

Neural Ensemble Decoding Reveals a Correlate of Viewer- to Object-Centered Spatial Transformation in Monkey Parietal Cortex

David A. Crowe,^{1,3} Bruno B. Averbeck,⁴ and Matthew V. Chafee^{1,2,3}

¹Department of Neuroscience, University of Minnesota Medical School, and ²Center for Cognitive Sciences, University of Minnesota, Minneapolis, Minnesota 55455, ³Brain Sciences Center, Veterans Affairs Medical Center, Minneapolis, Minnesota 55417, and ⁴Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London, London WC1N 3BG, United Kingdom

The parietal cortex contains representations of space in multiple coordinate systems including retina-, head-, body-, and world-based systems. Previously, we found that when monkeys are required to perform spatial computations on objects, many neurons in parietal area 7a represent position in an object-centered coordinate system as well. Because visual information enters the brain in a retina-centered reference frame, generation of an object-centered reference requires the brain to perform computation on the visual input. We provide evidence that area 7a contains a correlate of that computation. Specifically, area 7a contains neurons that code information in retina- and object-centered coordinate systems. The information in retina-centered coordinates emerges first, followed by the information in object-centered coordinates. We found that the strength and accuracy of these representations is correlated across trials. Finally, we found that retina-centered information could be used to predict subsequent object-centered signals, but not vice versa. These results are consistent with the hypothesis that either area 7a, or an area that precedes area 7a in the visual processing hierarchy, is performing the retina- to object-centered transformation.

Key words: parietal; object centered; neural ensemble; area 7a; time-resolved decoding; object based

Introduction

We have a singular and seamless perception of space suggesting a similarly singular neural representation of space within the brain. However, previous neurophysiological investigation in posterior parietal cortex has suggested that the brain constructs several representations of space concurrently. During visually guided eye movements, for example, different populations of posterior parietal neurons represent the direction of a saccade and/or the position of a saccade target in eye-centered (Mountcastle et al., 1981), head-centered (Andersen et al., 1985), body-, and world-centered (Snyder et al., 1998) spatial coordinates, demonstrating a multiplicity of spatial representation by parietal neurons. Parietal neurons can also represent spatial variables that are associated with spatial cognitive as opposed to sensorimotor function. For example, when monkeys are required to covertly traverse a path through a visual maze, the activity of parietal neurons modulates in time as the direction of the mental traversal changes in the absence of any physical movement or concurrent change in

the visual stimulus (Crowe et al., 2004, 2005). We found that parietal neurons represented spatial variables related to cognitive and not to sensorimotor function in the context of an object construction task as well. When monkeys were required to evaluate the position of one part of an object relative to others, largely distinct populations of posterior parietal neurons represented the spatial position of object parts in two reference frames. One population coded viewer-centered position defined relative to the midline of the viewer, and another population coded object-centered position defined relative to the midline of the reference object (Chafee et al., 2007).

The present study is motivated by the hypothesis that object-centered signals, like those we found during object construction, reflect a transformation of viewer-centered neural signals, because visual information enters the brain in a retina-centered coordinate system. We test the following predictions based on the above hypothesis. First, neural activity should code viewer-centered position before object-centered position. Second, the strength of viewer- and object-centered signals should be correlated across trials. Third, because object-centered representations depend on viewer-centered representations, viewer-centered information should predict subsequent object-centered information.

To test these predictions, we applied time-resolved linear discriminant analysis (LDA) to extract viewer- and object-centered positions from the activity of simultaneously recorded parietal neurons in monkeys performing the object construction task

Received Aug. 10, 2007; revised March 14, 2008; accepted April 7, 2008.

This work was supported by United States Public Health Service–National Institutes of Health Grants NS17413 and R24MH069675, Whitehall Foundation Grant 2005-08-44-APL, the Department of Veterans Affairs, and the American Legion Brain Sciences Chair. We thank Apostolos Georgopoulos for insightful and essential intellectual contributions to this work and for his generous support.

Correspondence should be addressed to David A. Crowe, Brain Sciences Center, Veterans Affairs Medical Center, 1 Veterans Drive, Minneapolis, MN 55417. E-mail: crowe009@umn.edu.

DOI:10.1523/JNEUROSCI.5105-07.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/285218-11\$15.00/0

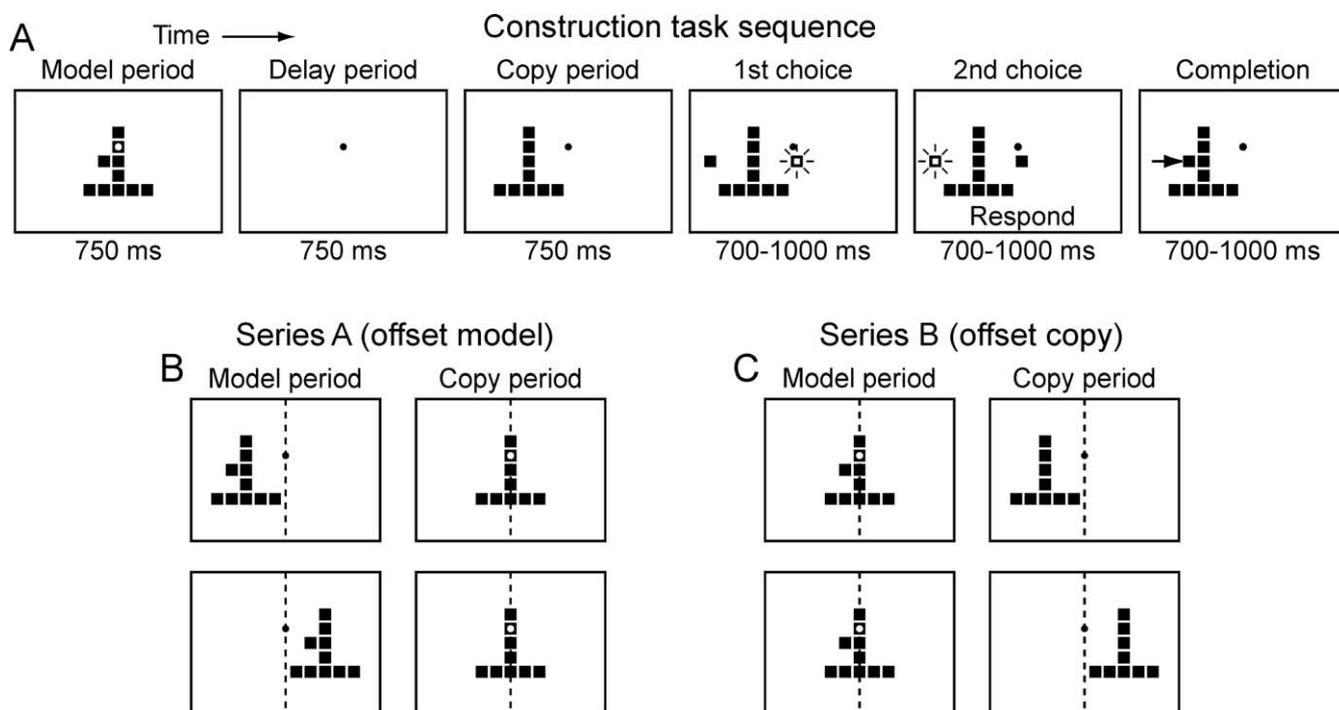


Figure 1. The object construction task. **A**, Task timeline. Objects were comprised of a variable configuration of squares. In each trial, two objects were shown in sequence, a model and a copy, separated by a delay. Monkeys localized the one critical square missing from the copy object and replaced it by selecting the closer of two flanking choice squares at the end of the trial. **B**, **C**, In different experimental series, the position of the model or copy object was offset randomly to the left or right of the fixation target, dissociating the horizontal position of the critical square in object- and viewer-referenced coordinates. **B**, Model object locations in series A were offset to the left or right randomly. **C**, Copy object locations in series B were offset to the left or right randomly.

(Chafee et al., 2005, 2007). LDA provides a concise measure of the information coded by neural activity. We took advantage of this to generate separate time courses of representation strength of viewer- and object-based positions. That, in turn, allowed us to assess the statistical dependence between these time courses. These analyses provided evidence consistent with the predictions above. We found that (1) neural activity represented viewer-centered spatial information before object-centered information, (2) the strength and accuracy of spatial representation in the two spatial frameworks was correlated across trials, and (3) that viewer-centered information predicted subsequent object-centered signals, but not vice versa.

Materials and Methods

Neural recording. We recorded the electrical activity of single neurons from area 7a in the posterior parietal cortex of two male rhesus macaques (4 and 6 kg) performing the object construction task. Neural activity was recorded using a 16 microelectrode Eckhorn Microdrive (Thomas Recording, Giessen, Germany). We advanced each electrode in the parietal cortex independently under computer control until we isolated the action potentials of ~20–30 neurons. This group of neurons constituted a neuronal ensemble, and we recorded the electrical activity of the neurons in the ensemble concurrently as monkeys performed a set of trials of the object construction task (below). As such, neural ensembles in this study are defined by sampling and not functional considerations, and in this sense are unlike the “cell assemblies” that Hebb (1949) defined as groups of synaptically connected and functionally synergistic neurons.

Action potentials were isolated on-line by a combination of waveform discriminators (MultiSpike detector; Alpha Omega Engineering, Nazareth, Israel) and time-amplitude window discriminators (DDIS-1; Bak Electronics, Mount Airy, MD). Two operators monitored the fidelity and stability of the action potential isolations during the experiment. Details of surgery, recording technique, and the locations of neural recording in area 7a of parietal cortex have been reported previously (Chafee et al., 2005, 2007). Care and treatment of the animals conformed to the Prin-

ciples of Laboratory Animal Care of the National Institutes of Health (NIH publication 86–23, revised in 1995). The Internal Animal Care and Use Committees of the University of Minnesota and the Minneapolis Veterans Affairs Medical Center approved all experimental protocols.

Object construction task. Two monkeys performed the object construction task (Fig. 1A). The monkeys were required to maintain their gaze fixated on a central target (within 1.5°) throughout each trial. Two objects were presented in sequence. Each object consisted of a collection of blue squares placed at various positions within a 5 by 5 grid. The first object constituted a model whose structure monkeys were required to reproduce. All model objects included, at a minimum, squares within the base row and central column of the grid, forming an inverted T-shaped frame. Unique model objects were constructed by the addition of either one or two squares at various locations in addition to the frame. The model object was visible for 750 ms (Fig. 1A, model period). After a delay (750 ms), a copy object was displayed, identical to the preceding model on that trial except that one square was missing. We refer to square that would be removed from the model object to produce the copy as the “critical square.” In the copy object, we refer to the location where a square was absent relative to the preceding model as the “missing critical square.” After the copy object was visible for 750 ms, a pair of choice squares was presented flanking the copy object. Choice squares were either located on opposite sides of the copy object at the same vertical position (horizontal choice array), or on the same side of the copy object in different vertical positions (vertical choice array), at random across trials. A short time after the two choice squares appeared (300–600 ms), one of them brightened for a period of 700–1000 ms (Fig. 1A, first choice). If the monkey pressed a response key during this interval, the brightly illuminated choice was animated to translate smoothly to the copy object (Fig. 1A, completion). If it did not press the response key during this time, the first choice returned to its original illumination and the second choice was made bright for 700–1000 ms. If the monkey pressed the response key when the second choice was brightly illuminated, it translated to the copy object. Monkeys were rewarded with a drop of juice if the completed object matched the configuration of the model object. The choice sequence was randomized across trials with

respect to whether the first or second choice was correct. The task required monkeys to perform spatial computations on objects without producing spatially variable motor output to report the result of those computations.

In two different experimental series, the horizontal position of either the model object (series A) (Fig. 1B) or the copy object (series B) (Fig. 1C) varied randomly across trials. The respective object was presented offset from the fixation target either to the left or right, at random. The offset was of a distance that displaced the object entirely into either the left or right visual hemifields (objects were 8° wide, and the center of the object was offset from the gaze fixation target by 4.2°).

Dividing ensembles into viewer-coding and object-coding subsets. We analyzed neuronal ensemble activity to decode a dichotomous spatial variable, side, relevant to the successful performance of the object construction task. Side refers to the position of the critical square present within the model or missing from the copy object, and is a factor with two levels, left and right. Side is defined in two spatial frames of reference concurrently. Viewer-centered side specifies whether the critical square was located to the left or right of the gaze fixation target. Object-centered side specifies whether the critical square was located to the left or right of the midline of the reference object. The critical square was located on the left and right side of the reference object at random across trials. The reference object was positioned to the left and right side of the gaze fixation target at random across trials. Therefore viewer-centered side and object-centered side were statistically independent variables.

As a preprocessing step in the decoding analysis, we performed a two-way ANCOVA to select subsets of neurons within each ensemble in which activity varied significantly as a function of the viewer-centered side and object-centered side of the critical square. Object-coding neurons were identified as those in which activity related significantly ($p < 0.01$) to the object-centered side, and not to the viewer-centered side or their interaction. Viewer-coding neurons were similarly identified as those in which activity related significantly to the viewer-centered side ($p < 0.01$), and not to the object-centered side or their interaction. Thus defined, object- and viewer-coding neurons comprised nonoverlapping populations. In the series A data, we used the firing rate within the entire model period as the dependent variable in the ANCOVA (in series A, the position of the model object varied). In the series B data, we used the firing rate within the entire copy period as the dependent variable (in series B, the position of the copy object varied). Two covariates were included in the ANCOVA model: baseline firing rate in the fixation period (before model onset) and the start time of the trial within the recording session. The above definitions of object- and viewer-coding neurons ensured that the two sets of neurons were nonoverlapping. We define a group of simultaneously recorded neurons with viewer- or object-related activity as a subset.

Decoding viewer-centered and object-centered sides (left or right) from viewer-coding and object-coding neurons. We decoded the time course with which neuronal subset activity represented viewer-centered side and object-centered side, to determine whether the strength of these two signals covaried in time. For that purpose, we applied LDA to the firing rates of each neuron in a subset measured within successive 100 ms bins throughout the construction trial (Johnson and Wichern, 1998; Averbeck et al., 2003; Chafee et al., 2005). In each time bin, LDA indicated the probability that neural subset activity coded either left or right relative to the viewer and also relative to the object. The results of the LDA analyses provided two concurrent decoding time series. One time series provided a quantitative measure of the strength with which subset activity represented the viewer-centered side of the critical square. The other time series provided a quantitative measure of the strength with which subset activity represented the object-centered side of the critical square.

LDA is a multivariate statistical technique. It classifies observations that are defined by a set of simultaneous measurements to one of a set of predefined categories. In our case, observations are 100 ms time bins within trials, each of which is defined by the set of firing rates observed in a group of simultaneously recorded neurons. Our analysis involved two categories, left and right, defined relative to the viewer or the object. We performed the classification with the Classify function in the Matlab Statistical Toolbox (The MathWorks, Natick, MA) using fivefold cross-

validation. Classify requires training and test data as input. We used a successive one-fifth of the trials as test data, and the remaining four-fifths of trials as training data, repeating the classification five times until all trials were included in the test data and were classified. LDA uses training data to compute the parameters of a set of discriminant functions, each defining one of the categories in the analysis. Each category is defined by a mean vector containing the average value of each discriminating variable across all observations in that category. In our case, the mean vector contained the mean firing rate in each viewer-coding or object-coding neuron within a given subset when the critical square was located left or right in the respective spatial framework. Categories are also defined by the covariance matrix of the discriminating variables, averaged across categories. The mean vector and covariance matrix provide the free parameters of a multidimensional Gaussian probability density function defining each category. Because the categories left and right were balanced in the design, equal prior probabilities for the categories were assumed in the analysis.

The training data were used to define the discriminant functions, and the classification was performed on the test data. For each test trial, we measured the firing rates of viewer- or object-coding neurons within a subset, and Classify compared this vector of firing rates to the mean vectors computed from the same subset defining the categories left and right in the training data, computing the posterior probability that the new (test) observation belonged to each category. The posterior probability is calculated by first computing the likelihood that either left or right was being represented by the neural activity, in the respective coordinate frame, and then dividing this value by the sum of the likelihoods for the two possibilities. This converts the two likelihoods to values that sum to one, which are the posterior probabilities. We classified the test trial as left or right depending on which category yielded the higher posterior probability. We tallied the number of times the classification was correct (across trials and across subsets) in decoding the side of the critical square relative to the viewer and relative to the object, within each time bin. The number of times the classification was correct provided a measure of the strength with which subset activity represented each variable. Treating viewer-coding and object-coding neurons as separate, simultaneously recorded subsets, and repeating the classification procedure in 100 ms bins in each subset produced the concurrent representation time courses for the object- and viewer-centered sides on which the subsequent analyses were based.

Figure 2 shows performance of the LDA analysis applied to a neuronal ensemble of 25 neurons, containing a subset of four viewer-coding cells and a subset of five object-coding cells, during the model period of an example trial in which the critical square was located to the right of the fixation point, but is on the left side of the object. Figure 2B shows the decoding time series for viewer-centered position and Figure 2C shows the decoding time series for object-centered position.

Correlating viewer- and object-coding signals. We measured the correlation between the two decoding time series representing the viewer-centered and object-centered sides of the critical square, using several methods. First, we computed the correlation coefficient between the maximum posterior probabilities for the correct object- and viewer-centered sides of the critical square obtained in each trial (across bins). Next, we performed a χ^2 test to assess the significance of the association of success or failure in correctly decoding the side of the critical square in the two coordinate frames on each trial (interpreting each trial as coding left or right in each framework based on the highest posterior probability across bins). Finally, we used a linear time-series regression analysis to quantify the degree to which the viewer-centered decoding time series could be used to predict the object-centered decoding time series, as described below.

In the regression analysis, we predicted the strength of object-centered signals using lagged viewer-centered signals. To do this, we assessed the variance of the residuals of two linear regression models. In the first, object-centered posterior probability was predicted by a five-lag autoregressive model:

$$O_t = \beta_1 O_{t-1} + \beta_2 O_{t-2} + \beta_3 O_{t-3} + \beta_4 O_{t-4} + \beta_5 O_{t-5} + \epsilon, \quad (1)$$

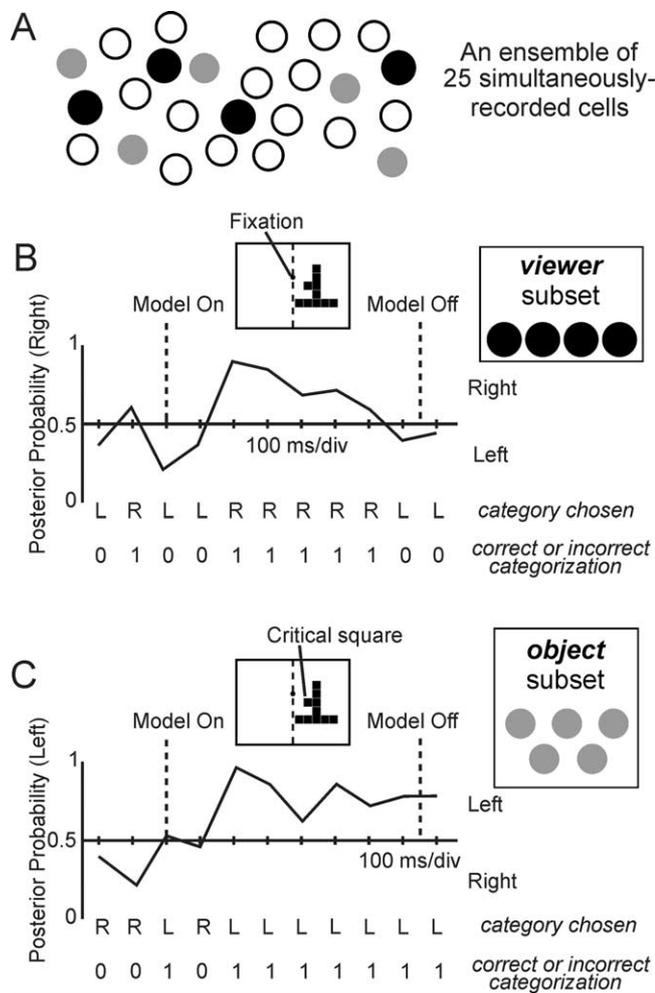


Figure 2. Performance of the LDA analysis on a single trial. **A**, An ensemble of simultaneously recorded cells, four of which were significantly related to the viewer-centered position of the critical square (black circles; viewer subset), and five of which were related to the object-centered position (gray circles; object subset). **B**, Posterior probability of classifying the critical square as right in viewer-centered coordinates (the correct classification). The line indicates a decoding time course. Each point represents the posterior probability, evaluated from a 100 ms bin of subset activity, that the critical square was located in the right half of viewer-centered space. In the LDA, a case is classified as the category with the highest posterior probability. Because there were two possible categories, and because posterior probabilities over all categories sum to one, if the posterior probability for right is >0.5 , the bin is classified as right. On this trial, the LDA consistently classifies the trial correctly beginning 200 ms after the model object is shown (model on). **C**, Here, the subset activity of five object-centered cells is used to classify the location of the critical square relative to the object (in this case, left). The line indicates the posterior probability that the critical square was located on the left side of the object.

where O_t is the object-centered posterior probability in the current bin, and O_{t-1}, \dots, O_{t-5} are the posterior probabilities in the preceding five bins. In the second regression, we added the viewer-centered posterior probabilities of the preceding five lags:

$$O_t = \beta_{1..5}O_{t-1..5} + \beta_6V_{t-1} + \beta_7V_{t-2} + \beta_8V_{t-3} + \beta_9V_{t-4} + \beta_{10}V_{t-5} + \epsilon, \quad (2)$$

where V_{t-1}, \dots, V_{t-5} are the viewer-centered posterior probabilities in the preceding five bins. We tested the significance of the addition of the viewer-centered bins by comparing the variances of the residuals obtained from the two models, using an F test, evaluated using k and $n-2k$ degrees of freedom, where k is the number of lags and n is the number of observations. Before the regressions, the time-series data were differenced to improve stationarity. Thus, the model we fit was an ARIX model (Ljung, 1999). This above analysis was repeated on a bin-by-bin basis

Table 1. Numbers of subsets and significant neurons included in time-course and correlation analyses using subsets defined by less-restrictive and more-restrictive criteria

	Model period (series A)		Copy period (series B)	
	Subsets	Neurons	Subsets	Neurons
Less-restrictive criterion				
Time-course analysis	45	140	18	146
Correlation analyses	29	119	14	118
More-restrictive criterion				
Time-course analysis	14	64	23	130
Correlation analyses	4	27	3	28

In the time-course analysis, subsets were restricted to those containing one (less restrictive) or three (more restrictive) significant viewer- or object-coding neurons. In the correlation analysis, subsets were restricted to those containing at least one significant viewer- and one significant object-coding neuron (less restrictive), or at least three significant viewer- and three significant object-coding neurons (more restrictive).

throughout the trial, providing us with a time-varying estimate of the linkage between viewer- and object-centered signals, as measured by the ability of one signal to predict the other. We repeated the above analysis with two additional variations. We assessed association in the opposite direction, i.e., evaluating the ability of object-centered signals to predict viewer-centered signals, and we also performed the analysis using only one instead of five lagged bins.

Selection of neurons and subsets. Decoding accuracy generally scaled with the number of neurons in each subset significantly related to viewer- and object-centered sides (see Fig. 4). The number of neurons with viewer- and object-centered signals that we could record simultaneously was limited by the size of the neural ensembles we could study at one time using the 16 electrode recording matrix (ensembles usually included 20–30 neurons). Typically, we encountered ensembles containing a small number of significant neurons. More rarely we encountered ensembles containing many significant neurons. In considering which ensembles to include in the decoding analysis, there was therefore a trade-off between the number of ensembles included and the number of significant neurons contained within each ensemble. In light of this trade-off, we performed two analyses. In the first, we included all ensembles containing a subset of at least 1 significant viewer- or object-coding neuron. This included a large fraction of the ensembles we recorded, and so provided a better estimate of the information coded by the “average ensemble” we were able to record. The information coded by these ensembles was necessarily noisier than obtained in our second analysis, which was restricted to a smaller number of ensembles in our database that included a subset of a minimum of three viewer- or object-coding neurons. We refer to these two criteria (at least one or at least three significant viewer- or object-coding neurons) as being less and more restrictive, respectively, and report the decoding results obtained using both criteria.

Neuronal database. We recorded the activity of 51 neural ensembles in series A (in which we varied the position of the model object), including a total of 1013 neurons. We recorded the activity of 18 ensembles in series B (in which we varied the position of the copy object), including a total of 504 neurons. These sets were nonoverlapping, so therefore our dataset includes electrophysiological recordings from a total of 1517 neurons. In series A, we analyzed neural activity during the model period. We varied the position of the model object in this series, and this allowed us to dissociate the viewer- and object-based sides of the critical square during the model period. In series B, we varied the position of the copy object, and so analyzed neural activity during the copy period to dissociate viewer- and object-based coding of the critical square missing from the copy object. Neural ensemble activity was recorded as the monkeys performed either 128 trials in series A or 160 trials in series B.

The numbers of subsets and neurons in them which met the more- and less-restrictive statistical criteria used to screen subsets for both the time course and correlation analyses (described above) are listed in Table 1. Using the less-restrictive criterion in the time course analysis, 33 of the subsets were recorded in monkey 1, and 30 were recorded in monkey 2. Using the less-restrictive criterion in the correlation analyses, 23 subsets were recorded in monkey 1, and 20 subsets were recorded in monkey 2.

Using the more-restrictive criterion in the time course analysis, 13 of the subsets were recorded from monkey 1 and 19 were recorded from monkey 2. Under this criterion in the correlation analyses, we recorded four subsets from monkey 1 and three from monkey 2.

Using ensembles to estimate network dynamics: testing significance of simultaneous activity using a bootstrap analysis. To determine the extent to which the results of the correlation analyses depended on the simultaneity of recording object- and viewer-centered signals, we randomly paired viewer- and object-centered decoding time series from different ensembles, recorded at different times (frequently on different days). This shuffling procedure broke the simultaneity between the activity of subsets of viewer- and object-coding neurons in each ensemble used to derive the viewer-centered and object-centered decoding time series. However, the shuffling did not alter the viewer- or object-centered decoding time series themselves or the firing rates of neurons on which these were based. We repeated this shuffling procedure 1000 times, and after each shuffling we computed the correlation between viewer- and object-centered signals, as well as the degree to which one could be used to predict the other, using the regression method, above. This provided a set of R^2 values obtained from the correlation or regression analyses, under conditions in which the two-signals could not influence one another because they were recorded at different times. We could then evaluate the proportion of R^2 values obtained by shuffling that was as large as or larger than the value we obtained from the original, unshuffled data. This percentage quantified the probability that the linkage we detected in the original analysis was spurious, because of either a sample of ensembles that was too small, variations in firing rate and neural representation that were time-locked to the behavioral events of the trial and therefore repeatable across experiments and days, or to other factors of the analysis that may have overestimated the degree of temporal correlation in viewer- and object-centered representation.

Results

In the construction task, monkeys had to localize and replace a critical square within a reference object. We randomly varied whether the critical square was located on the left or right side of the reference object, and whether the reference object was located to the left or right of the viewer. Therefore, object-centered left and right and viewer-centered left and right varied independently across trials, allowing us to decode the neural representation of the position of a single locus (corresponding to critical square) in these two spatial coordinate systems concurrently at each time point in the trial.

Figure 3 illustrates the activity of two parietal neurons in which firing rate varied as a function of the side of the critical square in viewer-centered (Fig. 3A–D) and object-centered (E–H) coordinates, respectively. Activity of the viewer-coding neuron was elevated during the model period when the critical square and model object were located to the left of the fixation target, regardless of whether the critical square was located on the left (Fig. 3A) or right (B) side of the model object with respect to its central, vertical axis (activity during model period of series A shown). Activity of the object-coding neuron, in contrast, was greater when the missing critical square was located on the right side of the copy object (Fig. 3F, H; arrow points to location of the missing critical square), regardless of whether the missing critical square was located in the left (Fig. 3F) or right (H) visual hemifields, and therefore, regardless of whether the critical square was located in the left and right halves of viewer-centered space (activity during copy period of series B shown).

Temporal order of spatial representation

By decoding the strength of neural representation of viewer- and object-based positions across a succession of time bins throughout each trial, we obtained evidence that neurons in parietal cortex

represented viewer-centered position before object-centered position (Fig. 4). The decoding time series in Figure 4 illustrate the percentage of trials (averaged across subsets and trials) in which the side of the critical square was classified correctly, referenced to the viewer (solid lines) or the object (dashed lines) within each bin. Regardless of whether we used subsets of at least one significant neuron (Fig. 4A), subsets of at least three significant neurons (Fig. 4B), or all significant neurons not recorded simultaneously (Fig. 4C), we found that during the model period the strength of the neural representation of viewer-centered position increased before the neural representation of object-centered position, immediately after the presentation of the model object.

Decoding accuracy increased with the inclusion of the activity of increasing numbers of neurons in which firing rate related significantly to the decoded parameter. For example, the mean posterior probability obtained by LDA decoding for critical square position increased when comparing the results of the analysis applied to subsets of one or more significant neurons (Fig. 4A), subsets of three or more significant neurons (Fig. 4B), or the entire population of significant neurons (Fig. 4C). Viewer-centered signals tended to decay after presentation of the model object, whereas object-centered signals tended to persist throughout most of the trial (Fig. 4A–C). This is relevant because the object-centered information was critical for task performance during the copy and choice periods, whereas the viewer-centered information was not.

We also decoded the viewer- and object-centered sides of the critical square missing from the copy object during the copy period in series B. Neurons included in this analysis were selected by virtue of exhibiting a significant relation in firing rate to critical square position during the copy period. In this analysis, we found that the representation of object-centered side persisted from the model period, and was stronger at all time points relative to the representation of viewer-centered side (Fig. 4D–F). This suggests that early in the trial, viewer-centered representation is earlier and stronger, whereas later in the trial this pattern is reversed.

The lag between viewer- and object-centered signals was present in the spike-rate time courses as well as the decoding time courses. We averaged spike-density functions from neurons with the strongest object- and viewer-centered signals ($p < 0.001$ for each factor in a two-way ANCOVA) across trials and neurons (Fig. 5A). Population activity coded the position of the critical square in viewer-centered coordinates first (at the divergence between thick solid and thick broken lines) before it coded position in object-centered coordinates (divergence between thin solid and thin broken lines). To compare the relative timing of viewer- and object-centered representation as measured by LDA decoding, we also cross-correlated the object and viewer posterior probability time series obtained from simultaneously recorded activity during the model period. The posterior probability is related to the information about the critical square in each coordinate system, and as such shows us when information about each coordinate system increased. The average (across trials and subsets) cross-correlation function (Fig. 5B) peaks when the object-centered decoding time course is shifted -100 ms relative to the viewer-centered decoding time course, indicating that the representation of viewer-centered position precedes that of object-centered position.

Correlation of viewer- and object-centered spatial representations

In addition to being offset in time, viewer- and object-centered signals were correlated in strength across trials. To reveal this, we examined the posterior probabilities obtained from the LDA analysis. Posterior probabilities are related to the strength of rep-

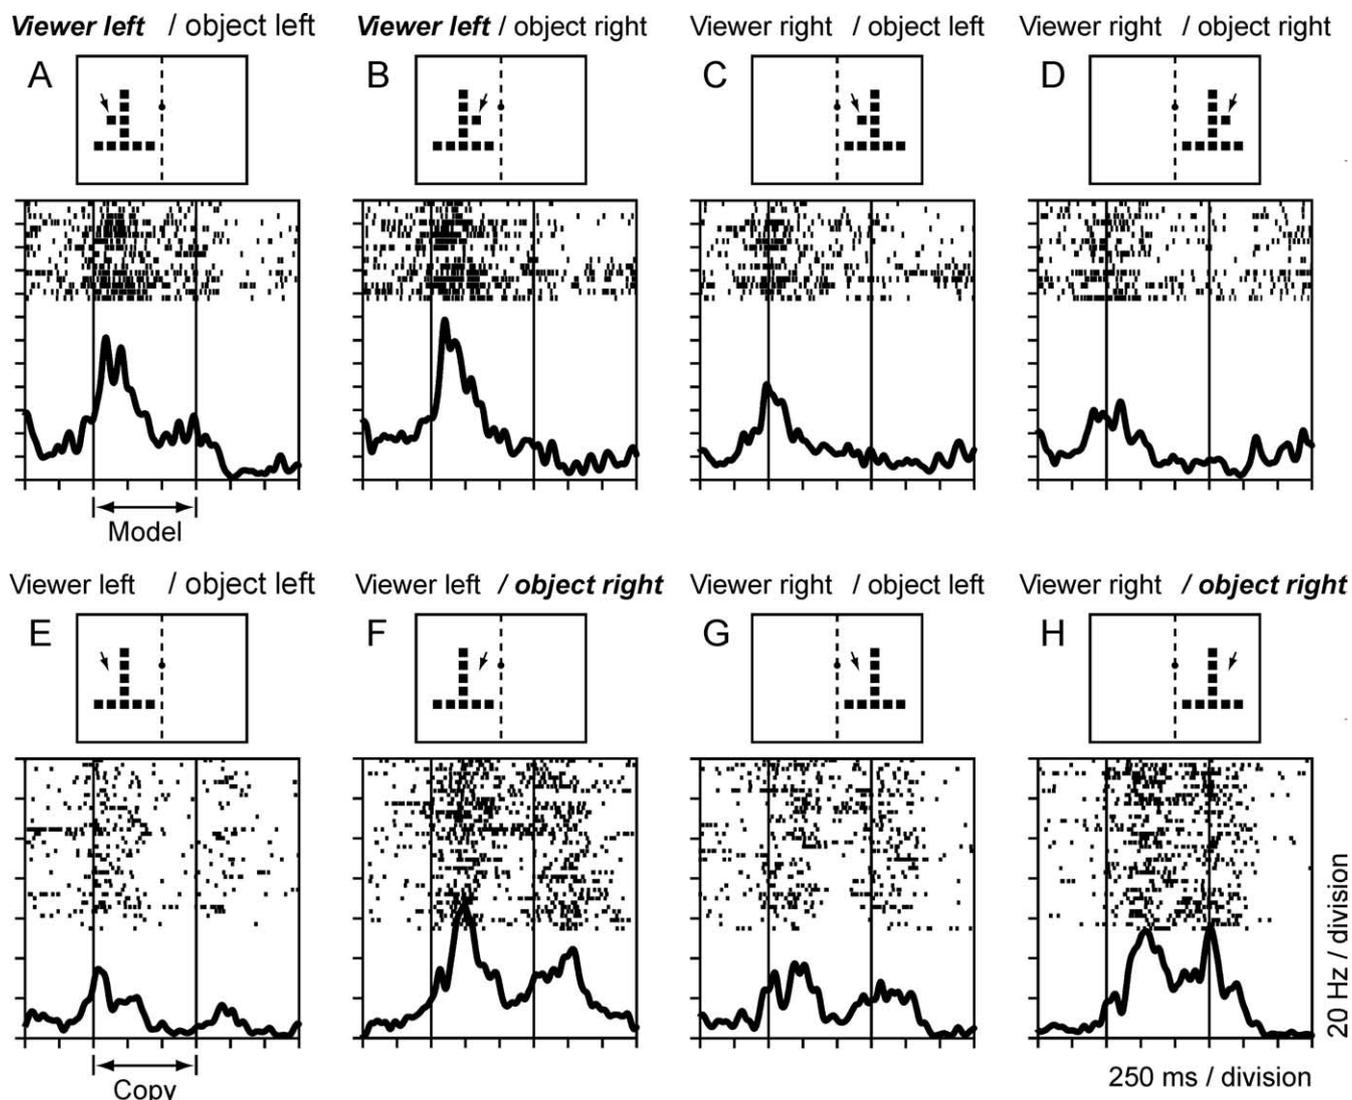


Figure 3. Two forms of spatial coding during the object construction task. *A–D*, Neural activity of an area 7a neuron was elevated during the model period when the critical square and model object were located to the left of the fixation target (viewer-centered coding). *E–H*, The firing rate of another area 7a neuron was greater when the missing critical square was located on the right side of the copy object, regardless of whether this object was presented to the right or left of the fixation target (object-centered coding). Arrows represent the position of the critical square (*A–D*) or missing critical square (*E–H*). *A, B, E, F*, Critical square on viewer left. *C, D, G, H*, Critical square on viewer right. *A, C, E, G*, Critical square on object left. *B, D, F, H*, Critical square on object right.

resentation. The higher the posterior probability, the more strongly the neural subset represents the side of the critical square in the chosen framework. For each trial, we found the maximum posterior probabilities for correct classifications: one indicating the certainty with which viewer-centered position was decoded, and the other indicating the certainty with which object-centered position was decoded. These probabilities in the two spatial reference frames were significantly positively correlated (Fig. 6). This was true during both the model period (Fig. 6*A, C*) and the copy period (Fig. 6*B, D*) of the trial, and it was true regardless of whether the analysis was based on subsets selected using either the less-restrictive (Fig. 6*A, B*) (36 subsets, 237 significant neurons) or more-restrictive (Fig. 6*C, D*) (7 subsets, 55 significant neurons) criterion. Using subsets defined by the less-restrictive criterion, the strengths of viewer-centered and object-centered representations were significantly correlated across trials in both the model period (Fig. 6*A*) ($r = 0.08$; $p < 10^{-5}$) and the copy period (Fig. 6*B*) ($r = 0.13$; $p < 10^{-9}$). The linear relation between the average viewer-centered and object-centered posterior probability was strongest using

fewer subsets each containing more (minimum three) significant neurons (Fig. 6*B, D*). In this case, the correlation coefficient between the two posterior probabilities was 0.14 during the model period (Fig. 6*C*) ($p = 0.001$) and 0.26 during the copy period (Fig. 6*D*) ($p < 10^{-9}$). Therefore, the correlation in strength between viewer- and object-centered representations was significant regardless of how the subsets were defined or the number of subsets included in the analysis.

We also assessed the association between overall success and failure in classification on a trial-by-trial basis in the two coordinate frames. In the less-restrictive case, we found that these measures were significantly associated during the model period ($\chi^2 = 3.0$; $p < 0.05$), but not during the copy period. Similarly, when basing the decoding on subsets defined with the more-restrictive criterion, we found that the outcomes (success/failure) of decoding viewer-centered and object-centered sides were significantly associated across trials ($\chi^2 = 5.90$; $p < 0.01$) during the model period, but not the copy period. Thus during the model period, in the case that LDA decoding yielded the incorrect answer for the viewer-centered side of the critical square on a given trial, it

tended to yield the incorrect answer for the object-centered side as well, regardless of whether subsets were defined by the less- or more-restrictive criterion.

Using the viewer-centered decoding time series to predict the object-centered decoding time series

We were interested in determining whether object-centered signals could be predicted from viewer-centered signals. Therefore, we performed a linear regression analysis that modeled the posterior probability in each bin of the object-centered time series as a linear function of the posterior probabilities in the preceding five bins of the viewer-centered time series. We controlled for the autocorrelation in the object-centered decoding time series by including the posterior probabilities in the preceding five bins of the object-centered time series in the model. First, we fit an autoregressive model that predicted the object-centered posterior probability in each bin of the decoding time series using the preceding five bins in the object-centered series only. We then tested the hypothesis that the addition of viewer-centered posterior probabilities in the same previous five bins would significantly improve our estimate of the object-centered representation beyond the estimate obtained with just the autoregressive terms. We did this analysis for each bin, starting at the first bin after the onset of the model object (where the five preceding bins were contained within the 500 ms pre-trial fixation period). In Figure 7, we plot the significance (p value) of the increase in variance in the object-centered posterior probability time series explained by the addition of the viewer-centered terms in the model, as a function of time throughout the trial (Fig. 7, thick lines). We found that during the model period, inclusion of the lagged bins of the viewer-centered decoding time series improved the fit by explaining a significantly larger proportion of variance in the object-centered time series, relative to the linear model excluding these terms (Fig. 7A) (note that thick line drops below the level of significance at $p < 0.05$ during the model period). In contrast, when this analysis was reversed, lagged object-centered information did not significantly predict viewer-centered signals (Fig. 7A) (thin line). This was true regardless of whether we used the less-restrictive criterion for subset inclusion (Fig. 7A) or more-restrictive criterion (Fig. 7C). These results were maintained when we included only one lagged bin in the analysis, allowing us to test whether the interaction between object- and viewer-centered time series was still evident when a shorter time window was examined (Fig. 7B,D).

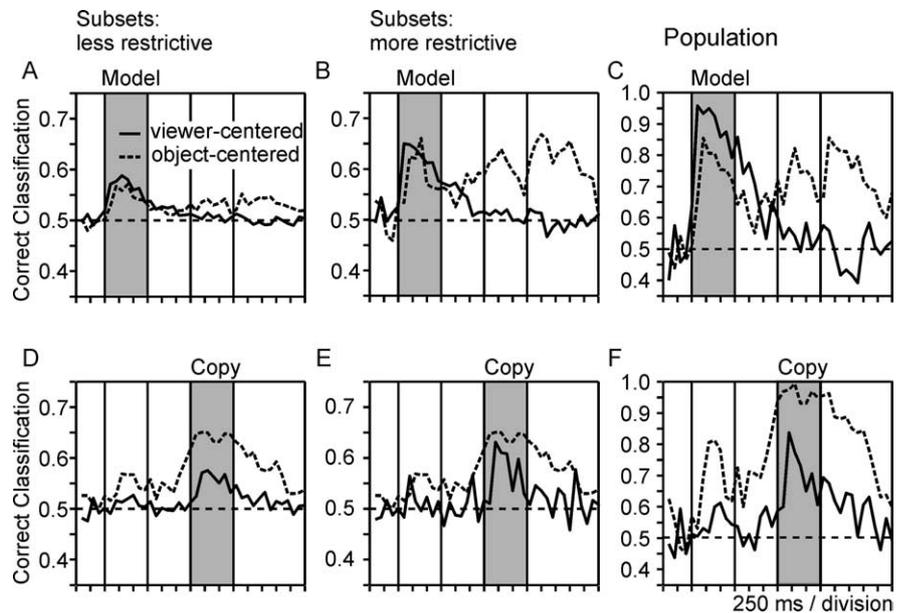


Figure 4. Time course of spatial representation in alternative frameworks. *A–F*, LDA was applied to 100 ms bins of subset (*A, B, D, E*) or population (*C, F*) activity in parietal cortex to classify each bin as representing right or left in viewer-centered (solid lines) or object-centered (dashed lines) space. The ordinate represents the percentage of bins correctly classified to these categories as defined by the positions and configurations of the objects presented. Neurons were included if their activity related significantly to horizontal position in one of the two reference frames during either the model period (*A–C*) or copy period (*D–F*) of the trial. *A*, Model period, less restrictive (viewer: 40 subsets with 86 cells; object: 34 subsets with 54 cells). *B*, Model period, more restrictive (viewer: 12 subsets with 50 cells; object: 4 subsets with 14 cells). *C*, Model period, total population (viewer: 86 cells; object: 54 cells). *D*, Copy period, less restrictive (viewer: 12 subsets with 26 cells; object: 18 subsets with 120 cells). *E*, Copy period, more restrictive (viewer: 3 subsets with 10 cells; object: 18 subsets with 120 cells). *F*, Copy period, total population (viewer: 26 cells; object: 120 cells). Note that for this analysis, the sum of viewer and object subsets is more than the total listed in Table 1, because a single ensemble often contained both object- and viewer-centered cells. Because the cells were selected as nonoverlapping groups, the total number of cells is the sum of the number of cells in the viewer and object groups.

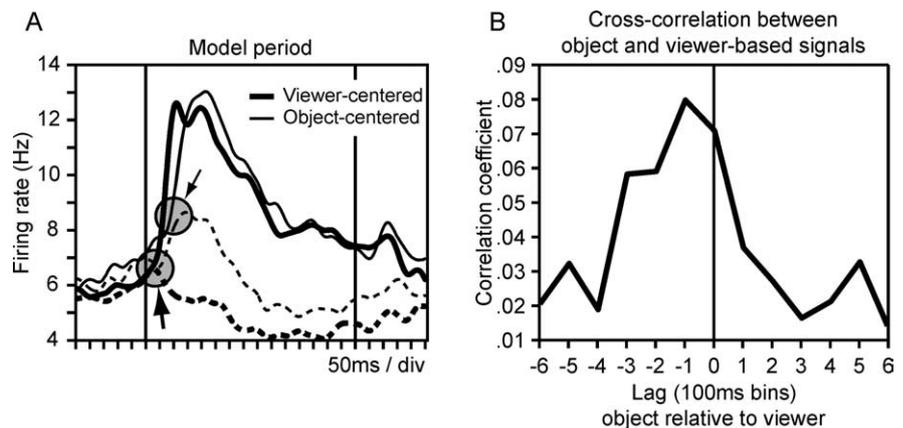


Figure 5. Confirmation of the order of spatial representation. *A*, Lines illustrate the average firing rate time course (spike-density functions; $\sigma = 20$ ms) of neural populations coding critical square side in either viewer-centered (thick lines; 30 neurons) or object-centered (thin lines; 50 neurons) space. Activity is plotted separately when the critical square was located on the preferred side (solid lines) or nonpreferred sides (dashed lines) of each spatial reference frame. The point in time at which population activity resolved the position of the critical square is indicated by the point at which population activity diverged on preferred versus nonpreferred trials. *B*, Average cross-correlation function of posterior probabilities of classifications in object- and viewer-centered space during the model period.

Correlation of viewer and object representation depends on simultaneous activity

Our hypothesis that viewer- and object-centered spatial representations are functionally linked is supported by the finding that fluctuations in the strengths of these representations are correlated over time, and that one decoding time series can be used to predict the other. If in fact the object-centered representation is

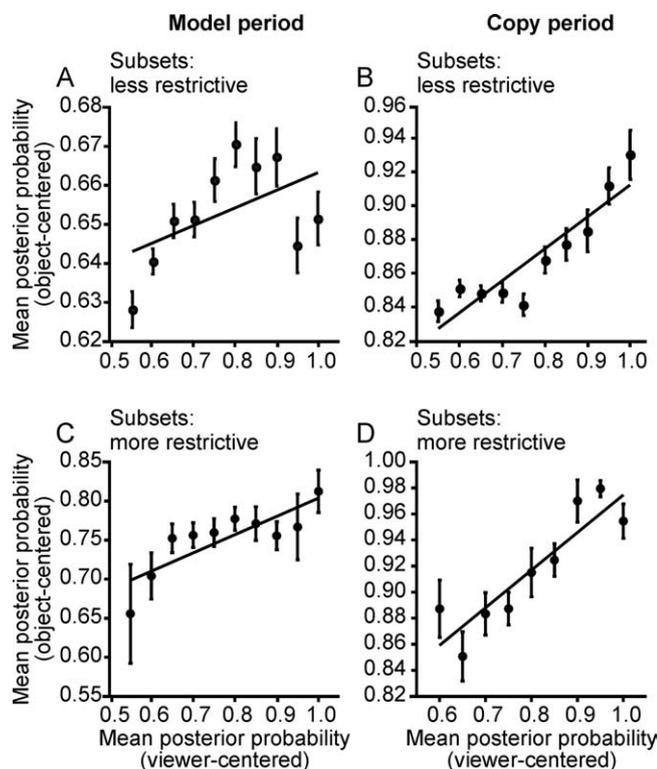


Figure 6. Correlation of representation strength of viewer- and object-centered positions across trials. Posterior probability associated with classification of the viewer-centered position is represented along the abscissa, divided into 10 equal bins in the range 0.55–1.0. Posterior probability associated with classification of object-centered position is represented along the ordinate. The mean object-centered probability is computed within each bin of viewer-centered probability. Error bars represent 1 SEM. **A**, Representation during model period, less-restrictive criterion (29 subsets of at least 1 viewer and 1 object cell, 119 cells). **B**, Representation during copy period, less-restrictive criterion (14 subsets of at least 1 viewer and 1 object cell, 118 cells). **C**, Representation during model period, more-restrictive criterion (4 subsets of at least 3 viewer and 3 object cells, 4 subsets of 27 cells). **D**, Representation during copy period, more-restrictive criterion (3 subsets of at least 3 viewer and 3 object cells, 28 cells).

produced by a transform applied to the viewer-centered representation, these correlations should only be present in the case that the two representations were decoded from the activities of simultaneously recorded neurons. We tested this prediction using a bootstrap analysis in which we compared the results of our correlation and regression analyses using both the original data in which the two time series were derived from simultaneously recorded neural activity, and shuffled data, in which the two time series were derived from neural activity recorded at different times.

We randomly paired viewer-centered and object-centered decoding time series from neural ensembles recorded at different times, and duplicated the analysis above quantifying the correlation of the strength of viewer- and object-based signals. We repeated this procedure 1000 times, noting the R^2 value obtained at each iteration. In this way, we used the same set of neural data, the same firing rates, and the same subsets of neurons in each ensemble used to generate the viewer- and object-centered decoding time series used in the analysis illustrated in Figure 6, but the condition of simultaneous recording across the neural subsets generating viewer- and object-centered decoding time series was broken. We found that in no iteration of this bootstrap analysis did the R^2 value in the nonsimultaneous case exceed that obtained in the simulta-

neous case ($p < 0.001$), either in the copy or model period, using either selection criterion.

We repeated this bootstrap procedure with the regression analysis, computing the increase in R^2 associated with inclusion of the viewer-centered independent terms in the linear model (Eq. 2). We compared R^2 values obtained in the nonsimultaneous recording bootstrap iterations to those obtained from the original data. We first performed this analysis using the less-restrictive criterion. We found that $<5\%$ of R^2 values from randomized sets were greater than those obtained with the original data set ($p < 0.05$), at each significant time point, using either one or five lagged bins. Furthermore, when we summed R^2 values over all significant bins, 0 and 0.5% of the randomized R^2 values surpassed the original ($p < 0.001$; $p < 0.01$) when we used one and five lagged bins, respectively. Finally, we repeated this analysis using the more-restrictive criterion. In this case, no R^2 values from randomized sets were greater than those obtained with the original data set, at each significant time point in the original data. These data show that the linkage of neural representations we observed required that viewer- and object-centered decoding time series were derived from simultaneously recorded neural activity.

Interaction between object- and viewer-centered representation

It is possible that neural activity relating significantly to the interaction between viewer- and object-centered positions may participate in the transformation of one to the other spatial representation. If neurons coding the interaction between the two spatial frameworks represent an intermediate representation, we would predict that neural activity should represent the viewer-centered position first, then the interaction between viewer- and object-centered positions, followed finally by the object-centered position. That order of representation can be seen in Figure 8. The representation time course of the interaction between viewer- and object-centered positions (green line) falls between the time courses of viewer- (blue) and object-centered (red) representation. Using the regression analysis above, we found that the viewer-centered time course significantly predicted the interaction time course, and that the interaction time course predicted the object-centered time course during the model period. Because interaction effects were much less prevalent in area 7a, this regression analysis was performed using the less-restrictive criteria above (subsets containing at least one interaction cell and one viewer- or object-centered cell, depending on the analysis). The decoding from this data set was more noisy than those reported for our main findings above, so we square-root transformed the posterior probabilities and then converted them to Z scores (means calculated within subsets). We found that these transformations had little effect when applied to the data set used to produce our main findings above.

Discussion

The hypothesis examined in this study is that object-centered spatial representation emerges within parietal cortex as the product of a transformation from a more fundamental, retinocentric representation of spatial position. We provide several pieces of evidence that are consistent with this model.

Functional relation between viewer- and object-centered signals

The primary findings we report are that (1) neural signals coding viewer-centered position lead signals coding object-centered po-

sition (Figs. 4, 5), (2) the representations of viewer- and object-centered position are correlated in strength across trials (Fig. 6), and (3) the viewer-centered decoding time series can be used to predict the object-centered time series, but not the reverse (Fig. 7). We also show a significant association between the outcomes (success or failure) of decoding the side of the critical square successfully in viewer- and object-centered coordinates, across trials. The temporal lag and correlation between viewer- and object-centered signals is consistent with a model in which the brain derives signals that represent object-centered position by transforming signals that represent viewer-centered position. This transform could take place within area 7a or within a more broadly distributed cortical network that includes area 7a.

In this study, we first used linear discriminant analysis to decode spatial information from bins of neuronal subset activity. We then measured the correlation between the time series of posterior probabilities produced by viewer- and object-centered decoding. This approach presented several advantages. The decoding step allowed us to correlate the information represented by subset neural activity (as quantified by the posterior probability), rather than correlating firing rates directly. This is an important distinction because firing rate and information are not equivalent. For example, Figure 5A shows that after presentation of the model object, the firing rate of object-coding neurons increases ~100 ms before the activity of this population begins to carry information about the object-centered position of the critical square (as reflected in the delayed separation in the firing rate of the population on preferred and nonpreferred trials). Second, LDA provides a concise measure of the representation of the neuronal subset taken as a whole (the posterior probability). This is in contrast to, for example, a group of measures obtained for each subset quantifying the correlation in firing rate between neurons taken two at a time.

Interaction between viewer- and object-centered representation

In parietal visual neurons that possess gain fields, firing rate varies as a multiplicative interaction between eye position and retinal stimulus position (Andersen et al., 1990). Artificial neural network models have demonstrated that hidden units that are sensitive to the interaction between eye position and retinal stimulus position are capable of transforming retina-centered representations of space in the input layer into head-centered representations of space in the output layer of the network (Zipser and Andersen, 1988). We were interested in whether neurons coding the interaction between the two spatial signals we studied (viewer- and object-centered positions) participate in the transformation of one signal into the other. Consistent with this pos-

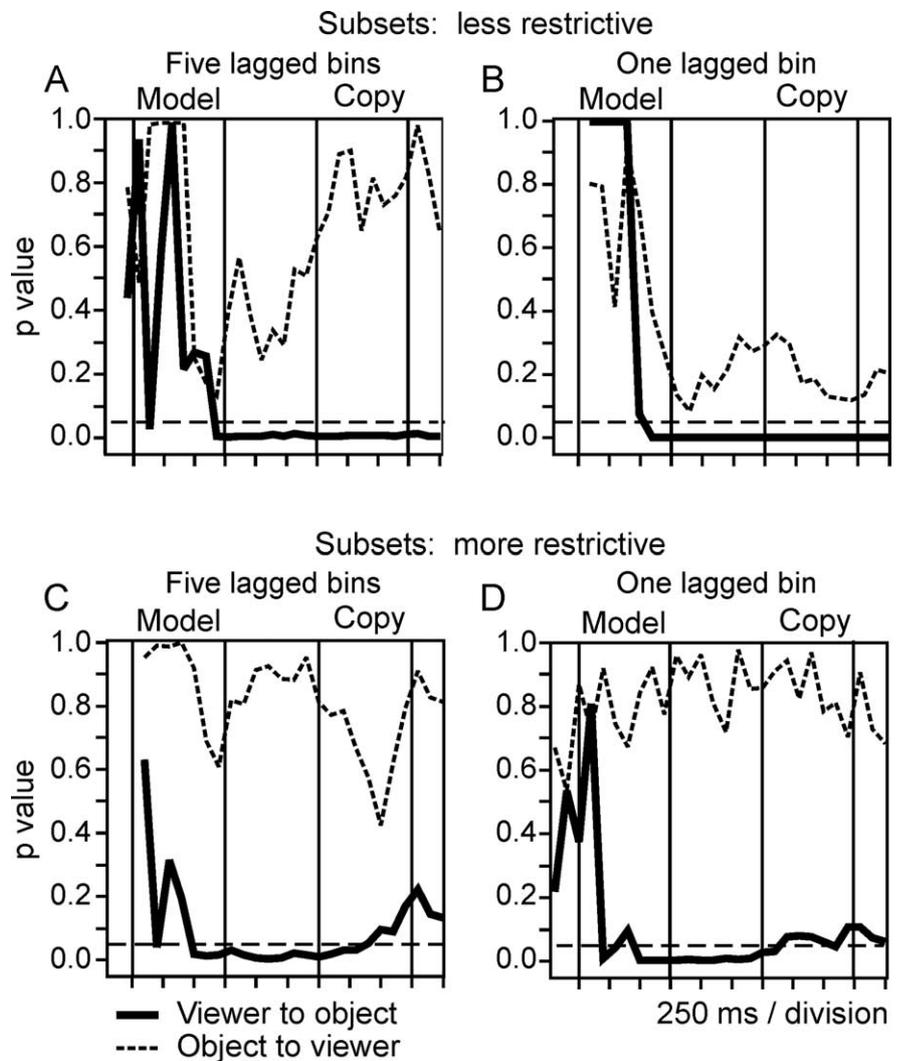


Figure 7. Predicting object representation using viewer representation. *p* values of the significance of predicting posterior probabilities of one reference frame using lagged information from the other reference frame. Solid lines are viewer signals predicting object signals. Dashed lines are object signals predicting viewer signals. *A, B*, Less-restrictive subsets. *C, D*, More-restrictive subsets. *A, C*, Five lagged bins (100 ms each) used in the prediction. *B, D*, One lagged bin used in the prediction. Subset sizes are as in Figure 6.

sibility, we found that neurons coding the interaction between the two factors were activated at a time point intermediate between the representation of viewer-centered and object-centered positions (Fig. 8). Furthermore, we found that the interaction posterior probability time course could be predicted by the viewer-centered time course and, in turn, could predict the object-centered time course.

Ability of subset activity to accurately capture network representation

Because we correlate temporal variation in the results of two parallel subset decoding analyses, our data quantify the correlation in the information coded by two neural populations over time. Our data do not (for the most part) quantify the temporal correlation in the spike trains of neurons. Our conclusions relate instead to the temporal interrelationship between two dynamic neural representations that coexist in posterior parietal cortex: coding position relative to the viewer and relative to a reference object. We consider that the neurons we happened to encounter during recording belonged to much larger populations engaged

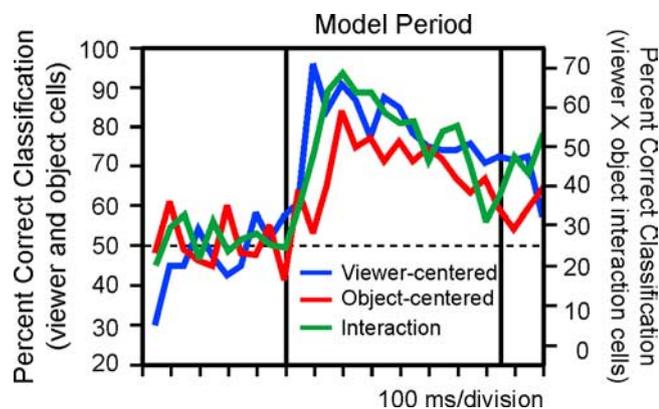


Figure 8. Time course of interaction between object- and viewer-centered representation. LDA was applied to 50 ms of population activity to classify each bin as representing right or left in viewer-centered (blue lines) or object-centered (red lines) space. The green line is the performance of LDA classification of the representation of the interaction between viewer- and object-centered positions. The left ordinate represents the percentage of trials correctly classified as left or right in these reference frames. In this case, chance classification was 50% (represented by the dashed line). The right ordinate represents the percentage of trials correctly classified to one of the four possible horizontal critical square locations (a result of the interaction of viewer- and object-centered square positions). In this case, chance classification was 25%. Neurons were included if their activity related significantly to horizontal position in viewer-centered space (blue line), object-centered space (red line), or their interaction (green line).

to sustain these neural representations. An important question therefore is the degree to which the few neurons we could record at one time could suffice to accurately capture temporal variability in the information coded by these larger populations. We found that decoding accuracy scaled with the number of neurons in which firing rate varied significantly with viewer- and object-centered positions that were included in the analysis (Fig. 4). The minimum number of neurons in a subset required to address the temporal relationship between the representation of viewer- and object-centered sides is two: one object-coding and one viewer-coding neuron recorded simultaneously. Although decoding accuracy for the side of the critical square in each spatial frame of reference was limited in this case, it was still above chance and sufficient to detect significant covariation in the representation of the two distinct spatial variables by the brain over time. The ability to detect a significant relationship between viewer- and object-centered representations over time when only one neuron of each type was present argues for (and not against) the strength of the relationship between these neural representations (as our estimate of viewer- and object-centered positions at each time point was noisier when a given ensemble contained fewer neurons coding in each framework).

Previous studies of object-centered spatial representation

The activity of neurons in the supplementary eye field (SEF) represents the object-centered direction of planned saccades (Olson and Gettner, 1995, 1999; Olson and Tremblay, 2000; Tremblay et al., 2002). Furthermore, the activity of single SEF neurons is often influenced by both eye- and object-centered saccade direction (Moorman and Olson, 2007); however, the temporal correlation in the neural signals coding direction in these alternative coordinate systems has not been assessed. A previous investigation of the neuronal representation of saccade direction in the lateral intraparietal area has indicated that parietal neurons code saccade direction in eye-centered and not object-centered coordinates, in a task dissociating these coordinate systems (Sabes et al.,

2002). Using a different task and recording in a different parietal area, we found that neurons in parietal area 7a code position relative to a reference object during the object construction task (Chafee et al., 2007), enabling the present examination of the functional relationship between simultaneously recorded viewer- and object-centered signals in parietal cortex.

Dependence of functional relation on simultaneity of recording

If viewer- and object-coding representations are functionally related, such that the object-centered representation is computed by a transform applied to the viewer-centered representation moment to moment, one would predict that the correlation between these two representations should only exist when viewer- and object-centered sides were decoded from simultaneously recorded activity. We compared the ability to predict the object-centered representation using viewer-centered representation under two conditions, one in which the two decoding time series were derived from simultaneously recorded activity, and one in which they were derived from activities recorded at different times. We found that the viewer-centered representation predicted the object-centered representation only when derived from the activity of simultaneously recorded neurons. This provides evidence in support of the hypothesis that the object-centered representation derives from a transform of the viewer-centered representation on a moment to moment basis. The directionality of this transform (viewer to object) is indicated by the finding that viewer signals predicted object signals but not the converse (Fig. 7). Other, potentially spurious sources of this linkage would not account for its dependence on simultaneity of neural activity and its directionality.

We have shown a neural correlate of a viewer- to object-centered spatial transformation in area 7a of the posterior parietal cortex. Considerable evidence from neuropsychology suggests that damage to parietal cortex causes a loss of object-centered representations, in the form of object-centered neglect (Farah et al., 1990; Driver and Halligan, 1991). In this case, patients often neglect information on the side of the object contralateral to their lesion, relatively independent of the location of the object in world-centered coordinate systems. The coexistence of viewer- and object-centered signals within parietal area 7a (Chafee et al., 2007), along with the lag and correlation in these signals we presently report, are consistent with parietal cortex playing an important role in transforming one spatial representation into the other.

References

- Andersen RA, Essick GK, Siegel RM (1985) Encoding of spatial location by posterior parietal neurons. *Science* 230:456–458.
- Andersen RA, Bracewell RM, Barash S, Gnadt JW, Fogassi L (1990) Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque. *J Neurosci* 10:1176–1196.
- Averbeck BB, Crowe DA, Chafee MV, Georgopoulos AP (2003) Neural activity in prefrontal cortex during copying geometrical shapes. II. Decoding shape segments from neural ensembles. *Exp Brain Res* 150:142–153.
- Chafee MV, Crowe DA, Averbeck BB, Georgopoulos AP (2005) Neural correlates of spatial judgement during object construction in parietal cortex. *Cereb Cortex* 15:1393–1413.
- Chafee MV, Averbeck BB, Crowe DA (2007) Representing spatial relationships in posterior parietal cortex: single neurons code object-referenced position. *Cereb Cortex* 17:2914–2932.
- Crowe DA, Chafee MV, Averbeck BB, Georgopoulos AP (2004) Neural activity in primate parietal area 7a related to spatial analysis of visual mazes. *Cereb Cortex* 14:23–34.
- Crowe DA, Averbeck BB, Chafee MV, Georgopoulos AP (2005) Dynamics

- of parietal neural activity during spatial cognitive processing. *Neuron* 47:885–891.
- Driver J, Halligan PW (1991) Can visual neglect operate in object-centred co-ordinates?: an affirmative single case study. *Cogn Neuropsychol* 8:475–496.
- Farah MJ, Brunn JL, Wong AB, Wallace MA, Carpenter PA (1990) Frames of reference for allocating attention to space: evidence from the neglect syndrome. *Neuropsychologia* 28:335–347.
- Hebb DO (1949) *The organization of behavior*. New York: Wiley.
- Johnson RA, Wichern DW (1998) *Applied multivariate statistical analysis*. Upper Saddle River, NJ: Prentice Hall.
- Ljung L (1999) *System identification, theory for the user*, Ed 2. Upper Saddle River, NJ: Prentice Hall.
- Moorman DE, Olson CR (2007) Combination of neuronal signals representing object-centered location and saccade direction in macaque supplementary eye field. *J Neurophysiol* 97:3554–3566.
- Mountcastle VB, Andersen RA, Motter BC (1981) The influence of attentive fixation upon the excitability of the light-sensitive neurons of the posterior parietal cortex. *J Neurosci* 1:1218–1225.
- Olson CR, Gettner SN (1995) Object-centered direction selectivity in the macaque supplementary eye field. *Science* 269:985–988.
- Olson CR, Gettner SN (1999) Macaque SEF neurons encode object-centered directions of eye movements regardless of the visual attributes of instructional cues. *J Neurophysiol* 81:2340–2346.
- Olson CR, Tremblay L (2000) Macaque supplementary eye field neurons encode object-centered locations relative to both continuous and discontinuous objects. *J Neurophysiol* 83:2392–2411.
- Sabes PN, Breznen B, Andersen RA (2002) Parietal representation of object-based saccades. *J Neurophysiol* 88:1815–1829.
- Snyder LH, Grieve KL, Brotchie P, Andersen RA (1998) Separate body- and world-referenced representations of visual space in parietal cortex. *Nature* 394:887–891.
- Tremblay L, Gettner SN, Olson CR (2002) Neurons with object-centered spatial selectivity in macaque SEF: do they represent locations or rules? *J Neurophysiol* 87:333–350.
- Zipser D, Andersen RA (1988) A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331:679–684.