Population coding of conditional probability distributions in dorsal premotor cortex

Joshua I. Glaser[1*], Matthew G. Perich[2], Pavan Ramkumar[1], Lee E. Miller[1-3], Konrad P. Kording[1-4]

1. Department of Physical Medicine and Rehabilitation, Northwestern University and Shirley Ryan Ability Lab, Chicago, IL, USA
2. Department of Physiology, Northwestern University, Chicago, IL, USA
3. Department of Biomedical Engineering, Northwestern University, Chicago, IL, USA
4. Department of Applied Mathematics, Northwestern University, Chicago, IL, USA

* Contact: j-glaser@u.northwestern.edu

Abstract:

Our bodies and the environment constrain our movements. For example, when our arm is fully outstretched, we cannot extend it further. More generally, the distribution of possible movements is conditioned on the state of our bodies in the environment, which is constantly changing. However, little is known about how the brain represents such distributions, and uses them in movement planning. Here, we recorded from dorsal premotor cortex (PMd) and primary motor cortex (M1) while monkeys reached to randomly placed targets. The hand's position within the workspace created probability distributions of possible upcoming targets, which affected movement trajectories and latencies. PMd, but not M1, neurons had increased activity when the monkey's hand position made it likely the upcoming movement would be in the neurons' preferred directions. Across the population, PMd activity represented probability distributions of individual upcoming reaches, which depended on rapidly changing information about the body's state in the environment.

Introduction:

To plan movements, we must incorporate knowledge of the state of our bodies within the current environment. For example, if we are standing in front of a wall, we cannot walk forwards; if our arm is fully outstretched, extending it further is not possible. Considerations like these make some movements more likely than others, resulting in probability distributions over possible movements (Fig. 1a). To understand everyday movement planning, it is essential to understand how the brain represents these probability distributions.

Several studies have investigated whether the brain represents probabilities during movement planning[1-4]. In most, subjects needed to decide to move in one of a small number of directions, and the probabilities of those choices were manipulated[1-3]. These studies have shown that neurons in several brain areas have higher firing rates when there is a greater probability of an upcoming movement planned in those neurons' preferred directions (PDs). Recently, we began to study how the brain represents a continuous probability distribution rather than probabilities of discrete movements[4]. We displayed a point cloud representing an uncertain target location for movement. When we increased the uncertainty, there was a broader recruitment of dorsal premotor cortex (PMd) neurons, suggesting that PMd activity can reflect a distribution of possible movements.

Still, there is a large gap between these previous experiments and the real world, which contains dynamically changing *conditional probability distributions*, i.e. probabilities dependent on the current state of the body in the environment. As the body moves, the probability distributions of possible upcoming movements change. If the brain is to make use of these conditional probability distributions, it must rapidly compute updated probability distributions. Are these conditional probability distributions represented in the motor cortex, and if so where? How does the population of neurons function to represent these rapidly changing probabilities?

Here, we recorded from PMd and primary motor cortex (M1) while macaque monkeys reached to a series of targets that were chosen approximately randomly within the workspace. At all times, the position of the hand relative to the borders of the workspace dictated a conditional probability distribution of possible upcoming target locations. Behaviorally, the latencies and trajectories of the monkeys' movements were affected by this distribution, suggesting that they used this information during movement planning. Critically, PMd, but not M1, reflected these conditional probability distributions of upcoming movements prior to individual reaches, suggesting that such distributions are incorporated by the planning areas of motor cortex when coordinating movement.

Results:

Experiment and behavior

To study conditional probability distributions about upcoming movements, we recorded from electrode arrays chronically implanted in PMd and M1 while conducting a random-target reaching experiment[5,6]. In the experiment (Fig. 1b), three monkeys reached sequentially to four targets, before receiving a reward. About 200 ms after the cursor reached a given target, a new target appeared, to which the monkey could reach immediately. Due to the borders of the workspace, upcoming targets were more likely to be presented approximately opposite of the current hand position (Fig. 1c). That is, if a monkey's hand just landed on a target on the right side of the screen, it was more likely that the next target (and therefore, movement) would be to the left of this current hand position. Therefore, probability distributions in this experiment were conditioned on the hand's current position in the workspace at the time of target presentation.

The statistics of target presentation were not completely random within the workspace; rather, targets were slightly more likely to be selected in a clockwise direction (Fig. 1d; see *Methods*). To summarize the dependence of upcoming target locations on the current hand position, we first found the angular position of the hand relative to the center of the workspace (Fig. 1e). We analyzed the distribution of Φ's: the angular differences between the current hand position vector and the upcoming movement vector (the vector that moves the hand to the target; Fig. 1e,f). A Φ of 180° signifies that the target was exactly opposite of the current angular hand position. Importantly, this distribution was centered at 156° rather than 180° because of the slight clockwise bias in target selection. Additionally, we can see that the farther the hand position was from the center of the workspace, the more likely the upcoming target was to be in the opposite direction (Fig. 1d). When in a position near the center, there is little information about the upcoming target direction (Fig. 1g). We aimed to determine the effect of the conditional probability distributions of upcoming movements on behavior and neural activity.

To determine whether these conditional probability distributions influenced behavior, we first analyzed movement trajectories. The reaches generally did not go straight from one target to the next; they had some curvature that was influenced by the target probabilities (Fig. 2a,b). Early in the reach, trajectories were biased toward the expected target direction, defined as the most probable direction given the distribution of target presentations (i.e., 156° relative to the angular hand position; Fig. 1f). Additionally, the initial reach directions were more biased toward the expected target direction than simply toward the center of the workspace (Supplementary Fig. 1a,b). Further, when the hand position was farther from the center (and the potential target distribution was more peaked) the magnitude of this bias was larger (Fig. 2b). These behavioral results suggest that the monkeys learned and accounted for the conditional probability distributions of possible upcoming targets when planning movements.

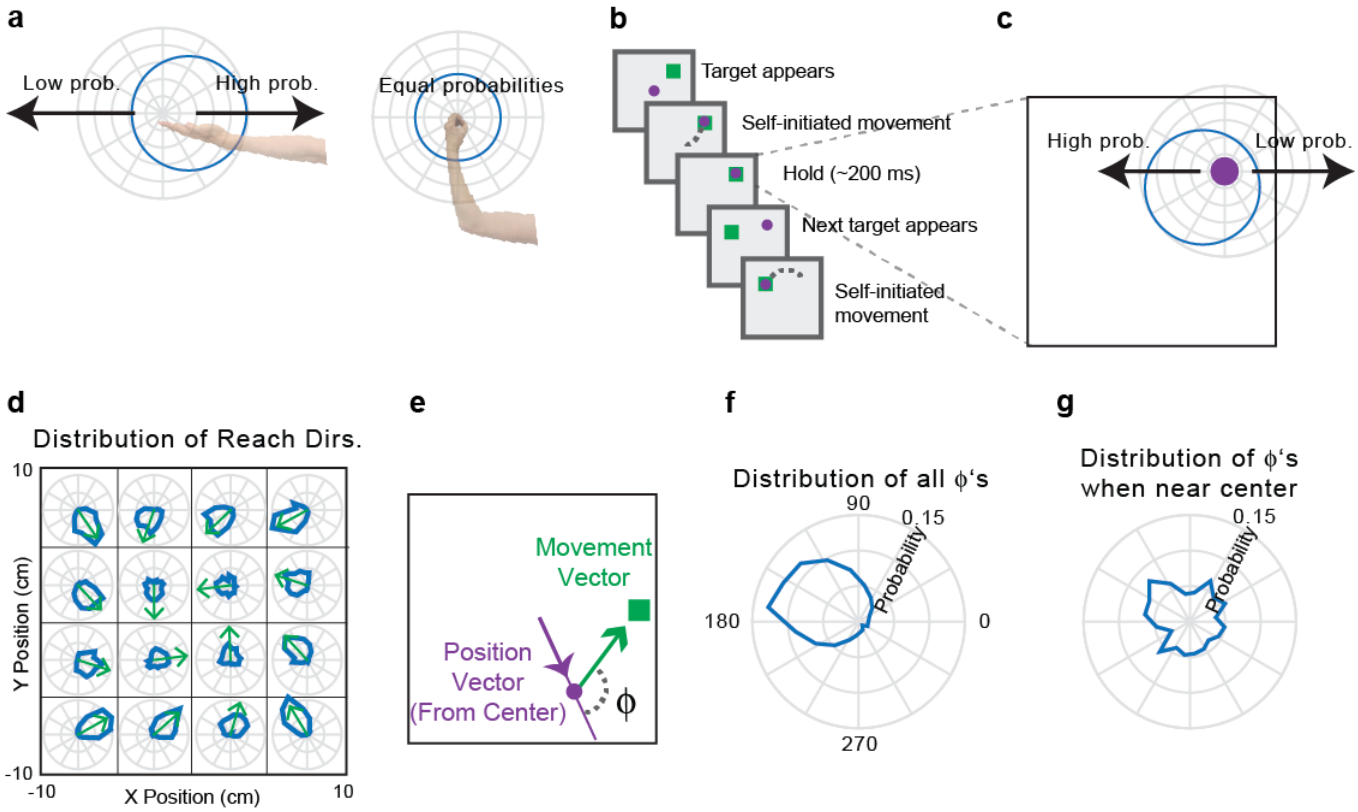

Figure 1. Experimental design and statistics.
(a) If the arm is outstretched, the only possible arm movements are back toward the body **(left)**. In other limb postures, it may be possible to move the arm in any direction **(right)**. In blue, circular probability distributions are shown for the possible upcoming movements based on the current arm posture. **(b)** Experimental design. The monkey makes sequences of four reaches, briefly holding within each target box before the next target appears. **(c)** The current hand position limits the range of possible locations of the next target, due to the borders of the workspace and target presentation algorithm. **(d)** The probability distributions of upcoming reach directions (blue) from different areas of space (x and y divided into quartiles). Green arrows point toward the circular means of the distributions. **(e)** Φ is the angular difference between the upcoming movement vector (the vector that brings the hand to the target) and the current angular hand position (relative to the center of the workspace). **(f)** The probability distribution of Φs from all hand positions. **(g)** The probability distribution of Φs from initial hand positions within 2cm of the center of the workspace.

We then analyzed how the conditional target probabilities affected movement latencies. We found shorter latencies when the target appeared close to the vector of the expected direction (Fig. 2c, Supplementary Fig. 1c; Monkey M, Pearson's r = 0.26, p < 1e-10; Monkey T, Pearson's r = 0.21, p < 1e-10; Monkey C, r = 0.045, p = 0.0018). Moreover, the distance of the hand from the center also affected the latency. For initial hand positions farther from the center (resulting in a tighter probability distribution), there was a larger latency difference between reaches to targets in expected and unexpected directions. (Fig. 2d, Supplementary Fig. 1d; Monkey M, p = 2.9e-7; Monkey T, p = 0.01; Monkey C, p = 0.089*)*. The monkeys' behaviors suggest that the motor system began movement preparation towards highly probable directions prior to target appearance.

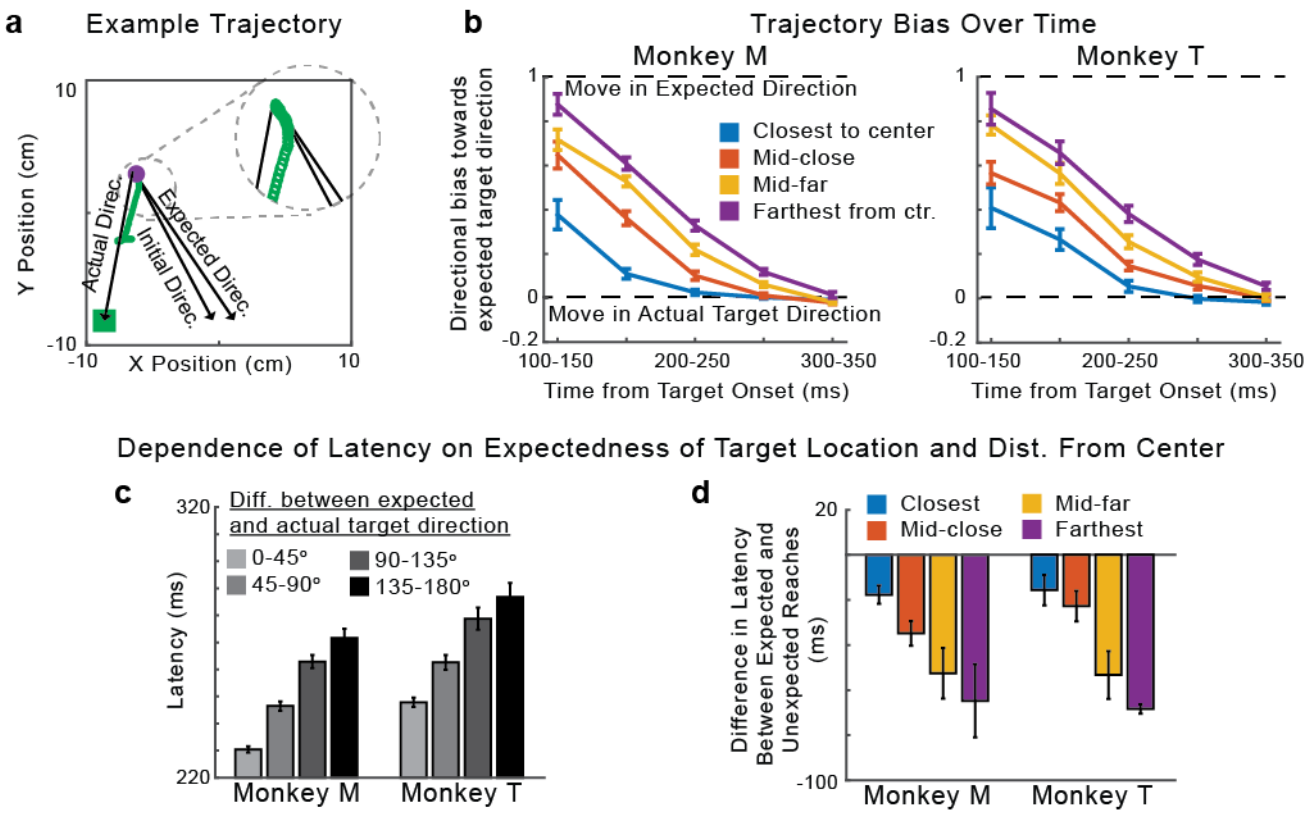

Figure 2. Behavior.
(a) An example trajectory. The initial direction of the reach (green) starts toward the expected direction of the target, given the current hand position. It later moves in the actual direction of the target. The inset shows an enlarged view of the beginning of the reach. **(b)** The median bias of the trajectory over time. A bias of 1 signifies that the direction of the trajectory is toward the expected target direction, while a bias of 0 signifies that the direction of the trajectory is toward the actual target direction. Negative biases signify movement away from the expected direction. Different traces are shown for hand positions at varying distances from the center of the workspace (divided into quartiles). Error bars are standard errors of the median. **(c)** The mean latency of reaches as a function of the angular difference between the actual and expected target directions. **(d)** The difference in mean latency between expected and unexpected reaches (expected minus unexpected), depending on the hand's distance from the center. "Expected" reaches are those that had an angular difference between the actual and expected target directions of less than 60°. "Unexpected" reaches had an angular difference of more than 120°. In panels c and d, error bars represent SEMs.

Single PMd neurons are modulated by the conditional probabilities of upcoming movements

Given that the conditional probability distributions about the potential upcoming movements affected behavior, we asked whether PMd represented this information in two monkeys. If PMd represents these probabilities, then we would expect neural activity preceding target onset to be modulated based on the anticipated possible target locations.

When observing peristimulus time histograms (PSTHs; Fig. 3a-c, Supplementary Fig. 2a-c), we found some "potential-response" (PR) neurons (nomenclature as in [7]). As expected for PMd neurons, these neurons' activity increased when a target was presented near their PDs. Crucially, PR neurons' activity was also modulated prior to target presentation by the range of possible upcoming movements. When the angular hand position was opposite these neurons' PDs (causing a higher probability that the upcoming target would be near the PD), pre-target activity increased. We also found "selected-response" (SR) neurons, whose activity was only modulated after target presentation. Thus, PSTHs suggest that a subset of PMd neurons is modulated by the probability of upcoming movements.

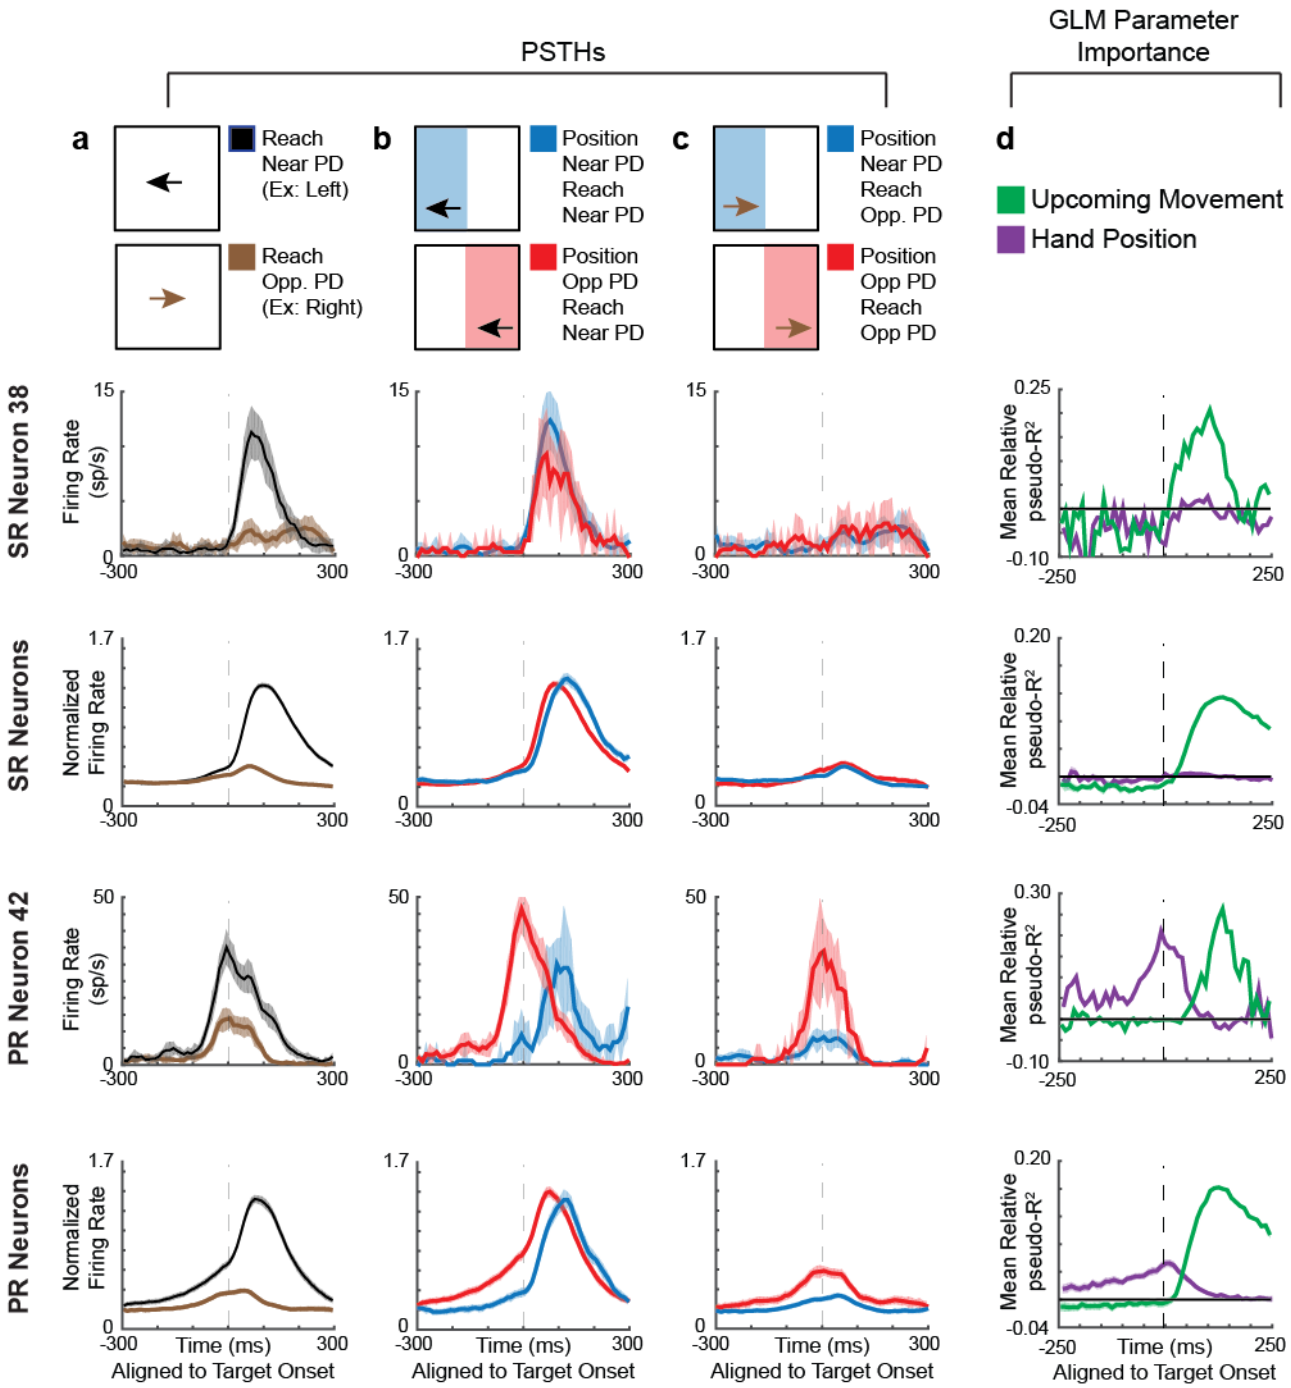

Figure 3: PMd PSTHs and GLM Results

First Row: A selected-response (SR) neuron. **Second Row**: Normalized averages of SR neurons. **Third Row**: A potential-response (PR) neuron. **Bottom Row**: Normalized averages of PR neurons. **(A-C)** Peristimulus time histograms (PSTHs) for PMd neurons, aligned to target onset. Shaded areas represent SEMs. **(A)** PSTHs of reaches near the preferred direction (PD; black) versus opposite the PD (brown). **(B)** PSTHs of reaches near the PD, with a starting hand position near the PD (lower probability of moving near the PD; blue) versus a position opposite the PD (higher probability of moving near the PD; red). **(C)** PSTHs of reaches opposite the PD, with a starting hand position near the PD (blue) versus a position opposite the PD (red). **(D)** We utilized a generalized linear model (GLM) to control for confounds in the PSTHs, including different distributions of starting positions, upcoming movements, and previous movements. Here, we show the importance of parameters in the GLM, across time, for PMd neurons. We show mean relative pseudo-R^2 over time, of the upcoming movement (green) and hand position (purple) covariates.

To analyze the factors contributing to neural activity more rigorously, we used a generalized linear modeling (GLM) approach. This approach can inform us whether the current hand position (and consequent probability distribution of upcoming target locations) significantly modulated neural activity above potential confounds related directly to the previous and upcoming movements. The GLM found that 13% (99/770) of neurons were PR and 42% (322/770) were SR neurons (see *Methods* for classification criteria). For both types of neurons, the upcoming movement covariate began to matter after target onset (Fig. 3d, Supplementary Fig. 2d). For PR neurons, but not SR neurons, the importance of the hand position covariate began to increase more than 200 ms prior to target onset, until target onset. The GLM analysis thus supports our PSTH results; prior to target presentation, PR neurons' activities are modulated by hand position, which determines the distribution of potential upcoming targets.

PMd population jointly represents the distribution of upcoming movements

An important question is how the neural population represents the probabilities about upcoming movements. We showed above that when the monkey's hand is in a position that makes an upcoming target more likely to appear near a PR neuron's PD, the neuron will have greater activity prior to target onset than it otherwise would (Fig. 3). This could be because the neural population activity is related to the statistical distribution of possible upcoming movements, conditioned on the current state. Alternatively, the neural population could be using some type of heuristic to determine the likely location of the next target (e.g., assuming the next target will always be toward the center).

To understand how the neural population activity relates to the conditional probability distributions of possible movements, we calculated the average activity (across neurons and reaches) as a function of the current angular hand position relative to each neuron's PD (Fig. 4a, Supplementary Fig. 3). Neural activity during the 100 ms prior to target appearance closely reflected the statistics of possible target locations (Fig. 4b; peak angles are not significantly different). This finding was robust to altering the correlation between the previous and upcoming movements, robust to resampling to create differing distributions of upcoming movements, and robust to subtracting out movement-related activity (Supplementary Fig. 3). Thus, when averaging across reaches and neurons, the population does represent the distribution of upcoming reaches.

How do neurons function together to create this distribution of upcoming reaches? It is possible that individual neurons reflect this distribution, and thus the population does as well. Alternatively, the distribution could be created only by many neurons working in concert. When we look at neurons with an upward preferred direction, we see that many individual PR neurons do not have maximal firing rates at hand positions corresponding to a maximum probability of moving upwards (Fig. 4c,d). Rather, these neurons have different "preferred positions", spanning many different areas at which upward movements are possible. However, when the activity of all these neurons is summed, the activity as a function of position closely matches (r=.94) the probability of an upward movement as a function of position. Conversely, we can look at reach PDs relative to preferred angular position (the angular hand position leading to peak activity). When we orient the preferred position to be down, we see that there are a wide range of reach PDs, reflecting possible upcoming movements from an initial downward position (Fig. 4e). The PDs of neurons are distributed approximately in proportion to how likely upcoming movement directions are (Fig. 4f). The population of PMd neurons works together to represent the probability distribution of available upcoming movements given the current hand position.

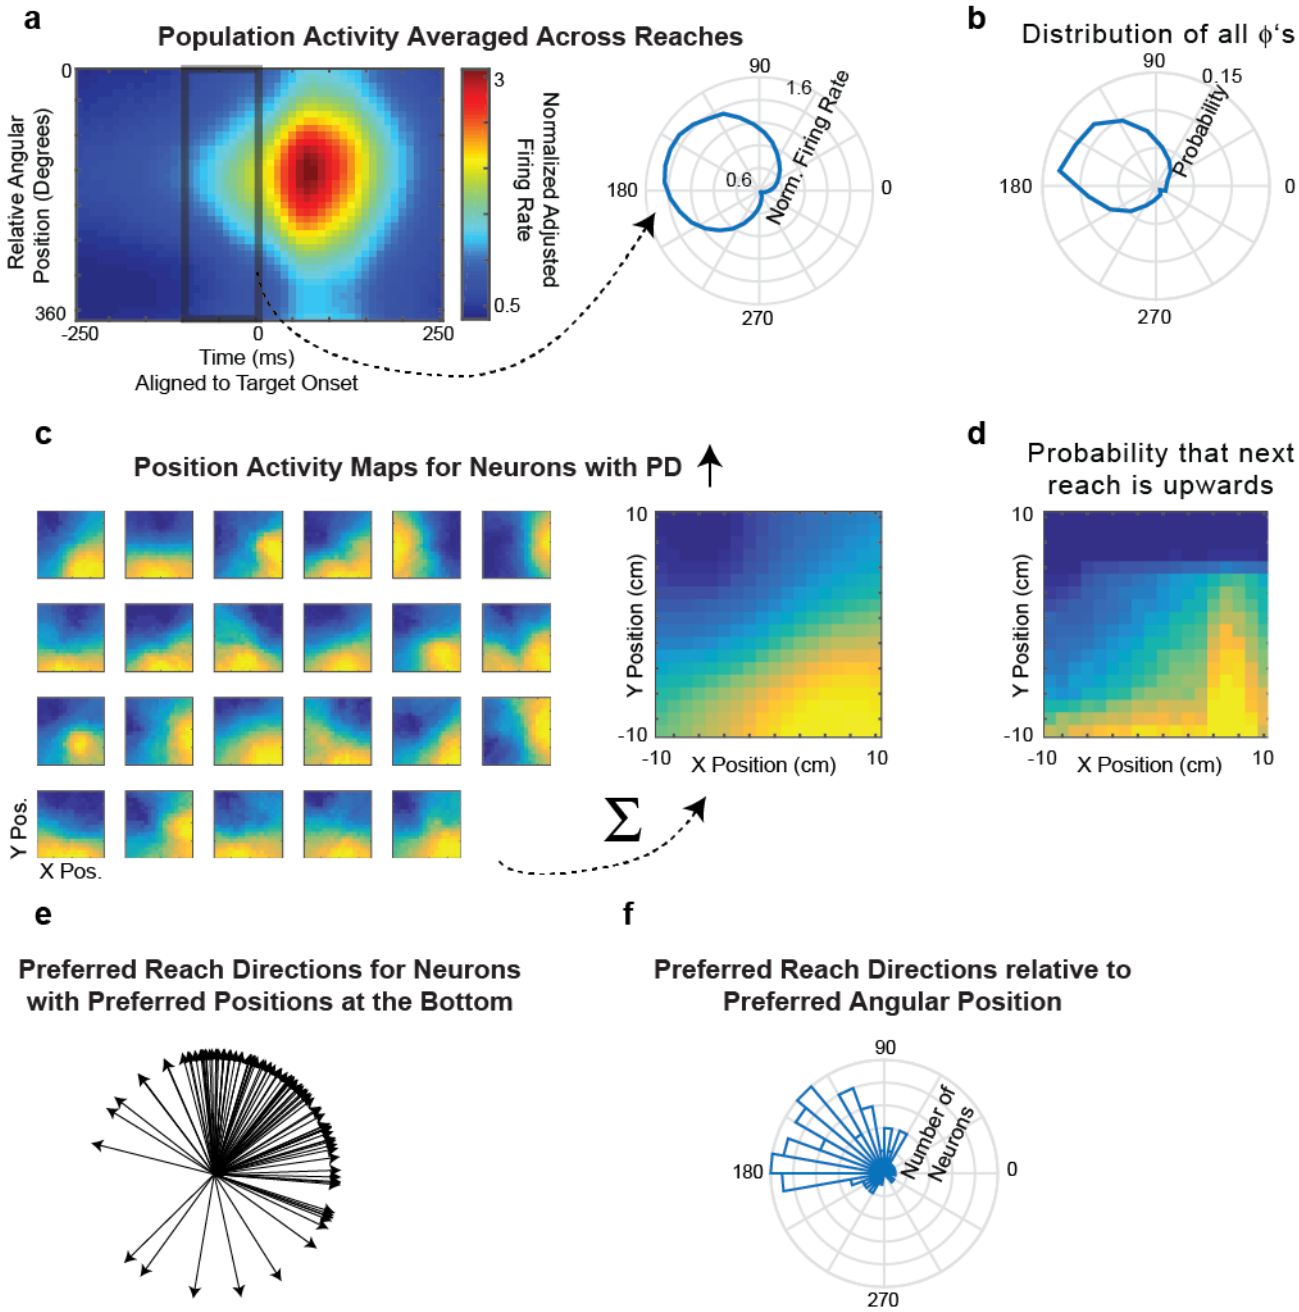

Figure 4: PMd population jointly represents the distribution of upcoming movements
a) **Left**: The average normalized firing rate of all PR neurons, over time, as a function of relative angular hand position. For each neuron, the relative angular position is the preferred direction of the neuron minus the angular hand position. Activity is normalized and averaged across all PR neurons. **Right:** The average normalized firing rate in the 100 ms prior to target onset, plotted as a function of the relative angular hand position. **b)** The distribution of upcoming movement directions relative to position angles. This is duplicated from Fig. 1f for easy comparison with panel a. **c) Left:** Position activity maps for example PR neurons with preferred movement directions that are oriented upwards. The position activity maps show the neurons' activity as a function of hand position (blue is low; yellow is high) from -100 to 50 ms surrounding target onset. **Right:** The sum of the position maps for all PR neurons, when their preferred directions are oriented upwards. **d)** A map showing the probability that the next movement will be upwards, as a function of initial hand position. **e)** Preferred reach directions for all PR neurons, when space is rotated so that their preferred hand position angle is oriented to be at the bottom (270°). **f)** A histogram of preferred reach directions of all PR neurons relative to their preferred angular hand position (the reach PD minus the preferred angular hand position).

PMd population activity represents the distribution of upcoming movements on single reaches

While we have shown that the PMd population represents conditional probability distributions averaged across trials, we also want to know what is occurring prior to single reaches. Because we recorded many neurons simultaneously, we can decode the monkey's intended movement prior to each reach. To do this, we first trained a naïve Bayes decoder to predict the reach direction (see *Methods*) during the time period 50-200 ms after target presentation. We then used this decoder (with firing rate rescaling due to differing firing rates before and after target presentation; see *Methods*) to estimate what movement the neural population was planning in the 100 ms prior to target presentation. Note that this method assumes that the PDs of neurons stayed the same between these two time periods. While PMd neurons are known to have different PDs during preparation and movement[8-10], both our time periods were during preparation. Thus, we believe it is reasonable to use knowledge about neurons' PDs after target onset to decode planning prior to target onset.

As expected, the planned reaches decoded prior to target onset were usually approximately opposite of the current angular hand position (Fig. 5a,b). This can be seen in example trials (Fig. 5a), where the pre-target decoded reach direction was to the left when the hand position was on the right, while the post-target decoded reach direction was toward the target. In general, the distribution (across reaches) of pre-target decoded reach directions relative to the angular hand position approximately represented the statistical distribution of target presentations determined by the current hand position (Fig. 5b,c; compare to Fig. 1f,g).

There are two explanations for PMd's apparent representation of the distribution of potential upcoming movement directions (Fig. 4a, Fig. 5b). One hypothesis is that PMd consistently preplans a specific reach prior to target presentation. That is, the PMd population does not actually represent a distribution of reaches on single trials, but averaging across trials yields the observed distribution. The alternative hypothesis is that PMd represents a distribution of possible movements prior to single reaches.

To distinguish between these two hypotheses, we looked at how single reach decoding depended on possible upcoming movements. The output of our probabilistic decoding method is a probability distribution reflecting the animal's movement intention encoded by the population for single reaches. Uncertainty in the population about the upcoming movement will make this distribution wider. We used the width of the decoded distribution to distinguish between the two hypotheses. If a single movement was being preplanned every reach, then the width of the decoded distributions should be approximately the same for every reach (Fig. 5d; Hypothesis 1). However, if PMd represents a distribution of possible movements, when there are fewer possibilities for upcoming target locations (i.e., when the hand position is farther from the center), the distribution should be narrower (Fig. 5d; Hypothesis 2). The data show that the distributions are narrower prior to target onset when the hand position is farther from the center (Fig. 5e, f, g; Monkey M, $p < 1e\text{-}10$; Monkey T, $p = 2.4e\text{-}6$). Thus, the decoded distributions suggest that PMd does represent a distribution of possible movements prior to single reaches.

Our decoding results provide insight into how neural activity prior to target onset influences individual upcoming reaches. The decoded pre-target reach directions were predictive of the monkeys' subsequent behavior. First, when the decoded direction was closer to the true target direction, reach latencies were shorter (Fig. 5h; Monkey M, $p < 1e\text{-}10$; Monkey T, $p = 6.3e\text{-}5$). Second, the initial direction of many reaches was initially biased toward the pre-target decoded reach direction (Fig. 5i; both monkeys $p < 0.05$ using bootstrapping). Additionally, the uncertainty of the decoded distributions influenced upcoming reaches. For targets in an expected upcoming direction, latencies were shorter when the width of the decoded distribution was narrower (Monkey M, $p < 1e\text{-}10$; Monkey T, $p = 2.3e\text{-}5$). These decoding results could provide a neural explanation for our observed latency and trajectory behavioral effects (Fig. 2). Overall, our decoding results provide insight into the expectations represented by the neural population prior to target onset.

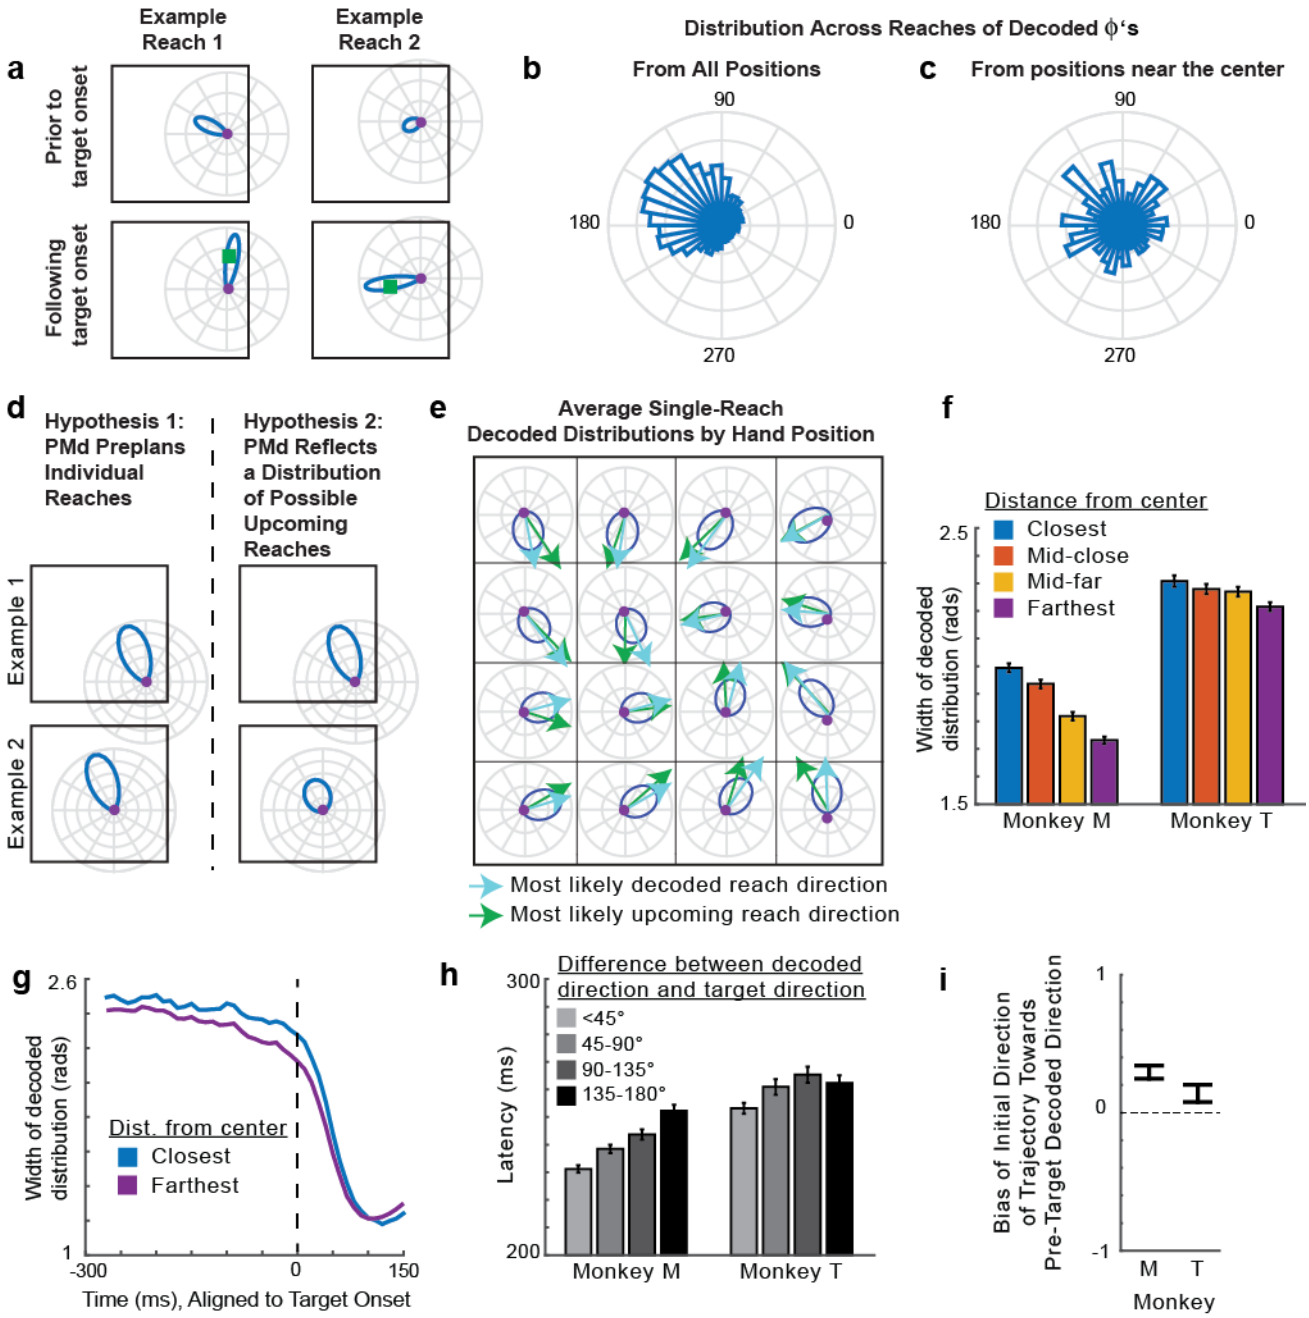

Figure 5: PMd population activity represents the distribution of upcoming movements– single reach decoding **(a)** The distribution of decoded reach directions (blue) from the population of PMd neurons for two example reaches (left and right panels), before and after target onset (top and bottom panels). As in Fig. 1b, the purple circle is the current hand position, and the green square is the target location. **(b)** The distribution across all reaches of pre-target decoded reach directions relative to the hand's angular position (this is Φ from Fig. 1, except with decoded reach angles instead of actual reach angles). **(c)** The distribution across reaches starting within 2 cm of the center, of pre-target decoded reach directions relative to the hand's angular position. **(d)** The predictions of two hypotheses (left and right), shown for two different hand positions (example 1 vs. example 2). **(e)** Average pre-target decoded reach direction distributions as a function of hand position. These distributions are constructed to have the average width and peak angle (according to circular statistics) of all reaches from the hand positions within the grid square. Light blue arrows point towards the circular means of these distributions. Green arrows point towards the most likely movement direction (from Fig. 1d). **(f)** The full width at half maximum (FWHM) of the pre-target decoded distribution as a function of hand distance from the center of the workspace. Distances are binned into quartiles. Error bars represent SEMs. **(g)** The width of the decoded distributions over time, for starting positions that are the closest (blue) and farthest (purple) from the center. **(h)** The latency of the reach as a function of the angular difference between the

pre-target decoded direction and the actual target direction. Error bars represent SEMs. **(i)** The bias of the initial trajectory of the reach (100-150ms from target onset) toward the pre-target decoded direction. A bias of 1 signifies the initial trajectory is toward the decoded direction, while a bias of 0 is toward the actual target direction. Negative values are away from the decoded direction. 95% confidence intervals, computed via bootstrapping, are shown.

Control: Visuomotor Rotation

In our task, the probabilities of upcoming target locations were determined by the current hand position. Thus, it is theoretically possible that the neural activity could only be modulated by position for some purpose other than representing upcoming movements[11,12]. As a control, we used data where the monkeys performed a visuomotor rotation (VR) learning task, which changed the probabilities of upcoming movements for the same position in the workspace. In this task, cursor feedback on the screen was rotated by 30° counterclockwise relative to the hand movement. That is, if a target were directly upwards on the screen, the monkey would now need to reach up and right to get there (Fig. 6a). The monkeys performed a block of random-target reaches with normal feedback (baseline), followed by a block of movements with the VR. In this task, the statistics of reach directions relative to the hand position on the screen rotate by 30° (Fig. 6b). Thus, for the same position in the workspace, the monkeys should plan movements 30° more clockwise than in the baseline condition. In fact, the monkeys did mostly learn to adapt their expectations of the upcoming movement directions, judging by their initial reach directions (Fig. 6c). We can therefore determine whether PMd activity changed to reflect the new movement probabilities at the same workspace positions.

Did PMd activity change to reflect the modified expected movement directions? Looking at individual neurons, the majority of PR neurons' preferred positions shifted counterclockwise in the VR task (Fig. 6d,e). This was expected, as the hand position most likely to result in a movement into a neuron's PD will rotate counterclockwise in the VR task. There was a wide range of changes across neurons, again demonstrating that distributions are represented across the population.

To look at population level changes, we decoded reach direction using activity prior to target onset, and compared the distributions of these decoded reaches in the baseline and rotation periods of the task (Fig. 6f,g). As the rotation was small, we looked only at hand positions not near the center, where monkeys had more information about the upcoming movement directions. In every session for both monkeys, the reach directions shifted counterclockwise relative to the workspace position. On average, there was a small, but significant, clockwise shift in decoding of about 10° for each monkey (Fig. 6d,e; $p < 0.05$ for both monkeys; for all decoded reaches see Supplementary Fig. 4a,b). Additionally, when we analyzed the data by looked at the average activity of PR neurons prior to target onset (as in Fig. 4a) rather than decoding, we found a similar shift in activity in every experimental session (Supplementary Fig. 4c,d). Thus, PMd activity prior to target onset is modulated by the probabilities of upcoming movements, not simply by hand position for some other purpose.

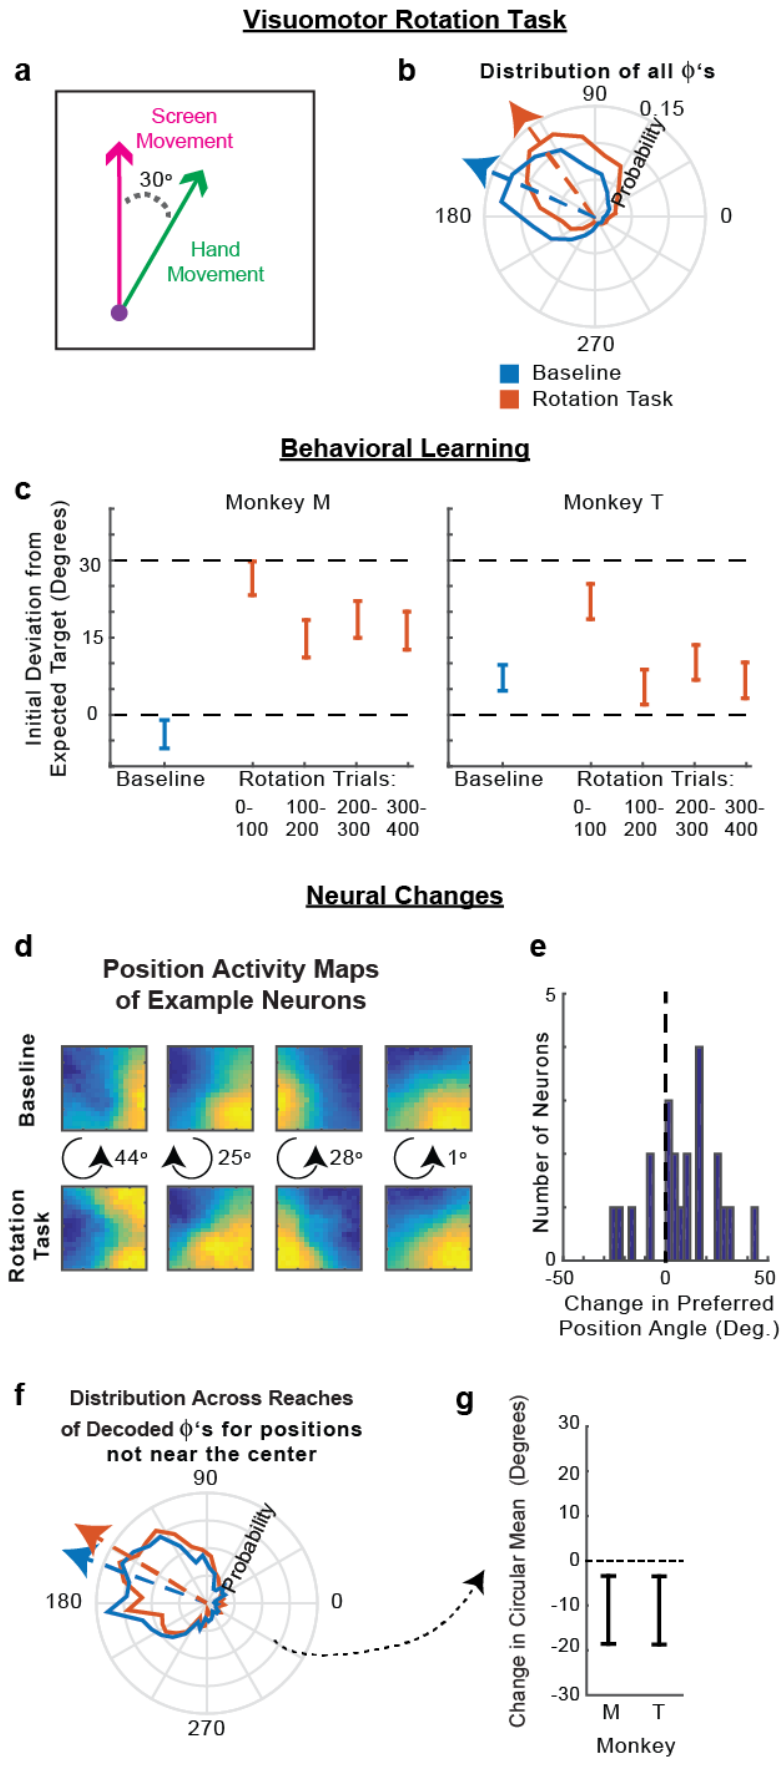

Figure 6: Visuomotor Rotation Control Task
(a) The visuomotor rotation (VR) task. Movements on the screen (in the workspace) are rotated 30° counterclockwise relative to the hand movement. **(b)** The distribution of hand movements relative to the position angle in the workspace, i.e. Φs, for the baseline task (blue) and the VR task (orange). Arrows point toward the circular means of the distributions. **(c)** The difference between the initial reach direction (100-150ms from target onset) and the expected movement direction during the baseline task (blue) and different periods during the VR task (orange). The expected movement direction was the most likely upcoming movement direction given the current workspace position and movement statistics (panel b). Positive values mean the initial reach direction was counterclockwise of the expected target direction, meaning the monkey had not adapted. **(d)** Position activity maps (activity as a function of position in the workspace) of example PR neurons in the baseline (top) and VR task in the second 2/3 of trials (bottom), as in Fig. 4c. In the middle, we show the direction (clockwise or counterclockwise) and magnitude of change of the preferred angular position. **(e)** The change in preferred angular position of all PR neurons (VR minus baseline). Positive means a counterclockwise shift. **(f)** The distribution of pre-target decoded reach directions relative to the hand's angular position (decoded Φs) for positions not near the center (greater than the median distance). Decoding from the VR task used the second 2/3 of trials. **(g)** The difference between the circular mean of the distributions of decoded Φs in panel d, between the baseline and VR tasks (VR minus baseline). Error bars represent 95% confidence intervals from bootstrapping.

M1 does not represent the conditional probability distribution of upcoming movements

To determine whether primary motor cortex (M1) also represents conditional probability distributions of upcoming movements, we ran the same set of analyses for M1 as we did for PMd. When we did a GLM analysis, we found that many neurons have significant modulation with hand position. However, these neurons did not respond to hand position in the same manner as the PMd neurons did. Instead, their activity prior to target onset occurred when the hand position was in the same direction as the neurons' PDs, rather than opposite it (Fig. 7, Supplementary Fig. 5). This activity could be explained by the end of the previous movement, since previous movements into the PD (which correspond to higher M1 activity), often result in angular hand positions in the same direction as the PD. Moreover, the effect of the hand position didn't ramp up as it did in PMd; rather, it appears to be a decreasing effect from the previous movement (Fig. 7, Supplementary Fig. 5). Thus, while M1 activity varies according to position (as in [13,14]), it likely does so in a way that reflects movement execution, rather than information about the upcoming movement.

Discussion

In order to plan everyday movements, we take into account the probability distributions of possible movements determined by the state of our body in the environment. Here, we have demonstrated that these conditional probability distributions influence behavior, specifically movement trajectories and latencies. Populations of neurons in PMd, but not M1, function together to represent these probabilities, even prior to individual reaches. We used a visuomotor rotation task to show that the effect was not simply a position dependent component of the firing rate.

Information that shapes conditional probability distributions of upcoming movements

Probability distributions can be conditioned on many sources of information. Here, we demonstrated that PMd represents probability distributions that were determined by the body's state in the task environment. When hand position within the environment changed, the probability distributions of upcoming movements changed. Likewise, changes in the task environment caused by the visuomotor rotation task altered the probability distributions. Although not a significant component of our study, another source of probability distributions is biomechanics. In the extreme, biomechanics limit the possible movements. Softer constraints may arise from biomechanical costs, e.g., the ease of movement, which can affect both choice of arm movements[15,16] and PMd activity[16]. While biomechanics may influence the representation within PMd of possible upcoming movements, the fact that PMd's probability distributions also changed when the task environment changed in the VR task, suggests our findings are not primarily due to biomechanical constraints.

In order to represent the probability distributions of upcoming movements, PMd needed to have information about the hand position in the workspace. Previous studies have shown that PMd activity is modulated by hand position, either to make the neurons' PDs compatible with the orientation of the arm[11], or to encode the relative position between the hand and eye[12]. Thus, there is evidence that PMd neurons have access to information about hand position. However, in our task PMd began to represent the possible upcoming movement while the current movement was still occurring and hand position was changing. Thus, it is more likely that PMd was making use of visual information about the current target, or an efference copy of the current movement command, to determine the probability distribution of the next movement. In everyday life, PMd likely uses a mixture of sensory, proprioceptive, and movement information to determine possible upcoming movements.

Representation of probability distributions across the population of neurons

There is much debate on how probabilities are represented in the brain. Some argue that probability distributions are represented across populations of neurons[17-19] (e.g. probabilistic population coding[18]), while others believe there is a temporal coding of uncertainty[20]. We showed that neurons were more strongly active in workspace locations from which movements into their preferred direction were more likely. Only when looking across neurons did the activity as a function of location closely match the task's movement probabilities. Also, when looking at neurons that were active at a nearby location, the distribution of PDs was proportional to the probability of upcoming movement directions from

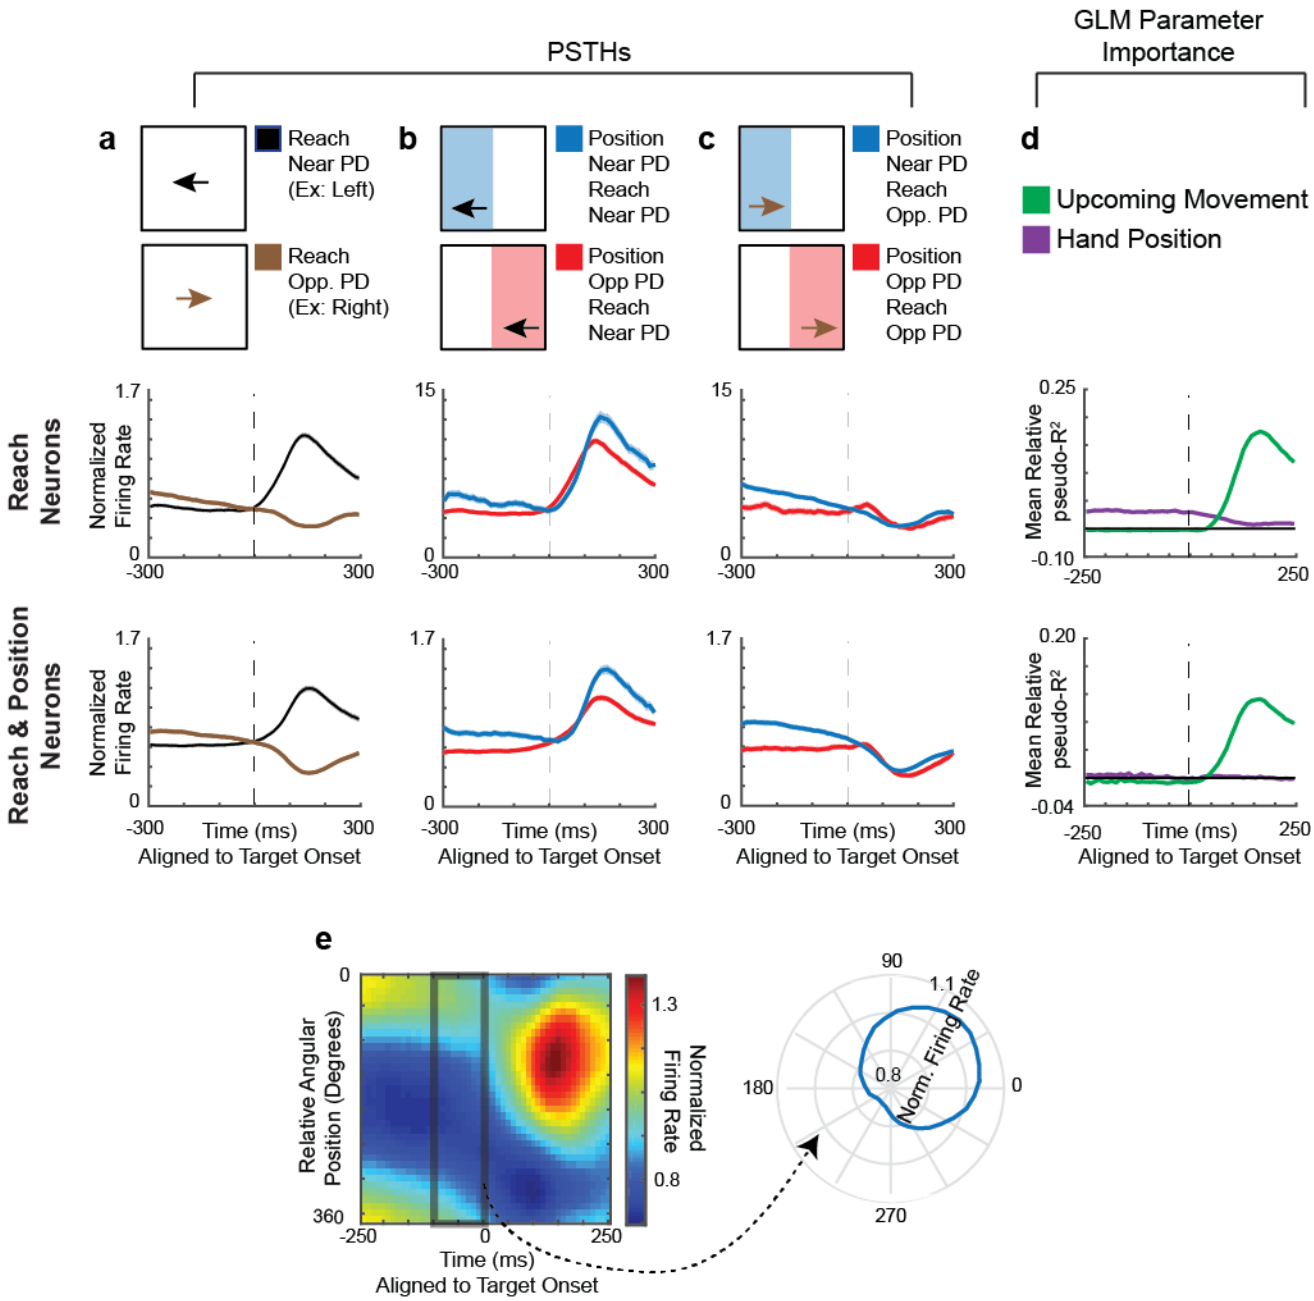

Figure 7: M1 does not reflect the probability of upcoming movements
(a-d) PSTHs and GLM results for M1 neurons. Columns have the same schematics as Fig. 3. **First Row of PSTHs**: Normalized averages of reach neurons, defined as those neurons significant for movement during the late period, but not position in the early period of the GLM. **Second Row**: Normalized averages of reach & position neurons, defined as those neurons significant for movement during the late period, and position in the early period of the GLM. Note that we did not use the same "SR/PR" nomenclature as PMd, because there was not evidence in the PSTHs of M1 neurons that position was used to represent potential upcoming movements. **(e)** Same schematic as Fig. 4a, but for M1 neurons. **Left**: The normalized average firing rate, as a function of time and relative angular position. Activity is averaged across all reach & position neurons. **Right:** The normalized average firing rate in the 100 ms prior to target onset, plotted as a function of the relative angular hand position.

that location. Our results suggest that coding of probability distributions across populations of neurons plays a central role in the movement decision process.

PMd's role in representing movement possibilities

Our study builds on much research about the role of PMd in planning upcoming movements[4,7,21-25]. Previous work has demonstrated that monkeys represent possible movements when selecting between alternatives[7,22-25] and when estimating the likely target location from visual cues[4]. These studies suggested that PMd can represent a probability distribution. Our work extends these findings by showing that PMd also represents dynamically changing probability distributions that are dependent on interactions between the body and the environment. Moreover, our work shows that PMd does not only represent probability distributions that are explicitly manipulated. Here, PMd represented probability distributions even in a standard reaching task that has been used in a variety of motor studies[5,6], where probabilities are usually considered to be irrelevant. The representation of conditional probability distributions of possible movements in PMd appears to be ubiquitous.

Methods

Behavioral Paradigms

Random-target Experiment

Three monkeys (Monkeys M, T, and C) performed a reaching task in which they controlled a computer cursor using arm movements (Fig. 1). Monkeys were seated in a primate chair while they operated a two-link planar manipulandum. Arm movements were constrained to a horizontal plane within a workspace of 20 cm x 20 cm. On each trial, the monkey consecutively reached to 4 targets (2 cm x 2 cm squares), with each new target appearing once the monkey reached the previous target. More precisely, once a target was reached, a new target was triggered 100 ms later, as long as the cursor remained on the target. The target appeared on-screen 96 ms after this trigger on average, due to delays from graphics processing and the monitor refresh rate. In accordance, the monkey was required to keep the cursor on the target for an additional 100 ms after a new target was triggered. Thus, in total there was a 200 ms hold period after landing on the target. This brief hold period forced the monkeys to decelerate as they approached the target, but was not so long that the monkeys completely stopped on the target. After a successful trial (4 successful reaches), the monkey received a liquid reward. The next trial started after a delay of one second with a new random target presentation.

Target locations were chosen to be 5 – 15 cm from the current target. Specifically, they were chosen as follows. 1) Randomly choose a distance between 5 and 15 cm, and an angle between 0° and 360° for the new target (relative to the current target). 2) If the new target falls outside of the workspace, add 90° to the angle and set the distance to be 5 cm. 3) Repeat step 2 until the target is in the workspace.

Many of the analyses are aligned to target onset. These experiments did not use a photodiode to determine the exact moment the target was displayed. Thus, in all analyses, the target onset time we used was the time the computer sent the target command plus the average delay time (96 ms).

In total, we recorded 8 sessions for monkey M, 6 sessions for monkey T, and 5 sessions for monkey C.

Visuomotor Rotation Experiment

In 3 sessions each, monkeys M and T performed a visuomotor rotation (VR) task following the baseline random-target task. The VR task was equivalent to the random-target task, with the exception that the movement vectors displayed on the screen (workspace) were rotated 30° counterclockwise relative to the hand movement vectors (as in Fig. 6a).

Neural Data Acquisition and Preprocessing

Monkeys M and T were implanted with 100-electrode Utah arrays (Blackrock Microsystems, Salt Lake City, UT) in dorsal premotor cortex (PMd). Monkeys M and C were implanted with Utah arrays in primary motor cortex (M1). See [4] for the location of the arrays in Monkeys M and T. Units were manually sorted with Offline Sorter (Plexon, Inc, Dallas, TX, USA). Only well-isolated individual units were included. Since we used chronically implanted arrays, it is likely that some neurons were recorded on multiple sessions and thus were not unique.

We only included neurons with firing rates of at least 2 spikes / second in either the early or late period. The early/late periods were defined as -100 to 50 and 50 to 200 ms from target onset, respectively. In PMd, this left us with 520 neurons from Monkey M and 250 neurons from monkey T. In M1, this left us with 352 neurons from Monkey M and 266 neurons from Monkey C.

Behavioral Analysis

Each trial consisted of 4 reaches. We did not include the first reach in any of our analyses, as this was preceded by a reward period without movement (rather than being in the midst of continual movement). We also only included successful movements to the target (those that

were held at the target for at least 200 ms). Behavioral data was combined across all sessions for each monkey.

Statistics of target presentation

We defined the angular position, ϕ_P, as the hand position (prior to movement) relative to the center of the workspace (Fig. 1). We defined ϕ as the angular difference between the upcoming movement direction (also the direction to the target), ϕ_T, and the angular position. That is, $\phi = \phi_T - \phi_P$.

Trajectory bias

We calculated whether the movement trajectory within a given time interval was biased toward the expected target direction, ϕ_E. The expected direction was the most likely direction of the next target given the current hand position, based on the distribution of ϕ's. So if ϕ^* is the value corresponding to the circular mean of the distribution of ϕ's, $\phi_E = \phi_P + \phi^*$. The bias of the movement trajectory within a given time interval was defined as follows. First, a movement direction, ϕ_M, was determined within that time interval based on the start and end hand position in that time interval. We calculated the bias, $B = \frac{\phi_M - \phi_T}{\phi_E - \phi_T}$, where the numerator and denominator were made to be in the interval of $[-180° \; 180°]$ prior to dividing. When the current movement direction is toward the expected direction, B will be near 1, and when the movement direction is toward the actual target direction, B will be near 0. B can also be negative when the movement direction is away from the expected direction. For the summary statistics of B, we used the median and standard error of the median, as B has outliers when dealing with circular variables. To calculate the standard error of the median, we used bootstrapping. Note that in Supplementary Fig. 1, we also calculated the bias toward the center, rather than the expected direction. This has the exact formulation as above, except $\phi^* = 180°$.

Latency effects

The latency of a reach was defined as the time from target onset until the movement surpassed a velocity of 8 cm / sec. Latencies greater than 6 standard deviations from the mean were excluded as outliers.

We computed the mean latency of movements as a function of the expectedness of the target location, which was defined as the difference between the target direction and expected direction: $|\phi_T - \phi_E|$. We calculated the Pearson's correlation between latency and the expectedness of the movement, and determined significance based on the p-value associated with the correlation (2-sided t-test).

We also analyzed differences in latencies between expected reaches (expected direction < 60° from target direction) and unexpected reaches (expected direction > 120° from target direction), based on the distance of the hand position from the center of the workspace. To test whether the latency of expected reaches decreased as a function of distance from the center more than unexpected reaches, we used linear regression to fit the latency of reaches as a function of distance from the center. We then did a 2-sided t-test to analyze whether the slope was less (more negative) for expected reaches.

Neural Data Analysis

As with the behavioral analyses, we only included successful reaches.

Determining PDs of Neurons

We determined the preferred movement direction (PD) of each neuron from 50 to 200 ms following target onset. Let $\boldsymbol{Y}$ be the vector of firing rates in that interval for every movement. It is possible that some of the neural activity during these time periods was related to the hand position, rather than the upcoming movement. Thus, we first aimed to remove the effect of any position-related signal that might bias the calculated PD. To do so, we created a smoothed map of how position was related to neural activity (see below), and then subtracted the position-related activity from the total activity (leaving activity due to movement). That is, if $\boldsymbol{\theta_P}$ is the vector of expected firing rates at given positions, then we fit the tuning curves to this "firing rate due to movement", $\boldsymbol{Y} - \boldsymbol{\theta_P}$. More specifically, we fit a von Mises function to relate the movement directions to the firing rate due to movement:

$$\boldsymbol{Y} - \boldsymbol{\theta_P} = \alpha \exp\left[\beta \cos(\boldsymbol{\phi_M} - \phi_M{}^*)\right]$$

where $\boldsymbol{\phi_M}$ is the vector of movement directions, and α, β, and $\phi_M{}^*$ are the parameters that we fit. $\phi_M{}^*$ is the PD of the neuron.

Preferred angular positions

We determined the preferred angular (hand) position of each neuron from -100 to 50 ms following target onset. We first aimed to remove the effect of any movement-related signal that might bias the calculated preferred angular position. To do so, we created a smoothed map of how movement was related to neural activity, and then subtracted the movement-related activity from the total activity (as in the previous section). We then fit a von Mises function to relate the angular positions to this residual. This fitting is identical to the fitting of PDs, except replacing movement directions with angular hand positions.

PSTHs

When plotting PSTHs of individual neurons, we plotted the mean firing rate across movements. The error bars on PSTHs are the standard error of the mean (SEM) across movements. When plotting the PSTHs averaged across neurons, we first normalized the mean firing rate (across time) for each neuron by dividing by the

maximum firing rate of the average trace (across all conditions). We then plotted the average of these normalized firing rates across neurons. Error bars are the SEM across neurons. All traces were smoothed using a 50 ms sliding window.

PSTHs were made for different categories of movements. For the PSTHs, movements near the PD were those that were within 60° of the PD. Movements opposite the PD were those greater than 120° from the PD. Hand positions opposite the PD were angular positions greater than 120° away from the PD. Hand positions near the PD were angular positions less than 60° away from the PD.

Smoothed maps of neural activity

For many aspects of the following neural data analysis, we computed smoothed maps of neural activity in relation to some variable (hand position, previous movement, or the upcoming movement). For instance, we created a map of how neural activity varied over all hand positions in the workspace, and a map of how neural activity varied in response to all upcoming movement vectors. For our maps, rather than assuming a parametric form, we non-parametrically estimated the average firing rate at each point in space using weighted k-nearest neighbor smoothing. The parameters were the number of nearest neighbors, k, and a decay parameter, d. As an example, for the movement variable (previous or upcoming), for each movement we found the k nearest movement vectors (based on Euclidean distance). We then averaged the firing rates associated with each of the k movements, but with each weighted proportional to its distance from the given movement vector to the d power.

The parameters we used for the generalized linear models (GLMs) were k = 20% of the data points, d=0. The parameters were found using cross-validation on a held out data set, in order to not inflate the number of significant neurons in the GLM analysis. The parameters used for plotting the position maps (Figs. 4 and 6) were k = 20% of the data points, d=-0.5, in order to show slightly higher resolution. The parameters we used at other times were k = 30% of the data points, d=-1. These parameters were found using cross-validation on the current data sets in order to create as accurate maps as possible. Importantly, all results were robust to a wide range of smoothing parameters.

For visualizing the position maps in Fig. 4c, they were rotated either 90, 180, or 270° so that the PD of that neuron (after the same rotation) was always upward (between 45° and 135° relative to horizontal), which facilitated the interpretation and comparison of the results.

Generalized Linear Model

To determine which variables were reflected in the neural activity, we used a Poisson Generalized Linear model (GLM). Let $\boldsymbol{Y}$ be a vector containing the number of spikes in the time interval we are considering, for every movement. It has size $m\ x\ 1$. We aimed to predict $\boldsymbol{Y}$ based on several factors. We used the hand position, the previous movement vector, the upcoming movement vector, the peak velocity of the upcoming movement, and a baseline term.

Overall, the model that generates the firing rate ($\boldsymbol{\lambda}$; also known as the conditional intensity function can be written as:

$$\boldsymbol{\lambda} = \exp\left(\boldsymbol{X\beta}\right)$$

where $\boldsymbol{\beta}$ is a vector of the weights for each covariate that we fit, and $\boldsymbol{X}$ is the matrix of covariates, which is z-scored before fitting. If there are j covariates, then $\boldsymbol{\beta}$ has size $j\ x\ 1$. $\boldsymbol{X}$ has size $m\ x\ j$. Note the use of an exponential nonlinearity to ensure that firing rates are positive. The model assumes that the number of spikes, $\boldsymbol{Y}$, is generated from the firing rate, $\boldsymbol{\lambda}$, according to a Poisson distribution.

For hand position, previous movement, and upcoming movement, the covariates were the smoothed maps relating each of these variables to neural activity. These maps essentially represent the expected firing rate for each variable by itself. Note that the previous and upcoming movement covariates were fit separately and do not need to have the same smoothed map (as PDs can be different during planning and movement [8-10]).

We fit the model weights to the data using maximum likelihood estimation. That is, we found $\boldsymbol{\beta}$ that was most likely to produce the true spike output (assuming spikes were generated from the firing rate in a Poisson nature). Critically, we used (5-fold) cross-validation, meaning that the model was fit to the data using one set of data (the training set), and model fits were tested with an independent set of data (the testing set). Similarly, when calculating the test set covariates for movement and position (described in *Smoothed maps of neural activity*), we only used k-nearest neighbors from the training set, to avoid overfitting.

To test whether an individual covariate significantly influenced neural activity, we first made sure that a simplified model with only that individual covariate had significant predictive power. To determine the value of a model fit, we used pseudo-R^2 [26,27], a generalization of R^2 for non-Gaussian variables, which is defined as:

$$R_D{}^2(\text{model}) = 1 - \frac{\log L(n) - \log L(\hat{\lambda})}{\log L(n) - \log L(\bar{n})}$$

where $\log L(n)$ is the log likelihood of the saturated

model (i.e., one that perfectly predicts the number of spikes), log $L(\hat{\lambda})$ is the log likelihood of the model being evaluated, and log $L(\bar{n})$ is the log likelihood of a model that uses only the average firing rate.

Then, in order to determine the importance of that covariate to the full model, we test whether the full model predicts neural activity significantly better than a model where that covariate is left out (reduced model). To compare the fits between the full model and reduced model, we used relative pseudo-R^2, which is defined as:

$$R_D{}^2(\text{model 1, model 2}) = 1 - \frac{\log L(n) - \log\ L(\hat{\lambda}_2)}{\log L(n) - \log\ L(\hat{\lambda}_1)}$$

where log $L(\hat{\lambda}_2)$ is the log likelihood of the full model and log $L(\hat{\lambda}_1)$ is the log likelihood of the reduced model.

To determine significance, we bootstrapped the fits to create 95% confidence intervals, and checked whether the lower bounds of these confidence intervals were greater than 0. Note that the pseudo-R^2 and relative pseudo-R^2 values can be less than 0 due to overfitting.

Neuron types

In PMd, we defined selected-response (SR) neurons as those that were significantly modulated by upcoming movement in the late period in the GLM, but were not significantly modulated by hand position in the early period. Potential-response (PR) neurons were significantly modulated by upcoming movement in the late period and by hand position in the early period. While our criteria for determining SR and PR neurons was different from [7], we used the same terminology due to the same perceived function. In the VR task, PR neurons were those that were significant during both the baseline and VR periods.

Population activity over time averaged across trials

For each neuron, we calculated the firing rate as a function of the relative angular position (Fig. 4). We defined the relative angular position as the difference between a neuron's PD and the hand's angular position (the PD minus the angular hand position). In the VR task (Supplementary Fig. 4c), the relative angular position was calculated relative to the neuron's PD in the baseline task. We then normalized each neuron by dividing by its mean firing rate, and then averaged the normalized activity across neurons. We then smoothed the activity for plotting using the parameters from the smoothed maps.

We also made several variants of the above plot (Supplementary Fig. 3). We made a plot where movements were only used if the angle between previous and upcoming movements was less than 90°. We made a plot in which the movements were resampled, so that the distribution of ϕ's was centered at 180°, rather than being off center. More specifically, we resampled from a von Mises distribution: $g(\phi) \propto \exp\left[\cos(\phi - 180°)\right]$. We also made a plot where we used firing rate residuals, where activity due to the upcoming movement was subtracted from the total activity, leaving residuals (that still depend on position).

To determine the relative angular position resulting in peak activity (in the 100 ms before target onset), we calculated the activity at 20 relative angular positions (evenly spaced from 0 to 360°), and calculated the circular means of the angles weighted by their activities. We determined whether the activity prior to target onset was related to the distribution of upcoming movements by testing whether the relative angular position resulting in peak activity was significantly different from 156 ° (the circular mean of the distribution of ϕ's). We created a 95% confidence interval of peak relative angular positions by bootstrapping over the set of neurons, and checked whether this overlapped with 156°.

Decoding

We aimed to determine the movement intention of the neural population in the 100 ms prior to target onset. As only including PR neurons would give us a small number of neurons per session for decoding, we expanded our criteria. While PR neurons were significant for hand position and upcoming movement with 95% confidence, here we included neurons that were significant at a level of 50% (the median pseudo-R^2 and relative pseudo-R^2 values were greater than 0). For comparing between the VR task and baseline, we required that neurons had positive median pseudo-R^2 values in both conditions. Additionally, as neurons changed from session to session, separate decoders were trained for each session.

We first fit tuning curves to each neuron during 50 to 200 ms after target onset, when the neurons' preparatory responses to different target directions was known. This was done using a von Mises function, as in *Determining PDs of Neurons*. We wanted to use these tuning curves to decode during the 100 ms prior to target onset. However, as firing rates were greatly different during these two time periods, we needed to rescale the tuning curves. To do so, we fit tuning curves to the future movement in the 100 ms prior to target onset. We then modified this tuning curve by giving it the preferred direction calculated after target onset. This essentially gives us rescaled versions of the tuning curves determined when the target is known.

Note that for decoding in the VR task, we still fit the initial tuning curve using activity after target onset in the baseline task. We then decoded using activity from before target onset in the VR task, using the rescaling

described above. We use this procedure to make the comparison meaningful.

Next, for each neuron, we found the likelihood of the number of spikes given all possible movement directions (in 1 degree increments). This was done by assuming the number of spikes during the time period is a Poisson random variable with a mean rate determined by value of the tuning curve at the angle being tested. If r_i is the number of spikes during the interval for neuron i, s is the angle, and $f_i(s)$ is the value of the tuning curve (the expected number of spikes) for neuron i at angle s:

$$P(r_i|s) = \frac{\exp\left[-f_i(s)\right] f_i(s)^{r_i}}{r_i!}$$

We assumed that neurons' activities were conditionally independent given the angle (a naïve Bayes decoder), and thus multiplied their probability distributions:

$$P(\boldsymbol{r}|s) \propto \prod_i P(r_i|s)$$

We can use Bayes rule to determine the likelihood of all the movement directions given the number of spikes of all neurons. Assuming a uniform prior, by Bayes rule:

$$P(s|\boldsymbol{r}) \propto P(\boldsymbol{r}|s)$$

Finally, we normalized $P(s|\boldsymbol{r})$ (so it was a probability distribution), and this term was the decoded distribution.

The decoded angle was the angle corresponding to the peak of the distribution (the maximum likelihood decoded angle). The width of the decoded distribution was the full width half maximum (FWHM) of the decoded distribution.

To calculate the width of the decoded distribution over time (Fig. 5g), we decoded using a 50 ms sliding window of neural activity. All methods were the same as above, just replacing the 100ms of activity prior to target onset with the given 50 ms of activity. This choice allowed us a better temporal resolution.

As we did decoding separately for each session, to do significance testing, we used a simple multilevel model analysis – specifically, a random intercepts model, where the baseline (intercept) can be different for every session. Thus, if there were 4 sessions, when looking at the width of the decoded distribution ($\boldsymbol{w}$) as a function of the distance from center ($\boldsymbol{d}$), we wrote the model as $\boldsymbol{w} = \beta_1\boldsymbol{I}_1 + \beta_2\boldsymbol{I}_2 + \beta_3\boldsymbol{I}_3 + \beta_4\boldsymbol{I}_4 + \beta_5\boldsymbol{d}$, where $\boldsymbol{I}_1$ through $\boldsymbol{I}_4$ are indicator variables for whether the values are from a given session. We looked at whether β_5 was significantly different from 0 using a 2-sided t-test. We used an equivalent approach to determine the significance of the relationship between latency and the difference between the decoded angle and the target angle. We also used an equivalent approach to determine the significance of the relationship between latency and the width of the distribution for expected reaches (expected reaches are those where the difference between the actual and expected target direction is less than 60°).

We calculated the bias of the initial trajectory toward the decoded angle, ϕ_D, equivalently to how we calculated the behavioral bias toward the expected target direction:
$B = \frac{\phi_M - \phi_T}{\phi_D - \phi_T}$.

To determine whether the distribution of decoded reach directions shifted from the baseline task to the VR task, we calculated the difference between the circular means of the distributions of "decoded Φs" (the decoded reach direction relative to the angular position). For significance testing, we bootstrapped this difference in circular means. More specifically, 1000 times, we resampled decoded reaches within the baseline task and calculated the baseline circular mean, and did the same thing for the VR task. This led to 1000 differences in circular means. We looked at the 95% confidence interval of this difference.

Acknowledgements

We would like to thank Brian Dekleva, Daniel Wood, Pat Lawlor, Mark Segraves, Hugo Fernandes, and Xuelong Zhao for helpful comments and discussions.

For funding, JG would like to thank NIH F31 EY025532 and NIH T32 HD057845. MP would like to thank NIH F31 NS092356 and NIH T32 HD07418. PR, LM and KK would like to thank NIH R01 NS074044.

References

1 Basso, M. A. & Wurtz, R. H. Modulation of neuronal activity in superior colliculus by changes in target probability. *The Journal of Neuroscience* **18**, 7519-7534 (1998).

2 Dorris, M. C. & Munoz, D. P. Saccadic probability influences motor preparation signals and time to saccadic initiation. *The Journal of Neuroscience* **18**, 7015-7026 (1998).

3 Yang, T. & Shadlen, M. N. Probabilistic reasoning by neurons. *Nature* **447**, 1075-1080 (2007).

4 Dekleva, B. M., Ramkumar, P., Wanda, P. A., Kording, K. P. & Miller, L. E. Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. *eLife* **5**, e14316, doi:10.7554/eLife.14316 (2016).

5 Serruya, M. D., Hatsopoulos, N. G., Paninski, L., Fellows, M. R. & Donoghue, J. P. Brain-machine interface: Instant neural control of a movement signal. *Nature* **416**, 141-142 (2002).

6 Flint, R. D., Lindberg, E. W., Jordan, L. R., Miller, L. E. & Slutzky, M. W. Accurate decoding of reaching movements from field potentials in the absence of spikes. *Journal of neural engineering* **9**, 046006 (2012).

7 Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* **45**, 801-814, doi:10.1016/j.neuron.2005.01.027 (2005).

8 Crammond, D. J. & Kalaska, J. F. Prior information in motor and premotor cortex: activity during the delay period and effect on pre-movement activity. *Journal of neurophysiology* **84**, 986-1005 (2000).

9 Johnson, M., Coltz, J., Hagen, M. & Ebner, T. Visuomotor processing as reflected in the directional discharge of premotor and primary motor cortex neurons. *Journal of neurophysiology* **81**, 875-894 (1999).

10 Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387-400 (2010).

11 Caminiti, R., Johnson, P. B., Galli, C., Ferraina, S. & Burnod, Y. Making arm movements within different parts of space: the premotor and motor cortical representation of a coordinate system for reaching to visual targets. *The Journal of neuroscience* **11**, 1182-1197 (1991).

12 Pesaran, B., Nelson, M. J. & Andersen, R. A. Dorsal premotor neurons encode the relative position of the hand, eye, and goal during reach planning. *Neuron* **51**, 125-134 (2006).

13 Georgopoulos, A., Caminiti, R. & Kalaska, J. Static spatial effects in motor cortex and area 5: quantitative relations in a two-dimensional space. *Experimental Brain Research* **54**, 446-454 (1984).

14 Wang, W., Chan, S. S., Heldman, D. A. & Moran, D. W. Motor cortical representation of position and velocity during reaching. *Journal of neurophysiology* **97**, 4258-4270 (2007).

15 Cos, I., Bélanger, N. & Cisek, P. The influence of predicted arm biomechanics on decision making. *Journal of neurophysiology* **105**, 3022-3033 (2011).

16 Wood, D. K., Chouinard, P. A., Major, A. J. & Goodale, M. A. Sensitivity to biomechanical limitations during postural decision-making depends on the integrity of posterior superior parietal cortex. *Cortex*, doi:http://dx.doi.org/10.1016/j.cortex.2016.07.005.

17 Kim, B. & Basso, M. A. A probabilistic strategy for understanding action selection. *The Journal of Neuroscience* **30**, 2340-2355 (2010).

18 Beck, J. M. *et al.* Probabilistic population codes for Bayesian decision making. *Neuron* **60**, 1142-1152 (2008).

19 Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215-234 (2008).

20 Deneve, S. Bayesian spiking neurons I: inference. *Neural computation* **20**, 91-117 (2008).

21 Thura, D. & Cisek, P. Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* **81**, 1401-1416 (2014).

22 Riehle, A. & Requin, J. Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *Journal of neurophysiology* **61**, 534-549 (1989).

23 Klaes, C., Westendorff, S., Chakrabarti, S. & Gail, A. Choosing goals, not rules: deciding among rule-based action plans. *Neuron* **70**, 536-548 (2011).

24 Bastian, A., Riehle, A., Erlhagen, W. & Schoner, G. Prior information preshapes the population representation of movement direction in motor cortex. *Neuroreport* **9**, 315-319 (1998).

25 Rickert, J., Riehle, A., Aertsen, A., Rotter, S. & Nawrot, M. P. Dynamic encoding of movement direction in motor cortical neurons. *The Journal of Neuroscience* **29**, 13870-13882 (2009).

26 Waldhör, T., Haidinger, G. & Schober, E. Comparison of R2 measures for Poisson regression by simulation. *Journal of Epidemiology and Biostatistics* **3**, 209-215 (1998).

27 Fernandes, H. L., Stevenson, I. H., Phillips, A. N., Segraves, M. A. & Kording, K. P. Saliency and saccade encoding in the frontal eye field during natural scene search. *Cerebral Cortex*, bht179 (2013).

Supplementary Figures

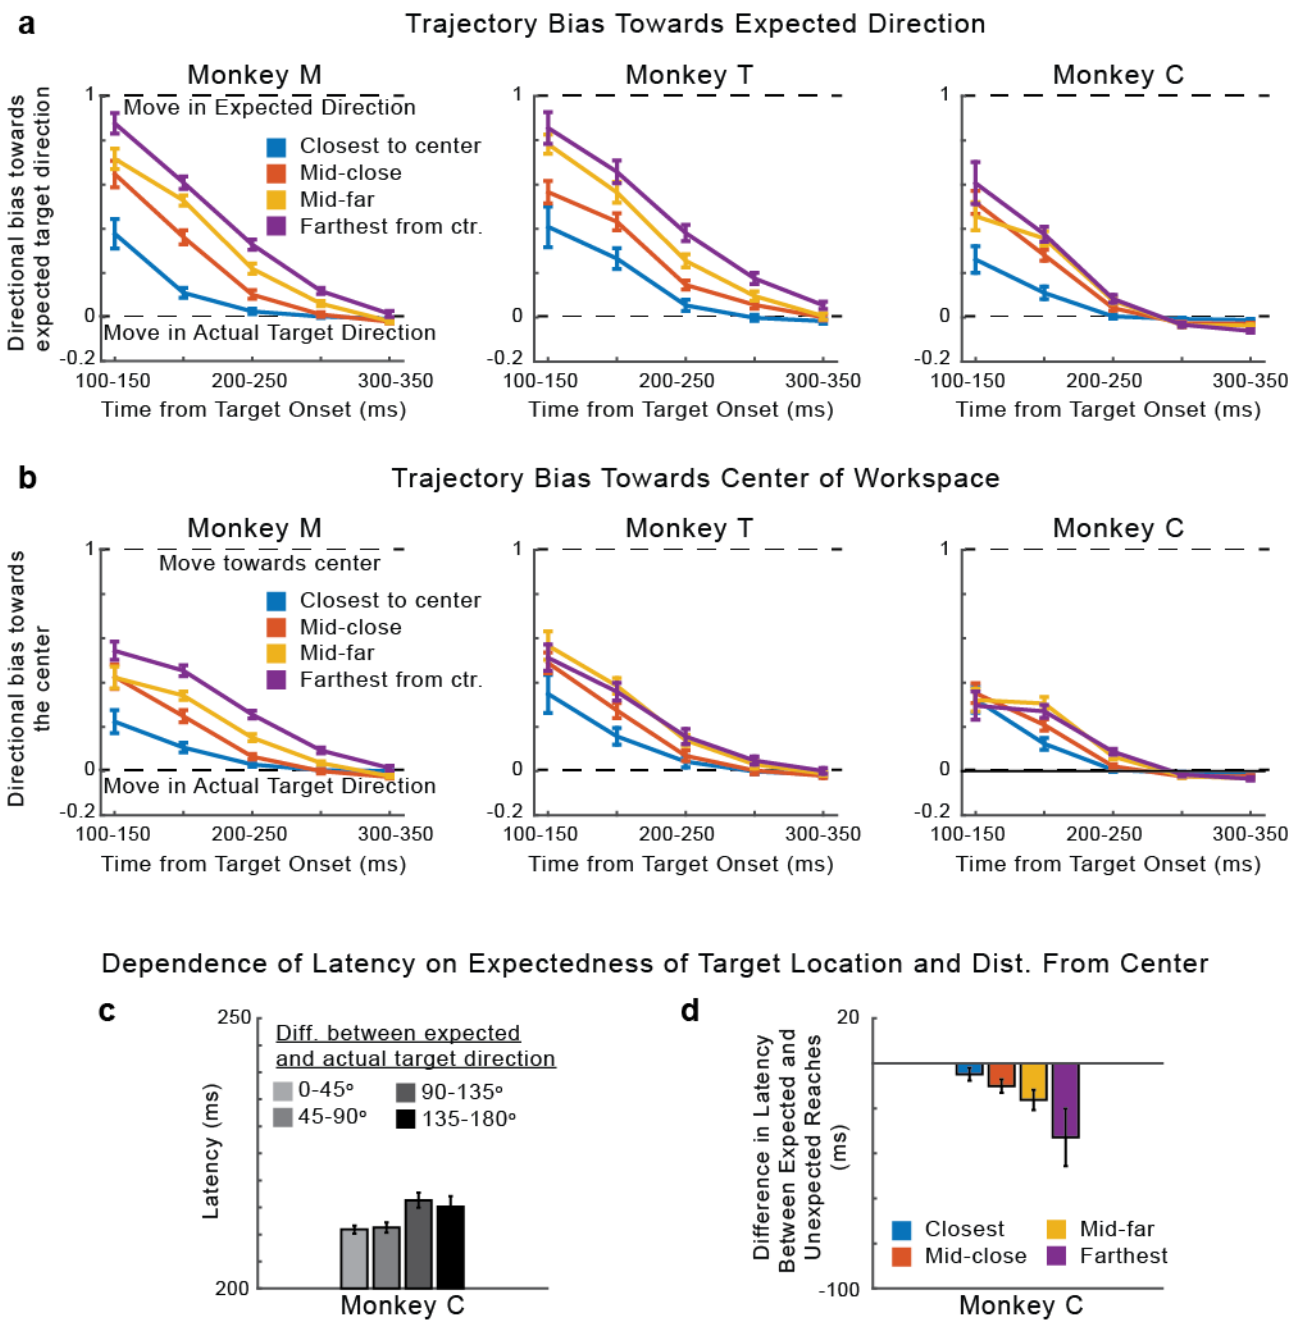

Supplementary Figure 1. Trajectory Biases, and Behavior for Monkey C
We compare **(a)** the biases in the reach direction towards the expected direction of the reach given the statistics of target presentation (identical to Fig. 2a) versus **(b)** the biases in the reach direction towards the center of the workspace. **(c&d)** Same as Fig. 2c&d, for Monkey C.

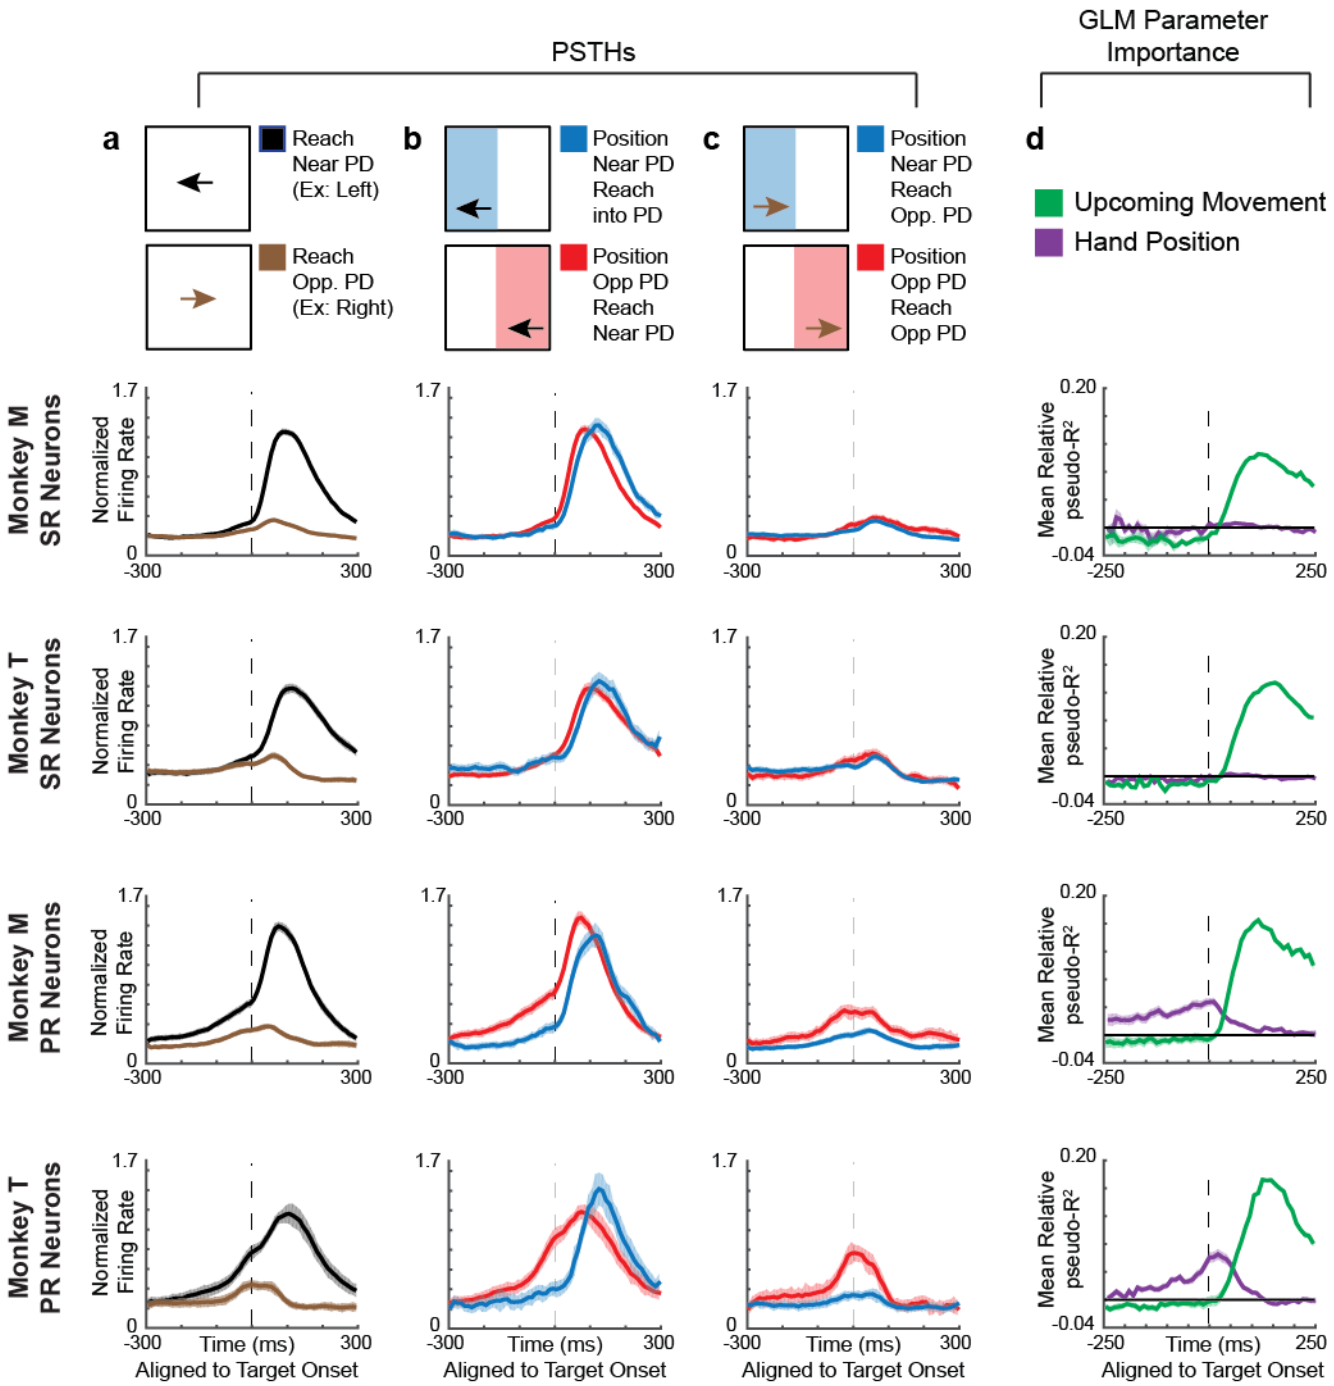

Supplementary Figure 2: PMd PSTHs and GLM results, for individual monkeys
Peristimulus time histograms (PSTHs) and GLM results for PMd neurons, for individual monkeys. All columns are the same as in Fig. 3. **First and Second PSTH Rows**: Normalized averages of selected-response (SR) neurons from Monkey M and Monkey T, respectively. **Third and Bottom Rows**: Normalized averages of potential-response (PR) neurons from Monkey M and Monkey T, respectively.

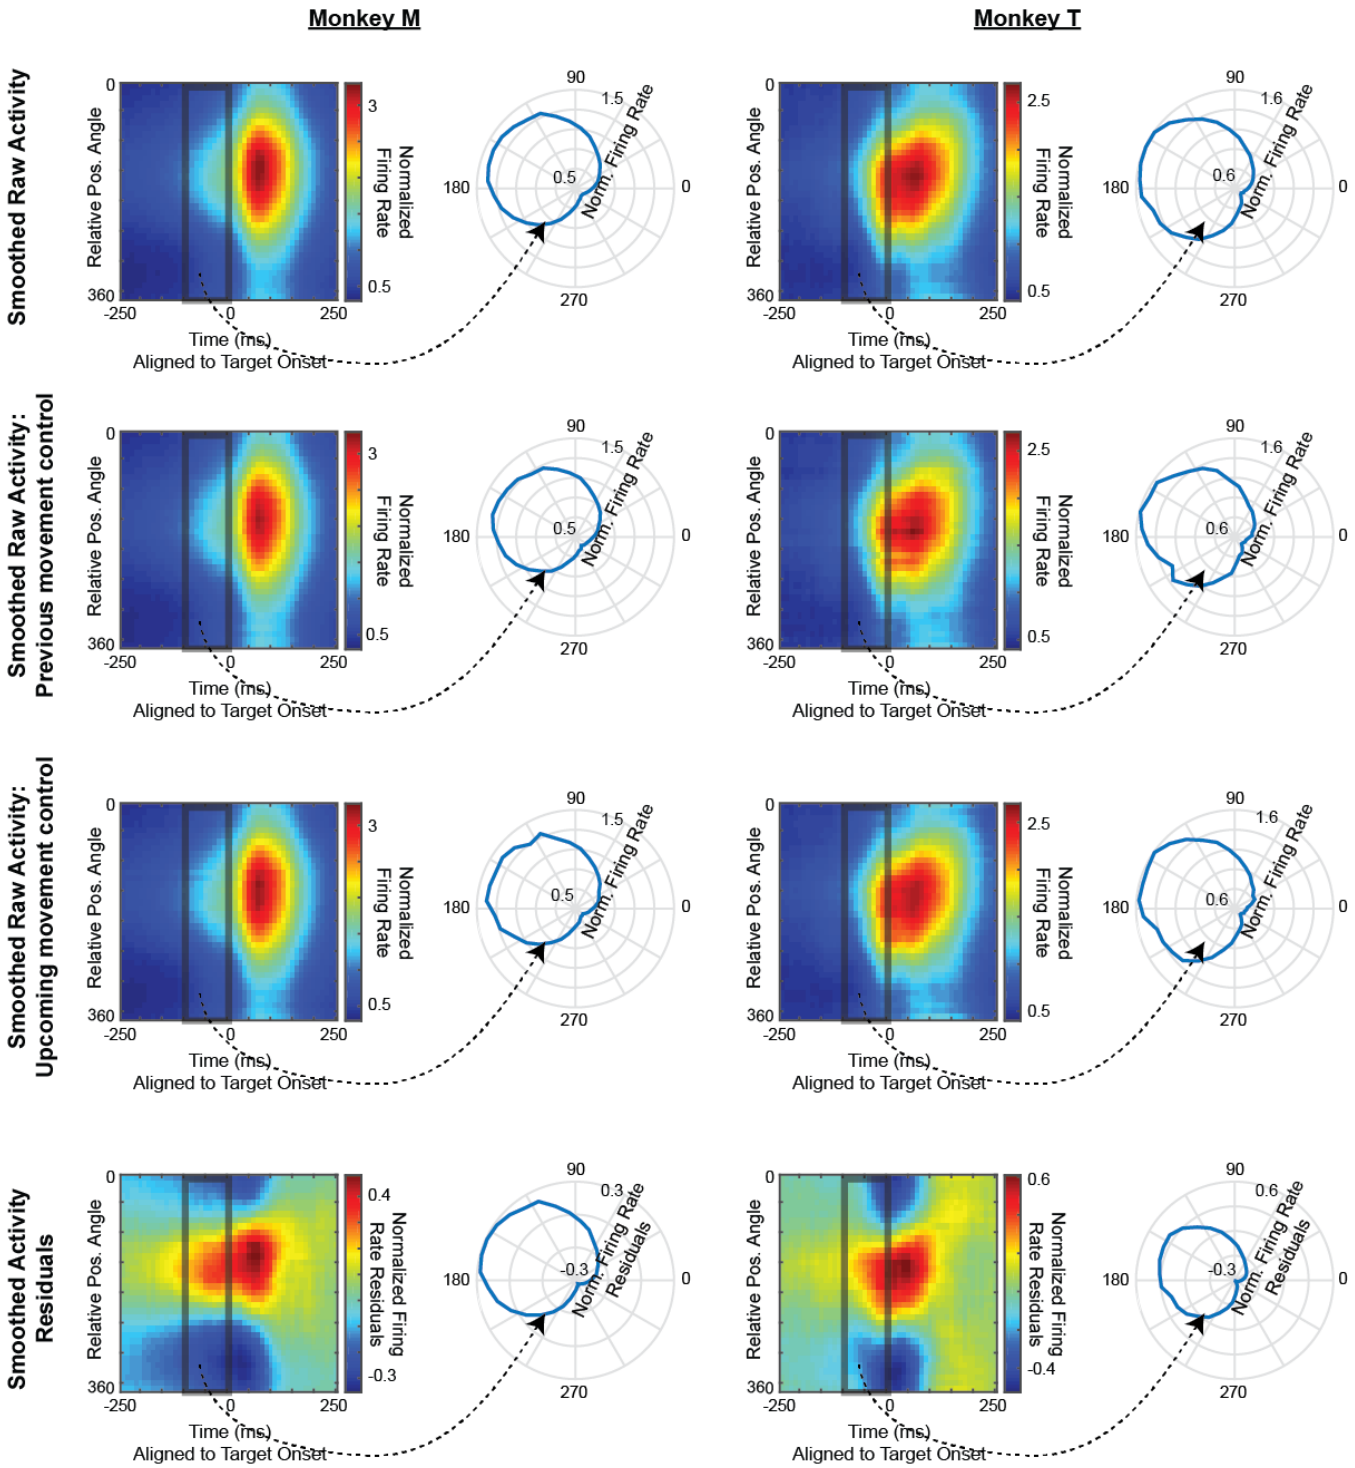

Supplementary Figure 3: PMd population activity represents the distribution of upcoming movements–accumulated across reaches: individual monkeys and controls.
As in Fig. 4a, all heat maps show normalized average activity over time as a function of relative angular position (the angular hand position relative to neurons' PDs). To the right of each heat map, normalized activity in the 100 ms prior to target onset is shown on a polar plot, as a function of the relative angular position. The **left** column is for monkey M, and the **right** column is for monkey T. Each row calculates activity in a different way. **Top row:** Normalized smoothed firing rate (as in Fig. 4a). **Second row:** Normalized smoothed firing rate, controlling for the correlation between the previous movement direction and upcoming movement direction. One potential concern with the main results is that the neural activity might be related to the previous movement itself, rather than the probability distribution of upcoming movements. It is possible we were only capturing a correlation between the previous and upcoming

movements, which tended to be oppositely directed. As a control, we examined neural activity in the infrequent cases when pairs of sequential reaches were in similar directions (less than 90° away from each other). **Third row:** Normalized smoothed firing rate, controlling for the correlation between the angular hand position and upcoming movement direction. We resampled reaches to create a distribution of reach directions relative to angular hand position that were centered on 180° (rather than 156°). This plot ensures that our main results were not simply caused by a correlation with the true upcoming movement (in which case the activity after resampling would become centered on 180°), but rather reflected the distribution of upcoming movements (in which case the activity after resampling would remain centered at 156°). **Bottom row:** Normalized smoothed firing rate residuals. Activity due to the upcoming movement is subtracted from the total activity, leaving residuals (that still depend on position). We normalize these residuals (divide by the mean firing rate of each neuron).

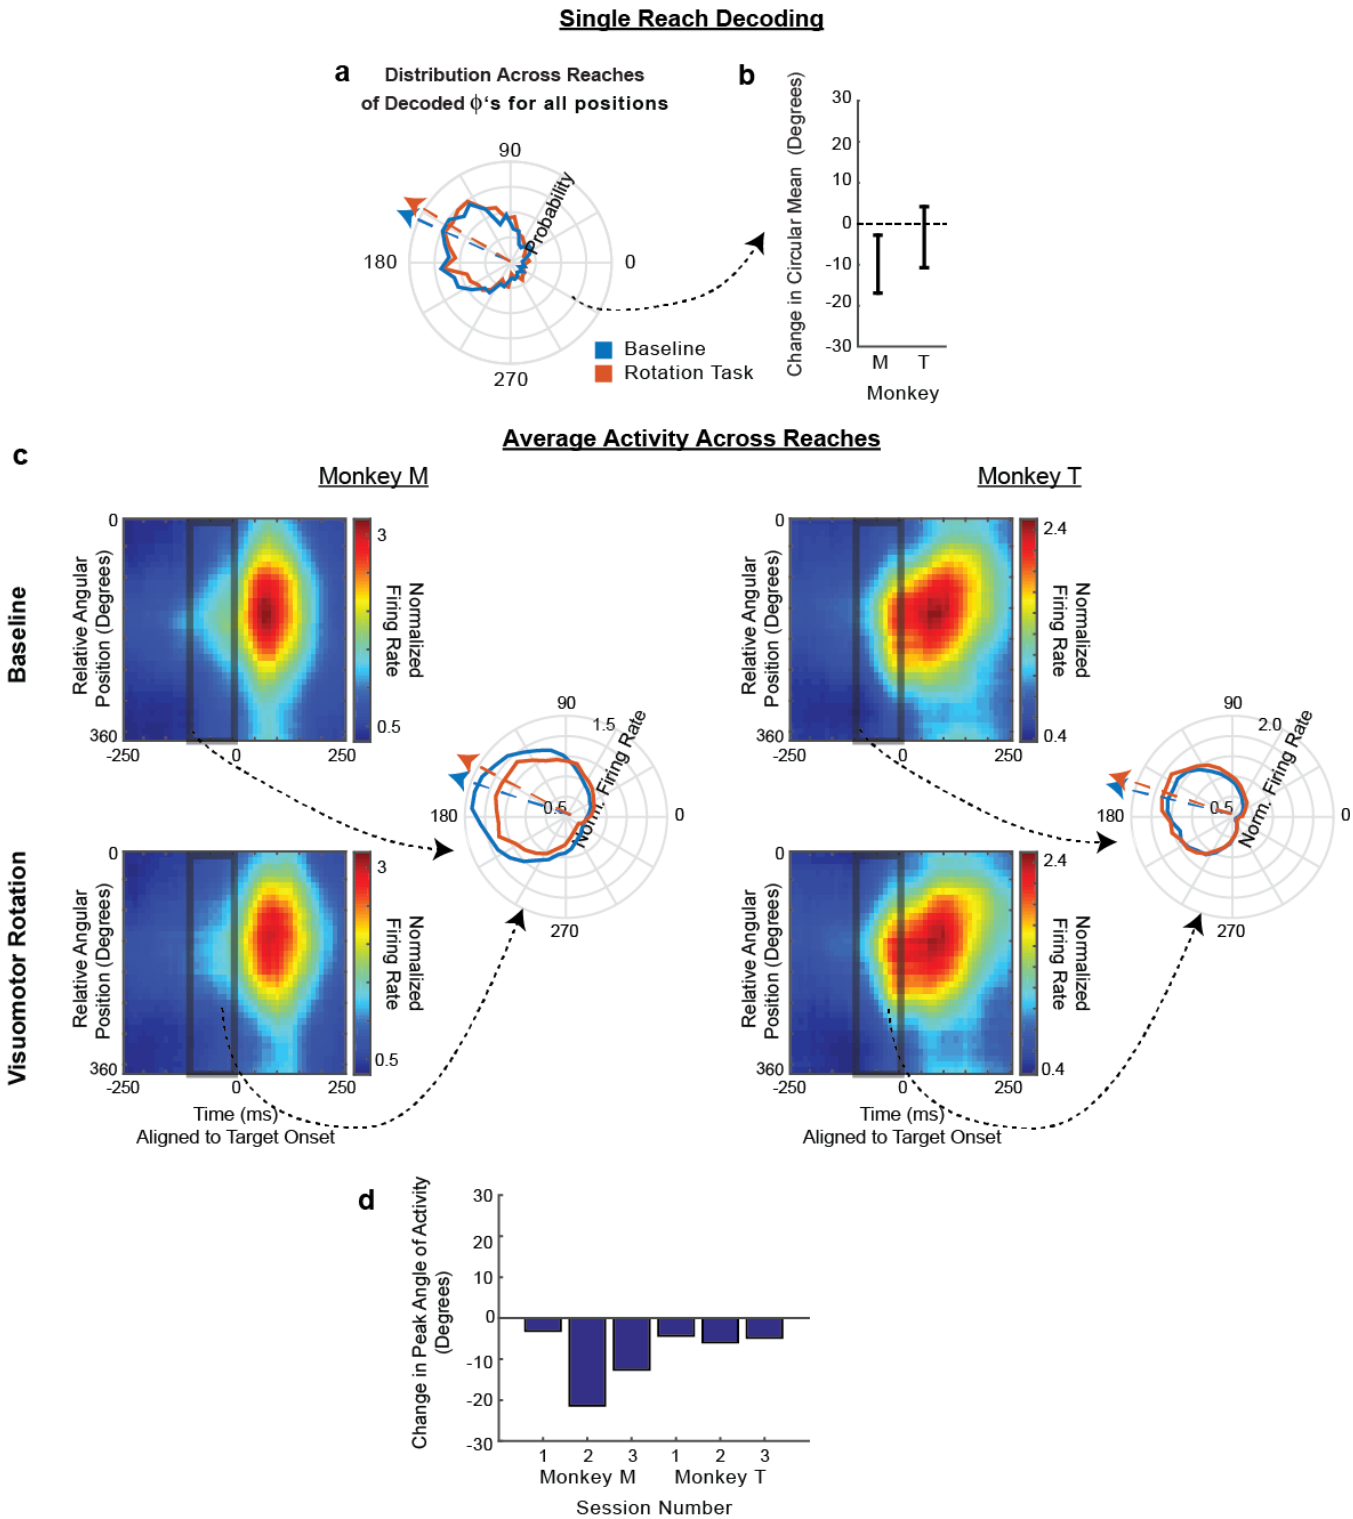

Supplementary Figure 4: Visuomotor Rotation Control Task

(a) The distribution of pre-target decoded reach directions relative to the hand's angular position (decoded Φs) for all positions. Arrows point toward the circular means of the distributions. **(b)** The difference between the circular mean of the distributions of decoded Φs in panel a, between the baseline and rotation tasks (rotation minus baseline). Error bars represent 95% confidence intervals from bootstrapping. **(c)** For each monkey, in both the baseline task and visuomotor rotation task, we display on the **left**: The average normalized firing rate of all PR neurons, over time, as a function of relative angular hand position. For each neuron, the relative angular position is the preferred direction of the neuron minus the angular hand position. Activity is normalized and averaged across all PR neurons. For each monkey, we display on the **right:** The average normalized firing rate in the 100 ms prior to target onset, plotted as a

function of the relative angular position (baseline in blue; rotation task in orange). Arrows point toward the peak angles of activity. **(d)** For each session, the difference between the angle corresponding to peak activity in panel c, between the baseline and rotation tasks (rotation minus baseline). In all panels, results from the VR task used the second 2/3 of trials.

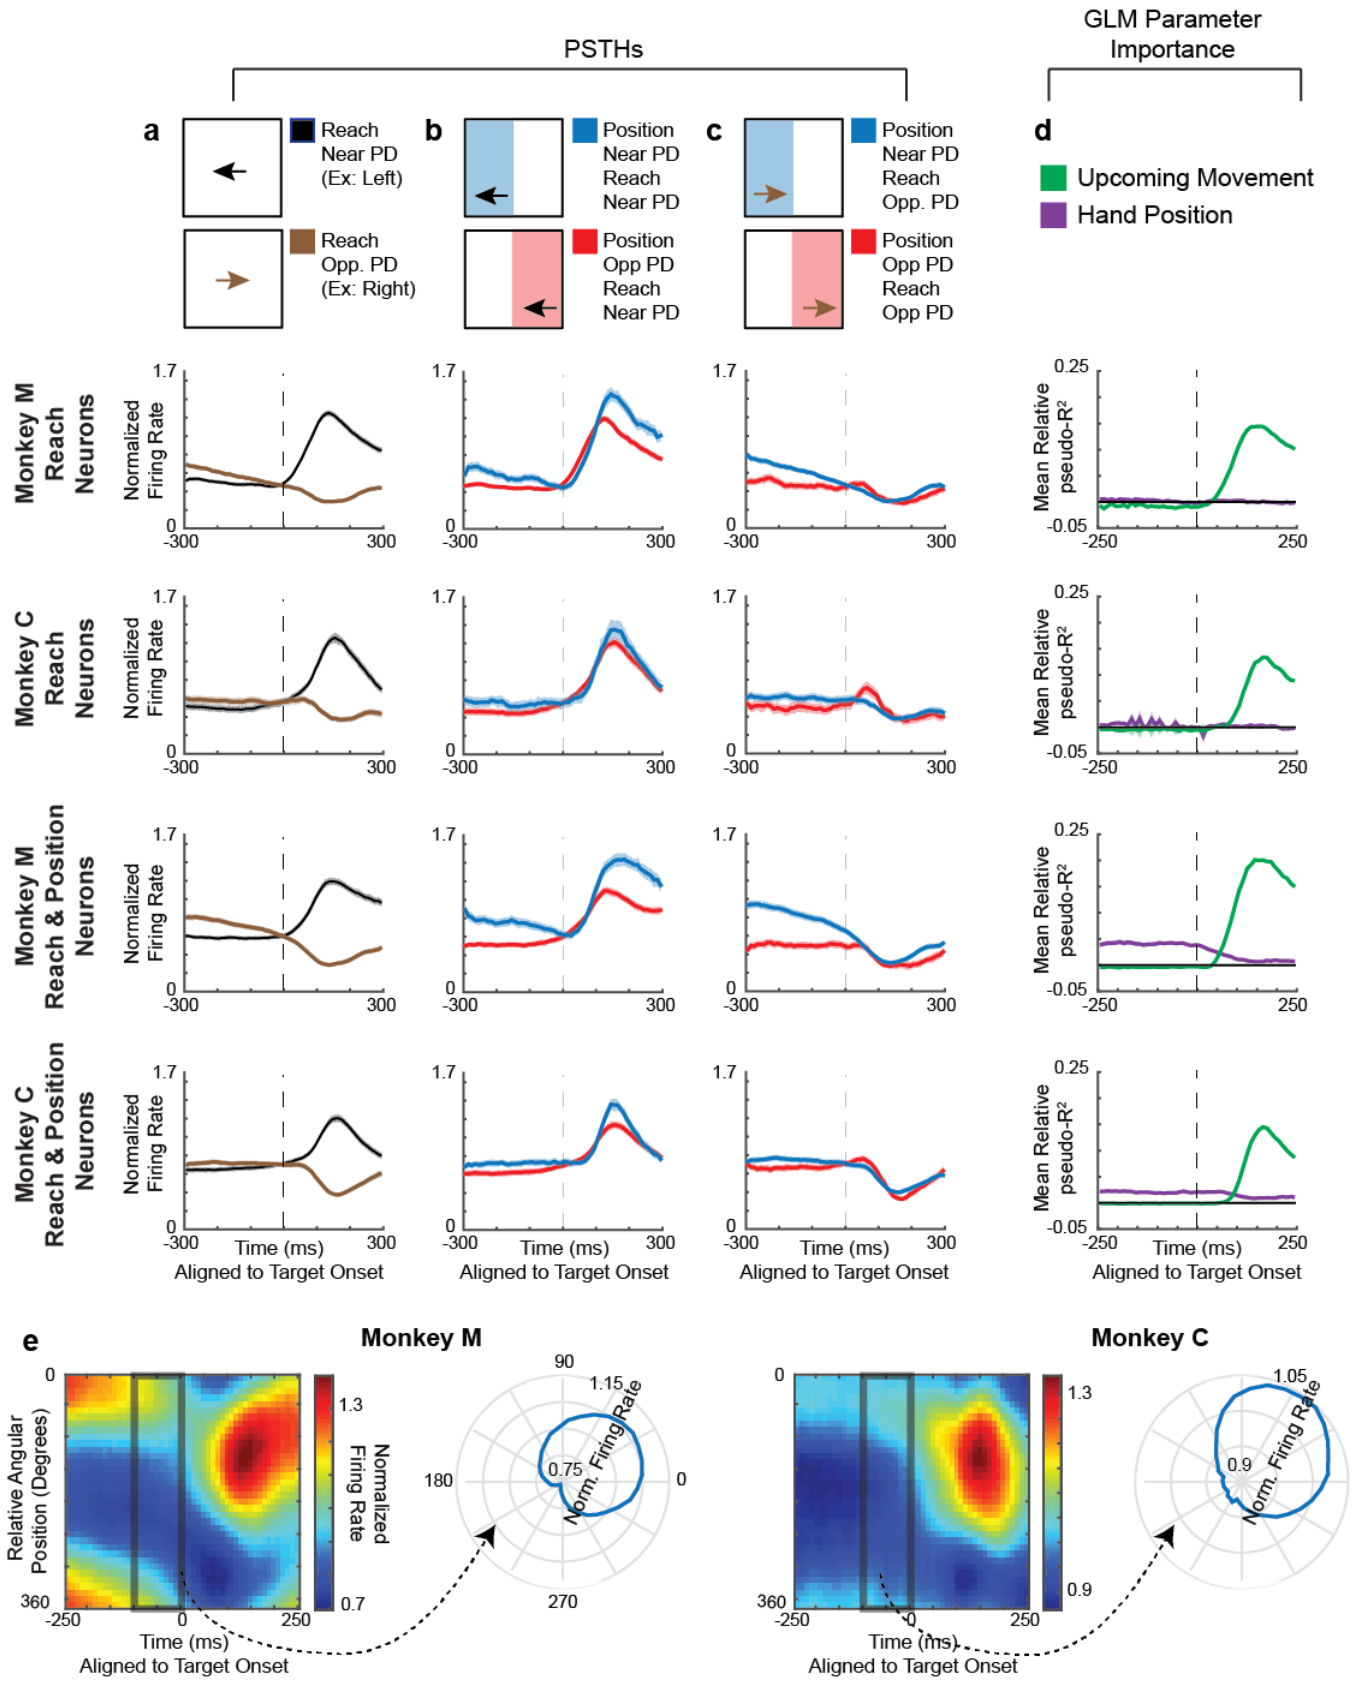

Supplementary Figure 5: M1 does not reflect the probability of upcoming movements, for individual monkeys (a-d) PSTHs and GLM results for M1 neurons. Columns have the same schematics as Fig. 3. **First and Second Rows of PSTHs**: Normalized averages of reach neurons, defined as those neurons significant for movement during the late

period, but not position in the early period of the GLM. **Third and Fourth Rows**: Normalized averages of reach & position neurons, defined as those neurons significant for movement during the late period, and position in the early period of the GLM. Note that we did not use the same "SR/PR" nomenclature as PMd, because there was not evidence in the PSTHs of M1 neurons that position was used to represent potential upcoming movements. **(e)** Same schematic as Fig. 4a, but for M1 neurons. **Left**: The normalized average firing rate, as a function of time and relative angular hand position. Activity is averaged across all reach & position neurons. **Right:** The normalized average firing rate in the 100 ms prior to target onset, plotted as a function of the relative angular position.