective), and the shape associated with the larger mean for the spike count distribution was deemed the preferred shape. We also conducted a two-way ANOVA (shape and occlusion as factors) to evaluate whether the responses to the two shape stimuli were significantly different; a significant main effect for shape or interaction ($p < 0.01$) was taken as evidence that the neuron was shape selective.

To determine how partial occlusion influenced the strength of shape selectivity for each neuron, we also computed shape selectivity based on occluded and unoccluded stimulus trials separately; we used the same methods as described above to compute AUC_{shape} but used only the occluded stimulus trials (percentage of visible area $< 100\%$; Fig. 1C) for the former and unoccluded stimulus trials (percentage of visible area, 100% ; Fig. 1C) for the latter. We used the Spearman’s rank correlation between AUC_{shape} for occluded and unoccluded stimuli across the IT population to assess the relationship between the two quantities. To confirm that our results were not dependent on our choice of metric, we also quantified shape selectivity (F_{shape}) by computing the F statistic based on a one-way ANOVA (shape as a factor) separately for unoccluded and occluded shape responses. As with the ROC analysis above, occluded responses were first z-scored within each occlusion level and then pooled across occlusion levels.

Influence of occlusion. To determine how partial occlusion modulates response magnitude, for each neuron and for each of the shapes tested (preferred and nonpreferred), we constructed two spike count distributions based on responses to occluded and unoccluded versions of the shape, and we computed the area under the ROC curve (AUC_{occl}). Here, responses to the different occlusion levels were pooled without any normalization. In this case the metric ranged from 0 to 1.0; $AUC_{occl} > 0.5$ signifies enhanced responses under occlusion, that is, stronger responses to occluded stimuli, and $AUC_{occl} < 0.5$ signifies suppression under occlusion. We also quantified preference for occlusion with a signed metric (F_{occl}) based on the ANOVA. The magnitude of F_{occl} was given by the F statistic based on a one-way ANOVA between responses to unoccluded and occluded stimuli; the sign was given by the sign of the difference between occluded and unoccluded responses. Thus, $F_{occl} > 0$ signifies that responses to occluded stimuli were stronger and $F_{occl} < 0$ signifies that responses to unoccluded stimuli were stronger; the magnitude of F_{occl} captures the strength of difference in responses between the two classes of stimuli. We also used the two-way ANOVA to evaluate whether response modulation by occlusion was statistically significant ($p < 0.01$, main effect of occlusion or interaction).

Simulations. We considered two models (segmented and unsegmented, see below, Results) for how responses to the shape and occluders may be combined. For simulations based on both models, simulated curves of mean responses to the unoccluded shape (S_{shape} ;

$shape = 1, \dots, 501$) were given by a Gaussian function f , of mean $\mu = 0$ and SD $\sigma_s = 100$ as follows:

$$R_{shape} = S_{shape} + E, \quad (1)$$

$$S_{shape} = 100 \times f(shape|\mu, \sigma_s), \text{ for } 1 \leq shape \leq 501 (\mu = 0, \sigma_s = 100) \quad (2)$$

$$E = f(0, \sigma_e). \quad (3)$$

Here, μ and σ_s define the tuning peak position and sharpness of the simulated response curve. We chose a bell-shaped Gaussian for describing the tuning curve based on the reasoning that similar shapes would evoke similar responses from IT neurons. The choice of number of shapes (501) and the SD of the Gaussian ($\sigma_s = 100$) were arbitrary, but this does not influence our results. Trial-to-trial variability for spike counts, E , was normally distributed with mean 0 and SD σ_e . We performed simulation runs at four different SD values ($\sigma_e = 0.025, 0.05, 0.1, 0.2$) to assess how variability of signal-to-noise in simulated responses affect our simulations (see below, Results).

For the segmented model, response to the combination stimulus, that is, partially occluded stimulus, was modeled as the weighted sum of the responses to the occluders and the occluded shape in isolation as in the following:

$$R_{combo} = S_{shape} \times k_{occl} + W \times R_{occluder} + E, \quad (4)$$

$$k_{occl} = 1 - A, \quad (5)$$

and

$$R_{occluder} = (1 + A). \quad (6)$$

Specifically, response to the occluded shape was the product of the responses to the same shape when unoccluded, S_{shape} , and an occlusion level, k_{occl} , which was titrated by the diameter, A , of the occluders. For these simulations, k_{occl} represents both the suppression from missing features because of occlusion and from normalization because of multiple stimuli. For each simulation run, we chose five values for the diameter of the occluding dots to produce levels of occlusion corresponding to the percentage of visible area = 99, 96, 90, 82, and 72% as in our experiments (see above, Visual stimuli; Fig. 1C). Response to the occluders, also a function of occluder diameter A , was scaled by a randomly chosen constant W in the range of $0.004 \leq W \leq 0.02$ for each simulation run. The magnitude of W determines whether a simulated neuron responds more strongly to occluded or unoccluded stimuli. We empirically chose the range for W so that a majority of simulated neurons responded more strongly to occluded stimuli as in our dataset.

For the unsegmented model, responses to the combination stimulus at each occlusion level were specified by the Gaussian function f with a randomly chosen μ and w as in the following:

$$R_{combo} = w \cdot 100 \times f(shape|\mu, \sigma_s) + E, \quad (7)$$

$$\text{for } 1 \leq shape \leq 501 (0 < w \leq 2, 0 < \mu \leq 500, \sigma_s = 100).$$

For each simulation run, w and μ were randomly chosen at each occlusion level, thus ensuring that there were no lawful relationships between responses to combination stimuli at the different occlusion levels.

On each simulation run, we constructed shape tuning curves at each of five occlusion levels (percentage of visible area = 99, 96, 90, 82, and 72%) by computing mean simulated responses for all 501 shapes across 100 repetitions for unoccluded and 20 repetitions for occluded shapes. The higher repeats for unoccluded stimuli matched the neuronal

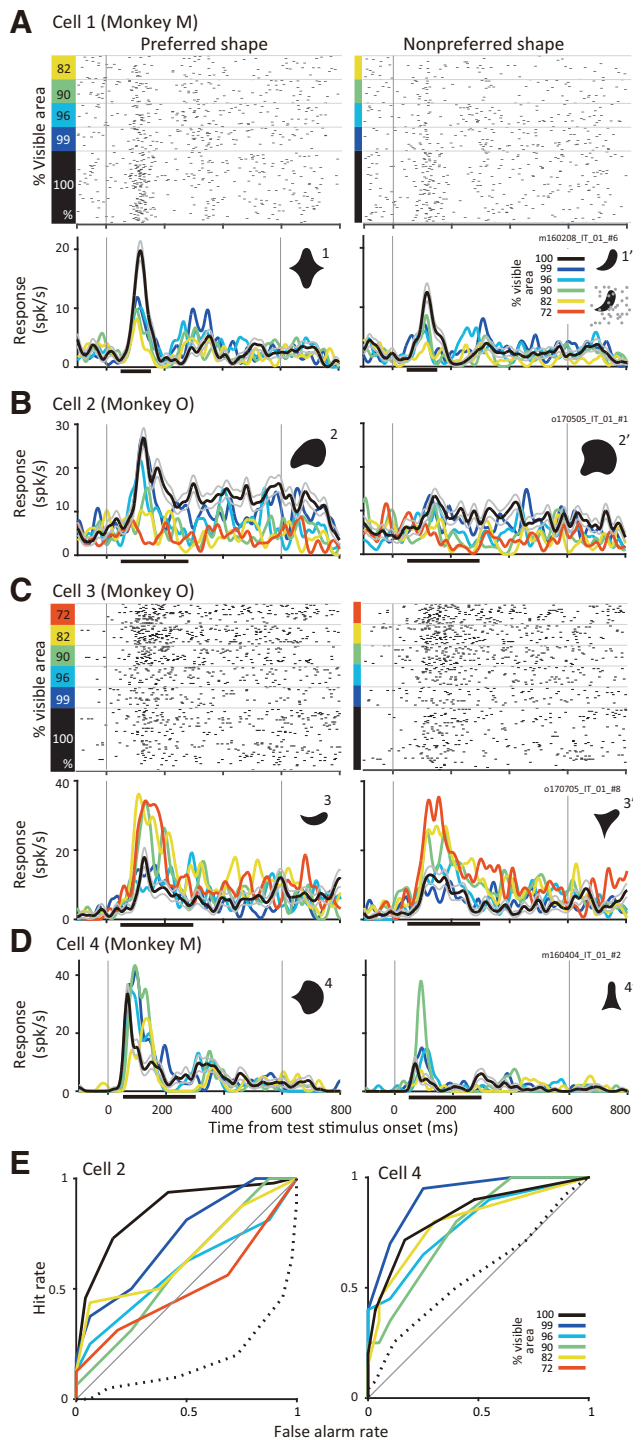


Figure 2. Responses of example IT neurons. **A**, Example neuron (Cell 1) that responds stronger to unoccluded stimuli and weaker to occluded stimuli. Responses to preferred and nonpreferred shapes (compare left and right) at different occlusion levels (colors) are shown by rasters (top) and PSTHs (bottom). PSTHs were smoothed with a Gaussian window ($\sigma = 10$ ms). Thin gray lines indicate SEM. Horizontal line along the x-axis denotes the time period for quantifying responses during the test epoch. Inset, discriminanda. **B**, Responses of another example neuron (Cell 2) to its preferred and nonpreferred shape at different levels of occlusion. Responses of this neuron also declined gradually with increasing partial occlusion. All other details as in **A**. **C**, **D**, Responses of two IT neurons, which responded preferentially to occluded stimuli. Responses to preferred and nonpreferred stimuli were enhanced when stimuli were partially occluded (compare black and colors). **E**, ROC curves for shape selectivity (solid lines) were constructed based on spike count distributions of responses to preferred and nonpreferred shapes at different occlusion levels (colors); ROC curves for occlusion preference (dotted line) were constructed based on responses to unoccluded and occluded

recordings. For the segmented model, this simply meant a random choice of W to vary how the occluders influence the neuronal response (Equation 4); for the unsegmented model, a random choice of w and μ for each occlusion level on each simulation run (Equation 7). We then chose a pair of shapes satisfying the near criterion [$\text{abs}(\text{shape1} - \text{shape2}) < 50$] or the far criterion [$\text{abs}(\text{shape1} - \text{shape2}) > 100$] for the near and far simulations, respectively (see below, Results), to see how the modulation of occlusion on the simulated responses depends on the similarity of stimuli. We then assessed AUC_{occl} and AUC_{shape} based on the responses to the chosen pair. We conducted 300 runs each, for the segmented and unsegmented models with near and far shape pairs, that is, a total of 1200 simulation runs at each of four different noise levels.

Statistical analysis

Two-way ANOVA was used to evaluate the influence of the two factors, occlusion and shape, on the responses of single neurons. An F statistic based on a one-way ANOVA was used to quantify the influence of shape or occlusion on neuronal responses. We used two-sample Kolmogorov–Smirnov tests to statistically examine individual differences between monkeys. To evaluate relationships between pairs of variables, Spearman’s rank correlation was computed. Mann–Whitney U test was conducted to evaluate whether there was a statistically significant difference in the baseline subtracted mean responses to test stimuli between match and nonmatch trials. In all statistical analysis, a p value < 0.05 was considered significant.

Results

To determine how IT neurons encode occluding and occluded objects, we studied the responses of well-isolated neurons as macaque monkeys were engaged in a sequential shape discrimination task under partial occlusion.

Experimental paradigm

On each trial, two stimuli, the reference and the test, were presented in sequence, and the animal’s task was to report whether the two shapes were the same or different by making a saccade to one of two choice targets (Fig. 1A). The test stimulus was partially occluded with a field of dots, and the level of occlusion was quantified as the percentage of the shape area that remained visible as the diameter of the dots was varied (percentage of visible area; Fig. 1C). In each session, two shapes were chosen from a standard stimulus set (Pasupathy and Connor, 2001; Fig. 1Bs) to serve as the test and reference stimuli. The two monkeys used in this study previously participated in our V4 and vlPFC studies, and as in those studies (Kosai et al., 2014, their Figs. 2–4; Fyall et al., 2017, their Fig. 1B); task performance was high for unoccluded stimuli (100% visible area) and decreased gradually as occlusion increased (Fig. 1D).

Enhancement and suppression of IT responses under occlusion

Past studies have demonstrated that partial occlusion may suppress or enhance responses of IT neurons to shape stimuli (Kovacs et al., 1995). Our results confirm these previous findings. Figure 2A,B show results from two representative IT neurons with responses that gradually declined with occlusion. Cell 1 (Fig. 2A) was one of 5 neurons recorded simultaneously during an experimental session and the shapes used as discriminanda

←

versions of the preferred stimulus. For Cell 2 (left), the ROC curve for shape selectivity gradually declined to the diagonal (gray line) consistent with a decrease in shape selectivity with occlusion; ROC curve for occlusion was below the diagonal consistent with a preference for unoccluded stimuli. For Cell 4 (right), ROC curves for shape remained above the diagonal at all levels of occlusion; ROC curve for occlusion lies just above the diagonal consistent with a mild preference for occluded stimuli.

(shapes 1 and 1'; Fig. 2A) were not customized to the preferences of this neuron. Nevertheless, this neuron responded more strongly to shape 1 than 1' at all levels of occlusion, and there was a statistically significant influence of shape on the responses of the neuron (two-way ANOVA, main effect for shape, $p < 0.001$). The level of occlusion also significantly modulated the responses of this neuron (two-way ANOVA, main effect for occlusion, $p < 0.001$); responses were strongest for the unoccluded versions of both shapes and declined gradually with occlusion. For Cell 2 (Fig. 2B) preferred and less preferred color and shape were chosen for the discriminanda. The preferred shape evoked greater responses, and the influence of occlusion was similar to Cell 1; responses were strongest for the unoccluded versions of both stimuli. A two-way ANOVA revealed statistically significant main effects for shape, occlusion, and the interaction between the two factors ($p < 0.001$). See Figure 1D (black dashed line) for the behavioral performance curve during this experimental session.

In contrast to the above examples, neurons in Figure 2C,D responded best to the occluded versions of the visual stimuli. Responses of Cell 3 were weak for the unoccluded versions of both shapes (black lines, both panels) and increased with increasing levels of occlusion. Notably, there was not much difference in the responses to the two shapes at any occlusion level. The two-way ANOVA revealed a significant influence of occlusion ($p < 0.001$) but not shape ($p = 0.50$) on the responses consistent with the hypothesis that this neuron simply carried information about the occlusion level. Corresponding behavioral curve across occlusion levels is shown in Figure 1D (gray dashed line). Cell 4 (Fig. 2D) also showed weaker responses to both shapes when they were unoccluded but responded more strongly to shape 4 than 4'. As a result, there was a statistically significant influence of shape on the responses of the neuron (two-way ANOVA, main effect for shape, $p < 0.001$). The responses of this neuron to both shapes increased with occlusion (two-way ANOVA, main effect for occlusion, $p = 0.035$), but there was no significant interaction of shape and occlusion (two-way ANOVA, $p = 0.407$).

To assess the influence of occlusion and of shape on neuronal responses across our dataset we computed two metrics. The influence of occlusion was evaluated by calculating the area under the ROC curve (AUC_{occl}) constructed from the spike count distributions based on responses to the occluded versus unoccluded stimuli (see above, Materials and Methods). ROC curves constructed from the spike count distributions for the preferred shape between occluded and unoccluded responses are shown in Figure 2E (dotted lines, Cell 2 and Cell 4). This metric ranged from 0.0 (prefers unoccluded) to 1.0 (prefers occluded). For Cells 1 and 2, which responded stronger to unoccluded stimuli, AUC_{occl} values were 0.32 and 0.17, respectively. For Cells 3 and 4, which responded stronger to occluded stimuli, they were 0.75 and 0.56, respectively. Shape preference was quantified by the area under the ROC curve (AUC_{shape}) constructed from the spike count distributions based on responses to the preferred versus the less preferred shape (see above, Materials and Methods). Examples of ROC curves constructed from the spike count distributions for the two shapes at different occlusion levels are shown in Figure 2E (colors). For Cell 2, shape selectivity gradually declined with increasing occlusion, although it remained moderately high at all occlusion levels for Cell 4. AUC_{shape} ranges from 0.5 (not shape selective) to 1 (strongly shape selective). For Cells 1–4 in Figure 2, AUC_{shape} pooled across all occlusion levels was 0.61, 0.71, 0.53 and 0.83, respectively.

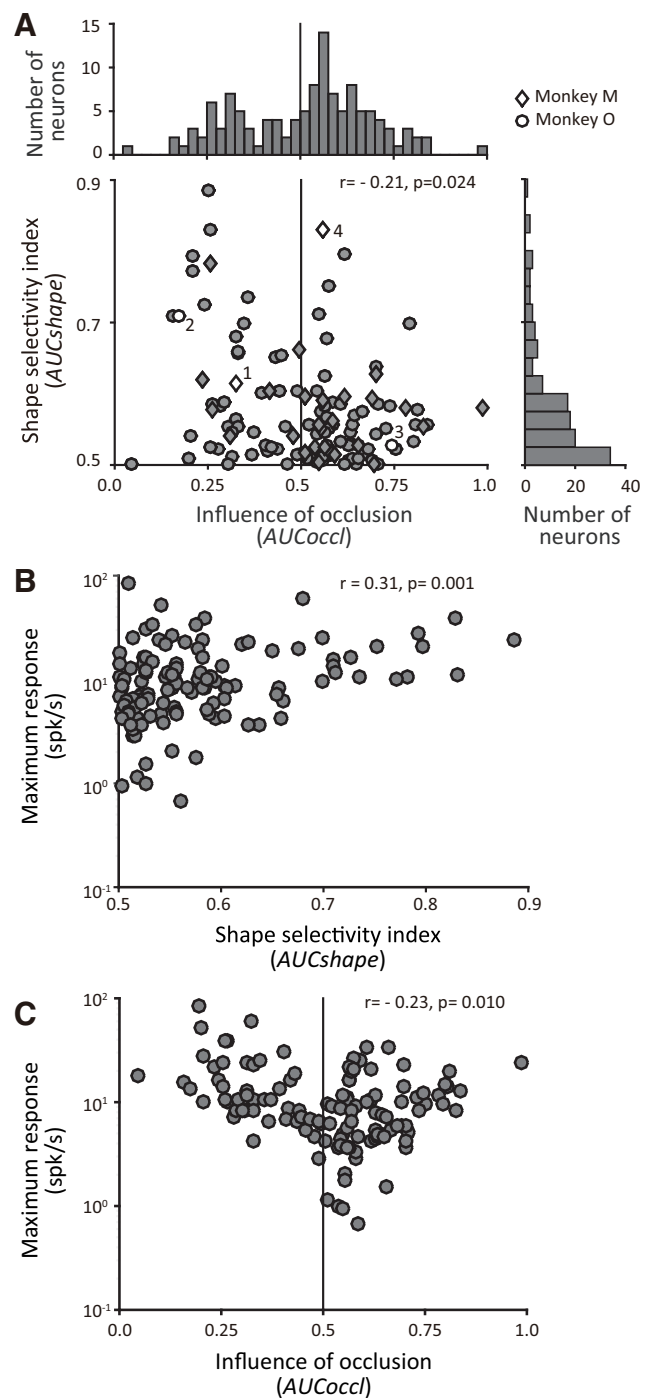


Figure 3. Influence of occlusion and shape selectivity across IT neurons. **A**, Scatter plot and marginal histograms quantify the relationship between the influence of occlusion on individual IT responses (x -axis, AUC_{occl}), and shape selectivity (y -axis, AUC_{shape}). The shape selectivity index (AUC_{shape}) ranged from 0.5 (unselective) to 1.0 (selective). The influence of occlusion (AUC_{occl}) ranged from 0 to 1.0: $AUC_{occl} > 0.5$ signifies enhanced responses under occlusion. Bin widths for vertical and horizontal bar graphs are 0.05 and 0.025. Diamond and circle markers indicate data from monkeys M and O. Data from example neurons shown in Figure 2 are denoted by open symbols and are identified by the corresponding cell number. AUC_{occl} and AUC_{shape} were computed based on the responses during the test epoch (see above, Materials and Methods). **B**, Relationship between shape selectivity (AUC_{shape}) and maximum response of individual neurons. Baseline subtracted maximum firing rates across unoccluded and occluded stimuli are shown (y -axis). **C**, Relationship between influence of occlusion (AUC_{occl}) and maximum response. To assess the influence of noise on observed trends, we computed Spearman's correlation coefficient between AUC metrics and maximum firing rate separately for odd and even trials. For AUC_{shape} versus maximum firing rate, $r_{\text{odd trials}} = 0.28$, $p = 0.002$; $r_{\text{even trials}} = 0.30$, $p = 0.000$; AUC_{occl} versus maximum firing rate, $r_{\text{odd trials}} = -0.22$, $p = 0.015$; $r_{\text{even trials}} = -0.23$, $p = 0.001$.

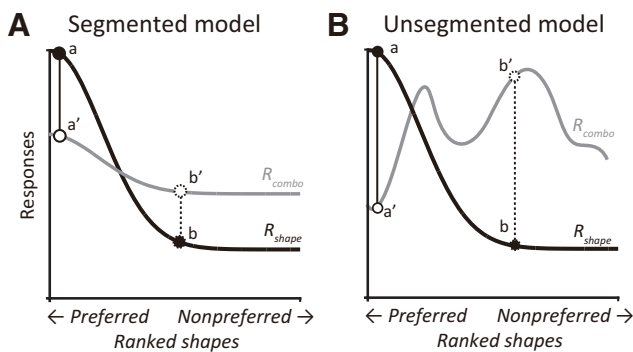


Figure 4. Schematic representation of the hypothetical models. **A**, Segmented model. Response tuning curve to unoccluded shapes (R_{shape}) is denoted by the black curve. When a shape is superimposed by an occluder, the combined response (R_{combo}) may decrease (compare a and a') or increase (compare b and b'), but the relationship between R_{shape} and R_{combo} can be described by a lawful linear equation (see above, Materials and Methods). **B**, Unsegmented model. Here too, responses may decrease (compare a and a') or increase (compare b and b') in the presence of occlusion, but no consistent lawful stimulus-agnostic relationship exists between R_{shape} and R_{combo} .

Figure 3 shows the spread of AUC_{occl} and AUC_{shape} across our dataset. In both monkeys we found some neurons with responses that declined with occlusion ($AUC_{occl} < 0.5$; monkey M = 8/28; monkey O = 40/93) and others with responses that increased with occlusion ($AUC_{occl} > 0.5$; monkey M = 20/28; monkey O = 53/93). We did not find any statistically significant difference in the distribution of AUC_{occl} between the two monkeys (two-sample Kolmogorov–Smirnov test, $p = 0.354$). We also found a broad range of shape selectivity values across neurons ($0.5 < AUC_{shape} < 0.9$).

In both animals we found that a majority of neurons (~60–70%) responded preferentially to occluded stimuli (right of vertical line in Fig. 3A). A few of these occlusion-preferring neurons were moderately shape selective with $AUC_{shape} > 0.6$ (Fig. 2, Cell 4), but the vast majority were weakly shape selective. Among neurons with $AUC_{occl} > 0.5$, the median shape selectivity was $AUC_{shape} = 0.55$. In contrast, among neurons that responded more strongly to unoccluded stimuli (left of vertical line in Fig. 3), shape selectivity was more uniformly distributed, and median shape selectivity was higher ($AUC_{shape} = 0.58$). As a result, there was a weak but statistically significant negative correlation between AUC_{occl} and AUC_{shape} (Spearman's $r = -0.21$, $p = 0.024$; Fig. 3A). Thus, consistent with the examples in Figure 2A–C, neurons that exhibited stronger shape selectivity tended to show declining responses with increasing levels of occlusion, and those with weaker shape selectivity were more likely to show increasing responses with increasing levels of occlusion. Neurons like Cell 4 in Figure 2D that were shape selective for unoccluded and occluded stimuli and showed responses increasing with occlusion were rare across our recorded population.

To consider the possibility that neurons that respond more strongly to occluded stimuli are in general weakly driven, Figure 3B,C plot baseline subtracted maximum firing rate of each neuron versus AUC_{shape} and AUC_{occl} . To minimize the influence of noise, we used a split-half analysis (Fig. 3). We found a very mild, but statistically significant correlation between maximum firing and both AUC values (for AUC_{shape} , Spearman's $r = 0.31$, $p = 0.001$; for AUC_{occl} , Spearman's $r = -0.23$, $p = 0.01$). This correlation was primarily because of a subset of weakly driven neurons that responded preferentially to occluded stimuli (right of vertical line in Fig. 3C). More broadly, our dataset

included many neurons ($n = 66$) with net maximum firing rate < 10 spikes/s. This weak magnitude of peak firing rate was not because of the size of the spike counting window, which was either 50–300 ms (for 85/121 neurons) or 50–150 ms (36/121 neurons), based on when responses were significantly above baseline (see above, Materials and Methods). It was also not because of the peripheral placement of test stimuli; maximum firing rate of neurons to reference stimuli (presented at fixation) was mildly weaker ($\Delta = 1.93$ spikes/s) than that at the peripheral test location, but this difference did not reach significance (two-sample t test $p = 0.16$). Weaker peak responses in our dataset were most likely because of the random choice of the discriminanda in our experiments (see above, Materials and Methods). This raises the possibility that many neurons may respond more strongly to unoccluded stimuli if a broader set of test stimuli was used.

It is also possible that the stronger responses to occluded stimuli in many IT neurons are the result of adaptation. Specifically, because the reference stimuli were always unoccluded, responses to unoccluded test stimuli may be weaker than occluded stimuli. To evaluate this possibility, we compared the responses to unoccluded test stimuli when the preceding reference stimulus was the same versus different. If adaptation was a significant factor, we expect weaker responses for the match versus nonmatch trials. For both preferred and nonpreferred unoccluded test stimuli, we did not find a statistically significant difference between match and nonmatch responses (Mann–Whitney U test, $n = 121$; preferred stimuli, $p = 0.873$; nonpreferred stimuli, $p = 0.916$). We also compared shape selectivity (AUC_{shape}) of IT responses between trials where the reference and test stimuli were a match versus nonmatch (and both were unoccluded stimuli). The strength of shape selectivity was not significantly different (Mann–Whitney U test, $p = 0.974$, $n = 121$).

Thus, consistent with previous studies, the results above indicate that the responses of most neurons in our dataset are dictated by one or both factors—the shape of the occluded stimulus and the level of occlusion. To determine whether the IT responses encoding the occluded shape and the occluding dots are separable and thus linearly decodable, next we considered two alternative models that embody separable and nonseparable encoding.

Hypothetical models

To delineate how the occluding and occluded stimuli dictate IT responses, we considered two models for how signals may be multiplexed. One possibility is that IT responses encode occluded and occluding object information in a joint but separable fashion, so they are linearly decodable. In this case, responses to the combined stimulus are directly related to the responses to the stimuli in isolation. Specifically, responses to the combination stimulus R_{combo} are given by the following:

$$R_{combo} = R_{shape} \times k_{occl} + W \times R_{occluder}. \quad (8)$$

The first term on the right-hand side of Equation 8 represents the responses to the occluded shape, equivalent to the unoccluded shape in isolation (R_{shape}) scaled by k_{occl} . The k_{occl} term represents the occlusion-level-dependent scaling in response, possibly caused by missing features of the occluded shape and normalization because of the presence of a second stimulus. The

second term on the right-hand side denotes the responses to the occluder presented in isolation, $R_{occluder}$, weighted by a scalar W to account for any normalization because of the presence of a second stimulus (the occluded shape). We refer to this model as the “segmented model” because the responses to the combination can be related to the stimulus responses in isolation. It is important to note that Equation 8 may take other forms; for example, R_{combo} could be given by a product instead of a sum, but the critical aspect is that the responses to the stimuli in isolation may be related directly to R_{combo} in a stimulus-agnostic fashion. This model also captures responses of neurons dictated simply by one or the other component—just the level of occlusion or the occluded shape—when either $R_{occluder}$ or R_{shape} is zero.

As an alternative to the segmented model, the occluded object and the overlapping occluders may be jointly encoded as a new complex object. In this case, R_{combo} is unrelated to R_{shape} and $R_{occluder}$ and is dictated solely by the features of the combined stimulus. We refer to this model as the “unsegmented model.” Figure 4A,B schematize the above two models. In both cases, partial occlusion may suppress or enhance responses (compare a and a’ versus b and b’) but the segmented model is based on a lawful relationship among R_{combo} , R_{shape} , and $R_{occluder}$ whereas the unsegmented model is not.

These two models can be easily distinguished, for example, on the basis of a separability index (Mazer et al., 2002) if we could study the responses of a wide variety of stimuli across many occlusion levels for each neuron. But this is not feasible given experimental constraints; our experiments only measure the responses to a pair of randomly chosen shapes at several occlusion levels. In this case, the separability index provides too lax a constraint. To determine how we might differentiate between these models based on these limited measurements, we conducted a series of simulations, which we describe next.

Simulation results

Figure 5 shows one instantiation of how an entire shape tuning curve may be modulated by different levels of partial occlusion (compare black vs other colors). For the segmented (Fig. 5A) but not the unsegmented model (Fig. 5B), there is a lawful relationship between the tuning curves across occlusion levels. Unlike in Figure 4B, we used a unimodal function for R_{combo} , but our choice embodied some key characteristics. For the unsegmented model, responses were locally smooth but globally nonmonotonic. So similar shapes (nearby on the x -axis) evoked similar responses even under occlusion, but the effect of a specific occlusion level on dissimilar shapes far apart on the x -axis was not predictable. Second, the effect of different occlusion levels was also nonmonotonic on neuronal responses of a single shape. We treated these tuning curves as a basis for the responses of IT neurons and then ran simulations where we chose a pair of shapes and measured how shape responses varied as a function of occlusion in accordance with the segmented and unsegmented models. Example ROC curves derived from simulated responses demonstrate how shape selectivity (given by the area under the ROC curve) gradually declines with increasing levels of occlusion for the segmented model (Fig. 5C), but no such lawful relationship is expected for the unsegmented model (Fig. 5D).

With these simulations we sought to determine whether we could differentiate between the two models given our experimental constraints. Specifically, our experiments measured responses of each neuron to two shapes at different levels of occlusion.

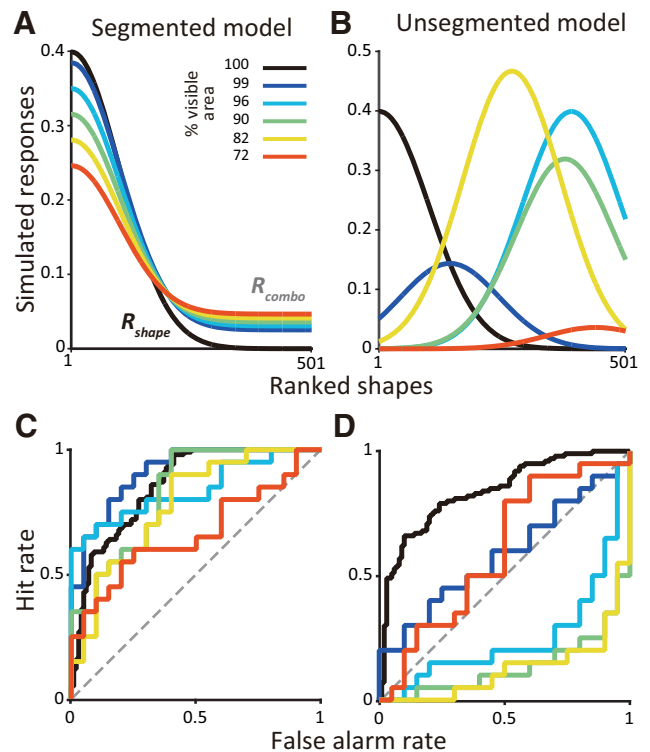


Figure 5. Examples of segmented and unsegmented model instantiations. **A**, Segmented model. Responses to unoccluded shapes (black curve) was modeled as a Gaussian function (see above, Materials and Methods). Responses to partially occluded shapes under five occlusion levels (R_{combo} , colored lines) are shown here for one simulation based on the choice of a single random weight (W ; see above, Materials and Methods). **B**, Unsegmented model. Responses to the unoccluded shapes are modeled as in **A**. For each simulation run, peak position and amplitude were randomly chosen for each occlusion level. So tuning curves across occlusion levels bear no lawful relationship. **C**, For the segmented model, the ROC curves based on spike count distributions for preferred and nonpreferred shapes reveal declining shape selectivity with increasing occlusion. **D**, For the unsegmented model, the ROC curve flips above or below the diagonal across occlusion levels, and the preferred shape and strength of shape selectivity varies dramatically across occlusion levels.

Because shapes were often chosen at random, they may be similar or different in terms of their features and evoked responses, and this could influence our ability to differentiate between the models. For example, if shapes were similar, the influence of occlusion on neuronal responses would be similar regardless of the underlying model, thus hampering our ability to differentiate between them. Our goal with these simulations then was to identify a population level metric that could differentiate between the two models regardless of whether the chosen shapes were similar or different.

To identify a diagnostic metric that could work regardless of the distance between the chosen shapes, we ran two types of simulations. In the shape-near simulation, to ensure similar responses for the chosen stimuli, we restricted the shapes to be no more than 50 shapes apart along the x -axis, that is, less than half an SD apart along the shape tuning curve of the simulated neuron (Fig. 5). In the shape-far simulations, the shapes were at least 100 shapes apart, that is, at least 1 SD apart, thus evoking disparate responses. Our results were very similar for other choices of shape-near and shape-far criteria (data not shown). From the simulated responses we calculated how shape preference (AUC_{shape}) depended on whether the two shapes were occluded and how the influence of occlusion (AUC_{occl}) depended on whether the shapes were preferred or nonpreferred.

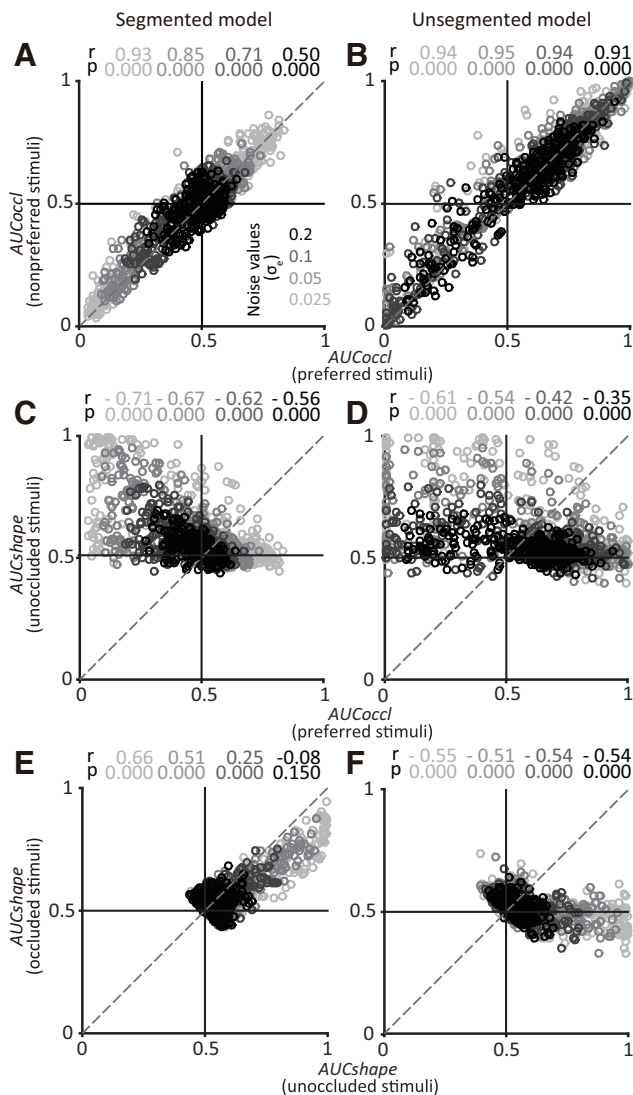


Figure 6. Results for shape-near simulations. For each simulation run we constructed shape tuning curves consistent with the segmented and unsegmented models, chose a pair of shapes <50 shapes apart, and computed AUC_{occl} and AUC_{shape} (see above, Materials and Methods). **A, B**, Relationship between the influence of occlusion (AUC_{occl}) on preferred (x-axis) and nonpreferred stimuli (y-axis) for segmented (**A**) and unsegmented (**B**) models. Each symbol corresponds to data constructed from one simulated unit ($n = 300$). Results for four different noise levels (grayscale) and corresponding Spearman's r and p values are shown. For both models, points lie along the diagonal indicating similar influence of occlusion on the responses of both shapes. **C, D**, Relationship between influence of occlusion (AUC_{occl}) on preferred stimuli (x-axis) and shape selectivity (AUC_{shape}) for unoccluded stimuli (y-axis). For both models, stronger shape selectivity (larger values along y-axis) is associated with decreased responses under occlusion (smaller values along x-axis). **E, F**, Relationship between shape selectivity (AUC_{shape}) for unoccluded (x-axis) and occluded stimuli (y-axis). Shape selectivity was positively correlated for the segmented model for all except the highest levels of noise and negatively correlated for the unsegmented model (compare **E** and **F**).

Figure 6A–F shows the results of the shape-near simulations for the segmented (A, C, E) and unsegmented models (B, D, F) at four different noise levels (grayscale). Because the preferred and nonpreferred shapes were close to each other along the x-axis, the influence of occlusion was similar (as expected) and there was a strong correlation between AUC_{occl} for the preferred and nonpreferred stimuli for the segmented (Fig. 6A) and unsegmented models (Fig. 6B). There was also a negative correlation between AUC_{shape} for the unoccluded stimuli and AUC_{occl} for

the preferred shape (Fig. 6C,D). In other words, when there was a strong preference for one of the shapes (i.e., $AUC_{shape} > 0.6$), then the responses tended to decline with occlusion. This was also consistent between the two models (Fig. 6C,D). What was strikingly different between the two models was the relationship between AUC_{shape} for occluded and unoccluded stimuli (Fig. 6E,F). For the segmented model, AUC_{shape} was positively correlated for occluded and unoccluded stimuli. Notably, the strength of the correlation declined with increasing noise, but a statistically significant positive correlation was evident for a range of noise levels. When the unoccluded responses were shape selective, that is, when $AUC_{shape} \gg 0.5$, shape selectivity declined with occlusion, that is, a majority of points are below the diagonal in Figure 6E because of the lawful relationship between shape tuning curves for occluded and unoccluded stimuli. On the other hand, for the unsegmented model there was no positive correlation between AUC_{shape} for occluded and unoccluded stimuli (in fact, the correlation was negative). The negative correlation was statically significant even at high levels of noise in the simulations. This is consistent with lack of a lawful relationship between occluded and unoccluded responses and points to a key metric that can be used to assess which model might best capture IT responses.

For the shape-far simulations (Fig. 7), the strong positive correlation between AUC_{occl} for preferred and nonpreferred stimuli disappears for both the segmented and unsegmented models (Fig. 7A,B). But a difference between the two models in terms of a positive correlation between AUC_{shape} for occluded and unoccluded stimuli persists even at the highest levels of noise (Fig. 7E,F). Thus, regardless of whether the chosen shapes are near or far, our simulations suggest that the relationship between AUC_{shape} for occluded and unoccluded stimuli across the population of IT neurons can differentiate between the underlying encoding models. Specifically, a positive correlation between AUC_{shape} for occluded and unoccluded stimuli would support a segmented model, whereas a negative correlation would be consistent with an unsegmented model.

IT population results: occlusion preference and shape selectivity

We next examined the relationship between AUC_{shape} and AUC_{occl} for preferred and nonpreferred stimuli for our IT population data (Fig. 8A,C,E). Our data reveal a strong positive correlation between AUC_{occl} for preferred and nonpreferred stimuli (Fig. 8A) and a moderate negative correlation between AUC_{shape} for unoccluded stimuli and AUC_{occl} for preferred stimuli (Fig. 8C). These results are most consistent with the shape-near simulations (with the narrow range for W) presented above (Fig. 6A,C,E). This makes sense given that the discriminanda were all 2D silhouettes and had the same luminance and chromatic contrast in any given session. A more diverse set of discriminanda may have produced results similar to the shape-far simulations. More important, we found that there was a statistically significant positive correlation between AUC_{shape} for occluded and unoccluded stimuli (Spearman's $r = 0.27$, $p = 0.003$; Fig. 8E), with AUC_{shape} for unoccluded shapes typically stronger than for occluded stimuli. Even when we considered individual occlusion levels, percentage of visible area = 99, 96, 90, and 82%, we found a statistically significant positive correlation between AUC_{shape} for unoccluded and occluded. For percentage of visible area = 72%, shape selectivity was weak and no significant correlation was detected. But even here, we did not find a negative correlation between AUC_{shape} for unoccluded

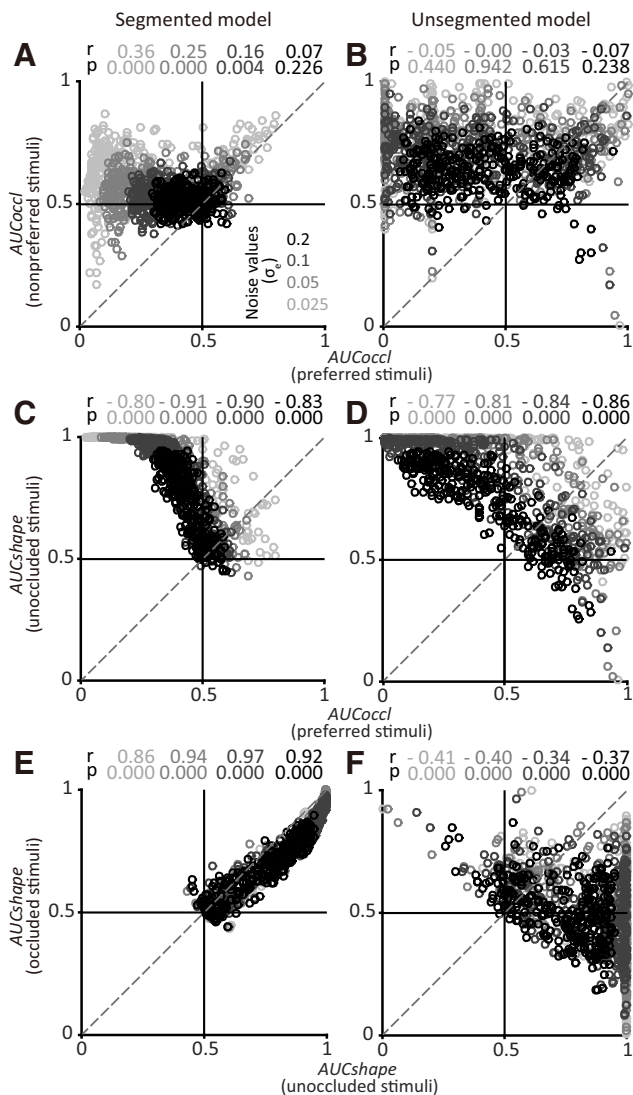


Figure 7. Population results for shape-far simulations. Simulation was performed for pairs of shapes >100 shapes apart (see above, Materials and Methods). All other details and format as in Figure 6. Results based on segmented (**A**, **C**, **E**) and unsegmented (**B**, **D**, **F**) models are shown. Unlike shape-near simulations, correlation between *AUCoccl* for preferred and nonpreferred stimuli was weak or absent for both models (compare Fig. 7*A,B* with Fig. 6*A, B*). Here again, shape selectivity between occluded and unoccluded stimuli was positively correlated only for the segmented model.

and occluded. We confirmed that these trends persisted when we used ANOVA-based metrics instead of AUC to measure shape selectivity (*Fshape*) and preference for occluded stimuli (*Foccl*; Fig. 8*B,D,F*). This pattern is most consistent with the segmented model and highly inconsistent with the unsegmented model and supports the hypothesis that IT neurons encode information about the occluders and the occluded shape independently. At one extreme (Fig. 8*C*, top left) are neurons that are highly shape selective for the unoccluded shapes, but not the occluders, and these neurons exhibit a decline in responses under occlusion. At the other extreme Fig. 8*C*, bottom right) are neurons that are not shape selective, but encode information about occluders. In between are neurons that carry information about both classes of stimuli, and their positions along the vertical and horizontal axes indicate the level of shape selectivity and the modulation by occlusion, respectively.

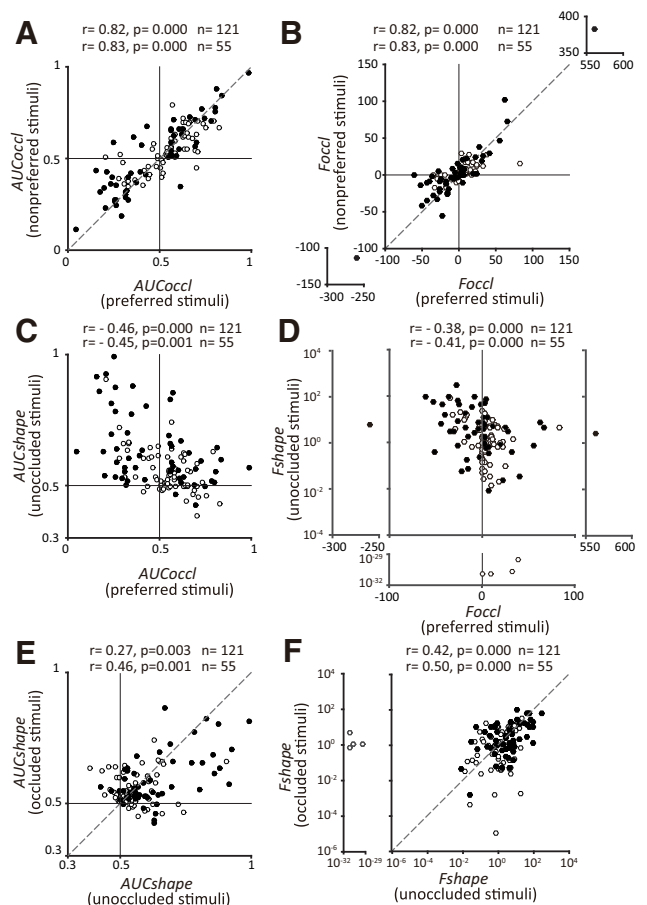


Figure 8. Population results for IT neurons. **A**, Relationship between the influence of occlusion (*AUCoccl*) on preferred (*x*-axis) and nonpreferred stimuli (*y*-axis). As with shape-near simulations of both models, a strong positive correlation was observed across the population. Filled symbols identify the more responsive subset of neurons defined as those with the baseline subtracted peak firing rate ≥ 10 spikes/s (Figure 3). Correlation values for all neurons ($n = 121$) and the more responsive subset ($n = 55$) are as shown. **B**, A positive correlation is also evident when the influence of occlusion is quantified by a signed *F* statistic (*Foccl*) based on a one-way ANOVA (see above, Materials and Methods). **C**, Influence of occlusion (*AUCoccl*) on preferred stimulus (*x*-axis) and shape selectivity (*AUCshape*) to unoccluded stimulus (*y*-axis) are related for IT neurons. Stronger shape selectivity was associated with neurons in which the responses to preferred shape declined with occlusion ($AUCoccl < 0.5$). **D**, Results were consistent when shape selectivity, and occlusion preference were measured with ANOVA-based metrics *Fshape* and *Foccl*. *Fshape* is plotted on a logarithmic scale. **E**, Shape selectivity (*AUCshape*) across the IT population for unoccluded stimuli (*x*-axis) and occluded stimuli (*y*-axis) showed a statistically significant positive correlation consistent with the segmented model. **F**, Results were consistent when shape selectivity was quantified using an ANOVA-based metric (*Fshape*). In all panels filled symbols identify the more responsive subset of neurons (**A**).

Role of stimulus color

In our task design, the occluded shape and occluders were presented in two different colors. For some sessions we customized the colors to the preferences of one of the recorded neurons, but for the majority we chose these colors at random (see above, Materials and Methods). Even so, stronger responses to the occluded stimuli could reflect a preference for the color of the occluders. Alternatively, weaker responses to occluded stimuli may reflect suppression by a nonpreferred occluder color rather than missing features from the occluded shape. This is especially possible in IT cortex, where previous studies have identified localized patches that are highly color selective (Komatsu et al., 1992; Lafer-Sousa and Conway, 2013). To test whether modulation by

occlusion relates to the color preference of neurons, we swapped the colors of occluded shape and occluding dots for 28 neurons. Figure 9A,B shows the responses of two example neurons. The neuron in Figure 9A responded more strongly to unoccluded than to occluded stimuli (compare solid and dotted lines) and also exhibited a strong preference in its responses to one of the shapes (compare left and right panels). This preference for shape and weaker responses under occlusion were observed regardless of the color of the shape or the occluder (compare blue and orange lines). In fact, for this neuron stimulus color does not appear to influence the responses at all. In contrast, responses of the neuron in Figure 9B were strongly influenced by color; responses were stronger when the shape was green and occluders were magenta. Despite this, the influence of occlusion and lack of shape preference were consistent across the color conditions; responses were always stronger for occluded than unoccluded stimuli (compare solid and dashed lines of the same color) and shape selectivity was weak (compare corresponding lines in the left and right panels). For ease of visualization, only one level of occlusion is depicted in these panels (percentage of visible area = 72%, dotted lines) but results were consistent across all occlusion levels.

The above examples are representative of our dataset. Across all neurons tested ($n = 28$), we found a strong correlation in the influence of occlusion (AUC_{occl}) in the responses to original and color-reversed stimuli (Fig. 9C): Spearman's correlation coefficient between AUC_{occl} for original and color-reversed stimuli was 0.70 ($p < 0.001$). Of the 28 neurons, only two showed a significant switch in the influence of occlusion between the responses to the original and color-reversed stimuli (data points identified with filled circle in the second quadrant; Fig. 9C). We also found no statistically significant correlation between shape selectivity (AUC_{shape}) and the relative area of the two shapes used in any experimental session, supporting the hypothesis that shape selectivity we observed was not based on the relative chromatic contrast of the two stimuli being compared (Fig. 9D). Together, these results suggest that although color can and does modulate responses of IT neurons, the preferential encoding of occluding or occluded stimuli reflects a fundamental encoding strategy that cannot be explained by the color of the stimuli.

Influence of occluder shape

Finally, we asked whether responses to occluders were dictated by the level of occlusion alone or if the shape of occluders was also critical. To address this question we targeted a subset of neurons that responded preferentially to occluded stimuli ($n = 9$) and compared responses to circular versus elliptical occluders at two occlusion levels (Fig. 10A). On a subset of trials occluders

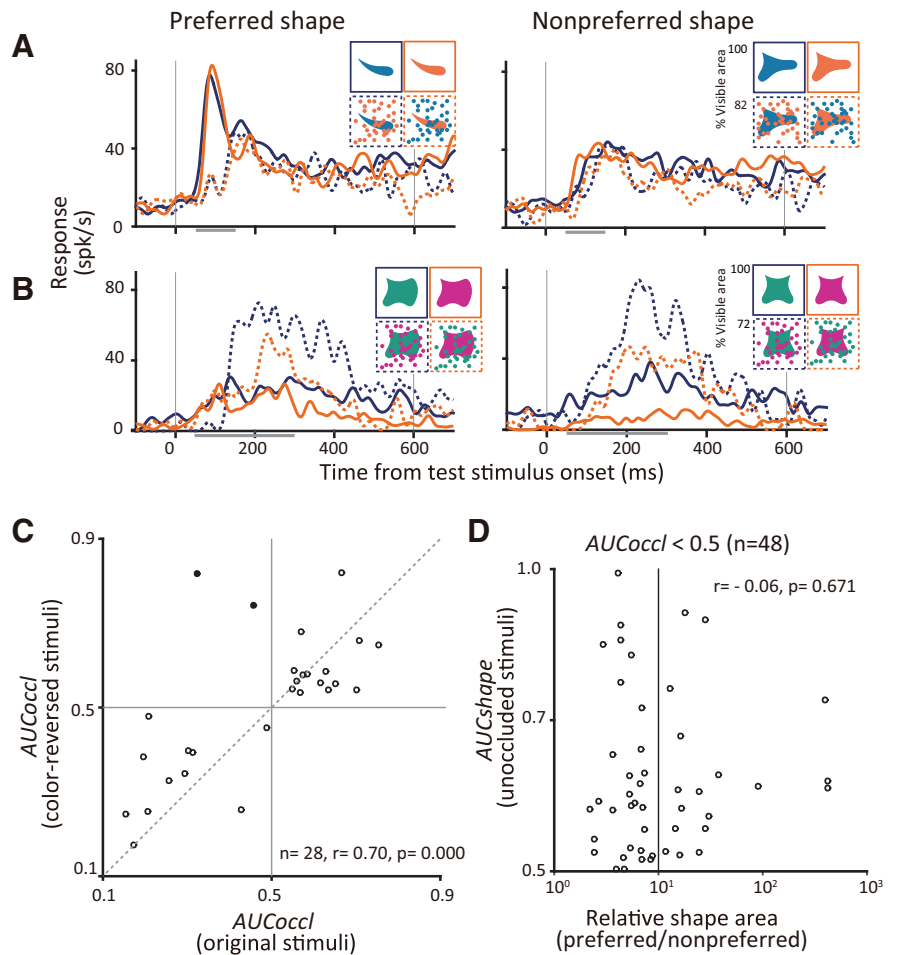


Figure 9. Effect of color on IT responses. **A, B**, Responses of two example neurons that responded preferentially to (**A**) unoccluded stimuli and (**B**) occluded stimuli. Responses to preferred (left panels) and nonpreferred (right panels) are shown for original (blue lines) and color-reversed (orange lines) stimuli. Solid and dotted lines show the responses to unoccluded (percentage of visible area = 100%) and occluded stimuli (percentage of visible area = 72%). Inset, images show corresponding visual stimuli. **C**, Influence of occlusion on original (x -axis) versus color-reversed stimuli (y -axis); $n = 28$ from monkey **D**. Spearman's correlation coefficient (r) between AUC_{occl} for original and color-reversed was $r = 0.70$ ($p < 0.001$). Filled dots (in the second quadrant) indicate neurons that showed a significant switch in the influence of occlusion with color reversal. **D**, To determine whether stronger responses for an unoccluded shape relate to the area of preferred color, shape selectivity (AUC_{shape}) is plotted as a function of the relative area of the shape stimuli (area of preferred stimulus/area of non-preferred stimulus). Only neurons that responded strongly to unoccluded stimuli are included ($n = 48$ from two monkeys). No correlation was observed (Spearman's correlation, $r = -0.06$, $p = 0.671$). The x -axis is plotted on a logarithmic scale.

were first presented alone, and the occluded stimulus was then turned on (Fig. 10B; see above, Materials and Methods). For all except one neuron (gray symbol), responses were similar for occluders with or without the occluded stimulus in the background (Fig. 10C). This was the case even when the unoccluded test stimulus evoked a statistically significant response above baseline when presented in isolation (eight of nine neurons; Figs. 2C, 10B). Thus, for this group of neurons, responses to the occluded stimuli were primarily dictated by occluders. The neuron denoted by the gray symbol was unusual in that it was both shape selective and responded preferentially to occluded stimuli; responses to occluders alone in this case were weaker than the responses to the occluded stimuli. We also found that responses were similar for circular and elliptical occluders for all except one neuron (Fig. 10D). To determine whether the neuronal activity reflects the occlusion level regardless of the shape of the occluders, for each neuron we constructed spike count distributions based on responses to circle and ellipse occluders and quantified the discriminability among the different occlusion

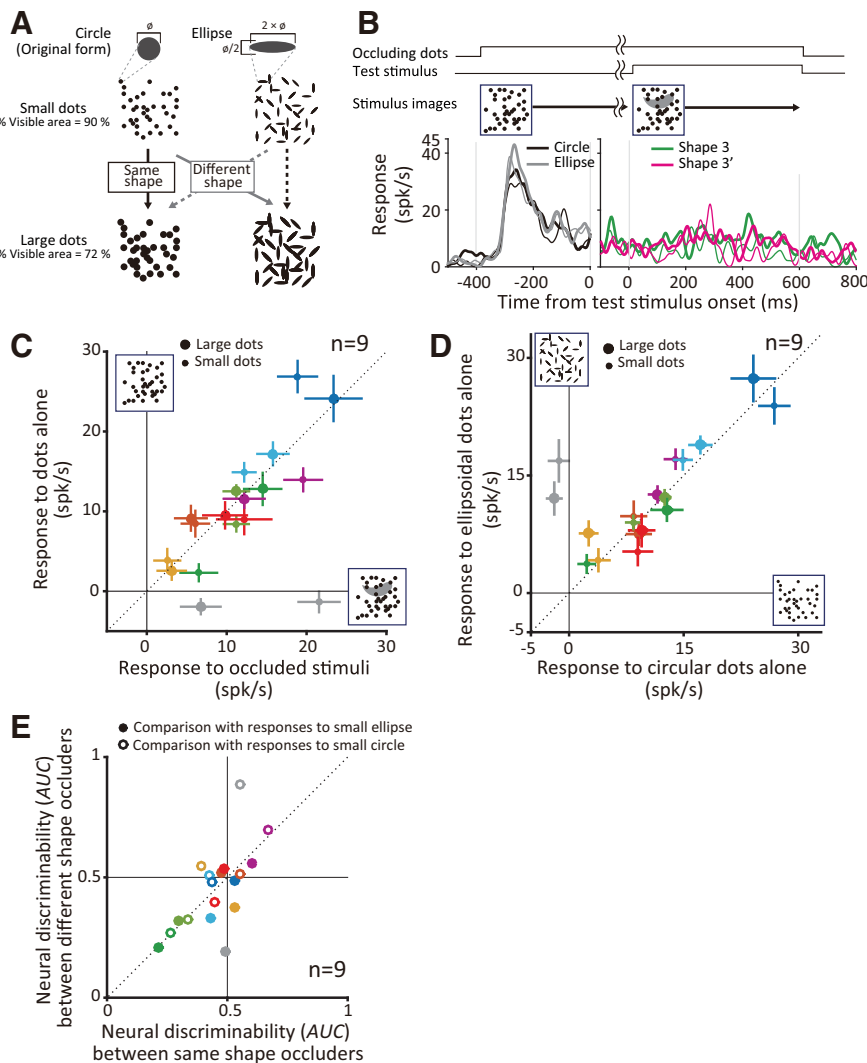


Figure 10. Influence of occluder shape on neuronal responses. **A**, Schematics of control stimuli. Ellipsoidal dots were presented at two different levels (percentage of visible areas, 72% and 90%) in addition to circular dots used in the main experiment. Major and orthogonal minor axes relate to the dot diameter as shown. The orientation of each ellipse was randomly chosen in steps of 45°. **B**, Schematic of temporal asynchrony (top) and example responses to circular and ellipsoidal occluding dots alone (compare black and gray in left PSTHs). Occluding dots were presented alone first followed by the occluded shape behind the occluders (right panel). Green and magenta lines indicate the responses to two shapes that were partially occluded with the preceding dots. Thick and thin lines in left and right panels denote; large occluding dots (percentage of visible area = 72%) and small occluding dots (percentage of visible area = 90%), respectively. Example neuron shown here is identical to Cell 3 in Figure 2C. **C**, Influence of temporal asynchrony on neuronal responses. Relationship between the responses to occluding dots with (x -axis) and without (y -axis) the occluded objects. Error bar indicates SEM. Large and small markers indicate responses to large and small occluders, respectively. Color identifies different neurons. **D**, Influence of occluder shape on neuronal responses. Scatter plot shows the relationship between the responses to circular dots alone (x -axis) and the responses to ellipsoidal dots alone (y -axis). The color code and format as in **C**. **E**, Comparison of response modulation by occluder size and shape. Open circles plot neural discriminability (AUC) between responses to small and large circles (x -axis) versus small circles and large ellipses (y -axis). Filled symbols plot AUC between responses to small and large ellipses (x -axis) versus small ellipses and large circles (y -axis). Points fall mainly along the diagonal suggesting that responses are primarily modulated by occluder area.

levels for same-shape occluders (Fig. 10E, x -axis) or different shape occluders (y -axis). The discriminability across occlusion levels was similar except for the neuron denoted by the gray symbol, indicating that the responses primarily reflect the occlusion level and not the shape of the occluders. Our dataset is small but the trends are highly consistent and suggest that the responses to the occluders observed in our study are dictated by the level of occlusion and not the shape or color of the occluding stimuli. Further confirmation would be needed with a larger dataset and greater diversity of stimuli.

Discussion

We investigated whether occluded and occluding stimuli are encoded in a linearly separable manner in the responses of individual IT neurons. Our results based on population metrics of our recorded neurons were highly consistent with simulated trends based on a model where the influence of occluders and the occluded shape on neuronal responses was separable. This was true both for IT neurons that were shape selective and responded preferentially to unoccluded stimuli and for those that responded preferentially to occluded stimuli and carried information about the level of occlusion. Our findings provide strong support for the hypothesis that IT responses provide a segmented representation of simple visual scenes containing partial occlusion.

Relationship to previous studies in IT cortex

Kovacs et al. (1995) previously investigated how IT neurons encode partially occluded objects during fixation. Consistent with our results, they too found that shape-selective IT neurons maintained the rank order of shape preference despite declining responses with increasing occlusion. Unlike our results, however, only a minority (~20% compared with our 60%) showed stronger responses to occluded than to unoccluded stimuli. This divergence from our results may relate to the differences in the experimental paradigm. Kovacs et al. conducted preliminary screening with the same eight unoccluded shapes that were later used for testing (Kovacs et al., 1995, their Materials and Methods). This may bias the studied population toward neurons that responded best to unoccluded stimuli, similar to what we observed in Kosai et al., (2014), where we too customized the discriminanda to the preferences of V4 neurons. More important, Kovacs et al. used occluders that filled the visual display and appeared before the occluded stimulus onset. This was unlike our experiments with small occluders that were presented synchronously with the occluded shape. Finally, adaptation from repeated stimuli is a concern with our experimental design, that is, responses to unoccluded stimuli may have been weaker in some neurons because of a preceding reference stimulus that was identical, although our statistical analyses confirm that there was no significant influence of adaptation at the population level.

The greater proportion of occlusion-preferring neurons in our dataset may also be a local trend. Our electrodes may have targeted an IT region with a majority of neurons selective for

textures (Tanaka et al., 1991; Komatsu and Ideura, 1993; Liu et al., 2004), and targeting a different chunk of IT cortex may reveal a greater preference for unoccluded stimuli as modularity in stimulus preferences is a documented feature of IT cortex (Kiani et al., 2007; Sato et al., 2013). Additional experiments would be needed to determine whether the greater proportion of occlusion-preferring neurons is a local or global trend. Regardless, the preferential encoding of occluded stimuli that we observed supports the previously proposed hypothesis that IT cortex plays a critical role in encoding scene context (Verhoef et al., 2015). Using fMRI in monkeys, Verhoef et al. demonstrated preferential activity in anterior IT cortex for scenes that included depth structure (context of near/far). Likewise, our control experiments support the idea that the occlusion-preferring neurons may encode the occlusion level, possibly in terms of the texture characteristics of the field of occluding dots rather than the specific color or shape of the occluders. This singular encoding of occlusion context is distinct from the multiplexed encoding of occluded shape and occlusion level evident in downstream vIPFC (Fyall et al., 2017) and may be important for the development of the latter signal (see below).

When multiple nonoverlapping stimuli are simultaneously presented in an IT receptive field, the evoked response may be modeled by a linear-weighted sum of the responses to the component stimuli (Zoccolan et al., 2005; McMahon and Olson, 2009; Sripathi and Olson, 2010; Bao and Tsao, 2018). Our analyses extend this model to overlapping stimuli. Our simulations suggest that IT responses are consistent with a linear-weighted sum of two components—the responses to the occluded shape and the occluders. Future studies should determine whether this model generalizes across a greater diversity of occluded and occluding stimuli and for more naturalistic occlusion.

V4-IT-PFC network for discriminating partially occluded stimuli

Our working hypothesis is that recurrent interactions between areas V4, IT cortex, and the vIPFC are critical to facilitate recognition under occlusion. This is supported by demonstrations of poor performance by feedforward models at recognizing partially occluded objects (Wyatte et al., 2012; Pepik et al., 2015), which improve with the inclusion of recurrent processing (O'Reilly et al., 2013; Tang et al., 2014), and the prevailing wisdom that critical computations of object recognition are mediated by feedback signals (Yuille and Kersten, 2006; Rust and Stocker, 2010; Kriegeskorte, 2015; Tang and Kreiman, 2017). We previously used the same behavioral task, stimuli, and animal subjects to probe the responses of neurons in V4 and the vIPFC (Kosai et al., 2014; Fyall et al., 2017). Together with those previous studies, our results here indicate that V4, IT, and vIPFC exhibit distinct response patterns that point to complementary roles in processing occluded shapes. In V4 the vast majority of neurons exhibit

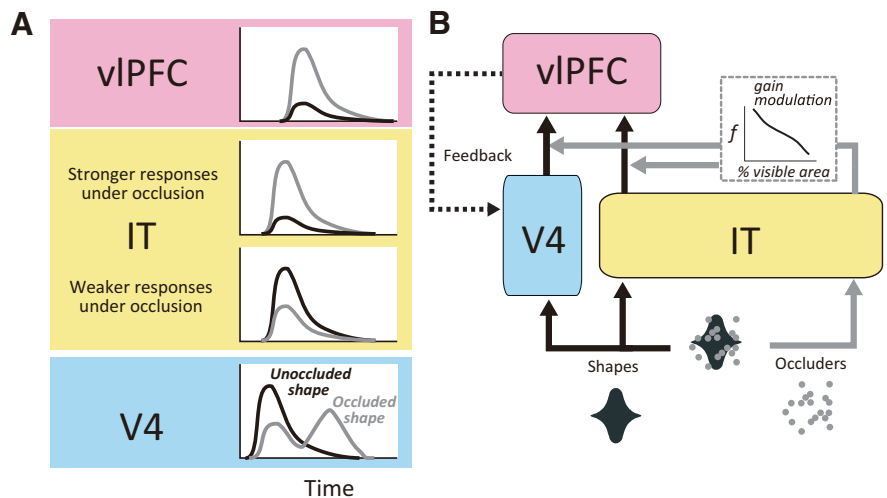


Figure 11. *A, B*, Hypothetical model of V4-IT-vIPFC interactions. Information flow in V4-IT-vIPFC circuit (*B*) and the associated response patterns in the different areas (*A*) are shown. Feedforward signals confer shape selective signals to V4 and IT (black arrows), whereas the occlusion levels are conveyed to IT (gray arrows). As a result, some neurons in V4 and IT cortex respond stronger to unoccluded (black) and weaker to occluded (gray) stimuli. Some other neurons in IT cortex respond stronger to occluded stimuli (compare PSTHs in *A*). The occlusion level signals from IT cortex could modulate the gain of shape selective signals from V4 and IT that arrive in vIPFC, giving rise to the vIPFC responses that are strong and shape selective under occlusion (see PSTHs for vIPFC in *A*). Feedback signals from vIPFC to V4 give rise to a second burst of responses that are stronger under occlusion (compare gray and black PSTHs for V4). Feedforward pathways conveying bottom up shape and occluder information are indicated with black and gray solid lines, respectively. Feedback from vIPFC to V4 is indicated with a dotted line.

an initial transient response peak that is strongly shape selective for unoccluded shapes and declines in strength and selectivity with occlusion. In vIPFC responses followed the opposite trend—weakest for unoccluded stimuli but stronger and more shape selective when partially occluded. Although some IT neurons were similar to those observed in V4 and some others to those in vIPFC, the majority primarily encoded the occlusion level, that is, responses increased in strength with occlusion, but they were not shape selective. Thus, a majority of IT neurons had response properties distinct from those in vIPFC and V4. Many V4 neurons (~30%) also exhibit a second transient peak that mimics vIPFC responses, that is, stronger and more shape selective under occlusion. These observations support the framework summarized in Figure 11. The initial burst of responses in V4 and IT carries information about the visual features of the occluded shape and the occluders [see peristimulus time histograms (PSTHs)]. The shape-selective signals from V4 and IT cortex and the occlusion level-dependent signals from IT cortex may multiply to give rise to the stronger and more shape-selective responses to occluded stimuli in vIPFC. These vIPFC signals may then feed back to V4 to give rise to a second burst of V4 responses with enhanced shape selectivity at intermediate levels of occlusion. This framework highlights a difference between visual and frontal representations, where the former is more closely aligned with visual stimuli, and the latter reflects task difficulty-dependent enhancement of visual information to facilitate task performance. The feedback from vIPFC may augment and clarify representations in V4 to facilitate discrimination as the second transient in V4 is associated with stronger shape selectivity (Fyall et al., 2017). Future perturbation studies will be needed to validate the proposed framework.

Segmented representation of visual scenes

Segmentation, the process of parsing a scene into component objects, is a critical aspect of scene perception. It is important not

just for object recognition but also for attentional selection (Treisman et al., 1983; Kahneman et al., 1992) and for our interactions with the world, for example, object manipulations, spatial navigation, and so on. Our results reveal that IT neurons encode a segmented representation of the component objects. We see no evidence that responses of IT neurons are selective for specific complex combinations of occluder and occluded shape information. One limitation of our study is that our results are based on a specific instantiation of the unsegmented model in which responses at different levels of occlusion were randomly chosen. In reality, this may not be the case, for example, if enhanced responses for occluded stimuli are because of features created at the junction of occluded and occluding contours. More generally, IT responses to stimuli at all levels of occlusion may be dictated by a feature encoding model that is agnostic to whether features are real or accidental. But we think this is unlikely given our prior results in V4 (Bushnell et al., 2011). It is also possible that IT neurons violate our assumption of smooth, continuous tuning, for example, at object category boundaries. Future studies are needed to update and expand our findings as more accurate encoding models of IT neurons are developed.

Building a segmented representation can be especially difficult if the visual scene includes partially occluded objects. In this case, the segmented representation is likely the culmination of computations that rely on contextual influences. Two fundamental properties of neurons in early visual cortex—collinear facilitation and flexible surround suppression—may provide critical building blocks for segmented object representations in later stages by enhancing the representation of collinear contours (Kapadia et al., 1995; Polat et al., 1998; Bakin et al., 2000; Bauer and Heinze, 2002) and contrasting texture regions (Blakemore and Tobin, 1972; Nelson and Frost, 1985; Knierim and van Essen, 1992; Kapadia et al., 1995; Levitt and Lund, 1997; Nothdurft et al., 1999; Cavanaugh et al., 2002; Coen-Cagli et al., 2015). These signals may enhance continuous contours at the expense of texture elements (Gheorghiu et al., 2014) and contribute to the emergence of object-based representations (Pasupathy et al., 2020). Thus, when there are sizable contrasts in color (as in our stimuli) or texture or long contours that bound objects, contextual modulations in early and midlevel processing stages may facilitate the enhanced representation of segmented objects in IT cortex. Future studies will need to investigate how strength of stimulus contrasts and the length of those boundaries influence the segmented representation of objects and facilitate recognition and scene understanding.

References

- Bakin JS, Nakayama K, Gilbert CD (2000) Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *J Neurosci* 20:8188–8198.
- Bao P, Tsao DY (2018) Representation of multiple objects in macaque category-selective areas. *Nature communications* 9:1774.
- Bauer R, Heinze S (2002) Contour integration in striate cortex. Classic cell responses or cooperative selection? *Exp Brain Res* 147:145–152.
- Blakemore C, Tobin EA (1972) Lateral inhibition between orientation detectors in the cat's visual cortex. *Exp Brain Res* 15:439–440.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Bushnell BN, Harding PJ, Kosai Y, Pasupathy A (2011) Partial occlusion modulates contour-based shape encoding in primate area V4. *J Neurosci* 31:4012–4024.
- Cavanaugh JR, Bair W, Movshon JA (2002) Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons. *J Neurophysiol* 88:2547–2556.
- Coen-Cagli R, Kohn A, Schwartz O (2015) Flexible gating of contextual influences in natural vision. *Nat Neurosci* 18:1648–1655.
- Fyall AM, El-Shamayleh Y, Choi H, Shea-Brown E, Pasupathy A (2017) Dynamic representation of partially occluded objects in primate prefrontal and visual cortex. *eLife* 6.
- Gheorghiu E, Kingdom FA, Petkov N (2014) Contextual modulation as de-texturizer. *Vision Res* 104:12–23.
- Holmes EJ, Gross CG (1984) Effects of inferior temporal lesions on discrimination of stimuli differing in orientation. *J Neurosci* 4:3063–3068.
- Horel JA, Pytko-Joiner DE, Voytko ML, Salsbury K (1987) The performance of visual tasks while segments of the inferotemporal cortex are suppressed by cold. *Behav Brain Res* 23:29–42.
- Kahneman D, Treisman A, Gibbs BJ (1992) The reviewing of object files: object-specific integration of information. *Cogn Psychol* 24:175–219.
- Kapadia MK, Ito M, Gilbert CD, Westheimer G (1995) Improvement in visual sensitivity by changes in local context: parallel studies in human observers and in V1 of alert monkeys. *Neuron* 15:843–856.
- Kiani R, Esteky H, Mirpour K, Tanaka K (2007) Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J Neurophysiol* 97:4296–4309.
- Knierim JJ, van Essen DC (1992) Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *J Neurophysiol* 67:961–980.
- Komatsu H, Ideura Y (1993) Relationships between color, shape, and pattern selectivities of neurons in the inferior temporal cortex of the monkey. *J Neurophysiol* 70:677–694.
- Komatsu H, Ideura Y, Kaji S, Yamane S (1992) Color selectivity of neurons in the inferior temporal cortex of the awake macaque monkey. *J Neurosci* 12:408–424.
- Kosai Y, El-Shamayleh Y, Fyall AM, Pasupathy A (2014) The role of visual area v4 in the discrimination of partially occluded shapes. *J Neurosci* 34:8570–8584.
- Kovacs G, Vogels R, Orban GA (1995) Selectivity of macaque inferior temporal neurons for partially occluded shapes. *J Neurosci* 15:1984–1997.
- Kriegeskorte N (2015) Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annu Rev Vis Sci* 1:417–446.
- Lafer-Sousa R, Conway BR (2013) Parallel, multi-stage processing of colors, faces and shapes in macaque inferior temporal cortex. *Nat Neurosci* 16:1870–1878.
- Levitt JB, Lund JS (1997) Contrast dependence of contextual effects in primate visual cortex. *Nature* 387:73–76.
- Liu Y, Vogels R, Orban GA (2004) Convergence of depth from texture and depth from disparity in macaque inferior temporal cortex. *J Neurosci* 24:3795–3800.
- Mazer JA, Vinje WE, McDermott J, Schiller PH, Gallant JL (2002) Spatial frequency and orientation tuning dynamics in area V1. *Proc Natl Acad Sci U S A* 99:1645–1650.
- McMahon DB, Olson CR (2009) Linearly additive shape and color signals in monkey inferotemporal cortex. *J Neurophysiol* 101:1867–1875.
- Miller EK, Gochin PM, Gross CG (1993) Suppression of visual responses of neurons in inferior temporal cortex of the awake macaque by addition of a second stimulus. *Brain Res* 616:25–29.
- Missal M, Vogels R, Orban GA (1997) Responses of macaque inferior temporal neurons to overlapping shapes. *Cereb Cortex* 7:758–767.
- Missal M, Vogels R, Li CY, Orban GA (1999) Shape interactions in macaque inferior temporal neurons. *J Neurophysiol* 82:131–142.
- Namima T, Yasuda M, Banno T, Okazawa G, Komatsu H (2014) Effects of luminance contrast on the color selectivity of neurons in the macaque area V4 and inferior temporal cortex. *J Neurosci* 34:14934–14947.
- Nelson JL, Frost BJ (1985) Intracortical facilitation among co-oriented, co-axially aligned simple cells in cat striate cortex. *Exp Brain Res* 61:54–61.
- Nothdurft HC, Gallant JL, Van Essen DC (1999) Response modulation by texture surround in primate area V1: correlates of “popout” under anesthesia. *Vis Neurosci* 16:15–34.
- O'Reilly RC, Wyatte D, Herd S, Mingus B, Jilk DJ (2013) Recurrent processing during object recognition. *Front Psychol* 4:124.
- Pasupathy A, Connor CE (2001) Shape representation in area V4: position-specific tuning for boundary conformation. *J Neurophysiol* 86:2505–2519.
- Pasupathy A, Popovkina DV, Kim T (2020) Visual functions of primate area V4. *Ann Rev Vision Sci* 6:363–385.

- Pepik B, Benenson R, Ritschel T, Schiele B (2015) What is holding back convnets for detection? *Lecture Notes in Computer Science*, pp 517–528. New York: Springer, Cham.
- Polat U, Mizobe K, Pettet MW, Kasamatsu T, Norcia AM (1998) Collinear stimuli regulate visual responses depending on cell's contrast threshold. *Nature* 391:580–584.
- Rust NC, Stocker AA (2010) Ambiguity and invariance: two fundamental challenges for visual processing. *Curr Opin Neurobiol* 20:382–388.
- Sato T, Uchida G, Lescroart MD, Kitazono J, Okada M, Tanifuji M (2013) Object representation in inferior temporal cortex is organized hierarchically in a mosaic-like structure. *J Neurosci* 33:16642–16656.
- Sripati AP, Olson CR (2010) Responses to compound objects in monkey inferotemporal cortex: the whole is equal to the sum of the discrete parts. *J Neurosci* 30:7948–7960.
- Tanaka K, Saito H, Fukada Y, Moriya M (1991) Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *J Neurophysiol* 66:170–189.
- Tang H, Kreiman G (2017) Recognition of occluded objects. In: *Computational and Cognitive Neuroscience of Vision* (Zhao Q, ed), Chap 3. Singapore: Springer.
- Tang H, Buia C, Madhavan R, Crone NE, Madsen JR, Anderson WS, Kreiman G (2014) Spatiotemporal dynamics underlying object completion in human ventral visual cortex. *Neuron* 83:736–748.
- Treisman A, Kahneman D, Burkell J (1983) Perceptual objects and the cost of filtering. *Perception and psychophysics* 33:527–532.
- Verhoef BE, Bohon KS, Conway BR (2015) Functional architecture for disparity in macaque inferior temporal cortex and its relationship to the architecture for faces, color, scenes, and visual field. *J Neurosci* 35:6952–6968.
- Wyatte D, Curran T, O'Reilly R (2012) The limits of feedforward vision: recurrent processing promotes robust object recognition when objects are degraded. *J Cogn Neurosci* 24:2248–2261.
- Yuille A, Kersten D (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cogn Sci* 10:301–308.
- Zoccolan D, Cox DD, DiCarlo JJ (2005) Multiple object response normalization in monkey inferotemporal cortex. *J Neurosci* 25:8150–8164.
- Zoccolan D, Kouh M, Poggio T, DiCarlo JJ (2007) Trade-off between object selectivity and tolerance in monkey inferotemporal cortex. *J Neurosci* 27:12292–12307.