# Flexible neural coding in sensory, parietal, and frontal cortices during goal-directed virtual navigation

Jean-Paul Noel[1, *], Edoardo Balzani[1, *], Eric Avila[1, *], Kaushik Lakshminarasimhan[1, 2], Stefania Bruni[1], Panos Alefantis[1], Cristina Savin[1, #], Dora E. Angelaki[1, #]

[1] Center for Neural Science, New York University, New York City, NY, USA
[2] Center for Theoretical Neuroscience, Columbia University, New York City, NY, USA
* Equal Contribution as co-first authors
# Equal Contribution as co-senior authors

Correspondence:
Dr. Dora E. Angelaki
Da93@nyu.edu

## Abstract

We do not understand how neural nodes operate within the recurrent action-perception loops that characterize naturalistic self-environment interactions, nor how brain networks reconfigure during changing computational demands. Here, we record local field potentials (LFPs) and spiking activity simultaneously from the dorsomedial superior temporal area (MSTd), parietal area 7a, and dorsolateral prefrontal cortex (dlPFC) as monkeys navigate in virtual reality to "catch fireflies". This task requires animals to actively sample from a closed-loop visual environment while concurrently computing latent variables: the evolving distance and angle to a memorized firefly. We observed mixed selectivity in all areas, with even a traditionally sensory area (MSTd) tracking latent variables. Strikingly, global encoding profiles and unit-to-unit coupling suggested a functional subnetwork between MSTd and dlPFC, and not between these are 7a, as anatomy would suggest. When sensory evidence was rendered scarce, lateral connectivity through neuron-to-neuron coupling within MSTd strengthened but its pattern remains fixed, while neuronal coupling adaptively remapped within 7a and dlPFC. The larger the remapping in 7a/dlPFC and the greater the stability within MSTd, the less was behavior impacted by loss of sensory evidence. These results highlight the distributed nature of neural coding during closed-loop action-perception naturalistic behaviors and suggest internal models may be housed in the pattern of fine-grain lateral connectivity within parietal and frontal cortices.

**Introduction**

Despite the closed-loop interplay between action and perception that exists in the real world, our understanding of sensory encoding and the neural architectures supporting goal-directed behavior is largely derived from artificial tasks segregating perception from action. This limited purview of the traditional approach has hindered our ability to understand neural coding for naturalistic and dynamic behaviors (see Gomez-Marin et al., 2014; Cisek & Kalaska, 2010, Pitkow & Angelaki, 2017, Yoo et al., 2021, for similar arguments).

First, severing the natural closed loop between action and perception disrupts the natural timing between sensory events and internal neural dynamics. In natural vision, for example, eye movements dictate the content, relative resolution, and timing of visual input. Recent work from active sensing (Schroeder et al., 2010; Yang et al., 2016) has shown that neural excitability in primary visual cortex (Barczak et al., 2019) and the anterior nuclei of the thalamus (Leszczynski et al., 2020) are enhanced at saccade offset – precisely when new observations are made. This enhancement is mediated by phase-reset of neural oscillations (Lakatos et al., 2008; Rajkai et al., 2008) caused by the shifting gaze. In turn, it is likely that physical movements of the eyes during sensory processing – an aspect absent in most laboratory, binary decision-making tasks – may not only enhance low-level encoding, but also favor certain channels of inter-area communication via local field potential (LFP) phase alignment or other coupling mechanisms (e.g., Jutras et al., 2013).

Second, the emphasis on tasks with poor dynamics, together with a technology-limited focus on studying one area at a time, has limited our ability to explore how within- and cross-area communication enables flexible behavior. This has possibly led to a degeneracy in the number of functions ascribed to each neural node. For instance, redundancy has notoriously been observed between parietal and frontal cortices with both areas showing similar properties during visual feature categorization (Swaminathan & Freedman, 2012), target selection (Suzuki & Gotlieb, 2013), visuo-spatial memory (Chafee & Goldman-Rakic, 1998) and working memory training (Olesen et al., 2003), among others (see Katsuki & Constantinidis, 2012, for a review). While this redundancy is certainly an adaptive feature (Moreno-Bote et al., 2014; Driscoll et al., 2017), the joint characterization of sensory, parietal, and frontal areas during tasks requiring a flow of cognitive operations typical of daily life (e.g., sensing, memorizing, acting) may offer the possibility to study how these regions differ, and how they interact.

Lastly, not only are single cognitive operations typically studied in isolation and in distinct and artificial temporal intervals, but they are also seldomly examined under different computational demands, e.g., from requiring the usage of incoming sensory evidence, vs. relying on pre-existing internal models. This focus on studying behaviors from a single computation perspective and requiring a fixed strategy is likely because asking over-trained animals to use diverse computational approaches in solving artificial tasks would require additional training on the second approach – a problem of over-fitting. This challenge is abated in naturalistic settings where fast across task generalization is possible (Noel et al., 2021).

To tackle these limitations, we have developed a task requiring monkeys to catch memorized fireflies in virtual reality (Lakshminarasimhan et al., 2020). This goal-directed virtual navigation task addresses many of the limitations of the traditional approach, while being strongly rooted in established and well-understood elements: motion detection (Newsome & Pare, 1988), optic flow processing (Britten, 2008), and navigation (Kravitz et al., 2011, Ekstrom et al., 2018). Animals first detect a briefly flashed target, much like the blinking of a firefly. Then, they use a joystick controlling linear and angular velocity to navigate via path integration to the location of this target from memory (**Figure 1A**). Importantly, the observers' eyes are free to sample from their sensory surroundings, and trials last on the order of 3-4 seconds. In turn, this task requires integration of sensory evidence over protracted time periods, allows for active sensing, and engages the real-world loop between observations, beliefs, actions, and states (**Figure 1A**). Critically, the task-relevant variable (i.e., the relative position of the self to the target) is latent. It must be dynamically computed, and thus offers a window into the mechanisms of internal computation. Furthermore, by decreasing the density of ground-plane optic flow elements, we can shift the nature of the computation being performed; from relying on incoming sensory evidence to relying on an internal model of the physics of the environment and control dynamics.

Here, we leverage this "firefly task" to simultaneously characterize the encoding profiles of sensory (dorsomedial superior temporal area, MSTd), parietal (area 7a), and frontal (dorsolateral prefrontal cortex, dlPFC) areas during

naturalistic, closed-loop visual behavior. We record local field potentials (LFPs) and single unit activity from this series of reciprocally connected areas (Andersen et al., 1990; Rozzi et al., 2006) involved in the processing of optic flow for self-motion, sensorimotor transformations, and belief formation/working memory (**Figure S1**). We observe that all areas probed – including MSTd, an area classically considered sensory – encode latent task variables. Surprisingly, global encoding profiles and unit-to-unit couplings suggest a functional subnetwork composed of dlPFC and MSTd (and not area 7a, as would be predicted from anatomy). When sensory information is reduced, changing the nature of the computation from sensory-driven to internal-model driven, area 7a and dlPFC dynamically remap their unit-to-unit connectivity, while MSTd coupling strengthens but it's fine-grain pattern remains unchanged. Strikingly, the degree to which MSTd remains stable, and the degree to which area 7a and dlPFC reconfigure correlates with changes in behavior. Overall, these findings highlight heterogenous subnetworks (not always predicted by anatomy) that adaptively reconfigure during changing computational demands to support goal-directed behavior in a dynamic environment.
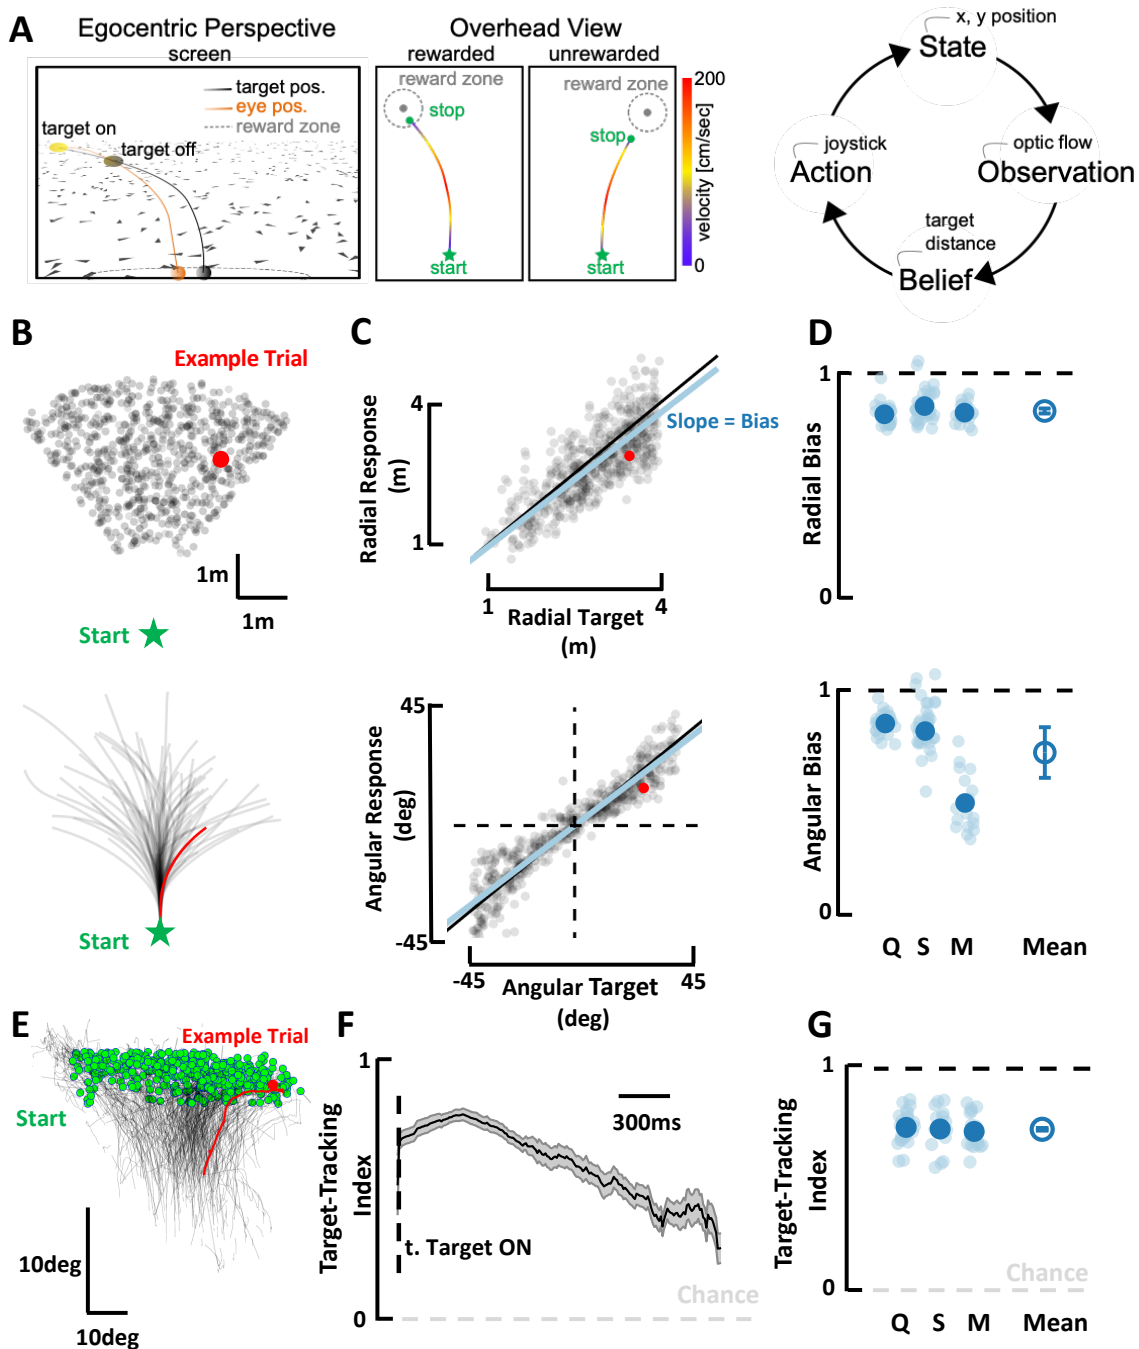
## Results

### Monkeys navigate to remembered targets employing natural strategies

Targets ("fireflies") were displayed for 300ms, on a ground plane composed of triangular optic flow elements that appeared transiently (250ms) with random orientations. The standard density of these optic flow elements was 5 elements/m$^2$. A single target was presented per trial. The firefly targets were uniformly distributed (across trials) within a radial range of 1 to 4m, and an angular disparity spanning from -40° (leftward) to 40° (rightward) of visual angle (**Figure 1B**). Performance feedback was provided at the end of each trial in the form of juice reward for correctly stopping within a reward boundary (see *Methods* for further detail).

Visualizing example trials shows that monkeys followed curvilinear trajectories and stopped near the latent targets (**Figure 1B**). To quantify their performance, we expressed the monkeys' trajectory endpoints and firefly locations in polar coordinates, with an eccentricity from straight-ahead ($\theta$) and a radial distance ($r$). **Figure 1C** shows radial (top; $r$ vs. $\tilde{r}$ slope = 0.95; R$^2$ = 0.55) and angular (bottom; $\theta$ vs. $\tilde{\theta}$ slope = 0.94; R$^2$ = 0.78) responses as a function of target location for an example session (omitting 18% of trials the animals aborted early or never stopped, see *Methods*). As shown by this example session, and confirmed across all three animals and all sessions (Monkey Q, 27 sessions; S, 38 sessions; M, 17 sessions; **Figure 1D**), monkeys were on average accurate, but tended to undershoot targets ($r$ vs. $\tilde{r}$ slope, Mean=0.842, 95%CI=[0.830, 0.854]; $\theta$ vs. $\tilde{\theta}$ slope, M=0.772, 95%CI=[0.732, 0.812]). Undershooting targets is, in fact, the optimal strategy given that uncertainties over one's own and the firefly's position scales with time and distance (see Lakshminarasimhan et al., 2018).

Over the course of individual trajectories toward the target, the monkeys' eye position moves downward and becomes less eccentric (**Figure 1E**), as if tracking the (invisible) target progressively becoming nearer and aligned with straight ahead (Lakshminarasimhan et al., 2020; also see Ilg & Thier, 1999, 2003). This behavior was quantified by deriving predictions for the binocular position of the observer's eyes, assuming the monkeys maintained fixation at the center of the target throughout the trial. Then, we expressed a target tracking index as the square root of the fraction of variance in the observed eye position that was explained by the prediction (see Lakshminarasimhan et al., 2020, for further detail). An index of 1 implies that the subject consistently looked at the center of the firefly, while a value of zero denotes a complete lack of correspondence between target and gaze locations. The target-tracking index was high at firefly onset, highest at firefly offset, and remained above chance throughout the course of the trial (**Figure 1F** shows an example session). Across all animals and sessions, the target-tracking index averaged 0.73 (95% CI = [0.72, 0.74]) during the first 1s of each trial, suggesting that all animals naturally fixated on the firefly, and attempted to maintain gaze on it even after the target had disappeared (**Figure 1G**).

Altogether, the animals can make use of the degrees of freedom afforded by the joystick to accumulate evidence over prolonged periods of time within a closed loop setting, and seemingly sample from their environment in an intelligent manner.

**Figure 1. Task and behavioral results A.** *Behavioral task. Left: Egocentric perspective of the monkey during the firefly task, emphasizing the ground plane optic flow elements (flashing triangles), the fact that the target disappears (target off), and eye-positions (following target position perfectly when the target is on, and then continuing on the correct path but deviating with time). Center: Overhead view of the monkey, starting location (green star), accelerating and progressively re-orienting to face the firefly, before de-accelerating and stopping at (left; rewarded) or near (right: unrewarded) the location of the now invisible target. Right: This task involves making observation of the sensory environment (composed of optic flow), using these observations to generate a dynamic belief (of the relative distance to target), and producing motor commands based on the current belief, which in turn updates the state of the environment (right). **B.** Top: spatial distribution of target positions across trials. Bottom: monkey trajectories. The red dot and trajectory are an example trial, maintained throughout the figure. **C.** Example session. Radial (top) and angular (bottom) endpoint (y-axis) as a function of target (x-axis). Gray dots are individual trials, black line is unity, and blue line is the regression between response and target (slope of 1 indicates no bias). **D.** Radial (top) and angular (bottom) bias (slope, =1 means no bias, >1 means overshooting, and <1 means undershooting) for every session (transparent blue circles) and monkey (x-axis, dark blue circles are average for each monkey, Q, S, and M). Rightmost circle is the average across monkeys and error bars represent +/- 1 S.E.M. Overall, monkeys are fairly accurate but undershoot targets, both radially and in eccentricity. **E.** Eye trajectories (green = eye position at firefly offset) for an example session, projected onto a 2-dimensional 'screen'. Eyes start in the upper field and gradually converge in the lower center (where the firefly ought to be when they stop). **F.** Target-tracking index (variance in eye position explained by prediction of fixating on firefly) for an example session as a function of time since firefly onset. **G.** Average target-tracking index within 1s for all sessions (light blue) and monkeys (dark blue) showing the monkeys attempt to track the invisible target.*

## Mixed selectivity across sensory, parietal, and frontal cortices

We recorded single-unit activity from MSTd (231 units), area 7a (3200 units), and dlPFC (823 units) in a total of 82 sessions while monkeys performed the firefly task (**Figure 2A** and **Figure S1**). Visual inspection of raw spike times suggested a putative coarse organization and area-specific features (e.g., high firing rates seemingly preceding stopping time in area 7a, black dots, **Figure 2B**), but also displayed fine detail differences across trials and areas. To quantitatively and agnostically account for single unit encoding properties during this naturalistic, closed-loop task, we fit spike trains to a Poisson Generalized Additive model our group has previously developed (P-GAM; see Balzani et al., 2020). Beyond capturing arbitrary non-linearities and handling collinear predictors (Dormann et al., 2013), this particular encoding model has the strong advantage of inferring marginal confidence bounds for the contribution of each feature to neural responses (see Balzani et al., 2020 for details). This property allows us to robustly identify a minimal set of task features that each neuron is responsive to, while circumventing the need for computationally demanding model selection (particularly given the number and time-varying nature of the variables in this naturalistic task). Indeed, using model recovery we demonstrate that the P-GAM will capture the simplest possible interpretation of the neural responses in the statistical regime of our data, i.e., even when neurons are weakly firing and predictors are correlated (see **Figure S2**).

In addition to discrete task events (e.g., time of target onset, movement onset and offset), continuous sensorimotor variables (e.g., linear and angular velocity and acceleration) and latent parameters (e.g., distance from origin and to target), we included in the encoding model elements of brain dynamics. These included the phase of LFP in different frequency bands (theta: 4-8 Hz; alpha: 8-12 Hz; beta: 12-30 Hz), and the coupling within (i.e., spike-history, 36 ms wide temporal filter) and between units, both within (36 ms wide temporal filter) and across cortical areas (600 ms wide temporal filters, **Figure 2C,** see *Methods*). In total the model had 17 (analog or digital) task inputs (**Figure 2C**), in addition to hundreds of potential lateral and across region connections (i.e., coupling filters). The fitted model accounts well for observed spiking activity (see raw and model reconstructed spiking activity of simultaneously recorded units in **Figure 2D**).

**Figure 2E** shows the tuning of an example MSTd (green), 7a (blue), and dlPFC (red) neuron. The striking feature of these example neurons is the fact that all units coded for a multitude of variables. For instance, each of these example neurons responded to the time of firefly and movement onset. The 7a neuron showed a strong dependency not only to sensorimotor and latent variables, but also to LFP phases (in all frequency ranges). Finally, and perhaps most surprisingly, the MSTd neuron encoded not only sensorimotor and eye-position variables, but also latent ones, such as the angular distance to the (invisible) target. This mixed selectivity was true not only of the example neurons, but of each population as a whole (**Figure 2F** shows 200 neurons from each area). In the *Supplementary Materials* (**Figures S3-S6**) we depict all the different types (i.e., high-dimensional clusters) of tuning functions that exist for the different task-variables in MSTd, area 7a, and dlPFC, thus creating what is akin to a set of 'words' composing the 'dictionary' available to each area in encoding task-relevant features (see *Methods* for further detail).

Interestingly, while MSTd, area 7a, and dlPFC showed mixed selectivity within this naturalistic path integration task, the likelihood of being tuned to a particular variable, or a set of variables, was not random. As one would expect from more traditional, open-loop tasks, neurons in area 7a were tuned to a greater number of sensorimotor variables (Kravitz et al., 2011), while units in dlPFC were tuned to a greater number of latent task-variables (Duncan, 2001), in this case the radial and angular distance from origin and to the target (**Figure 2G**). Spiking activity in area 7a showed a marked dependence on its local LFP phase (79% of neurons being conjunctively tuned to theta, alpha, and beta phases), a feature that was much less frequent in MSTd and dlPFC. **Figure 2H** confirms the above-mentioned observations, now depicted as the fraction of neurons tuned to each variable. In summary, encoding profiles during this closed action-perception loop largely followed general trends predicted in traditional open loop tasks, while also showing some stark novel features; the encoding of latent variables even in a traditionally considered sensory area (i.e., MSTd), and the very prominent dependence of spiking activity in area 7a (and not MSTd or dlPFC) on ongoing LFP phases.
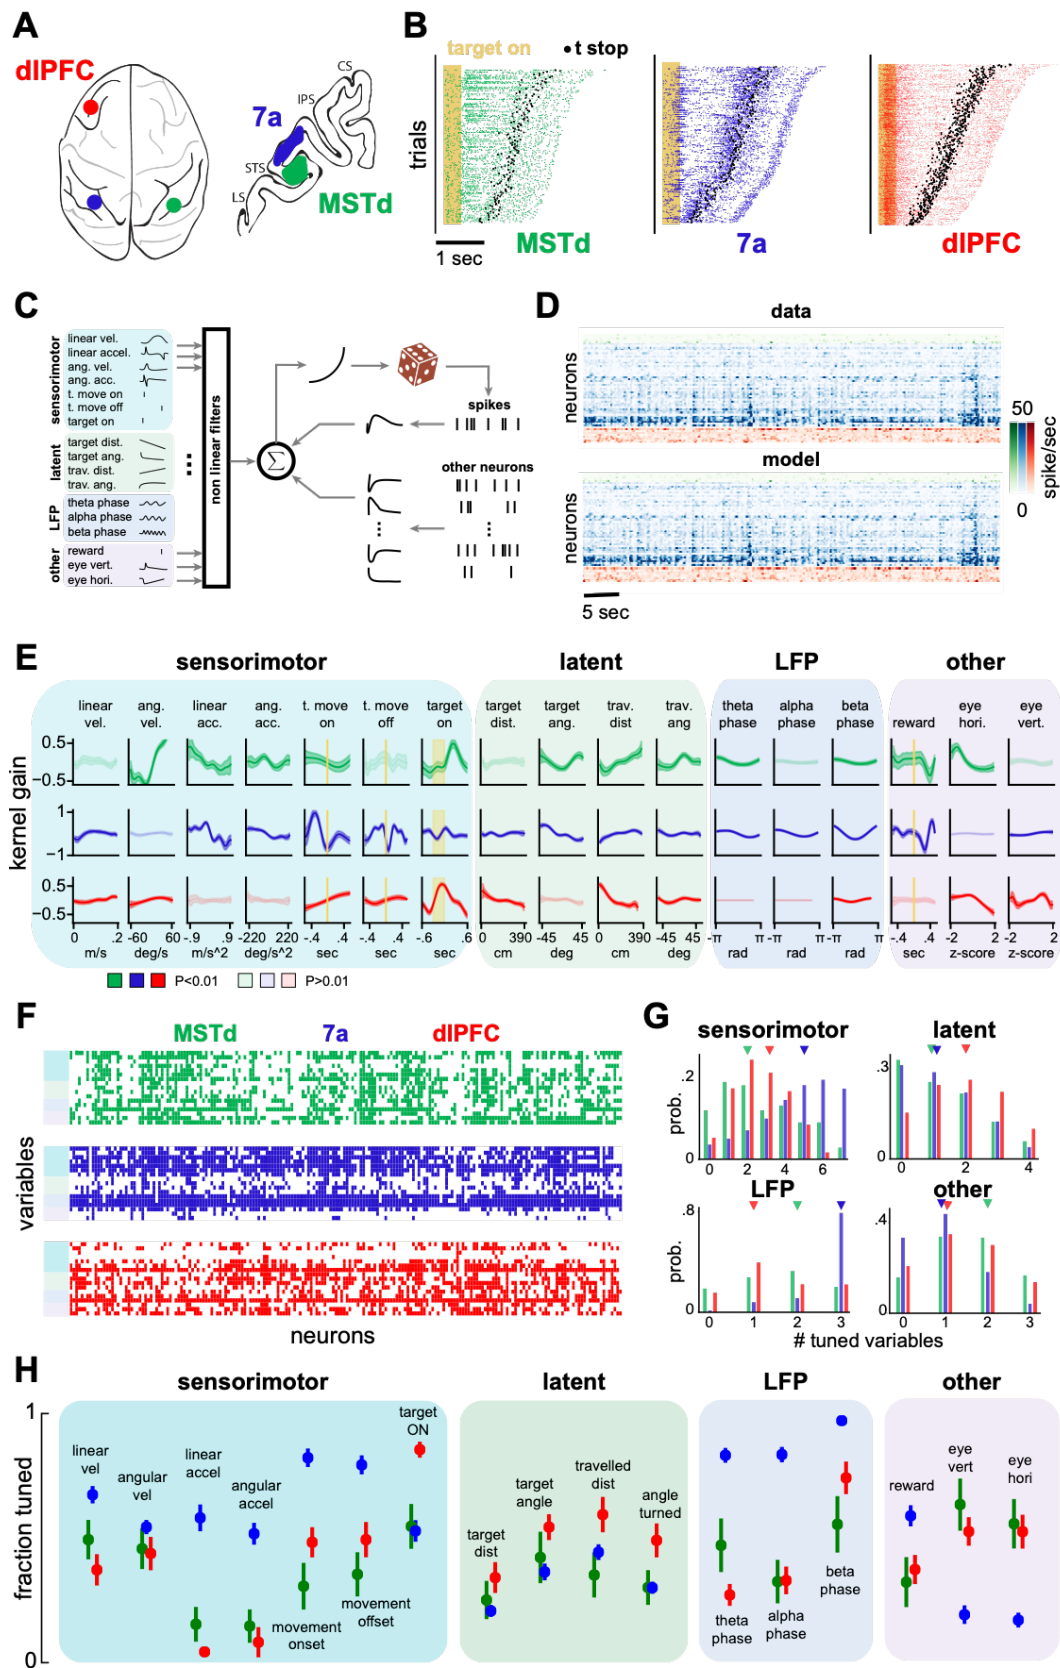
***Figure 2. MSTd, area 7a, and dlPFC encode a heterogeneous mixture of task variables. A.*** *Schematic of brain areas recorded.* ***B.*** *Raster plots of spiking activity in MSTd (green), 7a (blue) and dlPFC (red). Shaded yellow area represents the time of the target being visible, and black dots represent the timing of movement offset. Trials are sorted from shortest (bottom) to longest (top) duration of the trial.* ***C.*** *Schematic of the generalized additive model (GAM) used to fit spike trains. The task and neural variables used as input to the model were: linear and angular velocity and acceleration, time of movement onset and offset, time of firefly onset, distance and angle from origin and to target, time of reward, vertical and horizontal position of the eyes, and ongoing phase of LFP at theta, alpha, and beta bands.* ***D.*** *Top: random snippet of spiking activity of simultaneously recorded neurons (green = MSTd; blue = 7a; red = dlPFC). Bottom:*

*corresponding cross-validated prediction reconstructed from the P-GAM. The average cross-validated pseudo-R2 was 0.072 (see Cameron et al., 1997). The direct comparison of the goodness-of-fit afforded by this method and those in other reports is difficulted by differences in neural areas and tasks, but see Balzani et al., (2020) for evidence of the superiority of the P-GAM with respect to traditional GLMs in synthetic data. Similarly, Benjamin et al., (2018) recently contrasted GLMs and techniques from modern machine learning in fitting spike trains across primary motor (M1), primary somatosensory (S1), and hippocampus. Even though neural responses in M1 and S1 are arguably simpler than those in MSTd, 7a, and dlPFC, Benjamin et al. (2018) report cross-validated pseudo-R2 ranging from ~ 0.03 (for the GLM, and more than 2 times worse than our P-GAM in dlPFC) to ~0.09 (for neural networks, which is about equivalent to our method, particularly when considering the complexity of their areas and tasks, reaches). **E.** Example tuning functions to all the task variables and neural dynamics for a MSTd (green), 7a (blue) and dlPFC (red) neuron. Shaded areas are 95% posterior confidence intervals, and significant tuning (p<0.01) are opaque, while non-significant (p>0.01) tunings are transparent. **F.** Mask of significant tuning by variable (order is the same as in **E**, but vertical) and brain region for an example 200 units (x-axis). Significant tuning is colored, while non-significant is white. **G.** Number of variables a given neuron is tuned to, split by groups; sensorimotor, belief, LFP, and "other". Area 7a (blue) is tuned to more sensorimotor variables and to LFP phases than MSTd (green) or dlPFC (red). dlFPC is tuned to more internal variables, while MSTd (green) is seemingly tuned to eye-position. Upside-down arrow heads are medians. **H.** Fraction of neurons tuned in each area to a particular task variable (x-axis). Error bars are SEMs.*

## Cortical codes for path integration and vector-based navigation: asymmetric tiling of latent spaces

The crucial task-relevant variables in the firefly task are the relative angle and distance from self to target. These are latent, in that they are not directly observable by the monkey and must instead be dynamically computed based on immediate sensory experience. The encoding and updating of these spatial goals could take various forms. One suggested mechanism posits that the path toward spatial targets may be represented by activating a sequence of place cells (or more broadly, cells encoding self-position) from the target's location to the animal's current location (Ambrose et al., 2016), or vice-versa, from one's location to the goal location (Pfeiffer & Foster, 2013, also see Zhu et al., 2021 for an analogous phenomenon in gaze patterns, where sequential eye movements sweeping forward trace the trajectory from self to (sub-)goal). This mechanism is difficult to probe here, given that we were not recording in the hippocampal formation. A (non-mutually exclusive) alternative is a vectorial representation of the goal (Burgess & O'Keefe, 1996); an explicit encoding of the direction and distance to the target (also see Banino et al., 2018, for an intermediary solution, i.e., an artificial agent showing vector-based navigation using grid-like codes).

Indeed, Sarel and colleagues (2017) described units in the hippocampus (CA1) of bats that explicitly encoded the angular and radial distance to a target, even when this latter one was hidden (angular, 27%; radial, 16% of units). Here, we similarly observed units with a vectorial representation of spatial goals, even when targets are hidden and change position on a trial-by-trial basis. The angular distance to the hidden firefly was similarly frequent in MSTd (38.1%) and area 7a (34.1%), and approximately 10% more frequent than what is described in the bat's hippocampus. These neurons were twice as frequent in monkey dlPFC (53.6%) than bat hippocampus, and over half of the neurons recorded in dlPFC were tuned to the angular distance to target (**Figure 2H**). Similarly, the distance to target was most commonly encoded in dlPFC (30.6%), yet are still more frequent in area 7a (21.7%) and MSTd (25.1%) than what has previously been reported in the hippocampus (Sarel et al., 2017). Respectively 14.7%, 11.5%, and 20.5% of neurons in MSTd, 7a, and dlPFC, conjunctively encoded both the angular and radial distance to target – thus defining a particular spatial location in the environment relative to target. This is 1.53x, 1.53x, and 1.25x more frequent than what one would expect by chance (Chi-Square, all $X^2 > 9.9$, $p < 0.002$), suggesting that throughout sensory, parietal, and frontal cortices, the brain maintains a dynamic and explicit representation of (hidden) spatial goals relative to oneself.

The distribution of preferred angles and distances from origin (path integration) and to the target (vector-based navigation) are also informative. Each of these distributions spanned the entire range of angles and distances visited, demonstrating a full basis set covering the feature space (**Figure 3A,** distributions of preferred angle and distance from origin and to target in Monkey S).

Angles specifying a 0° offset from the heading at origin were over-represented (**Figure 3B**, top panel, proportions across all animals), while the 0° offset from target location was under-represented (**Figure 3A and B**, second row). Instead, particularly in area 7a and dlPFC, the distribution of preferred angles to target was bimodal (bimodality coefficients, null = 0.55, MST = 0.57, 7a = 0.89, dlPFC = 0.80, 7a & dlPFC p < 0.05), with many neurons showing a preference for ~45-90° away from target, either leftward or rightward. This pattern is in stark contrast with observations from the bat's hippocampus, where a goal angle of 0° is over-represented (Sarel et al., 2017). Speculatively, this reversed pattern of preferences between cortex and hippocampus may suggest

that the former plays a role in guiding heading toward the desired state (akin to an error or correction signal), while the hippocampus maintains this heading.

In terms of radial distance, MSTd, area 7a, and dlPFC all showed an over-representation of units coding for distances near to, as opposed to far from, the starting location (**Figure 3A and B**, 3rd row, < ~150cm). On the other hand, we observed a clear differentiation between cortical areas regarding their preferred distance to target. Paralleling observations from the bat's hippocampus (Sarel et al., 2017), area 7a showed a strong preference for nearby targets, though the absolute distances were very different. Approximately 50% of units coded for target distances within ~30cm (max = 400cm) in area 7a, while approximately 50% of units in the bat's hippocampus coded for distances within ~200cm (max = 1000cm). Neurons in MSTd and dlPFC responded primarily at intermediary and far distances to target (~200-400 cm from target; **3A and B**, 4th row).
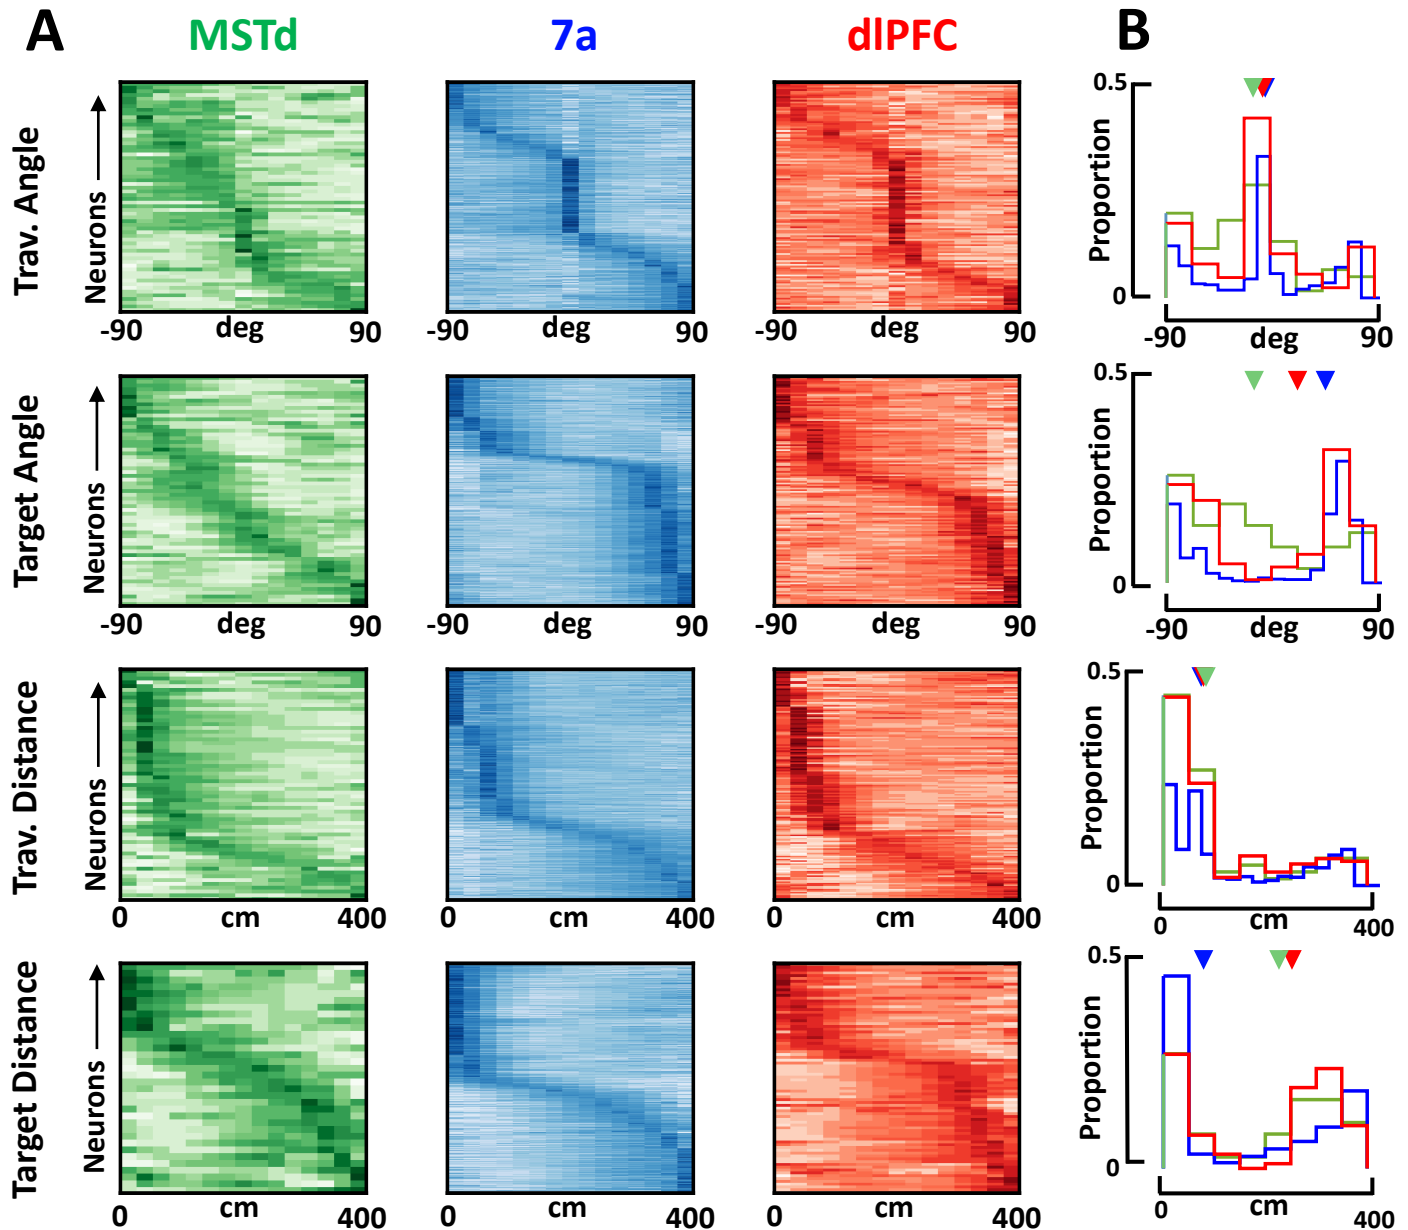


*Figure 3. Preferred angle and distance from origin and to target in MSTd, 7a, and dlPFC. **A.** Heatmaps showing neurons (y-axis) sorted by preferred angles from origin (top), angle to target (2nd row), distance from origin (3rd row), and distance to target (bottom row) for MSTd (green), 7a (blue) and dlPFC (red) in monkey S (data simultaneously recorded). Darker color indicates higher normalized firing rate. **B.** Histograms showing the probability of observing a given preferred angle or distance (order as in A) across all three monkeys. Inverted triangles at the top of each subplot indicate the median. Of note, however, medians may not appropriately summarize distributions in the case of bimodality.*

**Single-cell properties and unit-to-unit coupling suggest two distinct functional sub-networks**

To systematically examine how encoding of all task variables, and not solely a select few, varied across cortical areas, we employed a statistical clustering approach (see Minderer et al., 2019, for a similar approach).

First, we leveraged the knowledge of which variables each neuron was significantly tuned to (e.g., **Figure 2F**), and attempted clustering neurons based on this binary data, each neuron being defined by a size 17 binary vector (i.e., tuned or not tuned) corresponding to the 17 task variables defined in **Figure 2E**. **Figure 4A** shows the likelihood of neurons being tuned to any pair of variables (thresholded to a Jaccard Index > 0.5 for visualization, see *Methods*) for an example monkey. This visualization illustrates a differing pattern between MSTd and dlPFC on one hand, often being jointly tuned to eye position, target onset, and the latent variables (particularly in dlPFC), and area 7a on the other hand, being frequently jointly tuned to all sensorimotor variables and the phase of LFP bands. Most notably, clustering based on these 17-dimensional binary vectors (see *Methods*) showed that nearly all MSTd (89%) and dlPFC (94%) neurons were categorized as belonging within the same cluster (**Figure 4B,** cluster number 1), one that was defined by true mixed selectivity (**Figure 4C**, top left). In contrast, area 7a neurons appeared in three distinct clusters (**Figure 4B,** cluster numbers 1-3, respectively, 36%, 22%, and 31%). Cluster 2 was characterized by a strong selectivity for sensorimotor variables and firing in alignment with specific LFP phases (**Figure 4C**, top center), while Cluster 3 was characterized by a near complete lack of tuning to latent variables and eye position (**Figure 4C**, top right). Other cluster types existed, for instance composed of neurons selectively tuned to the ongoing phase in LFP bands but no other task variable (**Figure 4B and C**, Cluster 4), or driven also by motor onset and offset (**Figure 4B and C**, Cluster 5). These remaining clusters were, however, less common (1-5%). Together, this pattern of clustering results based on whether or not neurons were tuned to different task variables demonstrates a surprising commonality between MSTd and dlPFC, which are in turn different from area 7a.

Next, we questioned whether utilizing knowledge of the shape of these tuning functions (as opposed to simply whether or not tuning to a variable was significant) would dissociate between neurons in each area. Tuning functions for all task variables of a given neuron were stacked, then all neurons were projected onto a low dimensional manifold via Uniform Manifold Approximation and Projection (UMAP; McInnes et al., 2020), and clustered on this projection via DBSCAN (Density-Based Spatial Clustering of Applications with Noise; Ester et al., 1996). Area 7 and dlPFC neatly segregated from one another, while neurons from MSTd could be found along a continuum from area 7a to dlPFC. Notably, however, the centroid of MSTd was 6.49 times closer to the centroid of dlPFC than area 7a (**Figure 4D,** top row. Note that Becht et al., 2018, have shown UMAP to conserve global structure and thus allows for a meaningful interpretation of distances). Similarly, DBSCAN clustering highlighted 3 main clusters within this low dimensional projection of tuning functions. Cluster 1 (**Figure 4D,** second row) contained 81.5% of all dlPFC neurons, 58.9% of all MSTd neurons, and 3.6% of all neurons from 7a. Clusters 2 and 3 exclusively contained neurons from area 7a (**Figure 4D,** third and fourth row). Thus, just as the clustering based on whether or not neurons were tuned to different variables, clustering based on tuning shapes also highlighted a stronger similarity in the encoding properties of MSTd and dlPFC, as opposed to either of the former and area 7a (see **Figure S7** for UMAP projections color coded by the mutual information between neural responses and a particular task variable).
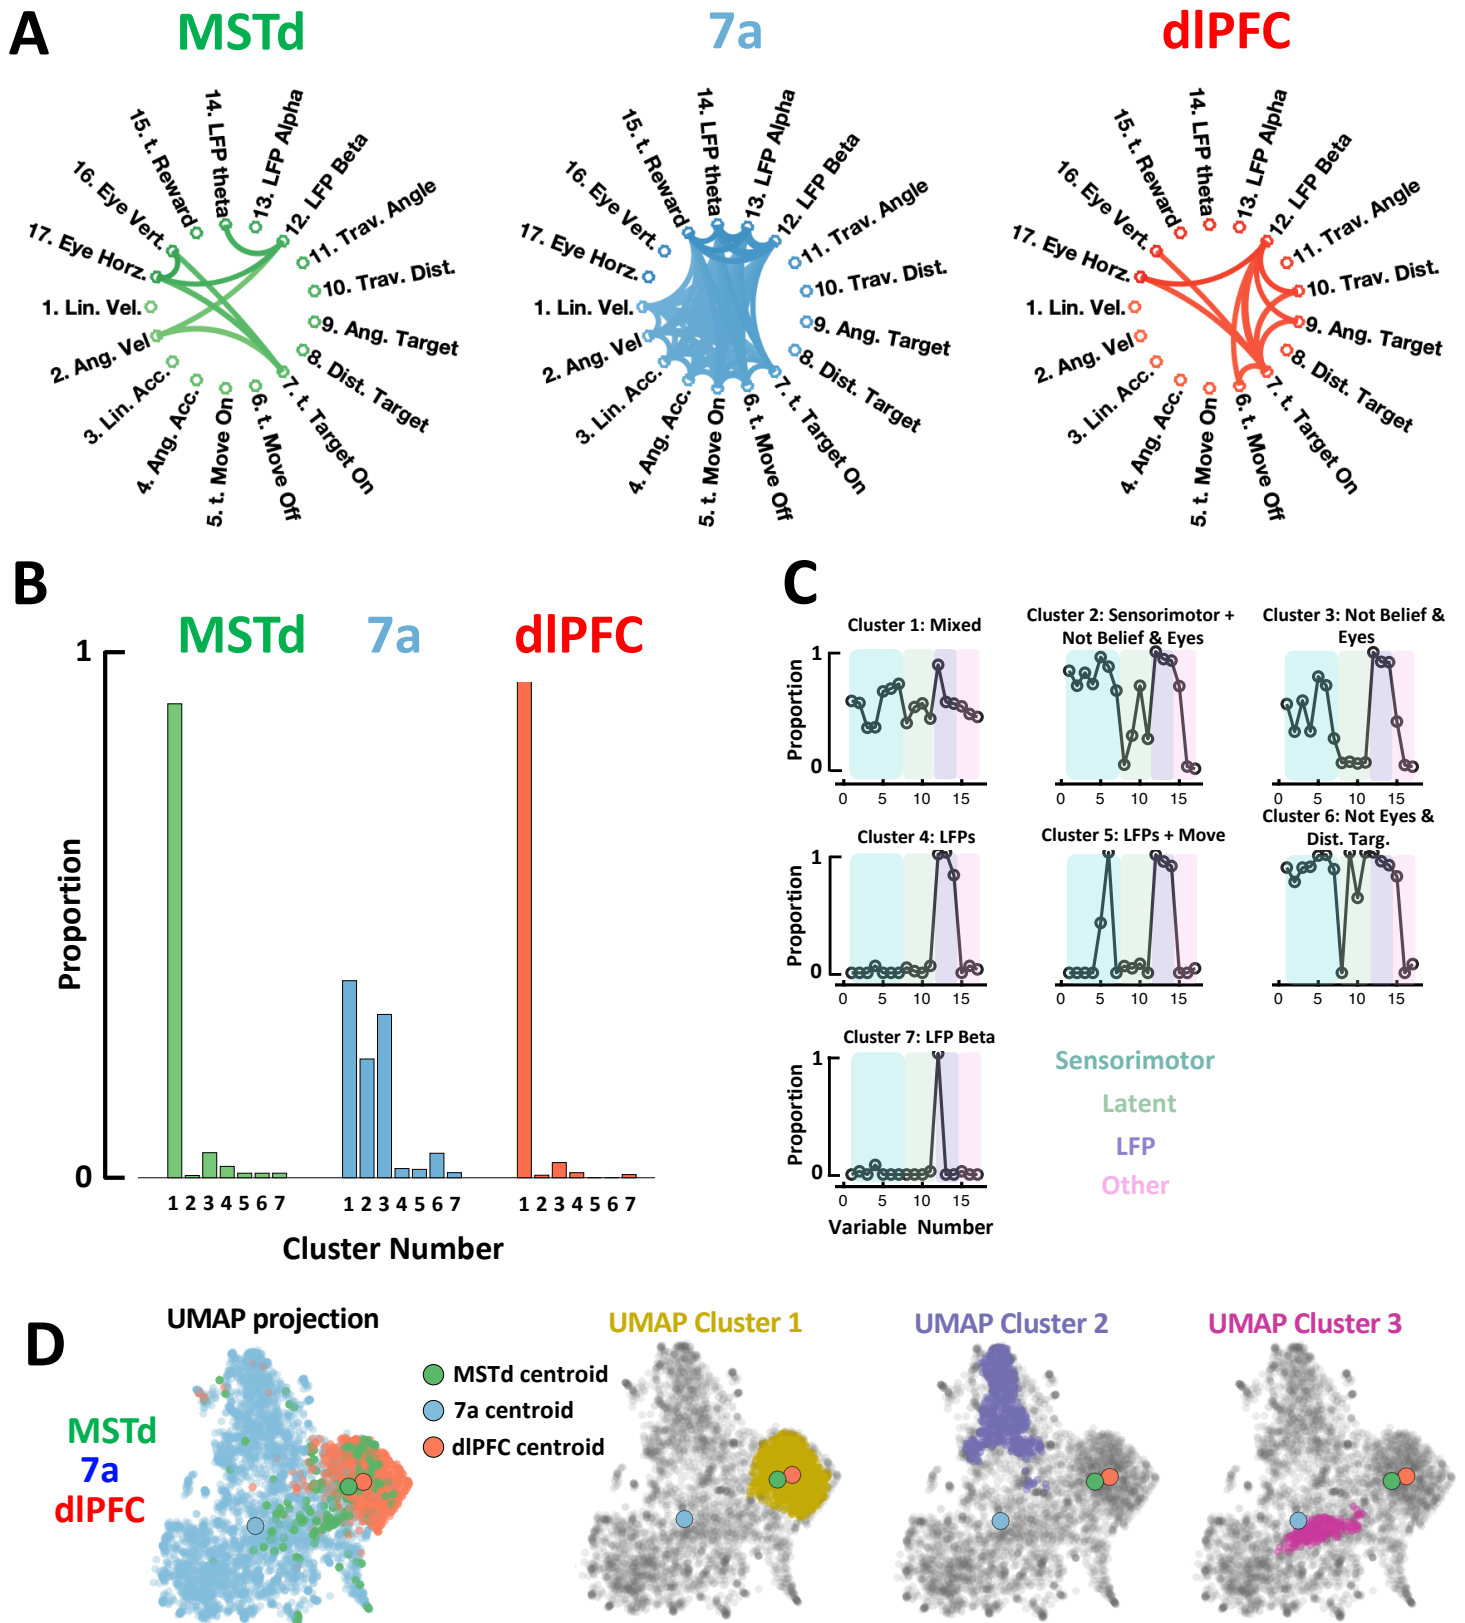
**Figure 4. Clustering of single unit tuning properties suggests a common functional role for MSTd and dlFPC. A.** *Circular graph illustrating which variables were commonly jointly tuned within a single cell in MSTd (green), area 7a (blue), and dlPFC (red), for an example animal. Thickness of the lines indicates frequency of joint tuning (masked to illustrated frequent pairs).* **B.** *Proportion of neurons being classified into distinct cluster (x-axis, 7 total) according to their profile of significant tunings.* **C.** *Fraction of neurons tuned to each of the 17 task variables (order is the same as in* **Figure 2**) *according to cluster.* **D.** *UMAP projection of the tuning function shapes, color coded by brain area (top) or DBSCAN cluster (second, third, and forth column).*

Given the surprising similarity in encoding profiles across MSTd and dlPFC, we next examined inter-area global synchronization and unit-to-unit coupling, in an attempt to question whether these two areas form an interacting sub-network. We examined inter-area coordination both from the standpoint of coarse grain LFP-to-spike phase locking, and via finer unit-to-unit couplings as estimated from the P-GAM.

As stated earlier (**Figure 2**), we found a very strong phase-locking of spiking in area 7a neurons to ongoing LFP phase within the same area, across theta, alpha, and beta bands (**Figure 5A**, blue self-connecting arrows). Phase-locking of spikes in dlPFC were more common in the higher (i.e., beta), rather than lower frequency bands, while phase locking in MSTd neurons remained constant across LFP bands. More importantly, spiking activity in some units was also well explained by the ongoing LFP phase in different regions (above and beyond the ongoing LFP phase in their local area). Most notably, 37% of neurons in dlPFC were tuned to the ongoing LFP phase within the beta range in 7a. An even greater proportion of units showed the reciprocal property, with 59% of neurons in area 7a being modulated by the ongoing phase of beta-band LFP in dlPFC. A considerable proportion of 7a units were also modulated by the ongoing phase of alpha-band LFP in MSTd (24%, **Figure 5A**). Globally, therefore, phase locking of spiking activity to LFP in different areas reflected the known anatomical connections, with reciprocal projections between MSTd and 7a, as well as between 7a and dlPFC (Andersen et al., 1990; Rozzi et al., 2006). Similarly, the putative "feedback" influence from dlPFC to area 7a (potentially reflecting the "belief" guiding motor action) was stronger than the putative "feedforward" influence of 7a onto dlPFC. Most noteworthily, however, there was no strong LFP-to-spike phase-locking between dlPFC and MSTd (the strongest influence between these being 12% of neurons in MSTd phase-locking to theta band LFP from dlPFC).

Finer grain unit-to-unit coupling was sparse, and within each area the probability of two units being functionally connected decreased as a function of the distance between neurons (**Figure S8**). The overall likelihood of two units within a given area being coupled changed as a function of brain area. There were more units coupled in MSTd (13%, corrected for distance between electrodes, see *Methods*) and area 7a (22%) than in dlPFC (7%, **Figure 5B**), potentially reflecting an increase in the dimensionality of the neural code (i.e., decrease in redundancy) as we move from sensorimotor areas to cognitive areas. More importantly, the across area unit-to-unit coupling did not reflect global anatomical trends, and instead supported the functional association between MSTd and dlPFC, as also suggested by the encoding profiles of these areas. Twelve percent of neurons in dlPFC were coupled to activity in MSTd, while 7% of neurons in MSTd were coupled to activity in dlPFC. Importantly, neither of these regions showed either a "feedforward" or "feedback" likelihood of being coupled with area 7a that exceeded 4% (**Figure 5B**).

Four arguments support the fact that unit-to-unit coupling as estimated by the P-GAM reflect a neural property beyond and orthogonal to global volume-conductance or spurious noise correlations. First, they significantly improve fits of the encoding model. **Figure 5C** shows the cross-validated pseudo-$R^2$ for the encoding model, including all task variables, LFP phases, and unit-to-unit coupling ("coupled model", y-axis), relative to a reduced model lacking the unit-to-unit coupling ("uncoupled model", x-axis. Also see **Figure 5D**). Second, the likelihood of two units within a given area being coupled was not random or purely depending on their physical distance, but instead varied as a function of tuning similarity. **Figure 5E** shows that the more similar the tuning functions were between two neurons, the more like these were to be coupled. The P-GAM was able to reflect this well-established "like-to-like connectivity" (Cossell et al., 2015; Chettih & Harvey, 2019), even for latent variables, and even in MSTd. Third, inter-area coupling filters did not show oscillatory behavior. Instead of sinusoidal patterns, we frequently observed gradual increases or decreases in the likelihood of the examined neuron to spike given the activity of another neurons (see **Figure S9** for a 'dictionary' of coupling filters across areas and their relative frequency). Lastly, the fact that our estimated cross-region unit-to-unit couplings are not likely to reflect a more global and unaccounted for correlation between areas is supported by the fact that while units phase-locked to the LFP in a different region were very likely to be phase locked to LFP in their own area (reflecting the utility of LFPs in coordinating spiking activity across areas), there was no significant change in the likelihood of coupled vs. uncoupled units to be phased locked to LFPs (**Figure S10**).

Altogether, both encoding profiles and fine grain unit-to-unit coupling across areas suggests that while the anatomical connectivity suggests a hierarchy from MSTd to area 7a to dlPFC, functionally and within this closed-loop virtual navigation task, MSTd and dlPFC form a subnetwork that is distinct from 7a.
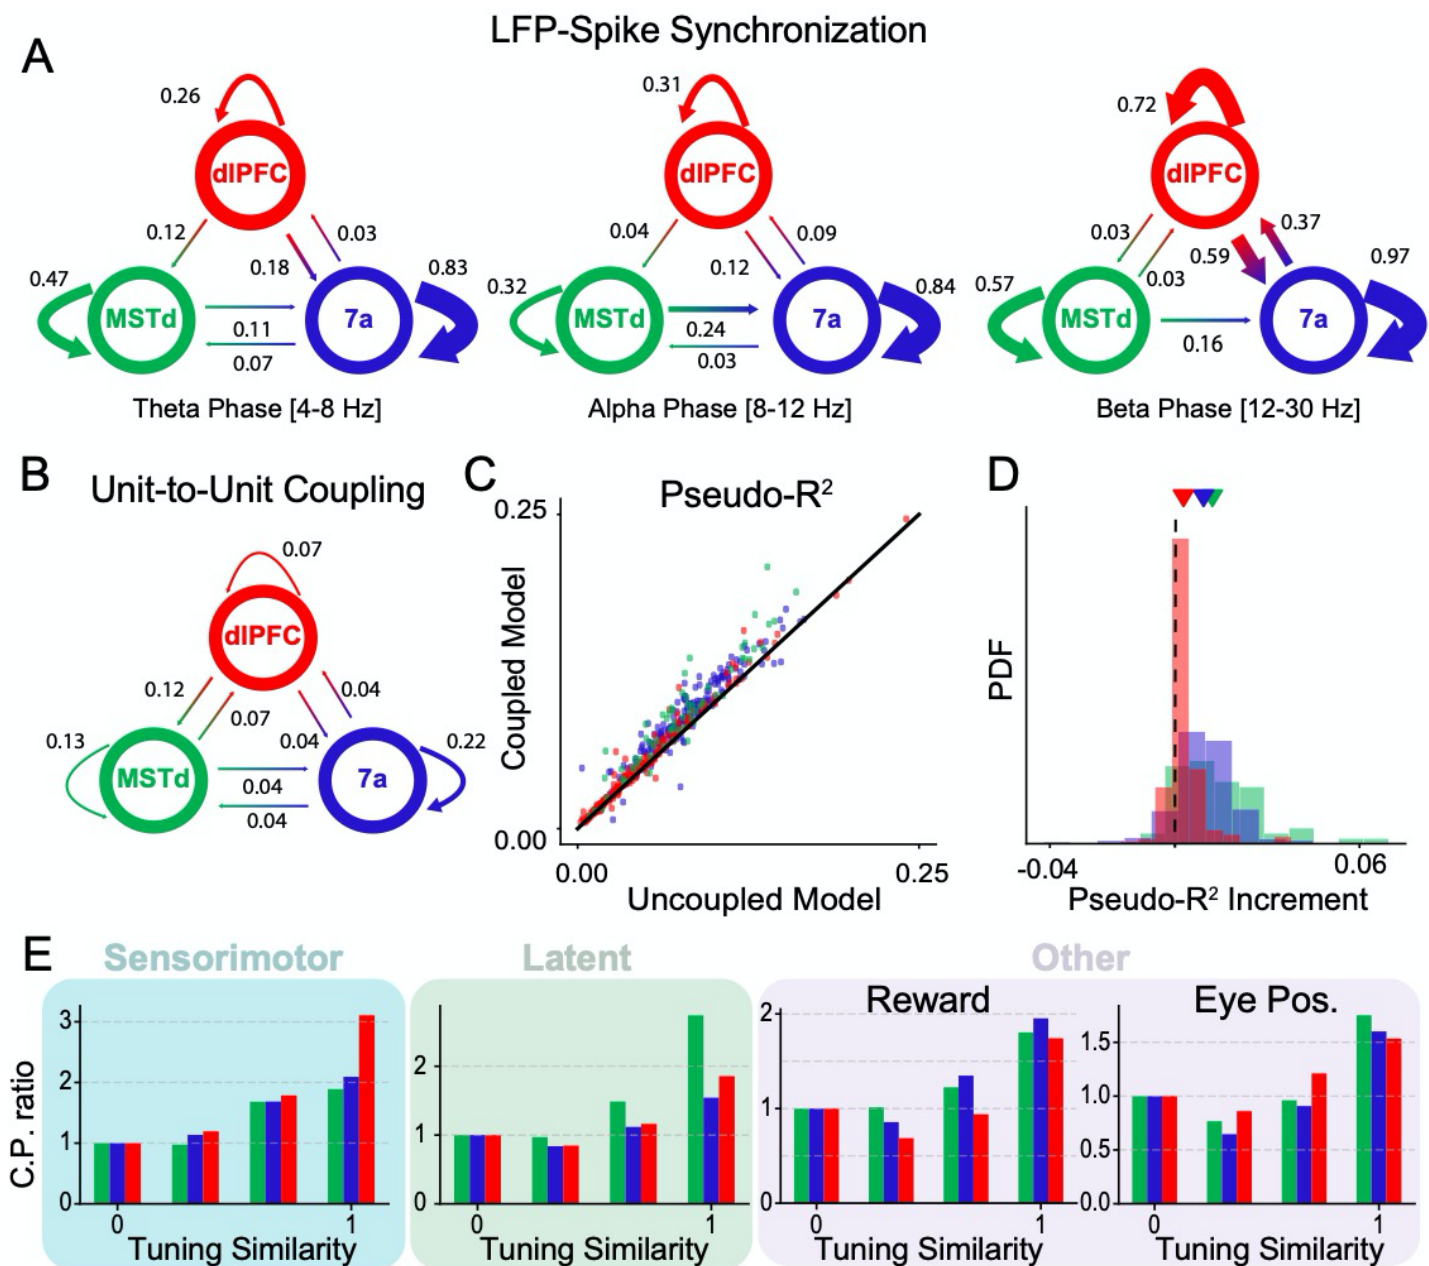
***Figure 5. LFP phase-spike and unit-to-unit coupling within and across MSTd, 7a, and dlPFC. A.** Fraction of neurons whose spiking activity is phase-locked to LFP phases within its own area (self-connecting arrows) and to that of other areas, in theta (leftmost), alpha (center) and beta (rightmost) bands. An arrow projecting from e.g., MSTd to 7a (center, 0.24), indicates that the neuron in area 7a is influenced by ongoing LFP phase in MSTd. Width of arrows and associated number indicate the proportion of neurons showing the particular property. **B.** As **A.** but for unit-to-unit coupling. An arrow projecting from e.g., MSTd to dlPFC indicates that the firing of a neuron in MSTd will subsequently influence spiking activity in dlPFC. **C.** Cross-validated pseudo-$R^2$ of the full encoding model (y-axis) and a reduced model without within and across area unit-to-unit coupling (x-axis). **D.** Probability distribution of the change in pseudo-$R^2$ from the reduced to the full model. **E.** Coupling probability (C.P.) between any two units given their tuning similarity to sensorimotor, latent, or other (reward and eye positions) variables. Tuning similarity is computed as the correlation in tuning functions, then these are discretized by their tuning similarity ($r^2 = [0-.25; 0.25-0.5; 0.5-0.75; 0.75-1]$) and averaged within each category (e.g., sensorimotor or latent) and tuning similarity bin. C.P. is expressed as a ratio, normalized to the bin with lowest tuning similarity (leftmost), such that a C.P. ratio of ~3 (e.g., sensorimotor variables in dlPFC) indicates that coupling is three times more likely given high vs. low tuning similarity.*

## Remapping of unit-to-unit coupling in 7a and dlPFC, and strengthening of connections in MSTd, during internal-model driven navigation

To further understand the role of each putative functional unit (MSTd and dlPFC on one hand, and 7a on another), we introduced an experimental manipulation changing the nature of the computation being performed. On a subset of randomly interleaved trials, the standard optic flow density (5 elements/m$^2$) was decreased by a

factor of 50. In this novel condition, the monkeys could not rely on incoming sensory evidence (Alefantis et al., 2021), and instead had to use their internal model of the mapping between joystick position and their velocity in the virtual environment (**Figure 6A**).

Compared to the high-density condition, during low-density optic flow monkeys further undershot targets (p = $5.32 \times 10^{-17}$), and their precision worsened (p = 0.005, **Figure 6B,** see Lakshminarasimhan et al., 2020, for a similar observation). This is precisely the pattern of results expected, given that observers are trying to maximize the overlap between their inferred position and that of the firefly. The observer's uncertainty of their own position grows over time/space (see **Figure 1C**), and thus their expectation of reward peaks at earlier distances under greater self-position uncertainty (see Lakshminarasimhan et al., 2018, for formal behavioral modelling and **Figure S11** for corroborative evidence showing that monkey's uncertainty over space scales quicker with time/space in the low- rather than high-density condition).
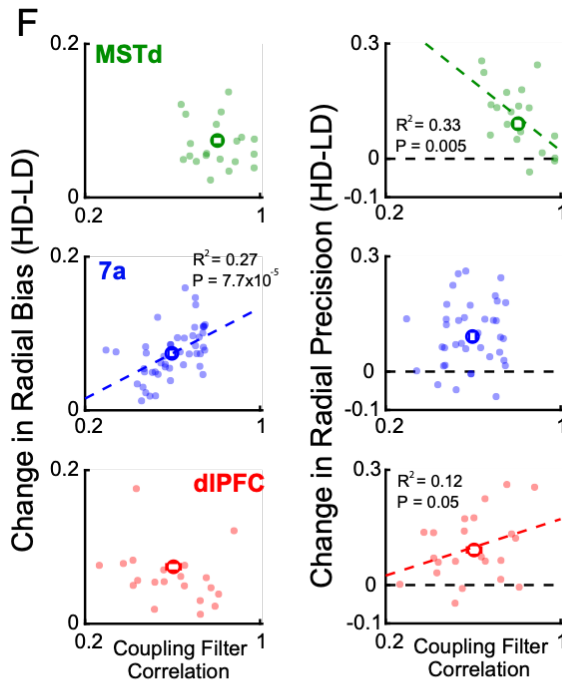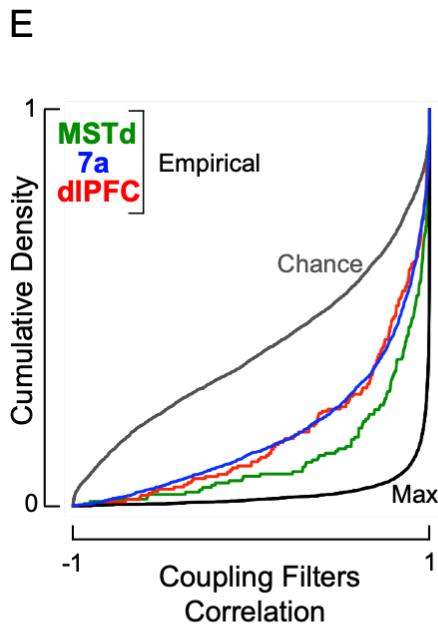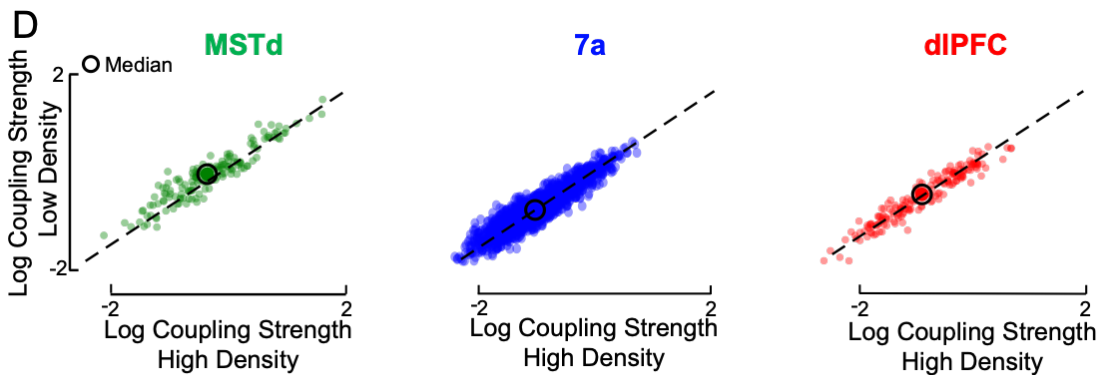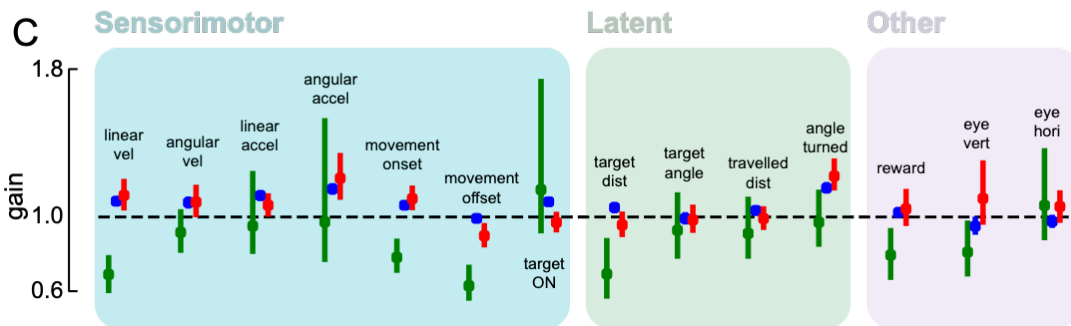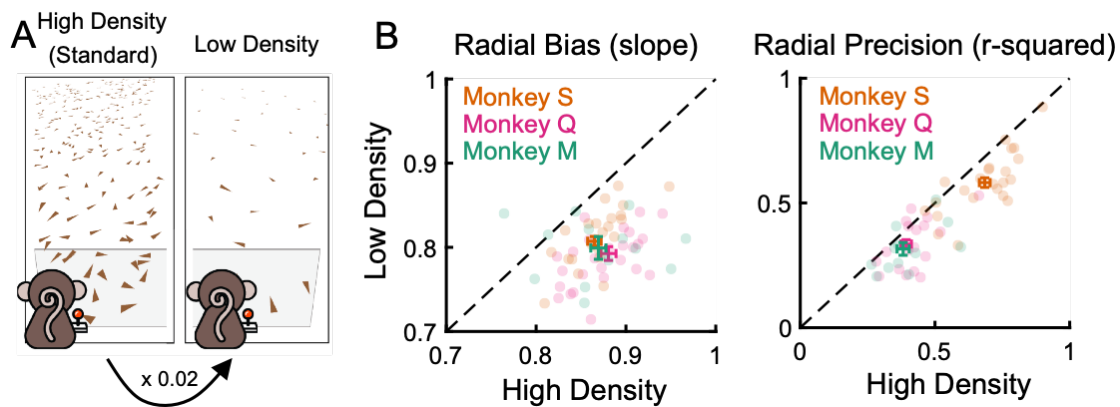
To quantify the effect of the change in optic-flow density (and presumably from sensory-driven to internal-model driven computation) on neural tuning curves, we computed the correlation between the tuning functions estimated by the P-GAM in each of these conditions. The slope of this correlation quantifies a change in gain in the neural responses (see **Figure S12** for example tuning functions in both optic flow conditions, as well as their gain modulation). While we observed significant gain modulation in all brain regions (all p <0.005), the effect size was considerable in MSTd (Cohen's d = 0.21) and negligible in the other areas (e.g., Cohen's d = 0.06 in area 7a). Responses were less prominent in MSTd (i.e., gains < 1) during the low-density condition, while the opposite was true for area 7a and dlPFC (**Figure 6C**). Interestingly, as if attempting to offset the impact of the unavailability of sensory evidence, the strength of coupling filters within MSTd increased during the low-density manipulation (p<0.005, paired t-test, Cohen's d= 0.51, **Figure 6D**). A similar observation has been reported in primary visual cortex, where the amplitude of LFP waves increases with decreased visual stimulus contrast, suggesting that lateral connections in sensory areas dominate when sensory evidence is weak (Nauhaus et al., 2009). In contrast, the strength of coupling filters within area 7a and dlPFC were unchanged in the low vs. high optic-flow density conditions (Cohen's d < 0.06, **Figure 6D**). The probability of unit-to-unit coupling, either within (all p>0.65) or across areas (all p>0.24, with one exception with negligible effect size, Cramer-V < 0.03), did not change between the different optic flow conditions.

Beyond the strength of neural responses, and the strength and probability of coupling filters, we may also index the shape of these coupling filters. As above, for all three areas we performed a correlation between the coupling filters in both conditions. **Figure 6E** shows the cumulative density function of the correlation coefficients for each area. In diametrical contrast to the coupling strength results, here we found that the parietal and frontal areas flexibly remapped from one condition to another, while MSTd remained fixed (Kruskal-Wallis, p = 0.002; post-hoc Holm-Sidak corrected, MSTd vs. 7a, p = 0.003; MSTd vs. dlPFC, p = 0.02; 7a vs. dlPFC, p = 0.76; **Figure 6E**). This is the first result where instead of MSTd and dlPFC showing a strong association (showing similar encoding profiles and across-area unit-to-unit coupling), we have area 7a and dlPFC showing a flexible remapping of unit-to-unit coupling, while MSTd coupling filters remain fixed yet strengthened. This latter result, with the traditionally considered sensory area showing inflexibility, is reminiscent of work suggesting that population coupling is invariant to stimulus conditions in visual cortex (Okun et al. 2015). Together with the current results, we may speculate that while sensory cortices generally show a fixed coupling pattern, higher-order areas flexibly reconfigure according to task demands. Estimating the flexibility or lack thereof of coupling filters across areas and density conditions was not possible due to their relatively low occurrence.

Finally, for sessions with both optic flow conditions and showing a minimum of 5 units coupled with another unit, we performed a correlation between the monkey's change in radial bias (slope, **Figure 6F**) and precision ($R^2$, **Figure 6E**) between low- and high-density environments, and the average stability of their unit-to-unit coupling (averaged within each brain region and session). Regarding radial bias, we found a strong correlation in area 7a ($R^2 = 0.27$, p = $7.7 \times 10^{-5}$; **Figure 6F**), indicating that the more the coupling filters changed in this area (i.e., the smaller the correlations coefficients on average), the less was accuracy impacted by the lack of sensory evidence. In other words, given that accuracy in trajectory endpoints were less accurate under the low rather than high optic-flow density condition (**Figure 6B**), this result suggests that the more units in 7a remapped, the less behavior suffered from the decreased availability of sensory evidence. There was no relation between the stability of unit-to-unit coupling in MSTd or dlPFC and changes in behavioral accuracy (all $R^2 < 0.05$ and all p > 0.14).

Regarding the precision of radial endpoints (i.e., the change in $R^2$ during high and low optic flow conditions), we found a negative correlation with the stability of unit-to-unit coupling in MSTd ($R^2 = 0.33$, $p = 0.005$), and a positive one with the stability in dlPFC ($R^2 = 0.12$, $p = 0.05$; **Figure 6G**). In other words, the more stable was the unit-to-unit coupling in MSTd, the less stable it was in dlPFC, the less was behavioral precision impacted by the loss of incoming sensory evidence. The individual nodes in the network formed by MSTd and dlPFC, thus, seemingly play different roles: the sensory area attempts to maximize a fixed signal, while dlPFC flexibly reconfigures according to task demands. There was no relation between the stability of unit-to-unit coupling in area 7a and changes in behavioral precision ($R^2 = 0.01$, $p = 0.26$).

Together, this pattern of results suggests that when incoming sensory evidence became scarce, neurons in the sensory region strengthened their lateral connections, but did not change their fine grain pattern of connectivity. On the other hand, area 7a and dlPFC remapped their fine grain unit-to-unit coupling, and the more they were able to do so, the less was behavioral accuracy (for 7a) and precision (for dlPFC) impacted by the loss of incoming sensory evidence.

**A** High Density (Standard)  Low Density

x 0.02

**B**

Radial Bias (slope)
Monkey S
Monkey Q
Monkey M
High Density
Low Density

Radial Precision (r-squared)
Monkey S
Monkey Q
Monkey M
High Density

**C**

Sensorimotor   Latent   Other

gain

linear vel, angular vel, linear accel, angular accel, movement onset, movement offset, target ON

target dist, target angle, travelled dist, angle turned

reward, eye vert, eye hori

**D**

MSTd   7a   dlPFC

O Median

Log Coupling Strength Low Density
Log Coupling Strength High Density

**E**

MSTd
7a
dlPFC

Empirical

Cumulative Density

Chance

Max

Coupling Filters Correlation

**F**

MSTd

7a   $R^2 = 0.27$   $P = 7.7 \times 10^{-5}$

dlPFC

Change in Radial Bias (HD-LD)

Coupling Filter Correlation

$R^2 = 0.33$   $P = 0.005$

$R^2 = 0.12$   $P = 0.05$

Change in Radial Precisioon (HD-LD)

Coupling Filter Correlation

**Figure 6. Behavioral and neural responses during low sensory evidence. A.** *Schematic showing the reduction in optic-flow density.* **B.** *Radial bias (slope describing the relation between target position and response) and precision ($R^2$ of the regression between target position and response) during high (x-axis) and low (y-axis) optic flow density conditions. Monkeys are in different colors, transparent circles are individual sessions, and error bars are means and S.E.M of each monkey. Overall, monkeys undershoot more and becomes less precise with less sensory evidence.* **C.** *Gain modulation of tuning functions, for all task variables and the three brain areas; MSTd, area 7a, and dlPFC. Gain is the slope describing the tuning functions in low vs. high density. Error bars are bootstrapped 95% CI.* **D.** *Unit-to-unit coupling strength for each brain region, and contrasting low (x-axis) and high (y-axis) optic flow density conditions. Coupling strength is quantified as the area under the curve of the coupling filters, and these are plotted on a log-scale for visibility.* **E.** *Cumulative density of the coupling filters correlation coefficient (r) between low- and high-density conditions, for all three brain regions. Chance and ceiling-levels were determined by permutation.* **F.** *Left: Correlation between the average coupling filter correlation within a given session (transparent circles), for each brain region, and the change in radial bias (high density – low density; positive values indicate more undershooting in the low-density condition). Larger, opaque circle is the average (y-axis = behavior; x-axis = coupling filter stability) across sessions. Only sessions with both density conditions and at least 5 coupled neurons were included. Right: As Left but for change in behavioral precision (quantified as the $R^2$ of the correlation between target location and response) and unit-to-unit coupling.*

## Discussion

The lion's share of our understanding linking spiking activity to goal-directed behavior originates from tasks that bear little resemblance to real world self-environment interactions. Major differences include the prevalence of closed loops between action and perception in natural behaviors, as well as the necessity to concurrently accomplish multiple sensory and cognitive operations over protracted periods of time (e.g., navigating while keeping in mind a target; see Sosa & Giacomo, 2021, for a recent review on the topic, yet with an exclusive focus on rodents and the hippocampal formation). To establish how the brain operates under these naturalistic conditions, and examine how neural dynamics change with changing computational demands, we recorded simultaneously from three cortical areas (MSTd, 7a, and dlPFC) while non-human primates detected a target, remembered its location within 2-dimensional space, and path integrated to it by actively sampling from and acting on their environment. We focus on the encoding profiles of single-units and describe four main findings.

First, we found that all task variables were present across all brain areas probed. This multiplexing, however, was patterned. The prefrontal cortex showed the greatest proportion of neurons tuned to latent variables, posterior parietal cortex showed a strong selectivity for sensorimotor variables and the ongoing phase of LFPs, and MSTd was prominently tuned to eye position. Most notoriously, even MSTd, an area typically considered a sensory node, showed a considerable number of neurons encoding latent variables, such as the angle or radial distance from origin and to target. In fact, the global encoding profile of MSTd was strikingly akin to that of dlPFC, suggesting this area is not purely sensory (see Mendoza et al., 2014 for a similar observation).

The second major finding relates to functional organization across the MSTd, 7a, dlPFC network. Anatomical tracer studies suggest a distributed hierarchy, with recurrent connections between MSTd and 7a, and then recurrent connections between 7a and dlPFC (Andersen et al., 1990; Rozzi et al., 2006). As expected from anatomy, here we observed that spiking activity in 7a is influenced by LFP fluctuations in MSTd (particularly in the alpha range). In turn, spiking activity in dlPFC and 7a are influenced by LFP fluctuations, respectively, in 7a and dlPFC (particularly in the beta range). In other words, we observe that the coarse channels of communication riding on LFPs are in line with anatomy. However, in contrast to an anatomy-based hypothesis that MSTd signals reach dlPFC through area 7a, we observe a likelihood of MSTd neurons being coupled to dlPFC, and of dlPFC neurons being coupled to MSTd, that is, respectively, 3 and almost 2 times as large as coupling among either of these areas with 7a. These results suggest that despite the lack of direct anatomical connections, MSTd and dlPFC may form a functional subnetwork within this closed-loop path integration task. The area or areas that functionally link MSTd and dlPFC is at present unknown, but it is not area 7a.

Third, sensory, parietal and frontal cortices re-organize differently during changing computational demands. More precisely, when sensory evidence became scarce, and thus the animal had to largely rely on their internal models, we uncovered what resembles a two-pronged process. First, the tuning curves of the sensory area were impacted by the lack of sensory evidence: decreasing their gain but not changing in overall shape. And as if attempting to compensate for this, the lateral connectivity within MSTd strengthened. However, the connectivity pattern – the shape of coupling filters – remained relatively fixed and did not adapt to the changing environment. On the other hand, the connectivity strength of coupling filters was unchanged within parietal and frontal areas, yet their patterns flexibly remapped. Strikingly, the less the sensory area remapped, and the more the prefrontal and parietal cortices did, the less was behavior impacted by the loss of sensory evidence. Speculatively, this two-pronged approach, with diverging patterns in the remapping of unit-to-unit coupling among "sensory" and

"cognitive" areas, may allow for a compromise between faithfully reflecting our sensory environment and not completely changing our interpretation of the world around us given small, stochastic, fluctuations. In other words, the "sensory" area, showing a fixed connectivity pattern (Okun et al. 2015) that is gain-modulated (Nauhaus et al., 2009) may allow maximizing our encoding capacity. On the other hand, it is possible that our internal models of the environment are housed in the precise connectivity within "cognitive" areas.

These findings and interpretation – while at a very different scale and relying on different modalities, i.e., spiking activity vs. oxygenation in vasculature – are reminiscent of observations from fMRI functional connectivity. Indeed, authors (Cole et al., 2013) have argued that the fronto-parietal network is a flexible hub for adaptive control whose coupling with sensory areas may change according to task goals (Chadick & Gazzaley, 2011; Zanto & Gazzaley, 2013). In fact, Ekman and colleagues (2012) showed that the degree of functional connectivity (in fMRI) between a cognitive 'core' (e.g., dlPFC) and task-relevant sensory 'periphery' (e.g., MSTd) is predictive of behavioral performance. Here we demonstrate an analogous effect (i.e., coupling between a task-relevant "sensory" area and dlPFC, as well as remapping of lateral connectivity within dlPFC and area 7a), at the level of unit-to-unit coupling and within a naturalistic and dynamic task.

Lastly, we found that MSTd, area 7a, and dlPFC all show evidence for vector-based coding of spatial goals (Sarel et al., 2017, also see Ekstrom et al., 2003, and Gauthier & Tank, 2018, for similar evidence in humans and rodents). Interestingly, we observed that area 7a showed a majority of neurons coding for locations near the origin or near the target. That is, 7a is likely involved in body-centered state transitions (e.g., from static to moving) and approaching behavior (see Medendorp & Heed, 2019 for a similar argument and Serino, 2019, for a review on the role of posterior parietal cortex in peri-personal space encoding). dlPFC, on the other hand (and to a lesser extent MSTd), showed a preponderance of units coding for locations near the origin and far from the target. Interestingly, human fMRI work (Howard et al, 2014) has suggested that the posterior hippocampus encodes path distance to goals, while the anterior hippocampus encodes a more abstract Euclidian distance. Together with the current findings, and the fact that anterior and posterior hippocampus respectively project to the prefrontal and parietal cortex (Strange et al., 2014), suggests that a circuit including the anterior hippocampus and prefrontal cortex may abstractly plan goal-directed trajectories, while a circuit including the posterior hippocampus and parietal cortex may compute the body-centric sensorimotor affordances.

The fact that all areas probed showed a preference for distances near (as opposed to far) from where the monkey started path integrating (i.e., the origin) may suggest a change in the computations undertaken early vs. late within a trajectory. Namely, early on observers may compute their ongoing distance from the origin, and then switch to computing distance not from origin, but from the target. This switch from first coding displacement to then coding location relative to a target has been previously suggested empirically (Gothard et al., 1996) and in recent computational accounts (Kanitscheider & Fiete, 2017). The full account will be best probed by future studies examining population dynamics within different periods of a single trajectory.

In conclusion, we demonstrate what may be broad principles of brain computation during closed action-perception loops: (i) mixed yet patterned single cell selectivity, (ii) fine grain connectivity between task relevant areas (i.e., MSTd, during our optic flow task) and pre-frontal cortex, and (iii) a dynamic strengthening (in "sensory" areas) and remapping (in "cognitive" areas) of lateral connectivity as our environments require a heavier weighting on internal models vis-à-vis sensory evidence (e.g., navigating our homes at night). Together, these insights demonstrate the power of studying brain function within a scenario akin to that it evolved to solve.

**Methods**

Animals and Animal Preparation

We recorded extracellularly from areas MSTd, area 7a, and dlPFC in 3 adult male rhesus macaques (Macaca Mulatta; 9.8-13.1kg). We collected behavioral and neural data from 27 recording sessions from monkey Q, 38 sessions from monkey S, and 17 sessions from monkey M. All animals were chronically implanted with a lightweight polyacetal ring for head restraint. Further, for acute recordings, animals were outfitted with a removable grid to guide electrode penetrations. For eye-tracking, a subset of monkeys (Monkey Q) were implanted with scleral coils (CNC Engineering, Seattle WA, USA), while eye tracking was performed using a video-based system (ISCAN Inc. Woburn, USA) in the remaining animals (Monkeys S and M). Monkeys were

trained using standard operant conditioning procedures to use a joystick in order to navigate in virtual reality (VR) and stop at the location of briefly presented targets, "fireflies". All surgeries and procedures were approved by the Institutional Animal Care and Use Committee at Baylor College of Medicine and New York University, and were in accordance with National Institutes of Health guidelines.

Experimental Setup

Monkeys were head-fixed and secured in a primate chair. A 3-chip DLP projector (Christie Digital Mirage 2000, Cypress, CA, USA) rear-projected images onto a 60 x 60 cm screen that was attached to the front of the field coil frame, ~ 30 cm in front of the monkey. The projector rendered stereoscopic images generated by an OpenGL accelerator board (Nvidia Quadro FX 3000G). To navigate, the animals used an analog joystick (M20U9T-N82, CTI electronics) with two degrees of freedom to control their linear and angular speeds in a virtual environment. This virtual world comprised a ground plane whose textural elements had limited lifetime (250 ms) to avoid serving as landmarks. The ground plane was circular with a radius of 70 m (near and far clipping planes at 5 cm and 40 m respectively), with the subject positioned at its center at the beginning of each trial. Each texture element was an isosceles triangle (base x height: 8.5 x 18.5 cm$^2$) that was randomly repositioned and reoriented at the end of its lifetime, making it impossible to use as a landmark. The maximum linear and angular speeds were fixed to 2 m/s and 90 deg/s. The standard density of the ground plane was 5.0 elements/m$^2$ (**Figures 1-5**). This density was reduced by a factor of 50 (0.1 elements/m$^2$ in **Figure 6**) in a manipulation condition, in order to force the animals to use their acquired internal model, rather than sensory evidence. All stimuli were generated and rendered using C++ Open Graphics Library (OpenGL) by continuously repositioning the camera based on joystick inputs to update the visual scene at 60 Hz. The camera was positioned at a height of 0.1 m above the ground plane. Spike2 software (Cambridge Electronic Design Ltd., Cambridge, UK) was used to record and store the timeseries of target and animal's location, animal linear and angular velocity, as well as eye positions. All behavioral data were recorded along with the neural event markers at a sampling rate of 833.33Hz.

Behavioral Task

Monkeys steered to a target location (circular disc of radius 20 cm) that was cued briefly (300ms) at the beginning of each trial. Each trial was programmed to start after a variable random delay (truncated exponential distribution, range: 0.2 – 2.0 s; mean: 0.5 s) following the end of the previous trial. The target appeared at a random location between –40 to 40 deg of visual angle, and between 1 and 4 m relative to where the subject was stationed at the beginning of the trial. The joystick was always active, and thus monkeys were free to start moving before the target vanished, or before it appeared. Monkeys typically performed blocks of 750 trials before being given a short break. In a session, monkeys would perform 2 or 3 blocks. Feedback, in the form of juice reward was given following a variable waiting period after stopping (truncated exponential distribution, range: 0.1 – 0.6 s; mean: 0.25 s). They received a drop of juice if their stopping position was within 0.6 m away from the center of the target. No juice was provided otherwise. Monkeys were first trained extensively, gradually reducing the size of the reward zone until their performance stopped improving.

Neural Recording and Pre-Processing

We recorded extracellularly, either acutely using a 24-channel linear electrode (100 µm spacing between electrodes; U-Probe, Plexon Inc, Dallas, TX, USA; MSTd in Monkey Q and S, and dlPFC in monkey M) or chronically with multi-electrode arrays (Blackrock Microsystems, Salt Lake City, UT, USA; 96 electrodes in area 7a in Monkey Q, and 48 electrodes in both area 7a and dlPFC in monkey S). During acute recordings with the linear arrays, the probes were advanced into the cortex through a guide-tube using a hydraulic microdrive. Spike detection thresholds were manually adjusted separately for each channel in order to facilitate real-time monitoring of action potential waveforms. Recordings began once waveforms were stable. The broadband signals were amplified and digitized at 20 KHz using a multichannel data acquisition system (Plexon Inc, Dallas, TX, USA) and were stored along with the action potential waveforms for offline analysis. Additionally, for each channel, we also stored low-pass filtered (-3dB at 250Hz) local-field potential (LFP) signals. For the array recordings, broadband neural signals were amplified and digitized at 30 KHz using a digital headstage (Cereplex E, Blackrock Microsystems, Salt Lake City, UT, USA), processed using the data acquisition system (Cereplex Direct, Blackrock Microsystems) and stored for offline analysis. Additionally, for each channel, we also stored low-pass filtered (-6 dB at 250 Hz) local-field potential (LFP) signals sampled at 500 Hz. Finally, copies of event

markers were received online from the stimulus acquisition software (Spike2) and saved alongside the neural data.

Spike detection and sorting were initially performed on the raw (broadband) neural signals using KiloSort 2.0 software (Pachitariu et al., 2016) on an NVIDIA Quadro P5000 GPU. The software uses a template matching algorithm both for detection and clustering of spike waveforms. The spike clusters produced by KiloSort were visualized in Phy2 and manually refined by a human observer, by either accepting or rejecting KiloSort's label for each unit. In addition, we computed three isolation quality metrics; inter-spike interval violations (ISIv), waveform contamination rate (cR), and presence rate (PR). ISIv is the fraction of spikes that occur within 1ms of the previous spike. cR is the proportion of spikes inside a high-dimensional cluster boundary (by waveform) that are not from the cluster (false positive rate) when setting the cluster boundary at a Mahalanobis distance such that there are equal false positives and false negatives. PR is 1 minus the fraction of 1 minute bins in which there is no spike. We set the following thresholds in qualifying a unit as a single-unit: ISIv < 20%, cR < 0.02, and PR > 90%.

Analyses

*Behavior*

The location of targets and monkey's endpoints were expressed in polar coordinates, with a radial distance (target = $r$, response = $\tilde{r}$) and eccentricity from straight ahead (target = $\theta$; response = $\tilde{\theta}$). On a subset of trials (~13%) animals stopped after <0.5m (range of targets, 1m to 4m). Similarly, on a subset of trials (~5%) animals did not stop during the course of a trial (max duration = 7 seconds). These trials were discarded before further behavioral analyses. As we have observed before (Lakshminarasimhan et al., 2018, 2020; Noel et al., 2020, 2021) a linear model with multiplicative gain accounted well for the observed data (average $R^2$ = 0.72). Thus, we used the slopes of the corresponding linear regressions as a measure of bias. Note that in this scheme a slope of one indicates no bias, while slopes smaller than one indicate undershooting (either in radial or angular distance).

To account for greater undershooting apparent in the low-density optic flow condition, we examined the scaling of uncertainty with distance to the target. Low- and high- density trials were discretized in equal size quartiles by radial distance (i.e., 4 different distances). Within each quartile, we computed the standard deviation of the radial responses. As expected, variance increased with distance. Importantly, we were particularly interested in whether the scaling of uncertainty with distance was greater in the low density vs. high density condition. Thus, we expressed the ratio between these (low density standard deviation divided by high density standard deviation), and examined this quantity as a function of distance. The fact that this ratio increases over space (from ~1 to ~1.2, **Figure S11**) shows that uncertainty over position scales quicker in the low rather than the high optic flow density condition.

*Poisson Generalized Additive Model*

The Poisson Generalized Additive Model (P-GAM; https://github.com/savin-lab) defines a non-linear mapping between spike counts of a unit $y_t \in \mathbb{N}_0$ (binned at 6ms temporal resolution) and a set of continuous covariates $x_t$ (angular and linear velocity and acceleration, angular and linear distance traveled, angular and linear distance to target, and LFP instantaneous phase across different frequency ranges), and discrete events $z_t$ (time of movement onset/offset, target onset, reward delivery, and the spike counts from simultaneously recorded units). In particular, the unit log-firing rate is modeled as a linear combination of arbitrary non-linear functions of the covariates,

$$\log \mu = \sum_j f_j(x_j) + \sum_k f_k * z_k \qquad \text{(Eq. 1)}$$

where * is the convolution operator, and the unit spike counts are generated as Poisson random variables with rate specified by (Eq. 1).

Input specific non-linearities $f(\cdot)$ are expressed in terms of flexible B-splines, $f(\cdot) \approx \boldsymbol{\beta} \cdot \boldsymbol{b}(\cdot)$ and are associated with a smoothness enforcing penalization term controlled by a scale parameter $\lambda_f$,

$$PEN(f, \lambda_f) = -\frac{1}{2}\lambda_f \boldsymbol{\beta}^T S_f \boldsymbol{\beta}, \quad S_f = \int b'' \cdot b''^T \, dx \qquad \text{(Eq. 2)}$$

The larger $\lambda_f$, the smoother the model. This penalization terms can be interpreted as (improper) Gaussian priors over model parameters, the resulting log-likelihood of the model takes the form,

$$\mathcal{L}(\mathbf{y}) = \log p(\mathbf{y}|\mathbf{x}, \mathbf{z}, \boldsymbol{\beta}) + \sum_f PEN(f, \lambda_f) \qquad \text{(Eq. 3)}$$

with $\mathbf{y} \in \mathbb{R}^T$ the spike counts of the unit, $\mathbf{x} \in \mathbb{R}^{J \times T}$ and $\mathbf{z} \in \mathbb{R}^{K \times T}$, $T$ the time points, $\boldsymbol{\beta}$ the collection of all B-spline coefficients and $p(\cdot)$ the Poisson likelihood. Both parameters $\boldsymbol{\beta}$ and the hyperparameters $\lambda$ are learned from the data by an iterative optimization procedure that switches between maximizing Eq. 3 as a function of the parameters, and minimizing a cross-validation score as a function of the hyperparameters (see Balzani et al., 2020 for further details).

The probabilistic interpretation of the penalization terms allowed us to compute a posterior distribution for the model parameters, derive confidence intervals with desirable frequentist coverage properties, and implement a statistical test for input variable inclusion that selects a minimal subset of variables explaining most of the variance. In **Figure S2** we demonstrate that this approach will appropriately select variables contributing to the spike trials of individual neurons. Further, in **Figure S13** we show that including all variables in the model (thousands of parameters, given the cell-to-cell coupling) does not outperform the selected model, which has an order of magnitude less parameters. The approach circumvents traditional model-comparison-based variable selection, which would be unfeasible in a fully coupled model with hundreds of covariates.

*Pseudo-R²*

Fit quality was assessed via the pseudo-$R^2$ on subset of held-out test trials (20% of the total trials, not used for inferring model parameters). Pseudo-$R^2$ is a goodness of fit measure that is suitable for models with Poisson observation noise (Cameron et al., 1997). The score is computed as,

$$\text{pseudo } R^2 = 1 - \frac{\mathcal{L}(\mathbf{y}) - \mathcal{L}(\hat{\mathbf{y}})}{\mathcal{L}(\mathbf{y}) - \mathcal{L}(\bar{\mathbf{y}})} \qquad \text{(Eq. 4)}$$

with $\mathcal{L}(\mathbf{y})$ the likelihood of the true spike counts, $\mathcal{L}(\hat{\mathbf{y}})$ the likelihood of the P-GAM model prediction and $\mathcal{L}(\bar{\mathbf{y}})$ the likelihood of a Poisson null-model (constant mean rate). It can be interpreted as the fraction of the maximum possible likelihood (normalized by the null-model) that the fit achieves. The score is 0 when the GAM fits are no more likely than the null model, 1 when it perfectly matches the data, and rarely can be negative (for test-set data, 3% of the recorded units) when overfitting occurs. In this latter case we excluded the unit from analysis. Empirically, the pseudo-$R^2$ is a stringent metric and ranges in values that are lower than the standard $R^2$ when both are applicable (Domencich et al., 1975) and depends on the firing statistics of the area recorded (such as mean firing rates and variance), the type of covariates, and their number (Benjamin et al., 2017). Our average score of 0.072 is about three times better than standard GLM performance even in areas such as primary somatosensory and motor areas (Benjamin et al., 2018).

*Clustering*

The circular graph in **Figure 4A** was plotted using circularGraph software (https://github.com/paul kassebaummathworks/circularGraph). Given two logical vectors (A and B, for variables x and y) in which '1' indicates the variable is encoded by the corresponding neuron, we calculated the Jaccard index:

$$J(A, B) = \frac{|A \cap B|}{|A| + |B| - |A \cap B|} \qquad \text{(Eq. 5)}$$

where $|A \cap B|$ indicates the neuron count that encodes both variables, $|A|$ and $|B|$ indicate neuron count encoding variable x and y, respectively. Variable pair which had a Jaccard index below 0.5 were masked away for visualization. For **Figures 4B, C,** clustering was performed using the Jaccard indices and employing a spectral

clustering algorithm (Shi & Malik, 2000) to find the $k$ eigenvectors to split the resulting distances $k$ ways. Different values of $k$ (>=3) resulted in conceptually identical findings.

For **Figures 4D, F,** all tuning function of a neuron were stacked together, so each neuron was represented by a vector of dimension $\sum n_i$ where $n_i$ is the dimension of the $i^{th}$ tuning function. Then, the high dimensional matrix (neurons x $\sum n_i$) was reduced to a matrix (neurons x $m$) by PCA projection, where $m$ is the number of PCs that explains >90% of the variance. Finally, the matrix is further projected onto a 2 dimensional manifold using UMAP (McInnes et al., 2020) and clustered using DBSCAN (Ester et al., 1996). This clustering method automatically determined the number of clusters on the basis of the spatial density of the UMAP projections.

For **Figure S7**, mutual information was computed as the difference of the entropy of spike counts H(Y), and the entropy of the counts conditioned on the stimulus H(Y|S),

$$I(Y,S) \; = \; H(Y) - H(Y|S) \quad \text{(Eq. 6)}$$

To estimate these quantities, we assumed that the counts $Y$ are Poisson distributed with rate parameter $\lambda$ equal to the mean firing rate of the unit. The stimulus was discretized and its distribution approximated as a binomial, while the conditional distribution of the counts given the stimuli $p(Y|S)$ is given by the P-GAM generative model. Finally, we computed $Y$ as $H(Y) = -\sum_y p(y) \log p(y)$, and $H(Y|S) = -\sum_s p(s) \sum_y p(y|s) \log p(y|s)$.

*Coupling*

Coupling filters (and the corresponding inferential statistics) were determined via the P-GAM. Within area coupling filters were set to a duration of 36ms, and across area filters were set to a duration of 600ms. For the within area coupling probability reported in **Figure 5B**, we corrected for the effects of unit distance (i.e., **Figure S8**) by first fitting a brain region-specific logistic regression. Specifically, we expressed coupling probability as a non-linear function of electrode distance as follows,

$$p(c = 1) = logit^{-1}(f(d)) \quad \text{(Eq. 7)}$$

With c a binary variable taking value 1 for significant coupling and 0 otherwise, d the electrode distance, and f a non-linear function expressed in terms of B-splines. Each brain area was fit independently, and the coupling probability in **Figure 5B** was set as the model prediction for a distance of 500um (**Figure S8**).

For **Figure 5E,** coupling probability was computed as the fraction of significantly coupled units within area pairs **(**p-value <0.001 covariate inclusion test**),** while tuning similarity was computed as $s = 1 - \left\| \bar{f} - \bar{g} \right\|_2$, where $\bar{f}$ and $\bar{g}$ denote a normalized version of the response functions $f$ and $g$ of two units, for a subset of inputs**.** Coupling strength (e.g., **Figure 6D**) was computed as the area under the filter. Coupling stability (e.g., **Figure 6E**) was computed as the correlation between coupling filters in the low- and high-density condition. Chance and ceiling levels were established by permutation. To relate the change in within area connectivity with the change in behavior caused by the decrease in optic flow signals, we computed the average coupling stability per area and session, given at least 5 connected units and a P-GAM pseudo-r^2 > 0.01.

## Acknowledgements

## Author Contributions

D.E.A, K.L., and E.A. conceived and designed the experiments; D.E.A., J.P.N., E.B., and C.S designed the data analyses and framed the hypotheses; E.A., and R.K., performed the surgeries; E.A., S.B., and P.A. collected

experimental data; J.P.N and E.B. performed data analyses; C.S. and E.B. contributed key analytical resources; J.P.N., E.B., C.S., and D.E.A. wrote the manuscript.

**Competing Interests Statement**

No author had any competing interests of any kind.

**References**

Alefantis, P., Lakshminarasimhan, K.J., Avila, E., Noel, J-P., Pitkow, X., Angelaki, D.E (2021). Sensory evidence accumulation using optic flow in a naturalistic navigation task. bioRxiv 2021.04.26.441532; doi: https://doi.org/10.1101/2021.04.26.441532

Ambrose RE, Pfeiffer BE, Foster DJ. Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. Neuron. 2016 Sep 7;91(5):1124-1136. doi: 10.1016/j.neuron.2016.07.047. Epub 2016 Aug 25. PMID: 27568518; PMCID: PMC6013068.

Andersen RA, Asanuma C, Essick G, Siegel RM. Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. J Comp Neurol. 1990 Jun 1;296(1):65-113. doi: 10.1002/cne.902960106. PMID: 2358530.

Balzani, E., Lakshminarasimhan, K., Angelaki, D., Savin C. Efficient estimation of neural tuning during naturalistic behavior. Advances in Neural Information Processing Systems, 33, 2020

Banino A, Barry C, Uria B, Blundell C, Lillicrap T, Mirowski P, Pritzel A, Chadwick MJ, Degris T, Modayil J, Wayne G, Soyer H, Viola F, Zhang B, Goroshin R, Rabinowitz N, Pascanu R, Beattie C, Petersen S, Sadik A, Gaffney S, King H, Kavukcuoglu K, Hassabis D, Hadsell R, Kumaran D. Vector-based navigation using grid-like representations in artificial agents. Nature. 2018 May;557(7705):429-433. doi: 10.1038/s41586-018-0102-6. Epub 2018 May 9. PMID: 29743670.

Barczak A, Haegens S, Ross DA, McGinnis T, Lakatos P, Schroeder CE. Dynamic Modulation of Cortical Excitability during Visual Active Sensing. Cell Rep. 2019 Jun 18;27(12):3447-3459.e3. doi: 10.1016/j.celrep.2019.05.072. PMID: 31216467; PMCID: PMC6598687.

Becht E, McInnes L, Healy J, Dutertre CA, Kwok IWH, Ng LG, Ginhoux F, Newell EW. Dimensionality reduction for visualizing single-cell data using UMAP. Nat Biotechnol. 2018 Dec 3. doi: 10.1038/nbt.4314. Epub ahead of print. PMID: 30531897.

Benjamin AS, Fernandes HL, Tomlinson T, Ramkumar P, VerSteeg C, Chowdhury RH, Miller LE, Kording KP. (2018). Modern Machine Learning as a Benchmark for Fitting Neural Responses. Front Comput Neurosci. doi: 10.3389/fncom.2018.00056. PMID: 30072887; PMCID: PMC60602699

Britten, K. H. (2008). Mechanisms of self-motion perception. Annual review of neuroscience 31, 389–410

Cameron, A.C. Windmeijer, F.A. (1997). An R-squared measure of goodness of fit for some common nonlinear regression models. Journal of Econometrics 77, 329-342.

Chadick JZ, Gazzaley A. Differential coupling of visual cortex with default or frontal-parietal network based on goals. Nature Neuroscience. 14: 830-2. PMID 21623362 DOI: 10.1038/nn.2823

Chafee, M.V. & Goldman-Rakic, P.S. Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. J. Neurophysiol. 79, 2919–2940 (1998).

Chettih SN, Harvey CD. Single-neuron perturbations reveal feature-specific competition in V1. Nature. 2019 Mar;567(7748):334-340. doi: 10.1038/s41586-019-0997-6. Epub 2019 Mar 6. PMID: 30842660; PMCID: PMC6682407.

Cisek P, Kalaska JF. Neural mechanisms for interacting with a world full of action choices. Annu Rev Neurosci. 2010;33:269-98. doi: 10.1146/annurev.neuro.051508.135409.

Cole MW, Reynolds JR, Power JD, Repovs G, Anticevic A, Braver TS. Multi-task connectivity reveals flexible hubs for adaptive task control. Nat Neurosci. 2013 Sep;16(9):1348-55. doi: 10.1038/nn.3470. Epub 2013 Jul 28. PMID: 23892552; PMCID: PMC3758404.

Cossell L, Iacaruso MF, Muir DR, Houlton R, Sader EN, Ko H, Hofer SB, Mrsic-Flogel TD. Functional organization of excitatory synaptic strength in primary visual cortex. Nature. 2015 Feb 19;518(7539):399-403. doi: 10.1038/nature14182. Epub 2015 Feb 4. PMID: 25652823; PMCID: PMC4843963.

Domencich, T. A., McFadden. (1975). Urban travel demand-a behavioral analysis. A behavioral analysis. North-Holland, Amsterdam.

Dormann CF, Elith J, Bacher S, et al. (2013). Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36: 27-46.

Driscoll LN, Pettit NL, Minderer M, Chettih SN, Harvey CD. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. Cell. 2017 Aug 24;170(5):986-999.e16. doi: 10.1016/j.cell.2017.07.021. Epub 2017 Aug 17. PMID: 28823559; PMCID: PMC5718200.

Duncan J. An adaptive coding model of neural function in prefrontal cortex. Nat Rev Neurosci. 2001 Nov;2(11):820-9. doi: 10.1038/35097575. PMID: 11715058.

Ekman, M., Derrfuss, J., Tittgemeyer, M., & Fiebach, C. J. (2012). Predicting errors from reconfiguration patterns in human brain networks. Proceedings of the National Academy of Sciences of the United States of America, 109(41), 16714–16719. https://doi.org/10.1073/pnas.1207523109

Ekstrom AD, Kahana MJ, Caplan JB, Fields TA, Isham EA, Newman EL, Fried I. Cellular networks underlying human spatial navigation. Nature. 2003 Sep 11;425(6954):184-8. doi: 10.1038/nature01964. PMID: 12968182.

Ekstrom, A., Spiers, H., Bohbot, V., and Rosenbaum, R. (2018). Human spatial navigation (Princeton University Press).

Ester, Martin, et al. "Density-Based spatial clustering of applications with noise (DBSCAN)." Proc. of the Second International Conference on Knowledge Discovery and Data Mining. 1996.

Gauthier JL, Tank DW. A Dedicated Population for Reward Coding in the Hippocampus. Neuron. 2018 Jul 11;99(1):179-193.e7. doi: 10.1016/j.neuron.2018.06.008. Epub 2018 Jun 28. PMID: 30008297; PMCID: PMC7023678.

Gothard KM, Skaggs WE, McNaughton BL. Dynamics of mismatch correction in the hippocampal ensemble code for space: interaction between path integration and environmental cues. J Neurosci. 1996 Dec 15;16(24):8027-40. doi: 10.1523/JNEUROSCI.16-24-08027.1996. PMID: 8987829; PMCID: PMC6579211.

Gomez-Marin A, Paton JJ, Kampff AR, Costa RM, Mainen ZF. Big behavioral data: psychology, ethology and the foundations of neuroscience. Nat Neurosci. 2014 Nov;17(11):1455-62. doi: 10.1038/nn.3812. Epub 2014 Oct 28. PMID: 25349912.

Gulli, R.A., Duong, L.R., Corrigan, B.W., Doucet, G., Williams, S. Fusi, S., Martinez-Trujillo, J.C. (2020). Context-dependent representations of objects and space in the primate hippocampus during virtual navigation. Nature Neuroscience. 23, 103–112

Hafting T, Fyhn M, Molden S, Moser M-B, Moser EI 2005. Microstructure of a spatial map in the entorhinal cortex. Nature 436: 801–806

Howard LR, Javadi AH, Yu Y, Mill RD, Morrison LC, Knight R, Loftus MM, Staskute L, Spiers HJ. The hippocampus and entorhinal cortex encode the path and Euclidean distances to goals during navigation. Curr Biol. 2014 Jun 16;24(12):1331-1340. doi: 10.1016/j.cub.2014.05.001. Epub 2014 Jun 5. PMID: 24909328; PMCID: PMC4062938.

Ilg and Thier 2003 Ilg UJ and Thier P. Visual tracking neurons in primate area MST are activated by smooth-pursuit eye movements of an "imaginary" target. J Neurophysiol 90: 1489–1502, 2003.

Ilg UJ, Thier P. Eye movements of rhesus monkeys directed towards imaginary targets. Vision Res. 1999 Jun;39(12):2143-50. doi: 10.1016/s0042-6989(98)00321-6. PMID: 10343796.

Jun EJ, Bautista AR, Nunez MD, Allen DC, Tak JH, Alvarez E, Basso MA. Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. Nat Neurosci. 2021 Jun 28. doi: 10.1038/s41593-021-00878-6. Epub ahead of print. PMID: 34183869.

Jutras, M.J., Fries, P., Buffalo, E.A. (2013). Oscillatory activity in the monkey hippocampus during visual exploration and memory formation. Proc. Natl. Acad. Sci. 110, 13144– 13149

Kanitscheider I, Fiete I (2017) Emergence of dynamically reconfigurable hippocampal responses by learning to perform probabilistic spatial reasoning. bioRxiv 231159. doi: 10.1101/231159.

Katsuki F, Constantinidis C. Unique and shared roles of the posterior parietal and dorsolateral prefrontal cortex in cognitive functions. Front Integr Neurosci. 2012 May 3;6:17. doi: 10.3389/fnint.2012.00017. PMID: 22563310; PMCID: PMC3342558.

Kravitz DJ, Saleem KS, Baker CI, Mishkin M. A new neural framework for visuospatial processing. Nat Rev Neurosci. 2011 Apr;12(4):217-30. doi: 10.1038/nrn3008. PMID: 21415848; PMCID: PMC3388718.

Lakatos P., Karmos G., Mehta A. D., Ulbert I., Schroeder C. E. (2008). Entrainment of neuronal oscillations as a mechanism of attentional selection. Science 320, 110–113. 10.1126/science.1154735

Lakshminarasimhan KJ, Avila E, Neyhart E, DeAngelis GC, Pitkow X, Angelaki DE. Tracking the Mind's Eye: Primate Gaze Behavior during Virtual Visuomotor Navigation Reflects Belief Dynamics. Neuron. 2020 May 20;106(4):662-674.e5. doi: 10.1016/j.neuron.2020.02.023.

Lakshminarasimhan KJ, Petsalis M, Park H, DeAngelis GC, Pitkow X, Angelaki DE. A Dynamic Bayesian Observer Model Reveals Origins of Bias in Visual Path Integration. Neuron. 2018 Jul 11;99(1):194-206.e5. doi: 10.1016/j.neuron.2018.05.040.

Leszczynski, M., Staudigl, T., Chaieb, L., Enkirch, S.J., Fell, J., Schroeder, C.E. (2020). Eye movements modulate neural activity in the human anterior thalamus during visual active sensing. bioRxiv 2020.03.30.015628; doi: https://doi.org/10.1101/2020.03.30.015628

Levi, A.J., Zhao, Y., Park, I.M. Huk, A.C. (2021). Sensory and choice responses in MT distinct from motion encoding. bioRxiv 2021.06.24.449836; doi: https://doi.org/10.1101/2021.06.24.449836

Masse NY, Hodnefield JM, Freedman DJ. Mnemonic Encoding and Cortical Organization in Parietal and Prefrontal Cortices. J Neurosci. 2017 Jun 21;37(25):6098-6112. doi: 10.1523/JNEUROSCI.3903-16.2017. Epub 2017 May 24. PMID: 28539423; PMCID: PMC5481944.

McInnes, Leland, John Healy, and James Melville. "UMAP: uniform manifold approximation and projection for dimension reduction." (2020).

Medendorp WP, Heed T. State estimation in posterior parietal cortex: Distinct poles of environmental and bodily states. Prog Neurobiol. 2019 Dec;183:101691. doi: 10.1016/j.pneurobio.2019.101691. Epub 2019 Sep 6. PMID: 31499087.

Mendoza-Halliday D, Torres S, Martinez-Trujillo JC. Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. Nat Neurosci. 2014 Sep;17(9):1255-62. doi: 10.1038/nn.3785. Epub 2014 Aug 10. PMID: 25108910; PMCID: PMC4978542

Minderer M, Brown KD, Harvey CD. The Spatial Structure of Neural Encoding in Mouse Posterior Cortex during Navigation. Neuron. 2019 Apr 3;102(1):232-248.e11. doi: 10.1016/j.neuron.2019.01.029. Epub 2019 Feb 13. PMID: 30772081; PMCID: PMC6642748.

Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A. Information-limiting correlations. Nat Neurosci. 2014 Oct;17(10):1410-7. doi: 10.1038/nn.3807. Epub 2014 Sep 7. PMID: 25195105; PMCID: PMC4486057.

Nauhaus I, Busse L, Carandini M, Ringach DL. Stimulus contrast modulates functional connectivity in visual cortex. Nat Neurosci. 2009 Jan;12(1):70-6. doi: 10.1038/nn.2232. Epub 2008 Nov 23. PMID: 19029885; PMCID: PMC2610236.

Newsome WT, Paré EB. A selective impairment of motion perception following lesions of the middle temporal visual area (MT). J Neurosci. 1988 Jun;8(6):2201-11. doi: 10.1523/JNEUROSCI.08-06-02201.1988. PMID: 3385495; PMCID: PMC6569328.

Noel JP, Caziot B, Bruni S, Fitzgerald NE, Avila E, Angelaki DE. (2021) Supporting generalization in non-human primate behavior by tapping into structural knowledge: Examples from sensorimotor mappings, inference, and decision-making. Prog Neurobiol. 2021 Jan 14:101996. doi: 10.1016/j.pneurobio.2021.101996.

Noel JP, Lakshminarasimhan KJ, Park H, Angelaki DE. Increased variability but intact integration during visual navigation in Autism Spectrum Disorder. Proc Natl Acad Sci U S A. 2020 May 19;117(20):11158-11166. doi: 10.1073/pnas.2000216117. Epub 2020 May 1. PMID: 32358192; PMCID: PMC7245105.

O'Keefe J, Burgess N. Geometric determinants of the place fields of hippocampal neurons. Nature. 1996 May 30;381(6581):425-8. doi: 10.1038/381425a0. PMID: 8632799.

O'Keefe J, Dostrovsky J 1971. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res 34: 171–175

O'Keefe J, Nadel L 1978. The hippocampus as a cognitive map. Clarendon, Oxford.

Okun M, Steinmetz N, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrsic-Flogel TD, Carandini M, Harris KD. Diverse coupling of neurons to populations in sensory cortex. Nature. 2015 May 28;521(7553):511-515. doi: 10.1038/nature14273. Epub 2015 Apr 6. PMID: 25849776; PMCID: PMC4449271.

Pachitariu M, Steinmetz M, Kadir M, Carandini M, Harris M. 2016. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv. DOI: https://doi.org/10.1101/061481

Pfeiffer BE, Foster DJ. Hippocampal place-cell sequences depict future paths to remembered goals. Nature. 2013 May 2;497(7447):74-9. doi: 10.1038/nature12112. Epub 2013 Apr 17. PMID: 23594744; PMCID: PMC3990408.

Pitkow X, Angelaki DE. Inference in the Brain: Statistics Flowing in Redundant Population Codes. Neuron. 2017 Jun 7;94(5):943-953. doi: 10.1016/j.neuron.2017.05.028

Rajkai, C., Lakatos, P., Chen, C. M., Pincze, Z., Karmos, G., & Schroeder, C. E. (2008). Transient cortical excitation at the onset of visual fixation. Cerebral Cortex, 18(1), 200–209

Rozzi S, Calzavara R, Belmalih A, Borra E, Gregoriou GG, Matelli M, Luppino G. Cortical connections of the inferior parietal cortical convexity of the macaque monkey. Cereb Cortex. 2006;16:1389–1417. doi: 10.1093/cercor/bhj076

Sarel A, Finkelstein A, Las L, Ulanovsky N. Vectorial representation of spatial goals in the hippocampus of bats. Science. 2017 Jan 13;355(6321):176-180. doi: 10.1126/science.aak9589.

Schroeder CE, Wilson DA, Radman T, Scharfman H, Lakatos P. Dynamics of Active Sensing and perceptual selection. Curr Opin Neurobiol. 2010 Apr;20(2):172-6. doi: 10.1016/j.conb.2010.02.010. Epub 2010 Mar 20. PMID: 20307966; PMCID: PMC2963579.

Serino A. Peripersonal space (PPS) as a multisensory interface between the individual and the environment, defining the space of the self. Neurosci Biobehav Rev. 2019 Apr;99:138-159. doi: 10.1016/j.neubiorev.2019.01.016. Epub 2019 Jan 24. PMID: 30685486.

Shi, J., and J. Malik. "Normalized cuts and image segmentation." IEEE Transactions on Pattern Analysis and Machine Intelligence. Vol. 22, 2000, pp. 888–905.

Sosa, M., Giocomo, L.M. Navigating for reward. (2021). Nat Rev Neurosci . https://doi.org/10.1038/s41583-021-00479-z

Strange, B. A., Witter, M. P., Lein, E. S., and Moser, E. I. (2014). Functional organization of the hippocampal longitudinal axis. Nat. Rev. Neurosci. 15, 655–669. doi: 10.1038/nrn3785

Suzuki, M. & Gottlieb, J. Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. Nat. Neurosci. 16, 98–104 (2013).

Swaminathan, S.K. & Freedman, D.J. Preferential encoding of visual categories in parietal cortex compared with prefrontal cortex. Nat. Neurosci. 15, 315–320 (2012).

Whitlock JR, Pfuhl G, Dagslott N, Moser MB, Moser EI. Functional split between parietal and entorhinal cortices in the rat. Neuron. 2012 Feb 23;73(4):789-802. doi: 10.1016/j.neuron.2011.12.028. PMID: 22365551.

Whitlock JR, Sutherland RJ, Witter MP, Moser MB, Moser EI. Navigating from hippocampus to parietal cortex. Proc Natl Acad Sci U S A. 2008 Sep 30;105(39):14755-62. doi: 10.1073/pnas.0804216105. Epub 2008 Sep 23. PMID: 18812502; PMCID: PMC2567440.

Whitlock JR. Posterior parietal cortex. Curr Biol. 2017 Jul 24;27(14):R691-R695. doi: 10.1016/j.cub.2017.06.007. PMID: 28743011.

Wirth, S., Baraduc, P., Planté, A., Pinède, S., Duhamel, J.-R. (2017). Gaze-informed, task-situated representation of space in primate hippocampus during virtual navigation. PLOS Biology. 15, e2001045

Yang, S. C. H., Wolpert, D. M. & Lengyel, M. Theoretical perspectives on active sensing. Current Opinion in Behavioral Sciences vol. 11 100–108 (2016).

Yoo SBM, Hayden BY, Pearson JM. Continuous decisions. Philos Trans R Soc Lond B Biol Sci. 2021 Mar;376(1819):20190664. doi: 10.1098/rstb.2019.0664. Epub 2021 Jan 11. PMID: 33423634; PMCID: PMC7815426.

Zanto TP, Gazzaley A. Fronto-parietal network: flexible hub of cognitive control. Trends in Cognitive Sciences. 17: 602-3. PMID 24129332 DOI: 10.1016/j.tics.2013.10.001

Zhu, S., Lakshminarasimhan, K., Arfaei, N., Angelaki, D.(2021). Eye Movements Reveal Spatiotemporal Dynamics of Active Sensing and Planning in Navigation. bioRxiv 2021.04.26.441482; doi: https://doi.org/10.1101/2021.04.26.441482
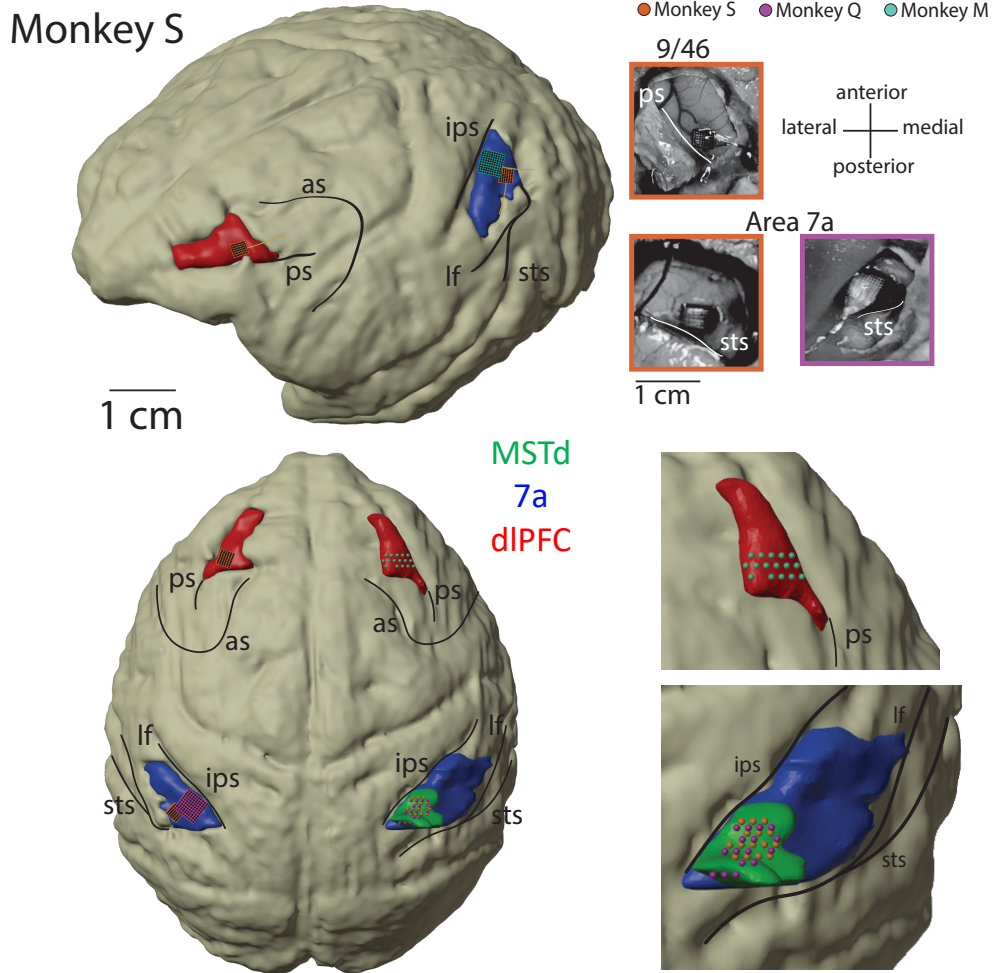
**Figure S1.** *Magnetic resonance imaging (MRI) reconstruction of recording sites and pictures from Utah array implants. 3D rendering of the brain is from monkey S, and all recording sites have been placed on the common reference (MSTd in green, area 7a in blue, dlPFC in red). Location of acute recordings are indicated by a sphere, color coded per animal (Monkey S, orange; Monkey Q, purple; Monkey M in green). Location of Utah arrays are indicated by squares. The pictures are framed in the color corresponding to the monkey. AS, arcuate sulcus; IPS, intraparietal sulcus; PS, principal sulcus; STS, superior temporal sulcus; LF, lateral fissure.*
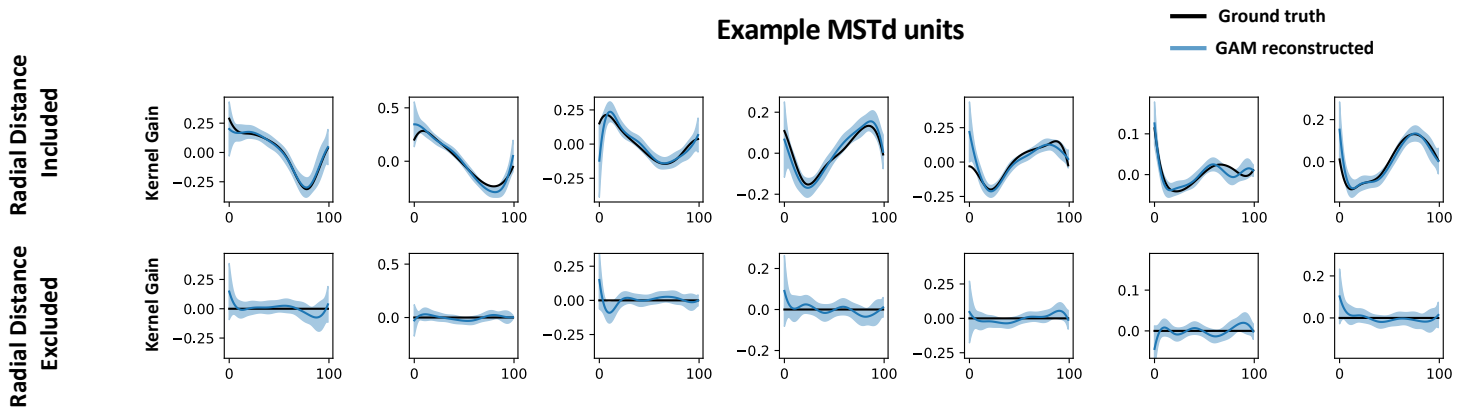
***Figure S2.*** *We generated synthetic spike trains according to real input statistics (e.g., linear and angular velocity profiles, correlated inputs) and tuning functions derived in MSTd by the P-GAM (all 17 variables, as in the main text). Next, we attempted to recover these tuning functions, while including or excluding (i.e., zeroing) the tuning to radial distance to target. Across all simulation (n = 1344), the P-GAM never excluded distance to target from its selected model when this variable had not been zeroed (0.0% false negative), and only once failed to exclude this variable from its final selection when the variable had been zeroed (0.0007% false positive). Top row are significant tuning functions to radial distance in MSTd (black) and their recovered shape by the P-GAM (blue) when not excluded. Bottom row are the same units, when radial distance (but not other variables) were zeroed. Columns show 7 example cells.*
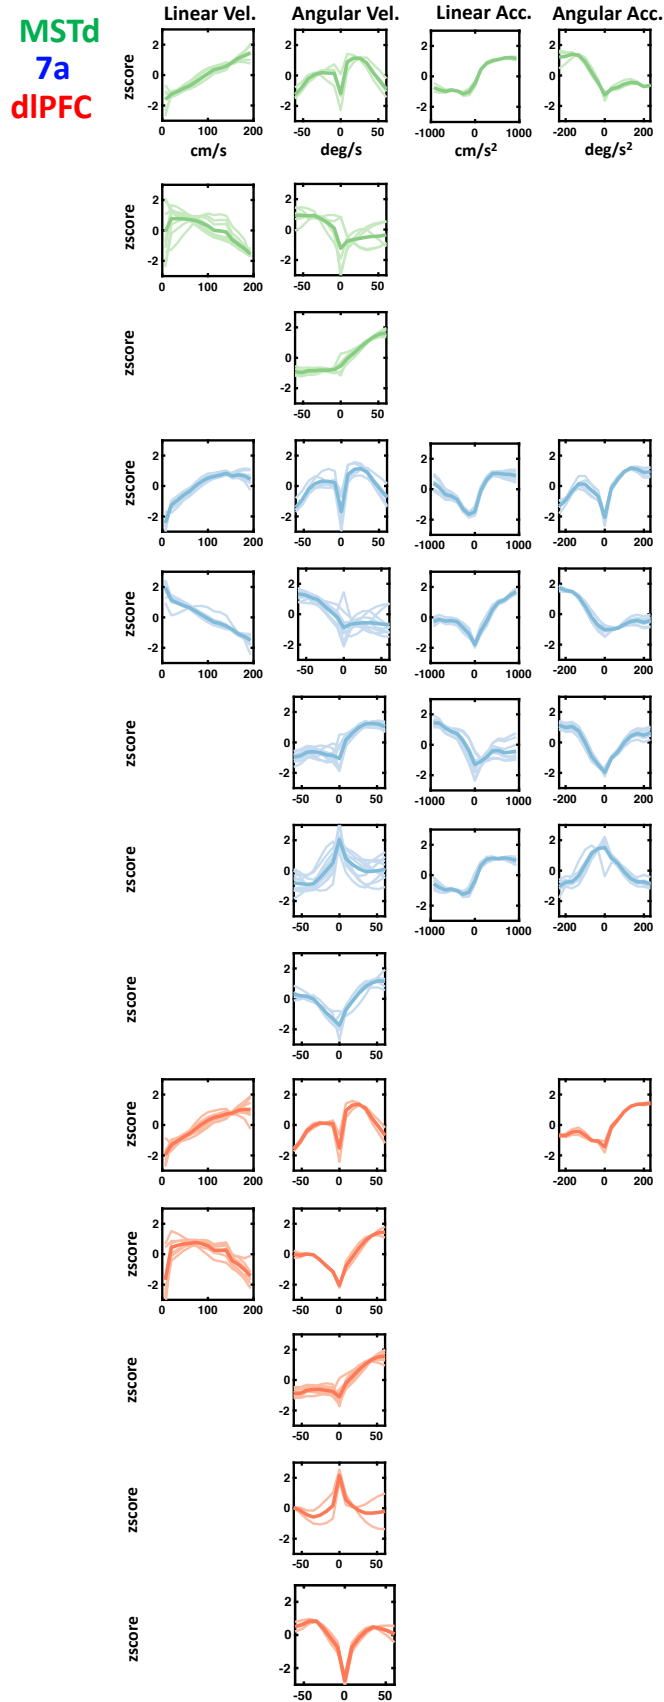
**Figure S3.** Sensorimotor tuning functions to continuous variables (linear and angular velocity and acceleration) encountered in MSTd (green), area 7a (blue), and dlPFC (red). Tuning function clusters were determined by DBSCAN. For each cluster, ten example tuning functions are plotted (semi-transparent and background), as well as the average of all tuning functions within the given cluster type (opaque and thicker line). Y-axis are firing rates (in Hz) z-scored and X-axis spans the state space of the particular variable, in cm/s for linear velocity, deg/s for angular velocity, cm/s$^2$ for linear acceleration, and deg/ s$^2$ for angular acceleration (see top row).
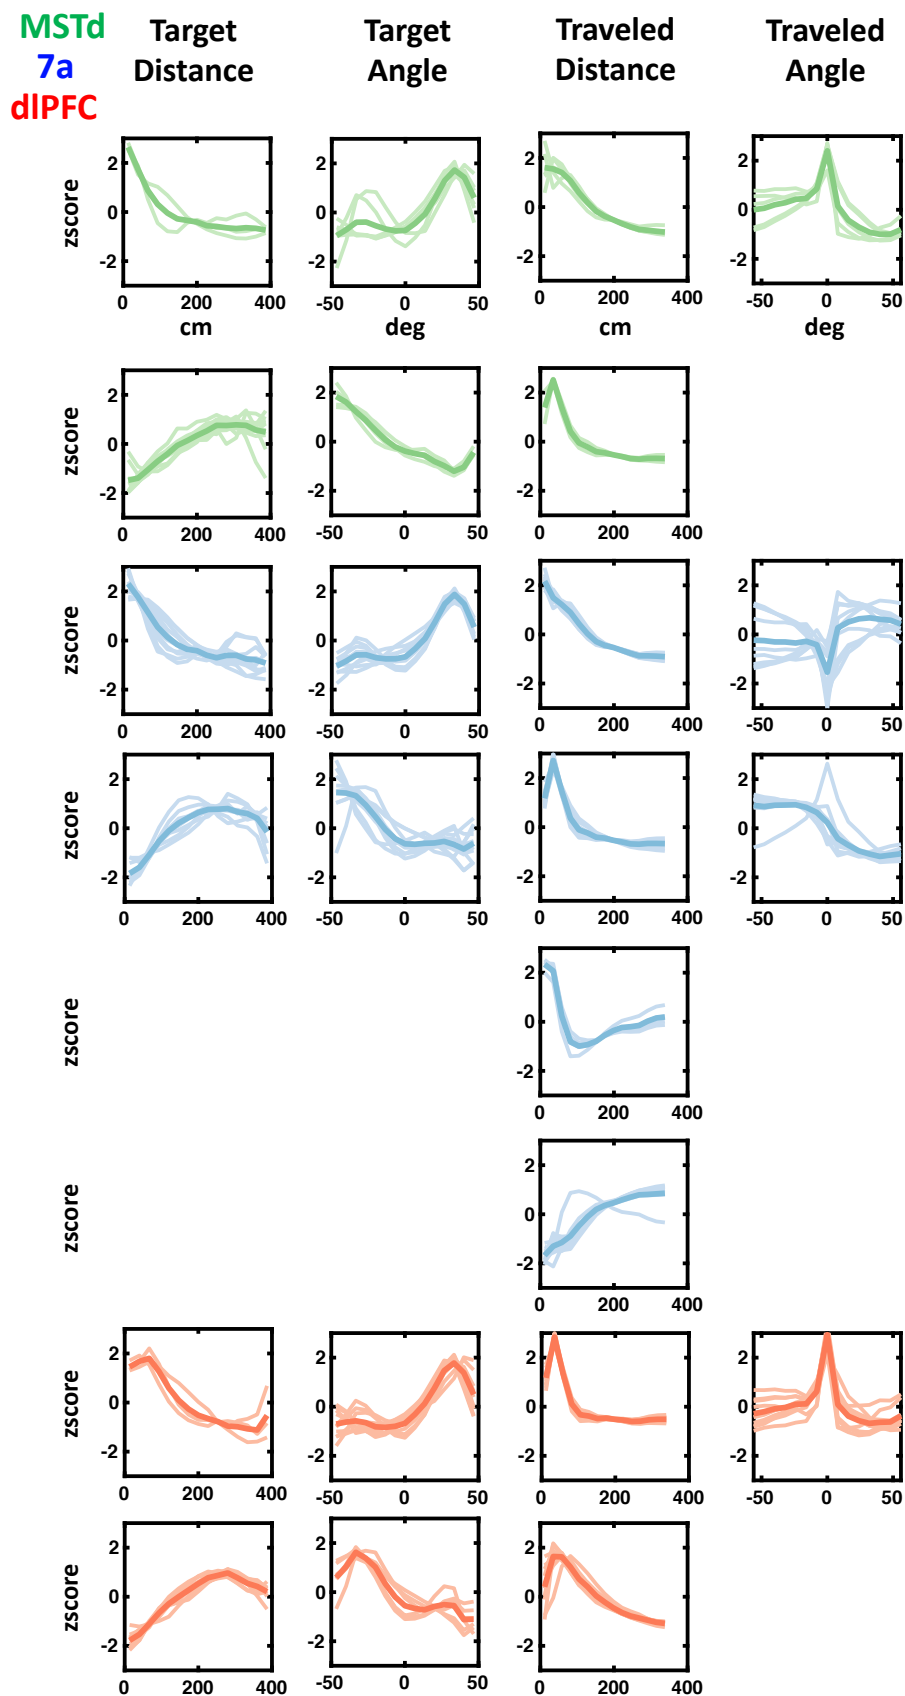
**Figure S4.** *Latent tuning functions (radial and angular distance to target and from origin) encountered in MSTd (green), area 7a (blue), and dlPFC (red). Tuning function clusters were determined by DBSCAN. As for Figure S3, ten example tuning functions are plotted per cluster (in semi-transparent), as well as the average of all tuning functions within the given cluster type (opaque and thicker line). Y-axis are firing rates (in Hz) z-scored and X-axis spans the state space of the particular variable, in cm for radial distances and degrees for angles (see top row).*
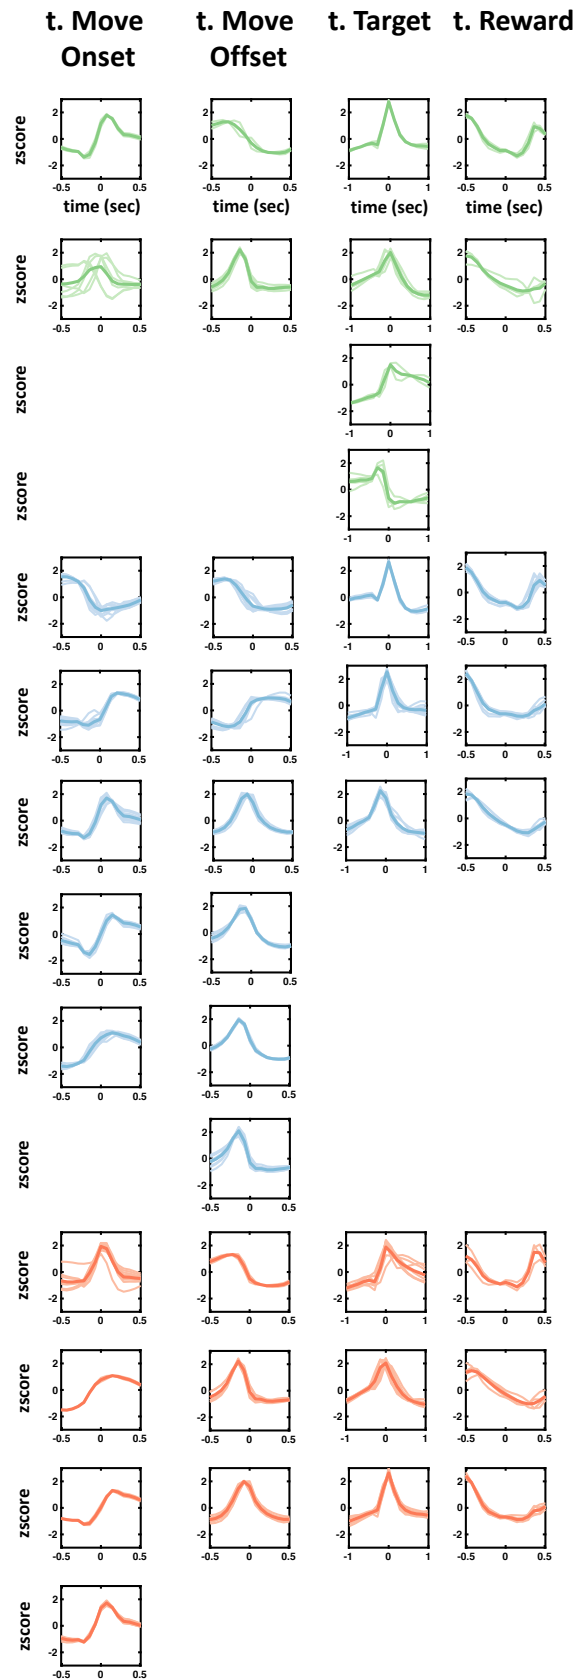
**Figure S5.** *Temporal kernels (time of movement onset and offset, time of target onset and reward presentation) encountered in MSTd (green), area 7a (blue), and dlPFC (red). Tuning function clusters were determined by DBSCAN. Ten example tuning functions are plotted per cluster (transparent), as well as the average of all tuning functions within the given cluster type (opaque and thicker line). Y-axis are firing rates (in Hz) z-scored and X-axis is time relative to the discrete event, in seconds (see top row).*
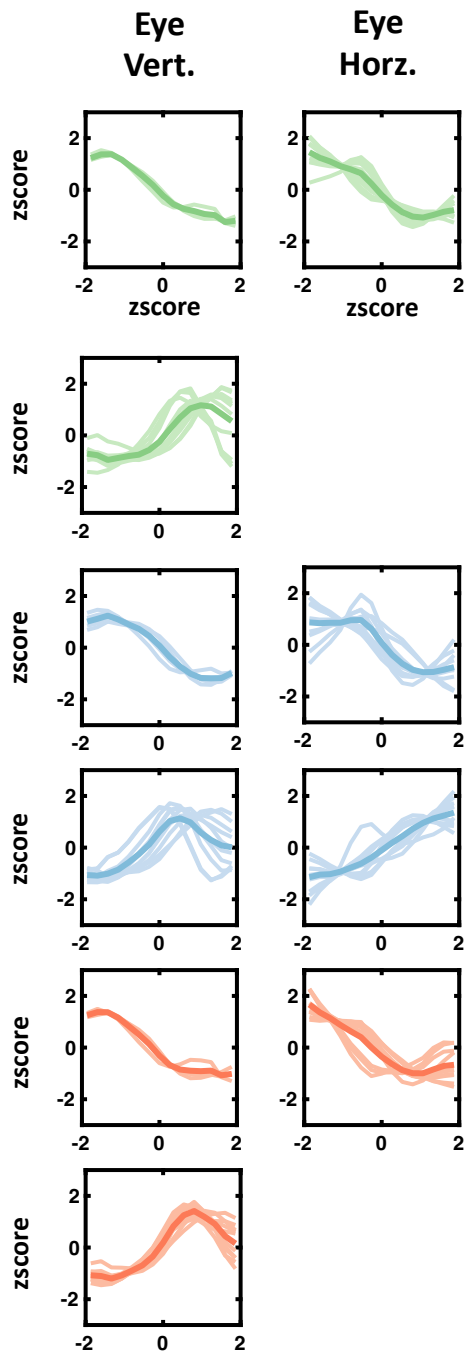
**Figure S6.** *Tuning functions to eye position variables (vertical and horizontal) encountered in MSTd (green), area 7a (blue), and dlPFC (red). Tuning function clusters were determined by DBSCAN. Ten example tuning functions are plotted per cluster (semi-transparent), as well as the average of all tuning functions within the given cluster type (opaque and thicker line). Y-axis are firing rates (in Hz) z-scored. and X-axis eye position, also in Z-score.*

**Figure S7.** *UMAP (McInnes et al., 2020) projection space color coded by mutual information with each task variable. Each subplot is normalized between its minimum and maximum mutual information (MI). Results demonstrate that the UMAP projection finds meaningful clusters – neurons with high MI for a particular variable clustering together - even though the algorithm is not know about task variables or MI. The plots also highlight the distribution in MI for the different variables; e.g., most neurons having an intermediary MI for angular acceleration, while a subset of neurons have a very strong MI for the timing of reward.*
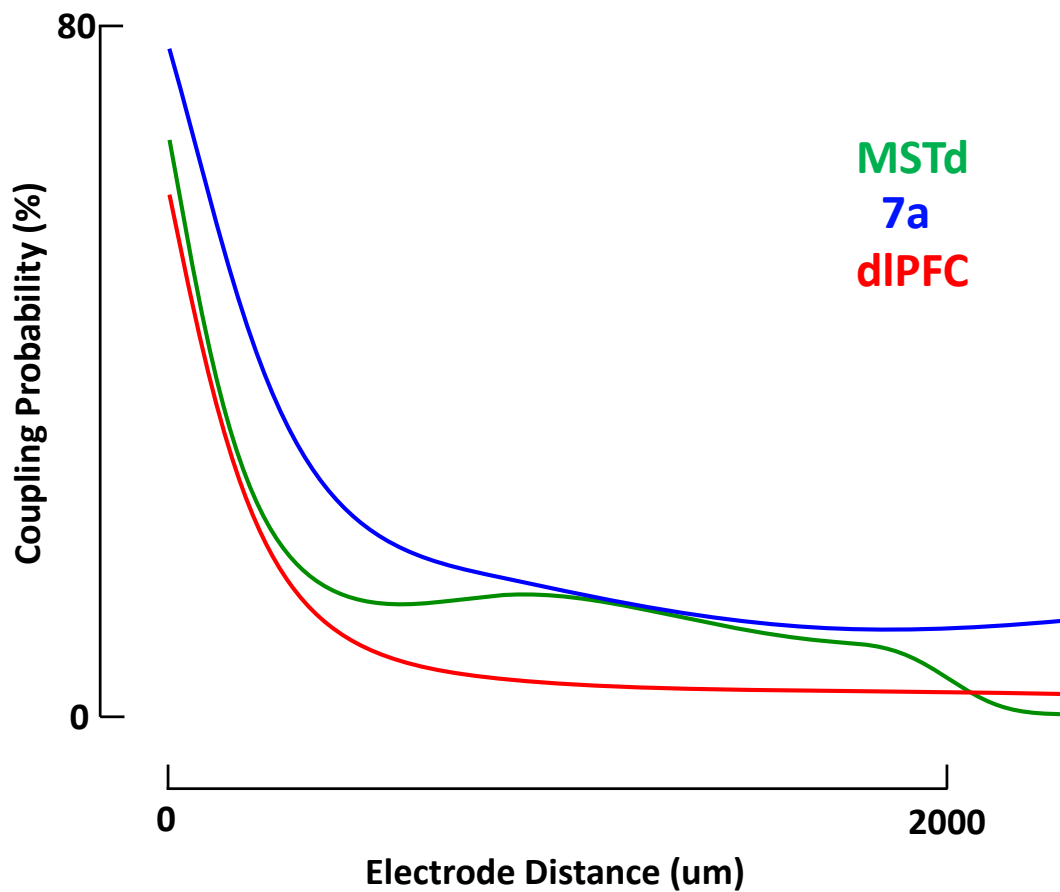
**Figure S8.** *Coupling probability (y-axis) between two neurons as a function of the distance between them (x-axis). This illustration is for data from a single monkey (Monkey S) with recordings in MSTd (green, number of coupling pairs = 1714), area 7a (blue, number of coupling pairs = 97711), and dlPFC (red, number of coupling pairs = 41665). As expected, neurons that are closer to each other are more likely to be coupled. Given this effect and that multiple recording techniques were used (with different spacing between electrodes, 400um in Utah array and 100um in linear probes), we used these estimates in correcting for coupling probability for a single distance (500um).*
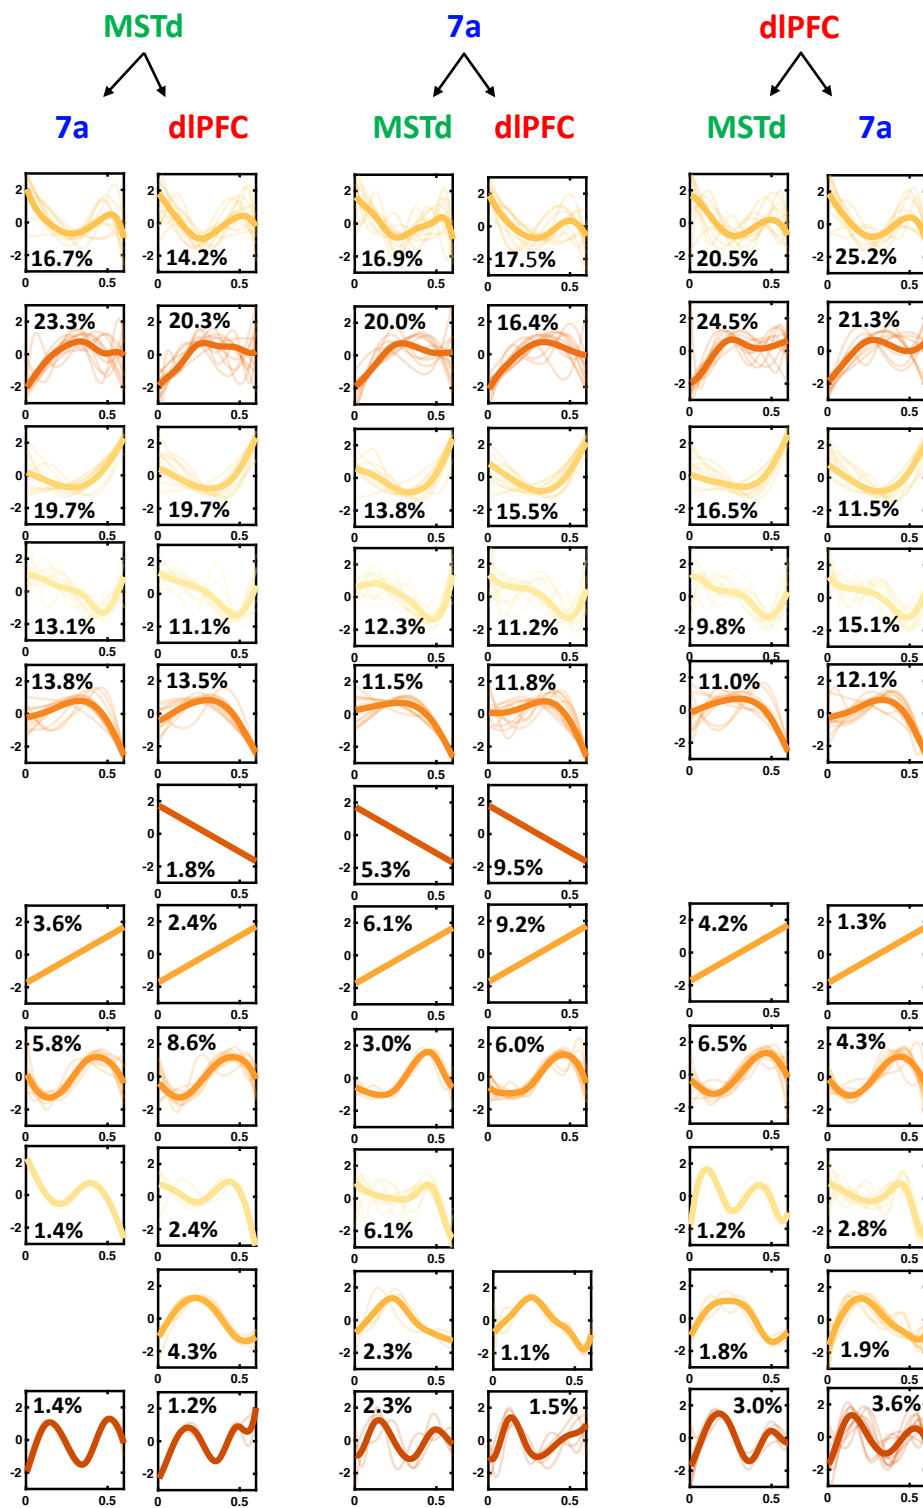
***Figure S9.*** *Coupling filters across cortical areas. Coupling functions were clustered by DBSCAN and are depictured from each of the possible emitters (top) to each of the possible receivers (bottom). The coupling filters are ordered from top to bottom, according to their maximal occurrence for any pair of emitter and receiver. The percentage within each box indicates what percent of coupling filters, for the particular emitter-receiver pair, that were of the particular kind. Ten examples (if available) are plotted in semi-transparent and in the background, and the average across all coupling filters of the particular type are plotted in opaque and with a thicker line. Percentages refer to the percent of coupling filters between an emitter and a receiver that were clustered within the given category. X-axis is a z-score of firing rate (in Hz) and Y-axis is time in seconds (maximum = 600ms). Importantly, these coupling filters are not sinusoidal.*
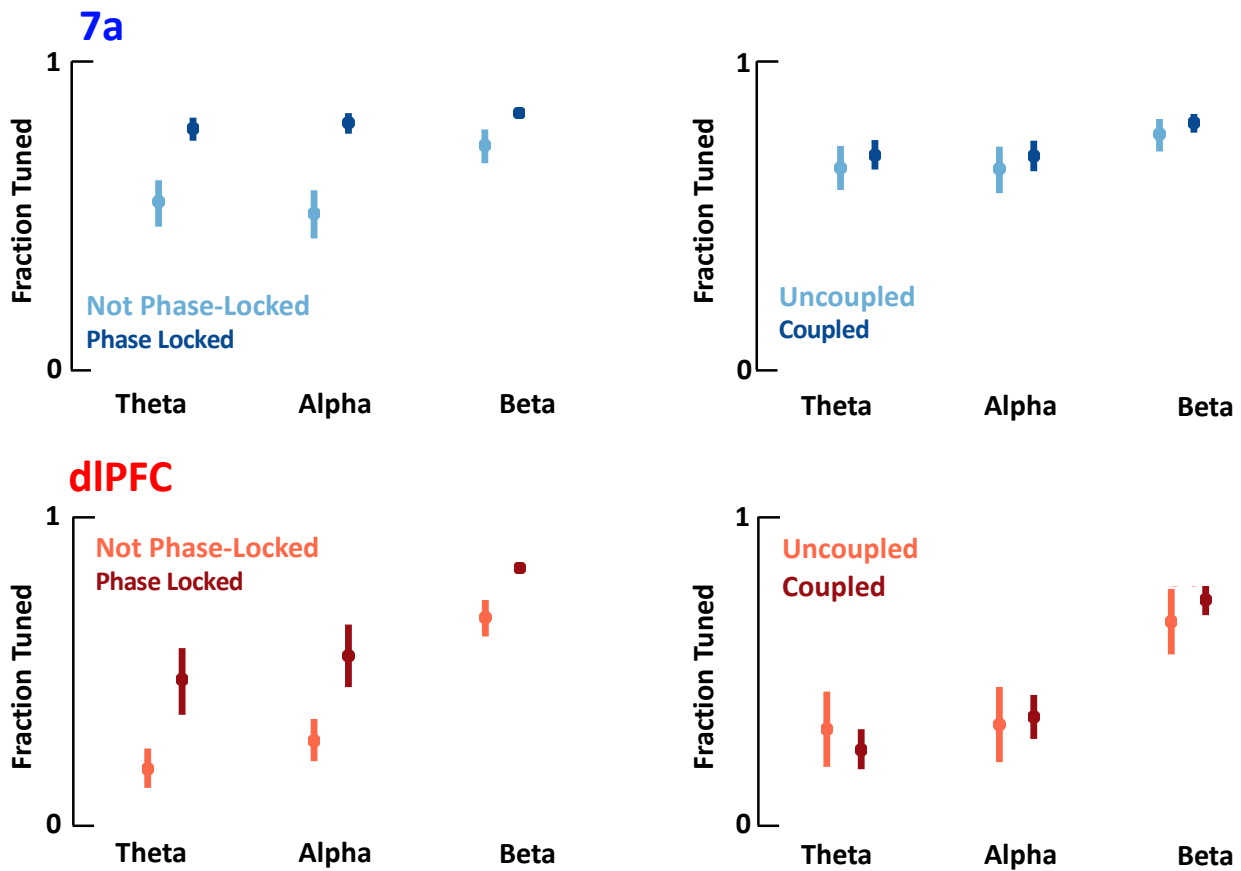
***Figure S10.*** *Fraction of neurons tuned (y-axis) to the phase of theta, alpha, and beta-band LFPs (x-axis) as a function of brain area (top = 7a, bottom = dlPFC) and whether or not the unit was phase-locked to LFP in the other area (left column) or coupled to a neuron in the other area (right column). As described in the main text, overall neurons in 7a were more likely to spike in phase with its own LFP phases. Neurons who were tuned to the LFP in a different area (dark blue), were significantly more likely to be tuned to the LFP in their own area (left column) than neurons that were not tuned to the phase of LFP in another area (light blue). On the other hand, whether units were coupled or not with neurons in another area (uncoupled in light blue and coupled in dark blue) did not impact their likelihood of being tuned to the ongoing phase of LFPs in their own area (right column). MSTd is not depicted, as there was no phase locking of dlPFC neurons to the ongoing LFP phase in MSTd. This analysis is based on 5929 coupling pairs in area 7a (blue) and 4190 coupling pairs in dlPFC (red).*
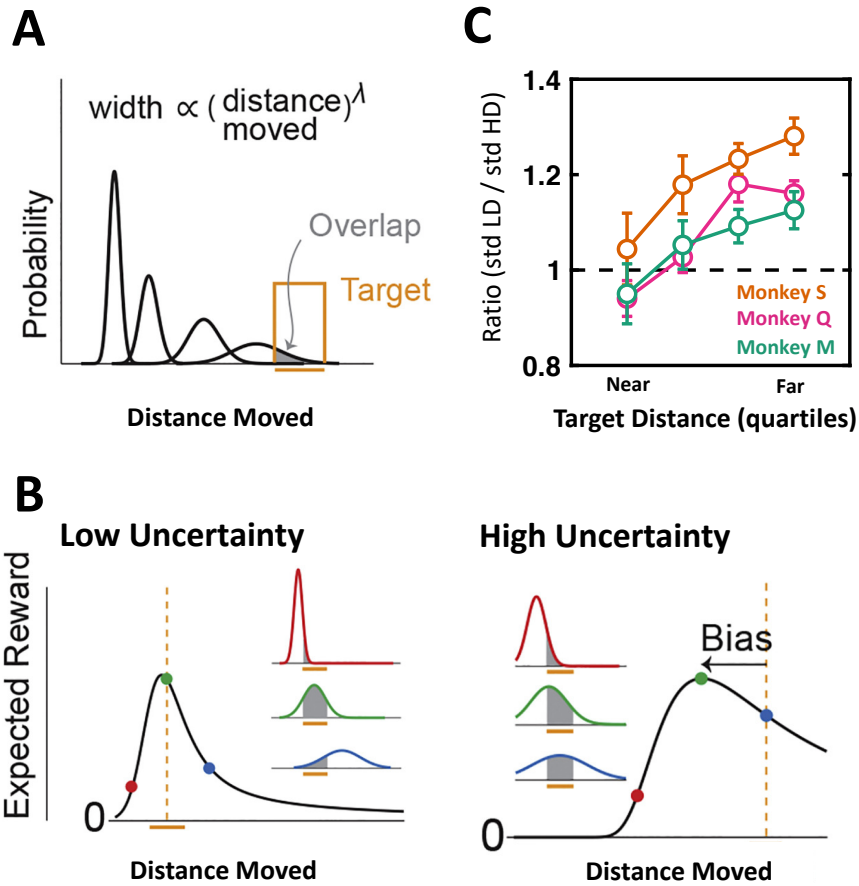
**Figure S11.** *Contrast of the behavioral uncertainty across high- and low-density optic flow conditions. A. Uncertainty over position is taken to increase as a power law with distance moved, and observers stop when the expected reward (the area under the curve between the overlap in self and target position) peaks (Lakshminarasimhan et al., 2018). B. Under this scenario, when uncertainty increases (contrast insets, red, green, blue), the expected reward peaks before reaching the target position. C. To establish whether greater undershooting in the low optic flow density condition, may be due to the effect postulated in A. and B. (and formally modeled in Lakshminarasimhan et al., 2018), we estimate the scaling of uncertainty with position in the low- and high-density conditions. Target distances were discretized in quartiles, and the standard deviation within each quartile was computed. The ratio between these standard deviations were computed (low density / high density), such that a value of '1' indicates equal uncertainty in both conditions. For all three monkeys (S, Q, and M, respectively orange, pink, and green), this ratio increased with distance, indicating that the scaling of uncertainty with distance is greater under the low-density conditions vis-à-vis the high-density condition. In turn, we would expect greater undershooting in the low-density condition.*
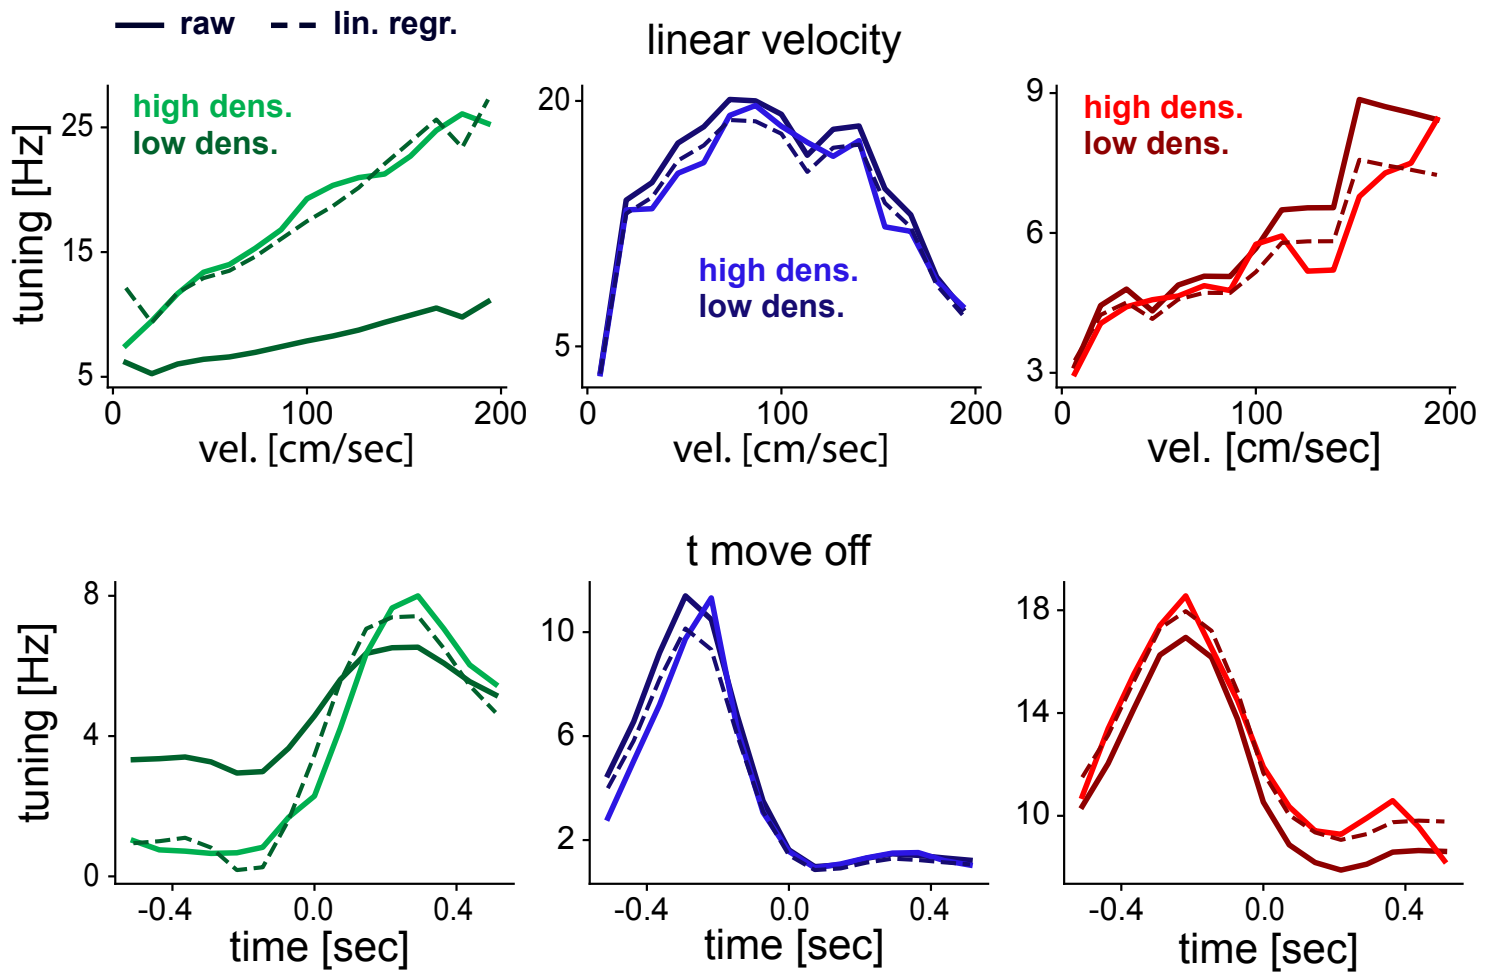
***Figure S12.*** *An example tuning function to linear velocity (top) and time from movement offset (bottom) for each neural area (Green = MSTd; Blue = Area 7a; Red = dlPFC) during high and (brighter color) low (darker color)) optic flow density conditions. In addition to the raw data (solid lines), we also demonstrate the re-scaling from low- and high-density condition (dashed lines) in order to estimate gain modulations. The examples are representative, in that they demonstrate no gain modulation in 7a and dlPFC, but a strong modulation in MSTd.*
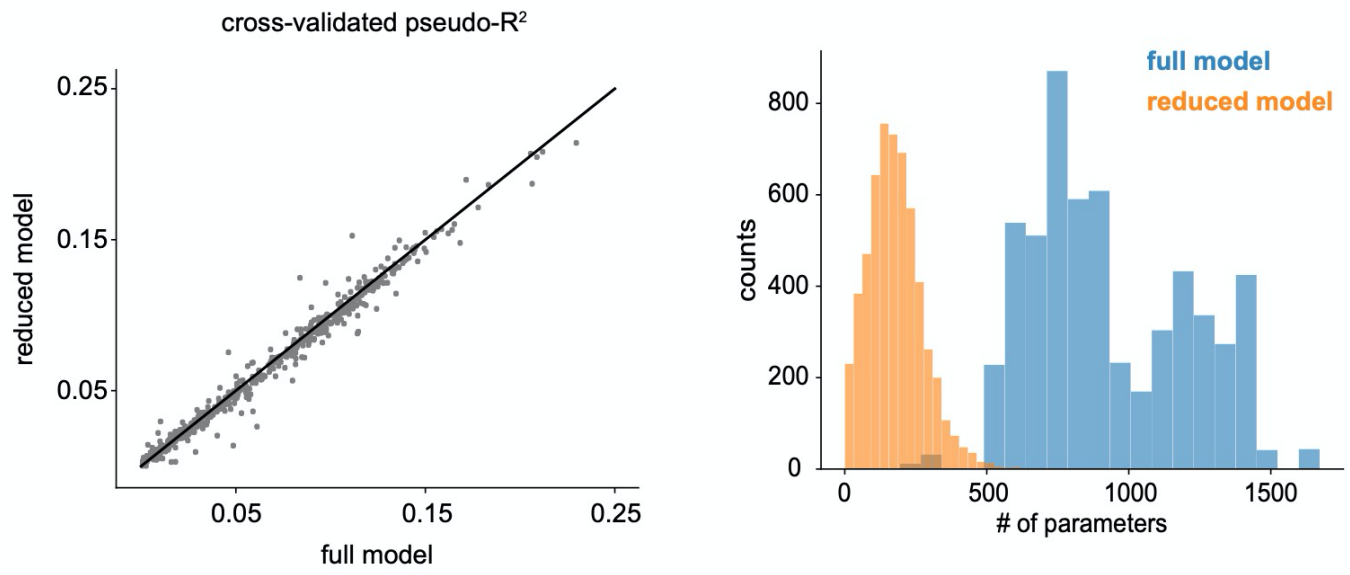
**Figure S13.** *P-GAM performance. Left: Cross-validated pseudo-R$^2$ for the reduced model (y-axis, including only variables deemed to significantly impact the unit's neural activity according to the P-GAM) vs. the full model (x-axis, including all parameters). Right: On average, the reduced model has 18.9% of the parameters in the full model (i.e., hundreds as opposed to thousands), yet it's ability to account for neural activity is comparable to the full model (left).*