

# The midbrain reticular formation in contextual control of perceptual decisions

## Highlights

- A novel task requires mice to remap visual stimuli to context-appropriate actions
- The midbrain reticular formation encodes abstract context that predicts flexible behavior
- Task-relevant visual responses emerge in the midbrain reticular formation with training
- Context-coding neurons and cortical inputs are spatially organized and aligned

## Authors

Jordan R. Shaker, Jeremy N. Schroeter, Daniel Birman, Nicholas A. Steinmetz

## Correspondence

nick.steinmetz@gmail.com

## In brief

Shaker et al. show that the midbrain reticular formation, a region traditionally associated with motor functions, contains abstract contextual signals that predict choices in a flexible decision-making task. The findings support distributed models of cognitive control, wherein conserved brainstem circuitry serves as a key node for guiding flexible behavior.

Article

# The midbrain reticular formation in contextual control of perceptual decisions

Jordan R. Shaker,<sup>1,2</sup> Jeremy N. Schroeter,<sup>1,3</sup> Daniel Birman,<sup>1,4</sup> and Nicholas A. Steinmetz<sup>1,5,\*</sup>

<sup>1</sup>Department of Neurobiology and Biophysics, University of Washington, Seattle, WA, USA

<sup>2</sup>Medical Scientist Training Program, University of Washington, Seattle, WA, USA

<sup>3</sup>Present address: Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

<sup>4</sup>Present address: Allen Institute for Neural Dynamics, Seattle, WA, USA

<sup>5</sup>Lead contact

\*Correspondence: [nick.steinmetz@gmail.com](mailto:nick.steinmetz@gmail.com)

<https://doi.org/10.1016/j.neuron.2026.04.003>

## SUMMARY

Flexibly responding to sensory cues is fundamental to animal behavior. However, growing evidence suggests that the relevant neural circuits have been incompletely resolved. We trained mice to apply an abstract rule that maps identical visual stimuli to opposing actions across contexts and investigated the role of the midbrain reticular formation (MRF) in this process. Large-scale recordings revealed that neurons in the MRF, along with several canonical decision-making areas, maintained persistent representations of task context in pre-stimulus activity. These representations predicted divergent population dynamics, putatively routing stimuli toward context-appropriate actions. The MRF was the only recorded region containing neuronal populations that both predicted lapses in contextual control pre-stimulus and exhibited contextually modulated premotor responses. The MRF also showed task-specific sensory plasticity. Context-coding neurons were aligned with cortical inputs, suggesting a spatial organization. These findings support the hypothesis that the MRF is a key node for setting and implementing abstract contextual states within the distributed circuitry for flexible perceptual decisions.

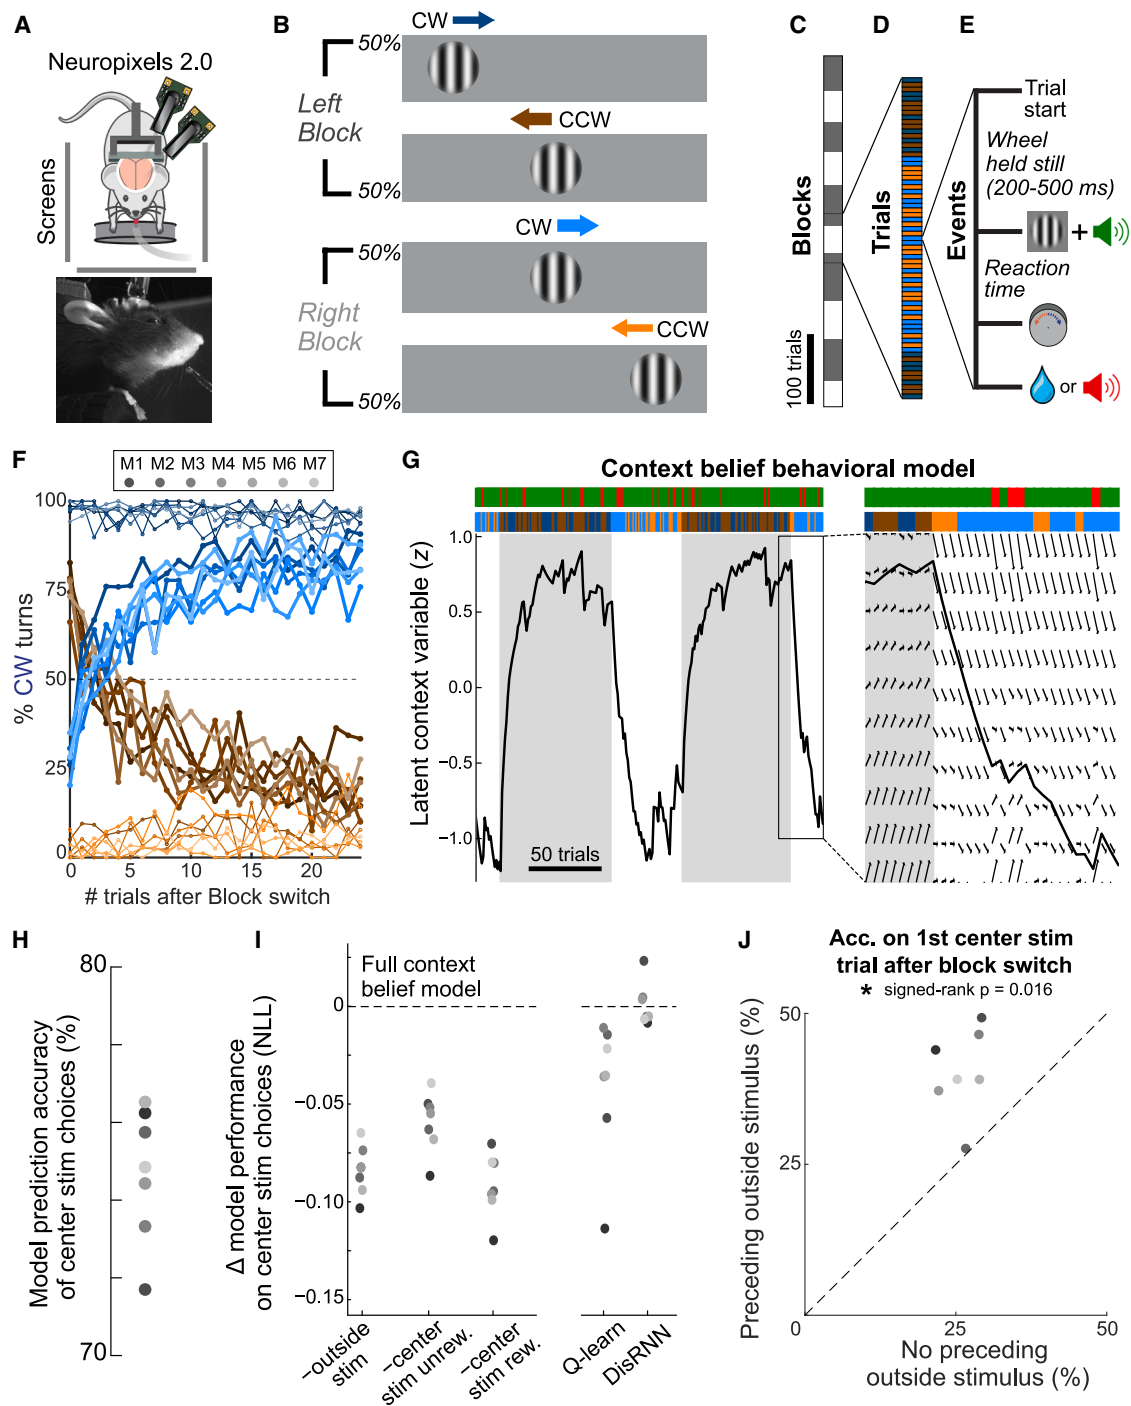
## INTRODUCTION

Our response to a warning siren differs greatly if we know it is part of a drill versus when it sounds off unexpectedly. This ability to flexibly adapt one's actions to identical sensory inputs based on learned rules is a hallmark of intelligent mammalian behavior.<sup>1</sup> The neural circuitry underlying aspects of flexible behavior appears widely distributed across the brain.<sup>2–8</sup> Indeed, brain-wide recordings in mice during perceptual decision-making tasks have revealed strongly task-linked neural responses in regions classically studied as being relevant for flexible decision-making and in areas not previously assumed to be part of decision-making networks.<sup>9–12</sup> Prominent among these is the midbrain reticular formation (MRF), where neurons respond to task sensory stimuli,<sup>9–12</sup> around the time of choices,<sup>9–11</sup> and in a manner modulated by expectations.<sup>12,13</sup>

An open question is whether similarly distributed sets of regions participate in cognitive control during the decision-making process. Generally, tasks requiring cognitive control involve abstract rules, which govern flexible responses to stimuli based on contextual information, as opposed to concrete rules, which describe fixed stimulus-response associations.<sup>8,14,15</sup> Foundational work posits that the prefrontal cortex (PFC) learns and maintains abstract rule representations that provide top-down bias signals to lower sensorimotor circuitry to enable flexible behavior.<sup>8,16–22</sup>

However, this hierarchical control framework is increasingly challenged by studies in both primates and rodents showing that abstract rules are similarly implemented by subcortical structures, particularly the dorsal striatum<sup>7,23–25</sup> of the basal ganglia, the mediodorsal nucleus (MD) of the thalamus,<sup>5,26</sup> and the superior colliculus (SC) of the midbrain.<sup>6,27–29</sup> Theoretical models incorporating these findings suggest that various cognitive control processes are mediated by simultaneous computations in distributed networks of regions,<sup>2,4,30,31</sup> supported by recurrent loops within cortico-basal ganglia-collicular-thalamic connections.<sup>32–35</sup> Importantly, these well-studied loops interconnect densely with other regions whose specific contributions to flexible behavior remain less defined.

One such structure, uniquely positioned within these recurrent loops, is the MRF. The MRF, a large, understudied, and poorly differentiated portion of the midbrain, possesses a strikingly integrative connectivity profile.<sup>36–38</sup> Distinct among subcortical regions, the MRF receives highly convergent input collaterals from all parallel output streams of the basal ganglia via the substantia nigra pars reticulata (SNr).<sup>39</sup> This basal ganglia convergence exists alongside direct inputs from broad cortical areas<sup>32,40–42</sup> and strong reciprocal connections with the SC.<sup>35</sup> Moreover, MRF stimulation has been shown to influence cortical dynamics through its thalamic projections,<sup>43</sup> which include cognition-related nuclei (notably, MD) and inhibitory controllers of the thalamus (thalamic reticular nucleus [RT] and



**Figure 1. Mice remap stimulus-action associations in a flexible decision-making task by maintaining an internal context-belief variable**

(A) Schematic of task performance and Neuropixels recordings. Mice sat with forepaws on a wheel in front of three screens where visual stimuli were presented. One or two Neuropixels 2.0 probes were inserted into target brain regions during behavior.

(B) Flexible stimulus-action remapping task schematic. There are 2 blocks (left and right), and within each block, there is a 50/50 probability of two Gabor stimulus locations:  $-80^\circ$  or  $0^\circ$  in the left block, and  $+80^\circ$  or  $0^\circ$  in the right block. The wheel must be turned CW or CCW such that the stimulus is moved to the middle of the two possible stimulus locations (i.e., a  $\pm 40^\circ$  turn). Colors of arrows indicate trial type: blue colors indicate CW wheel turns, and orange colors indicate CCW wheel turns; darker hues indicate the left block, and lighter hues indicate the right block.

(C) Block structure for an example session, representing 500 trials. Block lengths were randomized between 40 and 70 trials.

(D) Trial structure across blocks, showing randomized order and equal frequency of trials requiring CW and CCW responses.

(legend continued on next page)

the zona incerta),<sup>32,44–46</sup> as well as through its basal forebrain projections.<sup>47,48</sup> Early theories of MRF function, inspired by its unique anatomical arrangement and the powerful effects of electrical stimulation and lesions of the MRF on arousal and cortical synchrony,<sup>49</sup> emphasized its involvement in the higher-order setting of states for behaviorally relevant computations in the forebrain.<sup>38,50</sup> However, upon the discovery of neuromodulators, theories about the neural mechanisms promoting arousal have instead largely focused on nearby brainstem neuromodulatory nuclei,<sup>51–53</sup> leaving the majority of the MRF uncharacterized in behavioral tasks. Although small subregions of the MRF have been studied in motor control (e.g., saccades,<sup>54</sup> locomotion,<sup>55,56</sup> and action execution<sup>43</sup>), the role of intrinsic MRF neurons in cognitive control has not been tested.

In this study, we developed a novel decision-making task, in which mice learned to apply an abstract rule to flexibly remap stimuli to actions based on an inferred context. Large-scale electrophysiological recordings revealed that the MRF, alongside canonical decision-making circuits in the frontal cortex, basal ganglia, and SC, persistently represented this inferred context in their pre-stimulus activity. This pre-stimulus context signal was associated with divergent population dynamics underlying subsequent sensorimotor transformations. Notably, context representations in the MRF were predictive of flexible choice behavior and were accompanied by context-modulated premotor responses, a unique combination of features among all recorded areas. We also found that, after task training, MRF neurons developed robust visual responses specifically tuned to behaviorally relevant stimuli, indicating experience-dependent plasticity for adaptive sensorimotor behaviors. Within the MRF, context-coding neurons were enriched in the anterior aspect, aligned with the densest cortical inputs, suggesting a spatial organization of function. Together, these findings suggest that the MRF is a key participant in the network of regions that set and implement rule-based contextual states for flexible sensorimotor behavior, with a role in cognitive control that extends beyond previously ascribed motor functions.

## RESULTS

### Mice learn abstract rules and infer task context to flexibly remap stimulus-action associations

To investigate the neural basis of context-dependent sensorimotor remapping, we trained head-fixed mice to perform a novel

flexible decision-making task (Figures 1A–1E). On each trial, a single visual stimulus appeared at one of three possible screen locations: left (azimuth =  $-80^\circ$ ), center (azimuth =  $0^\circ$ ), or right (azimuth =  $+80^\circ$ ). Stimuli were large ( $7^\circ$  of visual angle), full-contrast Gabors, making them perceptually unambiguous to ensure that choices primarily reflected the application of task rules rather than challenges in sensory discrimination. Upon stimulus onset, mice were free to indicate their choice by turning a wheel clockwise (CW) or counterclockwise (CCW) with their forelimbs<sup>9,11,57</sup> (Figures 1A, 1B, and 1E). The task alternated between two block types, each lasting 40–70 trials (randomized), with no cues to indicate the transitions. In “left blocks,” stimuli appeared randomly on either the left or center screen with equal probability; in “right blocks,” stimuli appeared randomly on either the right or center screen (Figures 1B–1D). To achieve a correct response for reward, subjects had to choose the wheel turn direction that moved the stimulus toward the midpoint between the two possible locations for the current block (Figure 1B). The task thus had two trial categories: “outside stimulus” trials (left and right locations,  $-80^\circ$  and  $+80^\circ$  azimuth), which had fixed stimulus-response mappings, and “center stimulus” trials ( $0^\circ$  azimuth), which required flexible responses depending on the block. The task shares some similarities and key differences with previous context-dependent, rule-based perceptual decision-making tasks for rodents<sup>18,27,58–61</sup> (further described in STAR Methods).

Mice performed the task accurately, with consistent behavior across mice (Figure 1F). Performance on outside stimulus trials was above 90% (mean  $\pm$  standard error of the mean [SEM] accuracy on outside stimulus trials across mice =  $94\% \pm 0.9\%$ ,  $n = 7$  mice, 257 sessions). In contrast, performance on center stimulus trials after uncued block switches started below chance, indicating use of the previous block’s rule, and quickly increased to above chance ( $5.0 \pm 0.48$  trials to reach 60%,  $10.7 \pm 1.1$  trials to reach 75%), with saturating performance of  $79\% \pm 1.8\%$  toward the end of a block (last 20 trials before a block switch).

The mice’s high level of performance indicates that they inferred the task’s context-dependent sensorimotor mappings rather than relying on purely motor-biasing strategies. The equal probability of outside or center stimuli on each trial meant that rewarded actions (CW and CCW) were balanced throughout a session and independent of context (Figures 1B–1D). Thus, unlike conventional reversal learning paradigms in which subjects

(E) Timeline of a single trial. Trial start was uncued, and mice held the wheel still for 200–500 ms for the visual stimulus to appear. An auditory go cue played at the same time as the visual stimulus appearance. Mice were free to move the wheel upon presentation of the stimulus, and turns in the correct direction were rewarded with sucrose water drops, whereas turns in the incorrect direction resulted in a white-noise burst plus timeout.

(F) Percentage CW turns for each trial type as a function of number of trials after block switch. Data from 7 mice are plotted; percentage correct is compiled across all included sessions for a given mouse ( $n = 37 \pm 6$  behavioral sessions per mouse). Traces are colored by trial type according to the convention in (B).

(G) Example session latent context variable ( $z$ ) from the context-belief behavioral model fit to center stimulus choices of each subject. Trial outcome and type are displayed above the latent trace with color scheme as in (D) and (F). Inset: zoom-in on a subset of trials with vectors representing the latent update depending on trial type/outcome and current latent value.

(H) Context belief model accuracy in fitting each subject’s center stimulus choices (% correctly predicted choices).

(I) Contribution of different trial types to the context-belief model’s performance and comparison with Q-learning and DisRNN. All  $\Delta$ NLL are for fits to center stimulus choices. Model was constrained to freeze latent on outside stimulus, center stimulus unrewarded, or center stimulus rewarded trials, and deltas in performance are plotted for each subject. WLSL performance shown in Figure S1B.

(J) Percentage correct on center stimulus trials after a block switch, conditioned on whether the first trial post-switch was an outside or center stimulus.

See also Figure S1.

alternate between distinct actions,<sup>62</sup> mice could not have solved the task by simply biasing global action selection between blocks using exclusively model-free reinforcement learning strategies<sup>63</sup> (e.g., Q-learning<sup>64</sup>). Mice instead continuously and correctly updated their responses to identical center stimuli (Figure 1F), with the correct stimulus-action mapping determined by the experimenter-defined block identity. Because the block identity was not explicitly cued, mice could have inferred the current center stimulus mapping using two separate, but not mutually exclusive, strategies: (1) using the previous outside stimulus position (i.e., left or right screen) to deduce subsequent correct center stimulus responses and/or (2) applying reinforcement learning selectively to the center stimulus trials by integrating action-outcome history separately for this stimulus type. Adopting the first strategy would indicate that mice learned the task's abstract rules and structure, as it requires using the arbitrary location of one stimulus (outside) as contextual information to flexibly update the response to a different stimulus (center), a process independent of direct reward reinforcement. Adopting the second strategy would indicate that mice used a distinct form of model-based rule learning in which the stimulus itself defines the context for action-outcome reinforcement.<sup>63</sup> To disambiguate the extent to which mice used either or both of these strategies, we next built behavioral models fitted to each mouse's choices during the task.

Behavioral modeling suggested that mice tracked a single latent variable representing their belief about the current task context. To explore which strategies mice used in an unbiased manner, we first leveraged disentangled recurrent neural networks (DisRNN),<sup>65</sup> which are designed to discover parsimonious and interpretable behavioral models directly from data (Figures S1A–S1E). DisRNNs accurately predicted held-out choices on center stimulus trials (75.7% ± 0.3% accuracy) and, importantly, outperformed simpler, canonical reinforcement learning models that were specifically designed to solely update choices based on action-reward history (center stimulus-specific win-stay lose-switch (WSLS)<sup>66</sup> and center stimulus-specific Q-learning<sup>64</sup> models; see STAR Methods) (Figures 1I and S1B; mean ± SEM Δnegative log-likelihood [NLL] = −0.66 ± 0.02, −0.042 ± 0.01;  $p = 0.016, 0.016$ , Wilcoxon signed-rank test,  $n = 7$  mice). This indicates that mice used a more sophisticated strategy than solely integrating the action-reward history of center stimulus trials. Across all subjects, DisRNN consistently discovered a model using a single latent variable that evolved according to linear, observation-dependent updates (Figures S1C and S1D; STAR Methods) and mapped linearly to log choice probability (Figure S1E). This latent variable co-evolved with the task block structure (Figure S1C) and only influenced choices on center stimulus trials (Figure S1E), consistent with a representation of the subject's belief about the task context.

Mice updated their context belief using a combination of center stimulus action-reward history and the identity of outside stimuli. To determine which factors influenced the latent variable discovered by DisRNN, we defined a simplified form of the model it discovered as a set of equations that we call the “context-belief model.” This model achieved predictive performance on center stimulus choices indistinguishable

from DisRNN (Figures 1G–1I,  $p = 0.578$  vs. DisRNN, Wilcoxon signed-rank test). This model is given as:

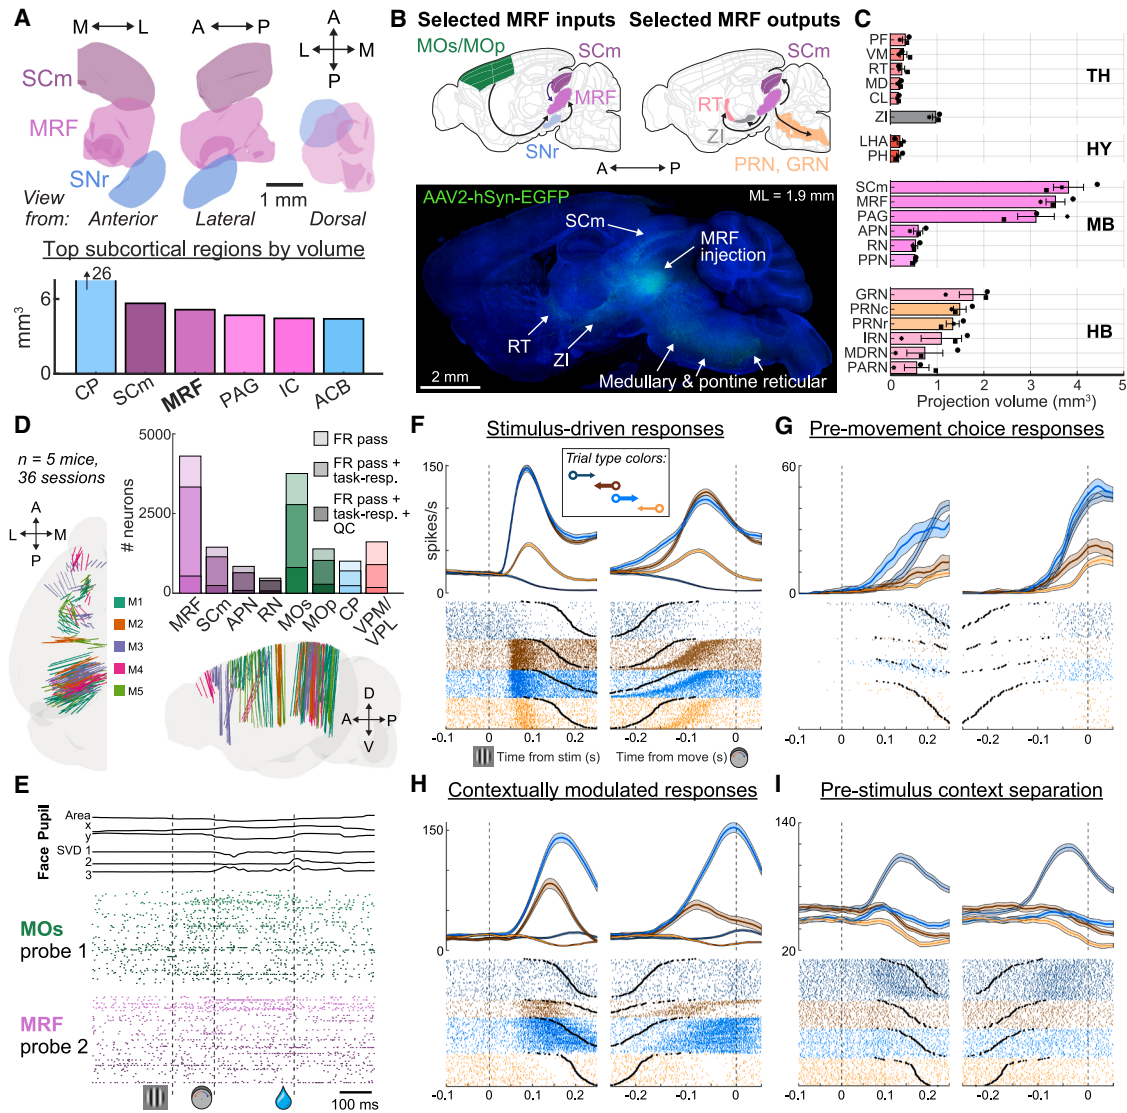
$$z_{t+1} = \alpha_o z_t + \beta_o, \quad \mathbf{o}_t = \begin{bmatrix} s_t \\ a_t \\ r_t \end{bmatrix} \quad \log \frac{P(a_t = CCW)}{P(a_t = CW)} = \rho z_t + \delta,$$

where  $z_t$  represents the subject's context belief, updated linearly with parameters  $\alpha$  and  $\beta$  that depend on the previous trial observation  $\mathbf{o}_t = [s_t, a_t, r_t]$  (stimulus, action, and reward; 8 possible configurations). Center stimulus choice probability follows a logistic function of  $z_t$  with parameters controlling sensitivity ( $\rho$ ) and bias ( $\delta$ ). To determine which trial types drove context-belief updates, we tested model performance after selectively removing the model's ability to update  $z_t$  from specific configurations of  $\mathbf{o}_t$ . This analysis revealed that mice used a combination of both strategies outlined above to infer the block context. Action-reward history from center trials contributed most to predictive performance (Figure 1I, ΔNLL = −0.059 ± 0.006, −0.091 ± 0.006 for unrewarded or rewarded center stimulus trials;  $p = 0.016, 0.016$ , Wilcoxon signed-rank test). Outside stimulus trials also significantly enhanced accuracy (Figure 1I, ΔNLL = −0.084 ± 0.005,  $p = 0.016$ ). Supporting these model-derived conclusions, we found that exposure to a single outside stimulus, unrewarded center stimulus, or rewarded center stimulus trial immediately after a block switch was sufficient to significantly raise accuracy on the subsequent center stimulus trial (Figures 1J, S1F, and S1G;  $p = 0.016$  for all comparisons). The magnitude of the behavioral improvement followed the hierarchy predicted by the model: rewarded center (Δ% accuracy = 22.3% ± 5.0%) > outside stimulus (14.3% ± 2.7%) > unrewarded center (10.0% ± 1.5%). As described above, the mice's reliance on outside stimulus identity to flexibly update center stimulus responses indicates that they learned the task's abstract rules and structure. Our behavioral model captured each mouse's implementation of this strategy through a latent variable representing belief about the current task context. In subsequent neural analyses, this variable is used as the measure for internal contextual state rather than the experimenter-defined block identity, unless specified otherwise.

In sum, mice in this task exhibited a form of flexible cognitive control that can be disambiguated from pure motor plans. We next considered the neural mechanisms underlying this capability.

### MRF neurons connect to decision-making circuitry and contain diverse responses to task events

The MRF, comprising the third-largest subcortical brain region by volume (Figure 2A, based on the Allen Common Coordinate Framework [CCF] v3 atlas<sup>67</sup>), has efferent and afferent connectivity with key nodes within the circuitry for flexible decision-making and motor control (Figure 2B). The MRF contains ~250,000 neurons in the mouse<sup>68</sup> and sits medial and dorsal to the SNr and ventral to the SC. Its cytoarchitectonic boundaries are poorly defined and generally delineated around more well-defined structures, including the red nucleus (RN) (embedded within the ventral-anterior aspect of the MRF) and superior cerebellar peduncle (courses through the ventral-posterior aspect of the MRF) (Figure 2A). To determine the outputs of the MRF, we



**Figure 2. Dense electrophysiological recordings of the MRF and connected regions during the flexible stimulus-action remapping task show diverse trial responses**

(A) Top: 3D renderings of the MRF along with adjacent SCm and SNr for reference. Bottom: top 6 subcortical brain regions by volume according to Allen CCF atlas.<sup>67</sup>

(B) Top: sagittal CCF atlas slices showing selected input and output connections of the MRF. Bottom: sagittal image of the MRF injection site and outputs from an example viral tracing experiment with AAV2-hSyn-EGFP injected into the MRF.

(C) Quantification of projection volumes (mean  $\pm$  SEM) for a subset of MRF outputs from AAV2-hSyn-EGFP injections in three mice (marked by square, circle, or star symbols; see Figure S2 for full list of injection sites). Imaging stacks were transformed to CCF, and projection volumes were quantified as the number of voxels passing threshold fluorescence within the CCF-defined region after background and artifact corrections (STAR Methods).

(D) 3D view of all Neuropixels 2.4 insertions from task recordings from dorsal (left) and lateral (bottom right). All insertions are projected onto the left hemisphere for visualization. Top right, number of neurons recorded from a set of 8 regions that have trial firing rates  $\geq 0.1$  sp/s (“FR pass”), that are responsive during the task (“task-resp.”), and/or that pass single-unit quality-control criteria (“QC”; STAR Methods).

(E) Example spike rasters during a dual-probe recording in the MRF and MOs, with task events overlaid and face video motion-energy SVDs, pupil position, and pupil area above.

(F–I) PETHs by trial type (filtered for non-repeat correct trials) for example MRF neurons. Spikes are aligned to stimulus onset on the left and wheel turn onset on the right. Top traces are trial-averaged PETHs with 1 SEM shaded, bottom are spike rasters sorted by trial type (color) and reaction time within each trial type (black dots indicate wheel turn times on the left plot and stimulus times on the right plot). Inset in (F): color legend for the 4 trial types as in Figure 1B (top to bottom: left stim, center stim left block, center stim right block, right stim). (F) MRF neuron with robust stimulus-driven responses (strong increase in firing rate after center stimuli with consistent latencies across trials; moderate increase and suppression for left and right stimuli). (G) Choice-responsive MRF neuron (pre-wheel turn firing rate higher for trials with CW choices than CCW choices). (H) MRF neuron showing contextual modulation of its center stimulus response (center stimulus response on right block > left block). (I) MRF neuron with pre-stimulus separation in firing rates by context (left block > right block).

See also Figures S2 and S3.

performed systematic brain-wide mapping of its axonal projections by injecting anterograde viral vectors (AAV2-hSyn-EGFP), followed by brain clearing<sup>69</sup> and atlas-registered light-sheet imaging ( $n = 3$  mice; **Figures 2B, 2C, and S2A–S2E**). Injection sites were predominantly within the MRF (75%–85% within the MRF) and targeted its intermediate and dorsal aspects (**Figures S2A and S2B**). MRF axons were found most prominently in the SC (particularly intermediate and deep layers, collectively the SCm), zona incerta, RT, parafascicular (PF), MD, ventromedial (VM), and central lateral (CL) thalamus, as well as parts of the pons and medulla. Notable smaller projections include cholinergic basal forebrain nuclei (nucleus of the diagonal band and substantia innominata) as well as the lateral and posterior hypothalamus (**Figures 2C and S2C–S2E**). These findings confirm, and comparatively quantify, prior reports of MRF output connections from the literature.<sup>32,35,43–48</sup> Taken together, the MRF's connectivity places it within the recurrent loops underlying flexible behavior,<sup>4,30,70</sup> as it receives inputs from the frontal and motor cortex,<sup>40–42</sup> SC,<sup>35</sup> and SNr<sup>32,39</sup> and sends outputs to the frontal and motor cortical-projecting thalamic nuclei, inhibitory thalamic controllers, and SC.

Based on MRF anatomy, connectivity, and prior decision-making literature, we performed 4-shank Neuropixels 2.0<sup>71</sup> recordings in the MRF, as well as targeted areas of the forebrain and midbrain ( $n = 5$  mice, 36 sessions [1–2 probes per session], 3,378 task-responsive, quality-controlled neurons, 2,444 neurons among the 8 regions in **Figure 2D**). We recorded from regions that strongly encoded decision variables in related tasks, including the secondary motor cortex (MOs), primary motor cortex (MOp), deep layers of the SCm, and the caudoputamen (CP).<sup>9–12,23,27,72</sup> We also recorded from two nearby midbrain regions, the anterior pretectal nucleus (APN) and RN, as well as the somatosensory parts of the thalamus, the ventral posteromedial and ventral posterolateral thalamus (VPM/VPL), as comparison regions. To reliably characterize the spatial coding of task variables within the MRF, recordings were targeted to densely tile the entire region (**Figure 2D**). During all electrophysiological recordings, we simultaneously recorded the subject's eye and front body using high-speed videography to assess the contributions of pupil dynamics and spontaneous body movements to neural activity (**Figures 1A and 2E**).<sup>73,74</sup>

MRF neurons had robust responses to task stimuli and the impending choice of the subject, as has been observed previously.<sup>9–12</sup> Peri-event time histograms (PETH) were constructed for each neuron recorded in the task (**Figures 2F–2I and S3A–S3L**). PETHs were split by the four trial types (left stimulus, center stimulus left block, center stimulus right block, and right stimulus; correct trials only) and were aligned to both stimulus onset and movement onset (**Figures 1B and 2F**; “trial type colors”). Inspection of trials sorted by reaction time indicated that some neurons had activity appearing to align better with stimulus onset and others with movement onset. For example, some MRF neurons had stimulus-aligned activity that was selective for one or two of the task stimuli (example neurons in **Figures 2F, 2H, 2I, and S3A–S3C**). Other MRF neurons contained premotor activity selective for the impending action (i.e., choice activity). This activity was aligned to movement onset and differed between CW and CCW movements, independent of

stimulus identity (e.g., **Figures 2G and S3B**). We noted that many neurons in the MRF and other recorded regions had complex responses, which appeared to be selective for multiple task variables (**Figures S3D–S3L**).

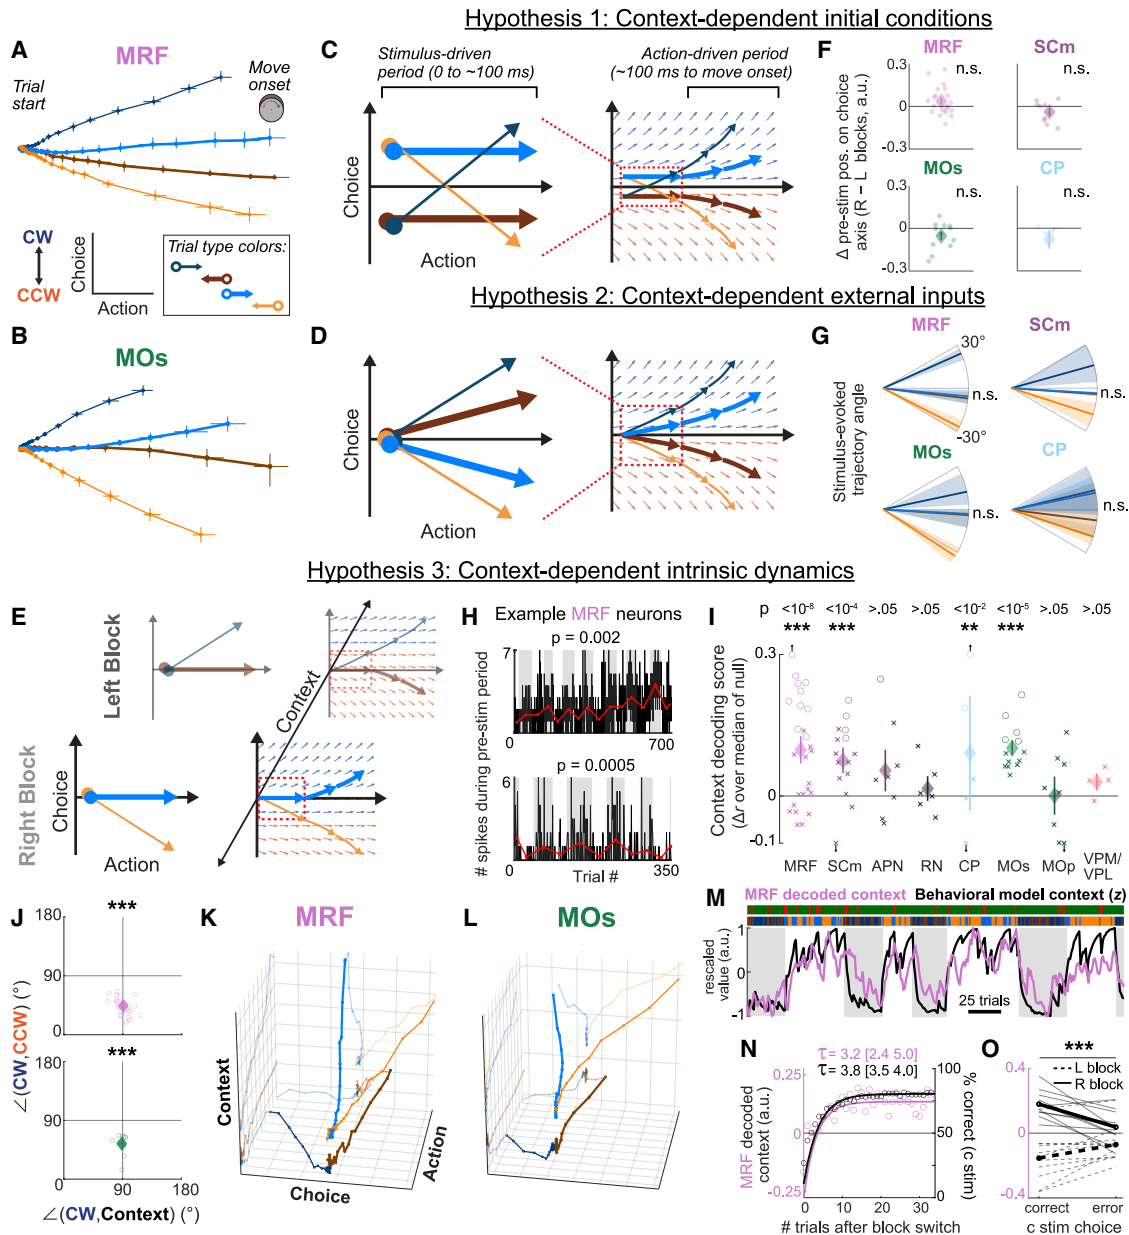
MRF neurons appeared to encode context both in pre-stimulus firing rates and in modulation of trial responses. Some MRF neurons differed in their activity level before stimulus onset, with consistently higher firing rates for trials of one block relative to the other (e.g., **Figures 2I and S3A**). We also observed that some MRF neurons' PETHs did not seem to arise from pure sensory-driven responses to the stimulus position (left, center, or right), premotor-driven responses to the direction of the action (i.e., choice, CW, or CCW wheel turn), or a combination of the two. Rather, it appeared that these neurons had stimulus and/or premotor activity that was modulated by context. For example, the neuron in **Figure 2H** displayed selective responses to center stimulus trials in a manner that varied between blocks (right block center trial response greater than left block). This difference could be due to context-dependent changes in visual-stimulus input or context-dependent changes in pre-movement action-selection dynamics (i.e., strong CW preference in right block, weak or no CW preference in left block). In the following section, we leverage a kernel-encoding model to disentangle the contributions of each task variable to these complex neuronal responses.

Having observed single MRF neurons that qualitatively appeared to code for distinct forms of context (before stimulus onset and in modulation of trial responses), we next sought to quantitatively assess these forms of context coding across brain regions.

### Pre-stimulus context representations in the MRF, MOs, SCm, and CP predict subsequent sensorimotor mapping

Choices on context-dependent center stimulus trials were decodable from a subset of regions (MRF, MOs, SCm, CP, and MOp), whereas stimulus identity was decoded from all recorded regions (**Figures S4A–S4C**). To quantify the time course of stimulus and choice information in each region, we decoded these variables from binned population activity using logistic regression. We assessed performance by comparing cross-validated (CV) classification accuracy to null distributions generated by shuffling task events (**Figures S4A–S4C; STAR Methods**). Stimulus identity was robustly decoded from all regions (one or more time bins with  $p < 0.001$ , Fisher's combined probability test across sessions, with Benjamini-Hochberg false discovery rate [FDR] correction for nine time bins); this decoding could arise from true stimulus representations and/or motor representations of the mouse's choice (which would distinguish left and right outside stimulus trials). Choices on center stimulus trials were decoded only from a subset of regions: MRF, MOs, SCm, CP, and MOp ( $p = 1.3 \times 10^{-12}$ ,  $3.4 \times 10^{-5}$ ,  $1.8 \times 10^{-10}$ ,  $6.5 \times 10^{-3}$ , 0.033;  $p > 0.05$  all other regions; **Figure S4C**). Given that this subset of regions evolved to distinct premotor states following identical center stimuli, we next investigated the potential mechanisms enabling this flexible sensorimotor transformation to occur.

A powerful approach for understanding the computations performed by neurons with diverse, mixed-selective



**Figure 3. MRF, MOs, SCm, and CP contain pre-stimulus context representations associated with remapping behavior and biasing of the neural dynamics underlying action selection**

(A and B) CV MRF (A) and MOs (B) trajectories in the 2D choice subspace (median across correct trials for each of the four trial types, standardized and combined across sessions; error bars are 95% CI). Trials start on the left and progress toward movement onset on the right. Trajectories span from  $-250$  to  $0$  ms aligned to movement onset in a sliding window of  $75$  ms bins with  $10$  ms step sizes.

(C–E) Schematics of three hypothesized mechanisms for context-dependent stimulus-action remapping. Each mechanism selectively biases one influence on choice subspace trajectories while holding others constant (initial conditions, external inputs, intrinsic dynamics) to drive center stimulus trials toward context-appropriate actions. (C) Context-dependent initial conditions: pre-stimulus position along the choice axis shifts toward the context-appropriate action (CCW in left blocks, CW in right blocks). The left plot is a zoomed-in view of the stimulus-driven period displayed within the red box on the right. During this period, stimulus input dynamics determine trajectories, and after this period, action-related dynamics determine trajectories (denoted by underlying flow field arrows). (D) Context-dependent external stimulus input: sensory-driven trajectories angle differently for center stimuli across contexts. (E) Context-dependent intrinsic dynamics: distinct action-selection dynamics for each context are set before stimulus onset and are determined by the position along a third context dimension that is orthogonal to the choice subspace.

(F) Difference in the average pre-stimulus positions between right blocks and left blocks in the choice subspace. Each circle represents a session, diamond represents the mean, and error bar is SEM. No significant differences between blocks were found in any region, using the signed-rank test.

(legend continued on next page)

responses in decision-making tasks<sup>75</sup> (Figures 2F–2I and S3) is to analyze their dynamics as a population projected onto a set of task-related dimensions.<sup>16,20,76</sup> Here, to understand how identical stimuli were transformed into distinct actions, we first examined how population dynamics evolved in a two-dimensional “choice subspace” as a function of the trial’s stimulus and context. This subspace was defined by identifying dimensions that separated pre-stimulus from pre-movement activity for each of the two possible actions (CW or CCW wheel turns), then rotating these dimensions such that one captures shared movement-related variance (“action” axis) and the other captures direction-specific variance (“choice” axis). Population activity on correct held-out trials was projected onto this subspace to create trajectories that were then standardized and averaged across sessions within each region (STAR Methods). We specifically focused on the regions from which choices on center stimulus trials could be decoded (MRF, MOs, SCm, CP, and MOP; Figure S4C), whose trajectories displayed a similar reliable structure separating task stimuli and choices (Figures 3A, 3B, and S5A–S5C) compared with the less strongly choice-driven activity in APN, RN, and VPM/VPL (Figures S5D–S5F). We observed that the initial separations of these trajectories from baseline appeared to be driven by sensory responses to the visual stimuli, as their divergent paths were stimulus-dependent and occurred with latencies matching those of stimulus decoding in each region (Figures S4A and S5G). Center stimulus trial trajectories then diverged toward contextually appropriate actions closer to movement onset (Figures 3A, 3B, and S5A–S5C).

We reasoned that the context-dependent sensorimotor transformation on center stimulus trials could be enabled by three potential mechanisms, reflecting three general ways a neural population’s state evolution can be controlled: initial conditions, external inputs, and intrinsic dynamics.<sup>77</sup> In the first hypothesized mechanism, context biases the initial conditions by shifting

the pre-stimulus population state toward one action along the choice axis, a bias that would be overridden by the strong sensory-driving input on outside stimulus trials (Figure 3C). In the second hypothesis, context modulates sensory input evoked by identical center stimuli to drive the initial activity toward the appropriate action (Figure 3D). The third hypothesis involves contextual shifts in the intrinsic dynamics, selecting for distinct actions within the choice subspace, such that center stimulus trials with the same initial conditions and external inputs evolve to the appropriate action. In this mechanism, action-selection dynamics are established before stimulus onset and are determined by the population’s position along a context dimension orthogonal to the choice subspace, similar to what has been observed in the primate PFC during context-dependent decision-making tasks<sup>16,20</sup> (Figure 3E). Thus, regions controlling context-dependent sensorimotor transformation would be expected to represent context before stimulus onset along this orthogonal context dimension.

All three hypothesized mechanisms could, in principle, solve the task. To test the computational feasibility of the dynamical mechanisms, we instantiated each in drift-diffusion models (DDMs) (Figures S5H and S5I; STAR Methods) and verified that all three achieved comparable performance to each other and to the behavior of the mice on the task (Figure S5J). We next sought to disambiguate which of these mechanisms was being implemented at the neural population level.

Pre-stimulus neural state positions and initial sensory-driven trajectories on center stimulus trials did not significantly differ by context (Figures 3F and 3G). To test the first hypothesized mechanism of contextual biases in initial conditions within the choice subspace (Figure 3C), we calculated the average difference in pre-stimulus positions between right and left blocks for each session and found no significant differences across sessions in any region ( $p > 0.05$ , signed-rank test; Figure 3F).

(G) Stimulus input angles across the four trial types (median  $\pm$  95% CI across all trials). Angles formed between the x axis and the line between averaged positions of 0–20 ms bins and 100–150 ms bins (aligned to stimulus onset) in choice subspace. No significant differences between center stimulus angles across blocks were found in any region, using the rank-sum test.

(H) Example MRF neurons with significant correlations ( $p < 0.05$ ) between pre-stimulus firing (300–50 ms before stimulus onset) and context. Black trace is trial-to-trial pre-stimulus spike count, block identity is indicated by shaded vs. unshaded patches, and red trace is the mean pre-stimulus spike count in each block.  $p$  values are derived from empirical null using pseudosessions (STAR Methods).

(I) Decoding of context from pre-stimulus firing (300–50 ms before stimulus onset). Diamonds are mean  $\pm$  SEM context decoding score (difference between true  $r$  relative to median of null) across sessions. Circles are individual sessions with significant decoding ( $p < 0.05$  against pseudosession null); x’s are sessions with non-significant decoding ( $p > 0.05$ ).  $p$  values per region are derived from FDR-corrected Fisher’s combined probability test.

(J) Angles between CW and CCW dimensions vs. CW and context dimensions for the MRF (top) and MOs (bottom). Individual sessions (circles) and mean  $\pm$  SEM across sessions (diamond with bars) are plotted. Whereas CW and CCW dimensions have significant overlap, CW, and context dimensions are nearly orthogonal ( $p < 0.001$ , signed-rank test).

(K and L) CV 3D trajectories of the MRF (K) and MOs (L) in a subspace spanned by the action/choice dimensions and a third context dimension. As opposed to the 2D choice subspace, center stimulus trajectories are separable across blocks starting at the pre-stimulus period.

(M) CV decoded context from MRF population activity (pink) and latent context variable  $z$  from behavioral model (black) on an example session, as in Figure 1G. Each trace was rescaled to the range  $[-1, 1]$  (min-max scaling for behavioral trace, robust sliding window scaling for neural trace; see STAR Methods).

(N) Projection of MRF activity onto the pre-stimulus context dimension (left y axis, pink) and percentage correct on center stimulus trials (right y axis, black) as a function of number of trials after a block switch. Context dimension activity was combined across blocks by flipping the sign of left block trials. Mean values across trials and sessions ( $n = 11$  sessions) for each point are plotted with exponential decay fit overlaid (STAR Methods). Time constants from exponential fits to context dimension activity and behavior are closely matching, with overlapping 95% CIs (bootstrapped).

(O) Average MRF context dimension activity on center stimulus trials as a function of the mouse’s ensuing choice (correct or error), split by block. Context dimension activity significantly represented the opposite block on error trials ( $p = 9.8 \times 10^{-4}$ , signed-rank test on overall percentage change between correct and error trials on each session). Thinner lines represent individual sessions (two per session, one for each block), bolded lines are mean across sessions. For (N) and (O), only sessions with  $p < 0.1$  for decoding of context from pre-stimulus activity were included.

See also Figures S4–S6.

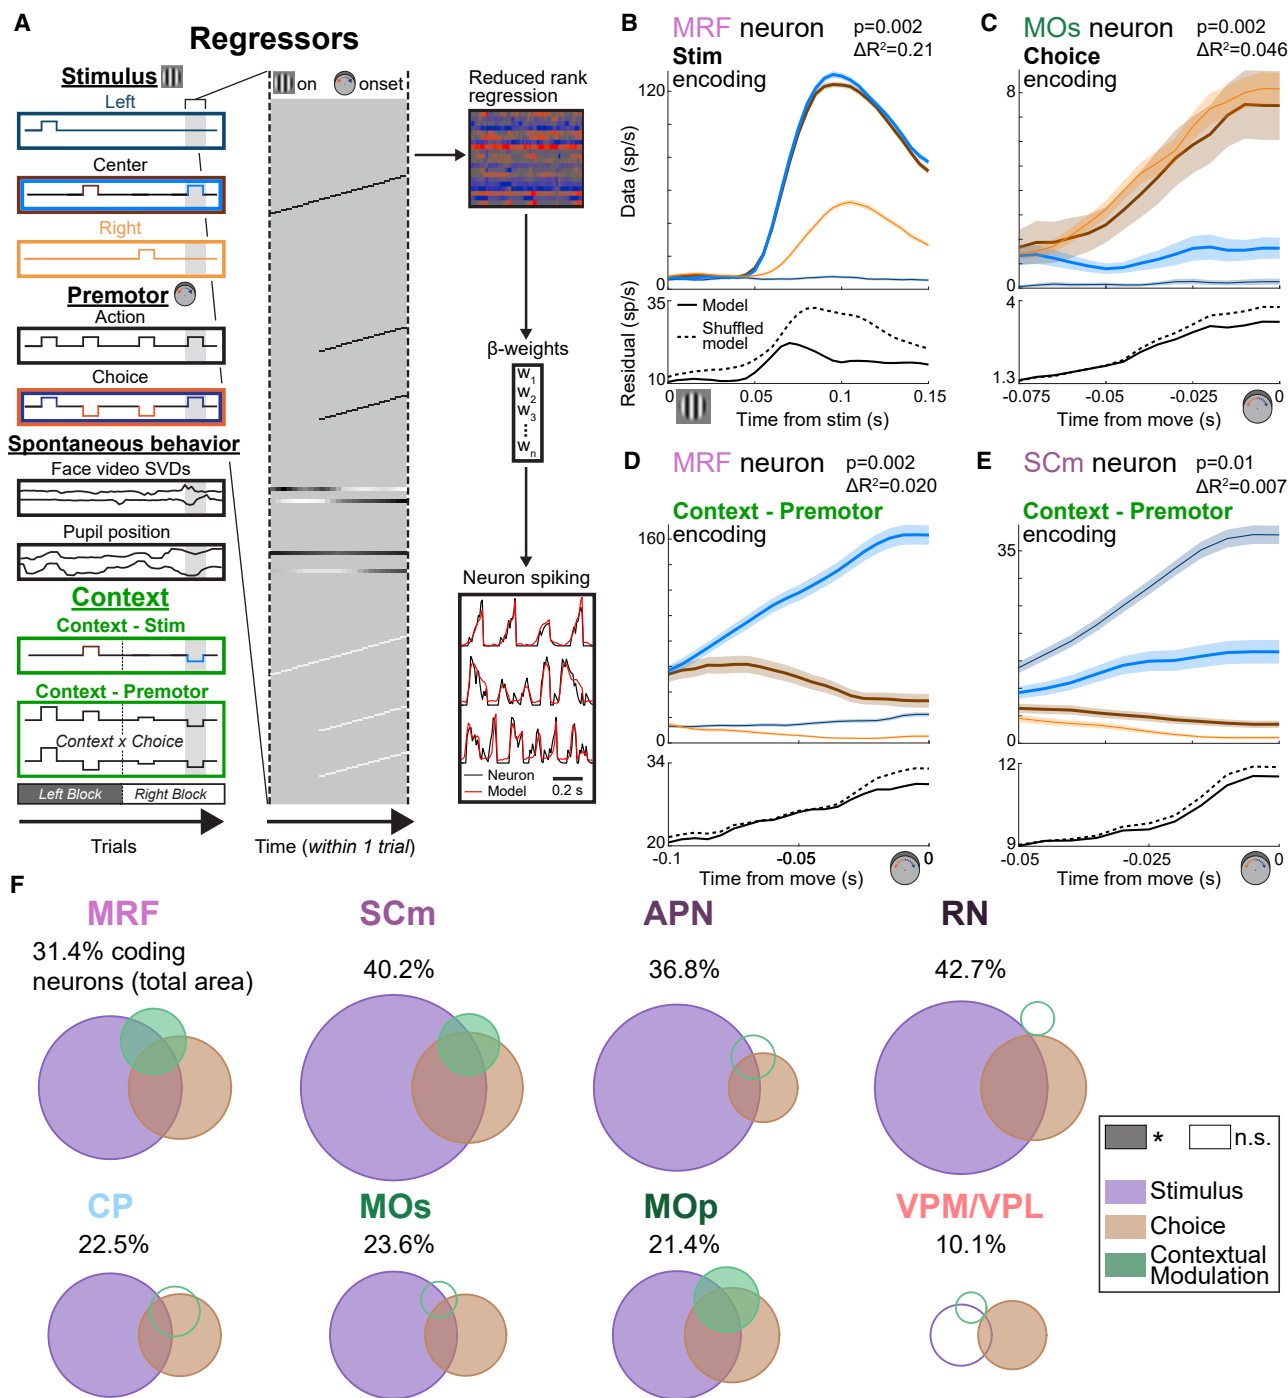
We next tested the second mechanism, contextual shifts in driving sensory inputs on center stimulus trials (Figure 3D), by calculating the angle of the initial trajectory from baseline for each trial type (formed by the vector between the average of positions of the pre-stimulus response bins (0–20 ms relative to stimulus onset) and post-response bins (100–150 ms). The left and right stimulus angles were consistently directed toward CW and CCW actions, respectively (MRF =  $24^\circ \pm 4.1^\circ$  left stimulus,  $-27^\circ \pm 4.9^\circ$  right stimulus; median  $\pm$  95% confidence interval [CI] across trials and sessions). The center stimulus input drove an intermediate trajectory angle, in between that of the left and right stimuli (MRF =  $-6.9^\circ \pm 4.1^\circ$  left block,  $-5.5^\circ \pm 4.9^\circ$  right block), with no significant differences in center stimuli between blocks in any region (Figure 3G,  $p > 0.05$ , rank-sum test). Given these results, we ruled out the first and second hypothesized mechanisms for stimulus-action remapping (Figures 3C and 3D).

MRF, MOs, SCm, and CP maintained robust representations of context in their pre-stimulus firing rates (Figures 3H, 3I, and S5K). Given that sensory-driven input and initial conditions showed no significant differences between blocks within the choice subspace, the late split of center trial trajectories into context-appropriate actions appeared to arise from context-dependent shifts in the internal dynamics guiding action selection. We hypothesized that these dynamics were set at baseline, before stimulus presentation, in a context-dependent manner (Figure 3E). We first examined whether pre-stimulus activity showed significant context-dependent differences in each region, outside of the 2D choice subspace. To control for spurious correlations<sup>78</sup> that may arise from slowly varying components in neural activity and the context variable, we used a conservative approach based on “pseudosessions,”<sup>79</sup> which served as empirical null distributions for each session’s context variable by randomizing its precise sequence while preserving overall temporal statistics (Figure S1H; STAR Methods). Trial-by-trial pre-stimulus firing rates in single neurons of the MRF, MOs, SCm, and CP were significantly correlated with context (Figures 3H and S5K; RN, APN, MOp, and VPM/VPL were below chance). Decoding context at the population level, there was a significant effect of region on the decoding score after accounting for random effects of subject and session ( $p = 0.011$ , linear mixed-effects model; see STAR Methods), and in post hoc testing, context could be decoded significantly from the same four regions, MRF, MOs, SCm, and CP (Figure 3I; Fisher’s combined probability test across sessions, with FDR correction for the 8 regions, MRF:  $p < 10^{-8}$ , MOs:  $p < 10^{-5}$ , SCm:  $p < 10^{-4}$ , CP:  $p < 10^{-2}$ ), and subtly from APN ( $p = 0.048$ ). The grouping of the MRF, MOs, SCm, and CP was further supported by the mixed-effects model, wherein their average decoding score was significantly higher than that of the other 4 regions grouped together (APN, RN, MOp, and VPM/VPL), after accounting for random effects of subject and session ( $p = 0.016$ , F test). The decodability of context in each session was independent of both the recorded neuronal population size (Figure S6A) and the degree of accuracy in behavioral performance (Figure S6B).

Pre-stimulus context was represented in a dimension orthogonal to the choice subspace and was associated with subsequent shifts in the dynamics underlying flexible sensorimotor

transformation (Figures 3J–3L and S5L–S5N). We tested the hypothesis that, in order to provide the appropriate contextual signal while minimizing interference with sensorimotor processing, pre-stimulus context signals were present in a dimension orthogonal to the choice subspace.<sup>16,20</sup> For each session, we calculated the angle between the pre-stimulus context and CW dimensions and compared it with the angle between the CCW and CW dimensions. The context dimension was orthogonal to the CW dimension across regions (Figures 3J and S5L;  $\angle(\text{context}, \text{CW}) = 91^\circ \pm 3^\circ, 89^\circ \pm 2^\circ, 94^\circ \pm 4^\circ, 90^\circ \pm 4^\circ$ , mean  $\pm$  SEM across sessions, in the MRF, MOs, SCm, and CP), whereas the CCW and CW dimensions were significantly more collinear ( $\angle(\text{CCW}, \text{CW}) = 44^\circ \pm 3^\circ, 54^\circ \pm 4^\circ, 42^\circ \pm 4^\circ, 61^\circ \pm 5^\circ$ ;  $p = 1.8 \times 10^{-5}, 2.4 \times 10^{-4}, 2.4 \times 10^{-4}, 0.063$ , signed-rank test comparing  $\angle(\text{context}, \text{CW})$  vs.  $\angle(\text{CCW}, \text{CW})$ ). We next projected population activity onto the context, action, and choice dimensions and visualized trajectories in this 3D subspace. In the MRF, MOs, SCm, and CP, trajectories from each block were separable before stimulus onset along the context dimension and exhibited context-appropriate dynamics in the action/choice dimensions (Figures 3K and 3L: MRF and MOs; Figures S5M and S5N: SCm and CP). These 3D trajectory visualizations are supported by the preceding statistical analyses: the absence of block-dependent differences in pre-stimulus position (Figure 3F) and stimulus-driven trajectory angle (Figure 3G) within the choice subspace, the significant decoding of center stimulus choices (Figure S4C) and pre-stimulus context in these regions (Figure 3I), and the orthogonality of the context dimension to the choice subspace (Figure 3J). Overall, the findings support a neural mechanism wherein context-dependent sensorimotor transformation operates through shifts in action-selection dynamics that are set by pre-stimulus context signals in the MRF, MOs, SCm, and CP.

Pre-stimulus context representations in the MRF and MOs were predictive of stimulus-action remapping behavior (Figures 3M–3O and S4D–S4M). To further establish the functional significance of pre-stimulus context representations, we analyzed how decoded context dimension activity (Figure 3M) evolved after block switches. If activity along this dimension truly reflected subjects’ internal belief about the current contextual rule, we hypothesized that its dynamics after a block switch would align with the dynamics of subjects’ choices on center stimulus trials (which require context integration). We found that, in the MRF, SCm, and CP, these dynamics indeed closely matched (Figures 3N, S4E, and S4G), with overlapping time constants of the average pre-stimulus context dimension activity (mean [bootstrapped 95% CI]  $\tau = 3.2$  [2.4, 5.0], 3.4 [1.9, 13.9], 2.8 [2.3, 4.3] trials) and center stimulus choices ( $\tau = 3.8$  [3.5, 4.0] trials). The context representation in MOs evolved more slowly, with a time constant (Figure S4F,  $\tau = 6.2$  [4.4, 9.1] trials) longer than that of the behavior. We next tested whether context dimension activity was predictive of context-dependent choices on individual trials. In the MRF and MOs, activity was relatively more like the opposite block when mice made errors on center stimulus trials (Figures 3O and S4H–S4K; median [interquartile range] percent shift toward the opposite block on error trials = 32% [19%–41%], 28% [15%–32%], 10% [–22%–18%], 21% [–23%–64%] for the MRF, MOs, SCm, and CP;



**Figure 4. Disentangling contributions of task stimuli, choices, and context to single-neuron responses with a kernel regression encoding model**

(A) Model schematic. Left: task events across 4 example trials aligned to stimulus onset or movement onset, as well as ongoing spontaneous behavior traces, throughout each trial. Middle: zoomed-in view of the regressors for one trial (marked by gray shading and red bracket), which are convolved with a temporal kernel aligned to either stimulus onset or movement onset. Spontaneous behavior regressors are not lagged. White indicates a regressor value of  $-1$ , black is  $+1$ , and gray is in between. Right: predictor matrix is reduced to 20 predictors using reduced-rank regression, where target variable is binned spiking from all neurons in a session. Ridge regression is then used to fit the binned spiking of one neuron.

(legend continued on next page)

$p = 9.8 \times 10^{-4}$ ,  $3.9 \times 10^{-3}$ , 1, 1, signed-ranked test;  $n = 11, 9, 5, 2$  sessions). This effect remained when specifically evaluating trials later into a block, after mice had reached saturating performance on center stimulus trials (Figures S4L and S4M; percent shift 15 or more trials after a block switch = 16% (0%–29%), 6% (0%–14%) for the MRF and MOs;  $p = 0.032, 0.039$ ), demonstrating that MRF and MOs context dimension activity predicted lapses in flexible behavioral performance beyond the rule-updating period after block transitions.

Uninstructed movements and pupil dynamics did not account for the neural context signals. Although mice were instructed to hold the wheel still during the pre-stimulus period used for this decoding (pre-stimulus wheel quiescence condition; Figure 1E; STAR Methods), we considered that, in principle, subthreshold movements and/or uninstructed movements elsewhere on the mouse's body could be correlated with context. Given the high degree of variance in neural activity that can be explained by body movements,<sup>73,74</sup> the interpretation of significant neural correlations with context may be confounded by possible context-aligned uninstructed movements. To sensitively detect whether movements were significantly aligned with context during the pre-stimulus period, we used logistic regression to decode trial-by-trial context from face video motion-energy singular value decompositions (SVDs)<sup>74</sup> (tracking face, body, and wheel movements, Figures 1A and 2E) and pupil position traces aligned to the 250 ms before stimulus onset (Figure S6C). Context could be decoded from uninstructed movements or pupil in only a minority of sessions (8/35), nearly half as many as those that had significant decoding from neural activity (15/35). Moreover, there was no relationship between the decoding scores from uninstructed movements/pupil positions and decoding scores from neural activity per session (Figure S6D). Just 2 sessions showed significant decoding from both sources, which approximates the expected 3.4 sessions of overlap if these variables were independent (Figure S6E).

The pre-stimulus context signal in the MRF was not a passive correlate of the task's stimulus or reward structure but rather required active engagement in the context-dependent behavior (Figures S6F and S6G). To rule out the possibility that bottom-up sensory inputs or reward history could give rise to pre-stimulus context signals independent of task engagement, we performed task-replay control experiments in which naive mice ( $n = 3$  mice, 6 sessions) were presented with an exact replay of task events from one session that had been performed by a trained mouse (Figure S6F). Context was not significantly decoded from pre-stimulus MRF activity in any naive replay session against pseudosession controls (Figure S6G;  $p > 0.05$  for all 6 sessions,  $p = 0.19$  Fisher's combined probability test),

in contrast to the matched trained session ( $p = 0.0005$ ,  $\Delta r = 0.42$  vs.  $0.064 \pm 0.024$  in naive sessions). Further supporting this interpretation, context representations were stable over time during the pre-stimulus window (Figure S5O) and were orthogonal to stimulus and premotor representations as measured during the post-stimulus epoch (Figures S5P and S5Q; see STAR Methods).

Having established a shared population-level mechanism for flexible sensorimotor transformation, we next sought to dissect single-neuron contributions to this mechanism. To do so, we analyzed the response properties of individual neurons to (1) identify whether sensory-evoked and/or premotor activity was modulated by context and (2) determine how distinct functional cell groups (i.e., those encoding stimulus, choice, context during the trial, and/or context before stimulus onset) are distributed within and across the recorded brain regions.

To disentangle and quantify the contributions of stimulus, premotor, and context variables to trial responses, we used a kernel regression approach to fit trial-by-trial single-neuron activity as a function of task-related events and behavioral measures<sup>9,80</sup> (Figure 4A). Context can be represented in trial responses of neurons by either modulating stimulus-aligned activity to the center stimulus (Figure 4A, "context-stim" predictor) or by modulating pre-movement (i.e., wheel-turn)-aligned activity (Figure 4A, "context-premotor" predictor), independent of the combinatory effects of the stimulus and choice on firing rates alone. We used a CV-nested approach, wherein the model was refit with either the stimulus, choice, or context predictors shuffled to identify the unique contributions of each variable against an empirical null distribution (Figures 4A–4E and S7A–S7J; STAR Methods). As the visual-stimulus kernel is shuffled through permutation of the stimulus labels (left, center, or right), visual-stimulus coding is defined here by stimulus-selective responses rather than nonspecific activation by all stimuli. Additionally, the model included face video motion-energy SVDs and pupil position traces as regressors to help account for any possible motor differences between blocks during the trial period (Figure 4A, "spontaneous behavior" predictors).

Choice coding was more widely distributed across the sampled regions, whereas neurons in the midbrain contained the highest proportions of visual-stimulus-encoding neurons (Figures 4F and S7B). Similar to previous studies,<sup>9–11</sup> we found choice-encoding neurons across the midbrain and forebrain, with significant (lower 95% CI above chance level; Figure S7B) coding in every region we recorded, and the highest proportions of choice neurons in the SCm, followed by the RN, MRF, MOp, MOs, and CP (range of 8.1%–14.9% among the six regions, relative to a chance level of 2%; Figures 4F and S7B). Choice

(B) PETH (trial-averaged with 1 SEM shaded) of a stimulus-encoding MRF neuron (top) and residuals from fits of the full model and the model with shuffled stimulus labels (bottom). The shuffle iteration with median  $R^2$  was used for plotting and  $\Delta R^2$  value compared with the full model;  $p$  value is derived from empirical null using shuffles.

(C–E) Same as (B) but for MOs, MRF, and SCm neurons, which encode choice, context-premotor, and context-premotor variables, respectively. PETHs are aligned to wheel movement onset.

(F) 3-Venn diagrams (estimated using Chow-Rodgers method<sup>81</sup>) of percentages of neurons across the 8 target regions encoding stimulus, choice, and/or trial context (i.e., contextual modulation: context-stim and context-premotor combined), defined as  $p < 0.02$  against the respective shuffled null. Filled circles indicate above-chance coding; empty circles indicate below chance coding (Figures S7B and S7C). Value below the region labels is the percentage of neurons in a region represented by the total area encompassed by the three circles.

See also Figures S7 and S8.

coding was lower (but still above chance) in the APN (5.8%) and VPM/VPL (5.6%). The visual stimulus was more prevalently encoded than choice (mean across all eight regions = 22.8% stimulus vs. 10% choice) and was also distributed, with significant coding in all midbrain and forebrain regions except for the VPM/VPL. Midbrain regions had a significantly higher proportion of visual-stimulus-coding neurons than the forebrain (mean 31.3% neurons in the midbrain, 19.1% in the forebrain;  $p = 2.4 \times 10^{-11}$ , chi-squared test), which may have contributed to the earlier decoding of visual stimuli in those regions at the population level (Figure S4A). In the midbrain and the forebrain, neurons exhibited above-chance levels of joint selectivity for stimulus and choice (Figures 4F, S8A, and S8B; 6.4% of neurons jointly encoding in midbrain, 95% CI = 5.1%–8.1%, 3.8% expected, based on independent probabilities; 4.2% (3.2%–5.3%), and 1.8% in the forebrain), consistent with potential roles in sensorimotor transformation.

A significant proportion of neurons in the MRF, SCm, and MOp had contextually modulated responses, which primarily influenced premotor rather than stimulus-related activity (Figures 4F and S7C). Neurons with contextually modulated responses after stimulus onset (also referred to as “trial context” hereafter, distinguishing it from the pre-stimulus context representations analyzed above) were rarer than stimulus- or choice-encoding neurons and were encoded in similar proportions across the MRF, SCm, and MOp (4.8%, 6.2%, 5.0% of total neurons; 14.0%, 13.5%, 18.9% of coding neurons). Subdividing the SCm into its intermediate (SCi) and deep (SCd) layers revealed similar coding and mixed-selectivity profiles across both layers (Figures S7H–S7J). Although MOs and CP neurons robustly encoded context before stimulus onset (Figure S5K), they did not exhibit significant contextual modulation of responses during the trial above chance levels in this analysis (Figures 4F and S7C). Notably, context almost solely modulated premotor, rather than stimulus-aligned, activity in single-neuron responses (Figure S7C). This provides further evidence for the dynamical mechanism we observed at the population level, with context influencing the evolution of premotor activity selecting for distinct actions (Figures 3E–3L) rather than altering early sensory representations of the visual stimulus (Figures 3D–3G).

Neurons coding for context before stimulus onset were largely distinct from those with contextually modulated responses during the trial (Figures S8A–S8E). Across all regions with significant pre-stimulus context coding (MRF, MOs, SCm, and CP; Figures 3I and S5K), the probabilities that neurons encoding pre-stimulus context also coded for the stimulus, choice, or trial context variables (and vice versa) did not exceed chance levels (Figures S8A–S8D). In contrast, trial-context-coding neurons were significantly more likely to also code for choice (and vice versa) across all regions encoding both variables (MRF, SCm, and MOp) (Figures S8B and S8C). This is consistent with our observation that the contextual modulation of single-neuron trial responses mostly affected premotor activity, as choice is a premotor-aligned variable (Figures 4A and S7C). Thus, rather than neurons with contextual differences in pre-stimulus firing rates developing modulated responses after stimulus onset, these responses instead appeared to arise from influences on a separate population of premotor neurons. This result corroborates

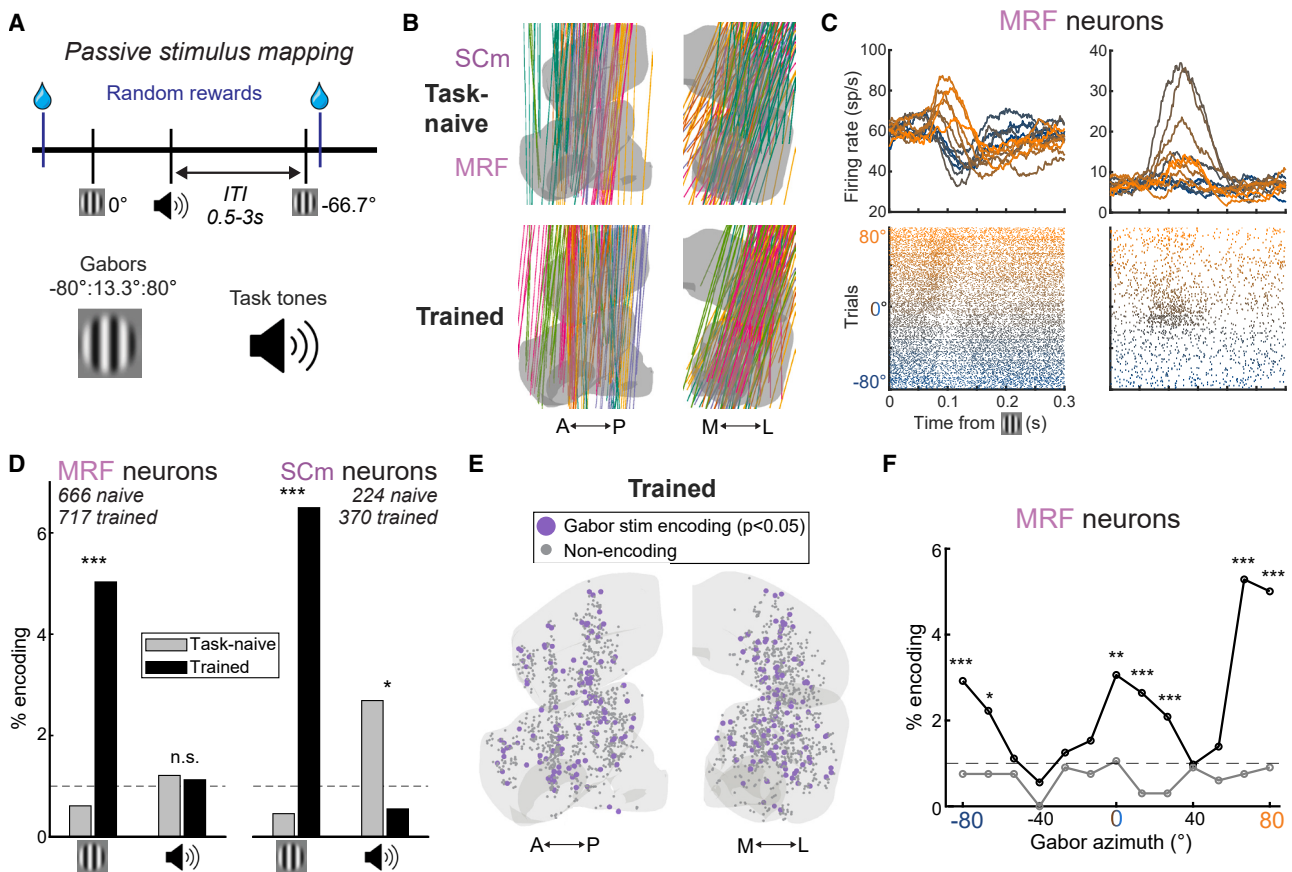
our finding at the population level that the pre-stimulus context dimension is orthogonal to the premotor choice dimensions (Figures 3J and S5L).

Having observed rapidly evolving, context-dependent dynamics in the MRF during flexible sensorimotor transformation, we next asked whether MRF activity displayed long-term contextual adaptations that would further support a role in this process. We thus investigated how MRF representations of stimuli evolved with task learning.

### The MRF develops responses to behaviorally relevant visual stimuli through task learning

The MRF and SCm exhibited enhanced responses to visual, but not auditory, task stimuli in a passive context following training. To assess the effects of task training on MRF stimulus responsiveness, we performed a separate set of recordings in task-naive mice, wherein a passive stimulus protocol was presented (Figure 5A). The identical protocol was also presented at the end of each recording session in trained mice. Stimuli present in the task (left, center, and right Gabors; go cue, miss noise, and reward tone) were produced in random order and timing, as well as being decoupled from any wheel movements. Gabor stimuli at intermediate azimuths (i.e., between  $-80^\circ$ ,  $0^\circ$ , and  $80^\circ$ ) were also displayed. Randomly timed sucrose rewards were given to ensure alertness. Similar to trained-mouse recordings, recordings in task-naive mice were targeted to densely tile the MRF to assess the spatial dependency of coding (Figure 5B). We found that there were significantly more neurons in the MRF of trained mice that responded to Gabor stimuli compared with task-naive mice (Figures 5C and 5D; 36/717 neurons with the FDR-corrected Zenith of Event-Based Time-Locked Anomalies [ZETA] statistic,<sup>82</sup>  $p < 0.01$  in trained mice, 4/666 neurons in task-naive mice,  $p = 9.5 \times 10^{-7}$ , chi-squared test). MRF Gabor stimulus encoding was approximately at chance level (0.6% coding neurons at a chance level of 1%) in task-naive mice, suggesting that stimulus responsiveness in the MRF occurs as a result of task training. In contrast to its responsiveness to visual stimuli, the MRF did not increase its responsiveness to auditory stimuli with task training: the reward tone was coded near chance in both task-naive and trained mice (Figure 5D; 8/717 trained, 8/666 task-naive,  $p = 0.88$ ). Similar to neurons in the MRF, those in the SCm also showed enhanced Gabor stimulus responsiveness with training (24/370 trained, 1/224 task-naive,  $p = 3.8 \times 10^{-4}$ ), along with a small decrease in reward tone responsiveness (2/370 trained, 6/224 task-naive,  $p = 0.029$ ). The training-induced visually responsive neurons were broadly distributed throughout the sampled extent of the MRF and SCm (Figure 5E), indicating that this plasticity was not confined to a small, discrete subregion.

The enhanced visual-stimulus responsiveness in the MRF and SCm was specific to behaviorally relevant task stimuli. Gabor stimuli were presented at 13 different azimuths, ranging from  $-80^\circ$  to  $80^\circ$ , during the passive stimulus protocol. Strikingly, MRF stimulus encoding was enhanced after training only at the azimuths at or near those present in the task (Figure 5F;  $p < 0.05$  at left  $-80^\circ$ , center  $0^\circ$ , and right  $80^\circ$  stimuli, as well as  $-66.7^\circ$ ,  $13.3^\circ$ ,  $23.7^\circ$ , and  $66.7^\circ$ , chi-squared test vs. task-naive). This same pattern of stimulus responsiveness was present in the



**Figure 5. Task-specific visual-stimulus representations emerge across MRF and SCm as a result of task training**

(A) Passive stimulus mapping protocol schematic. Task stimuli (visual: Gabors at  $-80^\circ$ ,  $0^\circ$ , or  $80^\circ$ , audio: go cue, reward tone, miss noise) and control stimuli were displayed in random order with random timing, and randomly timed sucrose rewards were given.

(B) Lateral (left) and anterior (right) 3D view of Neuropixels 2.4 recording insertions tiling SCm and MRF in task-naive (top) and trained (bottom) mice.

(C) PSTHs (trial-averaged with 1 SEM shaded) of two example visually responsive MRF neurons during passive stimulus presentation. Neuron on the left has a wide receptive field ( $-80^\circ$  to  $+80^\circ$ ) with varying responses across visual-stimulus azimuths (excitation in response to positive azimuths, inhibition in response to negative azimuths). Neuron on the right has a narrower receptive field selective for central stimuli.

(D) Percentage of neurons in the MRF (left) or SCm (right) encoding Gabor stimuli (all azimuths) or reward tone in task-naive and trained mice.  $p$  values are derived from the ZETA test.<sup>82</sup>

(E) Lateral (left) and anterior (right) 3D views of locations of Gabor stimulus-encoding and non-encoding neurons in the MRF and SCm. A 50- $\mu$ m jitter was added to the anterior-posterior (AP) and medial-lateral (ML) coordinates of neurons for visualization.

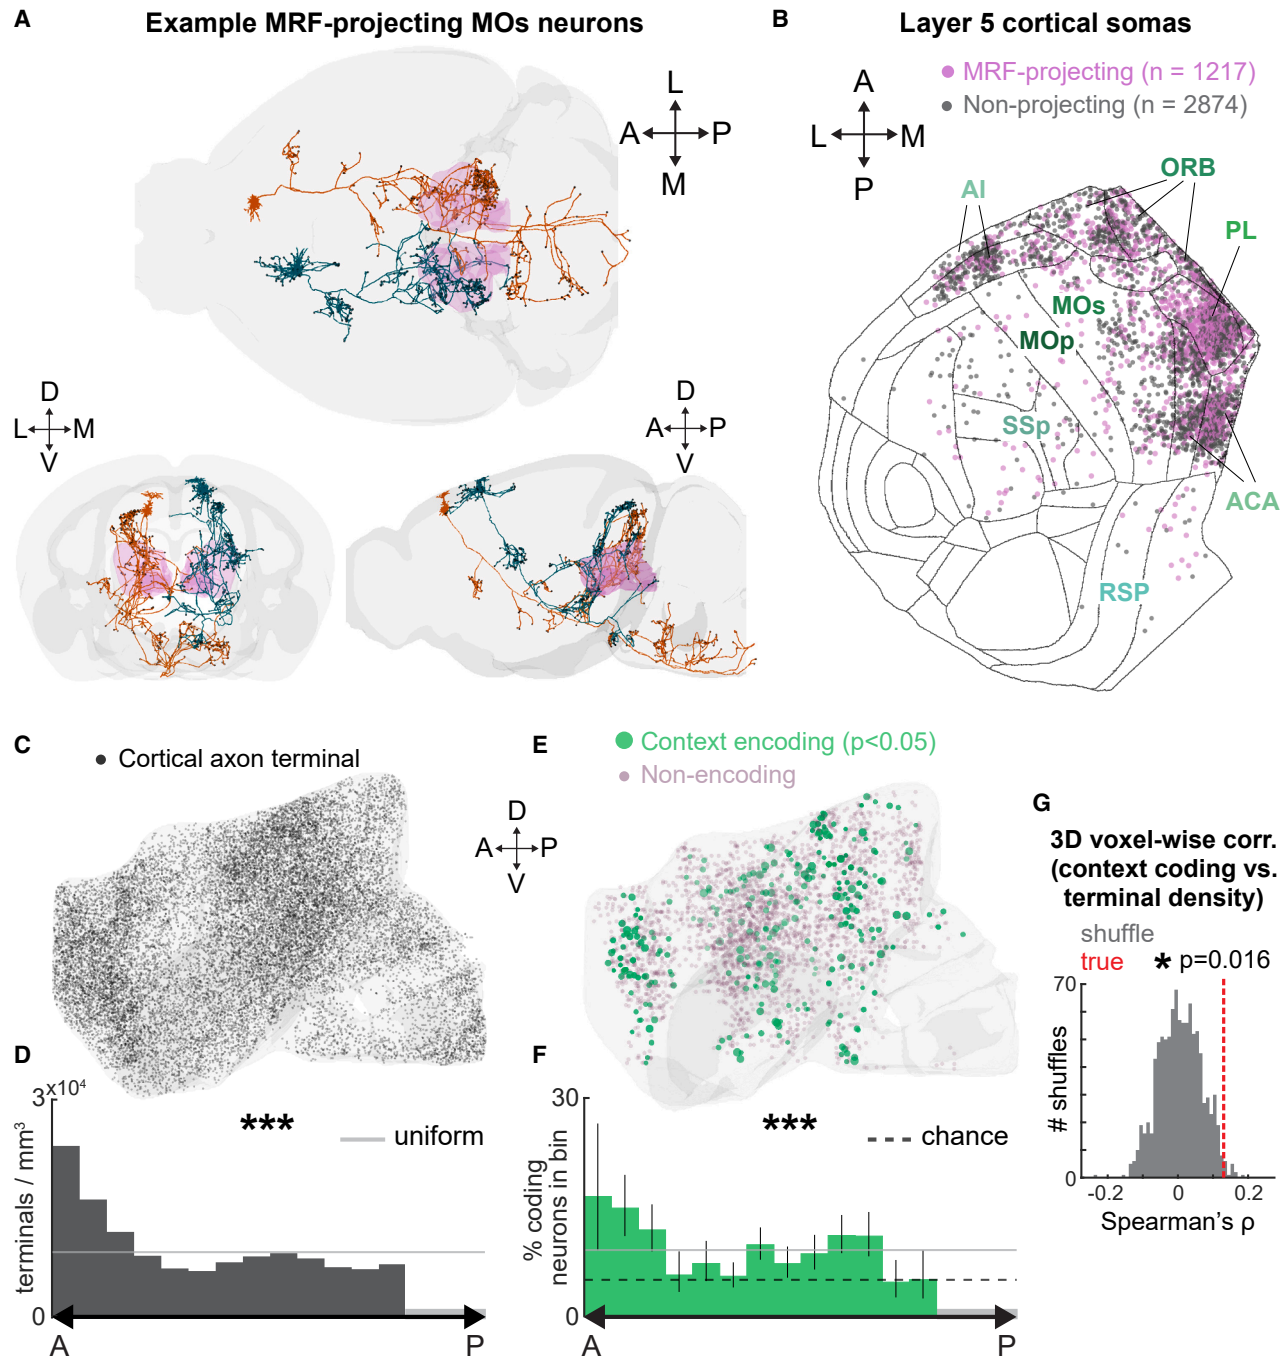
(F) Percentage of MRF neurons encoding Gabor stimuli split by azimuth in task-naive or trained mice. Above chance coding was only observed in trained mice, at or near task stimulus locations ( $-80^\circ$ ,  $0^\circ$ ,  $80^\circ$ ).  $p$  values are derived from chi-squared test between trained and task-naive mice at matched azimuths. See also [Figure S9](#).

SCm but not in forebrain regions, where the passively presented Gabor stimuli were coded near chance (Figures S9A and S9B; proportion coding [95% CI] = 0.95% [0.63%–1.4%] across forebrain regions, chance level of 1%). This finding corroborates the greater encoding of visual stimuli during the task in the midbrain compared with the forebrain (Figures 4F and S7B).

Behaviorally relevant visual-stimulus responsiveness in the MRF was not caused by stimulus-evoked movements. To test whether the passively presented stimuli evoked wheel turns or other overt movements in trained mice, we compared face video motion-energy and wheel rotary encoder traces before and after stimulus onset. We found that the Gabor stimuli did not significantly evoke movement responses except for the

right stimulus, which evoked small average changes in face and wheel movements (Figure S9C; 3.5% and 7.2% increases from pre- to post-stimulus period,  $p = 4.0 \times 10^{-5}$ ,  $9.5 \times 10^{-5}$ , signed-rank test). The left and center stimuli did not evoke movement responses (left:  $-2.5\%$ ,  $-5.8\%$  decrease,  $p = 0.36$ ,  $0.78$ ; center:  $-0.6\%$ ,  $-2.5\%$  decrease,  $p = 0.29$ ,  $0.25$ ). Thus, visual responses to behaviorally relevant stimuli during the passive task were likely not attributable to stimulus-evoked movements, particularly the enhanced responsiveness to center and left stimuli.

Thus far, we have shown that the MRF adapts its activity to represent a diversity of task variables during flexible decision-making. In particular, the key variable associated with flexibility



**Figure 6. Spatial distributions of cortical axon terminals and context-coding neurons in the MRF are correlated and enhanced in anterior MRF**

(A) Two example layer 5 MOs neuron morphologies (blue and orange structures) with several axon terminals (black spheres) in the MRF (pink). Data compiled from three open single-neuron morphology datasets.<sup>40–42</sup>

(B) Spatial locations of somas of layer 5 cortical neurons that project to the MRF ( $\geq 1$  axon terminal in the MRF). Each circle is a cell body; pink project to the MRF, gray do not project to the MRF. Soma locations are displayed on a cortical flatmap representation<sup>11</sup> (see Figure S10A for labeling of all regions), with most neurons residing in frontal cortical regions (Figure S10B).

(C) Lateral 3D view of the MRF with locations of cortical axon terminals plotted as spheres.

(D) Cortical axon terminal densities (terminals/mm<sup>3</sup>) across coronal slices of the MRF. Slice binning is 250- $\mu$ m wide with 150- $\mu$ m spacing. Significance stars indicate non-uniformity of distribution (STAR Methods, “spatial non-uniformity testing”).

(legend continued on next page)

in the task, pre-stimulus context, was also specifically found to be represented in canonical flexible decision-making regions interconnected with the MRF: MOs, SCm, and CP. We hypothesized that the similar dynamics and coding of context in these regions may reflect recurrent computations among them as a network. If so, we would expect that the parts of the MRF coding for pre-stimulus context should spatially align with the parts most strongly connected to the rest of this network.

### Context coding and cortical axon terminals spatially align in anterior MRF

Single-neuron morphological reconstructions displayed dense cortical inputs to the MRF, primarily originating from the frontal, motor, and somatosensory cortex. To analyze the degree and spatial distribution of cortical inputs to the MRF, we compiled single-neuron morphologies from three CCF-aligned open datasets.<sup>40–42</sup> We filtered the neurons in each dataset for those with their cell body in the neocortex and at least one axon terminal within the MRF (Figure 6A: example MOs neurons). The MRF received inputs from prefrontal (anterior cingulate, prelimbic, orbital), motor, retrosplenial, and somatosensory cortical regions (Figures 6B, S10A, and S10B; terminals from regions with 10 or more neurons projecting to the MRF were included for further analysis). Thirty percent of all layer 5 neurons included in the dataset projected to the MRF, with a small but significant variation in the proportion of projecting neurons between regions (Figure S10C;  $p = 1.7 \times 10^{-9}$ , chi-squared test, Cramér's  $V = 0.12$ ). In particular, the prelimbic cortex had a significantly higher proportion of layer 5 neurons projecting to the MRF, whereas the agranular insula and orbital cortices had lower proportions (36%, 23%, 24%; Bonferroni-corrected  $p = 3 \times 10^{-4}$ ,  $3 \times 10^{-3}$ ,  $8 \times 10^{-4}$ , post hoc tests with adjusted Pearson residuals).

The axon terminal projection patterns of cortical neurons into the MRF were significantly non-uniform. To characterize whether cortical neurons had a spatially biased or uniform targeting of the MRF, we binned the MRF into 250- $\mu\text{m}$  slices with 150- $\mu\text{m}$  spacing along the anterior/posterior (A/P), dorsal/ventral (D/V), and medial/lateral (M/L) planes and calculated the density of axon terminals in each slice (Figures 6C, 6D, and S10D). To assess non-uniformity, the distance of the density distributions from the expected densities based on the global mean was calculated for each plane (STAR Methods). We found that the distances of the spatial densities from a uniform distribution were significant along each plane, compared with a null distribution wherein terminal positions were randomly distributed throughout the MRF, with a high degree of non-uniformity in

the A/P plane and much weaker effect sizes in the D/V and M/L planes ( $p = 1 \times 10^{-4}$ ,  $1 \times 10^{-4}$ ,  $1 \times 10^{-4}$ ; Z score relative to null = 645, 17, 40 for the A/P, D/V, and M/L planes). The most prominent feature of the non-uniform cortical terminal distribution was the particularly high density in the anterior aspect of the MRF (Figures 6C, 6D, and S10D). Stratifying axon terminal distributions by cortical source region revealed broadly consistent spatial profiles across regions, with each exhibiting significant non-uniformity along all three spatial dimensions (Figure S11), supporting the pooling of cortical inputs across regions for terminal-organization analyses.

Pre-stimulus context coding showed a similarly non-uniform spatial distribution in the MRF (Figures 6E, 6F, and S10E). To determine the spatial distributions of neurons encoding pre-stimulus context, we performed the slice-binning procedure as before and calculated the proportion of coding neurons and the distance from a uniform distribution along each plane. We performed the same non-uniformity test against a null distribution, wherein coding vs. non-coding labels were permuted among all recorded neurons within the MRF. The distribution of pre-stimulus context-coding neurons was significantly non-uniform in the A/P and M/L planes, with a stronger effect size in the A/P than in the M/L plane ( $p = 1 \times 10^{-4}$ , 0.054,  $2 \times 10^{-3}$ ; Z score = 7.7, 1.8, 4.4 in the A/P, D/V, and M/L planes). Similar to the spatial distribution of cortical terminals, there was a particularly high proportion of pre-stimulus context neurons in the anterior aspect of the MRF (Figures 6E, 6F, and S10E).

Pre-stimulus context neurons and cortical axon terminals in the MRF were spatially aligned (Figure 6G). To test the alignment between coding and connectivity more sensitively across the volume of the MRF, we divided the MRF into 125- $\mu\text{m}$ -radius 3D voxels with 150- $\mu\text{m}$  spacing and calculated distances from uniform distributions for both the density of cortical terminals and proportion of pre-stimulus context neurons in each voxel. The voxel-wise context-coding proportions and terminal densities were then correlated using Spearman's rank correlation and compared with a null, wherein the labels for context coding vs. non-coding were randomly permuted among the recorded neuron locations. Across MRF voxels, we found a modest but statistically significant positive correlation between the proportion of pre-stimulus context-encoding neurons and the density of cortical axon terminals (Spearman's  $\rho = 0.13$ ,  $p = 0.016$  against a shuffle null; Figure 6G). This alignment was consistent with the observation that both context coding and cortical terminal density were enriched in the anterior MRF (Figures 6C–6F, S10D, and S10E).

(E) Lateral 3D view of the MRF with recording locations of pre-stimulus context-coding neurons ( $p < 0.05$ ; Figures 3H and S5K) and non-encoding neurons plotted as spheres. Sizes of context-coding neurons are scaled by  $-\log(p)$  value (derived from spike rate correlation with context vs. null), and transparency value is scaled by local density of context coding (calculated as the proportion of the nearest 10% of neurons that code for context). A 150- $\mu\text{m}$  jitter was added to AP and ML coordinates of neurons for visualization.

(F) Same as (D) but for proportion of neurons encoding pre-stimulus context ( $p < 0.05$ ). Error bars are 95% binomial proportion CIs; chance level is 5%. In (D) and (F), the three posterior-most bins are excluded from plot due to insufficient sampling of the recorded neurons (<50 neurons; these bins are still included in non-uniformity calculations, which takes sampling biases into account; see STAR Methods). See Figures S10D and S10E for distributions of terminals and context coding in D/V and M/L dimensions.

(G) True vs. null Spearman's rank correlation coefficients between voxelized (125- $\mu\text{m}$  radius and 150- $\mu\text{m}$  separation) context coding proportions and terminal densities throughout the MRF. Shuffles consisted of random permutations of context-coding neuron locations (STAR Methods).

See also Figures S10 and S11.

## DISCUSSION

Using a novel flexible decision-making task in mice, combined with behavioral modeling, large-scale electrophysiological recordings, and anatomical analyses (Figures 1 and 2), we found converging evidence that the MRF is a key node in the network of regions implementing contextual control of sensorimotor transformations. The MRF, alongside canonical decision-making regions, stably encoded an abstract task rule during the pre-stimulus period, a representation that appeared to establish a distinct dynamical state guiding subsequent action selection. Notably, this neural representation of context in the MRF was predictive of flexible choice behavior (Figures 3 and 4). The contextual adaptation of MRF activity during flexible behavior was further underscored by its development of robust responses that were specific to task-relevant stimuli after training (Figure 5). Finally, we found a spatial organization of structure and function in the MRF: frontal cortical inputs and context coding were spatially aligned and notably concentrated in the anterior MRF, distinguishing it from previously studied motor-related subregions in the posterior MRF<sup>43,54–56</sup> (Figure 6).

Our task design and recording strategy allowed us to test competing models of the circuitry involved in cognitive control during flexible behavior.<sup>2,4,14,83,84</sup> Several brain-wide surveys of perceptual decision-making tasks have revealed that, unlike the coding of movements and actions,<sup>73,74,85</sup> choice is represented in a narrower subset of regions, comprising multi-regional loops between the frontal cortex, basal ganglia, thalamus, midbrain, and hindbrain.<sup>9–11</sup> It has remained unclear whether such a pattern of coding would persist for cognitive variables that are abstracted from the motor output of the subject.<sup>14–18,27,28,58,60</sup> We targeted the MRF, which is among the most choice-selective regions in the brain<sup>9–11</sup> (Figure 4F), along with canonical decision-related areas, namely MOs,<sup>60,72,86</sup> SCm,<sup>6,27–29,87</sup> and CP,<sup>7,23–25</sup> during a task that required the use of an abstract task rule to contextually update sensorimotor associations.

Our findings argue against a hierarchical model positing a central executive in the frontal cortex<sup>3,8,22</sup> and instead support a model wherein cognitive control is more distributed and intertwined with systems for sensorimotor control.<sup>2,4,30,31</sup> A specific set of connected regions—MRF, MOs, SCm, and CP—maintained robust pre-stimulus representations of the abstract task context, which predicted subsequent population dynamics underlying flexible sensorimotor transformation (Figures 3E–3L, S5M, and S5N). Notably, this pre-stimulus representation of context was orthogonal to motor choice dimensions (Figures 3J and S5L) and appeared to set the initial conditions for the choice process, allowing the same sensory input to be mapped to different actions based on the prevailing rule. The absence of contextual modulation of early stimulus-evoked activity (Figures 3G and S7C) further supports this mechanism, suggesting that context primarily shapes the dynamics of action selection rather than early sensory processing. The orthogonality of the context and choice dimensions means that context activity can influence choice dynamics while still being separable from motor activity in downstream readout, analogous to the relationship between preparatory and move-

ment activity in motor control.<sup>88</sup> This orthogonal coding motif for flexible computation, previously identified within the primate PFC,<sup>16,20</sup> is thus shown here to operate across a distributed cortical-subcortical network, potentially supported by the recurrent connectivity between these areas (Figures 2B, 2C, and S2E).<sup>2,30,33–35</sup> The pathways through which this network exerts its contextual control remain to be determined but may include influences on processing in the primary sensory cortex. Notably, a recent study identified an orbitofrontal-to-primary auditory cortex pathway mediating context-dependent sensory modulation during a related flexible decision-making task.<sup>59</sup> A similar pathway could operate downstream of the network we identified here, as the MD thalamus, which is a source of primary input to the OFC,<sup>32</sup> is a major ascending target of both the MRF and SCm<sup>35,43,46</sup> (Figure S2) as well as a recipient of feedback-type (L6-originating) projections from MOs.<sup>89</sup>

Although the context-dependent population dynamics were remarkably similar across the MRF, MOs, SCm, and CP, the MRF was distinguished by a specific combination of coding properties. Its pre-stimulus context representation was, alongside the MOs, significantly predictive of mice's subsequent flexible choices on a trial-by-trial basis (Figures 3O and S4H–S4M). It is important to note, however, that our recordings may not have sufficiently sampled the most behaviorally relevant portions of the SCm,<sup>35,90</sup> and especially the CP,<sup>34,91</sup> potentially underestimating their degrees of context representation. Furthermore, the MRF was one of only three regions (including the SCm and MOp) to contain a distinct population of neurons whose premotor activity was directly modulated by context (Figures 4F and S7C). Thus, the MRF was the only recorded region containing both choice-predictive pre-stimulus context representations and contextually modulated premotor responses. This suggests that the MRF may uniquely play an integral role in both maintaining the contextual rule and in its downstream implementation in motor-related activity, in line with the MRF's integrative forebrain input connectivity<sup>36,38,39</sup> and robust ascending and descending outputs to both decision- and motor-related circuitry (Figures 2B, 2C, and S2E). Given that the neural populations underlying these functions were largely separate (Figure S8), an important future direction will be to dissect the roles of MRF neurons based on their inputs and projection targets.

Our operationalization of “context” as an implementation of abstract, learned rules governing sensorimotor transformations represents a specific form of cognitive control,<sup>4,14</sup> distinct from other contextual states studied in recent work utilizing similar large-scale recording approaches in head-fixed mice: for instance, motor priors,<sup>13</sup> task engagement,<sup>9</sup> and physiological need states governing switches between food versus water seeking.<sup>92</sup> Interestingly, each of these studies also reported that MRF neurons robustly correlated with behavioral context, among a subset of cortical and subcortical regions. In line with early theories of the MRF,<sup>38,50</sup> this raises the intriguing possibility that it serves as a general hub for integrating various signals related to the current behavioral state to shape sensorimotor processing appropriately. More broadly, the involvement of the MRF in diverse contextual processes underscores the perspective that models of cognitive control must take into account

interactions between cortical regions and the conserved subcortical circuits upon which they evolved.<sup>30,70,84,93</sup>

The MRF, which is not part of the classical visual pathway and does not receive direct retinal input,<sup>94</sup> has nevertheless repeatedly been found to contain robust visual responses during visual decision-making tasks.<sup>9,11,12</sup> What is the role of these visual responses in behavior? We found that MRF neurons in trained, but not task-naive, mice contained visual responses specifically tuned to the spatial locations of the task stimuli (Figures 5C, 5D, and 5F). Although previous work has shown visual responses emerging in trained mice during active behavioral contexts,<sup>9,11,12</sup> our results further demonstrate that these representations are tuned to task-relevant visual features, and they are present even in the absence of motor or reward contingencies during passive viewing. This task-specific tuning parallels the finding that neurons in the posterior MRF and adjacent pedunculopontine nucleus developed auditory responses specific to a task go-cue to trigger movement.<sup>43</sup> Together, these observations suggest a general, sensory-modality-aligned principle, whereby the MRF and other midbrain circuits become tuned to represent behaviorally important cues through learning, enabling context-appropriate sensorimotor transformations. In the task here, successful performance of the ambiguous center stimulus trials required integration of the stimulus location with the overarching contextual rule, and behaviorally, mice used the location of the instructive visual stimuli (i.e., the left or right outside stimuli) to update their internal context belief (Figures 1I and 1J). One possibility is therefore that the MRF visual responses do not only inform choices on the immediate trial but also drive the updating of contextual belief. This interpretation predicts that disrupting MRF visual responses should impair subjects' ability to update their internal context belief, which is a testable hypothesis for future work.

Finally, our work provides an initial framework for functional organization within the MRF's complex architecture. The discovery that both context-coding neurons and cortical axon terminals are spatially aligned and enriched in the anterior MRF (Figures 6C–6G, S10D, and S10E) suggests a functional gradient, distinguishing this zone from more posterior subregions historically implicated in specific motor functions.<sup>43,54–56</sup> A crucial next step will be to systematically parcellate the MRF by combining high-resolution connectivity mapping, molecular cell typing,<sup>95</sup> and functional characterization across diverse behaviors. Nevertheless, our results argue that the MRF serves as a previously unappreciated node within the circuitry for exerting contextual control of flexible sensorimotor behavior.

#### RESOURCE AVAILABILITY

##### Lead contact

Requests for further information and resources should be directed to and will be fulfilled by the lead contact, Nicholas A. Steinmetz ([nick.steinmetz@gmail.com](mailto:nick.steinmetz@gmail.com)).

##### Materials availability

This study did not generate new unique reagents.

#### Data and code availability

- Electrophysiology and behavioral datasets generated and analyzed in this study have been deposited at Figshare and are publicly available at <https://doi.org/10.6084/m9.figshare.31428575> as of the date of publication.
- Behavioral modeling code has been deposited at GitHub and is publicly available at <https://doi.org/10.5281/zenodo.19257765> as of the date of publication.
- Original code used to analyze the neural data has been deposited at GitHub and is publicly available at <https://doi.org/10.5281/zenodo.19226584> as of the date of publication.
- Any additional information required to reanalyze the data reported in this paper is available from the [lead contact](#) upon request.

#### ACKNOWLEDGMENTS

We would like to thank Ljivica Kolich for animal husbandry and behavioral training; Jovana Navarrete, Sam Golden, and Charles Zhou for light-sheet imaging and histology support; and Jeeseo Yoon, Julia Hopkins, and Rachel Yin for help with behavioral training. We thank Matthew Golub for early feedback on behavioral context models. We thank Kevin Miller for guidance in implementing the DisRNN model of this task. We also thank Zhiwen Ye and other members of the Steinmetz lab for useful discussions and feedback about this work. This work was supported by NIH grant F30EY035113 (to J.R.S.), a Washington Research Foundation postdoctoral fellowship (to D.B.), NIH training grant EY07031 (to D.B.), NSF grant CAREER #2142911 (to N.A.S.), the Pew Biomedical Scholars Program (to N.A.S.), and the Klingenstein Fellowship in Neuroscience (to N.A.S.).

#### AUTHOR CONTRIBUTIONS

J.R.S. and N.A.S. conceived and designed the study. J.N.S. performed behavioral modeling analyses and a subset of task-naive electrophysiological recordings. J.R.S. performed all other experiments and data analysis. D.B. helped develop experimental protocols and 3D visualizations. J.R.S. and N.A.S. revised the manuscript. J.R.S., J.N.S., and N.A.S. wrote the original draft manuscript. N.A.S. supervised and obtained primary funding for the project.

#### DECLARATION OF INTERESTS

The authors declare no competing interests.

#### DECLARATION OF GENERATIVE AI AND AI-ASSISTED TECHNOLOGIES IN THE WRITING PROCESS

During the preparation of this work, the authors used Claude 4 and Gemini 3 to help refine parts of the text written by the authors, with the goal of improving the language and readability of the manuscript. After any use of these tools, the authors reviewed and edited all outputs and take full responsibility for the content of the published article.

#### STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- [KEY RESOURCES TABLE](#)
- [EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS](#)
- [METHOD DETAILS](#)
  - Surgical procedures
  - Histology and imaging
  - Analysis of MRF projections with AAV tracing
  - Flexible rule-based visual decision-making task
  - Behavioral Analysis
  - Electrophysiological recordings

- Neural data analysis
- Drift-diffusion modeling of stimulus-action remapping mechanisms
- Analyses of cortical terminal and context coding spatial distributions in MRF
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
  - Significance testing with empirical null distributions
  - Pseudosessions

#### SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.neuron.2026.04.003>.

Received: October 6, 2025

Revised: February 27, 2026

Accepted: April 2, 2026

#### REFERENCES

1. Uddin, L.Q. (2021). Cognitive and behavioural flexibility: neural mechanisms and clinical considerations. *Nat. Rev. Neurosci.* 22, 167–179. <https://doi.org/10.1038/s41583-021-00428-w>.
2. Cisek, P. (2012). Making decisions through a distributed consensus. *Curr. Opin. Neurobiol.* 22, 927–936. <https://doi.org/10.1016/j.conb.2012.05.007>.
3. Gold, J.I., and Shadlen, M.N. (2007). The Neural Basis of Decision Making. *Annu. Rev. Neurosci.* 30, 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>.
4. Okazawa, G., and Kiani, R. (2023). Neural Mechanisms That Make Perceptual Decisions Flexible. *Annu. Rev. Physiol.* 85, 191–215. <https://doi.org/10.1146/annurev-physiol-031722-024731>.
5. Wolff, M., and Halassa, M.M. (2024). The mediodorsal thalamus in executive control. *Neuron* 112, 893–908. <https://doi.org/10.1016/j.neuron.2024.01.002>.
6. Basso, M.A.M.A., Bickford, M.E.M.E., and Cang, J. (2021). Unraveling circuits of visual perception and cognition through the superior colliculus. *Neuron* 109, 918–937. <https://doi.org/10.1016/j.neuron.2021.01.013>.
7. Ding, L., and Gold, J.I. (2013). The basal ganglia's contributions to perceptual decision making. *Neuron* 79, 640–649. <https://doi.org/10.1016/j.neuron.2013.07.042>.
8. Miller, E.K., and Cohen, J.D. (2001). An Integrative Theory of Prefrontal Cortex Function. *Annu. Rev. Neurosci.* 24, 167–202. <https://doi.org/10.1146/annurev.neuro.24.1.167>.
9. Steinmetz, N.A., Zatzka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576, 266–273. <https://doi.org/10.1038/s41586-019-1787-x>.
10. Chen, S., Liu, Y., Wang, Z.A., Colonell, J., Liu, L.D., Hou, H., Tien, N.-W., Wang, T., Harris, T., Druckmann, S., et al. (2024). Brain-wide neural activity underlying memory-guided movement. *Cell* 187, 676–691.e16. <https://doi.org/10.1016/j.cell.2023.12.035>.
11. International; Brain Laboratory, Angelaki, D., Benson, B., Benson, J., Birman, D., Bonacchi, N., Bougrova, K., Bruijns, S.A., Carandini, M., Catarino, J.A., et al. (2025). A brain-wide map of neural activity during complex behaviour. *Nature* 645, 177–191. <https://doi.org/10.1038/s41586-025-09235-0>.
12. Khilkevich, A., Lohse, M., Low, R., Orsolich, I., Bozic, T., Windmill, P., and Mrosovsky, T.D. (2024). Brain-wide dynamics linking sensation to action during decision-making. *Nature* 634, 890–900. <https://doi.org/10.1038/s41586-024-07908-w>.
13. Findling, C., Hubert, F., International; Brain Laboratory, Acerbi, L., Benson, B., Benson, J., Birman, D., Bonacchi, N., Buchanan, E.K., Bruijns, S., et al. (2025). Brain-wide representations of prior information in mouse decision-making. *Nature* 645, 192–200. <https://doi.org/10.1038/s41586-025-09226-1>.
14. Mansouri, F.A., Freedman, D.J., and Buckley, M.J. (2020). Emergence of abstract rules in the primate brain. *Nat. Rev. Neurosci.* 21, 595–610. <https://doi.org/10.1038/s41583-020-0364-5>.
15. Badre, D., Kayser, A.S., and D'Esposito, M. (2010). Frontal Cortex and the Discovery of Abstract Action Rules. *Neuron* 66, 315–326. <https://doi.org/10.1016/j.neuron.2010.03.025>.
16. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. <https://doi.org/10.1038/nature12742>.
17. Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature* 411, 953–956. <https://doi.org/10.1038/35082081>.
18. Rodgers, C.C., and DeWeese, M.R. (2014). Neural Correlates of Task Switching in Prefrontal Cortex and Primary Auditory Cortex in a Novel Stimulus Selection Task for Rodents. *Neuron* 82, 1157–1170. <https://doi.org/10.1016/j.neuron.2014.04.031>.
19. Durstewitz, D., Vittoz, N.M., Floresco, S.B., and Seamans, J.K. (2010). Abrupt Transitions between Prefrontal Neural Ensemble States Accompany Behavioral Transitions during Rule Learning. *Neuron* 66, 438–448. <https://doi.org/10.1016/j.neuron.2010.03.029>.
20. Aoi, M.C., Mante, V., and Pillow, J.W. (2020). Prefrontal cortex exhibits multidimensional dynamic encoding during decision-making. *Nat. Neurosci.* 23, 1410–1420. <https://doi.org/10.1038/s41593-020-0696-5>.
21. Reinert, S., Hübener, M., Bonhoeffer, T., and Goltstein, P.M. (2021). Mouse prefrontal cortex represents learned rules for categorization. *Nature* 593, 411–417. <https://doi.org/10.1038/s41586-021-03452-z>.
22. Miller, E.K. (2000). The prefrontal cortex and cognitive control. *Nat. Rev. Neurosci.* 1, 59–65. <https://doi.org/10.1038/35036228>.
23. Pasupathy, A., and Miller, E.K. (2005). Different time courses of learning-related activity in the prefrontal cortex and striatum. *Nature* 433, 873–876. <https://doi.org/10.1038/nature03287>.
24. Bissonette, G.B., and Roesch, M.R. (2015). Rule encoding in dorsal striatum impacts action selection. *Eur. J. Neurosci.* 42, 2555–2567. <https://doi.org/10.1111/ejn.13042>.
25. Sleezer, B.J., Castagno, M.D., and Hayden, B.Y. (2016). Rule Encoding in Orbitofrontal Cortex and Striatum Guides Selection. *J. Neurosci.* 36, 11223–11237. <https://doi.org/10.1523/JNEUROSCI.1766-16.2016>.
26. Rikhye, R.V., Gilra, A., and Halassa, M.M. (2018). Thalamic regulation of switching between cortical representations enables cognitive flexibility. *Nat. Neurosci.* 21, 1753–1763. <https://doi.org/10.1038/s41593-018-0269-z>.
27. Duan, C.A., Pagan, M., Piet, A.T., Kopec, C.D., Akrami, A., Riordan, A.J., Erlich, J.C., and Brody, C.D. (2021). Collicular circuits for flexible sensorimotor routing. *Nat. Neurosci.* 24, 1110–1120. <https://doi.org/10.1038/s41593-021-00865-x>.
28. Horwitz, G.D., Batista, A.P., and Newsome, W.T. (2004). Representation of an Abstract Perceptual Decision in Macaque Superior Colliculus. *J. Neurophysiol.* 91, 2281–2296. <https://doi.org/10.1152/jn.00872.2003>.
29. Peysakhovich, B., Zhu, O., Tetrack, S.M., Shirhatti, V., Silva, A.A., Li, S., Ibos, G., Rosen, M.C., Johnston, W.J., and Freedman, D.J. (2024). Primate superior colliculus is causally engaged in abstract higher-order cognition. *Nat. Neurosci.* 27, 1999–2008. <https://doi.org/10.1038/s41593-024-01744-x>.
30. Cisek, P., and Kalaska, J.F. (2010). Neural Mechanisms for Interacting with a World Full of Action Choices. *Annu. Rev. Neurosci.* 33, 269–298. <https://doi.org/10.1146/annurev.neuro.051508.135409>.
31. Wang, X.J. (2022). Theory of the Multiregional Neocortex: Large-Scale Neural Dynamics and Distributed Cognition. *Annu. Rev. Neurosci.* 45, 533–560. <https://doi.org/10.1146/annurev-neuro-110920-035434>.
32. Oh, S.W., Harris, J.A., Ng, L., Winslow, B., Cain, N., Mihalas, S., Wang, Q., Lau, C., Kuan, L., Henry, A.M., et al. (2014). A mesoscale connectome

- of the mouse brain. *Nature* 508, 207–214. <https://doi.org/10.1038/nature13186>.
33. Foster, N.N., Barry, J., Korobkova, L., Garcia, L., Gao, L., Becerra, M., Sherafat, Y., Peng, B., Li, X., Choi, J.-H., et al. (2021). The mouse cortico-basal ganglia-thalamic network. *Nature* 598, 188–194. <https://doi.org/10.1038/s41586-021-03993-3>.
34. Hintiryan, H., Foster, N.N., Bowman, I., Bay, M., Song, M.Y., Gou, L., Yamashita, S., Bienkowski, M.S., Zingg, B., Zhu, M., et al. (2016). The mouse cortico-striatal projectome. *Nat. Neurosci.* 19, 1100–1114. <https://doi.org/10.1038/nn.4332>.
35. Benavidez, N.L., Bienkowski, M.S., Zhu, M., Garcia, L.H., Fayzullina, M., Gao, L., Bowman, I., Gou, L., Khanjani, N., Cotter, K.R., et al. (2021). Organization of the inputs and outputs of the mouse superior colliculus. *Nat. Commun.* 12, 4004. <https://doi.org/10.1038/s41467-021-24241-2>.
36. Scheibel, A.B. (1984). The Brain Stem Reticular Core and Sensory Function. In *Comprehensive Physiology*, R. Terjung, ed. (Wiley), pp. 213–256. <https://doi.org/10.1002/cphy.cp010306>.
37. Scheibel, M., Scheibel, A., Mollica, A., and Moruzzi, G. (1955). Convergence and interaction of afferent impulses on single units of reticular formation. *J. Neurophysiol.* 18, 309–331. <https://doi.org/10.1152/jn.1955.18.4.309>.
38. Scheibel, M.E., and Scheibel, A.B. (1968). The Brain Stem Reticular Core—An Integrative Matrix. In *Systems Theory and Biology*, M.D. Mesarović, ed. (Springer), pp. 261–285. [https://doi.org/10.1007/978-3-642-88343-9\\_13](https://doi.org/10.1007/978-3-642-88343-9_13).
39. McElvain, L.E., Chen, Y., Moore, J.D., Brigidi, G.S., Bloodgood, B.L., Lim, B.K., Costa, R.M., and Kleinfeld, D. (2021). Specific populations of basal ganglia output neurons target distinct brain stem areas while collateralizing throughout the diencephalon. *Neuron* 109, 1721–1738.e4. <https://doi.org/10.1016/j.neuron.2021.03.017>.
40. Peng, H., Xie, P., Liu, L., Kuang, X., Wang, Y., Qu, L., Gong, H., Jiang, S., Li, A., Ruan, Z., et al. (2021). Morphological diversity of single neurons in molecularly defined cell types. *Nature* 598, 174–181. <https://doi.org/10.1038/s41586-021-03941-1>.
41. Winnubst, J., Bas, E., Ferreira, T.A., Wu, Z., Economo, M.N., Edson, P., Arthur, B.J., Bruns, C., Rokicki, K., Schauder, D., et al. (2019). Reconstruction of 1,000 Projection Neurons Reveals New Cell Types and Organization of Long-Range Connectivity in the Mouse Brain. *Cell* 179, 268–281.e13. <https://doi.org/10.1016/j.cell.2019.07.042>.
42. Gao, L., Liu, S., Gou, L., Hu, Y., Liu, Y., Deng, L., Ma, D., Wang, H., Yang, Q., Chen, Z., et al. (2022). Single-neuron projectome of mouse prefrontal cortex. *Nat. Neurosci.* 25, 515–529. <https://doi.org/10.1038/s41593-022-01041-5>.
43. Inagaki, H.K., Chen, S., Ridder, M.C., Sah, P., Li, N., Yang, Z., Hasanbegovic, H., Gao, Z., Gerfen, C.R., and Svoboda, K. (2022). A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* 185, 1065–1081.e23. <https://doi.org/10.1016/j.cell.2022.02.006>.
44. Kolmac, C.I., and Mitrofanis, J. (1998). Patterns of brainstem projection to the thalamic reticular nucleus. *J. Comp. Neurol.* 396, 531–543. [https://doi.org/10.1002/\(SICI\)1096-9861\(19980713\)396:4<531::AID-CNE9>3.0.CO;2-2](https://doi.org/10.1002/(SICI)1096-9861(19980713)396:4<531::AID-CNE9>3.0.CO;2-2).
45. Kolmac, C.I., Power, B.D., and Mitrofanis, J. (1998). Patterns of connections between zona incerta and brainstem in rats. *J. Comp. Neurol.* 396, 544–555. [https://doi.org/10.1002/\(SICI\)1096-9861\(19980713\)396:4<544::AID-CNE10>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1096-9861(19980713)396:4<544::AID-CNE10>3.0.CO;2-G).
46. Krout, K.E., Belzer, R.E., and Loewy, A.D. (2002). Brainstem projections to midline and intralaminar thalamic nuclei of the rat. *J. Comp. Neurol.* 448, 53–101. <https://doi.org/10.1002/cne.10236>.
47. Lee, A.M., Hoy, J.L., Bonci, A., Wilbrecht, L., Stryker, M.P., and Niell, C.M. (2014). Identification of a Brainstem Circuit Regulating Visual Cortical State in Parallel with Locomotion. *Neuron* 83, 455–466. <https://doi.org/10.1016/j.neuron.2014.06.031>.
48. Chen, Z., Liu, Y., Yang, Y., Wang, L., Qin, M., Jiang, Z., Xu, M., and Zhang, S. (2025). Whole-brain mapping of basal forebrain cholinergic neurons reveals a long-range reciprocal input-output loop between distinct subtypes. *Sci. Adv.* 11, eadt1617. <https://doi.org/10.1126/sciadv.adt1617>.
49. Moruzzi, G., and Magoun, H.W. (1949). Brain stem reticular formation and activation of the EEG. *Electroencephalogr. Clin. Neurophysiol.* 1, 455–473. [https://doi.org/10.1016/0013-4694\(49\)90219-9](https://doi.org/10.1016/0013-4694(49)90219-9).
50. Kilmer, W.L., McCulloch, W.S., and Blum, J. (1968). Some Mechanisms for a Theory of the Reticular Formation. In *Systems Theory and Biology*, M.D. Mesarović, ed. (Springer), pp. 286–375. [https://doi.org/10.1007/978-3-642-88343-9\\_14](https://doi.org/10.1007/978-3-642-88343-9_14).
51. McCormick, D.A. (1989). Cholinergic and noradrenergic modulation of thalamocortical processing. *Trends Neurosci.* 12, 215–221. [https://doi.org/10.1016/0166-2236\(89\)90125-2](https://doi.org/10.1016/0166-2236(89)90125-2).
52. Grady, F.S., Boes, A.D., and Geerling, J.C. (2022). A Century Searching for the Neurons Necessary for Wakefulness. *Front. Neurosci.* 16, 930514. <https://doi.org/10.3389/fnins.2022.930514>.
53. Lewandowski, M.H., Müller, C.M., and Singer, W. (1993). Reticular facilitation of cat visual cortical responses is mediated by nicotinic and muscarinic cholinergic mechanisms. *Exp. Brain Res.* 96, 1–7. <https://doi.org/10.1007/BF00230433>.
54. Graf, W.M., and Ugolini, G. (2006). The central mesencephalic reticular formation: Its role in space-time coordinated saccadic eye movements. *J. Physiol.* 570, 433–434. <https://doi.org/10.1113/jphysiol.2005.103184>.
55. Roseberry, T.K., Lee, A.M., Lalive, A.L., Wilbrecht, L., Bonci, A., and Kreitzer, A.C. (2016). Cell-Type-Specific Control of Brainstem Locomotor Circuits by Basal Ganglia. *Cell* 164, 526–537. <https://doi.org/10.1016/j.cell.2015.12.037>.
56. Ferreira-Pinto, M.J., Kanodia, H., Falasconi, A., Sigrist, M., Esposito, M.S., and Arber, S. (2021). Functional diversity for body actions in the mesencephalic locomotor region. *Cell* 184, 4564–4578.e18. <https://doi.org/10.1016/j.cell.2021.07.002>.
57. Burgess, C.P., Lak, A., Steinmetz, N.A., Zátka-Haas, P., Bai Reddy, C., Jacobs, E.A.K., Linden, J.F., Paton, J.J., Ranson, A., Schröder, S., et al. (2017). High-Yield Methods for Accurate Two-Alternative Visual Psychophysics in Head-Fixed Mice. *Cell Rep.* 20, 2513–2524. <https://doi.org/10.1016/j.celrep.2017.08.047>.
58. Tseng, S.-Y., Chettih, S.N., Arlt, C., Barroso-Luque, R., and Harvey, C.D. (2022). Shared and specialized coding across posterior cortical areas for dynamic navigation decisions. *Neuron* 110, 2484–2502.e16. <https://doi.org/10.1016/j.neuron.2022.05.012>.
59. Liu, Y., Xin, Y., and Xu, N.L. (2021). A cortical circuit mechanism for structural knowledge-based flexible sensorimotor decision-making. *Neuron* 109, 2009–2024.e6. <https://doi.org/10.1016/j.neuron.2021.04.014>.
60. Wu, Z., Litwin-Kumar, A., Shammash, P., Taylor, A., Axel, R., and Shadlen, M.N. (2020). Context-Dependent Decision Making in a Premotor Circuit. *Neuron* 106, 316–328.e6. <https://doi.org/10.1016/j.neuron.2020.01.034>.
61. Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R., et al. (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science* 346, 458–463. <https://doi.org/10.1126/science.1256573>.
62. Izquierdo, A., Brigman, J.L., Radke, A.K., Rudebeck, P.H., and Holmes, A. (2017). The neural basis of reversal learning: An updated perspective. *Neuroscience* 345, 12–26. <https://doi.org/10.1016/j.neuroscience.2016.03.021>.
63. Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704–1711. <https://doi.org/10.1038/nn1560>.
64. Rescorla, R., and Wagner, A. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In

Classical Conditioning II: Current Research and Theory (Appleton-Century-Crofts).

65. Miller, K.J., Eckstein, M., Botvinick, M.M., and Kurth-Nelson, Z. (2023). Cognitive Model Discovery via Disentangled RNNs. Preprint at bioRxiv. <https://doi.org/10.1101/2023.06.23.546250>.
66. Evenden, J.L., and Robbins, T.W. (1984). Win-Stay Behaviour in the Rat. *Q. J. Exp. Psychol. B* 36, 1–26. <https://doi.org/10.1080/14640748.408402190>.
67. Wang, Q., Ding, S.-L., Li, Y., Royall, J., Feng, D., Lesnar, P., Graddis, N., Naeemi, M., Facer, B., Ho, A., et al. (2020). The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* 181, 936–953.e20. <https://doi.org/10.1016/j.cell.2020.04.007>.
68. Erő, C., Gewaltig, M.-O., Keller, D., and Markram, H. (2018). A Cell Atlas for the Mouse Brain. *Front. Neuroinform.* 12, 84. <https://doi.org/10.3389/fninf.2018.00084>.
69. Park, Y.-G., Sohn, C.H., Chen, R., McCue, M., Yun, D.H., Drummond, G.T., Ku, T., Evans, N.B., Oak, H.C., Trieu, W., et al. (2018). Protection of tissue physicochemical properties using polyfunctional crosslinkers. *Nat. Biotechnol.* 37, 73–83. <https://doi.org/10.1038/nbt.4281>.
70. Cisek, P. (2021). Evolution of behavioural control from chordates to primates. *Philos. Trans. R. Soc. Lond., B Biol. Sci.* 377, 20200522. <https://doi.org/10.1098/rstb.2020.0522>.
71. Steinmetz, N.A., Aydin, C., Lebedeva, A., Okun, M., Pachitariu, M., Bauza, M., Beau, M., Bhagat, J., Böhm, C., Broux, M., et al. (2021). Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* 372, eabf4588. <https://doi.org/10.1126/science.abf4588>.
72. Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81, 179–194. <https://doi.org/10.1016/j.neuron.2013.10.020>.
73. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* 22, 1677–1686. <https://doi.org/10.1038/s41593-019-0502-4>.
74. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364, 255. <https://doi.org/10.1126/science.aav7893>.
75. Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature* 497, 585–590. <https://doi.org/10.1038/nature12160>.
76. Ebitz, R.B., and Hayden, B.Y. (2021). The population doctrine in cognitive neuroscience. *Neuron* 109, 3055–3068. <https://doi.org/10.1016/j.NEURON.2021.07.011>.
77. Vyas, S., Golub, M.D., Sussillo, D., and Shenoy, K.V. (2020). Computation Through Neural Population Dynamics. *Annu. Rev. Neurosci.* 43, 249–275. <https://doi.org/10.1146/annurev-neuro-092619-094115>.
78. Yule, G.U. (1926). Why do we Sometimes get Nonsense-Correlations between Time-Series?—A Study in Sampling and the Nature of Time-Series. *J. R. Stat. Soc.* 89, 1–63. <https://doi.org/10.2307/2341482>.
79. Harris, K.D. (2020). Nonsense Correlations in Neuroscience. Preprint at bioRxiv. <https://doi.org/10.1101/2020.11.29.402719>.
80. Park, I.M., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat. Neurosci.* 1710, 1395–1403. <https://doi.org/10.1038/nn.3800>.
81. Chow, S., and Rodgers, P. (2005). Extended Abstract: Constructing Area-Proportional Venn and Euler Diagrams with Three Circles. <https://kar.kent.ac.uk/14285/1/EulerStir.pdf>.
82. Montijn, J.S., Seignette, K., Howlett, M.H., Cazemier, J.L., Kamermans, M., Levelt, C.N., and Heimel, J.A. (2021). A parameter-free statistical test for neuronal responsiveness. *eLife* 10, e71969. <https://doi.org/10.7554/eLife.71969>.
83. Fine, J.M., and Hayden, B.Y. (2022). The whole prefrontal cortex is premotor cortex. *Philos. Trans. R. Soc. Lond., B Biol. Sci.* 377, 20200524. <https://doi.org/10.1098/rstb.2020.0524>.
84. Janacek, K., Evans, T.M., Kiss, M., Shah, L., Blumenfeld, H., and Ullman, M.T. (2022). Subcortical Cognition: The Fruit Below the Rind. *Annu. Rev. Neurosci.* 45, 361–386. <https://doi.org/10.1146/annurev-neuro-110920-013544>.
85. Wang, Z.A., Chen, S., Liu, Y., Liu, D., Svoboda, K., Li, N., and Druckmann, S. (2023). Not everything, not everywhere, not all at once: a study of brain-wide encoding of movement. Preprint at bioRxiv. <https://doi.org/10.1101/2023.06.08.544257>.
86. Finkelstein, A., Fontolan, L., Economo, M.N., Li, N., Romani, S., and Svoboda, K. (2021). Attractor dynamics gate cortical information flow during decision-making. *Nat. Neurosci.* 24, 843–850. <https://doi.org/10.1038/s41593-021-00840-6>.
87. Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161. <https://doi.org/10.1126/science.284.5417.1158>.
88. Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2014). Cortical activity in the null space: permitting preparation without movement. *Nat. Neurosci.* 17, 440–448. <https://doi.org/10.1038/nn.3643>.
89. Harris, J.A., Mihalas, S., Hirokawa, K.E., Whitesell, J.D., Choi, H., Bernard, A., Bohn, P., Caldejon, S., Casal, L., Cho, A., et al. (2019). Hierarchical organization of cortical and thalamic connectivity. *Nature* 575, 195–202. <https://doi.org/10.1038/s41586-019-1716-z>.
90. Thomas, A., Yang, W., Wang, C., Tipparaju, S.L., Chen, G., Sullivan, B., Swiekatowski, K., Tatam, M., Gerfen, C., and Li, N. (2023). Superior colliculus bidirectionally modulates choice activity in frontal cortex. *Nat. Commun.* 14, 7358. <https://doi.org/10.1038/s41467-023-43252-9>.
91. Peters, A.J., Fabre, J.M.J., Steinmetz, N.A., Harris, K.D., and Carandini, M. (2021). Striatal activity topographically reflects cortical activity. *Nature* 597, 420–425. <https://doi.org/10.1038/s41586-020-03166-8>.
92. Richman, E.B., Ticea, N., Allen, W.E., Deisseroth, K., and Luo, L. (2023). Neural landscape diffusion resolves conflicts between needs across time. *Nature* 623, 571–579. <https://doi.org/10.1038/s41586-023-06715-z>.
93. Sherman, S.M., and Usrey, W.M. (2021). Cortical control of behavior and attention from an evolutionary perspective. *Neuron* 109, 3048–3054. <https://doi.org/10.1016/j.neuron.2021.06.021>.
94. Morin, L.P., and Studholme, K.M. (2014). Retinofugal projections in the mouse. *J. Comp. Neurol.* 522, 3733–3753. <https://doi.org/10.1002/cne.23635>.
95. Yao, Z., van Velthoven, C.T.J., Kunst, M., Zhang, M., McMillen, D., Lee, C., Jung, W., Goldy, J., Abdelhak, A., Aitken, M., et al. (2023). A high-resolution transcriptomic and spatial atlas of cell types in the whole mouse brain. *Nature* 624, 317–332. <https://doi.org/10.1038/s41586-023-06812-z>.
96. Guo, Z.V., Hires, S.A., Li, N., O'Connor, D.H., Komiyama, T., Ophir, E., Huber, D., Bonardi, C., Morandell, K., Gutnisky, D., et al. (2014). Procedures for Behavioral Experiments in Head-Fixed Mice. *PLoS One* 9, e88678. <https://doi.org/10.1371/journal.pone.0088678>.
97. Ottenheimer, D.J., Hjort, M.M., Bowen, A.J., Steinmetz, N.A., and Stuber, G.D. (2023). A stable, distributed code for cue value in mouse cortex during reward learning. *eLife* 12, RP84604. <https://doi.org/10.7554/eLife.84604>.
98. Zátka-Haas, P., Steinmetz, N.A., Carandini, M., and Harris, K.D. (2021). Sensory coding and the causal impact of mouse cortex in a visual decision. *eLife* 10, e63163. <https://doi.org/10.7554/eLife.63163>.
99. Ye, Z., Bull, M.S., Li, A., Birman, D., Daigle, T.L., Tasic, B., Zeng, H., and Steinmetz, N.A. (2023). Brain-wide topographic coordination of traveling spiral waves. Preprint at bioRxiv. <https://doi.org/10.1101/2023.12.07.570517>.

100. Ertürk, A., Becker, K., Jähring, N., Mauch, C.P., Hojer, C.D., Egen, J.G., Hellal, F., Bradke, F., Sheng, M., and Dodt, H.-U. (2012). Three-dimensional imaging of solvent-cleared organs using 3DISCO. *Nat. Protoc.* *7*, 1983–1995. <https://doi.org/10.1038/nprot.2012.119>.
101. Shamonin, D.P., Bron, E.E., Lelieveldt, B.P.F., Smits, M., Klein, S., and Staring, M.; Alzheimer's Disease Neuroimaging Initiative (2014). Fast Parallel Image Registration on CPU and GPU for Diagnostic Classification of Alzheimer's Disease. *Front. Neuroinform.* *7*, 50. <https://doi.org/10.3389/fninf.2013.00050>.
102. Campbell, R., Blot, A., Rousseau, C., and Winter, O. (2025). Lasagna. GitHub. <https://github.com/SainsburyWellcomeCentre/lasagna>.
103. Faulkner, M.A. (2025). iblapps. GitHub. <https://github.com/int-brain-lab/iblapps/tree/master/atlas electrophysiology>.
104. Liu, L.D., Chen, S., Hou, H., West, S.J., Faulkner, M., International; Brain Laboratory, Economo, M.N., Li, N., and Svoboda, K. (2021). Accurate Localization of Linear Probe Electrode Arrays across Multiple Brains. *eNeuro* *8*. <https://doi.org/10.1523/ENEURO.0241-21.2021>.
105. LifeCanvas Technologies (2025). Passive Delipidation Protocol. <https://sites.google.com/lifecanvastech.com/protocol/delipidation/passive-delipidation-protocol>.
106. Bhagat, J., Wells, M.J., Harris, K.D., Carandini, M., and Burgess, C.P. (2020). Rigbox: An Open-Source Toolbox for Probing Neurons and Behavior. *eNeuro* *7*. <https://doi.org/10.1523/ENEURO.0406-19.2020>.
107. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* *21*, 1281–1289. <https://doi.org/10.1038/s41593-018-0209-y>.
108. Birman, D., Yang, K.J., West, S.J., Karsh, B., and Browning, Y.; Laboratory, the I.B., Siegle (2023). Pinpoint: Trajectory planning for multi-probe electrophysiology and injections in an interactive web-based 3D environment. *eLife* *12*. <https://doi.org/10.7554/eLife.91662.1>.
109. Banga, K., Benson, J., Bhagat, J., Biderman, D., Birman, D., Bonacchi, N., Bruijns, S.A., Buchanan, K., Campbell, R.A., Carandini, M., et al. (2024). Reproducibility of in vivo electrophysiological measurements in mice. *eLife* *13*. <https://doi.org/10.7554/eLife.100840.1>.
110. Rossant, C., Kadir, S.N., Goodman, D.F.M., Schulman, J., Hunter, M.L.D., Saleem, A.B., Grosmark, A., Belluscio, M., Denfield, G.H., Ecker, A.S., et al. (2016). Spike sorting for large, dense electrode arrays. *Nat. Neurosci.* *19*, 634–641. <https://doi.org/10.1038/nn.4268>.
111. Izenman, A.J. (1975). Reduced-rank regression for the multivariate linear model. *J. Multivariate Anal.* *5*, 248–264. [https://doi.org/10.1016/0047-259X\(75\)90042-1](https://doi.org/10.1016/0047-259X(75)90042-1).
112. Ratcliff, R., and McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Comput.* *20*, 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>.
113. Ye, Z., Shelton, A.M., Shaker, J.R., Boussard, J., Colonell, J., Birman, D., Manavi, S., Chen, S., Windolf, C., Hurwitz, C., et al. (2025). Ultra-high-density Neuropixels probes improve detection and identification in neuronal recordings. *Neuron* *113*, 3966–3982.e12. <https://doi.org/10.1016/j.neuron.2025.08.030>.
114. Harris, K.D., and Miller, K.J. (2023). Conditional Randomization Tests for Behavioral and Neural Time Series. Preprint at arXiv. <https://doi.org/10.48550/arXiv.2311.03554>.

## STAR★METHODS

### KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
<b>Bacterial and virus strains</b>		
AAV2-hSyn-EGFP	Addgene	Cat# 50465-AAV2
<b>Chemicals, peptides, and recombinant proteins</b>		
CM-Dil (red lipophilic dye)	Invitrogen/Thermo Fisher	Cat# V22888
Kwik-Cast silicone sealant	World Precision Instruments	Cat# KWIK-CAST
Dowsil 3-4680 dura gel	Dow Chemical	Cat# 3-4680
Metabond	Parkell	Cat# S380
<b>Deposited data</b>		
Electrophysiology and behavioral dataset	This paper	<a href="https://doi.org/10.6084/m9.figshare.31428575">https://doi.org/10.6084/m9.figshare.31428575</a>
<b>Experimental models: Organisms/strains</b>		
Mouse: C57BL/6J	The Jackson Laboratory	JAX Stock #000664; RRID: IMSR_JAX:000664
<b>Software and algorithms</b>		
Kilosort 2.5	Steinmetz et al. <sup>71</sup>	<a href="https://github.com/MouseLand/Kilosort">https://github.com/MouseLand/Kilosort</a>
<b>Other</b>		
Neuropixels 2.0 probes	IMEC	Available at <a href="https://www.neuropixels.org/">https://www.neuropixels.org/</a>

### EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

All experimental procedures were performed according to US National Institutes of Health guidelines for animal research, in strict accordance with protocol 4461–01 and approved by the Institutional Animal Care and Use Committee at the University of Washington.

Subjects were 18 male and female C57BL/6J mice (n=8 male, n=10 female) single-housed on a 12 hr light/dark cycle and aged 22–68 weeks at the time of recordings. Seven of the mice were trained on the task (n=3 male, n=4 female), and behavioral data from each of these mice was used for behavioral modeling. Five of these trained mice (n=2 male, n=3 female) were used for electrophysiological recordings during the task. Six mice (n=3 male, n=3 female) were not trained on the behavioral task and were used for electrophysiological recordings during passive stimulus mapping (Figures 5 and S9, “task-naive” group; task and control stimuli were shown in random order). Four of the trained mice (n=2 male, n=2 female) and two additional mice (n=2 female) that were trained on an early version of the behavioral task were used for passive stimulus mapping neural recording data only (“trained” group). Three mice (n=2 male, n=1 female) were not trained on the behavioral task and were used for electrophysiological recordings during task-naive replay experiments (Figures S6F and S6G; task stimuli and reward delivery of an example session were replayed with the exact same timing). All mice were the wild-type genotype (C57BL/6J). Electrophysiology experiments were performed during the light cycle. Mice were given free access to food in their home cages for the duration of the experiment. Mice trained on the behavioral task were water restricted for the duration of the experiments and maintained at above 80% of their baseline weight.<sup>96</sup> Task-naive mice were water restricted for one week prior to and during the electrophysiology experiment period.

### METHOD DETAILS

#### Surgical procedures

For all surgical procedures, mice were anesthetized under 3–4% isoflurane in an induction chamber first, then maintained at 1–2% isoflurane for the duration of the 1–2 hour procedure. Carprofen (5 mg/kg) was administered subcutaneously and lidocaine (2 mg/kg) was injected under the scalp for postoperative analgesia. The mouse was kept in a warmed recovery chamber until recovery from anesthesia before being returned to its home cage. Mice were allowed to recover in the home cage for at least 5 days before habituation to head fixation and task training. Carprofen (0.05 mg/ml) was given for 3 days in water after surgery.

#### Surgeries for electrophysiology

The headbar implantation and craniotomy surgical preparations were performed as previously described.<sup>9,97,98</sup> The scalp was shaved and further cleaned with hair removal cream before mice were transferred to a separate stereotaxic frame for surgery. Eye ointment (Alcon, Systane) was applied before surgery to prevent drying and body temperature was kept at 37°C with a far infrared

warming pad (Kent Scientific). The skin and periosteum connective tissue was cleared off the dorsal surface of the skull. A layer of cyanoacrylate (Vetbond Tissue Adhesive) was then applied to seal the junction between the exposed skull and cut skin. A 3D-printed plastic recording chamber was implanted on top of the skull with dental cement (C&B Metabond), and a custom-made steel head-plate was attached to the skull over the interparietal bone with Metabond for head fixation. To a protective layer over the skull, two separate thin layers of fast curing optical adhesive (NOA81, Thorlabs) were applied and cured until solidification with a 365 nm, 3 W output power UV Flashlight (LIGHTFE, UV301D).

After completion of task training, brief (15–30 min) surgery was conducted to perform 1 to 2 craniotomies for Neuropixels probe insertion. Briefly, following induction of anesthesia, small ( $2 \times 1.5$  mm ( $w \times h$ )) craniotomies were made over areas of interest with a handheld dental drill. The craniotomies were covered with a soft silicone gel (DOWSIL 3–4680) or Kwik-Cast, and the recording chamber was covered with a 3D-printed lid sealed with Kwik-Cast elastomer to protect craniotomies from debris.

### Viral injections

Viral injections ( $n=3$  mice) were performed with the Nanoject III (Drummond Scientific). Glass pipettes were pulled using the Narashige pipette puller PE-22 with 1.0 mm OD glass and the following settings: heater 70.5, sub magnet 20.9, main magnet 100.7. Pipettes were 5–6 mm in length and broken back to a tip diameter of  $\sim 30$  microns.

After performing a small ( $\sim 0.5 \times 0.5$  mm) craniotomy at the desired AP and ML coords, the pipette was driven down to and 100  $\mu$ m past the target injection depth. It was then drawn back 100  $\mu$ m to create a pocket for the injection. Injection volumes ranged from 50–75 nl with an injection speed of 1 nl/s. After the injection was completed, there was a 2 minute wait time, after which there were 2 sets of withdrawing 100  $\mu$ m and waiting 1 minute between each, then 5 sets of withdrawing 100  $\mu$ m waiting 30 seconds between each. The pipette was then fully withdrawn out of the brain over about 2 minutes.

Virus	Virus volume	Brain region/coordinates (mm from bregma)
AAV2-hSyn-EGFP (Addgene, viral titer $2 \times 10^{11}$ - $4 \times 10^{11}$ vg/ml (1:8-1:12 dilutions))	50-75 nl	<ul style="list-style-type: none"><li>● MRF intermediate anterior, (AP <math>-3.60</math>, ML <math>+1</math>, DV <math>3.25</math>)</li><li>● MRF intermediate posterior, (AP <math>-3.90</math>, ML <math>+1</math>, DV <math>3.0</math>)</li><li>● MRF posterior, (AP <math>-4.20</math>, ML <math>+0.5</math>, DV <math>2.9</math>)</li></ul>

### Histology and imaging

Animals were anesthetized with isoflurane. Mice were perfused intracardially with 0.9% phosphate buffer saline followed by 4% paraformaldehyde (PFA).

### Probe tracking

Post-hoc tracking of electrophysiological recording trajectories was performed as previously described.<sup>97,99</sup> Brains were extracted immediately following perfusion and post-fixed in 4% paraformaldehyde for 24 h. In preparation for light sheet imaging, brains were cleared using organic solvents following the 3DISCO protocol<sup>100</sup> (<https://idisco.info/>), with some modification. Briefly, on day 1 brains were washed 3X in PBS and dehydrated in a series of increasing MeOH concentrations (20%, 40%, 60%, 80%, 100%, 100%; 1 hr each) then incubated overnight for lipid extraction in 66% dichloromethane (DCM) in MeOH. On day 2 brains were washed 2X twice in 100% MeOH for 1 hr each, then bleached overnight in 5% H<sub>2</sub>O<sub>2</sub> in MeOH at 4°C. On day 3 brains were washed 2X in 100% MeOH, then final lipid extraction was accomplished in a series of DCM incubations (3 hr in 66% DCM in MeOH, 2X 100% DCM for 15 min each) before immersion in dibenzyl ether (DBE) for refractive index matching.

Brains were imaged on a light sheet microscope (LifeCanvas Technologies SmartSPIM or LaVision Ultramicroscope II) at least 3 days after completion of the iDISCO procedure. Brains were immersed in DBE in the imaging well secured in the horizontal position, and illuminated by one or two light sheets (100% width) imaging with the 1.6x objective ( $4 \times 4 \mu$ m resolution). Images were gathered from the dorsal surface of the brain to the ventral surface in 10  $\mu$ m steps in 488 nm (autofluorescence,  $\sim 30\%$  power) and 594 nm (Dil,  $\sim 10\%$  power) excitation channels. The raw TIF images were compiled into a single multi-image file with 10  $\mu$ m voxels, then spatially downsampled to 25  $\mu$ m voxels for transformation to the Allen CCF volume<sup>67</sup> using the Elastix algorithm.<sup>101</sup> CCF-transformed volumes were used to generate CCF fluorescent probe tract locations (pixel coordinates along the probe tract) using Lasagna.<sup>102</sup> Probe track CCF pixel coordinates (origin front, top, left) were transformed to bregma coordinates (origin bregma, x==ML, y==AP, and z==DV) in preparation for final integration with electrophysiology recordings using the International Brain Lab electrophysiology atlas graphical user interface (GUI).<sup>103</sup> Alignment between histological tracks and electrophysiological features was done for each shank. For this, we considered sorted spike firing rates and amplitudes, a correlation matrix between channels, and single neuron waveforms on each channel in relation to the estimated channel locations in Atlas space from the tracked probe. The recording sites were then aligned to the Atlas by manually identifying a warping such that recording sites were best fit to the electrophysiological characteristics of the brain regions (e.g. matching location of ventricles or white matter tracts with low firing activity bands). This procedure has been estimated to have a 70  $\mu$ m error.<sup>9,104</sup> Individual neuron locations were determined using the recording channel brain coordinates of each unit's maximum-amplitude waveform.

### AAV injections

To visualize whole-brain structure while preserving fluorescence of virally induced fluorescent proteins, we performed tissue clearing using stabilization under harsh conditions via intramolecular epoxide linkages to prevent degradation (SHIELD)<sup>69</sup> with the Clear+ passive clearing kit and protocol from LifeCanvas Technologies. The protocol is explained in detail by LifeCanvas Technologies<sup>105</sup>; here we briefly describe each step. Brains were extracted immediately following perfusion and post-fixed in 4% paraformaldehyde for 24 h. SHIELD preservation was performed by first incubating in fresh SHIELD OFF solution (25% DI H<sub>2</sub>O, 25% SHIELD-buffer solution, 50% SHIELD-Epoxy solution) for 24 hours at 4°C with shaking for 4 days. Following this incubation, brains were incubated in SHIELD ON Buffer at 37°C with shaking for 24 hours. To clear brains by delipidation, brains were incubated in delipidation buffer at 45°C with shaking for 7–8 days. Finally, brains were index matched at 37°C with shaking using 1:2 diluted EasyIndex solution for 24 hours, followed by 100% EasyIndex solution for 24–48 hours (until transparent).

Brains were mounted (horizontal orientation) onto custom 3D-printed holders by embedding into an EasyIndex agarose solution. Once brains were mounted in the agarose, the whole sample was immersed in EasyIndex overnight and imaged the following day. Brains were imaged on a light sheet microscope (LifeCanvas Technologies SmartSPIM) using two-sheet illumination (100% width) and the 3.6x objective (1.8x1.8 μm resolution). Images were gathered from the dorsal surface of the brain to the ventral surface in 2 μm steps in 488 nm (EGFP, ~15% power) and 639 nm (autofluorescence, ~70% power) excitation channels. The full resolution image stacks as well as CCF-transformed stacks (using procedure as described above) were then used for visualization and quantification (Figures 2B, 2C, and S2).

### Analysis of MRF projections with AAV tracing

All image processing and data analysis for anterograde AAV injections (Figures 2B, 2C, and S2) were performed using custom scripts written in MATLAB (R2020a, The MathWorks). The analysis pipeline was designed to quantify fluorescently labeled axonal projections from whole-brain light-sheet images registered to the 25-micron resolution Allen CCF.

### Signal correction

Raw signal from the 488 nm (GFP) channel was corrected for autofluorescence by subtracting the signal from the 639 nm (autofluorescence) channel. A robust linear regression was performed between the two channels on voxels below the 98th percentile of 488 nm intensity to determine a bleed-through coefficient. Additionally, bright, non-linear artifacts present in both channels were identified by thresholding the autofluorescence channel (experimenter-chosen, 99.7%–99.9%) and removed from the corrected signal volume.

### Injection site and projection mask generation

Injection sites were identified from the corrected signal volume. An initial intensity threshold was set manually using a GUI that visualized maximum intensity projections (MIP) of the volume and corresponding injection site masks for varying intensity thresholds. Thresholds were set to maximize inclusion of the site. To isolate the injection core from dense, nearby projections, the resulting binary mask was processed to retain only the single largest connected component (`bwconncomp` in MATLAB). The injection site mask was then visualized to manually verify its accuracy, and the initial intensity threshold was refined as needed.

A projection mask was generated from the volume after excluding the identified injection site, using a similar process as injection site identification. An intensity threshold was set using the GUI. This initial projection mask was filtered to remove small, disconnected objects smaller than 5 voxels (`bwareaopen` in MATLAB) to reduce noise. The projection mask was then visualized, and the threshold was refined as needed.

### Manual Mask Refinement

The projection mask for each dataset underwent a final refinement step using a custom GUI, in which artifacts were manually erased and faint projections that fell below the global threshold were added.

### Projection region labeling and quantification

The refined projection mask was used to quantify signal within each brain region. Each voxel in the projection mask was assigned to a CCF-defined region, from a manually selected list of 308 regions.<sup>11</sup>

For each region, we calculated the total volume of projecting axons (mm<sup>3</sup>) and the summed intensity of the corrected 488 nm signal. Injection site composition was calculated as the percentage of the total injection volume located within each region (i.e., # of voxels in region / total # of voxels in injection site mask).

### Flexible rule-based visual decision-making task

The behavioral task uses the same hardware as described previously.<sup>9,57</sup> Mice sat on a plastic apparatus with forepaws on a rotating wheel surrounded by three computer screens (Adafruit, LP097QX1) at right angles covering 270 × 70 degrees of visual angle (d.v.a.) (Figure 1A). Each screen was roughly 11 cm from the mouse's eyes at its nearest point and refreshed at 60 Hz. The screens were fitted with Fresnel lenses (Wuxi Bohai Optics, BHPA220-2-5) to ameliorate reductions in luminance and contrast at larger viewing angles near their edges, and these lenses were coated with scattering window film ('frostbite', The Window Film Company) to reduce reflections. The wheel was a ridged rubber Lego wheel affixed to a rotary encoder (Kubler 05.2400.1122.0360). A plastic tube for delivery of water rewards was placed near the subject's mouth. Experiments were run with Rigbox software<sup>106</sup> for MATLAB (MathWorks, Inc.).

To reduce baseline period movement and discourage the subject from choosing actions before stimulus onset, trials were only initiated after the subject had held the wheel still for a short interval (pre-stim quiescence condition; duration exponentially distributed

between 0.2–0.5 s on each trial; [Figure 1E](#)). At trial initiation, a visual stimulus was present on either the left screen ( $-80^\circ$  azimuth), the center screen directly in front of the mouse ( $0^\circ$ ), or on the right screen ( $80^\circ$ ) ([Figure 1B](#)). All stimuli were present at  $0^\circ$  altitude. As this task was designed to assess context-dependent abstract rule belief rather than perceptual discrimination, stimuli were designed to be highly salient and unambiguous: stimuli were full contrast square wave Gabor patches with orientation  $45^\circ$ , sigma 7 degrees of visual angle (d.v.a), and spatial frequency 0.1 cycles per degree. The patches were presented with a randomized initial phase and flipping phase by 180 degrees at a frequency of 2 Hz. Coincident with stimulus onset was a go cue (10 kHz pure tone, 0.1 s duration), after which movement of the wheel was coupled to movement of the visual stimulus on the screen, and mice were free to move the stimulus to the target location for a reward. Wheel turns in which the top surface of the wheel was moved to the subject's right (clockwise turns) led to rightward movements of stimuli on the screen, that is, a stimulus on the subject's left moved towards the central screen. A CW or CCW turn was registered when the wheel was turned by an amount sufficient to move the visual stimuli by 40 d.v.a. in either direction (i.e., the distance to the intermediate position between the two possible stimuli of a given block). To get the precise times of wheel movement onsets post-hoc, an algorithm was used to identify the earliest moment at which the wheel began detectably moving, described previously.<sup>9</sup> If the wheel turn was in the correct direction (CW for left stimulus/center stimulus Right Block; CCW for right stimulus/center stimulus Left Block), mice received a sucrose reward (with 25 ms delay) which was coincident with a reward tone (15 kHz, 0.1 s duration). If the wheel turn was in the incorrect direction (miss trial), a white noise burst was played (1.5 s duration). If the mouse did not turn the wheel sufficiently in either direction within the response window (10 s), the trial counted as a miss and a white noise burst was played. After feedback delivery, the inter-trial interval consisted of an epoch whose timing depended on the outcome (default 1.5 s after hit, 3 s after miss), during which the visual stimulus stayed on the screen at the location to which it was moved at the time of feedback delivery, followed by an epoch during which the stimulus turned off for a fixed 0.5 s. Then, the pre-stimulus quiescence condition of the next trial began, after which the trial would initiate and a new visual stimulus would appear.

The block structure of the task ([Figure 1C](#)) was predetermined by random draws from a uniform distribution between 40 and 70. Block lengths were only considered on non-repeat trials, i.e., the mouse had to complete the predetermined number of non-repeat trials (between 40 and 70) for the block to switch. To discourage guessing, miss trials were repeated until the correct action was chosen. Repeat trials were not included in the analysis of neural data. Mice performed the task until disengaged ( $\geq 4x$  reaction time difference from baseline), anywhere between 350 to 1000 trials. In each block, there was an equal probability of each of the two stimuli appearing (i.e., 50% left and 50% center stimuli in the Left Block, 50% right and 50% center stimuli in the Right Block), thereby preventing subjects from successfully performing the task using a block-dependent motor bias strategy.

### **Training protocol**

Mice were trained on the task in 3 stages. The first stage was intended to have mice learn to turn the wheel for reward and map distinct stimulus locations with opposing wheel turn directions. Stimuli were presented at either  $-40^\circ$  (left side of center screen) or  $40^\circ$  (right side of center screen) and had to be turned CW or CCW for reward respectively. Once mice achieved high accuracy and initiated movements rapidly ( $\geq 80\%$  performance on non-repeat trials, and with average reaction times at 500 ms or lower), the mice progressed to stage two (at experimenter's discretion).

The second stage of training was intended to have mice slowly generalize the rules mapping left stimuli to CW wheel turns and right stimuli to CCW wheel turns. In other words, mice had to learn to associate the relative positions of the possible stimuli that could appear on the screen with the correct action rather than their absolute positions. Both stimulus locations were gradually shifted to the right or left by the same amount as a function of the mouse's performance. Stimuli were shifted in 6.67 d.v.a increments from the starting positions of  $-40^\circ$  and  $40^\circ$  until they reached either  $0^\circ$  and  $80^\circ$  (full rightward shift) or  $-80^\circ$  and  $0^\circ$  (full leftward shift). Stimuli were shifted based on windowed performance of the mouse on non-repeat trials: 75% accuracy had to be achieved on both stimulus positions for the shift to occur. The performance window started at 200 trials and was decreased by the experimenter as mice more quickly improved their performance on different stimulus positions, down to a 50 trial window at the lowest. The performance window at the right and left extremes was always 3 times larger than the window for the intermediate positions (i.e., the window was 150 trials long for right and left extremes when the parameter was set to 50 trials). Once mice consistently shifted between the left and right stimulus position extremes in single training sessions, they progressed to the third stage.

As performance in stage two was typically very high in the intermediate positions ( $>90\%$ ), but lower for the center ( $0^\circ$ ) stimuli in the left and right extremes, the third stage quickly progressed through the intermediate positions, starting at a fixed length of 10 trials, while staying at the extreme positions for 200 or fewer trials. The intermediate stimulus position window was gradually decreased to 0 over the span of about 2 weeks such that only the left and right extremes remained (i.e. the Left and Right Blocks of the task). At this point, the number of trials it took for mice to improve their performance on center stimuli was monitored, and the length of the blocks was decreased accordingly by the experimenter, until mice were able to successfully switch between Left and Right Blocks with  $>70\%$  performance on center stimuli with block lengths of 50-60 trials. Mice were then ready for the final version of the task used in neural recordings, where block lengths were randomized between 40-70 trials.

Throughout the course of training, the following additional parameters were tuned at the experimenter's discretion to aid in learning of the task: wheel gain, balance of left and right stimulus trials (changed only temporarily when mice had a strong CW or CCW bias; never present during recording sessions), and/or inter-trial interval (defaults were 2 s total after hit, 3.5 s after miss; changed anywhere from 1 to 4 s). The total length of training from day one to neural recordings took between eight and twenty weeks. Roughly 40% of mice trained on the task made it to the final stage, with most mice failing during stage two. Two of the mice included in

behavioral plots (Figures 1F, 1H–1J, S1B, and S1E) who were proficient in the task did not have any electrophysiology recordings performed on them.

### Comparison with related behavioral paradigms

The task developed here shares key features with existing flexible sensorimotor decision-making paradigms for rodents, specifically the requirement to map sensory stimuli to different actions based on context.<sup>18,27,58–61</sup> However, the mechanisms for signaling and maintaining context differ from these prior studies in several ways.

Unlike tasks where the current rule is explicitly signaled by a sensory cue (e.g., a light or sound) on every trial, the context in our task is latent. It must be inferred and internally maintained without an explicit sensory instruction. By contrast, delayed-match-to-sample<sup>60,61</sup> or Pro/Anti<sup>27</sup> tasks require holding a specific sensory sample in short-term working memory.

Our design differs from other block-based rule inference tasks<sup>18,58,59</sup> mainly in the distribution of stimuli across blocks, and consequently, the mechanisms used by subjects to infer the block context. The task here employs a structure wherein the outside (non-reversing) stimuli are block-specific (Figure 1B). This allows the subject to infer the hidden state based on the spatial identity of the outside stimuli, in addition to utilizing reward feedback on center (reversing) stimuli (Figures 1I and 1J). By contrast, tasks wherein all possible stimuli are present in both blocks<sup>18,58,59</sup> require subjects to solely rely on reward feedback to infer context (or in some tasks, the presence of a specific set of cueing stimulus trials at the start of each block<sup>18</sup>).

### Passive stimulus presentation

After each recording session, we performed a passive stimulus presentation experiment (Figure 5A) while continuing to record from the same electrodes. This experiment was also displayed for recordings from task-naive mice. Mice were presented with stimuli from the task: full contrast gabors at  $-80^\circ$ ,  $0^\circ$ , and  $80^\circ$ , go cue beeps, reward tones, and white noise bursts. The gabor stimuli were also presented at azimuths in between those present in the final stage of the task ( $-80^\circ$  to  $80^\circ$  with  $13.33^\circ$  step increments in between). Gabor stimuli lasted for 0.5 seconds. Mice were also presented with control visual stimuli, which were not present in the task: expanding black discs (akin to looming stimulus) and contracting black discs. The contracting and expanding black stimuli shrunk or expanded in size between 0 and 90 d.v.a over 2 s. All stimuli were played at 1–3 s randomized intervals for 20 randomly interleaved repetitions each.

Visual receptive field mapping was performed using sparse noise stimuli with white squares of 8 d.v.a. edge length, positioned on a  $10 \times 36$  grid (some stimulus positions were located partially off-screen) on a black background. The stimuli were shown for 10 monitor frames (167 ms) at a time, and their times of appearance were independently randomly selected to yield an average rate of approximately 0.12 Hz.

### Task-naive replay experiments

To test whether pre-stimulus context signals in MRF could artifactually arise from the stimulus and reward structure of the task independent of active task engagement, we performed task replay control experiments in task-naive mice ( $n=3$  mice, 6 sessions). Neuro-pixels 2.4 probes were inserted bilaterally (2 probes per session) and spanned a range of MRF, including some recordings in the anterior context-coding-enriched portion (Figure S6F). In each session, naive mice were presented with an exact replay of the sequence of task events from the trained session that had exhibited the highest degree of MRF context decoding amongst all sessions. The replay preserved the precise timing and dynamics of visual stimuli, including gabor movement corresponding to the original wheel turns, as well as auditory go cues and reward delivery. Since the timing of all stimuli were determined by the original session being replayed, reward delivery was not contingent on wheel movements. Naive mice were water restricted for one week prior to and during the replay experiment period and licked following rewards in a manner that was generally similar to trained mice (however, this and other idiosyncratic motor behaviors were performed of the subject's own volition and therefore could not be precisely matched). The identical preprocessing and context decoding analysis pipelines used for trained task sessions (see "Population decoding of task events" section below) were then applied to decode context from pre-stimulus MRF activity in the naive replay sessions.

### Behavioral Analysis

For both neural and behavioral modeling analyses, we sought to only analyze sessions wherein the subject was sufficiently tracking context to successfully remap center stimulus actions. Sessions also needed to have enough trials for adequately powered statistical analysis. Thus, session inclusion criteria based on behavioral performance were as follows:  $\geq 200$  non-repeat trials, 50-trial sliding window of performance at some point crossing 60% correct on non-repeat center stimulus trials and 80% correct on non-repeat outside stimulus trials for both Left and Right Blocks. Sessions used for behavioral modeling include those during the final stage of task training prior to Neuropixels recording as well as recording sessions. Trials with no wheel turn response were excluded from analysis. Behavioral models were trained and evaluated on both repeat and non-repeat trials. Raw behavioral performance was always evaluated only on non-repeat trials.

### Video Monitoring

For electrophysiology experiments, the subject was illuminated with infrared light (850 nm, CMVision IR30) and eye and face movements were monitored (Figure 1A). The right eye was monitored with a camera (FLIR CM3-U3-13Y3M-CS) fitted with a zoom lens (Thorlabs MVL7000) and long-pass filter (Thorlabs FEL0750), recording at 70 fps. Face movements were monitored with another camera (FLIR CM3-U3-13Y3M-CS, zoom lens Thorlabs MVL16M23, long-pass filter Thorlabs FEL0750) directed at a  $2 \times 3$  cm mirror reflecting the left side of the face, recording at 70 fps. Licks were detected from the face video by thresholding the average intensity of an ROI centered between the lips and the lick spout, calculated for every frame.

### Measurement of pupil position and face video motion energy SVDs

We measured pupil area and position (Figures 1A, 4A, and S6C) from the high-zoom videos of the subject’s eye, using DeepLabCut.<sup>107</sup> In approximately 50 training frames randomly sampled from a subset of sessions, 4 points spaced at 90° around the pupil were manually identified, and the network was trained with default parameters. Frames were then manually annotated focusing on frames with errors, and the network re-trained. The pupil area was taken to be the area of an ellipse with major and minor axis lengths given by the distances between opposite pairs of detected points. Pupil position was defined as the center point between the dorsal/ventral and medial/lateral annotated points. The traces were passed through a median filter with a filter order of 5. Two recordings were excluded from these analyses owing to video quality that was unusable ( $n=2$  out of 36). We measured video motion energy SVDs from the low-zoom videos of the frontal aspect of the subjects, which included the face, arms, and part of the torso of the mice. Motion energy SVDs were calculated using Facemap ([www.github.com/MouseLand/FaceMap](https://www.github.com/MouseLand/FaceMap)), with the top 16 principal components (PCs) used for analyses.<sup>74</sup> Pupil traces and face video SVDs were centered to have a mean of 0 and scaled to have a median absolute deviation of 1 before kernel regression fitting or decoding.

### Disentangled RNN

We implemented the DisRNN as in Miller et al.,<sup>65</sup> with minor modifications (Figures S1A–S1E). Briefly, the DisRNN is a recurrent neural network architecture designed to discover sparse and interpretable cognitive models directly from data. The DisRNN achieves this in two ways: (1) the model is trained to minimize both a supervised loss, and a complexity penalty which encourages parsimony in the learned model; (2) each element of the model’s latent state is updated through a separate feed-forward network. At each timestep the model receives as input the current latent state  $\mathbf{z}_t$  as well as the observation of the previous trial  $\mathbf{o}_t=[s_t, a_t, r_t]$  where  $s_t$ ,  $a_t$ , and  $r_t$  are the stimulus (-1, 0, 1), choice (-1, 1), and reward (0, 1) received on trial  $t$  respectively. The elements of these vectors are passed in parallel through noisy channels referred to as information bottlenecks. Given input  $x$ , the outputs of these bottlenecks are samples  $\hat{x} \sim \text{Normal}(mx, \sigma)$  where  $m$  and  $\sigma$  are learned parameters. Each bottleneck  $b$  receives a penalty for deviating from the unit Normal:

$$L_b = \sum_t D_{KL}(\mathcal{N}(m_b x_{b,t}, \sigma_b) \parallel \mathcal{N}(0, 1)) = \frac{1}{2} ((m_b x_{b,t})^2 + \sigma_b - \log \sigma_b - 1)$$

where  $D_{KL}$  is the Kullback-Leibler divergence (KLD). For a given bottleneck, this penalty is exactly 0 when  $m=0$  and  $\sigma=1$ . In that case the bottleneck’s output will be samples from the unit Gaussian, independent of the input, and no information will pass through the bottleneck. Samples from these bottlenecks are then passed through parallel multilayer perceptrons (MLP). Each MLP is responsible for updating a single element of the latent state by outputting an update value  $w_t^i$  and candidate element  $u_t^i$ . These are combined to update the latent state:

$$\begin{aligned} \begin{bmatrix} w_t^i \\ u_t^i \end{bmatrix} &= \text{MLP}^i \left( \begin{bmatrix} \hat{\mathbf{z}}_t \\ \mathbf{o}_t \end{bmatrix} \right) \\ z_{t+1}^i &= (1 - w_t^i) \hat{z}_t^i + w_t^i u_t^i \\ \mathbf{z}_{t+1} &= (1 - \mathbf{w}_t) \odot \hat{\mathbf{z}}_t + \mathbf{w}_t \odot \mathbf{u}_t \end{aligned}$$

The elements of this latent state are then passed through a second set of “latent” bottlenecks to encourage sparsity in the latents themselves. The state is then concatenated with the stimulus  $s_{t+1}$ , and passed through a final “choice” MLP which outputs the log-odds of a counter-clockwise choice on the next trial given the stimulus.

$$\log \frac{P(a_{t+1} = CCW)}{P(a_{t+1} = CW)} = \text{MLP}_{\text{choice}} \left( \begin{bmatrix} \hat{\mathbf{z}}_{t+1} \\ s_{t+1} \end{bmatrix} \right)$$

Where  $\sim$  denotes that the state has gone through the latent bottleneck.

During fitting the model weights are adjusted to minimize the negative log-likelihood (NLL) given the predicted choice probabilities and true behavior, as well as the complexity cost associated with the bottlenecks. After fitting, bottlenecks corresponding to elements of the latent state and observations which are necessary to predict future actions have parameters  $m \approx 1$  and  $\sigma < 1$ , meaning that information is allowed to pass through. During evaluation, all bottlenecks are deterministic and output  $m_b x_t$ .

We found that a 5 dimensional latent state was large enough to accommodate the solutions found across these values. The update MLPs all had 3 layers of width 5, and the choice MLP had 2 layers of width 3. The relative importance of accuracy or simplicity in the fitted model can be controlled through changing the weight  $\beta$  of the bottleneck term.

$$L = \text{NLL}(a_t, P(a_t)) + \beta \sum_{b \in \text{bottlenecks}} L_b$$

For our hyperparameter sweep, we found solutions for each subject at the following KL weights  $[0, 5 \times 10^{-4}, 1 \times 10^{-3}, 3 \times 10^{-3}, 1 \times 10^{-2}, 1]$ . We sought to optimize parsimony and performance, therefore, the model for each subject was chosen based on the highest KLD weighting before model performance dropped.

Model solutions were consistent across subjects, using a single latent variable that evolves according to a set of linear, observation-dependent updates. The linearity arises from the bottlenecks associated with  $\mathbf{o}_t$  being the only ones left open after fitting (Figure S1D). For example, if a single latent  $z^j$  was used after fitting with open update bottlenecks  $z^j, s_t, a_t, r_t$  then the update rule for  $z^j$  would depend on both its current value and the trial observation, making its set of update rules potentially non-linear and observation-dependent. However, after fitting, models consistently only had a single latent variable with open update bottlenecks  $s_t, a_t, r_t$ ; therefore,  $u_t^j$  and  $w_t^j$  are only dependent on the trial observation and each of the update rules are linear in  $z^j$ .

### Context belief model

We represent the subject's latent context-belief variable  $z_t$  (Figure 1G) as a linear function of its previous value, with parameters conditioned on the previous trial's observation  $\mathbf{o}_t$ . Each trial observation  $\mathbf{o}_t$  consists of the stimulus presented  $s_t$ , the action taken by the mouse  $a_t$ , and the reward received  $r_t$ .

$$z_{t+1} = \alpha_{\mathbf{o}_t} z_t + \beta_{\mathbf{o}_t}$$

Since  $\mathbf{o}_t$  has 8 possible values (possible combinations of stimulus, action, and reward outcome), there are 8  $\alpha$  and  $\beta$  parameters to fit. The predicted probability of taking a CCW wheel turn on trial  $t$  is given by the logistic function applied to a linear transformation of the latent variable.

$$P(a_t = \text{CCW}) = \sigma(\rho z_t + \delta)$$

Before fitting, model parameters governing the updates of the latent variable were initialized at the identity function. All model parameters were fit to minimize the NLL on center stimulus trials. In the equation below  $a_t$  is 0 for leftward wheel turns and 1 for rightward wheel turns.

$$\mathcal{L}_{\text{NLL}} = - \sum_{\text{sessions}} \sum_{t=1}^T \mathbf{1}(s_t = 0) [a_t \log P(a_t = \text{CW}) + (1 - a_t) \log P(a_t = \text{CCW})]$$

### Comparison models: Center stimulus-specific Q-Learning and Win-Stay-Lose-Switch

The simple Q-learning and win-stay-lose-switch models here were included as baselines to establish whether mice relied solely on action-reward integration from center stimulus trials; they do not characterize the full capacity of reinforcement learning (RL) approaches for this task. More complex RL agents (for example, those with expanded state representations incorporating model-based knowledge of block structure) could in principle be designed to accurately model the mice's behavior in this task. However, the design space of such models is large and requires assumptions about which task features the animal represents. We therefore used DisRNN as a discovery tool to identify the model class that best captured the mice's strategy in a data-driven manner, and compared this against the simpler baselines below.

The Q-learning agent (Figures 1I and S1B) consists of two variables  $Q(a=\text{CCW})$  and  $Q(a=\text{CW})$  which capture recent rewards associated with making counter-clockwise and clockwise wheel turns respectively, on center stimulus trials. Each Q-value is updated according to the following equation, but only when its associated action was selected:

$$Q_{t+1}(a_t) = (1 - \alpha)Q_t(a_t) + \alpha r_t \quad Q_{t+1}(a \neq a_t) = Q_t(a \neq a_t)$$

The predicted probability of taking a counterclockwise wheel turn on trial  $t$  is given by the scaled difference between the two Q-values, passed through the logistic function.

$$P(a_t = \text{CCW}) = \sigma(\beta(Q(\text{CCW}) - Q(\text{CW})))$$

Model parameters  $\alpha$  and  $\beta$  were fit to minimize the NLL on center stimulus trials. Note that by fitting exclusively to center stimulus trials, the Q-learning agent received stimulus location as an implicit input.

The win-stay-lose-switch (WSLS) strategy predicts subjects will react to an unrewarded action on a center stimulus trial by selecting the opposite action on the subsequent center stimulus trial. On rewarded center stimulus trials, the strategy predicts subjects will repeat their rewarded action. To compare WSLS performance to other models in terms of NLL, probabilities of 0.99 and 0.01 were used for the predicted action and opposing action respectively.

To cross-validate model performances for each behavioral model, for each mouse we made 5 random partitions of sessions into training, validation, and testing sets, each containing approximately 70%, 15%, and 15% of a subject's sessions respectively. For each subject, the model with the best test set performance was chosen for all downstream analyses.

### Electrophysiological recordings

During recording sessions (Figures 2D and 5B), mice were headfixed and seated on a plastic platform with forepaws on a rubber wheel. Recordings were made using Neuropixels 2.4 electrode arrays (4 shanks, 5120 sites),<sup>71</sup> which have 384 selectable recording sites. Probes were mounted to a dovetail and affixed to a steel rod held by a micromanipulator (uMP-4, Sensapex Inc). For later electrode track localization within the brain, probes were coated with a red lipophilic fluorescent dye (CM-DiI, ThermoFisher Vybrant V22888) by holding 3  $\mu\text{l}$  in a droplet on the end of a micropipette and painting the probe shanks. Probe entry coordinates and angles were selected using Pinpoint software.<sup>108</sup> Medial/lateral and anterior/posterior coordinates were zeroed at bregma, and the depth

axis was zeroed at the surface of the brain at the craniotomy. In each session, one or two probes were advanced through the surface of the brain to their final position at 4–5  $\mu\text{m/s}$ . Probes were driven 100–150  $\mu\text{m}$  past the final depth, then withdrawn the same distance, to relieve tension on the brain and reduce motion of the brain relative to the probe. Electrodes were allowed to settle for around 10 min before starting recording. Recordings were made in internal reference mode using the ‘tip’ reference site on shank 0, with a 30 kHz sampling rate. Recordings were repeated at different locations on each of multiple subsequent days, with at least 100  $\mu\text{m}$  of distance and/or 5 degrees of probe pitch between probe insertions to allow for reliable distinguishing of probe tracks post hoc. Additional craniotomies were performed as needed.

### Neural data analysis

Electrophysiological recordings were automatically spike sorted with Kilosort2.5<sup>71</sup> (<https://github.com/MouseLand/Kilosort>). Sorted units were curated using automated quality control<sup>109</sup>: exclusions were based on noise floor violations (the estimated proportion of spikes that were missed because they fell below the noise level of the recording), and refractory period violations (the estimated proportion of spikes arising from the non-primary neuron). Quality control accuracy was assessed by manually reviewing a subset of the data using the phy GUI<sup>110</sup> (<https://github.com/kwikteam/phy>). For population decoding analyses (Figures 3, S4A–S4C, and S5), we included multi-unit activity in population firing rates (only the spike floor metric and a decoding period firing rate cutoff were applied). For all analyses of neural data during task performance, repeat trials were excluded. For analyses focusing on the task trial period (i.e., after stimulus onset and before movement onset) and/or the evolution of behavioral performance (Figures 3J–3O, 4A–4F, S4A–S4M, S5A–S5G, S5L–S5Q, S7A–S7G, and S8A–S8C), we filtered out trials with reaction times less than 75 ms and greater than 400 ms so as to exclude impulsive “flinches” and responses reflecting disengagement or delayed processing.

For visualizing neuron firing rates in peri-event time histograms (PETHs) aligned to task events (Figures 2F–2I), neural activity was binned at 0.015 s and smoothed with a causal half-Gaussian filter with standard deviation 0.05 s. For the kernel regression analysis (Figures 4 and S7), each neuron’s activity was binned at 0.01 s, smoothed with a causal half-Gaussian filter with standard deviation 0.025 s, and baseline soft normalized by dividing the trial spike rates by the baseline firing rate (–300 to 0 ms before stimulus onset) + 1 spikes/s. PETHs constructed for visualization included correct trials only, but both correct and error trials were included in the kernel regression analysis.

For the kernel regression analysis (Figures 4 and S7), to be discussed below, neurons were included in the analysis if their firing rate was significantly modulated during the task (Figure 2D, “task-resp.”). Units were first excluded if their trial firing rate (analysis window between stimulus and movement onsets) was less than 0.1 spikes/s (Figure 2D, “FR pass”). Then, a set of six statistical tests was used to detect changes in activity during various task epochs and conditions: (1) Wilcoxon signed-rank test between trial firing rate and pre-stimulus rate (defined in period –0.3 to 0 s relative to stimulus onset on each trial); (2) signed-rank test between stimulus-driven rate (firing rate between 0 and 0.125 s after stimulus onset) and pre-stimulus rate; (3) signed-rank test between pre-movement rates (–0.1 to 0 s relative to movement onset) and pre-stimulus rate; (4) Wilcoxon rank-sum test between pre-movement rates on CW choice trials and those on CCW choice trials; (5) rank-sum test between trial firing rates on Left Blocks vs. Right Blocks; (6) Kruskal-Wallis test between left, center, and right stimulus-driven firing rates. A neuron was considered active during the task, or to have detectable modulation during some part of the task, if any of the p values on these tests were below a Bonferroni-corrected alpha value (0.15/6=0.025). The lenient alpha=0.15 threshold was chosen due to the coarse nature of the tests which make them relatively insensitive to transient activity and thus increasing the likelihood of false negatives.

### Population decoding of task events

We used cross-validated logistic regression to decode stimuli, all choices, and center stimulus choices from trial firing rates of neural populations (Figures S4A–S4C). Units were included as predictors if they passed the noise floor quality metric (described above) and had an average trial firing rate  $\geq 0.2$  spikes/s. The refractory period metric was not used as multi-unit activity is acceptable for population decoding. Sessions were included in the decoding analysis if  $\geq 3$  units fell in the region of interest and passed the above metrics. For choice decodings, a sliding 100 ms window from –250 ms to 0 ms in 25 ms increments aligned to movement onset was used to bin neural firing rates. For stimulus decoding, a sliding 50 ms window from –50 ms to 150 ms in 25 ms increments aligned to stimulus onset was used to bin neural firing rates. Logistic regression was run separately on each time bin to get the time course of decodability of stimulus and choice. Elastic net regularized (alpha=0.5) five-fold nested cross-validated logistic regression was used, wherein trials were split into train and test sets, and within the train set, the optimal lambda value was picked using five-fold cross validation (using the package `cvglmnet` for MATLAB [https://hastie.su.domains/glmnet\\_matlab/index.html](https://hastie.su.domains/glmnet_matlab/index.html)). Performance of the decoding was measured as the balanced accuracy of predicted vs. actual stimulus or choice labels for a given session and region. To account for slight instability in the decoding performance due to random cross-validation partitions, the entire decoding procedure for each variable was run 10 times, and mean balanced accuracies were taken across the 10 decodings at each time bin.

To decode pre-stimulus context (Figure 3I), we used a similar procedure as the stimulus and choice decoding, except that linear regression was used to predict the continuous latent context variable, and only the baseline period before stimulus onset was considered. Thus, units with an average pre-stimulus (300 to 50 ms before stimulus onset) firing rate  $\geq 0.05$  spikes/s and passing the noise floor metric were included, and only one time bin (300 to 50 ms before stimulus onset) was used for decoding. Sessions were included in the decoding analysis if  $\geq 3$  units fell in the region of interest and passed the above metrics. Performance of the decoding was measured as the Pearson correlation coefficient ( $r$ ) between the predicted and actual context variable for a given session

and region.  $r$  was used rather than the coefficient of determination ( $R^2$ ) to assess the correspondence in temporal structure between predicted and actual traces, independent of scaling or offset differences potentially induced by drift and/or regularization. The entire decoding procedure was run 100 times, and the mean balanced accuracy was taken across the 100 decodings.

To obtain null distributions for decoding performance, the decoders were refit to 2000 shuffles for the pre-stimulus context variable or 250 shuffles of choice and stimulus variables. Decoding scores were then defined as the difference between the true decoding performance ( $r$  for continuous latent context variable; balanced accuracy for categorical choice and stimulus variables) relative to the median of its null distribution. Shuffles were performed for stimulus, choice, and context variables as described in the “Kernel regression” (stimuli and all choices) and “Pseudosessions” (pre-stimulus context and center stimulus choices) sections below.

To determine whether brain region significantly predicted pre-stimulus context decoding at the neural population level, we fit a linear mixed-effects model with region as a fixed effect, random intercepts for subject and for session nested within subject, and context decoding score as the response variable. An analysis of variance (ANOVA) testing for fixed effects revealed a significant effect of region ( $F(7,77)=2.6$ ,  $p=0.018$ ). Pre-stimulus context decoding performance was not dependent on the number of neurons included in the decoder (Figure S6A), behavioral performance on the decoded session (Figure S6B), nor uninstructed movements/pupil position during the pre-stimulus period (Figures S6C–S6E, described below).

### State space trajectory analyses

To establish how neural populations may solve the stimulus-action remapping problem posed by the task, we projected neural population activity onto dimensions that encode Action (movement initiation), Choice (turn direction), and Context (Figures 3A, 3B, 3K, 3L, S5A–S5F, S5M, and S5N). This approach allowed us to visualize the evolution of neural population activity across different trial types in a low-dimensional state space to understand how populations transition between actions in different contexts in response to identical sensory stimuli.

We defined neural populations based on neuron locations within anatomical regions aligned to the Allen CCF atlas, including the eight regions we specifically targeted in our electrophysiological recordings: MRF, SCm, APN, RN, MOs, MOp, CP, and VPM/VPL. For each region, to ensure representations of population trajectories were adequately powered, we only included sessions with at least 10 neurons meeting our quality control criteria. The criteria selected for neurons that had sufficient firing rates before stimulus onset and/or within a trial, as well as filtering out neurons with excess missing spikes (pre-stimulus or trial firing rate exceed 0.05 spikes/s; noise floor metric passed).

We first defined three dimensions using binned population firing rates as predictors and logistic regression to classify CW actions from baseline (i.e., the pre-stimulus period), CCW actions from baseline, and baseline activity between Left and Right Blocks (Context dimension). Specifically, the CW and CCW dimension decoders used pre-stimulus neural activity (-0.3 to -0.05 s relative to stimulus onset) and activity prior to a CW or CCW turn (stimulus onset to movement onset) to decode whether trials belonged to the baseline or pre-movement periods. The Context dimension decoder only used pre-stimulus activity (-0.3 to -0.05 s relative stimulus onset) to decode whether trials belonged to the Left or Right Block. To ensure the neural trajectories represented successful implementation of task rules, we included only trials where the mouse made the correct choice in the decoders and subsequent projection onto decoding dimensions.

For each dimension, we applied L2-regularized logistic regression using a nested cross-validation approach. The outer loop performed 25-fold cross-validation across trials to define dimensions on each training set and project only the held-out test data. Within each of these folds, we used a 5-fold inner cross-validation loop to select the optimal regularization parameter ( $\lambda$ ). To address any class imbalances, we used inverse class frequency weighting on each outer loop iteration.

We then defined the Action dimension as the sum of the weights  $\beta_{CW} + \beta_{CCW}$  and the Choice dimension as the difference  $\beta_{CW} - \beta_{CCW}$ . This rotation aligns the subspace such that the Action axis captures the shared variance associated with movement initiation, while the Choice axis captures the direction-specific variance. Since regression coefficients can have arbitrary sign flips between sessions, we standardized the orientation of dimensions by applying sign corrections when needed. Specifically, based on train set projections, we used the convention that pre-movement activity was greater than pre-stimulus activity along the Action dimension, and that Right Block activity exceeded Left Block activity along the Context dimension.

To assess the alignment of the Context dimension with the choice-related dimensions (Figures 3J and S5L), we calculated the angle between the Context and CW dimensions on each session. For comparison, we also calculated the angle between the CW and CCW dimensions on the same sessions (paired comparison using signed-rank test). Angles were calculated using the formula

$$\theta = \arccos\left(\frac{\mathbf{u} \cdot \mathbf{v}}{|\mathbf{u}| |\mathbf{v}|}\right),$$

where  $\mathbf{u}$  and  $\mathbf{v}$  are the vectors representing each dimension and  $\theta$  is the angle between them.

For trajectory visualizations and subsequent trajectory-based analyses, we implemented additional inclusion criteria based on decoding performance. Sessions were only included if they met minimum decoding accuracy thresholds: action decoding (classifying pre-stimulus from pre-movement trial periods irrespective of choice) cross-validated accuracy  $\geq 70\%$  and choice decoding (classifying CW from CCW choices based on pre-movement activity) accuracy  $\geq 60\%$  ( $n=21$ , 9, 3, 4, 12, 5, 5, 2 sessions in MRF, SCm, APN, RN, MOs, MOp, CP, VPM/VPL). These thresholds helped ensure that included sessions were properly representative of how choice dynamics unfolded in a given region.

After defining Action, Choice, and Context dimensions for each session, we calculated population trajectories by binning spike data using 75 ms sliding windows with 10 ms step sizes aligned to both stimulus onset (0 to 250 ms) and movement onset (-250

to 0 ms) and projecting the binned activity onto each dimension. To avoid overfitting, we only projected activity on held-out trials from our cross-validation procedure, ensuring all trials were represented through the k-fold partitioning described above.

To combine trajectories across sessions for each brain region, we first applied robust z-scoring (subtracting the median and dividing by  $1.4826 \times$  the median absolute deviation) to each session's trajectories independently across all trials and timepoints, bringing them into comparable units. We then calculated the median value of each of the four trial types for each time bin across all trials and sessions to obtain average trajectories for each brain region.

To test different hypothesized mechanisms for stimulus-action remapping, we quantified two key features of these trajectories:

1. Pre-stimulus positions in choice subspace (Figure 3F): We calculated the position along the Choice dimension at baseline ( $[-75\ 0]$  ms bin before stimulus onset) for Left and Right Blocks separately. This allowed us to test whether there was a context-dependent motor bias in the population state prior to stimulus presentation. Differences between blocks were calculated per session, and signed-rank tests were used to test for significance in each region.
2. Sensory-evoked trajectory angles (Figure 3G): We measured the angle formed between trajectory positions at baseline (median position amongst 0-20 ms bins relative to stimulus onset) and post-stimulus response (median position amongst 70-150 ms bins relative to stimulus onset) in the Action-Choice plane using the four-quadrant inverse tangent ("atan2" in Matlab). Angles were calculated for individual trials, and then the median angle per trial type was computed across trials. We used rank-sum tests for each region to determine whether the initial trajectories of center stimuli followed different angles depending on the context.

To assess the possibility that the pre-stimulus context signal reflected a decaying sensory or motor response from the previous trial, we projected population activity onto the Context dimension throughout the pre-stimulus window ( $-0.5$  to  $-0.05$  s relative to stimulus onset). The signals on each session were normalized by robust z-scoring, then flipping the sign of one block's context values such that the context representations from both blocks could be combined, with a positive value representing activity that is more in the direction of the current block's context. Activity along the context dimension remained stable across this epoch in MRF, MOs, SCm, and CP (Figure S5O), consistent with a maintained internal state rather than a lingering sensory or motor response. Note that relative to the pseudosession-controlled analysis in Figure 3I, in which MRF, SCm, MOs, and CP had similar magnitude of context representation, in Figure S5O the MOs and CP appear to have a larger signal. The pseudosession analysis acts as a kind of signal-to-noise ratio calculation, so this apparent difference presumably arises because the pseudosession control produces a wider distribution in MOs and CP than in MRF and SCm, which is consistent with the visibly larger variance (i.e. noise) across sessions and across time in this plot. Larger signals in pseudosession distributions can result from contributions of artifactual factors such as recording drift, so the pseudosession-controlled analysis in Figure 3I is the more interpretable measure of signal magnitude.

To further confirm that the context signal was distinct from task-related sensory or premotor representations, we calculated angles between the Context dimension and a Stimulus dimension separating left vs. right outside stimulus responses (0–150 ms post-stimulus) as well as a Choice dimension separating CW vs. CCW wheel turns (75 ms pre-movement onset). The Context dimension was nearly orthogonal to both the Stimulus and Choice dimensions in MRF, MOs, SCm, and CP (Figures S5P and S5Q), whereas the Stimulus and Choice dimensions exhibited significant overlap in comparison (potentially consistent with a sensorimotor transformation performed within these populations of neurons). Thus, the neural subspace encoding pre-stimulus context is separable from those representing sensory inputs and motor commands.

#### **Trial-by-trial dynamics of Context dimension activity and prediction of center stimulus choices**

To investigate the relationship between the pre-stimulus context signal and behavior, we analyzed the trial-by-trial projection of neural activity onto the Context dimension (as defined in the "Population decoding of task events" section above). For each session, this continuous signal was first smoothed using a causal Gaussian filter ( $\sigma=3$  trials). Then, to combine data across sessions for regional analyses, the signal from each session was normalized by first subtracting its median value and then linearly scaling the result such that the full range of the original data mapped to the interval  $[-1, 1]$ . For the visualization of the example decoded context trace in Figure 3M, this normalization was applied within a sliding window of 150 trials.

To compare the dynamics of this neural context signal with behavioral adaptation after a block switch (Figures 3N and S4D–S4G), we aligned trials to the first trial of each new block. For each session, we combined data across blocks by inverting the sign of the Context dimension projection values on Left Block trials. We then computed the average projection value for each trial number after the switch, across all included sessions for a given brain region. An exponential rise function of the form  $f(t) = A - (A - O) \cdot e^{-\frac{t-1}{\tau}}$ , where  $t$  is the trial number, was fit to these averaged data to extract the time constant  $\tau$ . The initial offset ( $O$ ) was fixed to the value at trial 1, and the asymptote ( $A$ ) was fixed to the mean value of the last 10 trials of the analysis window (trials 26–35 post-switch), leaving  $\tau$  as the sole free parameter. This function was implemented in MATLAB using the 'fit' function with the nonlinear least squares fitting method, where  $O$  and  $A$  were passed as fixed 'problem' parameters. A similar procedure was used to fit the behavioral data, using the percentage of correct choices on center stimulus trials. We computed 95% confidence intervals for  $\tau$  using a bootstrap procedure ( $n=1000$  resamples with replacement of trials).

To test whether the pre-stimulus context signal was predictive of choices on individual center stimulus trials (Figures 3O and S4H–S4M), we separated trials by block (Left vs. Right) and outcome (correct vs. error). For each session, we calculated the mean projection value onto the Context dimension for each of the four resulting groups. To quantify the relationship between the context

signal and choice, we calculated the percentage shift of the context signal on error trials toward the mean value associated with the opposite block's correct trials. A signed-rank test was used to determine if this percentage shift across sessions was significantly different from zero. This analysis was also performed on a subset of trials occurring 15 or more trials after a block switch to assess whether the context signal predicted behavioral lapses outside the initial rule-updating period.

### Task alignment of spontaneous movements and pupil dynamics

A potential confound in the neural coding of pre-stimulus context is context-dependent uninstructed movements or pupil position changes. These could drive neural activity that is correlated with context rather than the activity representing the context variable itself. In particular, if mice systematically shifted their pupil position before stimulus onset based on the current block, this could alter the retinal input for identical stimulus positions across blocks, potentially providing an alternative mechanism to solve the task (but note that we did not detect differences in initial stimulus responses between blocks in the choice subspace, [Figure 3G](#), nor in individual neuronal responses, [Figures 4F and S7C](#)).

We used a sensitive approach to detect any context-dependent uninstructed movements or pupil position changes by decoding context from face video motion energy SVDs and pupil position traces ([Figure S6C](#)). The predictor matrix was composed of the mean values during the pre-stimulus period (-300 to -50 ms relative to stimulus onset) on each trial of the top 16 principal components of face motion SVDs (capturing facial, whisker, front of body, and forelimb/wheel movements) and pupil x and y positions. This predictor matrix was regressed against the context variable using ordinary least squares (OLS) regression.

Performance was calculated on held-out trials using 5-fold cross validation. To account for variability in random data partitions on the results, the entire decoding procedure was run 25 separate times per session, and the final performance was calculated as the mean score across all 25 runs.

To assess statistical significance, null distributions were formed by running the same decoding procedure 1000 times using context variables from pseudosessions. Context could be significantly decoded ( $p < 0.05$ ) from the predictor matrix of uninstructed movements and pupil positions in 8 out of 35 sessions. To compare decoding performances from uninstructed movements with those from neural activity ([Figures S6D and S6E](#)), if multiple neural populations were decoded in a session (from different regions), the population with the lowest p-value and highest score were used to represent a given session. There was no significant correlation between the decoding score ( $\Delta r$  relative to null) for context predicted from uninstructed behaviors versus context predicted from neural activity ([Figure S6D](#),  $r = -0.08$ ,  $p = 0.66$ ). Furthermore, only 2 out of the 15 sessions with significant neural decoding of context showed significant uninstructed behavioral decoding ([Figure S6E](#)).

### Kernel regression model for single neuron encoding of task variables

To identify and disentangle the encoding of choice, stimulus, and contextual modulation of trial responses in single neuron firing rates ([Figures 4 and S7](#)), we fit a kernel regression model.<sup>9,80</sup> The firing rate of each neuron in this analysis is described as a linear sum of temporal filters aligned to task events. This analysis was restricted to the time periods stimulus onset and movement onset on each trial. The predicted firing rate  $r_n(t, i)$  for neuron  $n$  on timepoint  $t$  and trial  $i$  is given as:

$$\begin{aligned}
 r_n(t, i) = & K_{n, c_i}^s(t - S_i) \\
 & + Z_i \cdot K_n^{\bar{s}}(t - S_i) \cdot \mathbb{1}(c_i = 0) \\
 & + K_n^m(t - M_i) \\
 & + D_i \cdot K_n^d(t - M_i) \\
 & + Z_i \cdot K_n^{\bar{m}}(t - M_i) \\
 & + D_i \cdot Z_i \cdot K_n^x(t - M_i) \\
 & + \sum_{e=1}^2 \alpha_{e, n} \cdot E_e(t) \\
 & + \sum_{f=1}^{16} \beta_{f, n} \cdot F_f(t) \\
 & + \gamma_n \cdot Z_i \\
 & + \delta_n \cdot t
 \end{aligned}$$

- $K_{n, c_i}^s$  represents the Stimulus kernel function of stimulus  $c$  on trial  $i$  for neuron  $n$ .  $c$  represents the three visual stimulus types (left, center, right).  $S_i$  is the time of stimulus onset on trial  $i$ .
- $K_n^{\bar{s}}$  represents the Context-Stim kernel for neuron  $n$ . This indicator function  $\mathbb{1}(c_i = 0)$  takes a value of 1 only when  $c_i = 0$ , ensuring that this kernel only applies on center stimulus trials.
- $K_n^m$  represents the Action kernel for neuron  $n$ .  $M_i$  represents the time of movement onset on trial  $i$ .
- $K_n^d$  represents the Choice kernel for neuron  $n$ .  $D_i$  represents the direction of wheel movement (-1 for CCW, +1 for CW) on trial  $i$ .

- $K_n^m$  represents the Context-Premotor kernel for neuron  $n$ .  $Z_i$  represents the value of the latent context variable (Figure 1G) on trial  $i$ .
- $K_n^x$  represents the interaction kernel between the Context-Premotor and Choice kernels.
- $\alpha_{e,n} \cdot E_e$  represents the Pupil predictor for pupil position axis  $e$  (2 axes: dorsal/ventral and medial/lateral) and neuron  $n$ . The variable  $E$  represents the pupil position, and  $\alpha$  is its corresponding fitted parameter.
- $\beta_{f,n} \cdot F_f$  represents the Face predictor for face motion SVD component  $f$  (top 16 SVD components used) and neuron  $n$ . The variable  $F$  represents the face motion SVD, and  $\beta$  is its corresponding fitted parameter.
- The Pupil and Face predictors were not lagged, which has been found to maximize variance explained in neural activity.<sup>74</sup>
- $\gamma_n \cdot Z_i$  represents the Context-Baseline predictor for neuron  $n$  on trial  $i$ .
- $\delta_n \cdot t$  represents the Time predictor for neuron  $n$  at timepoint  $t$ .

To fit each kernel ( $K^s$ ,  $K^{\tilde{s}}$ ,  $K^m$ ,  $K^d$ ,  $K^{\tilde{m}}$ ,  $K^x$ ), we constructed a Toeplitz predictor matrix of size  $T \times l$ , in which  $T$  is the total number of time bins and  $l$  is the number of lags required for the kernel. The predictor matrix contains diagonal stripes starting each time an event occurs and 0 otherwise (Figure 4A). Stimulus-aligned kernels were supported over the window 0 to 0.4 s relative to stimulus onset, and movement-aligned kernels were supported over the window -0.25 to 0 s relative to movement onset. These predictor matrices along with the Pupil, Face, Context-Baseline, and Time predictors were horizontally concatenated to yield a global prediction matrix  $\mathbf{P}$  of size  $T \times L$  containing all predictor kernels. Rate vectors of all  $N$  neurons in a given session were horizontally concatenated to form  $\mathbf{R}$ , a  $T \times N$  matrix. If kernels extended before stimulus onset (relevant for trials with reaction times less than 0.25 s) or after movement onset (relevant for trials with reaction times less than 0.4 s) on a given trial, these rows were removed from  $\mathbf{P}$  and  $\mathbf{R}$ .

To prevent noisy and overfit kernels, we implemented reduced-rank regression,<sup>9,111</sup> which allows regularized estimation by factoring the kernel matrix  $\mathbf{K}$  into the product of a  $L \times j$  matrix  $\mathbf{B}$  and a  $j \times N$  matrix  $\mathbf{W}$ , minimizing the total error:  $E = \|\mathbf{R} - \mathbf{PBW}\|^2$ . The  $T \times j$  matrix  $\mathbf{PB}$  consists of a set of ordered temporal basis functions that can be linearly combined to estimate the neuron's firing rate over the whole training set and which results in the best possible prediction from any rank  $j$  matrix. To estimate each neuron's kernel functions, we generated the reduced rank predictor matrix  $\mathbf{PB}$  for  $j=20$  and estimated the weights  $\mathbf{w}_n$  to minimize the squared error  $E_n = \|\mathbf{r}_n - \mathbf{PBw}_n\|^2$  with elastic net regularization (using the package `cvglmnet` for MATLAB [https://hastie.su.domains/glmnet\\_matlab/index.html](https://hastie.su.domains/glmnet_matlab/index.html)) with parameters  $\alpha=0$  and using ten-fold cross-validation to determine the optimal value for  $\lambda$  for each neuron. The kernel functions for neuron  $n$  could then be unpacked from the  $L$ -length vector obtained by multiplying the first  $j=20$  columns of  $\mathbf{B}$  by  $\mathbf{w}_n$  (i.e.  $\mathbf{k}_n = \mathbf{B}_{1:L, 1:20} \mathbf{w}_n$ ).

To determine whether a given kernel significantly contributed to a neuron's firing activity, we generated empirical null distributions for each neuron using a shuffling approach. The kernel matrix was refit while shuffling either the Stimulus, Choice, or Context kernels. The Stimulus kernel was shuffled by randomly permuting the stimulus position labels of  $c$  within their respective blocks (i.e., center and Left stimuli shuffled with Left Blocks, center and right stimuli shuffled within Right Blocks). The Choice kernel was shuffled by permuting contiguous runs of identical choices (e.g., a set of three consecutive CW turns constitutes one such run) in  $D$ , thereby preserving the local temporal statistics of choices while disrupting their overall sequence across the session. The Context kernels were shuffled using the context representations generated from pseudosessions (described in detail in the Pseudosessions section below) in place of the true context trace in  $Z$ . Shuffles were performed 500 times, with only one kernel being shuffled at a time in each iteration. The Context-Stim and Context-premotor kernels were shuffled both together and separately to assess the total effect of contextual modulation on trial firing rates as well as its effects on stimulus and premotor-related activity specifically. Thus, a total of  $500 \times 5 = 2500$  shuffled kernel matrices were fit per neuron. We used the same five-fold cross-validation (splits performed across trials) for both the true and shuffled kernel regressions, and the cross-validated variance explained (coefficient of determination) was used as the test statistic (p-values and effect sizes were derived as described in the "Significance testing with empirical null distributions" section).

### Quantification of stimulus encoding during passive task presentation

To robustly assess the encoding of passively presented visual and auditory stimuli (Figures 5 and S9), we used the Zenith of Event-based Time-locked Anomalies (ZETA) test<sup>82</sup> for each neuron and stimulus type (Matlab code, <https://github.com/JorritMontijn/ZETA>). The ZETA test is parameter- and binning-free, generating a cumulative distribution of spike times relative to stimulus onset as a function of time and measuring deviation from a linear baseline. The ZETA test was applied using a 1 second window after stimulus onset, with 60 trials per stimulus type (for visual stimuli, each stimulus azimuth was considered separately, and each of 3 contrasts were considered together (0.25, 0.5, 1)). P values were obtained by comparing to a null distribution generated by jittering stimulus-onset times ( $n=500$  jitters). To broadly assess whether a neuron encoded visual stimuli irrespective of azimuths (Figure 5D), we controlled the false discovery rate (FDR) by applying the Benjamini-Hochberg procedure to the p-values from each azimuth. We then used the minimum FDR value among azimuths to determine stimulus encoding.

### Drift-diffusion modeling of stimulus-action remapping mechanisms

To examine whether each of our three hypothesized mechanisms (context-dependent shift in action dynamics, pre-stimulus motor bias, and shifting center stimulus input angle) could theoretically solve the stimulus-action remapping problem posed by the task, we instantiated each mechanism using drift-diffusion models (DDMs)<sup>112</sup> (Figures S5H-S5J).

We implemented a stochastic DDM for a one-dimensional decision variable representing position in state space along a CCW to CW axis. For each mechanism, we simulated 1000 trials with simulation parameters chosen to qualitatively match observed neural dynamics. Simulations were run for 200 ms using a timestep (dt) of 0.01 s. The decision variable at each timestep was updated according to:

$$v(t + 1) = v(t) + \eta(t + 1) + \alpha(v(t) - \theta)$$

where  $\eta$  represents Gaussian random noise (standard deviation=0.1),  $\theta$  is a mechanism-specific threshold parameter that determines the directional bias to the accumulation process, and  $\alpha$  (0.02) scales the influence of the current state on future updates.

The initial state of the decision variable for each simulation was drawn from a normal distribution with mechanism-specific mean starting value and standard deviation of 0.3. At the first timestep after  $t=0$ , a stimulus-specific input was applied, representing the sensory input to the system.

For the biasing of action selection dynamics mechanism (Figures 3E and S5H–S5J; Mechanism 3), we kept the starting value constant at 0 for all trial types, applied stimulus-specific inputs ([1 0 0 -1] for left, center Left Block, center Right Block, and right stimuli respectively), and set thresholds  $\theta$  to induce opposite dynamics between blocks ([1 1 -1 -1]; dynamics biased toward CCW action in Left Block and CW in Right Block). For the biasing of pre-stimulus position mechanism (Figures 3C and S5H–S5J; Mechanism 1), we varied the starting values ([-0.3 -0.3 0.3 0.3]) to bias initial positions toward CCW in Left Blocks and CW in Right Blocks, while keeping the threshold parameter at 0 and using the same stimulus-specific inputs. For the biasing of sensory-driven stimulus input mechanism (Figures 3D and S5H–S5J; Mechanism 2), we kept starting values and thresholds at 0 for all trial types, but modified the stimulus inputs ([0.75 -0.375 0.375 -0.75]) to induce trajectories toward CCW or CW actions for center stimuli in left or Right Blocks.

### Analyses of cortical terminal and context coding spatial distributions in MRF

#### Preprocessing of single neuron morphology open datasets

We integrated three open datasets of fully reconstructed single neuron morphologies<sup>40–42</sup> (example MOs neurons in Figure 6A), using custom analysis code in Matlab to analyze the files in a SWC format (<https://swc-specification.readthedocs.io/en/latest/swc.html>). The datasets primarily sampled cortical neurons, therefore our analyses focused solely on these cells. We only considered soma and axon terminal locations. To determine whether a point was an axon terminal, we filtered for points whose type was axon (column 2 “Type” == 2) and which weren’t a “parent” to any other sample (no entry in column 7 “Parent”, i.e. no other points branched from this point). As all coordinates were aligned to the CCF, we downsampled them to 10  $\mu\text{m}$  resolution and filtered for cortical neurons from regions with  $\geq 10$  cells projecting to the CCF-defined midbrain reticular nucleus. Neurons originated from prelimbic, infralimbic, anterior cingulate, orbital, primary and secondary motor, retrosplenial, primary and supplemental somatosensory, agranular insular, and frontal pole regions of the cortex (Figures 6B and S10A–S10C).

#### Non-uniformity testing of cortical terminals

To test whether cortical axon terminals exhibited non-uniform distribution within the MRF (Figures 6D, S10D, and S11B), we developed a statistical approach comparing observed terminal density distributions against a uniform null hypothesis. We first identified all axon terminals within the MRF from cortically-projecting neurons. For each spatial dimension (A/P, D/V, and M/L), we created histograms of terminal density in MRF slices using a sliding window with bin width of 250  $\mu\text{m}$  and step size of 150  $\mu\text{m}$ . Terminal density was calculated as the number of terminals within each slice divided by the MRF volume (in  $\text{mm}^3$ ) within that slice.

To test distance from a uniform distribution in each spatial dimension, we used a metric which computed the sum of squared deviations of the observed number of terminals in each slice from the expected number of terminals based on a uniform density. The expected density based on a uniform distribution was calculated as the total number of terminals divided by total MRF volume. This density was then multiplied by the volume of a given slice to get the expected terminal count. The advantage of calculating deviations between observed and expected counts versus observed and expected densities is that windows with higher volume are given more weight in the calculation (this especially important for more uneven sampling of recording probe locations, as described next for pre-stimulus context coding spatial distribution).

To assess the probability that the observed distances from uniform arose by chance, we formed a null distribution of randomly positioned cortical axon terminals in MRF. For each of 10000 permutations, we randomly sampled positions from all possible MRF voxels to create a surrogate terminal distribution with the same number of terminals as in the original dataset. For each permutation, we calculated terminal densities and distances from uniform using the same sliding window and sum of squared deviation approaches. This provided an empirical null distribution of deviation scores. This test was performed both for the complete terminal dataset and for a subsampled version containing one randomly selected terminal per neuron to account for potential bias from the non-uniform branching patterns of neurons with many MRF terminals. In each case, cortical terminal distributions were significantly non-uniform across all three spatial dimensions (all terminals:  $p=1 \times 10^{-4}$ ,  $1 \times 10^{-4}$ ,  $1 \times 10^{-4}$ ; 1 terminal per neuron:  $p=1 \times 10^{-4}$ ,  $1 \times 10^{-4}$ ,  $2 \times 10^{-3}$  in A/P, D/V, and M/L planes).

#### Non-uniformity testing of pre-stimulus context coding

To determine whether neurons encoding pre-stimulus context showed spatial organization within the MRF (Figures 6E and S10E), we performed non-uniformity testing similar to that applied to cortical terminals. Neurons were classified as encoding pre-stimulus context if they exhibited significant correlation ( $p < 0.05$  against pseudosession null) between their pre-stimulus firing rate and task

context (Figures 3H and S5K). The spatial distribution of these context-encoding neurons was assessed using the same sliding window approach (bin width of 250  $\mu\text{m}$ , step size of 150  $\mu\text{m}$ ) applied separately along anterior-posterior, dorsal-ventral, and medial-lateral dimensions.

For each MRF slice, we calculated the proportion of context-encoding neurons relative to all recorded neurons. The distance from a uniform distribution was again computed as the sum of squared deviations between the observed and expected counts of coding neurons in a slice. The expected uniform density was the overall fraction of neurons in MRF encoding pre-stimulus context, and the expected count was this density multiplied by the number of recorded neurons in a slice. To assess the probability that the observed distances from uniform arose by chance, we performed 10000 permutations. In each permutation, we randomly reassigned the “context-encoding” label among all recorded neurons while maintaining their spatial positions, then recalculated the proportion of context-encoding neurons and distances from uniform in each slice.

#### **Voxelization and spatial alignment of cortical terminals and pre-stimulus context coding**

To directly compare the spatial distributions of cortical terminals and pre-stimulus context-encoding neurons (Figure 6G), we implemented a voxelization approach that partitioned the MRF into overlapping spherical voxels (radius 125  $\mu\text{m}$ , grid spacing 150  $\mu\text{m}$ ). These voxels were composed of varying numbers of CCF-defined subvoxels (10  $\mu\text{m}^3$ ) depending on MRF topography around the voxel. To avoid edge artifacts, we excluded volumes containing fewer than 50 CCF-defined MRF subvoxels.

For each voxel, we calculated: (1) density of cortical axon terminals (terminals per  $\text{mm}^3$ ), and (2) proportion of neurons encoding pre-stimulus context. Voxels with no recorded neurons in them were excluded. As before, we calculated global expected terminal densities by dividing the total number of cortical terminals in MRF by total MRF volume. We calculated global expected coding proportions by dividing the total number of coding neurons by the total recorded neuron count in MRF. Voxel-wise expected terminal and coding neuron counts were then calculated by multiplying the expected densities by the volume or total number of recorded neurons within each voxel. We then computed the distance from uniform for each voxel using the sum of squared deviation of observed counts from expected counts as described above.

To assess the degree of spatial alignment between cortical terminals and pre-stimulus context coding, we computed Spearman’s rank correlation coefficient ( $\rho$ ) between the distances from uniform of terminals and context-coding neurons across voxels. Statistical significance was determined using the same permutation approach for context coding neurons as described above, wherein the “context-encoding” label was randomly reassigned amongst all recorded neurons 1000 times. Voxel context coding proportions, deviations from uniform, and Spearman’s  $\rho$  were recalculated for each permutation to form a null of  $\rho$  values. This approach preserved the spatial autocorrelation structure of the recorded neurons while testing the specific relationship between terminals and context coding.

To verify that the spatial alignment of cortical axon terminals and pre-stimulus context coding units in MRF did not arise from a preferential recording of cortical axons themselves, rather than MRF somas, we calculated the “spatial footprint” of each MRF unit’s waveform across recording sites. Small spatial footprints (<20  $\mu\text{m}$ ) have been shown to be reliable electrophysiological markers of axonal recordings.<sup>113</sup> We found no significant differences in the distributions of footprints from all units in MRF versus those that coded for pre-stimulus context (median footprint size 59  $\mu\text{m}$ , 57  $\mu\text{m}$ ;  $p=0.16$  t-test), with equally low proportions of small footprint units in each group (1.3%, 1.0%). Thus, the large majority of recorded units in MRF, including those coding for pre-stimulus context, presumably arose from somatic, rather than axonal, activity.

## **QUANTIFICATION AND STATISTICAL ANALYSIS**

The sample sizes for task recordings ( $n=5$  mice;  $n=36$  recording sessions;  $n=2444$  task-responsive, quality-controlled neurons falling within the 8 target regions; Figure 2D), task-naive passive stimulus mapping ( $n=6$  mice,  $n=21$  sessions,  $n=1681$  neurons; Figure 5B), trained passive stimulus mapping ( $n=6$  mice,  $n=40$  sessions,  $n=4988$  neurons; Figure 5B), and task-naive replay experiments ( $n=3$  mice,  $n=6$  sessions,  $n=258$  quality-controlled neurons falling within MRF; Figure S6F) were not determined with a power analysis.

In all figures, \*  $p<0.05$ ; \*\*  $p<0.01$ , \*\*\*  $p<0.001$ . Non-parametric tests were used unless otherwise stated (Wilcoxon rank-sum or Kruskal-Wallis test for unpaired data; Wilcoxon signed-rank or Friedman test for paired data). All statistical tests were two-sided. To calculate 95% confidence intervals, we performed bootstrap resampling 1,000 times with an alpha of 0.05 using the ‘bootci’ function in MATLAB (bias-corrected and accelerated percentile method), unless otherwise stated. An “above chance” effect is defined here as having its confidence interval bounds at a given alpha lie outside of the false discovery rate for a given statistical test. For comparing regions in single neuron-based analyses, each neuron was considered a replicate (Figures 4, 5, S5K, and S7–S9). For comparing regions in the neural population-level analyses in Figures 3F, 3I, 3J, S4A–S4C, S5L, and S5O–S5Q, each included session was considered a replicate. For comparing regions in the neural population-level analyses in Figures 3A, 3B, 3G, 3K–3L, 3N, 3O, S4D–S4M, S5A–S5F, S5M, and S5N, sessions were compiled by region, and each trial was considered a replicate.

#### **Significance testing with empirical null distributions**

For all analyses in which empirical null distributions were generated (Population decoding: Figures 3I, S4A–S4C, and S6; Kernel regression: Figures 4, S7, and S8; Spatial distributions of context coding and axon terminals in MRF: Figures 6D–6G, S10D, and

S10E), p-values were defined as the percentile of the actual test statistic  $T(X)$  within the null ensemble  $\{T(X')\}$  such that if  $T(X)$  is larger than  $M$  of  $N$  random shuffles, then  $p=(1+N-M)/(1+N)$ .<sup>114</sup> Effect sizes were quantified as the difference between the test statistic  $T(X)$  relative to the median of its null distribution  $\{T(X')\}$ .

### Pseudosessions

A key potential pitfall in examining relationships between an autocorrelated behavioral variable like task context and neuronal firing rates is the risk of finding “nonsense correlations”: since both variables are time-series with slowly varying components and high autocorrelation, they often covary by chance, leading to excess false positives with standard statistical tests.<sup>78,79</sup> To address this confound, we designed the task with randomized block lengths (Figures 1C and S1H), allowing us to create “pseudosessions” using the same generating distribution that produced the true block sequence. Pseudosessions preserve the temporal statistics of the original context time series while changing its precise sequence.

We generated pseudosessions using the same random distribution from which the block sequence was drawn (random uniform distribution between 40 and 70 non-repeat trials per block; first block between 1 and 70) to create “pseudoblock” sequences that served as an empirical null distribution for the context variable. When trials in the true session were filtered for reaction times, we applied identical filtering to pseudosessions: reaction times were assigned to each pseudosession trial by sampling with replacement from trials of matching type (i.e., left stim, center stim Left Block, center stim Right Block, right stim).

To obtain shuffled context representations (Figure S1H), i.e. the latent context variables from the context belief behavioral model applied to pseudosessions, we used the fitted weights for each mouse with pseudosession trial sequences. The model’s output probability distribution was sampled based on previous observations and latent values to predict each trial’s choice. Choices were fed back to update the latent value and predict subsequent choices. Following the behavioral task rule, trials with incorrect choices were repeated.