Supplementary materials for: Executive control processes underlying multi-item working memory

Antonio H. Lara & Jonathan D. Wallis

Supplementary Figure 1

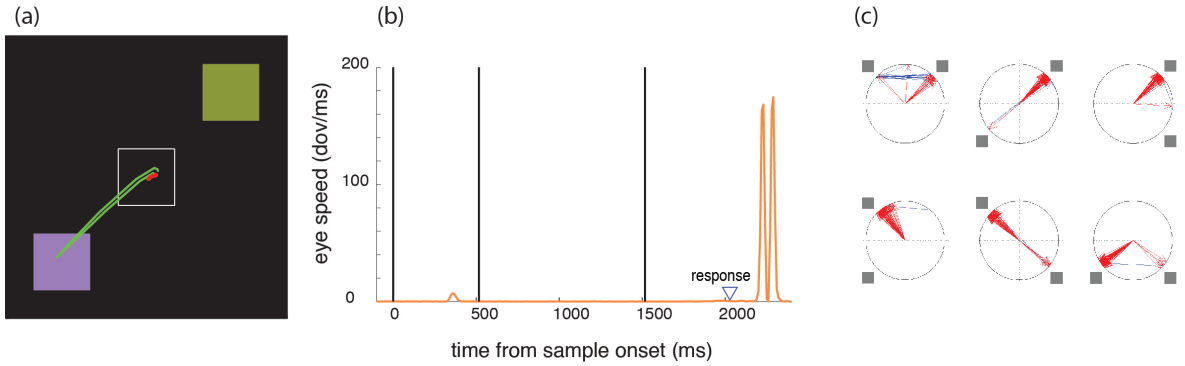

Supplementary Figure 1. Behavioral measures of covert attention. We analyzed subjects' eye movements and were able to detect microsaccades made during the sample period. We used the filtered eye speed to determine the occurrence and direction of the microsaccades. We used a threshold crossing detection algorithm that detected when the eye speed exceeded two standard deviations of the eye speed during the initial fixation (when no visual stimuli were present on the display). **(a)** Example of a single trial where items were presented in the upper right and lower left locations. The animal makes a microsaccade towards the upper right item during the sample period (red trace). For this trial, the test item appeared in the lower left and the animal made a full saccade toward that location after he was no longer required to hold fixation (green trace). **(b)** Eye speed during the course of the trial. During the sample period, there is a small, brief increase in speed corresponding to the microsaccade. The blue triangle indicates the time when the animal made his response and the large peaks in speed correspond to a saccade to the test item. **(c)** Pattern of microsaccades for all trials in this session. Each plot illustrates a specific two-item configuration, with the position of the items indicated by the gray squares. Red arrows indicate the angular direction of the first microsaccade. The blue lines indicate the angular direction of the second microsaccade (if there was one). There was a clear pattern evident. The subject favored microsaccades to the top left item. If no item was present in the top left, he favored microsaccades to the top right, and if there no items on the top of the screen he favored microsaccades to the bottom left. Such a pattern was consistent with the subject covertly attending to the items, starting with the item in the top right and moving counter-clockwise through the items. Similar stereotyped patterns of covert attention have also been reported in monkeys performing visual search tasks (Buschman & Miller, 2009). Note that microsaccades were only made during the sample epoch; they were absent during the delay. Thus, we cannot be sure how attention was allocated during the delay period.

Supplementary Figure 2

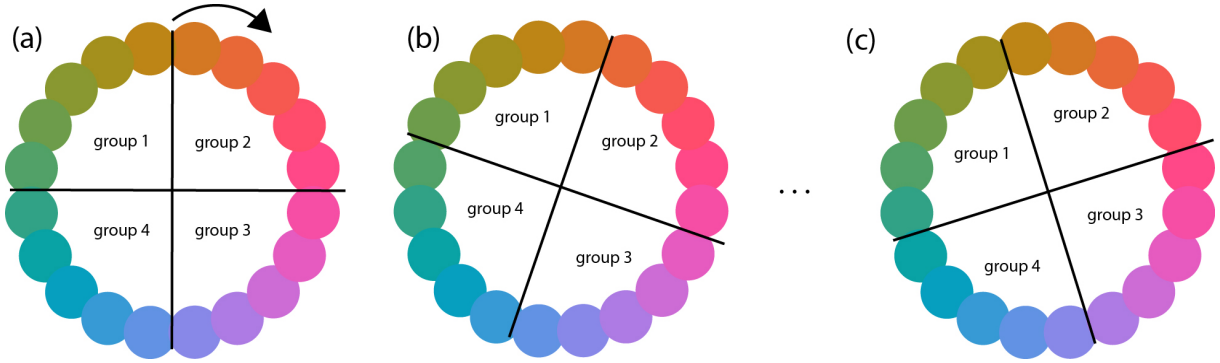

Supplementary Figure 2. Color grouping procedure. In order to give the sliding-ANOVA analysis the greatest chance of detecting color-related information, we grouped the 20 colors in the four groups of five colors per group. For each neuron, we grouped colors in four groups of five colors per group **(a)** and calculated the amount of information encoded by the neuron using the percentage of explained variance (PEV) measure. **(b)** We changed the colors that comprised each group by rotating the group boundaries clockwise by one color and calculated the PEV once more. **(c)** We repeated this procedure until we had traversed all the color space. The false discovery rate for this procedure, calculated by assessing the proportion of neurons that reached the criterion during the fixation period, was 4.9%.

Supplementary Figure 3

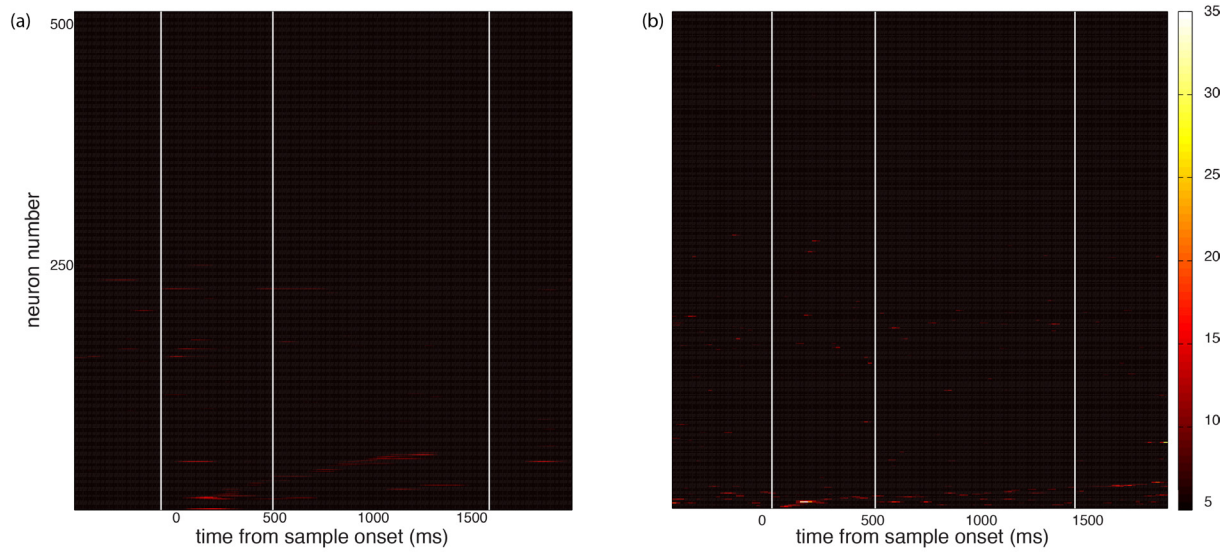

Supplementary Figure 3. Additional color selectivity measures. **(a)** To examine whether neural responses to colors could be fit by a von Mises distribution we performed a linear regression that was modified for circular data: $r = \beta_0 + \beta_1 \sin\theta + \beta_2 \cos\theta$, where r is the neuron's firing rate, and θ is the angle of the color in our color space (Figure 1). The plot shows the percentage of variance in neuronal firing rates that could be explained by the model for each neuron at each point in the trial. Each row represents the selectivity a single neuron and neurons are sorted based on the time they show significant color selectivity. Vertical white lines indicate the onset of the sample stimulus, the beginning of the delay and the onset of the test stimulus. We defined a neuron to be color tuned if the regression model significantly fit the data for three consecutive time bins. Using this criterion we found that only 13% (65/507) of neurons were color-tuned. Similar to our sliding PEV analysis, in order to correct for multiple comparisons, we calculated the number of neurons for which the p-value of the regression was larger than 0.005 for three consecutive time bins during the fixation period. This resulted in a false discovery rate of 4.3%. **(b)** To examine whether the incidence of color selectivity depended on the spatial tuning of the neuron, we restricted our sliding ANOVA analysis (as described in Supplementary Figure 2) to trials in which the color was shown in each neuron's preferred location. Color tuning remained evident in only a small minority of neurons (45/507 or 9% of neurons). Thus, despite using a variety of analysis methods, PFC neurons consistently showed little encoding of color.

Supplementary Figure 4

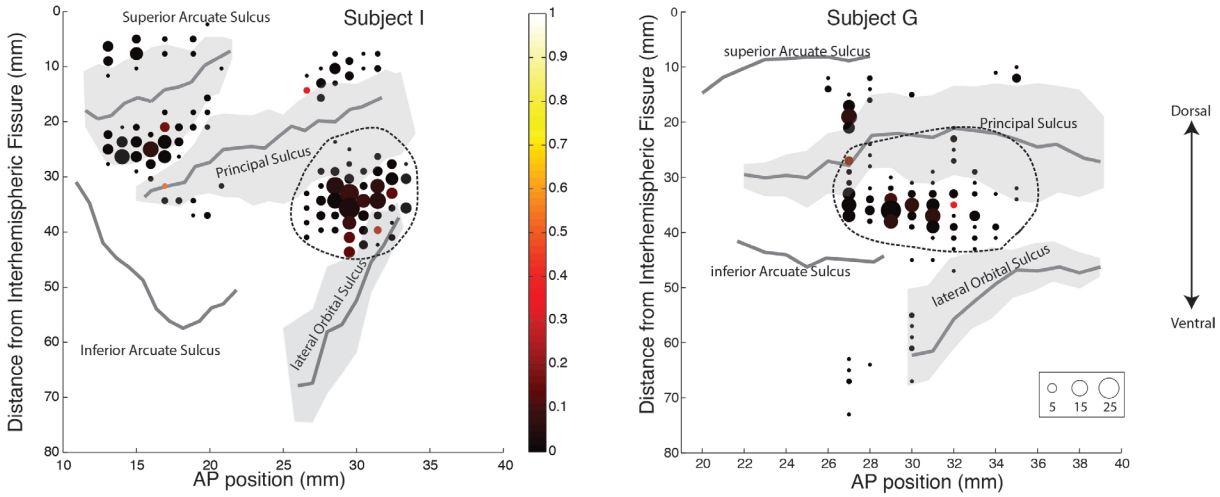

Supplementary Figure 4. Recording Locations. To ensure that the lack of color encoding in PFC was not due to the fact that we had confined our recordings to the ventrolateral PFC, we extended our recordings in to the dorsolateral prefrontal cortex, orbitofrontal cortex and the frontal eye fields. The plots show flattened cortical representations illustrating recording locations from the two subjects, and the proportion of color-selective neurons in each location (as indicated by the color of each data point). Gray shading indicates the position of a sulcus. The anterior–posterior (AP) position is measured relative to the interaural line. The circle dashed line demarcates the recording locations of the neurons in the main analysis. Despite recording an additional 323 neurons from these areas, we did not see a significant number of color-selective neurons.

Supplementary Figure 5

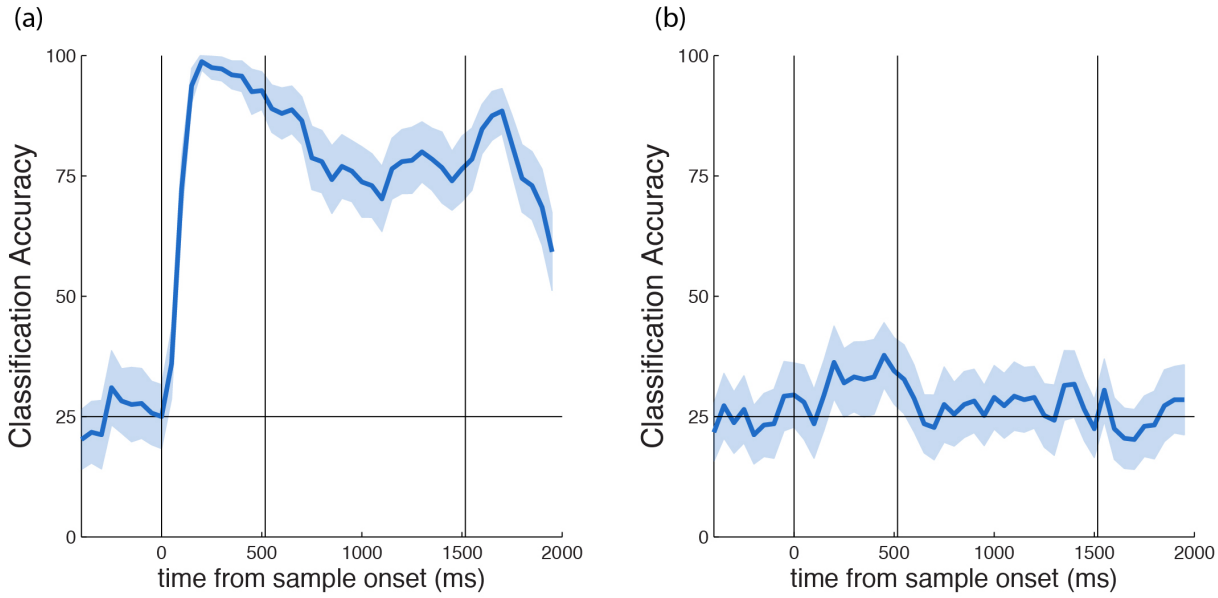

Supplementary Figure 5. Decoding of spatial and color information using a linear classifier. We used a neural-population decoding algorithm (Meyers et al, 2008; Meyers et al, 2012) that uses a maximum correlation classifier to classify multivariate patterns of activity across the population. By determining how accurately the classifier predicts which experimental conditions are present, one can determine how much information about the experimental conditions is encoded by the neuronal population. We focused our analysis on the single item trials and we decoded the location and the grouped color (see Supplementary Figure 2) of the items using the entire population of PFC neurons. The decoding algorithm used a cross-validation procedure in which the data was randomly split into 10 separate subsets. Nine of these subsets were used to train the classifier and the remaining set was used as a testing set. This cross-validation procedure was repeated 10 times to ensure that each set was used as part of the training and testing sets at least once. This generated 10 classification estimates, which were then averaged to obtain a more robust single classification estimate. This whole procedure was repeated 20 times, each time randomly selecting different subsets of training and testing sets. The results from the 20 repetitions were then averaged and are shown separately for **(a)** location and **(b)** color decoding. Shading indicates the standard error of the mean for the 20 separate repetitions of the decoding scheme. The horizontal line at 25% indicates chance performance. Vertical lines indicate the onset of the sample, delay and test epochs, respectively. The results from this decoding analysis strongly support the results from the demixed PCA analysis. We were able to decode spatial information robustly throughout the trial and at a level approaching 100% accuracy during the sample epoch (chance performance was 25%). In contrast, we could barely decode any information about color: it briefly exceeded chance during the sample epoch and returned to chance levels throughout the delay.

Supplementary Figure 6

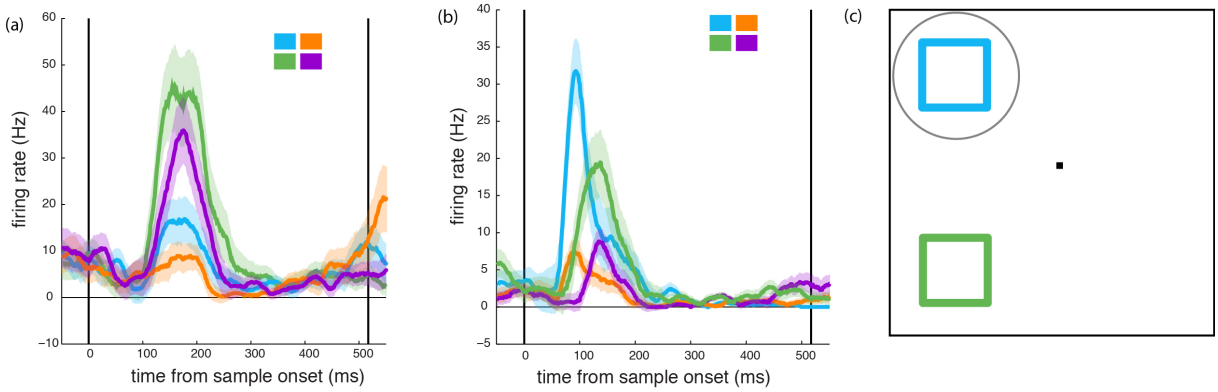

Supplementary Figure 6. Two item regression coding scheme. Illustration of how the x_3 regressor was determined. **(a)** Spike density histogram showing the firing rate of a single neuron for each of the four locations during the one-item trials. The color of the line corresponds to the location indicated by the legend. This neuron responded strongest to item in the lower left and weakest to items in the upper right. Vertical lines denote the time when the item was presented on the screen. **(b)** Response to single items presented at the four locations for a different neuron. This neuron responded more to items on the left. **(c)** Schematic of a two-item trial. In this trial items were presented in the upper and lower left. Using our behavioral estimate of attention, we determined that for this trial the subject first shifted attention to the item in the upper left (denoted by the gray circle). When these items were presented in isolation, the neuron in (a) fired with an average peak rate of 17-Hz to the item in the upper left location (blue trace) and 42-Hz to the item in the lower left (green trace). Thus, x_1 = 42 and x_2 = 17. The attended location was the neuron's less preferred location of the two and so $x_3 = x_2 - x_1 = (17 - 46) = -29$. The neuron in (b) had a peak one-item response to items on the upper left that was higher (x_1 = 19) than the one-item response to items on the lower left (x_2 = 9). Thus, attention was directed to the neuron's more preferred location, and so we would say that $x_3 = x_1 - x_2$ or (19 - 9) = 10. The upshot of this coding scheme is that the beta for x_3 will be positive when the neuron's firing rate correlates with the location indicated by our behavioral indices of covert attention and negative when the neuron's firing rate correlates with the alternative location.

Supplementary Figure 7

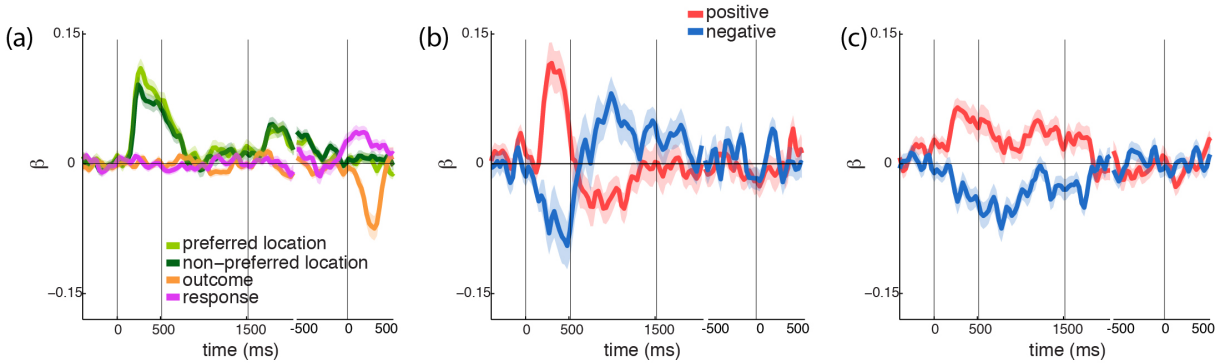

Supplementary Figure 7. Regression model with outcome and response parameters. To examine whether there was any additional variance in the firing rate of individual neurons (over and above the information we had already decoded about space) that could predict whether or not the animal would correctly detect the change in color we added two additional parameters to the regression model: whether firing rate predicted the subject's response ('change' or 'no change') and the subject's performance (trial outcome: correct or incorrect). **(a)** β values for all neurons corresponding to the preferred location, non-preferred location, trial outcome and subject's response respectively. There was virtually no variance explained by either response or outcome until the test epoch. **(b)** and **(c)** show the β_3 values separated according to whether the initial β_3-value was positive or negative and whether the β_3-value switched sign (b) or stayed constant through the trial (c). The regression was performed on spikes aligned to the onset of the sample up to 600-ms before subjects made their response (shown on the left hand side of the plot from time -500-ms to 2100-ms); and using spikes aligned to the time when animals made their response until 500-ms after the response (shown on the right hand side of the panel). Vertical lines indicate onset of the sample, delay and test epochs and the subjects' behavioral response, respectively. Only trials with a reaction time >600-ms were used in this analysis to avoid contamination of response related activity with test-epoch related activity. In practice, this excluded <1% of trials. We also note that this analysis does not necessarily contradict the fact that we were able to detect an effect on the precision of stored information in the local field potential of those electrodes containing spatially selective neurons (Figure 7d, 7e 8). First, the LFP averages activity across many thousands of neurons and so may be more sensitive at detecting the effect of PFC neuronal firing on behavior. More importantly, however, the single neuron firing rate and the LFP are not necessarily directly related to one another. It is the timing of spikes relative to one another that increases the power of the LFP, not their overall rate. Thus, the firing rate of PFC neurons could remain unchanged, and yet the increase in coherence in spikes communicating with posterior sensory areas would give rise to increased power in certain frequency bands of the LFP.

Supplementary Figure 8

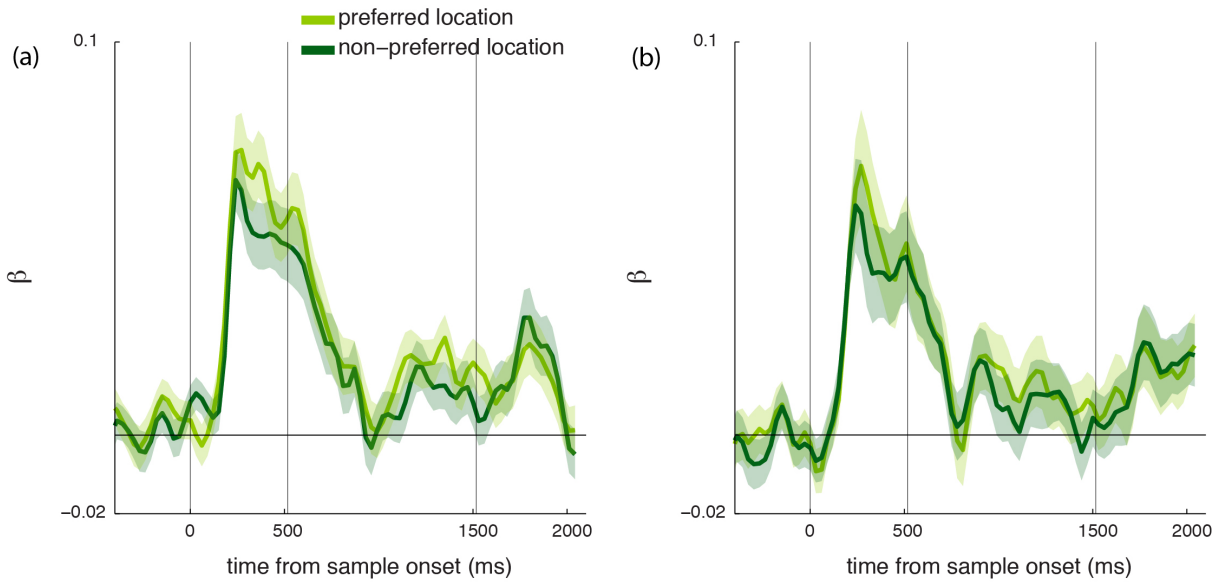

Supplementary Figure 8. Regression model without the attention parameter. The mathematical formulation of the biased competition model (Desimone & Duncan, 2005) is that a neuron's response to two items is the sum of the response to either item presented alone, weighted according to which item is currently being attended. To directly test whether this model predicted PFC neuronal responses, we removed the attention parameter from our original model (β_3), sorted the trials according to whether the first attended stimulus was the neuron's **(a)** preferred or **(b)** unpreferred stimulus (i.e. whether β_3 would have been positive or negative) and then examined the dynamics of β_1 and β_2. If the biased competition model were applicable to PFC, then β_1 should be greater than β_2 when the first attended stimulus is the neuron's preferred stimulus. In other words, the neuron's firing rate to the two item trials should be better predicted by the response to the preferred stimulus alone when the subject is attending to the preferred stimulus. In contrast, β_2 should be greater than β_1 when the first attended stimulus is the neuron's unpreferred stimulus. In other words, the neuron's firing rate to the two item trials should be better predicted by the response to the unpreferred stimulus alone when the subject is attending to the unpreferred stimulus. In fact, β_1 and β_2 encoding did not depend on whether the first attended stimulus was the neuron's preferred or unpreferred stimulus. This supports our original conclusion. PFC spatial signals consist of a static encoding of the spatial location of the stimuli and a dynamic modulation associated with the attentional locus. These findings are consistent with recent studies of attentional effects in PFC, which have emphasized dynamic modulation of spatial selectivity that does not necessarily reflect the original receptive field. For example, analysis of PFC activity at the population level during visual search shows that there is a rapid global allocation of resources to the target that is largely independent of whether the target is in the neuron's preferred or unpreferred location (Kadohisa et al., 2013). Similarly, analysis of single neurons in the frontal eye fields have shown that there is a dynamic increase in neural activity that correlates with the locus of attention and is largely independent of the neuron's receptive field (Buschman & Miller, 2009). Thus, PFC attentional mechanisms are distinct from those in earlier visual areas.