Supplementary information for
Generalized associative representations in parietal cortex

Jamie K. Fitzgerald, David J. Freedman, John A. Assad

Behavioral controls

Neuronal activity that resembles shape-pair selectivity may arise artifactually from behavioral factors that are similar for associated shape-pairs but distinct for non-associated shape-pairs. We controlled for this possibility by examining the influence of behavioral factors, including the accuracy of the animals' performance, microsaccadic eye movements, and mean eye position within the fixation window.

Activity of LIP neurons is affected by reward expectation and/or attention[1,2]. If the animals performed systemically better or worse for particular pairs of shapes, this may have produced differences in expected reward or attention that resemble shape-pair selectivity. In fact, the animals' accuracy was close to 90% correct for all six shapes (**Fig. 1c** of paper), making it unlikely that differences in attention/expected reward could have produced systematic differences in neuronal firing between shape-pairs. Nonetheless, for each neuron we asked whether variations in the animals' accuracy among the six sample-shape stimuli could account for differences in firing rate. Since accuracy is determined over multiple trials, we measured mean accuracy (accuracy = correct trials / total trials, excluding fixation breaks or premature touch-bar releases) throughout the recording session for each neuron; this was done separately for each of the six sample shapes. We also calculated the mean neuronal firing rates separately for each of the six sample shapes, in each of the four trial-analysis epochs. We performed a Spearman's correlation between the six mean firing rates and the corresponding six measures of mean accuracy. No more than 2.5% of the neurons had a significant correlation between mean firing rate and mean accuracy, in any of the six trial epochs ($P < 0.01$; fixation: 4/161, sample: 1/161, delay: 3/161, test: 0/161).

We also examined the effects of fixational eye metrics (microsaccadic eye movements and mean eye position). Eye movements/position within the fixation window could directly modulate the firing of LIP neurons[3]. Differences in these behavioral variables between pairs of associated shapes could masquerade as shape-pair selectivity. We measured the mean horizontal and vertical eye position throughout each of the four analysis epochs, and we also calculated the mean eye speed and number of microsaccadic eye movements in each epoch. We used the methods of Herrington et al.[3] to detect microsaccades.

We examined the effects of eye movement/position variables using a multiple linear regression analysis in which the three shape pairs and behavioral variables are taken as predictors of single-trial spike rates. (Unlike mean accuracy, eye movement/position variables can be measured on a trial-by-trial basis). Our approach was to ask whether the shape-pair predictor still had a statistically significant effect on firing rate after accounting for the other predictors in the model[4, 5].

Nature Neuroscience doi:10.1038/nn.2878

We started by examining a model in which *only* the associated shape-pairs are included as predictors:

$$Y = \beta_0 + \beta_1 P_a + \beta_2 P_b + \varepsilon \qquad \text{(Equation 1)}$$

Here Y is the spike rate within an epoch. β_0 is the constant coefficient. P_a and P_b are binary "dummy" variables which encode the pair membership of the sample shape: pair *a* ($P_a = 1$, $P_b = 0$), pair *b* ($P_a = 0$, $P_b = 1$), or pair *c* ($P_a = 0$, $P_b = 0$). β_1 and β_2 are the fitted coefficients of the shape-pair dummy variables, and ε is a random error term. If either of the partial F-statistics (equivalent to t^2) for the two shape-pair dummy variables was statistically significant ($P < 0.01$, Bonferroni corrected), a neuron was considered shape-pair selective. The percentages of neurons that were significantly modulated by shape-pairs are shown in Table 1. For comparison, Table 1 also shows the percentages of shape-pair selective neurons calculated by nested ANOVA (see main text). The somewhat lower percentages for the regression analysis are expected because Bonferroni correction is in general prone to false negatives.

Table 1	Percentage of shape-pair selective neurons ($P < 0.01$, $n = 161$ neurons)			
	Fixation	Sample	Delay	Test
Nested ANOVA	3%	64%	63%	37%
Partial F-test (Eq. 1)	2	54	58	32

We then added individual behavior variables to the regression model, as follows:

$$Y = \beta_0 + \beta_1 P_a + \beta_2 P_b + \beta_3 B + \varepsilon \qquad \text{(Equation 2)}$$

Here B is the behavioral variable under consideration, and β_3 is its corresponding fitted coefficient. We again determined the effect of individual predictor variables by calculating the partial F-statistic for each variable. The F-test determines whether each predictor is individually statistically significant after accounting for the other predictors in the model. The percentages of neurons that were still significantly modulated by shape-pairs are shown in Table 2 (for comparison, the top row of Table 2 shows percentages of cells that were significantly modulated in the model that only included shape-pairs). While accounting for individual behavioral variables reduced the percentages of cells that had a significant effect of shape-pairs, the reductions were generally small (< 18%).

Nature Neuroscience doi:10.1038/nn.2878

Table 2	Percentage of shape-pair selective neurons, after accounting for behavioral variables (partial F-test, $P < 0.01$)			
	Fixation	Sample	Delay	Test
Shape-pair selective, without considering behavioral metrics (Eq. 1)	2%	54%	58%	32%
Shape-pair selective, taking into account mean eye *x* position (Eq. 2)	3	54	57	32
Shape-pair selective, taking into account mean eye *y* position (Eq. 2)	2	51	48	26
Shape-pair selective, taking into account mean eye speed (Eq. 2)	2	55	57	33
Shape-pair selective, taking into account microsaccade number (Eq. 2)	2	52	57	33
Shape-pair selective, taking into account all above variables (Eq. 3)	2	48	45	26

Additionally, we combined all of these variables into a single model to calculate the number of neurons with significant effects of shape-pair after accounting for all variables together:

$$Y = \beta_0 + \beta_1 P_a + \beta_2 P_b + \beta_3 E_x + \beta_4 E_y + \beta_5 E_s + \beta_6 M + \varepsilon \qquad \text{(Equation 3)}$$

Here E_x is mean eye x position throughout the trial epoch, E_y is the mean eye y position, E_s is the mean eye speed, and M is the number of microsaccades. The results are similar to considering each behavioral variable individually: the reduction in pair-selective neurons was no greater than 23% (bottom row of Table 2).

We next asked how many neurons were significantly modulated by behavioral factors. This gives an upper bound of the number of cells for which behavioral factors may have contributed artifactually to the shape-pair selectivity. However, behavioral variables may also modulate neuronal activity in manners independent of the shape-pair selectivity or in ways that obscure pair selectivity. Given LIP's role in oculomotor planning, we expected to see some modulations by eye metrics, but no more than 21% of neurons were modulated. In contrast, up 58% of neurons were selective for shape-pairs.

Table 3	Percentage of neurons with a significant effect of behavior, after accounting for shape-pair (partial F-test, $P < 0.01$)			
	Fixation	Sample	Delay	Test
Mean eye *x* position	9%	12%	11%	10%
Mean eye *y* position	13	16	21	20
Mean eye speed	11	9	8	16
Number of microsaccades	4	8	5	11

Nature Neuroscience doi:10.1038/nn.2878

Controls for potential planning of saccadic eye movements

Saccadic eye movements and saccade planning affect the firing of LIP neurons[6]. Is it possible that saccades or saccade planning could have produced apparent shape-pair selectivity? This seems highly unlikely, for two reasons. First, our behavioral task required the animals to fixate, and the animals always signaled shape associations by a manual response, never by saccades. Second, the animals were never trained to associate particular shape-pairs with particular parts of visual space or directions of eye movement, and there was never any spatial bias in where we presented different shape-pairs during animal training or neuronal recording.

Nonetheless, there are two potential concerns about planned saccades. First, while all of the shapes were centered at the same location (the receptive field center) and subtended the same spatial extent at their widest/longest dimension, different shapes have different features – vertices, edges, etc. – that might attract planned saccades to different locations for different shapes. In principle, this could result in the animals' planning saccades to slightly different locations for different shapes, which could mimic shape selectivity. However, our main result in the paper is that LIP neurons responded selectively to *pairs* of associated shapes. The two associated shapes in a pair were chosen arbitrarily (and were paired differently for the two animals), so it is highly unlikely that any potential plan to saccade to a particular feature at a particular location would be the same for both shapes in a pair.

However, it is possible that the animals could have used some sort of mnemonic strategy in which different pairs of shapes became associated with plans to saccade in different directions. For example, whenever the animal was presented with associated shapes A or A', it might plan to saccade to the left but when presented with shapes B or B' it would plan to saccade to the right. This would be a rather odd strategy, but we cannot rule it out *a priori*. However, both animals were initially trained without a fixation requirement (although we were still monitoring eye position in many of those early training sessions), and we never observed the animals making saccades in different directions for different shape-pairs, even though they were free to do so at that point. Rather, if allowed to move their eyes, the animals would, quite naturally, always direct gaze at the location where the shapes were presented.

Nonetheless, if after fixation training the animals somehow developed a tendency to *plan* different saccades for different shape-pairs (perhaps as a mnemonic strategy), this might be reflected in the *actual* saccades that the animals typically make when they break fixation at the ends of trials. We thus analyzed saccades following correct trials.

We examined the effects of post-trial saccades using the multiple linear regression framework described above. We detected post-trial saccadic eye movements using the criteria of Pack et al.[7], and we examined first saccades occurring within 500 ms from the end of correct trials (although varying these parameters had very little effect on the outcome of the analyses). Because match and non-match trials had different cognitive demands and durations, we separated the data into match and non-match trials.

First, we asked whether the post-trial saccadic direction and amplitude influenced pair-selectivity. We calculated the saccade endpoints relative to the center position of the shape location. We

Nature Neuroscience doi:10.1038/nn.2878

considered neurons for which each sample pair had at least five trials in which a saccade occurred within 500 ms of the inter-trial interval. The vast majority of neurons met this criterion.

We used Equation 2 in which *B* is the saccadic eye-movement metric, and we calculated the partial F-statistic to ask whether the shape-pair predictor variables were statistically significant after accounting for the influence of eye-movement metrics. For comparison, we calculated the percentage of neurons showing shape-pair selectivity using a model in which the shape-pairs are the only predictors (Equation 1). We again only considered the trials with one or more post-trial saccades, and we separated the data into match and non-match trials. Because both of these manipulations reduce statistical power, fewer neurons had significant shape-pair selectivity than in the analyses in which all trials were analyzed.

Accounting for the post-trial saccade metrics did not change the percentages of cells that had a significant effect of shape-pairs by more than 8% (Table 4). Notably, in some cases taking into account eye-movement metrics *increased* the percentage of shape-pair selective neurons in the test epoch. We assume this was because shape-pair selectivity was occasionally partially obscured by neuronal activity related to planning upcoming eye movements that occurred at the end of the trial. In addition, the results were nearly identical if we considered absolute rather than signed saccade angles, relative to the location where the shapes were presented (the RF location).

Table 4	Percentage of shape-pair selective neurons, after accounting for post-trial saccadic metrics (partial F-test, $P < 0.01$)			
	Fixation	Sample	Delay	Test
Match trials (n = 129)				
Shape-pair selective, without considering saccade metrics (Eq. 1)	0%	31%	27%	29%
Shape-pair selective, taking into account saccade direction (Eq. 2)	1	31	27	27
Shape-pair selective, taking into account saccade amplitude (Eq. 2)	1	31	27	31
Non-match trials (n = 125)				
Shape-pair selective, without considering saccade metrics (Eq. 1)	2	34	36	14
Shape-pair selective, taking into account saccade direction (Eq. 2)	2	33	34	14
Shape-pair selective, taking into account saccade amplitude (Eq. 2)	2	33	33	14

We also asked whether the shape-pair selectivity was influenced by *whether or not* the animal made a saccade after the trial end, regardless of the metrics of the saccade. (Almost all cells were included in the analysis, except for a few cells for which the animal made a saccade within 500 ms of

Nature Neuroscience doi:10.1038/nn.2878

the end of *every* trial, and thus the regression was not full rank). As in the previous analysis, we saw very little change in the percentage of shape-pair selective neurons, no more than 5% (Table 5).

Table 5	Percentage of shape-pair selective neurons, after accounting for whether or not there was a saccade within 500 ms post-trial (partial F-test, $P < 0.01$)			
	Fixation	Sample	Delay	Test
Match trials (n = 160)				
Shape-pair selective, without considering whether or not a saccade occurred (Eq. 1)	1%	39%	39%	43%
Shape-pair selective, taking into account whether or not a saccade occurred (Eq. 2)	1	37	39	44
Non-match trials (n = 153)				
Shape-pair selective, without considering whether or not a saccade occurred (Eq. 1)	3	44	48	21
Shape-pair selective, taking into account whether or not a saccade occurred (Eq. 2)	3	43	48	22

We also used the regression framework to ask whether the post-trial saccadic eye movements are significant predictors of trial firing rates after accounting for shape-pair selectivity. After accounting for shape-pair selectivity, we found no more than 3% of neurons were selective for the metrics of post-trial saccades or the presence or absence of post-trial saccades. Importantly, this does *not* mean that the LIP neurons in our sample were *not* affected by saccadic eye movements; we separately tested every neuron with the eight-direction memory delayed saccade task and found that the vast majority of neurons had delay activity that was indeed selective for the direction of the upcoming saccade. Rather, any such planning for post-trial saccades did not have much impact on neuronal firing during correct trials in the shape-pair association task.

Nature Neuroscience doi:10.1038/nn.2878

Neuronal selectivity during the test period

For the main analyses, we examined selectivity for the sample shape-pair during the sample, delay, and test periods. But during the test period, one can also ask about selectivity for the test stimulus itself, as well as persistent selectivity for the previously presented sample stimulus.

We calculated test-period firing rate over a 220 ms window starting 80 ms after the onset of the test stimulus. For trials with reaction times < 300 ms, we included spikes only until the time of the manual response. (The allowable reaction-time window was 0–650 ms following the test-stimulus onset, but the average reaction time on match trials for the two animals was 368 ± 48 (s.d.m.) ms and 366 ± 51 ms. For both animals the reaction time was > 275 ms for 99% of trials.) The monkeys had to maintain fixation until the manual response, so the animals were always fixating throughout the period in which we calculated test-period firing rates.

We performed a three-way analysis of variance (ANOVA) to examine whether neuronal responses are selective for (1) the pair identity of the test shape, (2) the pair identity of the sample shape, and (3) whether the trials ended with a match or non-match. 70/161 (43%) neurons showed a main effect of test shape-pair and 52/161 (32%) neurons showed a main effect of the previously presented sample shape-pair (ANOVA, $P < 0.01$). 28/161 (17%) of neurons showed main effects of both the sample shape-pair and test-shape pair. Concomitant information about sample and test shape-pairs is presumably important during the test period, because this is when the animals must decide whether or not the sample and test shape belong to an associated shape-pair.

In contrast, very few LIP neurons explicitly encoded whether the trial outcome was a match or non-match: only 16/161 (10%) neurons showed a main effect of match/non-match outcome. Interestingly, previous categorization studies have shown that many neurons in prefrontal cortex do encode match/non-match trial outcomes[8]. The lack of match/non-match effects in parietal cortex suggests that associative signals in parietal cortex may be distinct from those in prefrontal cortex, and not simply an echo of signals sent back from prefrontal cortex.

Nature Neuroscience doi:10.1038/nn.2878

Error-trial analysis

The primary goal of our study was to ask whether LIP neurons signal learned associations between pairs of shape stimuli. We did not set out to specifically relate the firing of LIP neurons to the animals' behavioral ability to associate shape stimuli. The relationship between neuronal firing and behavior is best revealed in threshold tasks in which the animals often have different decisions/outcomes in response to the same stimulus, allowing neuronal firing to be directly related to behavioral responses on a trial-by-trial basis. We did not use a threshold task, but our animals still made some errors in the association task, typically less than 10–15% of trials. While we did not expect much statistical power to detect differences between correct and error trials, we nonetheless attempted to compare neuronal responses between correct and error trials, as follows.

We calculated the discriminability of shape-pairs separately for error trials and for correct trials, using a receiver operating characteristic (ROC) to calculate the probability that an ideal observer could correctly identify whether the firing rate on one trial corresponded to one shape-pair or another[9]. If LIP responses are related to behavioral outcome, the discriminability could be lower on error trials than on correct trials. As a first pass, we compared the shape-pair that elicited the strongest neuronal response ("preferred" pair) with the shape-pair that elicited the weakest neuronal response ("null" pair).

Because there were far fewer error trials than correct trials, we needed to first equate statistical power for the two trial types. For each cell and each trial epoch (sample, delay, and test) we thus matched the number of correct and error trials by randomly selecting a subset of the correct trials equal to the number of error trials for that same cell and epoch ("subsampling"). We used the remaining correct trials to define the preferred and null pairs for each cell, by epoch (this procedure eliminates any selection bias in defining the preferred and null shape-pair). For example, if a cell had 87 correct trials and 9 error trials for a shape-pair A and 90 correct trials and 10 error trials for shape-pair B, we subsampled 9/87 correct trials from pair A and 10/90 correct trials from pair B for the ROC calculation, and we used the remaining 78 (A) and 80 (B) correct trials to independently determine the preferred and null shape-pair. For each cell, we repeated the subsampling procedure 1,000 times for each epoch, yielding a distribution of 1,000 ROC values for correct and error trials. For each shuffle, we required at least three error trials from each of the preferred and null pairs, and we required at least 500/1000 shuffles for which this threshold was met to include a cell. Different numbers of neurons met this criteria by epoch (sample: $n = 54$, delay: $n = 64$, test: $n = 61$), because we defined the preferred and null pairs separately for each epoch.

We then compared the ROC value from the error trials to the ROC value from the correct trials. If the error ROC was greater or less than 97.5% of the subsampled correct trials, it was considered significantly different. This procedure corresponds to a two-tailed comparison with $P < 0.05$.

Supplementary Fig. 1 plots each neuron's error-trial ROC value against the mean of the distribution of subsampled correct-trial ROC values. This figure thus directly compares how well the cells could discriminate the shape-pairs on correct trials compared to how well the cells could discriminate the shape-pairs on error trials. The sample, delay, and test epochs are plotted separately. Filled symbols indicate cells for which the error trial ROC values were significantly different than the subsampled correct trials (by the two-tailed permutation test described in the previous paragraph; $P < 0.05$).

Nature Neuroscience doi:10.1038/nn.2878

The ROC values tended to be lower for error trials than correct trials: the distribution of ROC values for error trials was significantly lower than that for correct trials during the delay epoch (paired t-test; $P = 0.003$).

In addition, while there was a small number of neurons that showed a significant difference in ROC between correct and error trials, among those neurons the error-trial ROC was usually *less than* the subsampled correct-trial ROC (see filled symbols in Supplementary Fig. 1). This asymmetry was statistically significant for the delay epoch, for which we had the highest statistical power to make the comparison (10/64 cells with lower rather than higher ROC on error trials, versus 1/64 with higher ROC; chi-square test, $P = 0.0024$).

Thus, despite the low statistical power due to the small number of error trials, the ability to discriminate the sample shape-pairs based on neuronal firing rates tended to be lower on error trials than on correct trials. While this finding must be interpreted cautiously, one possibility is that reduced shape-pair selectivity in LIP is causally related to subsequent errors in the animals' performance of the task.

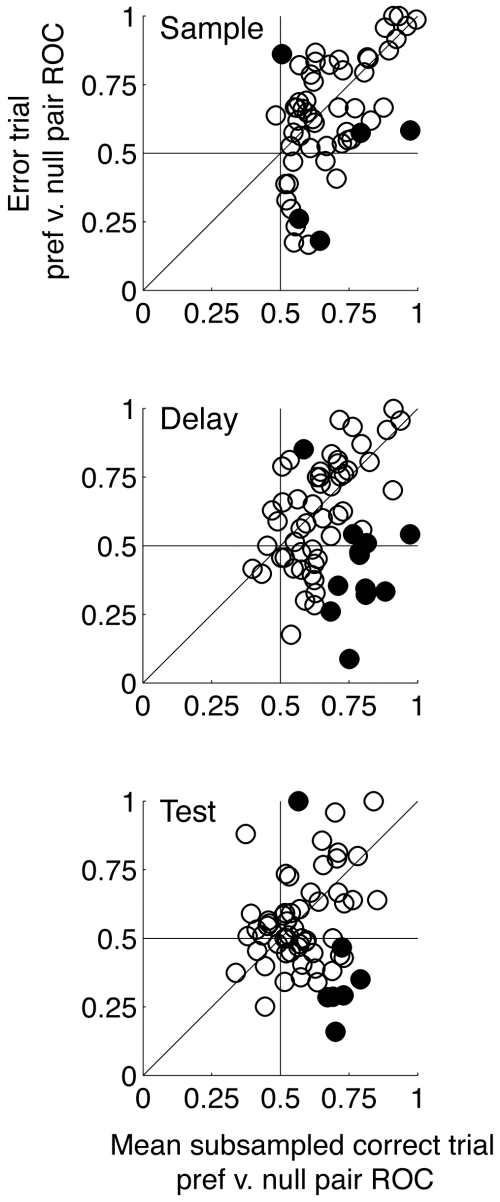

Supplementary Figure 1. Error and correct trial preferred versus null pair ROC values

Nature Neuroscience doi:10.1038/nn.2878

References

1. Platt, M.L. & Glimcher, P.W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233-238 (1999).
2. Maunsell, J.H. Neuronal representations of cognitive state: reward or attention? *Trends Cogn Sci* **8**, 261-265 (2004).
3. Herrington, T.M.*, et al.* The effect of microsaccades on the correlation between neural activity and behavior in middle temporal, ventral intraparietal, and lateral intraparietal areas. *J Neurosci* **29**, 5793-5805 (2009).
4. Chatterjee, S., Hadi, A.S., & Price, B. *Regression analysis by example* (John Wiley & Sons, Inc., New York, 2000).
5. Zar, J.H. *Biostatistical analysis* (Prentice Hall, Upper Saddle River, NJ, 1999).
6. Andersen, R.A. & Buneo, C.A. Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* **25**, 189-220 (2002).
7. Pack, C.C., Gartland, A.J. & Born, R.T. Integration of Contour and Terminator Signals in Visual Area MT of Alert Macaque. *J Neurosci* **24**, 3268-3280 (2004).
8. Freedman, D.J., Riesenhuber, M., Poggio, T. & Miller, E.K. Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* **291**, 312-316 (2001).
9. Green, D.M. & Swets, J.A. *Signal detection theory and psychophysics* (John Wiley, Oxford, England, 1966).

Nature Neuroscience doi:10.1038/nn.2878