Coding of latent variables in sensory, parietal, and frontal cortices during virtual closed-loop 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 and coordinate within the recurrent action-perception loops that characterize naturalistic self-environment interactions. Here, we record single-unit spiking activity and local field potentials (LFPs) 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 continuous latent variables: (i) the distance and angle they have travelled (i.e., path integration), as well as (ii) the remaining distance and angle to the memorized firefly location (i.e., a hidden spatial goal). We observed a patterned mixed selectivity, with the prefrontal cortex most prominently coding latent variables, parietal cortex coding for sensorimotor variables, and MSTd most often coding for eye-movements. Interestingly, however, even the traditionally considered sensory area (i.e., MSTd) tracked latent variables, demonstrating vector-coding of hidden spatial goals and path integration. Further, global encoding profiles and unit-to-unit coupling (i.e., noise correlations) suggested a functional subnetwork composed by MSTd and dlPFC, and not between these areas and 7a, as anatomy would suggest. We show that the greater the unit-to-unit coupling between MSTd and dlPFC, the more the animals' gaze position was indicative of the ongoing location of the hidden spatial goal. We suggest this MSTd-dlPFC subnetwork reflects the natural task strategy wherein monkeys continuously gaze toward the location of the (invisible) target. Together, these results highlight the distributed nature of neural coding during closed action-perception loops and suggest that fine-grain subnetworks may be dynamically established to reflect (embodied) task strategies.

Introduction

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

First, severing the closed loop that naturally exists between action and perception disrupts the timing that exists between sensory events and internal neural dynamics. In natural vision, for example, eye movements dictate the content, relative resolution, and timing of visual input. Indeed, recent work from active sensing (e.g., 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 likely mediated by phase-reset of neural oscillations (Lakatos et al., 2008; Rajkai et al., 2008) caused by the shifting gaze. In turn, physical movements of the eyes during sensory processing – an aspect absent in most laboratory, binary decision-making tasks – may not only enhance low-level visual encoding but may 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 across-area communication enables flexible behavior. This has possibly led to a degeneracy in the number of functions ascribed to each neural node. For instance, notorious redundancy has 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.

To tackle these limitations, we have developed a task requiring observers to catch memorized “fireflies” in virtual reality (Lakshminarasimhan et al., 2018, 2020; Noel et al., 2020, 2021). This goal-directed virtual navigation task addresses many of the limitations of the traditional approach, while remaining rooted in well-established 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 memorized location of this target (**Fig. 1A**). Importantly, the observers’ eyes are free to sample from their sensory surroundings, and trials last on the order of 2-3 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 environmental states (**Fig. 1A**). Importantly, the critical task variables are latent. Namely, observers must dynamically estimate their self-location (by accumulating velocity flow vectors) and that of the target (i.e., a hidden spatial goal). This dynamic computation of latent variables may offer a window into the mechanisms of internal computation (e.g., see Lakshminarasimhan et al., 2020).

Here, we leverage this “firefly task” to simultaneously characterize the encoding profiles of single units in sensory (dorsomedial superior temporal area, MSTd), parietal (area 7a), and frontal (dorsolateral prefrontal cortex, dlPFC) areas during closed-loop goal-directed behavior (**Fig. 1 - figure supplement 1**). 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 suggested a functional subnetwork composed by dlPFC and MSTd (and not area 7a, as would be predicted from anatomy). We suggest this MSTd-dlPFC subnetwork reflects the natural task strategy wherein monkeys continuously gaze toward the location of the (invisible) target.

Results

Monkeys navigate to remembered targets employing natural task 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 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 eccentricity spanning from -40° (leftward) to 40° (rightward) of visual angle (**Fig. 1A, B**). 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 (**Fig. 1B**). To quantify their performance, we expressed the monkeys' trajectory endpoints and firefly locations in polar coordinates, with an eccentricity from straight-ahead (θ) and a radial distance (r). **Figure 1C** shows radial (top; r vs. $\tilde{r}$ slope = 0.90; R^2 = 0.55) and angular (bottom; θ vs. $\tilde{\theta}$ slope = 0.95; 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 three animals and all sessions (Monkey Q, 27 sessions; S, 38 sessions; M, 17 sessions; **Fig. 1D**), monkeys were on average accurate, but tended to undershoot targets (r vs. $\tilde{r}$ slope, Mean=0.892, 95%CI=[0.860, 0.914]; θ vs. $\tilde{\theta}$ slope, M=0.792, 95%CI=[0.762, 0.842]). Undershooting targets is, in fact, the optimal strategy given that uncertainties over one's own and the firefly's position scales supra-linearly with time and distance travelled (see Lakshminarasimhan et al., 2018).

Most interestingly, this task affords observers with degrees of freedom that may reveal their innate task strategy. For instance, they are not required to fixate, and instead we may use their eye movements as an indication of their internal beliefs (see Lakshminarasimhan et al., 2020). Indeed, over the course of individual trajectories toward the target, the monkeys' eye position moves downward and becomes less eccentric (**Fig. 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; Adams et al. 2015). 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 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 (**Fig. 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. This suggest that, while not a task requirement, animals developed the strategy of fixating on the visible firefly and then attempted to maintain gaze on it even after the target had disappeared (**Fig. 1G,** see Lakshminarasimhan et al., 2020, for a full description of this phenomena and demonstration that pursuing the invisible firefly with our eyes results in improved task performance).

Altogether, the animals accumulated optic flow velocity signals (Alefantis et al., 2022) into an evolving (and largely accurate) estimate of self-position, and a relative estimate of the position to the target, over prolonged periods of time. In addition, they seemingly sampled from their environment in an intelligent manner, tracking the (invisible) target with their eyes in what is arguably an embodied mnemonic strategy.

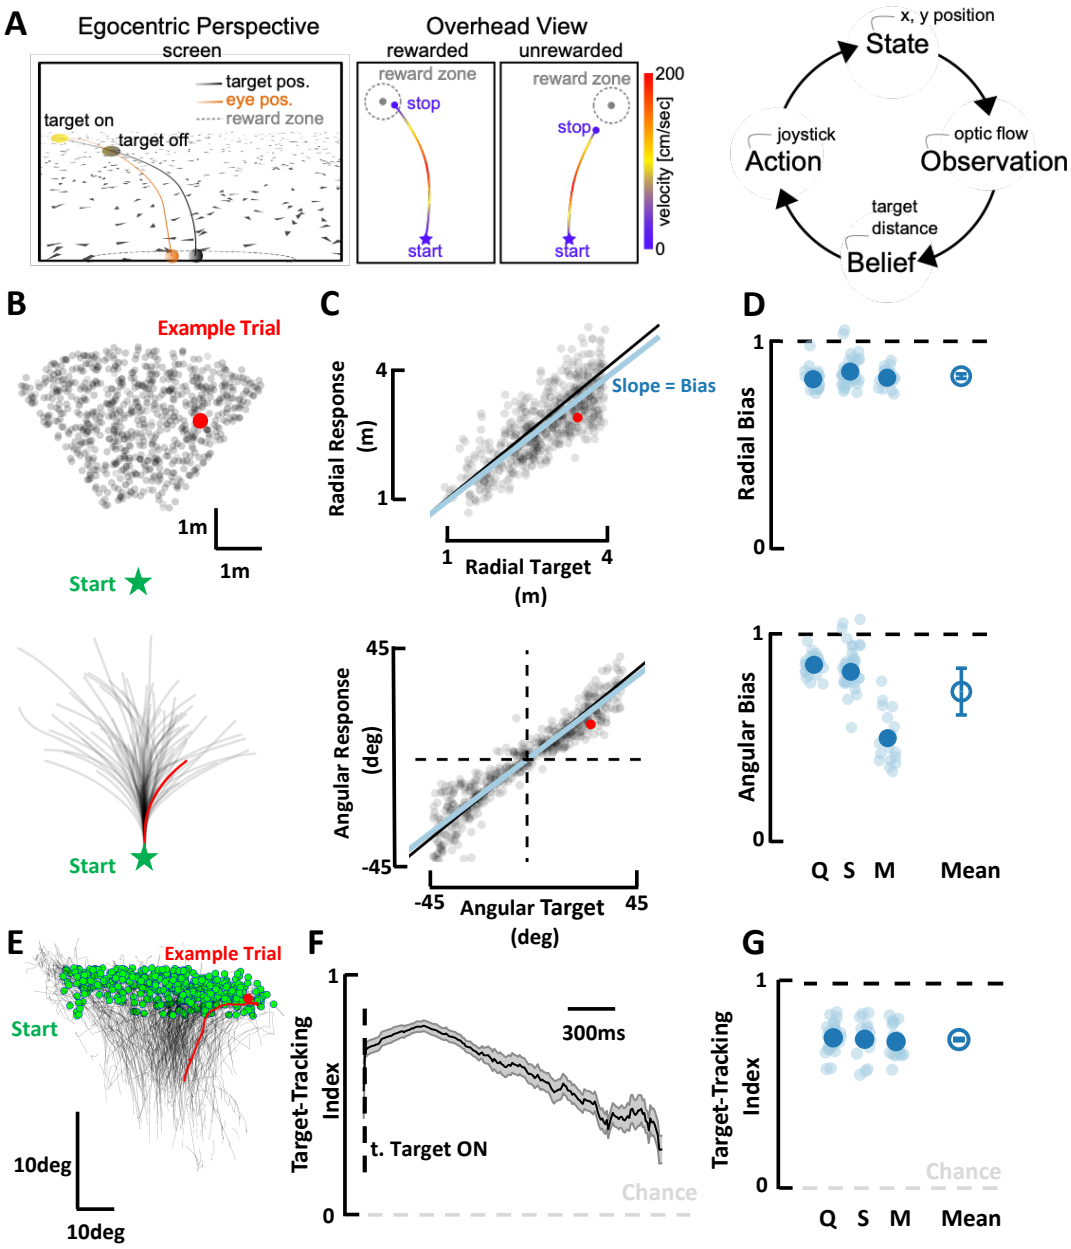

Figure 1. Task and behavioral results A. *Behavioral task. Left: Egocentric perspective of the monkey during the firefly task, emphasizing the ground plane optic flow elements (flashing triangles), the fact that the target disappears (target off), and eye-positions (following target position perfectly when the target is on, and then continuing on the correct path but deviating with time). Center: Overhead view of the monkey, starting location (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.*

Patterned 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) across 82 sessions while monkeys performed the firefly task (**Fig. 2A** and **Fig. 1 - figure supplement 1**).

To quantitatively account for the encoding properties of neurons within this task – wherein no two trials are alike (see **Fig. 2B** for rasters demonstrating heterogeneity in spike times and trial durations) – we fit spike trains to a Poisson Generalized Additive Model (P-GAM). The P-GAM we fit was purposely developed to efficiently and robustly estimate encoding profiles during closed-loop and naturalistic tasks such as the firefly task (see Balzani et al., 2020). Namely, beyond capturing arbitrary non-linearities and handling collinear predictors (Dormann et al., 2013), this 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 identify a minimal set of task features that each neuron is responsive to (setting statistical significance to $p < 0.01$), while circumventing the need for computationally demanding (an often unstable) model selection procedures – particularly given the large number and time-varying nature of the variables in this task. Indeed, beyond the numerical results in Balzani et al., 2020, demonstrating the ability to recover the ground-truth in artificial data, here we show 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 **Fig. 2 - figure supplement 1).**

Given our interest in examining neural dynamics at the single-cell level, in addition to discrete task events (e.g., time of target onset, as well as movement onset and offset), continuous sensorimotor variables (e.g., linear and angular velocity and acceleration) and latent features (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 causal unit-to-unit coupling filters 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, **Fig. 2C,** see *Methods*). These coupling filters capture noise correlations, i.e., the signal-independent relationship between one neuron firing and the likelihood that this same (i.e., spike history) or another neuron (i.e., coupling filter) will fire at a given time delay (see Hart & Huk, 2020). In total the encoding model had a 17 (analog or digital) task inputs (**Fig. 2C**), in addition to hundreds of potential lateral and across region coupling filters (see **Fig. 2 - figure supplement 1B** for the distribution of the number parameters in the full and reduced encoding models, and **Fig. 2 - figure supplement 1C** for demonstration that while the reduced model, i.e., that having only the set of parameters deemed to significantly impact firing rate, had a fifth of the parameters present in the full model, its cross-validated ability to account for observed spiking was untouched). The fitted model accounts well for observed spiking activity (see raw and model reconstructed spiking activity of simultaneously recorded units in **Fig. 2D**, mean psueudo-$R^2 = 0.072$).

Figure 2E shows the empirical (black) and P-GAM reconstructed (colored) firing rate to task variables in an example MSTd, 7a, and dlPFC neuron. The striking feature of these examples is that all units responded to, or were modulated by, a multitude of task variables. For instance, each of these example neurons responded at the time of firefly onset. The 7a neuron showed a strong dependency not only to sensorimotor variables, but also to LFP phases (in all frequency ranges). Finally, and perhaps most surprisingly, the MSTd neuron reflected not only sensorimotor and eye-position variables, but also latent ones, such as the angular distance to the (invisible) target. We used the P-GAM to factorize the contribution of different elements of the task to a neuron's overall firing rate (i.e., perform credit assignment). The example neurons in **Figure 2E** are opaque (contributes) or transparent (does not contribute) for different task variables, according to the P-GAMs estimate of their factorized contribution. Thus, there may be cases where there are evoked responses (e.g., time of movement onset in the dlPFC example neuron, **Fig. 2E** bottom) yet the neuron is not tuned to this variable, likely due to the correlated nature of input statistics during a naturalistic task (i.e., this neuron appears tuned to linear acceleration, which correlates with the time of movement onset).

The fraction of neurons tuned to different task-variables demonstrated a patterned mixed selectivity (**Fig. 2F).** Namely, the fraction of neurons tuned to sensorimotor variables was highest in area 7a (linear and angular velocity, linear and angular acceleration, time of movement onset and offset, all $p < 0.01$; Cohen's d linear acceleration $d = 0.2$, all the rest $d > 0.76$), while most neurons in MSTd coded for eye position (58.8% and 63.3%

respectively for vertical and horizontal eye position). Interestingly, a large fraction of dlPFC neurons also coded for eye position (dlPFC v. 7a, vertical and horizontal eye position, $p < 6.4 \times 10^{-17}$, all Cohen's $d > 3$), putatively reflecting the fact that the gaze may reflect the internal belief of firefly position (Lakshminarasimhan et al., 2020). Neurons in area 7a showed a very strong dependency to the phase of ongoing LFP fluctuations, a fact that was less observed in dlPFC or MSTd (all $p < 7.1 \times 10^{-8}$, all $d > 1.02$), particularly in the alpha and theta frequency range. Lastly, and most relevant to the task of path integrating to the location of invisible fireflies, we observed that the greatest fraction of neurons encoding for both path integration (i.e., travelled distance and angle turned) and the remaining distance and angle to the target (i.e., spatial goal) were in dlPFC (ranging from 30.6% to 57.5%). Interestingly, while 7a had more neurons coding for path integration than MSTd (all $p < 1.6 \times 10^{-5}$, all Cohen's $d > 0.54$), the latter area had a greater fraction of neurons coding for the latent distance and angle to target than area 7a (all $p < 6.1 \times 10^{-5}$, all Cohen's $d > 0.41$). This may suggest that while dlPFC is critically involved in estimating the relative distance between self and target, 7a may be preferentially involved in the process of path integration, while somewhat unexpectedly, MSTd may play an important role in keeping track of the latent spatial goal. In the next section we take a deeper look into the tunning functions for latent variables, the critical task variables. See **Fig. 2 - figure supplement 2** for a demonstration of the stability in the estimates of the fraction of neurons tuned.

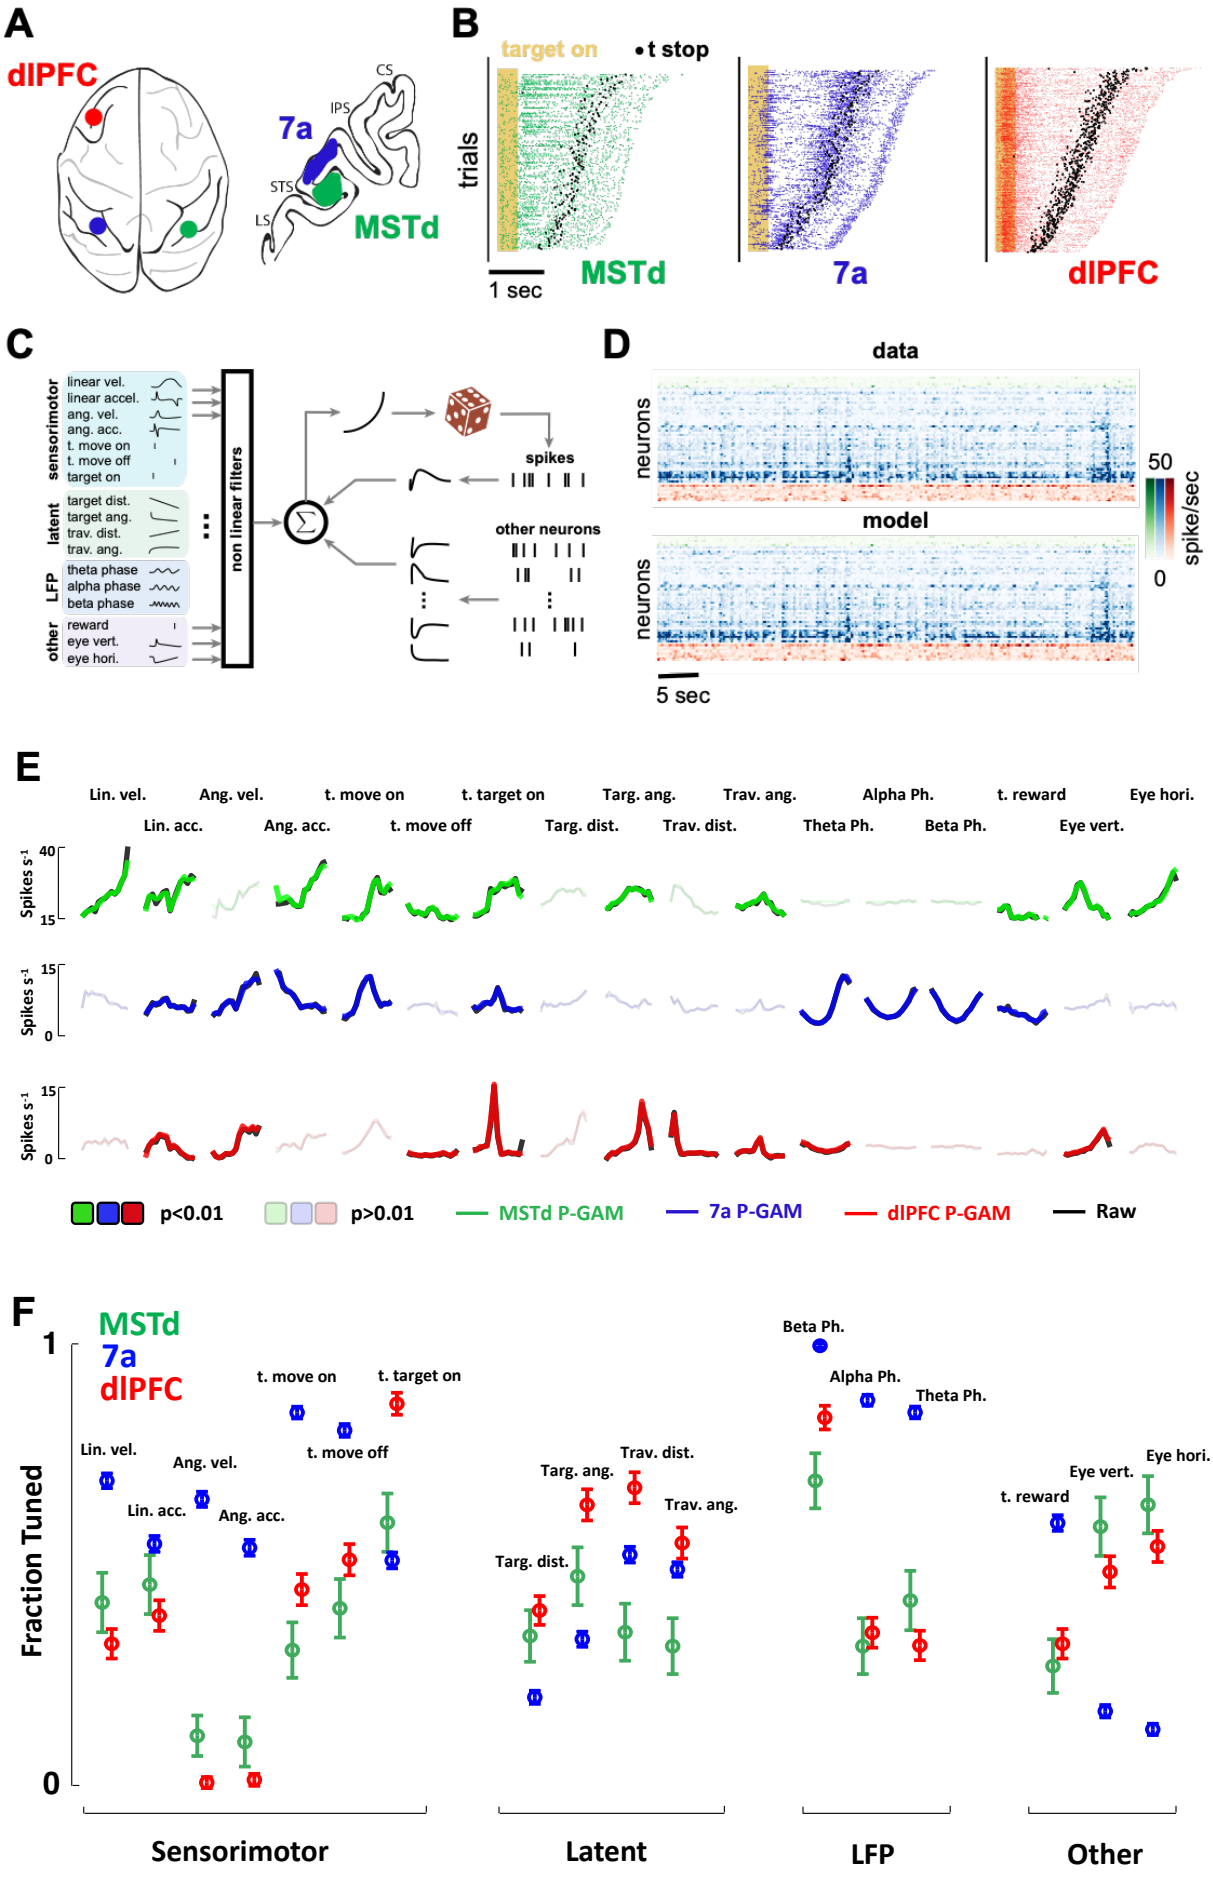
A
dlPFC
7a
MSTd
CS
IPS
STS
LS
B
target on
t stop
trials
1 sec
MSTd
7a
dlPFC
C
sensorimotor
linear vel.
linear accel.
ang. vel.
ang. acc.
t. move on
t. move off
target on
latent
target dist.
target ang.
trav. dist.
trav. ang.
LFP
theta phase
alpha phase
beta phase
other
reward
eye vert.
eye hori.
non linear filters
spikes
other neurons
D
data
neurons
model
neurons
50
0
spike/sec
5 sec
E
Lin. vel.
Lin. acc.
Ang. vel.
Ang. acc.
t. move on
t. move off
t. target on
Targ. dist.
Targ. ang.
Trav. dist.
Trav. ang.
Theta Ph.
Alpha Ph.
Beta Ph.
t. reward
Eye vert.
Eye hori.
Spikes s-1
40
15
0
p<0.01
p>0.01
MSTd P-GAM
7a P-GAM
dlPFC P-GAM
Raw
F
MSTd
7a
dlPFC
Fraction Tuned
1
0
Sensorimotor
Latent
LFP
Other

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 Poisson generalized additive model (P-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, the authors reported 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.*** *Responses from an example MSTd, 7a, and dlPFC neuron (black), aligned to temporal task variables (e.g., time of movement onset and offset), or binned according to their value in a continuous task variable (e.g., linear velocity). Colored (respectively green, blue, and red for MSTd, 7a, and dlPFC) lines are the reconstruction from the reduced P-GAM. The responses are opaque (p<0.01) or transparent (p>0.01), according to whether the P-GAM kernel for the specific task variable is deemed to significantly contribute to the neuron's firing rate).* ***F.*** *Fraction of neurons tuned to the given task variable, according to brain area. Error bars are 99% CIs, and thus non-overlapping bars indicate a pair-wise significant difference.*

Cortical codes for path integration and vector-coding of spatial goals

Beyond the frequency with which we observe neurons tuned to the distance or angle from origin (i.e., path integration) and to target (i.e., vector coding of spatial goals), we may expect the distributions of preferred distances and angles to also be informative. **Figure 3A** shows that each of these distributions spanned the entire range of angles and distances visited, demonstrating a full basis set covering the feature space (distributions shown are for a single animal, Monkey S).

Across all animals, angles specifying a 0° offset from the heading at origin were over-represented (**Fig. 3B**, top panel, proportions across all animals), while the 0° offset from target location was under-represented (**Fig. 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 vector-coding of hidden spatial goals exists and a goal angle of 0° is over-represented (Sarel et al., 2017). Speculatively, this reversed pattern of preferences between cortex (here) and hippocampus (Sarel et al., 2017) may suggest that while the former plays a role in guiding heading toward the desired state (akin to an error or correction signal), 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 (**Fig. 3A and B**, 3rd row, < ~150cm). On the other hand, we observed a clear differentiation between cortical areas regarding their preferred distance to target. Area 7a showed a strong preference for nearby targets with approximately 50% of units coding for targets within ~30cm. Neurons in MSTd and dlPFC, on the other hand, responded primarily at intermediary and far distances from the target (~200-400 cm from target; **Fig. 3A and B**, 4rd row). The preference for distances further from the target in MSTd and dlPFC seemingly concords with their frequent tuning to eye position, and the fact that the eyes attempt to pursue the hidden target (i.e., as if these areas were involved in planning and sampling from a distance).

Lastly, to depict the full shape of the kernels encoding the distance and angle to the origin and from the target, we performed statistical clustering of these kernels, separately for each area and task variable. **Figure 3C** shows the mean and 10 example tunning functions for the most frequently present cluster within each area. The most notorious difference between areas is that while MSTd and dlPFC prefer distances far from the target, area 7a responds mostly to distances near it (**Fig. 3C**, bottom row). In the *Supplementary Materials* (**Fig. 3 - figure supplements 1-4**) 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.

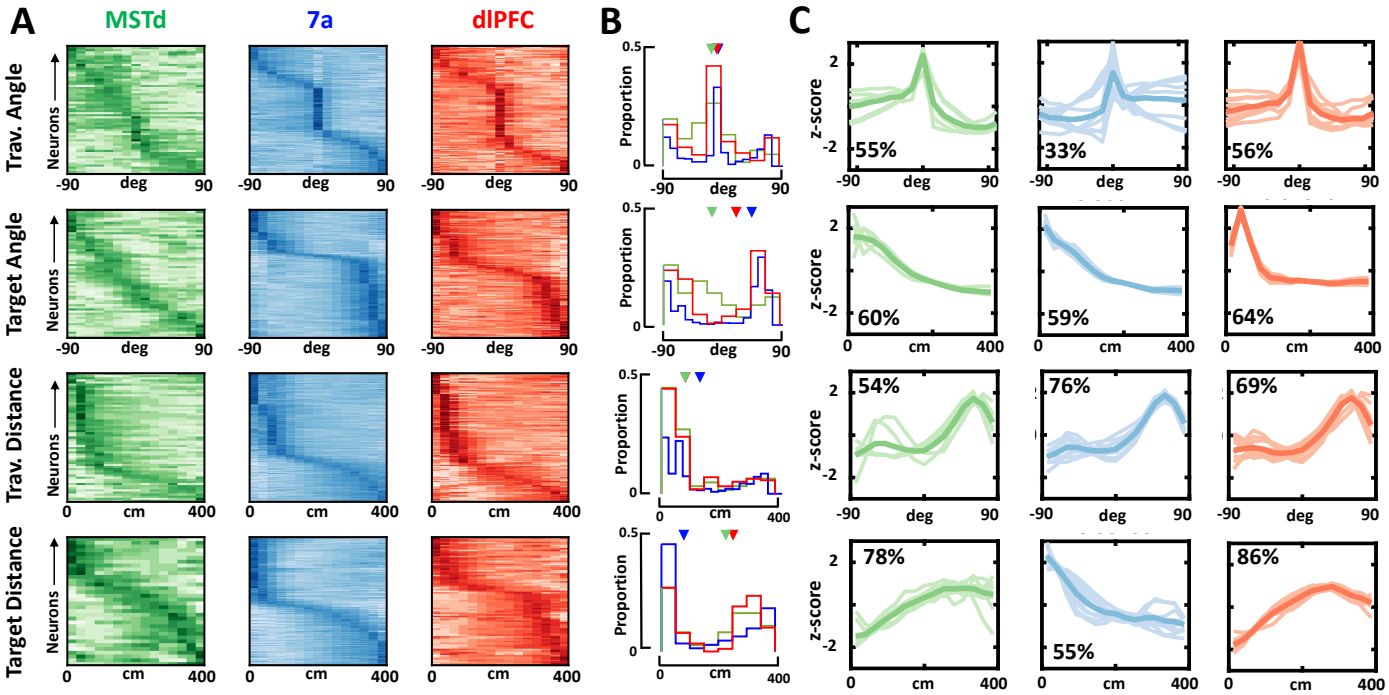

Figure 3. Preferred angle and distance from origin and to target in MSTd, 7a, and dlPFC. A. *Heatmaps showing neural responses (y-axis) sorted by preferred angles from origin (top), angle to target (2nd row), distance from origin (3rd row), and distance to target (bottom row) for MSTd (green), 7a (blue) 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.* ***C.*** *We clustered the kernels driving the response to angle or distance to origin and from the target. Here we show 10 representatives from each cluster (thin and transparent lines), as well as the mean of the cluster as a whole (ticker line). The inset quantifies the percentage of tunings within a particular area and for the particular variable that were deemed to belong within the cluster depicted (the most frequent one).*

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

Both the fraction of neurons tuned to different task variables (**Fig. 2**) and the distribution of preferred angles and distances from origin and to target (**Fig. 3**) show a surprising commonality between MSTd and dlPFC. In turn, to systematically examine how the encoding of all task variables, and not solely a select few, varied across cortical areas, we employed a statistical clustering approach (see e.g., Minderer et al., 2019, for a similar approach).

First, we leveraged the knowledge of which variables each neuron was significantly tuned to (e.g., **Fig. 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 2**. This approach (see *Methods* for detail) showed that nearly all MSTd (89%) and dlPFC (94%) neurons were categorized as belonging within the same cluster (**Fig. 4A,** cluster number 1), one that was defined by true mixed selectivity (**Fig. 4B**, top left). In contrast, area 7a neurons appeared in three distinct clusters (**Fig. 4A,** 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 (**Fig. 4B**, top center), while Cluster 3 was characterized by a near complete lack of tuning to latent variables and eye position (**Fig. 4B**, 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 (**Fig. 4A and B**, Cluster 4), or driven also by motor onset and offset (**Fig. 4A and B**, Cluster 5). These remaining clusters were, however, less common (~1-5%). Together, this pattern of clustering results based on whether neurons were tuned to different task variables demonstrated a surprising commonality between MSTd and dlPFC, which are in turn different from area 7a.

Subsequently, we questioned whether utilizing knowledge of the shape of these tuning functions (as opposed to simply whether tuning to a variable was significant or not) 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 (**Fig. 4C,** 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 (**Fig. 4C,** 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 (**Fig. 4C,** third and fourth row). Thus, just as the clustering based on what variables were neurons tuned to, 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 **Fig. 4 - figure supplement 1** for UMAP projections color coded by the mutual information between neural responses and a particular task variable).

Given the similarity in encoding profiles between MSTd and dlPFC, we next examined inter-area global synchronization and unit-to-unit coupling to question whether these two areas formed 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 (see Hart & Huk, 2020) as estimated from the P-GAM (Balzani et al., 2020).

Spiking activity in some units was well explained by the ongoing LFP phase in different regions (above and beyond the ongoing LFP phase in their local area, see **Fig. 2F**). Most notably, 37% of neurons in dlPFC were tuned to the ongoing LFP phase within the beta range in area 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%, **Fig. 4D**). Globally, therefore, phase locking of spiking activity to LFP phase in different areas reflected 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). Interestingly, 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. The lowest frequency we indexed (theta, 4-8hz) seemed to be reserved for putative “feedback” influence of dlPFC onto both MSTd (12%) and area 7a (18%).

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 (**Fig. 4 - figure supplement 2A**). 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%, **Fig. 4E**), potentially reflecting an increase in the dimensionality of the neural code (i.e., decrease in redundancy) as we move from sensorimotor 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 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% (**Fig. 4E**).

Four arguments support the fact that unit-to-unit coupling as estimated by the P-GAM reflect a property of neural coding beyond global volume-conductance. First, they significantly improved fits of the encoding model. **Figure 4D** 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). Second, 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 **Fig. 4 - figure supplement 2B** for a ‘dictionary’ of coupling filters across areas and their relative frequency). Third, while units phase-locked to the LFP in a different region were 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 (**Fig. 4 - figure supplement 2C**). Lastly, and most importantly, the likelihood of two units being coupled was not random or purely depending on their physical distance, but instead varied as a function of their tuning similarity. **Fig. 4 - figure supplement 2D** shows that the more similar the tuning functions were between two neurons, the more likely 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) within a naturalistic closed-loop task, even for latent variables, and even in MSTd.

Altogether, both the global encoding profiles and fine grain unit-to-unit coupling across areas (**Fig. 4**) suggested the presence of a MSTd-dlPFC functional subnetwork within this closed-loop virtual navigation task. Given the high probability to encountering neurons in MSTd and dlPFC tuned to the distance and angle from the target, as well as to eye position (**Fig. 2**), and given the animals' tendency to keep track of the firefly location with their eyes (**Fig. 1**, and Lakshminarasimhan et al., 2020), we hypothesized that this MSTd-dlPFC network may reflect the monkey's mnemonic strategy in navigating to hidden targets.

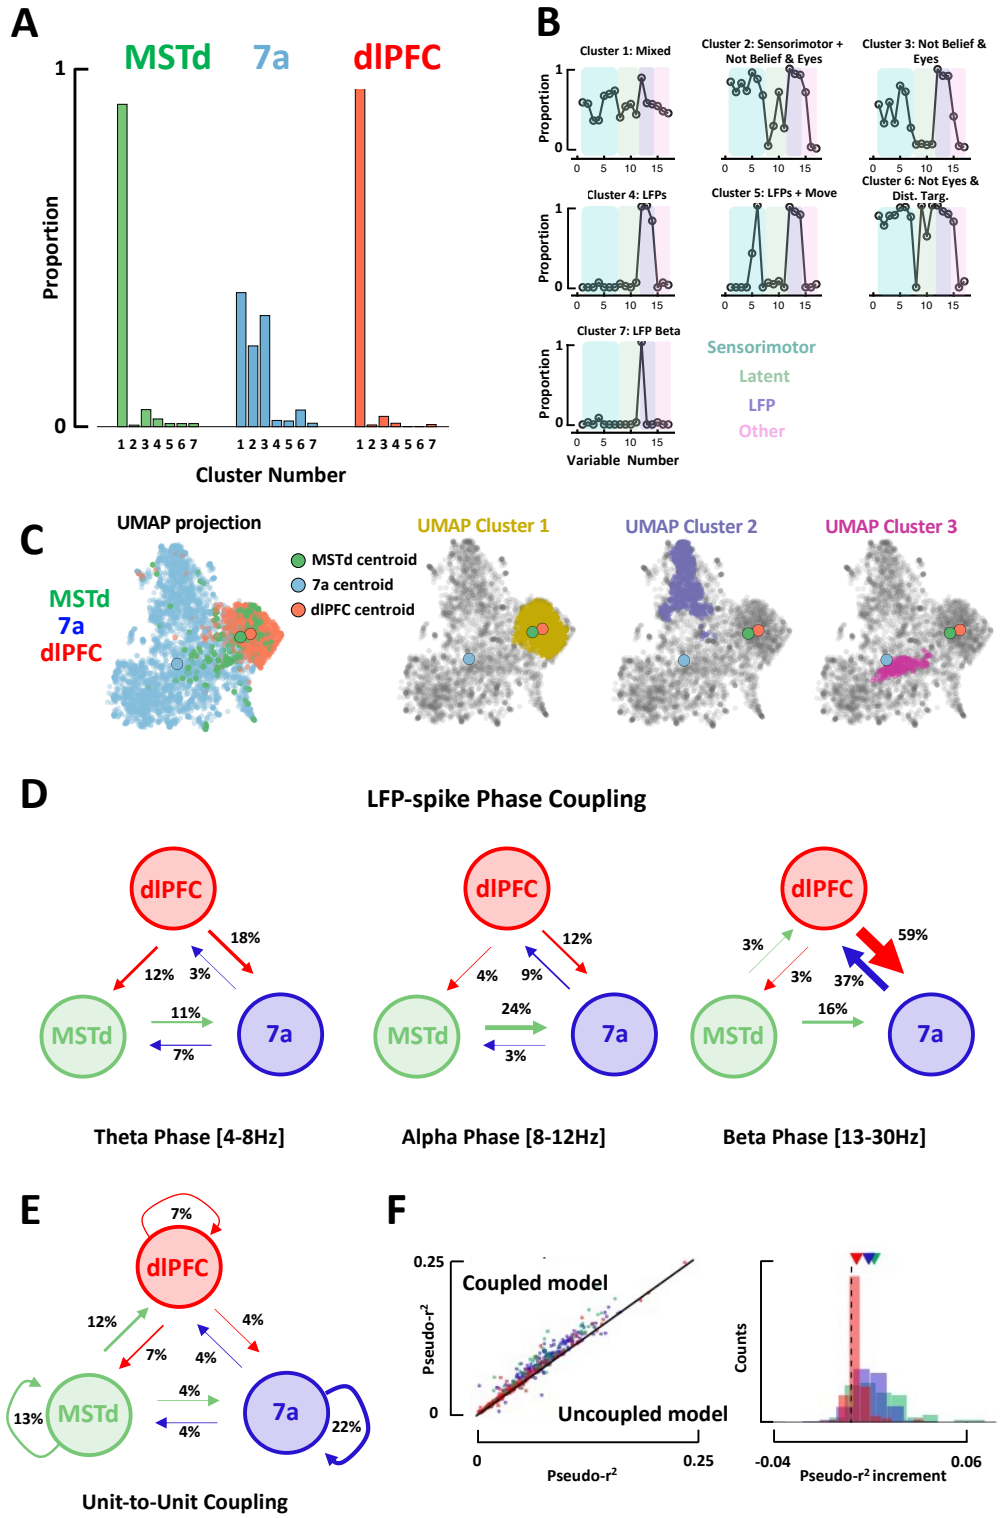

Figure 4. Global single-unit encoding profiles and unit-to-unit coupling properties suggests a common functional role for MSTd and dlFPC. A. *Proportion of neurons being classified into distinct cluster (x-axis, 7 total) according to which task parameters they were significantly tuned to.* ***B.*** *Fraction of neurons tuned to each of the 17 task variables (order is the same as in* ***Figure 2C, E, F****) according to cluster.* ***C.*** *UMAP projection of the tuning function shapes, color coded by brain area (first column) or DBSCAN cluster (second, third, and forth column).* ***D.*** *Fraction of neurons whose spiking activity is phase-locked to LFP phases in 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.* ***E.*** *As* ***D.*** *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.* ***F.*** *Left: 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). Right: Probability distribution of the change in pseudo-R^2 from the reduced to the full model.*

MST-dlPFC coupling reflect the animals' strategy of tracking hidden targets with their eyes

The monkeys' gaze reflected their belief about the firefly location. Indeed, within a session and across trials, the mean absolute difference between a monkey's eye position and where they ought to be looking if they were perfectly tracking the firefly correlated with steering endpoint error (r^2 ± standard deviation across all datasets, 0.36 ± 0.22; for shuffled data, 0.16 ± 0.09, p = 5.6 x 10^{-3}, paired t-test; see Lakshminarasimhan et al., 2020, for the original description of this effect). Further, this relationship also held across sessions, with better target tracking correlating with less bias (slopes closer to 1, see **Fig. 1C**), particularly in the angular domain (r^2 = 0.43, p = 0.004; radial: r^2 = 0.26, p = 0.04; **Fig. 5A**). Thus, we may question whether the likelihood of observing unit-to-unit coupling (defined within a session but not within a trial) relates to session-by-session changes in target tracking and/or steering performance.

For well-fit sessions (mean pseudo-r^2 > 0.05) with at least 2 units in each of 2 different areas, we computed the monkey's target tracking index (averaged across the entire trial and then across trials), as well as the probability of units being coupled with others within and across areas, as estimated by the P-GAM. As shown in **Figure 5B** (top right)**,** we observed a strong association between the fraction of units showing MST-to-dlPFC coupling and the monkeys' ability to track the hidden firefly with their eyes (r^2 = 0.63, p = 0.003). This association, with more coupling indexing better tracking of the firefly, was also true for the "feedback" projection from dlPFC-to-MSTd (r^2 = 0.41, p = 0.035, **Fig. 5B,** bottom left), as well as for dlPFC-to-dlPFC coupling (r^2 = 0.43, p = 0.047, **Figure 5B,** bottom right). The other sender-receiver pairings, including MSTd-MSTd, did not show a correlation with the ability of the monkey to track the hidden firefly with their eyes within a session (all p > 0.08). As shown in **Figure 5 - figure supplement 1**, there was no correlation between the unit-to-unit couplings within a session, and either radial or angular biases (7a-7a coupling v. angular bias p = 0.06; all other p > 0.12). Overall, therefore, the functional subnetwork between MSTd and dlPFC (**Fig. 4**) seemingly reflects the animals' strategy in keeping track of the hidden target with their eyes.

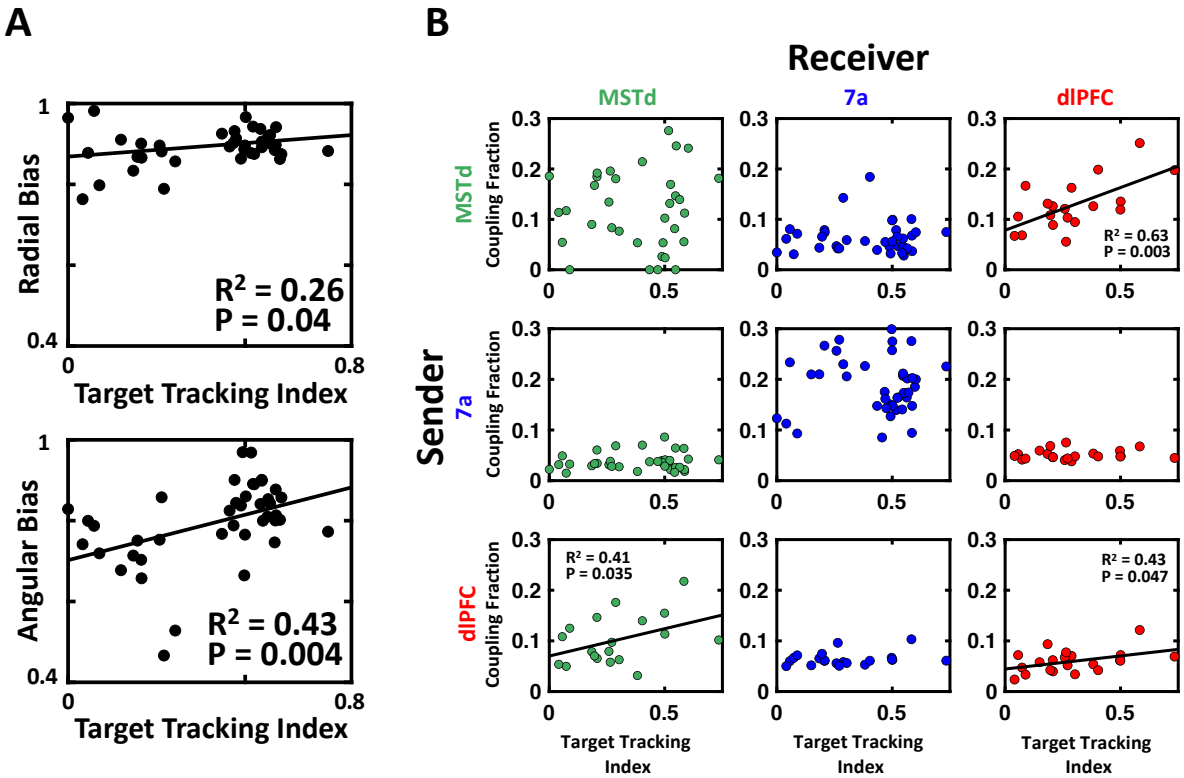

Figure 5. Increased MSTd-dlPFC coupling correlated with an increased likelihood of animals keeping track of the invisible fireflies with their eyes. A. *Correlation between the target tracking index (x-axis, i.e., tracking the hidden target with their eyes) and the radial (top) or angular (bottom) bias (defined as the slope relating responses and targets, as in Figure 1). Only sessions included in the neural analysis in Panel B, were included in Panel A.* ***B.*** *Correlation between the target tracking index (x-axis) and the fraction of neurons coupled within the "sender" region (y-axis). The diagonal shows within area couplings (MSTd-MSTd, 7a-7a, dlPFC-dlPFC), while off-diagonals show across area couplings. R^2 and p-values are shown as insets for the significant correlations.*

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. To examine how the brain operates and coordinates among multiple neural nodes under these naturalistic conditions 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 focused on the encoding profiles of single-units and described four main findings.

First, we found that all task variables were present across all brain areas probed. This multiplexing, however, was not random. 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. To the best of our knowledge, the striking difference between the posterior parietal node and other areas (here MSTd and dlPFC) vis-à-vis their dependency on the LFP phase has not been previously reported and may have been acutely evident here given the natural timing between sensory inputs, motor outputs, and ongoing neural dynamics that exists within this closed-loop setting. MSTd was prominently tuned to eye position (Komatsu & Wurtz., 1988; Nader et al., 2009). Most notoriously, even MSTd, an area traditionally considered a sensory node, showed a considerable number of neurons encoding latent variables, such as the angle or radial distance from origin (i.e., path integration) and to target (i.e., vector coding of spatial goals). In fact, the global encoding profile of MSTd was strikingly akin to that of dlPFC, suggesting this area need not be purely sensory (see Mendoza et al., 2014 for a similar argument).

Second, 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, Gauthier & Tank, 2018, and Marigold & Drew, 2017, for similar evidence in humans, rodents, and cats). 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 posterior parietal cortex may compute 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 single 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 third major finding relates to the functional organization across MSTd, area 7a, and dlPFC, and again suggests that navigating to a remembered target may involve two interdependent but separable computations; (i) estimating own's own location relative to a starting location (i.e., a landmark), and (ii) estimating the position of the target and discounting self-location from target location. Indeed, 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 observed that the coarse channels of communication riding on LFPs are in line with anatomy. Further, the strong tuning of area 7a to sensorimotor variables and the strong tunning of both area 7a and dlPFC (but not MSTd) to the distance and angle from origin, suggest that this pathway may be primarily involved in path integration as opposed to the vector coding of spatial goals. In addition to this MSTd-7a-dlPFC pathway, we also observed significant coupling between MSTd and dlPFC, independent from 7a. Namely, the likelihood of MSTd neurons being coupled to dlPFC, and of dlPFC neurons being coupled to MSTd, was 3 and almost 2 times as large as the unit-to-unit 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 navigate-to-target task.

Lastly, we examined the putative functional role of the MSTd-dlPFC subnetwork. As our group has previously reported, monkeys naturally track with their eyes the location of the hidden goal during navigation (Lakshminarasimhan et al., 2020). Given this task strategy, as well as the prominent tunning of both MSTd and dlPFC to eye position and the distance and angle to the hidden target, we hypothesized that the MSTd-dlPFC subnetwork may reflect the monkeys' innate (i.e., not imposed during training) task strategy. Indeed, we found that the more units were coupled between MSTd and dlPFC (either "feedforward" or "feedback"), the better monkeys tracked the hidden goal with their eyes. Thus, we suggest that the dynamic coupling between MSTd and dlPFC may reflect the innate task strategy monkeys adapted in pursuing the goal location with their eyes. In future experiments it will be interesting to test for causality within this network, by either demanding observers to fixate elsewhere (Lakshminarasimhan et al., 2020), or by directly perturbing this network, for instance by micro-stimulation.

In conclusion, we demonstrate what may be broad principles of neural computation during closed action-perception loops: (i) mixed yet patterned single cell selectivity, (ii) coding of latent variables even in areas traditionally considered as purely sensory, and (iii) differential coarse (e.g., LFP-spike phase alignment) and fine grain connectivity between task relevant areas. Most notoriously, here we indexed the presence of significant noise correlations between MSTd and dlPFC, independently from 7a. We suggest that this fine grain coupling between sensory and pre-frontal areas may reflect (embodied) task strategies.

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 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 density of the ground plane was 5.0 elements/m^2. 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 from the center of the target. No juice was provided otherwise.

Neural Recording and Pre-Processing

We recorded extracellularly, either acutely using a 24-channel linear electrode array (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 to facilitate real-time monitoring of action potential waveforms. Recordings began once the 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 chronic 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 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 = θ; response = $\tilde{\theta}$). On a subset of trials (~13%) animals stopped within 0.5m of the origin (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, whereas slopes smaller than one indicate undershooting (either in radial or angular distance).

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 $y_t \in \mathbb{N}_0$ (here binned at 6ms temporal resolution) and a set of continuous covariates $\boldsymbol{x}_t$ (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 $\boldsymbol{z}_t$ (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 λ_f,

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

The larger λ_f, the smoother the model. This penalization terms can be interpreted as Gaussian priors over model parameters. The resulting log-likelihood of the model takes the form,

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

with $\boldsymbol{y} \in \mathbb{R}^{\boldsymbol{T}}$ being the spike counts of the unit, $\boldsymbol{x} \in \mathbb{R}^{J \times T}$ being the continuous task variables, $\boldsymbol{z} \in \mathbb{R}^{K \times T}$ being the discrete task events, T being the time points, $\boldsymbol{\beta}$ being the collection of all B-spline coefficients, and $p(\cdot)$ the Poisson likelihood. Both parameters $\boldsymbol{\beta}$ and the hyperparameters $\boldsymbol{\lambda}$ 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).

Most importantly, 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 2 – supplement 1** we demonstrate that this approach will appropriately select variables contributing to the spike trials of individual neurons. Further, we show that including all variables in the model (hundreds to 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-R^2

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,

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

with $\mathcal{L}(\boldsymbol{y})$ being the likelihood of the true spike counts, $\mathcal{L}(\hat{\boldsymbol{y}})$ being the likelihood of the P-GAM model prediction, and $\mathcal{L}(\bar{\boldsymbol{y}})$ being 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 substantially 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., 2018). 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

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. For **Figures 4A, B,** 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 **Figure 4C** all tuning functions 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 ith 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 was 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. Clusters depicted in **Figure 3C** and **Figures 3 – supplement 1-4** were equally determined by DBSCAN.

For **Figure 4 – supplement 1**, 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),

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

To estimate these quantities, we assumed that the counts Y are Poisson distributed with rate parameter λ 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,

$$H(Y) = -\sum_y p(y) \log p(y), \text{ and } H(Y|S) = -\sum_s p(s) \sum_y p(y|s) \log p(y|s) \qquad \text{(Eq. 7).}$$

Coupling Filters

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 4E**, we corrected for the effects of unit distance (i.e., **Figure 4 – supplement 2A**) by first fitting a brain region-specific logistic regression. Specifically, we expressed coupling probability as a non-linear function of electrode distance as follows,

$$p(c = 1) = logit^{-1}(f(d)) \qquad \text{(Eq. 8)}$$

With c being a binary variable taking value 1 for significant coupling and 0 otherwise, d being the electrode distance, and f being a non-linear function expressed in terms of B-splines. Each brain area was fit independently, and the coupling probability in **Figure 4E** was set as the model prediction for a distance of 500um.

Data and code Availability

Data and code are available at: https://osf.io/eybc8/files/.

Acknowledgements

The authors thank Jing Lin and Jian Chen for programming the experimental stimulus. We also thank Roozbeh Kiani for his surgical expertise during the Utah array implantations. The work was funded by a 1U19 NS118246, 1R01 NS120407, and 1R01 DC014678 from NIH to D.E.A, as well as by 1R01MH125571 from NIH, the National Science Foundation under NSF Award No. 1922658, and a Google faculty award to C.S.

Author Contributions

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

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

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.

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.

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.

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.

Hart E, Huk AC. (2020). Recurrent circuit dynamics underlie persistent activity in the macaque frontoparietal network. Elife;9:e52460.

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.

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.

Komatsu H, Wurtz RH. (1988). Relation of cortical areas MT and MST to pursuit eye movements. I. Localization and visual properties of neurons. J Neurophysiol; 60(2):580–603

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

Marigold DS, Drew T. (2017). Posterior parietal cortex estimates the relationship between object and body location during locomotion. Elife. Oct 20;6:e28143.

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

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. (2014). *Nat Neurosci*.;17(10):1410-7. doi: 10.1038/nn.3807. Epub 2014 Sep 7. PMID: 25195105; PMCID: PMC4486057.

Nadler JW, Nawrot M, Angelaki DE, DeAngelis GC. (2009). MT Neurons Combine Visual Motion with a Smooth Eye Movement Signal to Code Depth-Sign from Motion Parallax. Neuron, 63

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.

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

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. (2000). Normalized cuts and image segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence. Vol. 22, pp. 888–905.

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).

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.

Supplementary Figure Legends

Figure 1 - supplement figure 1. Recoding sites. 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 2 - supplement figure 1. P-GAM controls. A. Feature inclusion in the "reduced" model. 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 the 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) was zeroed. Columns show 7 example cells. **B. Number of parameters in the "full" and "reduced" model.** The full model is the model with all parameters, before any selection is applied. The reduced model is the model where only parameters relating to features that significantly contribute to a neuron's spiking are kept. By design, the reduced model ought to have an equal or lower number of parameters than the full model. The full model number of parameters varies, according to how many simultaneously recorded neurons there were in a recording session, given that we include unit-to-unit coupling terms. **C. Cross-validated pseudo-R^2 in the full and reduced model.** While the reduced model has on average 18.9% the number of parameters as in the full model **(B)** its ability to account for neural activity is comparable to the full model.

Figure 2 - supplement figure 2. Stability in the fraction of neurons tuned to different task variables. To test the stability in the P-GAMs estimates, we divided our datasets in even and odd trials. We fitted separate P-GAMs and plot the fraction of neurons tuned (y-axis) to each task variable (x-axis) according to brain area (MSTd, 7a, and dlPFC) and whether these were odd of even trials. Error bars are 99%CI. Estimates in the fraction of neurons tuned were very stable.

Figure 3 - supplement figure 1. 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. 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 3 - supplement figure 2. 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 3 - supplement figure 3. 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 3 - supplement figure 4. 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 4 - supplement figure 1. 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 4 – supplement figure 2. A. 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). **B.** 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. 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. **C.** 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 the unit was phase-locked to LFP in the other area (left column) or coupled to a neuron in the other area (right column). 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). **D.** 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 (r2 = [0-.25; 0.25-0.5; 0.5-0.75; 0.75-1]) and averaged within each category (e.g., sensorimotor) 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.

Figure 5 - supplement figure 1. Correlations between fraction of neurons coupled in a "sender" region and steering behavior. Correlations between the radial (**A**) or angular (**B**) bias within a session (defined as the slope relating responses and targets, as in Figure 1) and the fraction of neurons coupled within and across areas. The diagonals show within area couplings (MSTd-MSTd, 7a-7a, dlPFC-dlPFC), while off-diagonals show across area couplings. There was no correlation between fraction of units coupled and steering behavior.

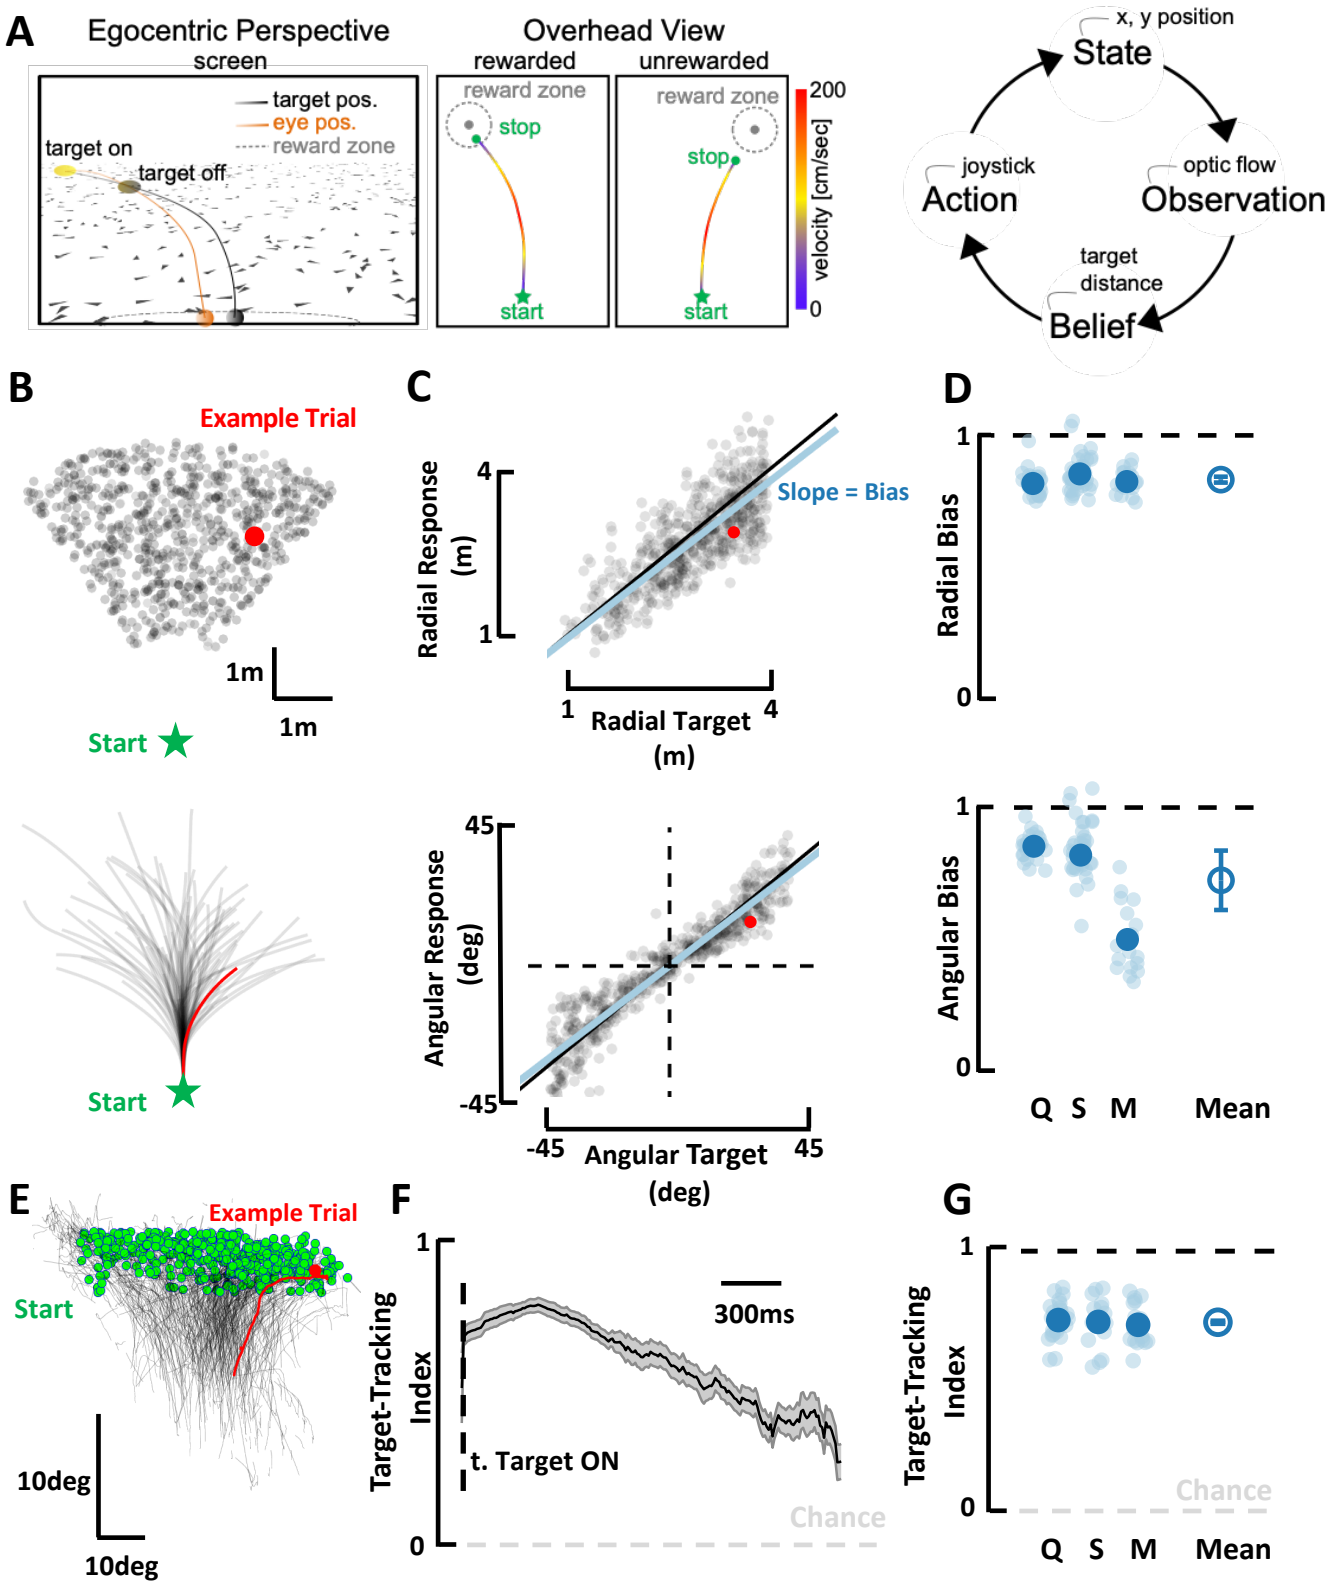
A
Egocentric Perspective
screen
target pos.
eye pos.
reward zone
target on
target off
Overhead View
rewarded
unrewarded
reward zone
stop
start
200
velocity [cm/sec]
0
x, y position
State
joystick
Action
optic flow
Observation
target
distance
Belief
B
Example Trial
1m
1m
Start
C
Radial Response
(m)
4
1
Slope = Bias
1
Radial Target
(m)
4
D
Radial Bias
1
0
Start
Angular Response
(deg)
45
-45
-45
Angular Target
(deg)
45
Angular Bias
1
0
Q S M Mean
E
Example Trial
Start
10deg
10deg
F
Target-Tracking
Index
1
0
300ms
t. Target ON
Chance
G
Target-Tracking
Index
1
0
Chance
Q S M Mean

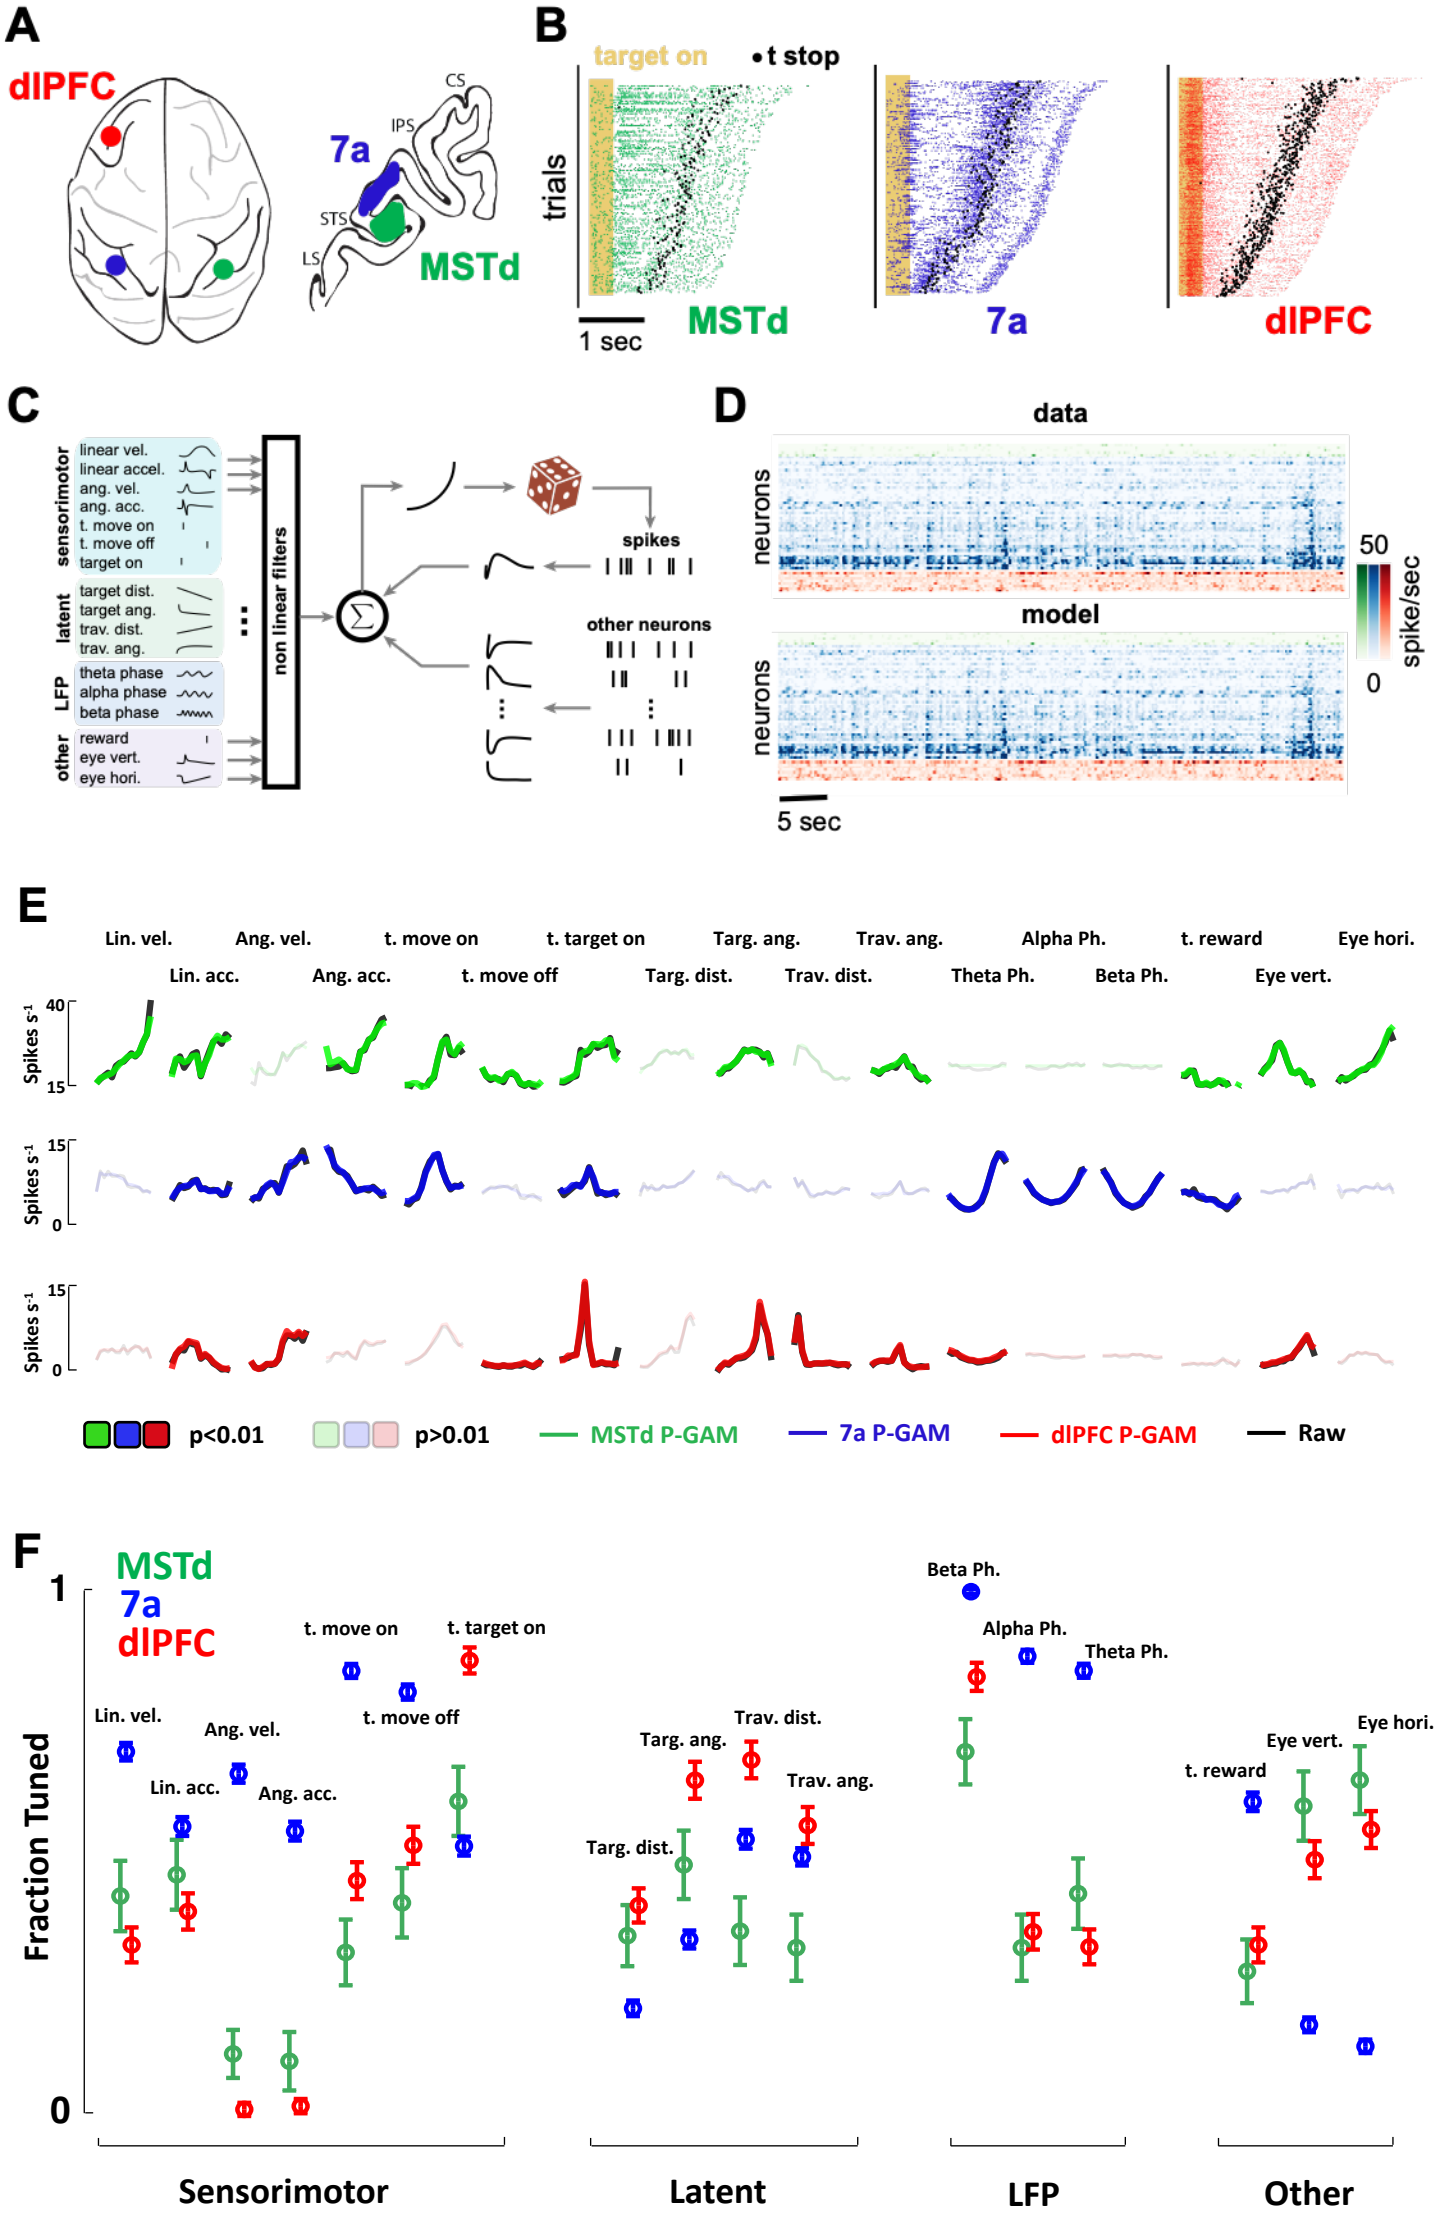
A
dlPFC
7a
MSTd
CS
IPS
STS
LS
B
target on
t stop
trials
1 sec
MSTd
7a
dlPFC
C
sensorimotor
linear vel.
linear accel.
ang. vel.
ang. acc.
t. move on
t. move off
target on
latent
target dist.
target ang.
trav. dist.
trav. ang.
LFP
theta phase
alpha phase
beta phase
other
reward
eye vert.
eye hori.
non linear filters
spikes
other neurons
D
data
neurons
model
neurons
50
0
spike/sec
5 sec
E
Lin. vel.
Lin. acc.
Ang. vel.
Ang. acc.
t. move on
t. move off
t. target on
Targ. dist.
Targ. ang.
Trav. dist.
Trav. ang.
Theta Ph.
Alpha Ph.
Beta Ph.
t. reward
Eye vert.
Eye hori.
Spikes s-1
40
15
15
0
15
0
p<0.01
p>0.01
MSTd P-GAM
7a P-GAM
dlPFC P-GAM
Raw
F
MSTd
7a
dlPFC
Fraction Tuned
1
0
Lin. vel.
Lin. acc.
Ang. vel.
Ang. acc.
t. move on
t. move off
t. target on
Targ. dist.
Targ. ang.
Trav. dist.
Trav. ang.
Beta Ph.
Alpha Ph.
Theta Ph.
t. reward
Eye vert.
Eye hori.
Sensorimotor
Latent
LFP
Other

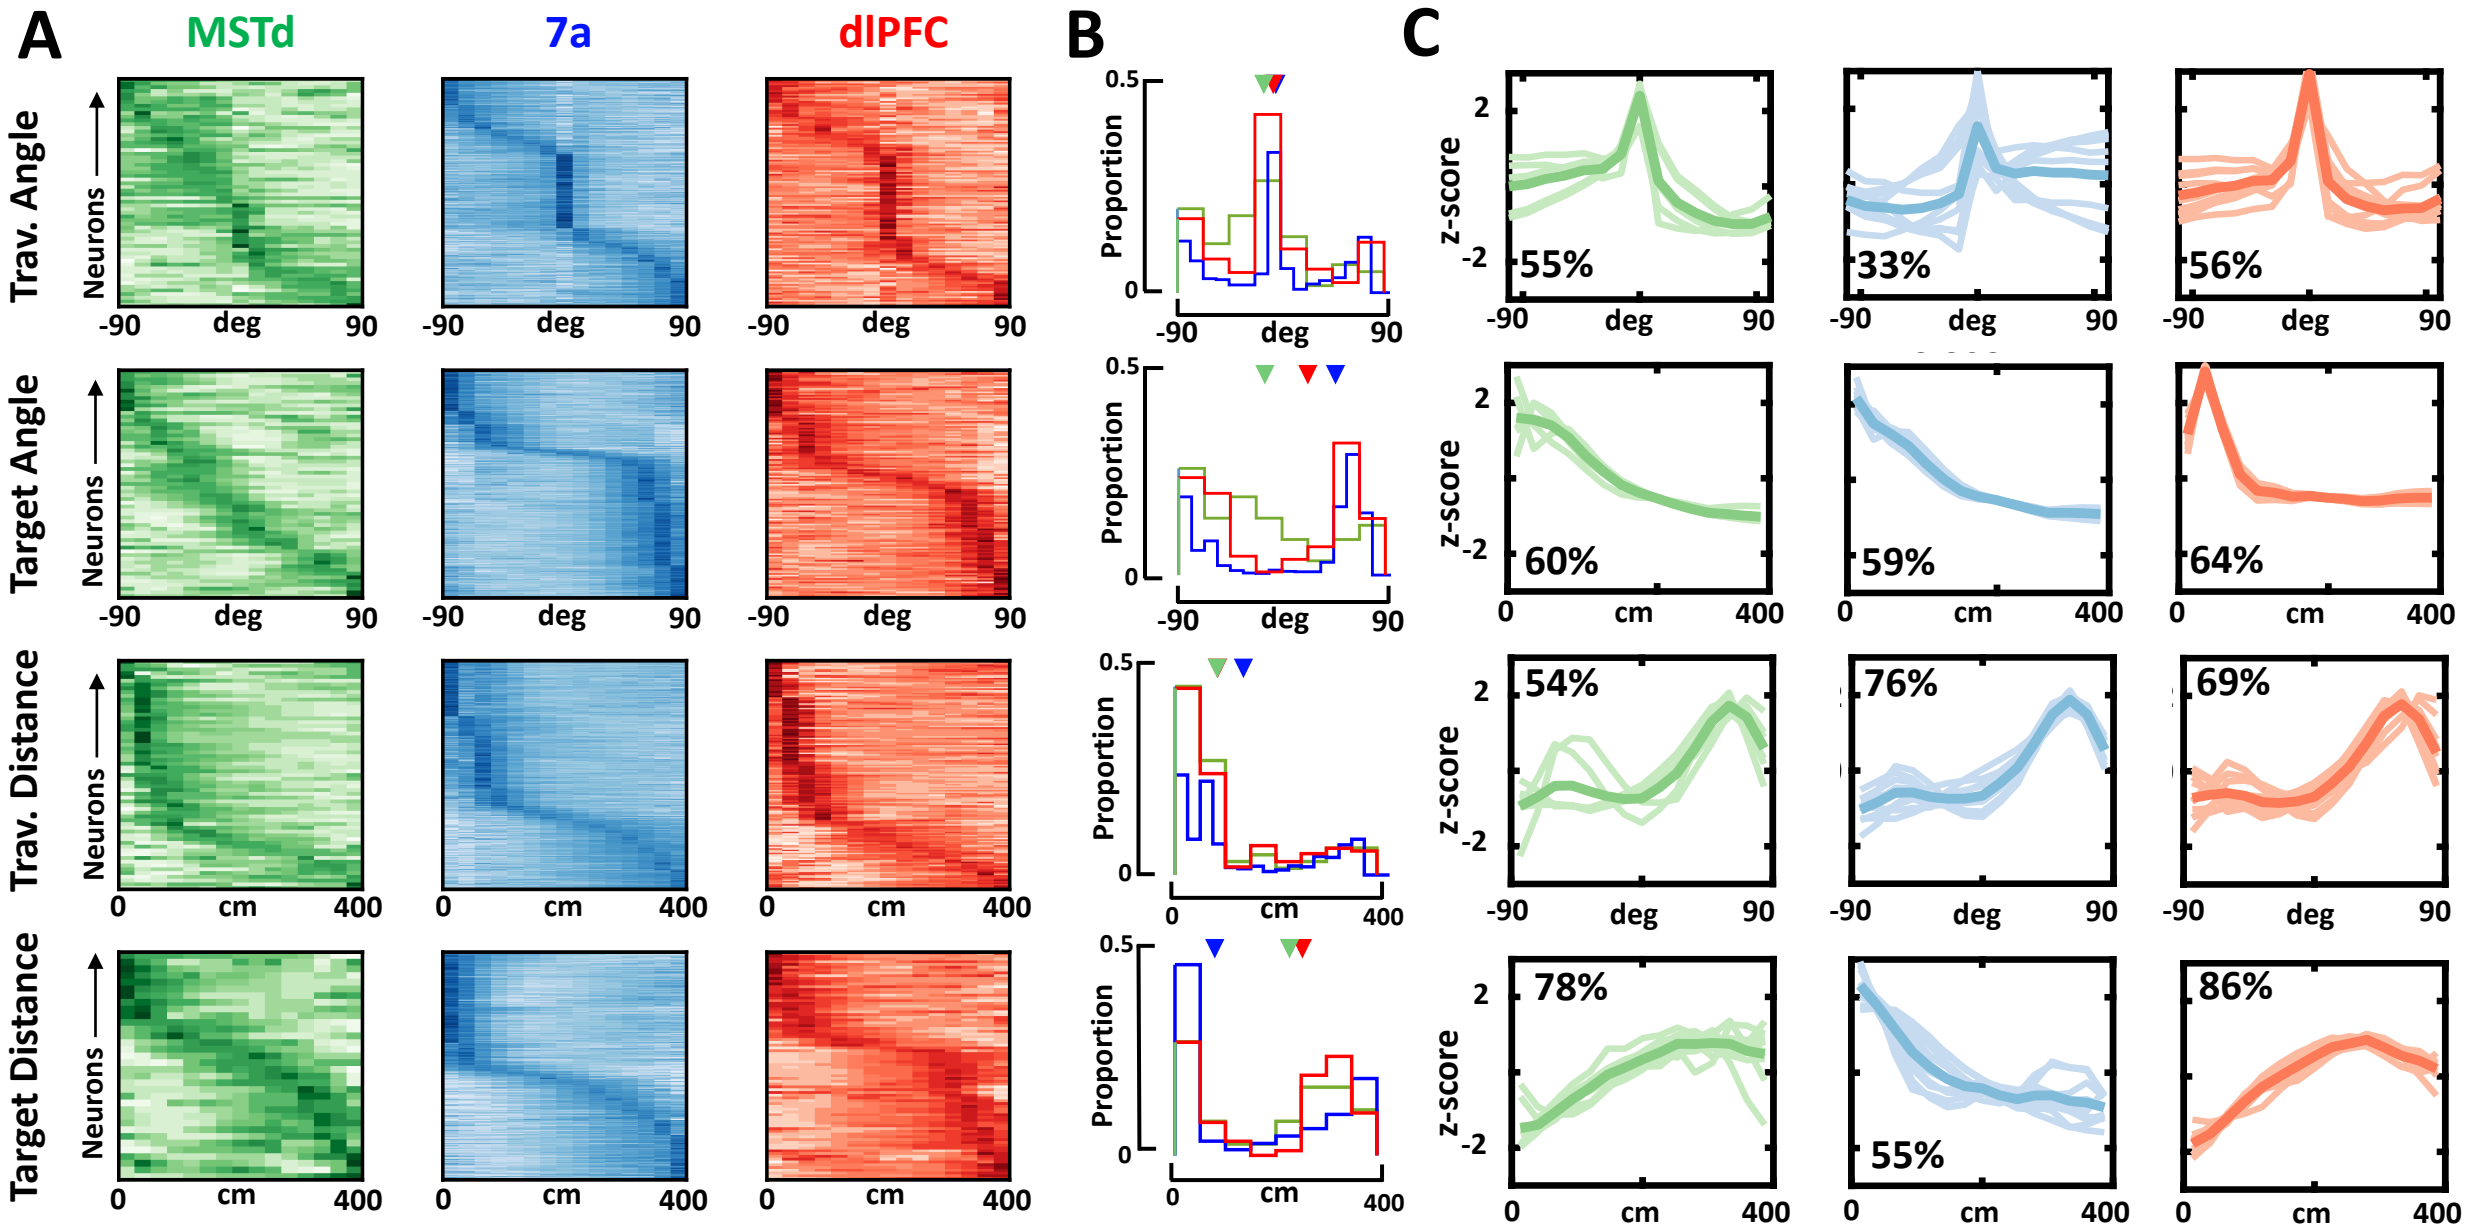
A
MSTd
7a
dlPFC
B
C
Trav. Angle
Target Angle
Trav. Distance
Target Distance
Neurons
Proportion
z-score
-90
deg
90
0
cm
400
0.5
2
-2
55%
33%
56%
60%
59%
64%
54%
76%
69%
78%
55%
86%

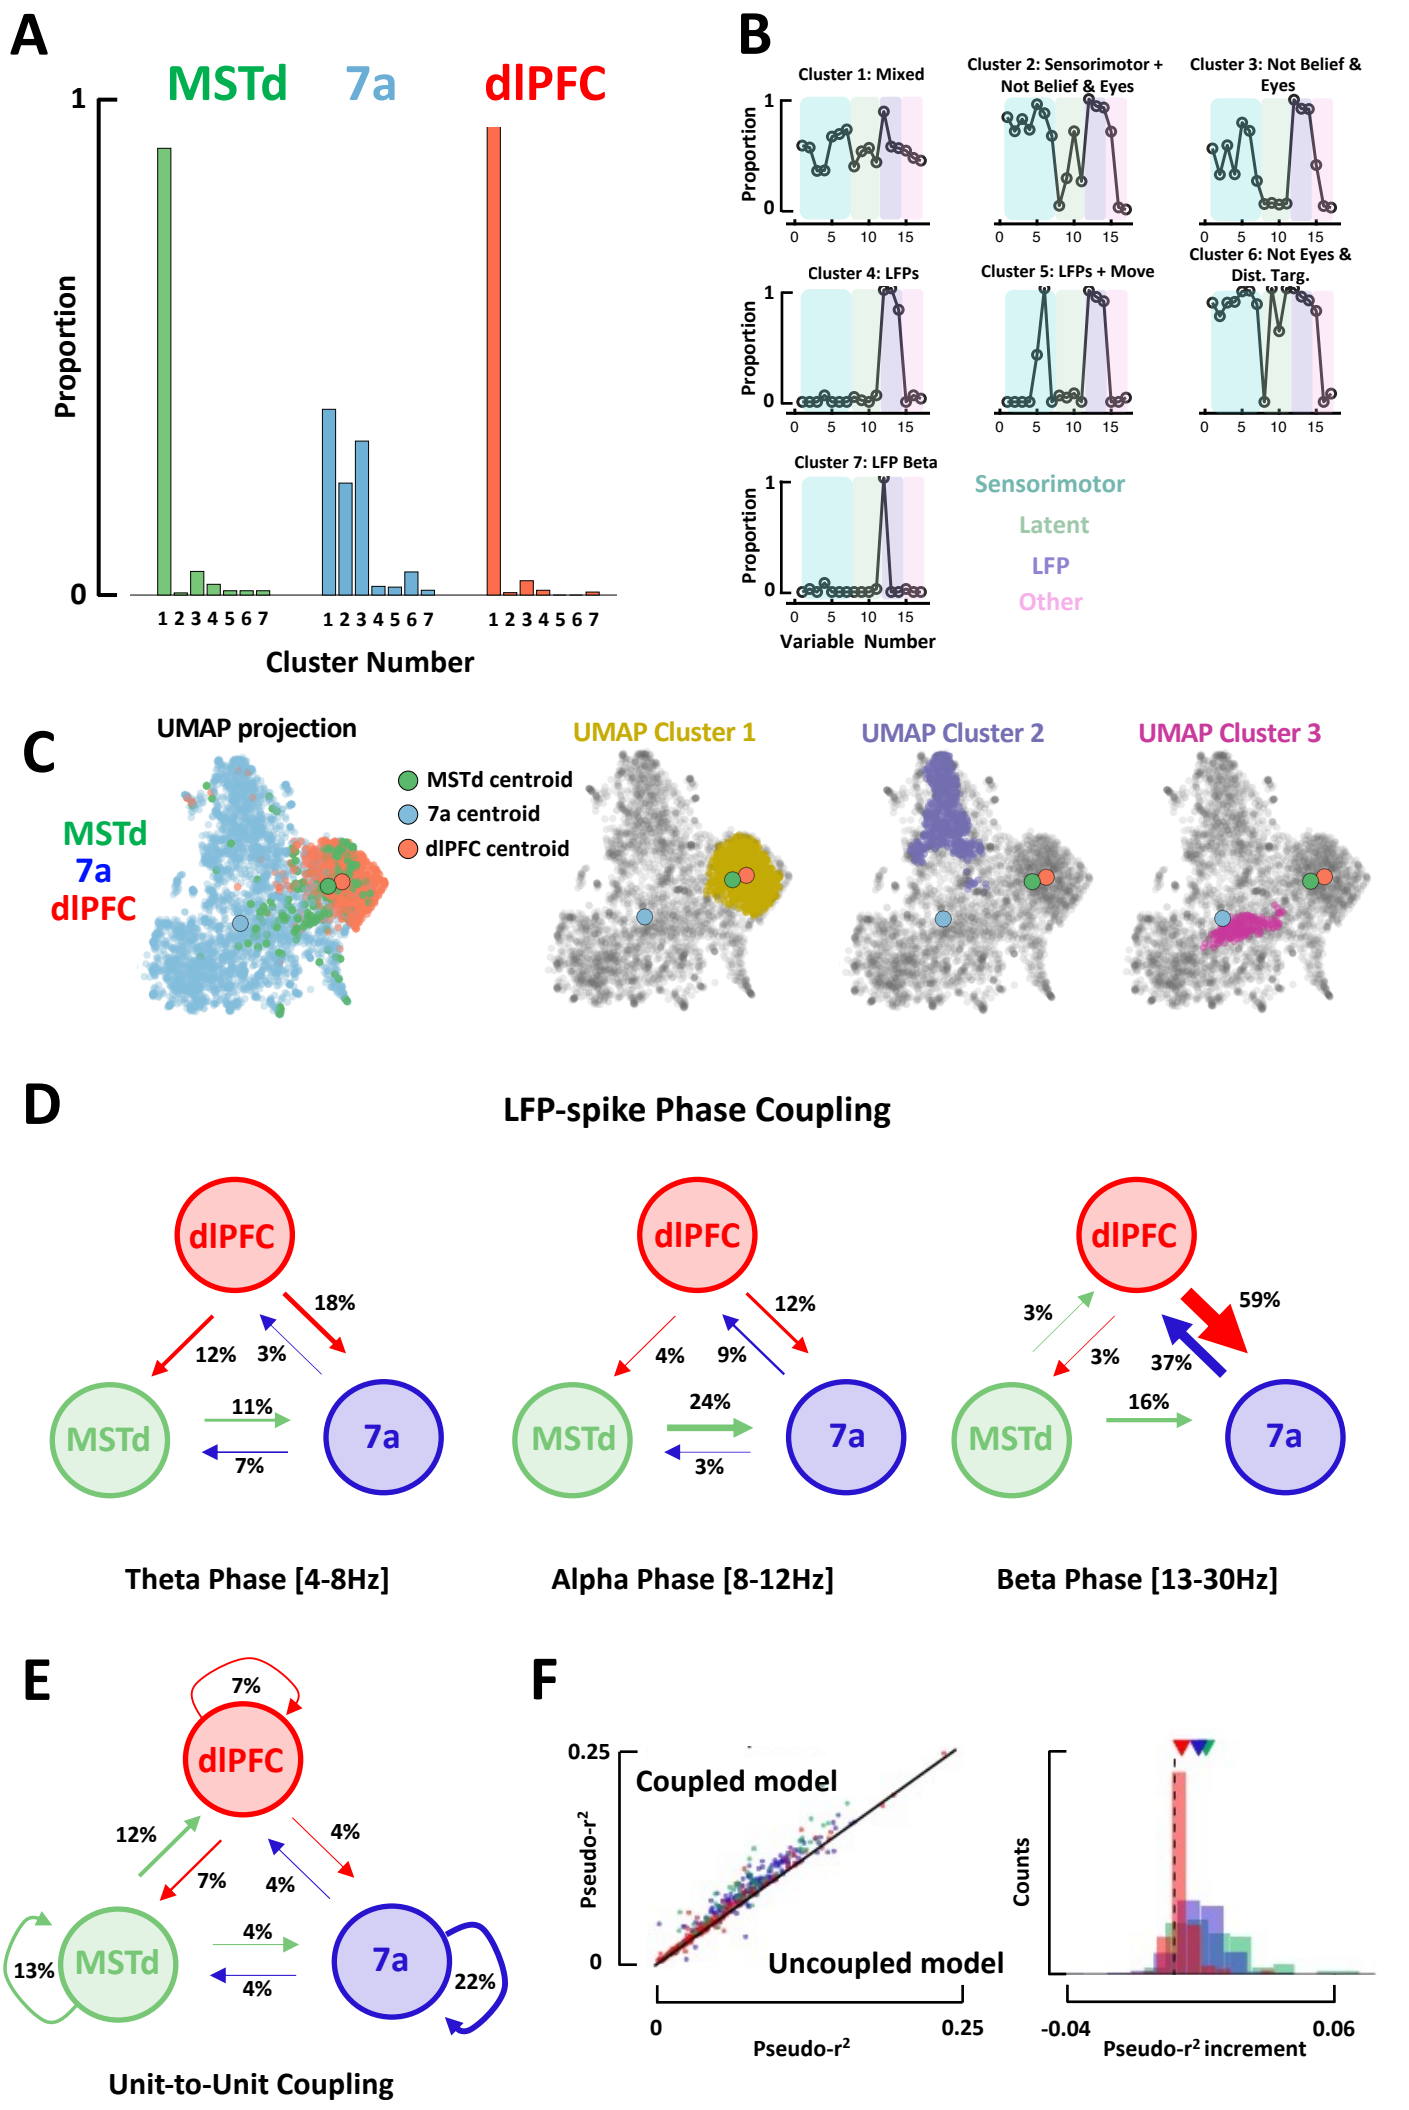
A
MSTd
7a
dlPFC
Proportion
1
0
1 2 3 4 5 6 7
Cluster Number
B
Cluster 1: Mixed
Cluster 2: Sensorimotor + Not Belief & Eyes
Cluster 3: Not Belief & Eyes
Cluster 4: LFPs
Cluster 5: LFPs + Move
Cluster 6: Not Eyes & Dist. Targ.
Cluster 7: LFP Beta
Proportion
0 5 10 15
Variable Number
Sensorimotor
Latent
LFP
Other
C
UMAP projection
MSTd
7a
dlPFC
MSTd centroid
7a centroid
dlPFC centroid
UMAP Cluster 1
UMAP Cluster 2
UMAP Cluster 3
D
LFP-spike Phase Coupling
dlPFC
MSTd
7a
18%
12%
3%
11%
7%
12%
4%
9%
24%
3%
3%
59%
3%
37%
16%
Theta Phase [4-8Hz]
Alpha Phase [8-12Hz]
Beta Phase [13-30Hz]
E
7%
12%
7%
4%
4%
13%
4%
4%
22%
Unit-to-Unit Coupling
F
0.25
Coupled model
Pseudo-r²
0
Uncoupled model
0
0.25
Pseudo-r²
Counts
-0.04
0.06
Pseudo-r² increment

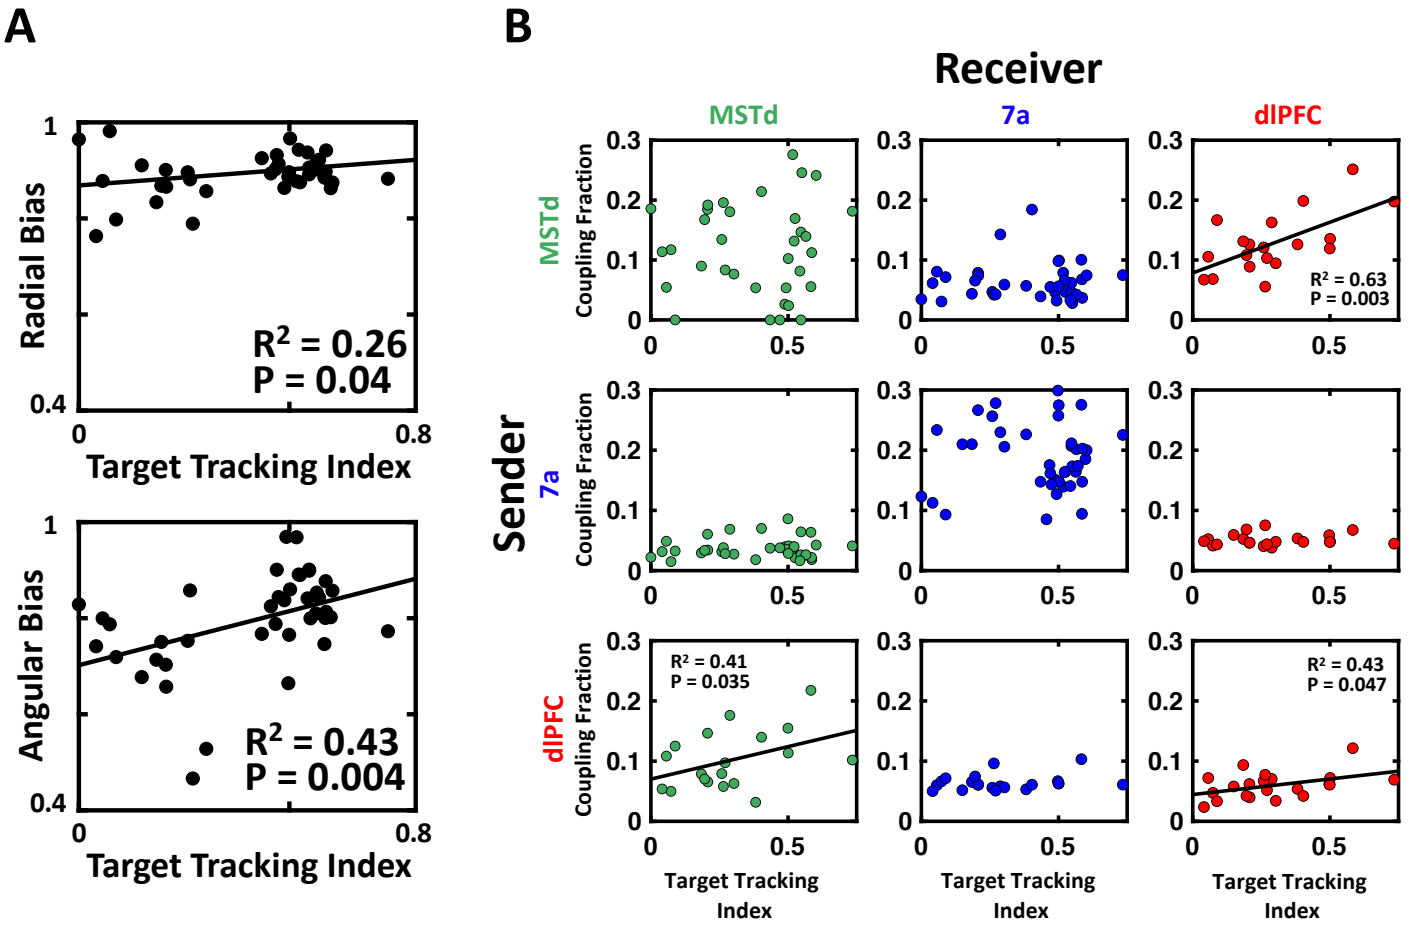
A
Radial Bias
1
0.4
0
0.8
Target Tracking Index
R2 = 0.26
P = 0.04
Angular Bias
1
0.4
0
0.8
Target Tracking Index
R2 = 0.43
P = 0.004
B
Receiver
MSTd
7a
dlPFC
Sender
MSTd
7a
dlPFC
Coupling Fraction
0.3
0.2
0.1
0
0.5
R² = 0.63
P = 0.003
R² = 0.41
P = 0.035
R² = 0.43
P = 0.047
Target Tracking Index