

Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex

David A. Crowe,^{1,5} Bruno B. Averbeck,² and Matthew V. Chafee^{3,4,5}

¹Department of Biology, Augsburg College, Minneapolis, Minnesota 55454, ²Laboratory of Neuropsychology, National Institute of Mental Health—National Institutes of Health, Bethesda, Maryland 20892, ³Department of Neuroscience and ⁴Center for Cognitive Sciences, University of Minnesota, Minneapolis, Minnesota 55455, and ⁵Brain Sciences Center, Veterans Affairs Medical Center, Minneapolis, Minnesota 55417

We characterized the temporal dynamics of population activity in parietal cortex of monkeys as they solved a spatial cognitive problem posed by an object construction task. We applied pattern classification techniques to characterize patterns of activity coding object-centered side, a task-defined variable specifying whether an object component was located on the left or right side of a reference object, regardless of its retinocentric position. During a period in which the value of object-centered side, as defined by task events, remained constant, parietal cortex represented this variable using a dynamic neural code by activating neurons with the same spatial preference in rapid succession so that the pattern of active neurons changed dramatically while the spatial information they collectively encoded remained stable. Furthermore, if the neurons shared the same spatial preference, then their pretrial activity (measured before objects were shown) was correlated to a degree that scaled as a positive linear function of how close together in time the neurons would be activated later in the trial. Finally, we found that while parietal cortex represented task-critical spatial information using a dynamic neural code, it simultaneously represented task-irrelevant spatial information using a stationary neural code. These data demonstrate that dynamic spatial representations exist in parietal cortex, provide novel insight into the synaptic mechanisms that generate them, and suggest they may preferentially encode task-critical spatial information.

Introduction

It is often assumed that the brain represents information using a stationary code in which each single neural representation corresponds to a single pattern of firing rates over neurons. However, the brain could represent a singular item of information dynamically, by generating an orderly sequence of activity patterns over time (Mazor and Laurent, 2005; Meyers et al., 2008; Buonomano and Maass, 2009). To explore whether dynamic codes of spatial information exist in primate parietal cortex, characterize their relation to spatial cognitive function, and elucidate their neural mechanisms, we analyzed the temporal dynamics of neural activity in parietal cortex of monkeys solving a spatial cognitive problem posed by an object construction task. The construction task required monkeys to assemble a copy of a model object consisting of a variable configuration of identical components. To perform

the task correctly, monkeys had to evaluate object structure (configuration), compute the object-centered locations of components, and compare incomplete objects they were constructing to a model object stored in working memory—all forms of spatial cognition with clear neural correlates dissociated from sensorimotor processing in parietal cortex (Chafee et al., 2007; Crowe et al., 2008).

In the present study, we characterized patterns of population activity representing the value of a covert spatial cognitive variable, object-centered side, specifying whether a single critical component was located on the left or right side of an object regardless of its location in viewer-centered space. Defined as such, object-centered side carried a single bit of task-critical spatial information that remained constant over time within the trial. Monkeys had to compute the value of side to make the correct choice at the end of each trial, but side was uncorrelated with stimulus features (such as object form or position) and motor parameters (such as response direction or timing) during task performance.

We sought to contrast two distinct mechanisms of neural representation. If the representation of side were mediated by a stationary code, then the distribution of firing rates over neurons should remain stable over the period of time that the value of side, as defined by task events, did not change. If instead the representation of side were mediated by a dynamic code, then the distribution of firing rates over neurons would change in an orderly

Received Feb. 20, 2010; revised March 31, 2010; accepted June 27, 2010.

This work was supported by United States Public Health Service, National Institutes of Health Grants R01 MH077779 and R24 MH069675, Whitehall Foundation Grant 2005-08-44-APL, the Department of Veterans Affairs, and the American Legion Brain Sciences Chair. We thank Apostolos Georgopoulos for his intellectual contribution to and steadfast support of this work. We thank Adam Johnson for helpful discussions, Bagrat Amirikian and Sofia Sakellari for insightful comments regarding data analysis, and Dale Boeff and Dean Evans for excellent technical support.

Correspondence should be addressed to Matthew V. Chafee, Department of Neuroscience, University of Minnesota, Brain Sciences Center (11B), Veterans Administration Medical Center, 1 Veterans Drive, Minneapolis, MN 55417. E-mail: chafee001@umn.edu.

DOI:10.1523/JNEUROSCI.0954-10.2010

Copyright © 2010 the authors 0270-6474/10/3011640-14\$15.00/0

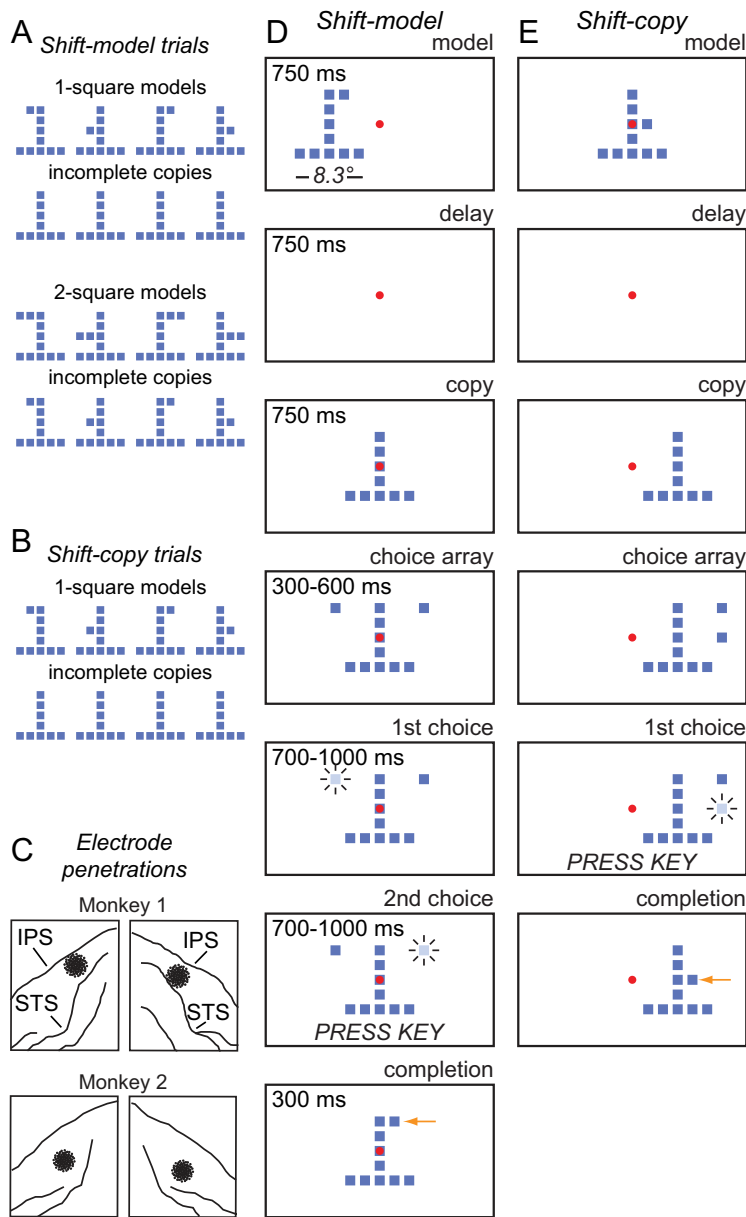


Figure 1. Visual stimuli and event sequence in the object construction task and neural recording locations in posterior parietal area 7a. **A, B**, Objects presented during the construction task. Each object consisted of an arrangement of squares (1.4° on a side) positioned within a 5×5 grid of possible locations. All objects included, at a minimum, squares within the central column and base row of the grid forming an object frame in the form of an inverted ‘T’. Model objects consisted of this frame plus either one or two additional squares placed at various locations. In each trial, the copy object was identical to the preceding model except that a single square had been removed, which we refer to as the missing ‘critical square.’ **A**, Model and copy objects presented on shift-model trials. When two additional squares were present in the same row of the model, only the outermost was removed to produce the copy object. **B**, Model and copy objects presented in shift-copy trials. **C**, Electrode penetrations on the surface of the inferior parietal gyrus in area 7a (IPS, Intraparietal sulcus; STS, superior temporal sulcus). **D**, Event sequence on shift-model trials (time advances downward, duration of each epoch indicated at upper left of each panel). The model object was presented randomly offset to the left or right over trials, whereas the copy object was always presented centered on the gaze fixation target. Choice stimuli were presented in a horizontal or vertical array at random (horizontal choice array is shown). **E**, Event sequence on shift-copy trials. The model object was presented centrally, whereas the copy object was presented randomly offset to the left or right of the gaze fixation target (vertical choice array is shown).

and systematic fashion during a period of time when the information collectively coded by that activity did not change.

In this study we do the following: (1) provide evidence that spatial representations are mediated by dynamic neural codes in parietal cortex; (2) show that dynamic codes represent task-critical but not task-irrelevant spatial information; (3) identify

specific neuronal cohorts activated in sequence to mediate dynamic spatial codes; (4) analyze correlation in pretrial activity of these neurons to provide novel insight into the cellular mechanisms that underlie the dynamic representation; and (5) characterize the neural representation of time combined with the representation of space in parietal activity.

Materials and Methods

Object construction task. We trained two male monkeys (*Macaca mulatta*, 4 and 6 kg) to perform an object construction task (Fig. 1D,E) that required them to construct a copy of a model object (Chafee et al., 2005, 2007; Crowe et al., 2008). The stimuli are similar to those used to demonstrate object-based spatial neglect in humans following damage to posterior parietal cortex (Driver et al., 1994).

Monkeys initiated trials of the construction task by directing their gaze at a central fixation target and had to maintain gaze within 1.5° of this target for the duration of the trial (eye position was monitored with an infrared eye tracking system; ISCAN). After 500 ms of central fixation, we presented a model object for 750 ms (Fig. 1D,E, model) at different positions relative to the gaze fixation target. Model objects consisted of various spatial configurations of identical square components arranged on a 5×5 square grid. Square components subtended 1.4° , and the object grid subtended 8.3° of visual angle. All model objects included a central column and base row of squares present within the grid forming an inverted ‘T’ configuration that we refer to as the object frame, in addition to which one or two additional squares were placed at various positions (Fig. 1A,B). A delay period (750 ms) ensued after the disappearance of the model object. We then presented an incomplete copy object (Fig. 1D,E, copy), at different positions relative to the gaze fixation target. The copy object was identical to the preceding model object on each trial except that a single square, which we refer to as the missing ‘critical square,’ had been removed. The construction task required monkeys to compare the configuration of the copy object, while it was visible, to the configuration of the model object stored in working memory to localize the missing critical square. Monkeys then had to add a square at this position to the copy object to ‘construct’ or reproduce the configuration of the model object presented earlier in the trial. Replacement of the missing critical square was accomplished via a forced choice mechanism. After the copy object had been visible for a period of 750 ms, we presented a pair of choice squares flanking the copy object. The choice squares were arrayed either horizontally (Fig. 1D,

choice array) or vertically (Fig. 1E, choice array) at random over trials. The monkey selected one of the two choice squares to add to the copy object by timing when it depressed a single response key (a pedal it pressed with its left foot) in relation to a randomized choice sequence as follows. After a random period (300–600 ms), we increased the brightness of one choice square at a time in random order for a period of

between 700 and 1000 ms each. The choice square that was bright at the time that the monkey pressed the response key automatically translated in a horizontal direction inward to join the copy object, producing a new configuration. The direction of the required motor response did not vary over trials—the timing of the response relative to the choice sequence determined whether the trial was successful or not. If the monkey pressed the pedal at the right time it selected the correct choice square, the addition of which to the copy object reproduced the model configuration (in which case the monkey was rewarded with a 0.1 ml drop of juice). If the monkey pressed the pedal at the wrong time it added the incorrect choice square to the copy object, and the new object configuration that resulted did not match that of the preceding model object (in which case the trial terminated without reward). The choice sequence was randomized, so that on one half of trials the first choice in the sequence was the correct choice (Fig. 1E, 1st choice), whereas in the remaining half of trials the second choice was correct (Fig. 1D, 2nd choice). The timing of the motor response therefore was not predictable in advance of the choice sequence.

On shift-model trials (Fig. 1D), we randomly shifted the center of the model object to the left or right of the gaze fixation target so that the model object (8.3° wide) was presented entirely in the left or right visual hemifield, and then the copy object was subsequently presented centered on the gaze fixation target. On shift-copy trials (Fig. 1E), we presented the model object centered on the gaze fixation target and then randomly shifted the horizontal position of the copy object relative to the gaze fixation target. We used the same two eccentric object locations to record the activity of all neurons without adjusting the locations of the reference objects to fall at the receptive field centers of individual neurons (as we typically recorded the activity of between 20 and 30 neurons simultaneously). We selected the two fixed eccentric locations of the reference objects to place the objects as close to the fovea as possible, while still presenting them entirely in the right or left visual hemifield, to minimize the effect of reduced visual acuity at more peripheral locations and make it easier for monkeys to analyze object structure without looking at the objects directly. Each neuronal ensemble was recorded during the performance of either shift-model or shift-copy trials. Shifting the position of reference objects relative to the gaze fixation target (which delimited the midline of eye, head, and body-centered spatial frameworks) allowed us to dissociate a spatial variable referring to the position of the critical square, which we refer as “side” in two distinct spatial frameworks. Object-centered side specified whether the critical square was located on the left or right side of the object with respect to its intrinsic midline. Viewer-centered side specified whether the critical square was located to the left or right of the gaze fixation target.

Set of model objects included. We presented monkeys with a variety of model object configurations (Chafee et al., 2005) but restricted the present analysis to a simple subset in which either only one square was present in addition to the object frame (Fig. 1A, B, 1-square models), or two additional squares were present in the same row and on the same side of the object (Fig. 1A, 2-square models). The monkey was able to anticipate the position of the critical square as soon as these model objects were presented, because each configuration included only a single square that was ever removed to produce the copy object on a given trial. When one square was present in addition to the frame in the model object, we removed this one square to produce the incomplete copy (Fig. 1A, B, compare 1-square models to corresponding incomplete copies). When two squares were present in addition to the frame in the same row and on the same side of the model object, we removed the outermost of the two squares to produce the incomplete copy (Fig. 1A, compare 2-square models to corresponding incomplete copies). We have shown previously that population activity codes the position of the critical square from the model period onward in the trial when this set of model objects was presented (Chafee et al., 2005), allowing us to examine the neural representation of the object-centered side of the critical square over an extended period of time within the trial.

Neural recording. We recorded neural activity using a multielectrode matrix containing 16 independently movable, glass-coated, platinum–tungsten metal microelectrodes (Thomas Recording). Typically, we would advance the electrodes individually until we had isolated the action potentials of between 20 and 30 neurons and then record the elec-

trical activity of the individually isolated neurons within this ensemble as the monkey performed a set of 128 or 160 trials of the object construction task. We isolated the action potentials of single neurons using either a waveform template matching system (Multi-Spike Detector, Alpha Omega Engineering) or time–amplitude window discriminators (DDIS-1, Bak Electronics). After we had recorded neuronal activity during a complete trial set, we would then advance the electrodes to study another neural ensemble. Additional details of surgery and recording technique can be found in our prior reports (Chafee et al., 2005, 2007). Care and treatment of the animals conformed to the Principles of Laboratory Animal Care of the National Institutes of Health (NIH) (NIH publication no. 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.

Neural database. The present results are based on the activity of 69 simultaneously recorded neural ensembles containing a total of 1601 area 7a neurons. We recorded 51 ensembles containing 1097 neurons from two monkeys in shift-model trials (Fig. 1D) and 18 ensembles containing 504 neurons from one monkey in shift-copy trials (Fig. 1E). Locations of neural recording in area 7a are provided in Figure 1C.

Data analysis. We evaluated whether neuronal activity during the model or copy period of the task reflected the side of the critical square in either object-centered or viewer-centered coordinates. For that purpose we performed a two-way ANCOVA in which the factors were the object-centered side (left or right relative to the object midline) and the viewer-centered side (left or right relative to the gaze fixation target) of the critical square. Because the critical square was located on the left or right side of the reference object at random and the reference object was presented to the left or right of the gaze fixation target at random, object-centered side and viewer-centered side were statistically independent. For shift-model trials, we evaluated the influence of object- and viewer-centered sides on neuronal activity during the model period. For shift-copy trials, we evaluated the influence of these factors on neuronal activity during the copy period. Firing rates during the premodel fixation period and the elapsed time since the start of data collection were included as covariates in the ANCOVA.

Using pattern classification to decode the side of the critical square from neuronal population activity. We represented the pattern of population activity on each trial as a vector consisting of the firing rates of individual neurons and applied a pattern classification analysis (Klecka, 1980; Johnson and Wichern, 1998; Averbeck et al., 2003; Chafee et al., 2005, 2008) to classify each population activity pattern as coding left or right, effectively decoding side from population activity. We performed the classification using fivefold cross validation. On each of five iterations, we divided trials into a training set (4/5 trials) and a test set (1/5 trials) (Fig. 2A). We used neural activity on the training trials to define the classification functions, which in this case were multivariate normal probability densities modeling the distribution of activity vectors within the two groups of trials when the critical square was left or right. To compute the classification functions, we further divided training trials into two groups on the basis of the side of the critical square and then averaged the firing rate of each neuron within each group of training trials. This provided the mean population activity vectors (patterns) associated with the neural representation of left and right. To measure the degree of spread of activity patterns around each mean, we computed the pooled within-groups covariance matrix. The mean population activity vectors on left and right trials and the pooled covariance matrix provided the free parameters of the two multivariate normal probability density distributions we used to model the distribution of activity patterns on left and right trials.

We then classified the population activity pattern measured on each test trial as coding either left or right. We computed the distance in rate space between the population activity vector on that test trial and the mean activity vectors coding left and right in the training data. We then converted these two distances to posterior probabilities under the assumption that the distribution of activity patterns on left and right trials was multivariate normal and classified the test trial to the side (left or right) associated with the greater posterior probability. We then repeated the procedure selecting the next 1/5 trials to serve as test trials and the remaining 4/5 trials to serve as training trials until all trials had been

A Five-fold cross-validation

ITERATION 1	ITERATION 2	ITERATION 5
Trial 1 ○ TEST	● TRAINING	● TRAINING
Trial 2 ○ TEST	● TRAINING	● TRAINING
Trial 3 ○ TEST	● TRAINING	● TRAINING
Trial 4 ● TRAINING	○ TEST	● TRAINING
Trial 5 ● TRAINING	○ TEST	● TRAINING
Trial 6 ● TRAINING	○ TEST	● TRAINING
Trial 7 ● TRAINING	● TRAINING	● TRAINING
Trial 8 ● TRAINING	● TRAINING	● TRAINING
Trial 9 ● TRAINING	● TRAINING	● TRAINING
Trial 10 ● TRAINING	● TRAINING	● TRAINING
Trial 11 ● TRAINING	● TRAINING	● TRAINING
Trial 12 ● TRAINING	● TRAINING	● TRAINING
Trial 13 ● TRAINING	● TRAINING	○ TEST
Trial 14 ● TRAINING	● TRAINING	○ TEST
Trial 15 ● TRAINING	● TRAINING	○ TEST

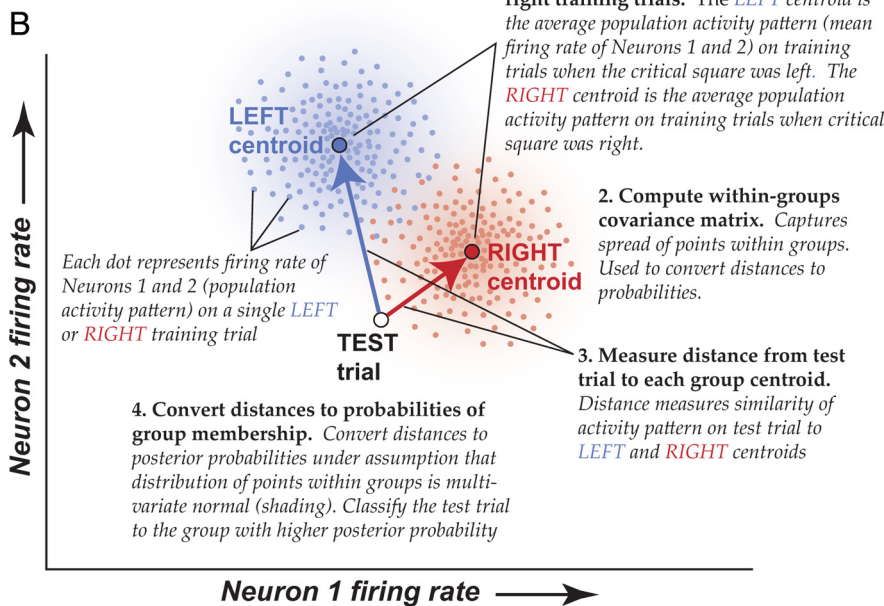


Figure 2. Graphical illustration of pattern classification analysis applied to a simplified case in which side is decoded from a population containing two neurons. **A**, Cross-validation procedure. Trials are divided into a test set (1/5 trials) and a training set (4/5 trials). Classification functions are defined by neural activity on training trials and used to classify each test trial as coding left or right. The classification is repeated five times, treating successive 1/5 trials as test trials until all are classified. **B**, Graphical illustration of classification based on a two-dimensional (e.g., two neuron) case. Each trial is represented by a point in the two-dimensional rate space with coordinates equal to the firing rates of neurons 1 and 2. Points are colored according to the object-centered side of the critical square. The separation of the group centroids, relative to the within-groups covariance, indicates the strength of the population signal coding side.

classified. The proportion of trials correctly classified by this procedure measured the strength of the population signal coding side. We performed the analysis assuming equal prior probabilities for left and right. We used the classify function in the Matlab Statistics Toolbox to perform this analysis (The MathWorks).

To provide a conceptual illustration of the classification analysis, we address the hypothetical case of decoding object-centered side from the pattern of activity in a population containing two neurons. The two-neuron case has the graphical advantage that the pattern of activity in the network can be expressed by a point in a two-dimensional rate space, the coordinates of which comprise a two-element vector corresponding to the firing rates of neurons 1 and 2 (Fig. 2B). The same principles apply if scaled up to a larger population of neurons, in which case the activity state of the network would be represented by a point in a multidimensional space with dimensions equal to the number of neurons, and the coordinates of the point along each axis would be given by a vector comprised of the firing rate of each neuron.

The first step in the classification procedure is to determine the average activity vectors associated with left and right trials in the training data. In the two-dimensional case, this can be visualized by plotting the activity vector of the population on each training trial as a point in the rate space and color coding the points according to the side of the critical square (Fig. 2B), producing one cloud of points associated with left (blue), and another cloud of points associated with right (red). The two clouds of points are separated to some degree in the rate space, indicating that the pattern of activity in the population differed as a function of side. The average activity vector associated with left in the training data is the center (centroid) of the blue cloud, with coordinates equal to the firing rate of neurons 1 and 2 averaged across all left trials. Likewise, the average activity vector associated with right is the center of the red cloud. The spread of points within the left and right trial clusters relative to their respective centroids is quantified by the within-groups covariance matrix, which measures variability in the activity pattern on trials with the same value of side or, more specifically, the degree to which firing rates of all possible pairs of neurons deviate together from their group centroid. The capacity of the brain to discriminate left from right based on the pattern of activity in this population is a function of the distance between the two group centroids relative to the within-groups covariance. If the centroids are widely separated in the rate space relative to the within-groups covariance, the population signal coding side is relatively strong and decoding accuracy will be correspondingly accurate. If the two distributions of activity patterns largely overlap or the distance between centroids is small relative to the covariance, the population signal coding side is weak and decoding accuracy will approach chance classification.

To illustrate how test trials are classified, we plot the position of the activity vector on a test trial in the rate space of the population (Fig. 2B, TEST trial) and determine the distance from this point to each group centroid computed from the training data (Fig. 2B, red and blue arrows). Distance between two points in the rate space is a measure of the dissimilarity between the activity vectors associated with the points. Consequently, the activity vector on the test trial is most similar to the nearest group centroid (Fig. 2B, RIGHT centroid). After converting the distances between the activity vector on the test trial and each group centroid to posterior probabilities, we then classify the test trial to the group, left or right, associated with the greater posterior probability. We performed the classification assuming equal covariance matrices within each group, which has been shown to extract all of the available information (Averbeck and Lee, 2006).

Contrasting stationary and dynamic neural codes. To determine whether the neural code for side was stationary or dynamic in time, we performed a time-resolved decoding analysis (Fig. 3). We divided each test trial into a sequence of 100 ms time bins (which we refer to as test bins). We then classified the population activity vector in each test bin as coding either left or right (Fig. 3A). This produced a sequence of classifications for each test trial, only some of which were correct (Fig. 3B). The proportion of test trials correctly classified as left or right in each test bin of the test data provided a measure of the strength of the representation

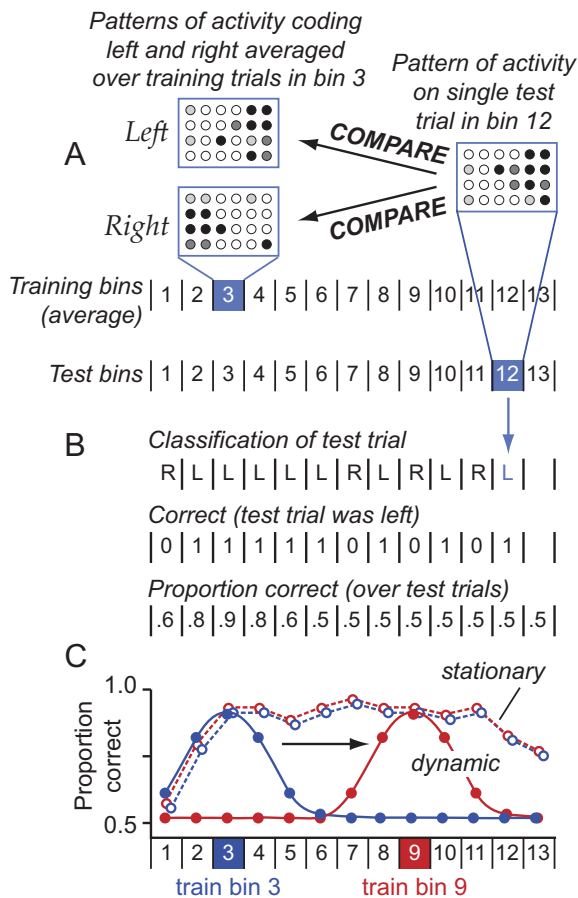


Figure 3. Time-resolved decoding analysis based on pattern classification. **A**, Grids represent a hypothetical population of neurons; firing rate is indicated by shading (dark, higher rate). The classifier is trained by computing the average activity patterns coding left and right (as well as the covariance matrix across neurons) using firing rates measured within a single time bin on training trials (bin 3 is the training bin in this example). The pattern of activity found in each 100 ms test bin of the test trial is classified based on its relative similarity to the two patterns coding left and right in the training bin. For example, the activity pattern in test bin 12 is most similar to the activity pattern coding left in training bin 3. The activity pattern in bin 12 is therefore classified as coding left. **B**, This produces a sequence of left/right classifications for each test trial that are recoded as correct (1) or error (0) depending on whether the decoded value of side (left or right) in each time bin matched the actual side of the critical square on that test trial. The proportion of test trials correctly classified in each time bin provides a decoding time course for the training bin used. **C**, The shapes of decoding time courses generated using different time bins to train the classifier indicate whether patterns of activity coding side are stationary or dynamic. Blue lines indicate decoding time courses obtained using bin 3 to train the classifier. Red lines indicate decoding time courses obtained using bin 9 to train the classifier. If parietal cortex employed dynamic patterns of activity to code side, each decoding time course would exhibit a peak that remained aligned to the training bin as it was stepped through the trial (solid symbols and lines). If parietal cortex employed stationary patterns of activity to code side, the decoding time courses would largely overlap regardless of the training bin used (open symbols, dashed lines).

of side by population activity at each time point (Fig. 3B). To perform the above classification, we used firing rates measured within a single 100 ms time bin on training trials (which we refer to as the training bin) to compute the mean population activity vectors associated with left and right and the within-groups covariance matrix, thereby defining the classification functions and training the classifier (Fig. 3A). Using classification functions based on population activity measured in a single time bin in the training data, we defined the population activity patterns associated with the neural representation of left and right at a single time point within the trial. For each individual training bin in the training data, we generated an entire decoding time course quantifying the proportion of test trials correctly classified within each test bin throughout the trial

(Fig. 3A, B). This compared population activity patterns at all time points in the test data to population activity at a single time point in the training data. Fluctuations in the accuracy of decoding over time bins in the test data reflected the degree to which the activity patterns associated with left and right in each time bin of test trials resembled those associated with left and right in the single training bin on training trials. This in effect provided a measure of the lifespan of the specific population activity patterns coding side in the training bin. During times in the test data when decoding accuracy was comparatively high, population activity patterns coding left and right resembled those coding left and right in the training bin. When decoding accuracy was comparatively low, population activity patterns coding left and right on test trials no longer resembled those coding left and right within the training bin. Having generated an entire decoding time course for a given training bin, we advanced the training bin and repeated the procedure, generating a new decoding time course.

The shapes of the resulting decoding time courses generated by successive training bins allowed us to contrast stationary and dynamic population codes for side. We could expect that for the period of time that parietal cortex used stationary population codes to represent left and right, decoding time courses generated from different training bins would essentially overlap, because in a stationary representation the same population activity patterns would be present to code left and right in each time bin. Consequently, which training bin had been selected to train the classifier would have little effect. Overall decoding accuracy might fluctuate to reflect variation in the strength with which population activity coded side at each time point, but peaks and valleys in the decoding time courses generated from different training bins should align (Fig. 3C, dashed lines). On the other hand, for the period of time that parietal cortex used dynamic population codes for side, we could expect that decoding time courses generated from different training bins would tend to exhibit a clear peak in accuracy that remained aligned to each training bin. As the training bin was stepped through the trial, a new peak in decoding accuracy would emerge at each step, aligned to the new training bin, so that the peak in decoding accuracy would travel through the trial to follow the training bin at each step (Fig. 3C, solid lines). The rise and fall in decoding accuracy to either side of each training bin would be evidence that patterns of population activity coding left and right in that training bin were short lived. The rising phase of the decoding time course leading up to the training bin would reflect the emergence of a new pair of activity patterns in the population associated with the neural representation of left and right. The subsequent falling phase of the decoding time course after that training bin would reflect the disappearance of those activity patterns in the population. The staggered peaks in decoding accuracy associated with successive training bins in this way reflect a succession of population activity patterns briefly emerging and then fading to carry the same spatial information over time throughout the trial (Fig. 3C, solid lines), specifically whether the critical square was located on the left or right side of the reference object. For this analysis, we included either all 504 neurons recorded during shift-copy trials or, alternatively, all 1097 neurons recorded during shift-model trials in the population without preselection.

Identifying neurons engaged to represent side in each training bin. Population decoding accuracy relies on reliable differences in the firing rates of single neurons as a function of the decoded variable. Given a particular decoding time course at the population level, there should then exist within the population a subset of single neurons that exhibited a difference in firing rate as a function of side with a similar time course. To identify which subset of neurons provided neural activity that most directly accounted for the shape of a given population decoding time course, we first subtracted the mean firing rate of each neuron on left and right trials in each time bin. This produced an activity difference time course measuring variation in the strength of the signal coding side in that neuron. We then computed the correlation coefficient between the activity difference time course for that neuron and each population decoding time course. If the correlation was significant ($p < 0.0001$), we assigned the neuron to the training bin that generated the decoding time course, as the correlation provided evidence that the neuron was one of those contributing most directly to the population activity patterns cod-

ing side in the training bin. We describe neurons assigned to a particular time bin by this procedure as being significantly engaged in the representation of side at that time point. Neurons could be assigned to multiple time bins by this procedure. For the purpose of assigning neurons to time bins, we again included either all 504 neurons recorded during shift-copy trials or, alternatively, all 1097 neurons recorded during shift-model trials in the population for the correlation analysis without preselection.

Measuring correlated noise in pretrial activity. We evaluated the degree to which spontaneous firing rate was correlated over trials in pairs of neurons that had been assigned to a time bin in the dynamic representation of side by the above procedure. To assess correlation in activity, we measured firing rates during a pretrial period spanning the 500 ms of central gaze fixation before the presentation of the model object on each trial. We used pretrial activity before objects were shown and neurons were engaged to represent side or other trial-specific information to ensure that correlations were not due to neurons having similar responses to task events. This measure of correlation is often referred to as noise correlation to distinguish it from signal correlation, which would be correlations due to changes in rates brought about by task-related responses. We compared the number of neuron pairs with significantly correlated pretrial activity when the two neurons were recorded simultaneously and at different times. (To correlate pretrial activity in pairs of neurons that were not recorded simultaneously, we matched trials according to the repetition of a given combination of model and copy objects presented within the trial set.)

We evaluated correlation in the pretrial activity of pairs of neurons that had been assigned to the same time bin as a function of their spatial preference for the side of the critical square. We also evaluated the correlation in pretrial activity of pairs of neurons that were assigned to different time bins, defining the bin lag as the number of 100 ms steps between the time bins containing the two neurons. In the latter analysis, we excluded pairs of neurons that were assigned to both time bins at nonzero lags (to constrain the analysis to pairs in which one neuron was engaged in the representation of side before or after the other neuron at a given bin lag) and included neuron pairs in which the correlation in pretrial activity was significantly positive ($p < 0.05$).

Decoding time. To characterize temporal information present within the population signal coding object-centered side, we applied pattern classification analysis to decode the current time within the trial (specifically, the number of the current time bin; the bins are numbered 1–40 from the start of the trial) from the population activity pattern in each bin (chance classification was 1/40 or 0.025 correct). For this decoding analysis, we limited the population to cells that exhibited a significant correlation between their activity difference time course and a population decoding time course for object-centered side at $p < 0.0001$. Furthermore, to reduce the possibility that our estimate of time information in population activity was contaminated by neural signals coding the retinocentric positions of objects (which varied between model and copy periods and, therefore, as a function of time), we excluded neurons from the population if their activity related ($p < 0.1$) to the viewer-centered side (left or right relative to the gaze fixation target) of reference objects during either the model or copy periods (based on ANCOVA with viewer-centered side as a factor).

Results

We recorded neural activity in parietal cortex of monkeys performing an object construction task and applied pattern classification techniques to population activity to decode the value of a binary spatial cognitive variable, object-centered side, specifying whether a critical object component was present on (or missing from) the left or right side of a reference object. An example of the activity of a single neuron in area 7a coding object-centered side is illustrated in Figure 4. Activity increased during the copy period when the critical square was missing from the right side of the copy object relative to the midline of the object (Fig. 4*A,B*), regardless of whether the critical square (and reference object) was located to the left (Fig. 4*A*) or right (Fig. 4*B*) of the gaze fixation target and consequently regardless of whether the critical

square fell in the left or right half of viewer-centered space. In a two-way ANCOVA (with object-centered and viewer-centered side as factors), the firing rate of this neuron during the copy period varied significantly as a function of object-centered side ($F_{\text{object}} = 56.40$, $p < 0.001$) and not viewer-centered side ($F_{\text{viewer}} = 2.67$, $p = 0.104$) or the interaction ($F_{\text{inter}} = 2.13$, $p = 0.146$). Therefore, this neuron encoded “right” with respect to the object in a way that could be statistically dissociated from retinocentric (viewer-centered) position, a physiological property common to many parietal neurons active during this task (Chafee et al., 2007). Our prior work has provided evidence that this neural signal does not reflect the form or position of the objects displayed, the position of an attended visual stimulus, or the direction of a planned saccade (Chafee et al., 2005), but appears instead to encode spatial information needed to solve the cognitive problem posed by the object construction task.

The value of object-centered side, as defined by the model and copy objects shown and quantified by our decoding analysis, remained fixed throughout the duration of each trial of the construction task. In the present study, we sought to determine whether parietal cortex represented the same value of object-centered side (left or right) on each trial using a stationary neural code, in which case the distribution of firing rates over neurons would remain relatively stable, or a dynamic neural code, in which case the pattern of population activity would continuously evolve.

We present evidence in the following order that: (1) parietal cortex generated a dynamic neural representation to represent the same fixed value of object-centered side over time; (2) distinct groups of neurons with the same object-centered spatial preference were activated in rapid succession to mediate the dynamic code; (3) the neurons recruited into this process exhibited correlated pretrial activity that predicted the relative timing of their sequential activation; (4) the temporal pattern of population activity (dynamic or stationary) varied as a function of whether the coded spatial information was relevant or irrelevant to the current cognitive goal; and finally (5) population activity patterns integrated both spatial and temporal information.

Parietal cortex generates dynamic neural representations of space

To visualize how neural firing rates in the population of neurons coding object-centered side were changing over time during the trial, we first measured the firing rate of 96 neurons coding object-centered side every 100 ms and then projected the resulting population activity vectors into a lower dimensional space defined by the first three principal components of population activity to produce neural trajectories illustrating how the distribution of firing rates over neurons was changing over time. Neural trajectories on object-left and object-right trials were separated in the principal components space (supplemental Fig. 1, available at www.jneurosci.org as supplemental material) as predicted by the selection of neurons coding object-centered side. However, the neural trajectories also exhibited prominent loops around the model and copy periods of the trial when the distribution of firing rates across neurons was changing rapidly, suggestive of a dynamic code. Had the population coded left and right using stationary patterns of activity, the projection of activity into the principal component analysis space would have produced two points or restricted clusters of points on right and left trials, rather than the extended trajectories shown (supplemental Fig. 1, available at www.jneurosci.org as supplemental material).

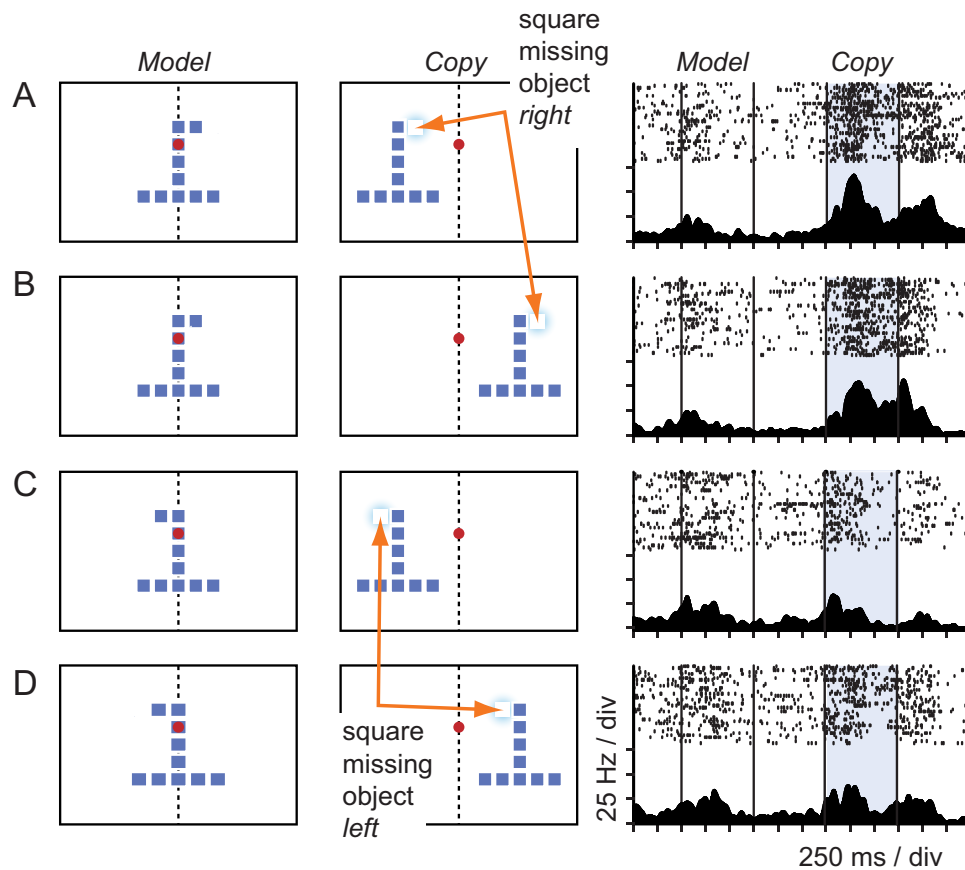


Figure 4. Activity of a single area 7a neuron coding the object-centered side of the missing critical square during the copy period. Data are from shift-copy trials: model objects were centered on the gaze fixation target, copy objects were randomly offset to the left or right of the gaze fixation target over trials. **A, B**, On trials in which the critical square was located on the right side of the model (left panel), the critical square was missing from the right side of the copy object (middle panel, orange arrows). The neuron increased its firing rate during the copy period in this case (right panel) to a comparable degree whether the copy object was presented in the left (**A**) or right (**B**) visual hemifield. Note that the retinal coordinate of the missing critical square, as indicated by the orange arrows, was never occupied by a part of the model object because it appeared at a different location in the display. The trials included in the rasters included a mixture of model objects in which the critical square was located in either the top or middle row of the object grid on the right side of the object. **C, D**, On trials in which the critical square was located on the left side of the model (left panel), the critical square was missing from the left side of the copy object (middle panel; orange arrows). The neuron was less active during the copy period in this case despite the fact that the same copy objects were presented at the same retinocentric locations as when the missing square was on the right, preferred side of the object.

To more directly measure the spatial information coded by population activity patterns and evaluate the possibility that the population code for object-centered side was dynamic in time, we employed a time-resolved pattern classification analysis (Fig. 3) to decode the value of object-centered side (left or right) using patterns of neural activity measured in successive 100 ms time bins. Using a single time bin to train the classifier, we generated a decoding time course measuring the degree to which activity patterns coding left and right in each time bin of test trials resembled those coding left and right within the single training bin on training trials.

Using this analysis, we found that decoding time courses exhibited a peak in accuracy that remained aligned to the training bin as it was stepped through the model or subsequent copy periods (Fig. 5A,C,D, each vertical dashed line indicates the training bin used to generate the decoding time course of matching color). This provided evidence that the neural code for side was dynamic during these task periods (Fig. 3C, solid lines). The rise in decoding accuracy leading up to each training bin reflected the emergence of a new pair of activity patterns in the population coding left and right. The fall in decoding accuracy after each training bin indicated the disappearance of these activity patterns. The emergence of new peaks in decoding accuracy each

shifted forward in time as the training bin advanced provided evidence that the population continuously generated new patterns of activity to represent the same spatial information (the side of the critical square) over time.

The results illustrated in Figure 5 were obtained by applying the decoding analysis to neural activity on shift-copy trials. We obtained similar results applying the analysis to neural activity recorded on shift-model trials, although the tendency for the peak to track the training bin later in the trial was less prominent during the copy period (supplemental Fig. 2, available at www.jneurosci.org as supplemental material). For both data sets, animations revealed a traveling wave in decoding accuracy that followed the training bin (at the vertical line) as it was swept through the trial (supplemental movies 1 and 2 for shift-copy and shift-model data, respectively, available at www.jneurosci.org as supplemental material). When the training bin was positioned within the delay period, decoding time courses had a broader profile and essentially overlapped regardless of the training bin used, indicating slower population dynamics (Fig. 5B; supplemental Fig. 2B, available at www.jneurosci.org as supplemental material).

Two sources of variability in the population representation contributed to the temporal pattern of population decoding we

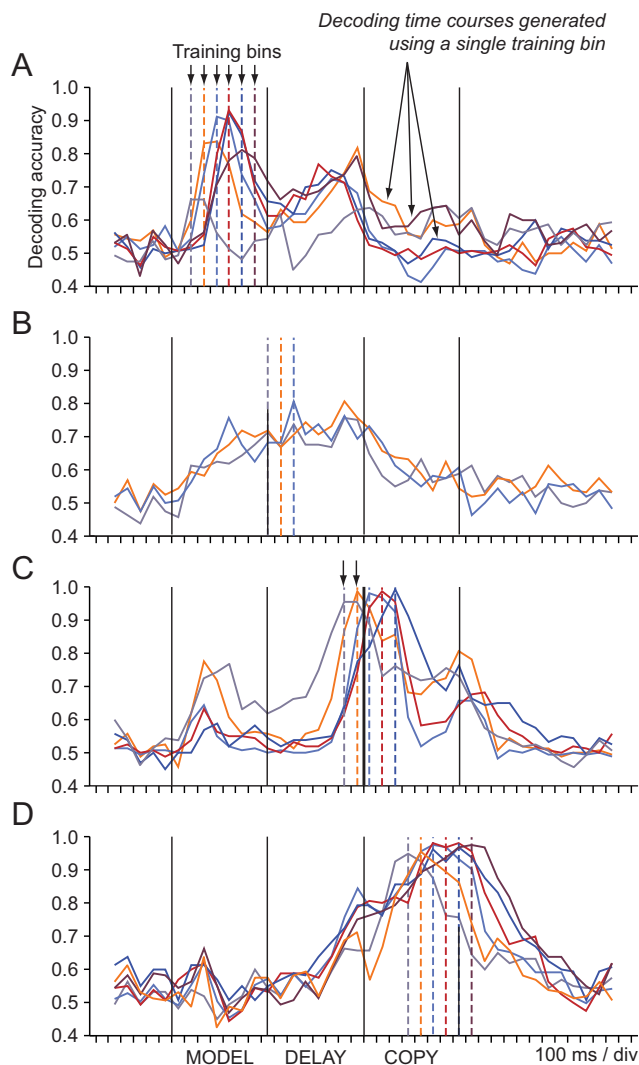


Figure 5. Population decoding time courses for object-centered side. Each line illustrates the decoding time course produced using a single time bin to train the classifier (the time of the training bin for each time course is indicated by the dashed vertical line of the same color). Fluctuations in each time course measure the varying degree to which activity patterns coding left and right at all time points resembled those coding left and right in the training bin. Neural data are from shift-copy trials. **A–D**, Groups of decoding time courses generated when the training bin was within (or near to) the model period (**A**), delay period (**B**), early copy period (**C**), or late copy period (**D**) of the trial. Peaks in decoding accuracy newly emerging just before presentation of the copy object are indicated by vertical arrows in **C**.

obtained. The first was variation in the strength with which a specific pair of activity patterns found within a single training bin was associated with left and right over time within the trial (as indicated by the rise and fall in decoding accuracy before and after each training bin as described above). The second was variation in the strength with which population activity represented side at different times in the trial overall, as indicated by variation in the height of the peak in decoding accuracy obtained in the decoding time course generated by each successive training bin (Fig. 5A) (for example, peak decoding accuracy increases as the training bin is stepped through the model period). To measure variation in the overall strength of population representation more directly, we generated a decoding time course indicating the proportion of test trials correctly classified in each time bin when the classifier was trained using neural data on training trials in the same time bin (supplemental Fig. 3, available at www.jneurosci.org).

This analysis shifted the training bin along with the test bin at each step, allowed any neurons to participate in the population representation at each time point, and so measured variation in the strength with which any activity patterns in the population were associated with the value of side in each time bin. This procedure therefore captured fluctuation over time bins in the overall strength of representation of side by the population taken as a whole. We found that although the population representation so defined varied in strength over time, the representation of side supported by the specific activity patterns found in each training bin fluctuated more rapidly, as indicated by a faster rise and fall in decoding accuracy in the time course generated from each training bin relative to the overall population representation, even during the model period (supplemental Fig. 3A, available at www.jneurosci.org as supplemental material). Therefore, the population representation was dynamic. New patterns of activity rapidly emerged to mediate the population representation of side at successive time points, even during periods in which the overall strength of population representation was also changing, but at a slower rate.

Rapid population dynamics were prevalent around the time that visual objects were shown, raising the possibility that they reflected application of the decoding analysis to transient visual signals in parietal cortex evoked by presentation of the objects. However, on shift-copy trials a new peak in decoding accuracy emerged at the end of the delay period during which the neural representation of side was relatively stationary, just before the copy object appeared (Fig. 5C, arrows identify peaks in decoding accuracy in time bins preceding the onset of the copy object). As this population dynamic occurred in advance of the visual input, it was internally initiated and not visually driven.

Distinct groups of neurons are activated in rapid succession to mediate the dynamic code

Decoding of object-centered side at the population level (Fig. 5) depended on the degree to which the firing rate of single neurons differed reliably on right versus left trials. To identify the individual neurons that most directly contributed to the newly emerging patterns of population activity coding side in each training bin of the decoding analysis, we subtracted the firing rate for each neuron on right and left trials and correlated the resulting activity difference time course with the population decoding time course generated using a given time bin to train the classifier. The procedure identified neurons in which the difference in firing rate as a function of side followed a time course that closely resembled the time course of decoding accuracy at the population level associated with a given training bin (Fig. 6).

For example, the difference in firing rate as a function of side in neuron A (Fig. 6C, top, black line) followed a time course that was significantly correlated ($p < 0.0001$) with the population decoding time course generated when using bin 13 to train the classifier (Fig. 6C, top, red line). We assigned neuron A to bin 13 on that basis. The difference in firing rate as a function of side in neurons B and C in contrast (Fig. 6C, middle and bottom, black lines) followed time courses that were significantly correlated ($p < 0.0001$) with the population decoding time course generated when using activity in bin 19 to train the classifier (Fig. 6C, middle and bottom, blue lines). We therefore assigned neurons B and C to bin 19. The activity difference time course averaged over neurons assigned to the same time bin (Fig. 6D, black lines) closely matched the decoding time course generated when that time bin was used to train the classifier (Fig. 6D, orange lines).

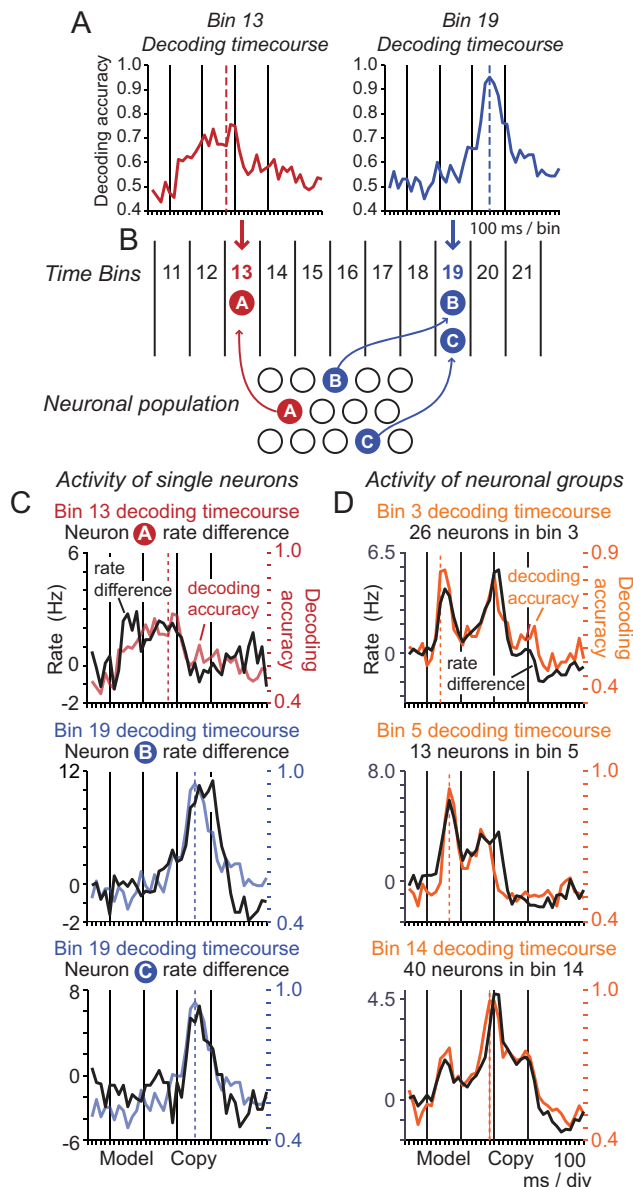


Figure 6. Assigning neurons to time bins based on correlation between firing rate and decoding accuracy. **A**, Population decoding time courses generated using bins 13 and 19 to train the classifier. **B**, Neurons were assigned to the training bin that had produced a given decoding time course if decoding accuracy over time and firing rate over time were significantly correlated ($p < 0.0001$). **C**, Correlation between decoding accuracy and firing rate. Black lines in each panel indicate the difference in firing rate of a single neuron on preferred and nonpreferred side trials. The activity difference time course of neuron A (top, black line) closely resembled the decoding time course produced using bin 13 to train the classifier (top, red line). Neuron A was assigned to bin 13 on that basis. The activity difference time courses of neurons B and C (middle and bottom, black lines) closely resembled the decoding time courses produced using bin 19 to train the classifier (middle and bottom, blue lines). Neurons B and C were therefore assigned to bin 19. **D**, Similarity between the average activity difference time course of neurons assigned to the same time bin (black lines), and the population decoding time course generated using the same time bin to train the classifier (orange lines). Examples using three training bins are illustrated (the number of the training bin relative to the start of the trial and the number of neurons assigned to the training bin are shown above each panel). div, Division.

The assignment of a neuron to a training bin by this procedure indicated not that the neuron was only active to code side at that single time point, but rather that the neuron was one of those driving fluctuations in decoding accuracy over the entire trial in the particular pattern observed when the time bin in question was

selected to train the classifier. Consequently we could conclude that the neuron was one of the those participating in the population activity patterns coding side in that training bin as measured by the decoding analysis. (Although single neurons tended to exhibit a strong signal coding side in the training bin, consistent with the tendency for the peak in population decoding accuracy to occur at that time, they typically exhibited differences in firing rate as a function of side at other time points in the trial as well.) (Fig. 6).

After assigning neurons to time bins, we could graphically illustrate the changing pattern of activity parietal cortex generated to represent a fixed value of object-centered side over time (Fig. 7). Filled circles within each rectangular grid represent the subset of a population of 196 parietal neurons that was significantly engaged to code object-centered side in a given training bin (circles are arranged left to right in the order of the time of their engagement in the representation). Parietal cortex generated a succession of activity patterns to carry the value of side over time during a period in which the value of side remained either left or right throughout (Fig. 7). New groups of neurons were recruited into the dynamic representation of side just before the presentation of the copy object on both shift-copy (Fig. 7, filled arrows) and shift-model trials (supplemental Fig. 4, available at www.jneurosci.org as supplemental material, filled arrows), reflective of an internally initiated process. Animations of the results of this analysis over successive time steps reveal the rapid succession of activity patterns generated in parietal cortex to represent the same item of spatial information over time (supplemental movies 3 and 4 illustrate population dynamics on shift-copy and shift-model trials, respectively). Interestingly, the first time of engagement of neurons was uncorrelated with their visual response latency (supplemental Fig. 5, available at www.jneurosci.org as supplemental material). Neurons were typically activated shortly after the presentation of the model object, but their activity often did not differentiate between left and right trials until later in the trial.

Neurons exhibit correlated pretrial activity that predicts the relative timing of their sequential activation

To further elucidate aspects of the synaptic organization of the network generating dynamic representations in parietal cortex, we analyzed correlations in pretrial activity between pairs of neurons as a function of their spatial preference for side and time of engagement. Pairs of neurons that were recorded simultaneously and assigned to the same time bin more frequently exhibited significantly positively correlated pretrial activity ($p < 0.05$) if they shared the same spatial preference for side (Fig. 8A, blue bar) than if they had opposite spatial preferences (Fig. 8A, red). To confirm that correlated pretrial activity was only detected in pairs of simultaneously recorded neurons, we repeated the correlation analysis using pairs of neurons recorded at different times and detected significant correlations in pretrial activity at the expected 5% false positive rate given the 0.05 significance level of the test (Fig. 8A, Non-simultaneous).

We next evaluated pretrial activity correlation in pairs of neurons that had been assigned to different time bins, defining the bin lag of each neuron pair as the number of steps between the time bins containing the two neurons (Fig. 8B). We found that if the pair of neurons shared the same spatial preference for side, then their pretrial activities were correlated to a degree that scaled as a positive linear function of how close together in time the two neurons would be activated later in the trial during the dynamic

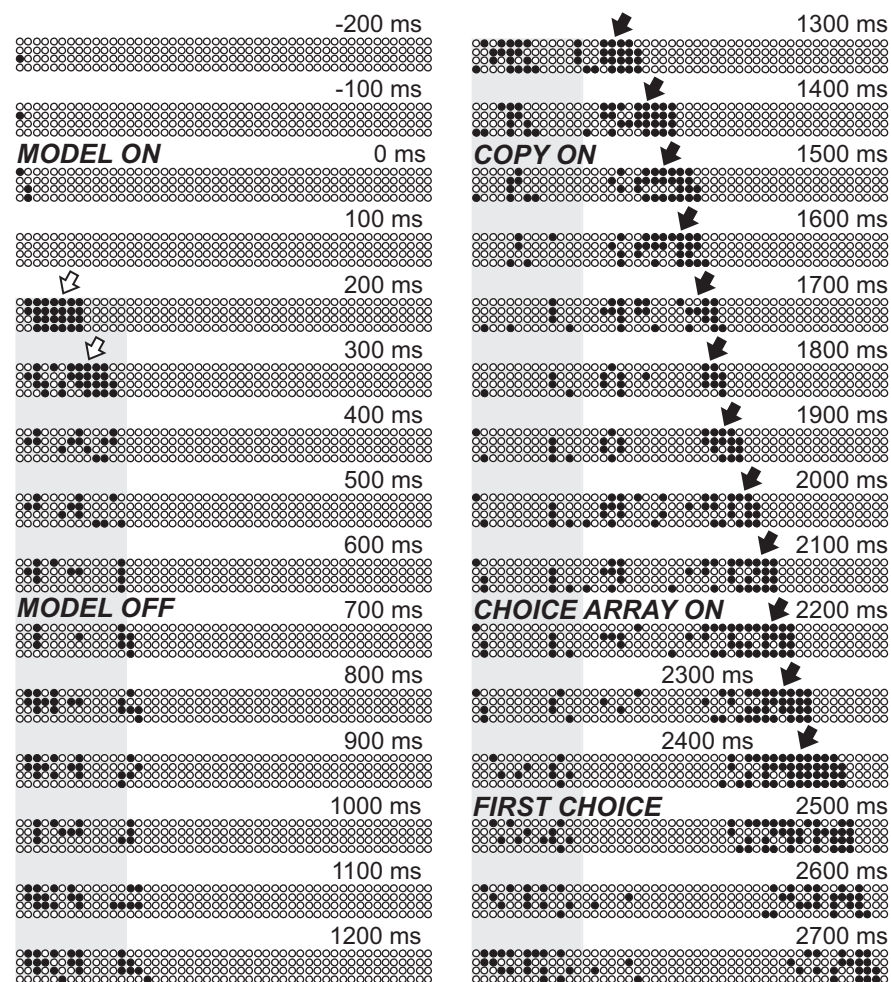


Figure 7. Parietal cortex represented each value of side (left or right) by activating a sequence of neuronal groups with the same spatial preference in rapid succession (data from shift-copy trials). Filled circles within each grid illustrate the subset of a population of 196 neurons that were significantly engaged to represent object-centered side in a given time bin (time is indicated relative to model onset). Circles are arranged left to right within each grid in order of the time of their first engagement in the trial (the position of neurons within the grid does not reflect their anatomical location). Open arrows indicate neuronal recruitment following presentation of the model object. Filled arrows indicate neuronal recruitment beginning before presentation of the copy object. The vertical gray shading indicates the set of neurons activated during the model period.

representation of side (or in other words, the strength of correlation was a decreasing linear function of bin lag) (Fig. 8C, blue line).

This was particularly interesting because the degree of correlation in spontaneous activity was observed during a period in the trial before neural activity was modulated by the objects shown and before side was defined, so that the correlation structure among neurons was predictive of the relative timing of neuronal engagement during the dynamic representation before the fact—before the dynamic representation unfolded and neurons were sequentially engaged to represent side later in the trial. Furthermore, we did not find a linear relation between activity correlation and bin lag in the case that two neurons had opposite spatial preferences for side (Fig. 8C, red line). This is evidence that the pattern of correlation in pretrial activity was functionally related to the spatial information processing required by the task. The slope of the regression line fit to the correlation-lag data was significantly different from zero for neuron pairs with the same ($p < 0.0005$) but not opposite ($p = 0.362$) spatial preferences (Fig. 8C, asterisks indicate a significant difference in pretrial ac-

tivity correlation as a function of spatial preference by t test at $p < 0.05$).

Parietal cortex selectively employs a dynamic code to represent task-critical but not task-irrelevant spatial information

We contrasted population dynamics associated with the representation of two different kinds of spatial information on shift-model trials derived from the same visual stimulus (the model object): the object-centered side of the critical square (left or right relative to the object midline) and the viewer-centered side of the critical square (left or right relative to the gaze fixation target). We found that rapid population dynamics in the model period were much more prominent in the neural population representing the object-centered side (Fig. 9A) than the viewer-centered side (Fig. 9B) of the critical square. Specifically, whereas decoding time courses for object-centered side exhibited a clear peak aligned to the training bin (Fig. 9A, arrows), decoding time courses for viewer-centered side did not—the time courses were broader (Fig. 9C, width measured at 60% of the peak), spanned the model period, and essentially overlapped regardless of the training bin (Fig. 9B, arrows).

We evaluated the tendency of the time of the peak in decoding accuracy to follow the training bin by plotting one time against the other (Fig. 9D). The slope of the resulting linear fit was significantly nonzero when decoding the object-centered side (Fig. 9D, solid lines; $p < 0.005$), but not the viewer-centered side (Fig. 9D, dashed line; $p = 0.17$) of the critical square. These observations suggest more stationary patterns of population

activity coded viewer-centered side during the model period. The overall accuracy of decoding (although not necessarily the temporal dynamics shown here) can depend on the number of neurons in the population and the strength of the signal they carry relating to the decoded variable. To demonstrate that differences in neuron number or signal strength did not account for differences in the population dynamics observed, we repeated the analysis after controlling for these variables. Significant differences in the temporal dynamics of population activity coding object-centered and viewer-centered side during the model period were still evident (supplemental Fig. 6, available at www.jneurosci.org as supplemental material).

We found, interestingly, that on shift-copy trials parietal neurons represented viewer-centered side dynamically during the copy period (supplemental Fig. 7F, available at www.jneurosci.org as supplemental material). Because the choice array moved with the copy object on shift-copy trials, the viewer-centered position of the critical square during the copy period was correlated with the position of the correct choice the monkey had to select (and therefore viewer-centered position was task-relevant

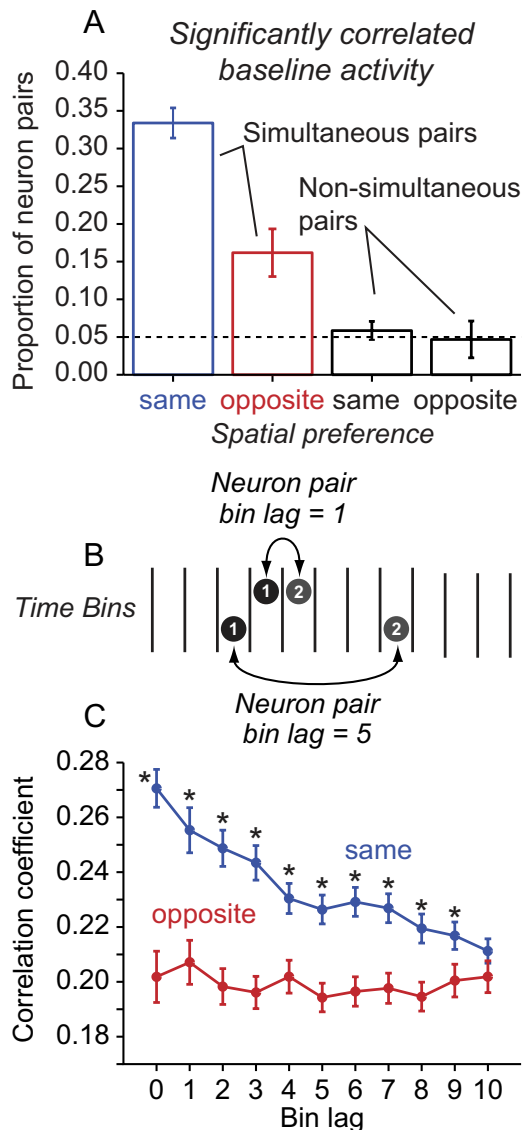


Figure 8. Correlation of pretrial activity in pairs of neurons as a function of their spatial preference for side and the relative time of their sequential engagement. The correlation analysis was restricted to pairs of neurons with positively correlated pretrial activity ($p < 0.05$) that were assigned to a time bin in the neural representation of object-centered side by the procedure illustrated in Fig. 6. **A**, Pretrial activity correlation in pairs of neurons assigned to the same time bin. A larger proportion of neuron pairs exhibited correlated pretrial activity if they shared the same spatial preference for object-centered side (blue bar) than if they had opposite spatial preferences (red bar), but only if the two neurons were recorded simultaneously. **B**, Definition of bin lag. We defined the bin lag of a pair of neurons as the number of steps between the time bins they were assigned to (neuron pairs of bin lag 1 and 5 are shown). **C**, Mean correlation coefficient for pretrial activity in pairs of neurons as a function of bin lag. For neurons sharing the same spatial preference for side (blue line), the mean pretrial activity correlation coefficient was a linear function of bin lag (significantly nonzero slope; $p < 0.0005$). For neurons with opposite spatial preferences (red line), pretrial activity correlation was not related to bin lag (slope not significantly different from zero; $p = 0.362$). Asterisks indicate bin lags where the mean pretrial activity correlation coefficients for pairs with same and opposite spatial preferences were significantly different (t test, $p < 0.05$).

under these conditions). Population dynamics did not appear to be a function of the coordinate frame used to define the coded spatial information, but instead reflected whether the spatial information coded by activity was task relevant in the sense of being correlated or uncorrelated with the spatial choice the monkey had to make.

Population activity patterns integrate spatial and temporal information

We employed pattern classification to decode the number of the current time bin (bins are numbered 1–40 from the start of the gaze fixation period) from the pattern of population activity in that bin. We included neurons in the population if they were assigned to a time bin in the dynamic neural representation of object-centered side (Fig. 6) and if their activity did not significantly reflect the retinocentric side of the critical square during either model or copy periods ($p > 0.1$ for viewer-centered side in an ANCOVA). Decoding accuracy for bin number peaked in the 45–50% correct range during the model, copy, and choice periods (Fig. 10) (chance classification is 1/40 or 2.5% correct in this analysis). Decoding accuracy for bin number started to climb just before the onset of model and copy objects (Fig. 10), suggesting that the representation of temporal information was anticipatory in nature to some degree and was not driven solely by the presentation of visual stimuli.

Discussion

We characterized the temporal dynamics of population activity in parietal area 7a of monkeys performing an object construction task. We sought to determine whether the neural representation of a spatial cognitive variable (object-centered side) was mediated by a stationary pattern of activity or by a sequence of activity patterns during a period of time in which the value of the coded variable did not change. If the representation was stationary, we would predict that the distribution of firing rates over neurons would not change as long as the information coded by that activity (side) did not change either. If the representation was dynamic, we would predict instead that side would be represented by an orderly sequence of activity states, even in the case that the value of side remained constant. We report here that area 7a neurons represented a single, fixed value of side by generating an orderly sequence of population activity patterns over time (Figs. 5, 7, supplemental Figs. 2, 4, and supplemental movies 1–4, available at www.jneurosci.org as supplemental material). Thus, side was represented dynamically, and one neural representation corresponded to a sequence of activity states in area 7a. For example, on trials when side was right, parietal cortex represented “right” by briefly activating distinct groups of neurons that shared the same spatial preference for right in a reliable order, each for a brief interval of time.

Dynamic representations reflect task relevance of coded spatial information

We found that population dynamics in parietal cortex varied as a function of the relevance of coded spatial information to the current behavioral objective. Across trial periods and task conditions we found that when side was relevant to the task, it was represented dynamically (Fig. 9A, supplemental Figs. 6C, 7F), defining relevance as being correlated with the side of the correct choice in a given spatial framework or during a particular task period. In contrast, when side was irrelevant to the task, the population representation of side was stationary (Fig. 9B, supplemental Figs. 6D, 7A). In this regard, it is interesting that in a match-to-category task (Meyers et al., 2008), the population representation of object category (task-relevant information) was more dynamic than the representation of object identity (task-irrelevant information) in inferotemporal cortex. The agreement between that study and the present one could point to a common principle of cortical information processing by which

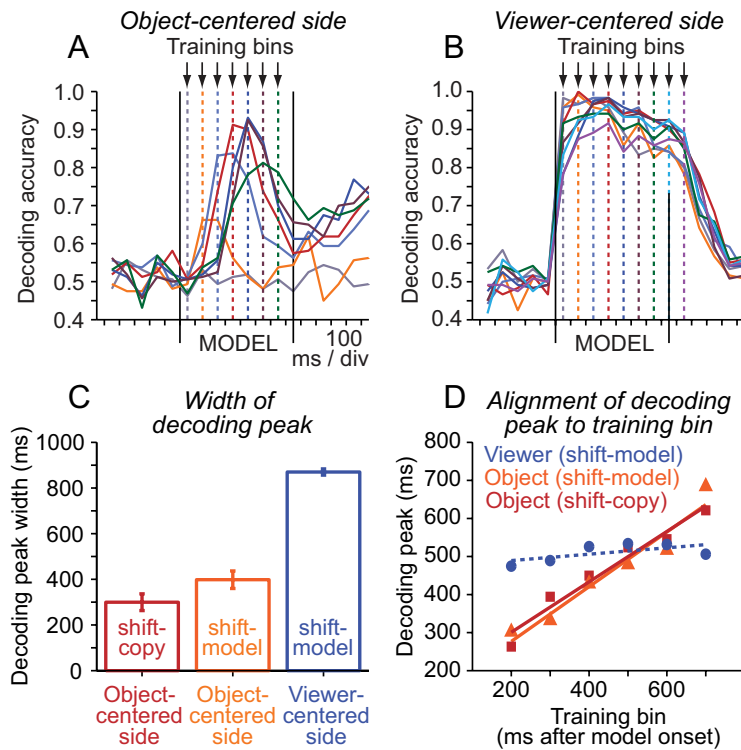


Figure 9. Neural populations coding different types of spatial information exhibit distinct temporal dynamics. **A**, Dynamic representation of the object-centered side of the critical square during the model period on shift-copy trials. Time courses of decoding accuracy exhibit a clear peak that is aligned with each training bin (indicated by the vertical dashed line of same color). **B**, Stationary representation of the viewer-centered side of the critical square during the model period on shift-model trials (left or right relative to the gaze fixation target). Time courses of decoding accuracy generated using different training bins in the model period essentially overlap. **C**, Average width of peaks in decoding accuracy (measured at 60% of peak amplitude). Error bars indicate \pm SEM. Peaks were significantly broader (t test; $p < 0.05$) when decoding viewer-centered side than object-centered side on both shift-copy and shift-model trials. **D**, Quantification of the tendency of the peak in decoding accuracy to track the training bin. The relationship between the time of the peak in decoding accuracy and the training bin is significantly linear when decoding the object-centered side of the critical square during the model period in both shift-model (orange line) and shift-copy (red line) trials (slopes of regression lines significantly different from zero; $p < 0.005$), but not when decoding the viewer-centered side of the critical square during the model period (blue line; slope of regression line not significantly different from zero; $p = 0.17$).

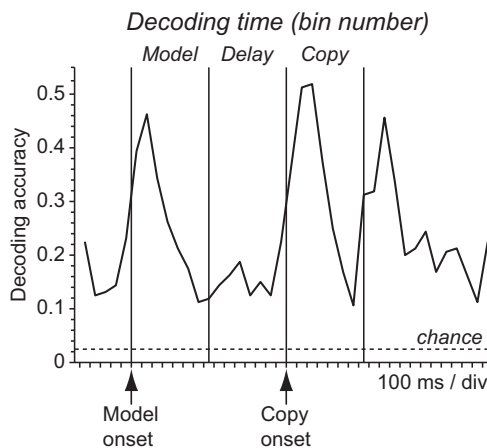


Figure 10. Accuracy of decoding the current time bin number (bins are numbered 1–40 starting at the beginning of the gaze fixation period) based on the pattern of population activity observed within that bin (chance decoding is 1/40 or 2.5% correct). Neurons were included in the population if they were assigned to a time bin in the representation of object-centered side by the procedure shown in Fig. 6 and if their activity was not significantly related to viewer-centered side during either the model or copy periods ($p > 0.1$).

rapid population dynamics mediate task-critical cognitive processing.

Relation between correlated activity and the timing of neuronal activation

We were able to identify the subgroups of neurons that were engaged at each time point during the dynamic representation (Figs. 6, 7), allowing us to discover a systematic relation between the correlated spontaneous activity of neurons and the relative timing of their activation in the dynamic representation. In pairs of neurons with the same spatial preference for side, the strength of correlation in pretrial activity predicted how close in time the two neurons would be sequentially engaged to represent side later in the trial (Fig. 8C). Neurons tend to exhibit correlated spontaneous activity to the degree that they exhibit similar tuning to features such as movement direction in primary motor cortex (Lee et al., 1998), line orientation in primary visual cortex (Kenet et al., 2003), or place in hippocampus (Wilson and McNaughton, 1994). By demonstrating a relation to the timing of neuronal activation, the present data provide new insight into the functional organization of the network generating dynamic representations. Neurons that will be activated closer together in the sequence of recruitment share a stronger synaptic drive. This relationship could reflect a number of different patterns of connectivity between neurons. Perhaps the simplest is a network in which each group of neurons participating in the dynamic representation provides the strongest synaptic drive to the next group in the

sequence, driving the network along a trajectory through the state space.

Relation of population dynamics to receptive or movement field properties

The decoding analysis identified a succession of population activity patterns that coded the value of object-centered side (left or right). The information carried by the activity of these neurons did not reflect sensory or motor parameters of the task, because object-centered side was uncorrelated with both the form and the retinocentric position of reference objects as well as the direction or timing of the required motor response (Chafee et al., 2005, 2007). In a prior study (Chafee et al., 2005), we provided evidence that the spatial selectivity that neurons demonstrated during object construction was not congruent with the location or shape of their retinocentric visual receptive or movement fields as tested in sensorimotor control tasks. The activity of >80% of parietal neurons that varied significantly as a function of the position of the missing critical square in the copy object did not vary significantly as a function of the position of a spot visual stimulus presented at the same retinocentric locations in control tasks, even though monkeys were required to either direct covert attention or plan a delayed saccade toward those visual stimuli (Chafee

et al., 2005). Of the remaining minority of neurons, many exhibited incongruent spatial preferences as a function of task context, suggesting that spatial selectivity during object construction was independent of the locations of retinocentric receptive fields tested in the control tasks. Consistent with this, we found that when we shifted the retinocentric position of reference objects, many parietal neurons coded the object-centered position of missing squares in a manner that could be statistically dissociated from retinocentric position (Fig. 4) (Chafee et al., 2007). Finally, the time of neuronal recruitment in the representation of object-centered side was temporally independent from receptive field properties in the sense of being uncorrelated with the visual response latency of neurons in the task (supplemental Fig. 5, available at www.jneurosci.org as supplemental material).

Dynamic representations in other neural systems

In the olfactory system, presentation of an odor briefly activates neurons in a temporally extended sequence that codes the identity of the stimulus (Laurent et al., 1996; Wehr and Laurent, 1996; Laurent et al., 2001, 2002). In the isolated nervous system of the leech, stochastic population dynamics predict which of two alternative motor patterns the network will generate in response to a single stimulus (Briggman et al., 2005, 2006). In monkeys, population activity in prefrontal cortex is dynamic during the delay period of working memory tasks, switching spontaneously between different activity states (Seidemann et al., 1996) or exhibiting temporal trends (Chafee and Goldman-Rakic, 1998; Romo et al., 1999) in which time and stored spatial information exert separable influences on firing rate (Machens et al., 2010)—a dynamic successfully modeled by networks in which units are simultaneously “tuned” along both feature and time dimensions (Singh and Eliasmith, 2006). Similarly, in the lateral intraparietal area, neurons carry spatial and temporal information, coding either elapsed time (Leon and Shadlen, 2003) or the time of anticipated events (Janssen and Shadlen, 2005), in addition to the direction of a forthcoming saccade. In the present study, rather than exhibiting relatively gradual changes in firing rate and coded information, parietal cortex generated a rapid succession of activity patterns to code side around the time that task-critical stimuli were presented and dynamics slowed during the delay period. Furthermore, the pace of population dynamics varied as a function of the behavioral relevance of the coded information, and the order of neuronal activation was reflected by the strength of correlation in spontaneous activity. Finally, although our data indicate that different groups of neurons are recruited at generally the same time within each trial, this does not require that the action potentials of the neurons be precisely timed or synchronous (Abeles et al., 1993; Aertsen et al., 1996).

Computational utility of dynamic representations

Theoretical studies have shown that artificial networks can code information using sequences of activity states rather than the steady-state patterns of activity associated with attractor dynamics (Maass et al., 2002; Buonomano and Maass, 2009). Transient inputs perturb such systems in an informative way and produce temporally extended responses. By incorporating time into the neuronal representation and allowing for sequences of states to become informative, the brain may increase the state space it has available to separate neural representations (Laurent, 2002; Buonomano and Maass, 2009). Another computational advantage of dynamic coding could derive from a property evident in our data, namely that the active neural group provided not only spatial but also temporal information (Fig. 10). Further-

more, dynamics were not uniform over the trial, being faster during the model and copy periods (Fig. 5*A,C*) and slower during the intervening delay period (Fig. 5*B*), so that the amount of temporal information carried by population activity was focused around the time of task-critical events (Fig. 10). Therefore dynamic coding could serve to synchronize cognitive information processing operations to a stereotypic sequence of sensorimotor events.

Our data provide evidence that temporally dynamic neural representations of space: (1) exist in parietal area 7a; (2) preferentially represent task-relevant spatial information; (3) are mediated by the rapid recruitment of distinct neuronal groups; (4) are predicted by patterns of correlation in spontaneous activity before they occur; and (5) incorporate temporal along with spatial information.

References

- Abeles M, Bergman H, Margalit E, Vaadia E (1993) Spatiotemporal firing patterns in the frontal cortex of behaving monkeys. *J Neurophysiol* 70:1629–1638.
- Aertsen A, Diesmann M, Gewaltig MO (1996) Propagation of synchronous spiking activity in feedforward neural networks. *J Physiol Paris* 90:243–247.
- Averbeck BB, Lee D (2006) Effects of noise correlations on information encoding and decoding. *J Neurophysiol* 95:3633–3644.
- 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.
- Briggman KL, Abarbanel HD, Kristan WB Jr (2005) Optical imaging of neuronal populations during decision-making. *Science* 307:896–901.
- Briggman KL, Abarbanel HD, Kristan WB Jr (2006) From crawling to cognition: analyzing the dynamical interactions among populations of neurons. *Curr Opin Neurobiol* 16:135–144.
- Buonomano DV, Maass W (2009) State-dependent computations: spatio-temporal processing in cortical networks. *Nat Rev Neurosci* 10:113–125.
- Chafee MV, Goldman-Rakic PS (1998) Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J Neurophysiol* 79:2919–2940.
- 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, Averbeck BB, Chafee MV (2008) Neural ensemble decoding reveals a correlate of viewer- to object-centered spatial transformation in monkey parietal cortex. *J Neurosci* 28:5218–5228.
- Driver J, Baylis GC, Goodrich SJ, Rafal RD (1994) Axis-based neglect of visual shapes. *Neuropsychologia* 32:1353–1365.
- Janssen P, Shadlen MN (2005) A representation of the hazard rate of elapsed time in macaque area LIP. *Nat Neurosci* 8:234–241.
- Johnson RA, Wichern DW (1998) Applied multivariate statistical analysis. Saddle River, NJ: Prentice Hall.
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A (2003) Spontaneously emerging cortical representations of visual attributes. *Nature* 425:954–956.
- Klecka WR (1980) Discriminant analysis. Newbury Park, CA: Sage Publications.
- Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci* 3:884–895.
- Laurent G, Wehr M, Davidowitz H (1996) Temporal representations of odors in an olfactory network. *J Neurosci* 16:3837–3847.
- Laurent G, Stopfer M, Friedrich RW, Rabinovich MI, Volkovskii A, Abarbanel HD (2001) Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu Rev Neurosci* 24:263–297.
- Lee D, Port NL, Kruse W, Georgopoulos AP (1998) Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J Neurosci* 18:1161–1170.
- Leon MI, Shadlen MN (2003) Representation of time by neurons in the posterior parietal cortex of the macaque. *Neuron* 38:317–327.

- Maass W, Natschläger T, Markram H (2002) Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput* 14:2531–2560.
- Machens CK, Romo R, Brody CD (2010) Functional, but not anatomical, separation of “what” and “when” in prefrontal cortex. *J Neurosci* 30:350–360.
- Mazor O, Laurent G (2005) Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* 48:661–673.
- Meyers EM, Freedman DJ, Kreiman G, Miller EK, Poggio T (2008) Dynamic population coding of category information in inferior temporal and prefrontal cortex. *J Neurophysiol* 100:1407–1419.
- Romo R, Brody CD, Hernández A, Lemus L (1999) Neuronal correlates of parametric working memory in the prefrontal cortex. *Nature* 399:470–473.
- Seidemann E, Meilijson I, Abeles M, Bergman H, Vaadia E (1996) Simultaneously recorded single units in the frontal cortex go through sequences of discrete and stable states in monkeys performing a delayed localization task. *J Neurosci* 16:752–768.
- Singh R, Eliasmith C (2006) Higher-dimensional neurons explain the tuning and dynamics of working memory cells. *J Neurosci* 26:3667–3678.
- Wehr M, Laurent G (1996) Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature* 384:162–166.
- Wilson MA, McNaughton BL (1994) Reactivation of hippocampal ensemble memories during sleep. *Science* 265:676–679.