Neuron, Volume *68*

Supplemental Information

Cortical Preparatory Activity: Representation of Movement or First Cog in a Dynamical Machine?

M M. Churchland, J P. Cunningham, M T. Kaufman, S I. Ryu, and K V. Shenoy

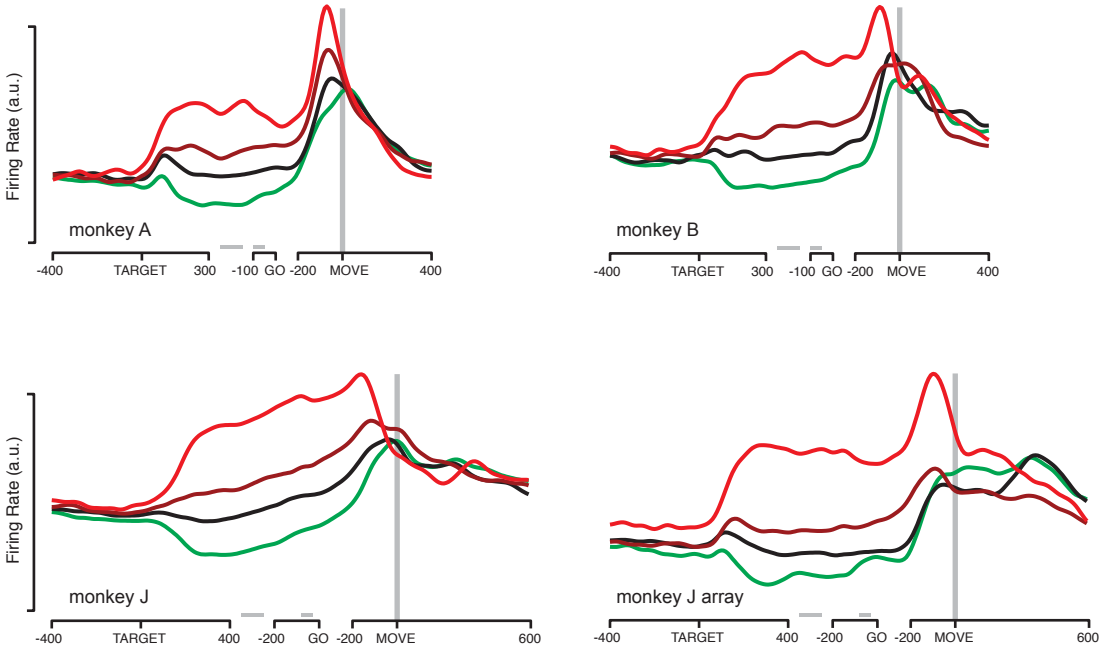

Supplemental figure 1

Supplemental Figure 1. Population PSTHs. Each panel plots the average firing rate (across all neurons) for the condition with the strongest preparatory response (*red*), the condition with the weakest preparatory response (*green*) and two intermediate conditions (*intermediate shades*). Such conditions were found separately for each neuron. They were then averaged after normalizing each neuron's activity by its maximum rate (results were similar without normalization). Analysis was restricted to include only neurons with robust preparatory tuning *and* robust peri-movement tuning (see *Experimental Procedures*, same restriction criteria as for Figure 4). Despite this restriction, peri-movement tuning appears quite weak in the population averages, especially by the time of movement onset. This is because averaging was segregated by each neuron's preparatory preference, which on average bore little resemblance to its peri-movement preference. What would otherwise be strong peri-movement tuning is thus averaged out.

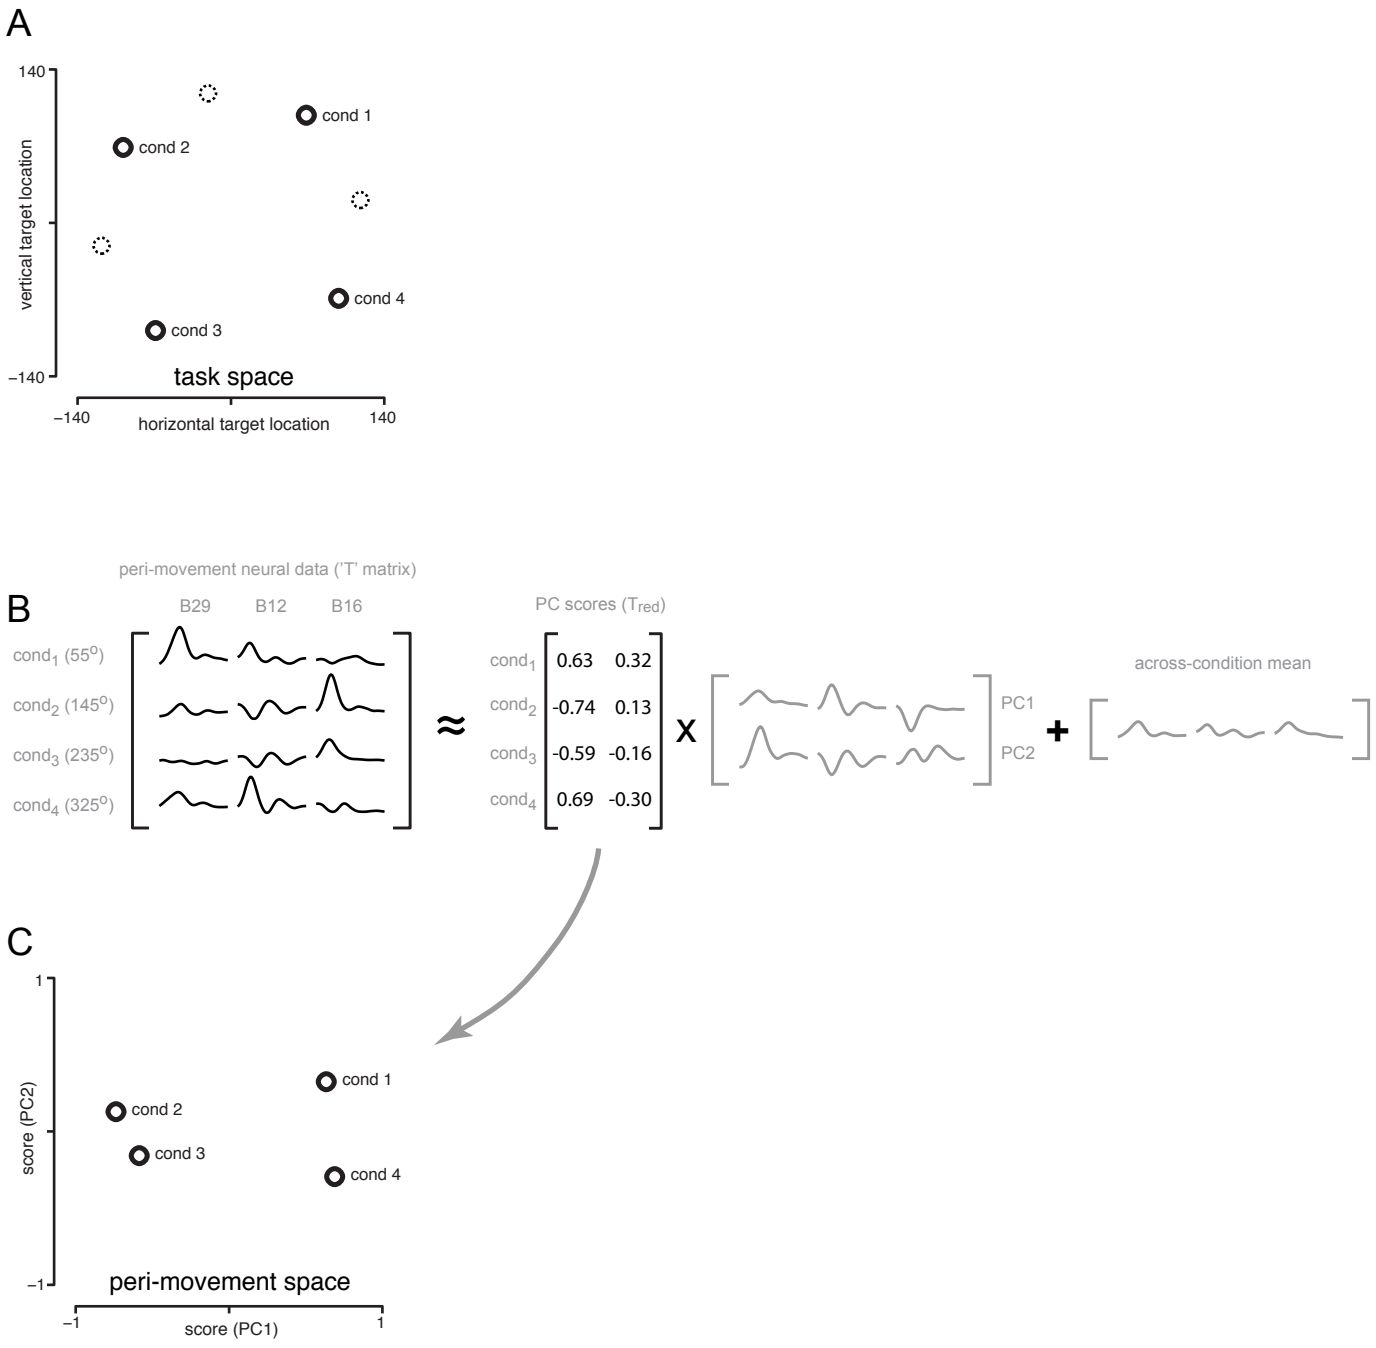

Supplemental Figure 2

Supplemental Figure 2. Illustration of how PCA is used to construct the peri-movement space. For illustration, analysis is restricted to data from four conditions (4 reach directions at the greater distance and instructed speed) and from three neurons (B29, B12 and B16). ***A.*** The four conditions are plotted in a traditional task space based on target location. For reference, the other three conditions at that distance are shown with dotted symbols. For this illustration these are not considered further. ***B.*** Illustration of the matrix 'T' containing peri-movement responses, and of how PCA is applied. The matrix T (*left*) is of size 4×3*t*, where *t* is the number of measured time-points. PCA approximates this matrix as the product of two matrices: a 4×2 matrix of scores (T_{red}) and a 2×3*t* matrix of PCs (*gray*). This product, added to the across-condition mean (a 1×3*t* vector, *gray*) provides an excellent approximation (not shown) to T. Critically, the PC scores now provide the location of each condition in a 2-D space. Each location indicates how the two PCs need to be weighted so that their sum approximates the relevant row of T. ***C.*** Locations of the four conditions in the peri-movement space, as determined by the PC scores. The peri-movement space is certainly related to the task space in *A*; the PC scores will correlate with target location. Yet the condition locations in *C* are a non-linear transformation of those in *A*. It is thus possible for the PD to perform better in one space than in the other.

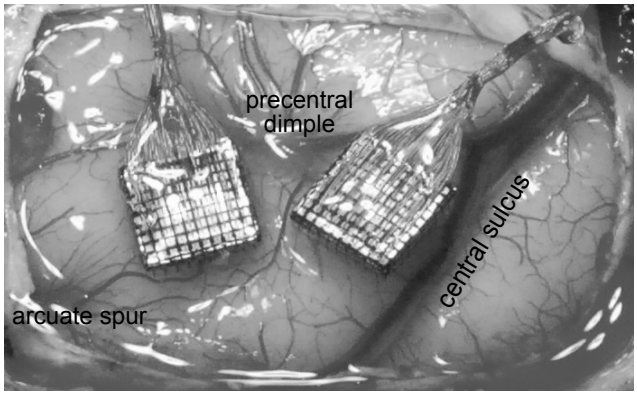

Supplemental Figure 3

Supplemental Figure 3. Array locations in monkey J: the first in caudal PMd (*left*) the second in surface M1 and perhaps extending into caudal PMd (*right*). The photograph was taken immediately prior to array implantation. Previous single-electrode penetrations were largely within the area spanned by the two arrays, but also progressed deeper in the central sulcus and the lateral bank of the precentral dimple. Microstimulation of this region produced movements primarily of the upper arm and shoulder, occasionally of the wrist and arm, and rarely of the leg and arm.

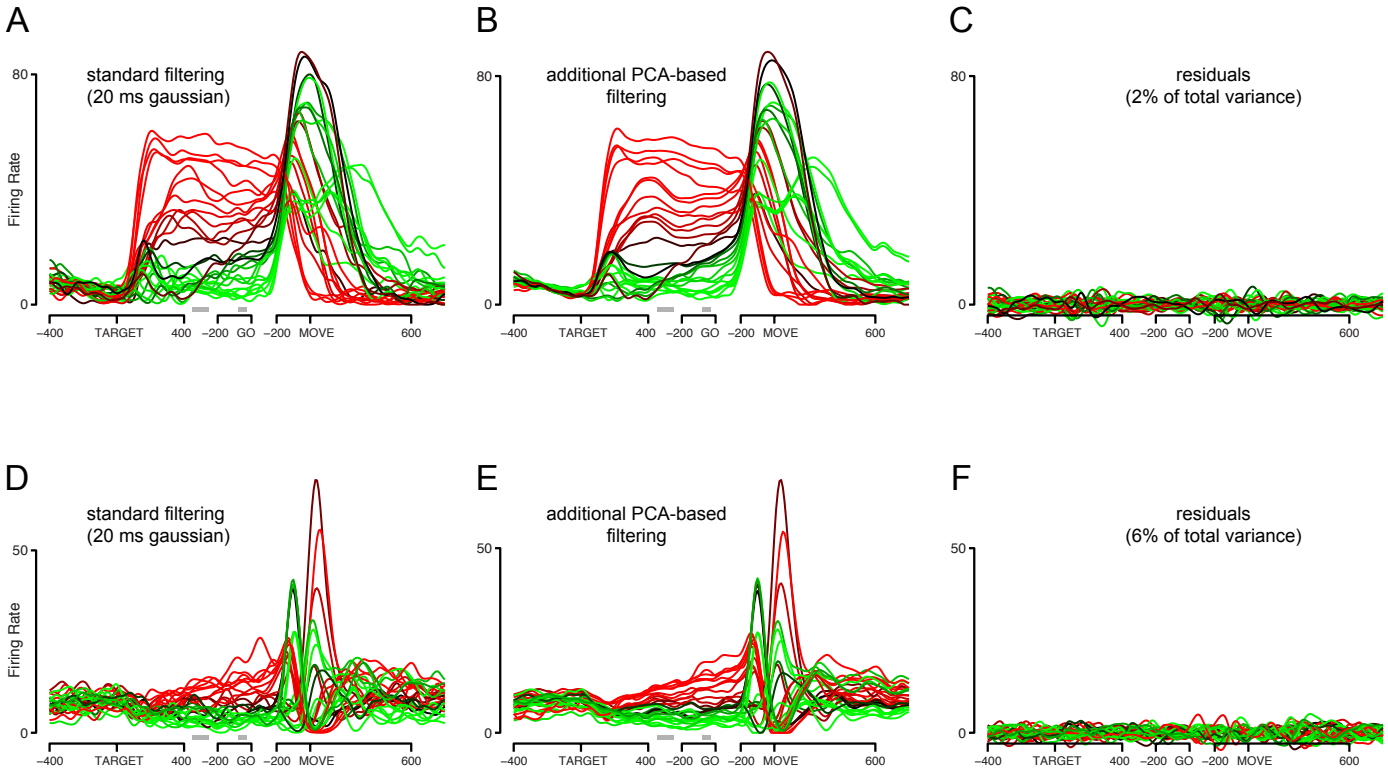

Supplemental Figure 4

Supplemental Figure 4. Illustration of how PSTHs were smoothed for improved visualization. The top and bottom rows show data for neurons J36 and J142. From left to right, panels show PSTHs computed with standard filtering, PSTHs computed using additional filtering (see below), and the residuals (the difference between the standard-filtered version and the version with additional filtering).

For the PSTHs in panels *A&D*, we computed the mean firing rate as a function of time in the usual way. We convolved each spike with a Gaussian kernel of width 20 ms, and averaged across all trials recorded for a given condition. The mean firing rate was computed three times: time-locked to target onset, the go cue, and movement onset. To allow the eye to visually follow an individual condition's response throughout the trial, the firing rates for these three epochs were spliced together. Gray bars at the bottom of each panel show the region over which we interpolated/smoothed the firing rate between two epochs. Note that the segment that is time-locked to the go cue actually occurs earlier than the go cue itself. This was done because firing rates typically changed rapidly between the go cue and movement onset, and we did not want the same change in firing rates to be represented twice.

Panels *B&E* were computed using an additional smoothing step. Instead of using a temporal filter (which exploits the fact that different times cannot be arbitrarily different from one another) we used principal component analysis (PCA) to exploit the fact that different conditions cannot be arbitrarily different from one another (note that this is a different use of PCA from that employed elsewhere in this study, and has an unrelated goal). To apply smoothing, we compiled a *cxt* data matrix, where each row contained the response of that neuron for one condition across all times. We then decomposed this matrix into its principal components (PCs), and then reconstructed the data using the first six principal components. This procedure preferentially discards high-frequency events that are unique to one of the 27 conditions (and thus more likely to result from sampling error). Something similar could of course be obtained by using a broader temporal filter, but that could come at the cost of losing real high-frequency aspects of the response. An advantage of our method is that it does not remove high-frequency aspects of the response if they are shared among a number of conditions (see especially the bottom row).

Panels *C&F* plot residuals. For example, each trace in *C* is the difference between the corresponding traces in *A* and *B*. The fact that the residuals are small, and largely uncorrelated either with one another or with themselves over time, indicates that the smoothing method is indeed preferentially discarding sampling noise. For comparison, across all neurons for monkey N, the average residual was 7% (and was typically lower for recordings with higher signal-to-noise). Thus, the smoothing method impacts the data minimally, and largely discards exactly those features (those that are uncorrelated across time or conditions) most likely to be noise.

A known drawback of temporal smoothing is that one can no longer interpret the similarity of firing rates at nearby times as an indication that a given feature is real; smoothness has been imposed, and can no longer be used to judge signal-to-noise. By analogy, a straightforward drawback of our method is that response similarity across conditions can no longer be used to judge whether a feature is real. Figure 3*F* of the main text provides a good example: the small undulations of firing rate with time during the delay period are shared across conditions. Although they may indeed be real, they are just

as likely to result from sampling error (especially because firing rate is low at that time). Their consistency across conditions is not necessarily an indication of their reliability.

To simplify interpretation, in this study we used the additional smoothing method only for visualization purposes (i.e., for plotting PSTHs). This aids visualization both directly, and by minimizing the confusion that results when many lines cross repeatedly due to small amounts of noise. All other analyses were based on conventionally-smoothed data (as in *A,D*) without the additional smoothing step. That said, the results of all analyses were virtually identical regardless of whether additional smoothing was first applied.

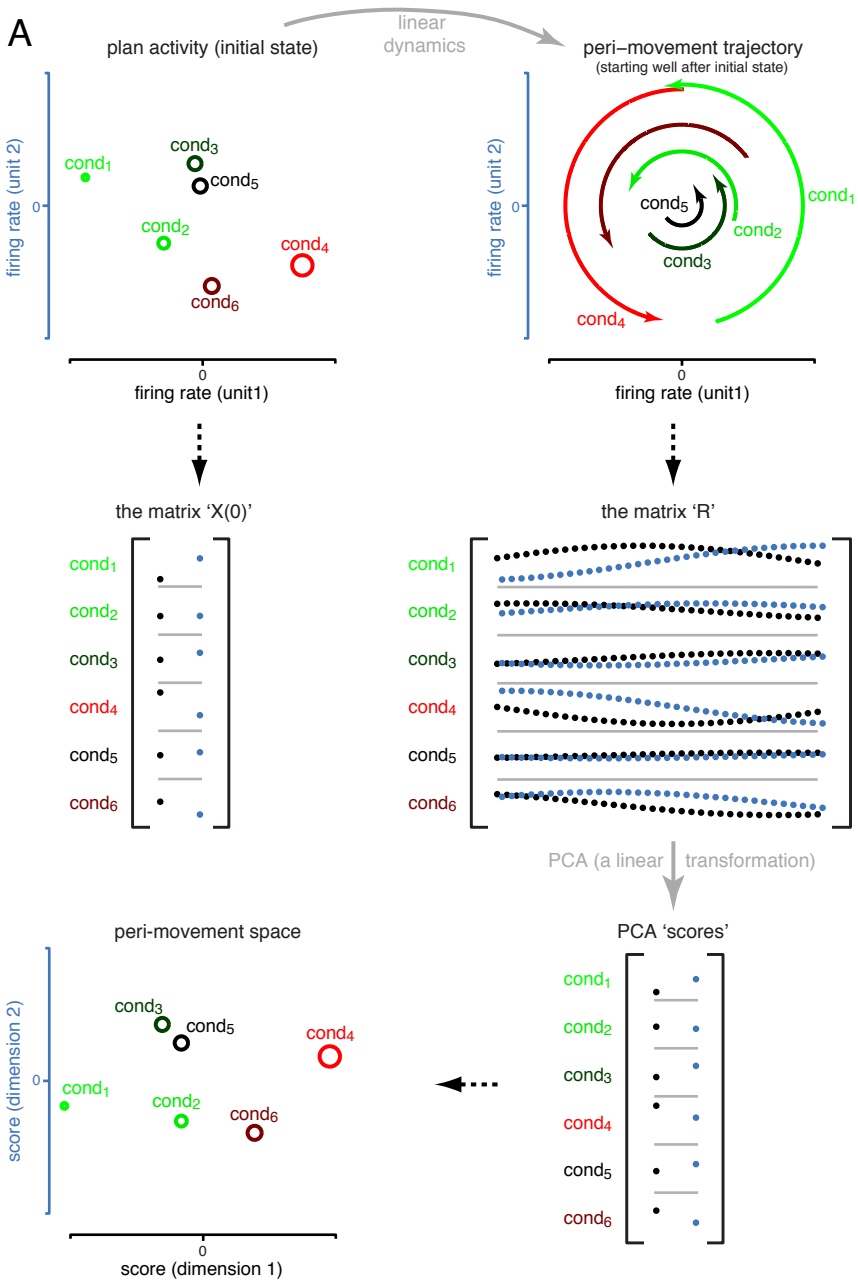

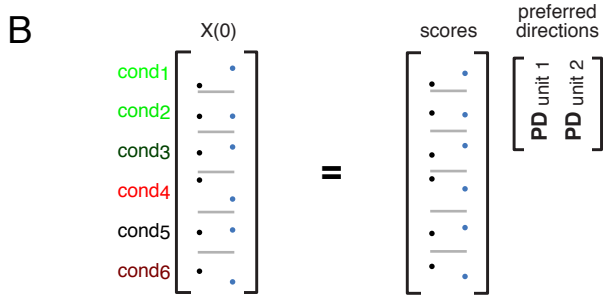

Supplemental Figure 5

Supplemental Figure 5. Illustration of why, for a linear dynamical system, preparatory tuning ought to be captured by a PD in peri-movement space. This illustration is based on the simulated data from Figure 8 of the main text. For visualization, only the first 6 conditions are plotted. A more formal explanation can be found in the supplemental text. ***A.*** Illustration of how we compiled matrices. The height of each dot indicates the value of that element. The matrix X(0) captures the initial state for all conditions and both units, and is of size $c \times n$ where c is the number of conditions and n is the number of neurons (6×2 in the present case). The matrix R captures the evolution of peri-movement activity for both units and is $c \times nt$, ($6 \times 2t$ in the present case) where t is the number of time-points. Applying PCA to R yields the scores, a $c \times k$ matrix where k is the reduced dimensionality (2 in the present case). Those scores give the location of each condition in the k-dimensional peri-movement space. Because the scores are a linear transformation of R, which is a linear transformation of X(0), the scores will be a linear transformation of X(0). That transformation will be invertible so long as information has not been lost (so long as the dynamics don't involve a rapid decay to zero) and so long as k is large enough (k=2 is sufficient in the present case but k=1 would not be). ***B.*** Illustration of why the above facts imply that preparatory tuning can be captured by a PD in the peri-movement space. As discussed above, X(0) and the principal component scores are related by a linear transformation. Each column of the $k \times k$ matrix that captures that transformation defines a preferred direction in the peri-movement space (in the present case there are two 2-dimensional PDs). The initial rate of unit 1 for a given condition can then be found by multiplying the PD for unit 1 by the scores for that condition (i.e., the location of that condition in the peri-movement space). This is analogous to having a preferred direction in velocity space, except the space is defined by the principal component scores rather than directly via a behavioral measurement (see Supp. Fig. 2).

Supplemental Text: *explanation of how a PD in peri-movement space can capture preparatory activity for a time-varying linear dynamical system (please see Supp. Fig 5 for a less formal illustration of the derivation below).*

Consider a dynamical neural system, with the firing rate of each neuron represented by the state $\mathbf{x}(t)$, a $1\times n$ vector. Suppose the dynamics of this system are captured by,

$$\mathbf{x}_{cond}(t+1) = \mathbf{x}_{cond}(t)W(t)$$

where $\mathbf{x}_{cond}(t)$ is the evolution of the state for a particular condition, and $W(t)$ is an $n\times n$ time-varying matrix. If we have c conditions we can write,

$$X(t+1) = X(t)W(t)$$

where X(t) has dimensions $c\times n$. Now consider the $c\times nt$ matrix,

$$R = [X(1), X(2)\ ), X(3)\dots\ ]$$

and the $n\times nt$ matrix,

$$F = [W(0), W(0)W(1), W(0)W(1)W(2)\dots\ ].$$

We can write

$$R = X(0)F$$

We can apply singular value decomposition and reduce the dimensionality of the l.h.s to k. (If the dimensionality of the system is $> k$, the equality becomes approximate; if the dimensionality is $<= k$, then it remains exact):

$$USV^t = X(0)F$$

Where U is $c\times k$, S is $k\times k$, and V^t is $k\times nt$. The dimensionality reduction lies in the fact that $k < c$.

Provided F is full rank, we can take its pseudoinverse, $F^\dagger$, and we know that $FF^\dagger = I$ (where I is $n\times n$). Thus,

$$USV^tF^\dagger = X(0)$$

Let us set the $k\times n$ matrix C to be equal to $V^tF^\dagger$. Then,

$$(US)C = X(0).$$

Note that every column of X(0) is the preparatory response of one neuron for all conditions (i.e., that neuron's 'tuning'). That column is a linear combination (determined by the corresponding column of coefficients in C) of the columns of US. Put another way, each neuron's tuning can be captured by a PD (whose direction is given by the coefficients of C) in the space described by US. US is the reduced dimensional ($c\times k$) representation of R ($c\times nt$). That is, it is the principal component 'scores' that result when applying PCA to R. This reduced-dimensional representation of R is identical to the reduced-dimensional representation of T (T is the matrix, defined in the main text, that was used to derive the 'peri-movement space'). For both R and T, each row gives, for one condition, the activity of every neuron at all peri-movement times. The only difference between R and T is that in R, we first index through all neurons for time one, whereas in T we first index through all times for neuron one. Thus, each neuron's preparatory activity (each column of X(0)) can be accounted for by a linear combination of the columns of US (which is reduced-dimensional representation of T). As illustrated in Figure 8 and Supp. Fig. 5, this can be though of as explaining the preparatory activity of a given neuron via a PD in peri-movement space, defined by the reduced-dimensional representation of T or R.

A side point is that, in referring above to US as the reduced-dimensional representation of R (i.e., the PC scores) we are assuming that the columns of R are of zero mean (in which case US is indeed identical to the principal component scores). R will have columns of zero mean if X(0) has columns of zero mean. The implication is that, because we are using PCA (which subtracts the mean from every column) to create the reduced-dimensional space, the PD in the reduced-dimensional space can only account for preparatory activity of zero mean (across conditions). This is naturally handled by the model in eqn. 1 of the main text, where the non-zero mean firing rate of each neuron is captured by the offset b_o. The PD need only account for the difference in rate among conditions, and for not the overall average firing rate.

Supplemental Experimental Procedures

Task design and behavioral control

Our basic methods have been described previously (Churchland et al., 2006c). Briefly, three adult male monkeys (monkeys A,B and J, *Macaca mulatta*, ~10 kg) sat in a customized chair with head restraint and performed the task on a fronto-parallel screen. The hand and eye were tracked optically (accuracy of 0.35 mm and ~1°, 60 and 240 Hz). On each task trial, a central touch spot was held (for at least 400 ms) and a reach target then appeared. This target initially jittered slightly (2 mm SD) in place, during which time the monkey could prepare his reach but was not yet allowed to move his hand. Following a variable (~0-1000 ms, the exact range varied across monkeys) delay, jitter ceased and the touch spot disappeared, providing the 'go cue' for the monkey to execute his reach. Monkeys rapidly learned that the jittering target could not be struck, and the hand was typically held very steady until the go cue. Allowable reaction times were 100-500 ms (speed task) and 120-900 ms (maze task). Median reaction times were typically ~300 ms. Juice reward was delivered after the target was held for 300 ms (speed task) or 700 ms (maze task).

For the speed task, green and red targets instructed 'slow' and 'fast' reaches (Churchland et al., 2006b). We employed 7 target directions and two target distances (see Figure 1*C*). Reaches were typically quite straight, with a small amount of curvature for fast reaches. For the maze task, 1/3 of the conditions involved conventional straight reaches with no barriers and 1/3 of the conditions required the reach to curve around one or more virtual barriers. These curve-requiring maze configurations were repeated for the final 1/3 of conditions, in the presence of additional distractor targets that could not be reached (e.g., figure 1*D*). Reaches were often very similar for the single-target and target-amid-distractor versions of the same maze, but this was not always the case, and these conditions are thus analyzed separately.

For the speed task, successful reaches were highly stereotyped (Churchland et al., 2006a, fig. 1), making it sensible to compute the mean firing rate across trials. For the maze task, most successful reaches were self-similar, but outliers were occasionally present. Subsequent analysis was thus restricted to the subset of trials (typically ~90% of all trials) where the horizontal and vertical velocity profiles correlated with those of a reference 'proto-trial' with $r > 0.9$.

Neural and EMG recordings

A total of 310 single-electrode isolations (64, 74, and 172 single-unit isolations for monkeys A, B and J) were recorded from caudal PMd, surface M1 and sulcal M1. For monkeys A and B, the surface locations of the penetration entry points are shown in (Churchland et al., 2006b). For these two monkeys, initial penetrations were guided by prior MRI. For monkey J this was deemed unnecessary. In all cases the dura was at some point reflected and the location of penetrated area confirmed relative to cortical landmarks (e.g., Supplemental Figure 3).

EMG data were collected as described previously (Churchland et al., 2006c). EMG records were filtered, rectified, smoothed and averaged before further analysis. Recordings were made from six muscle groups (deltoid, biceps brachii, triceps brachii,

trapezius, latissimus dorsi, and pectoralis). The triceps were minimally active for monkeys A and J and were not recorded. Recordings were often made from multiple heads of the same muscle, yielding a total of 6 (monkey A), 11 (monkey B) and 8 (monkey J) EMG recordings.

Computing the mean firing rate versus time

Delay-period duration and reaction time were variable, making it impossible to compute a single trace of the mean firing rate versus time. We thus computing the mean firing rate separately locked to target onset, the go cue, and movement onset. For visual presentation (where one wishes to follow a trace through different epochs) we interpolated over the gaps between the three epochs (gray bars in Figures 2 and 3). The location and duration of the gaps was chosen to minimize discontinuities in firing rate. This resulting single trace of firing rate, locked to all three events, was also useful for the analysis in Figure 4E, where we needed to sweep time over a greater extent than is spanned by a single epoch. All means were computed after smoothing spike trains with a 20 ms Gaussian. For figures showing the firing rates of single neurons, further smoothing was provided using a novel method (Supplemental Figure 4), making it much easier to visually follow the firing rate for individual conditions. Additional smoothing was used for visualization purposes and not for the central analyses (although results were virtually identical if it was used).

Estimating the upper limit of PD performance

Even under the assumption that a PD (rather than a non-linear model) is appropriate, perfect performance of the PD is expected only if 1) one is using the right type of space, 2) one is using the right dimensionality, and 3) the data is noise-free. To estimate the degree to which the latter two factors degrade performance, we employed a PCA-based space derived from the *preparatory* activity of all neurons (T was $c \times n$ with each row containing the mean preparatory activity for every neuron). On the assumption that this must approximate the right space (whatever one neuron is tuned for, it ought to be shared with at least some of the other neurons) the resulting performance (gray traces in Figure 6*A-D*) provides a rough upper limit on performance. This T-matrix also allows us to estimate the dimensionality of preparatory activity: depending on the dataset, 7-10 dimensions were required to capture 90% of the data variance.

Churchland, M.M., Afshar, A., and Shenoy, K.V. (2006a). A central source of movement variability. Neuron *52*, 1085-1096.

Churchland, M.M., Santhanam, G., and Shenoy, K.V. (2006b). Preparatory activity in premotor and motor cortex reflects the speed of the upcoming reach. J Neurophysiol *96*, 3130-3146.

Churchland, M.M., Yu, B.M., Ryu, S.I., Santhanam, G., and Shenoy, K.V. (2006c). Neural variability in premotor cortex provides a signature of motor preparation. J Neurosci *26*, 3697-3712.