nature neuroscience

Article

https://doi.org/10.1038/s41593-026-02514-7

Information formatting in neuronal populations predicts differential functional contributions to behavior

Received: 28 February 2025

Accepted: 13 October 2026

Published online: 10 November 2026

Check for updates

Douglas A. Ruff [1] ✉, Sol K. Markman [1,2], Jason Z. Kim [3] & Marlene R. Cohen [1] ✉

Neurons throughout the brain encode essentially everything, so why have distinct areas? We proposed that areas format different information sources differently, implying their computational role. Here we compared the middle temporal (MT) area, involved in visual motion perception, and the dorsolateral prefrontal cortex (dlPFC), implicated in decision-making and reward. In a task that required monkeys to combine motion and reward information, both regions encoded both motion and reward, but formatted differently. Area MT represented them separably, whereas the dlPFC combined them in a decision-related format. A recurrent network model trained to reproduce that formatting predicted a distinct effect of stimulating those areas. Consistent with predictions, MT stimulation biased choices whereas dlPFC stimulation produced 'winner-take-all' outcomes. Our findings suggest that a modular structure flexibly reformats information to enable complex behaviors.

Animals capable of complex behaviors tend to have more distinct brain areas than simpler organisms, and artificial networks that perform many tasks tend to self-organize into modules[1–3]. This suggests that different brain areas serve distinct functions supporting complex behavior. Many experiments demonstrate a functional segregation across different brain areas and pathways[4], yet the responses of individual neurons reflect multiple sensory, cognitive and/or motor processes. A recent set of studies indicates that essentially anything that an animal senses, knows or does can be decoded from neural activity in any brain area[5–7].

The goal of our study was to understand what the representation of information from different sources can tell us about the role of a neuronal population in behavior. For example, in the visual cortex, MT neurons are tuned for visual motion direction[8–12] and their responses are modulated by reward expectation and other cognitive processes[13–16]. However, the presence of signals related to both motion and reward does little to identify the role of these neurons in behavior. The same distributions of tuning and modulation are consistent with multiple ways of formatting information in the population (which is sometimes called representational geometry or neuronal population geometry[17,18]). Neuronal population geometry refers to the representation of stimulus, task or motor information in a high dimensional space where the response of each neuron represents one dimension[19–26]. In such a space, neuronal population responses to a circular variable like motion direction trace out a circle. Multiplicatively scaling the responses of individual neurons to simulate, for example, attention or reward expectation[27,28], can stretch and/or shift that circle. A shift normal to the plane of the circle would mean that motion and reward information are represented separably, whereas scaling the circle or a shift in the plane of the circle would imply different formatting, in which motion and reward representations are intertwined.

The different possibilities arise because even identically tuned neurons have heterogeneous modulation by cognitive processes[29]. Simulating this heterogeneity by adding some randomness in the

[1]Department of Neurobiology and Neuroscience Institute, University of Chicago, Chicago, IL, USA. [2]Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA. [3]Department of Physics, Cornell University, Ithaca, NY, USA. ✉e-mail: douglas.ruff@gmail.com; marlenecohen@uchicago.edu

https://doi.org/10.1038/s41593-026-02514-7

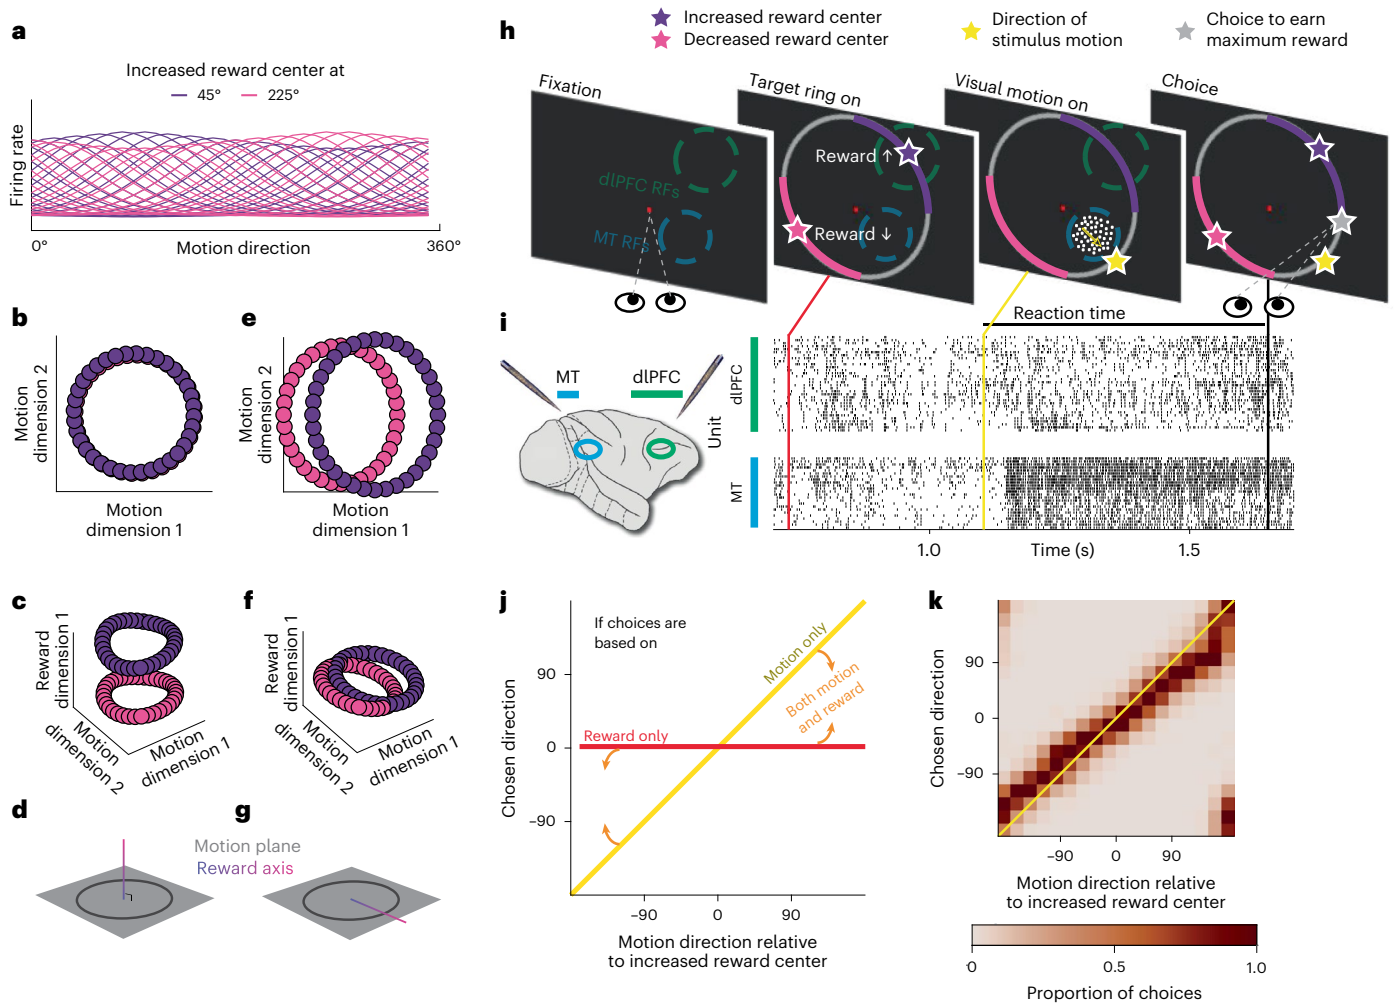


Fig. 1 | Population formatting, task and behavior. a, Simulation demonstrating that neural population formatting cannot be inferred from single neurons. We plotted tuning functions from simulated neurons with preferred visual motion directions that are uniformly distributed across 360° of motion in 10° steps and with gain that depends on the difference between their preferred direction and the direction of maximum expected reward (here either 45° or 225°). The modulation for each neuron was drawn from a single distribution (Methods). **b**, Example population representation of motion direction and reward expectation from the simulated population in **a** (using one representative random draw from the distribution of reward modulations), where motion direction and reward information are encoded separably (so that they are in different dimensions of the space where the response of each neuron is one dimension). To illustrate, the population responses are projected into the two dimensions that encode the most visual motion direction information. **c**, Same as **b**, with a third dimension that accounts for the most reward information. Colors denote the two reward conditions depicted in **a**. **d**, Schematic showing population formatting where motion direction and reward expectation are represented separably (different, or orthogonal, dimensions). **e**, Example population representation of motion direction information and reward information from the simulated population in **a** using a different random draw from the same distribution of reward modulations as in **b**, picked to illustrate that this random draw produced information about motion direction and reward expectation that is combined (occurring in the same dimensions of neuronal population space). Conventions are as in **b**. **f**, Same as **e**, with a third dimension that accounts for the most reward information. Conventions are as in **c**. **g**, Schematic showing population formatting where reward expectation and motion direction are represented jointly (same dimensions). **h**, Continuous motion estimation task schematic. Monkeys indicated their choice by looking at a point on the target ring that corresponded to a direction judgment. They were rewarded according to the accuracy of their estimate of the direction of a dynamic random dot display (actual motion direction depicted by a yellow star; a gray star indicates choice that would result in the maximum reward for this trial), which was scaled by a factor indicated by the color of that portion of the target ring (Methods and Extended Data Fig. 1; pink- and purple-colored stars signify the center of these two possible regions; no stars were shown to the monkeys). The motion direction and reward condition were independently randomly interleaved from trial to trial and the target ring colors associated with increased or decreased reward varied in blocks of trials with uncued changes. **i**, Recordings of populations of neurons in MT and dlPFC using multielectrode probes. In the raster plot, each tick represents an action potential during a single example trial and each row is a multiunit site in the dlPFC (top) or area MT (bottom). The *x* axis represents time (target ring onset, red; visual motion stimulus onset, yellow; onset of eye movement indicating choice, black). **j**, Schematic depicting psychometric functions when choices are based on visual motion direction (yellow), reward condition (red) or a combination (orange arrows signifying choices that fall in between the yellow and red lines). **k**, Behavior from 27 experimental sessions from 2 monkeys (18 from monkey 1; 41,106 total choices) demonstrating that the monkeys based their choices on a combination of motion direction and reward information.

amount of modulation by reward expectation among neurons with identical tuning for motion direction (Fig. 1a and Methods) can produce population representations of motion direction and reward expectation that are either separable (encoded in different dimensions in a space where the response of each neuron is one dimension; Fig. 1b–d) or combined (encoded in the same dimensions; Fig. 1e–g). This simulation demonstrates that the difference between separable and combined population formatting is not reliably inferable from single-neuron tuning properties alone because all responses were drawn from the same distributions of single-neuron responses. Instead, the population

formatting arises from how responses to motion and reward expectation are coordinated across the population.

Previous studies have found that the formatting of information from different sensory, cognitive or motor sources varies across brain areas, task and species[18–27]. Artificial networks that are trained to perform many tasks or complex tasks tend to self-organize into modules that likely format information differently[1,2]. The goal of our study was to use causal manipulations to link population formatting to the role of a neuronal population in behavior.

Results

New insights from a continuous, information-integration task

Our study hinged on having a behavioral task that (1) requires decision-makers to combine information from multiple sources, so that we could link population formatting to behavior, and (2) has a rich enough behavioral output to test hypotheses about the relationship between neuronal responses, causal manipulations and behavior. We therefore trained monkeys to combine two information sources in a continuous estimation task, which provides richer behavioral outputs when compared to typical binary choice tasks[30–32].

Two adult rhesus monkeys (*Macaca mulatta*, both male, 8 and 10 kg) combined sensory information (motion direction) with cognitive information (the reward magnitude associated with different choices) to make continuous choices. They indicated their choice on each trial by making an eye movement to a target ring. The amount of fruit juice reward that they earned depended on the accuracy of their motion direction estimate and a scaling factor associated with particular choices. Colored portions of the target ring indicated which choices would be associated with increased, decreased or unchanged scaling of the reward amount (Fig. 1h,i and Extended Data Fig. 1). We referred to the center of the colored regions as the increased reward or decreased reward center and different locations of those regions (which varied from trial to trial) as reward conditions. During each session, the high reward center was one of two opposite locations chosen randomly from trial to trial and the low reward center was 180° opposite. We chose the two locations based on the tuning of the recorded neurons, so there were different pairs of locations during different sessions. The monkeys demonstrated that they integrated motion and reward information by making choices that varied smoothly with the direction of the visual stimulus and biasing choices toward larger expected rewards (Fig. 1j,k) (choices toward the high reward region were typically associated with faster reaction times; Extended Data Fig. 1f).

Motion and reward information are formatted differently in visual and frontal cortex

While the monkeys performed the continuous estimation task, we recorded simultaneously from a few dozen neurons in two brain areas thought to contribute differently to the decision computation (Fig. 1i). Neurons in area MT are critical for motion perception[8,9,31] and neurons in the dlPFC are part of a network involved in decision-making, reward processing and eye movement planning[33–35]. Consistent with observations that sensory, cognitive and motor information is encoded throughout the brain[6,7,36,37], motion and reward information were readily apparent in the population activity of both area MT and the dlPFC (Fig. 2) and dlPFC units tended to exhibit more nonlinear mixed selectivity for motion and reward than MT units[17,38] (Extended Data Fig. 2a).

However, motion and reward information were formatted differently in the two areas. Motion information in both areas was represented in a low-dimensional manner (Extended Data Fig. 2b,c). However, in area MT, population representations of motion and reward tended to be separable. This means that information about them occupied distinct dimensions of a neuronal population space where each dimension represented the response of one neuron (Fig. 2a). In contrast, in the dlPFC, motion and reward tended to be represented jointly, such that they occupied overlapping dimensions in neuronal population space (Fig. 2b). This difference in neuronal population formatting (Figs. 2c,d) means that motion and reward information is encoded independently in area MT, whereas, in the dlPFC, the same population activity can indicate different motion directions in different reward conditions. These differences in formatting make this a good system in which to investigate questions about how formatting is related to the contribution of a brain area to behavior.

An artificial neural network to generate causal hypotheses

We sought to use causal manipulations (for example, measuring the impact of perturbing neuronal activity on behavior) to understand the functional implications of differently formatted population representations. From the existing literature, it was difficult to generate hypotheses about how the outcomes of causal experiments would relate to information formatting. Some previous studies have attempted to link neuronal population formatting to information coding or theoretical ideas about how that information could, in principle, be read out[39–41]. However, to our knowledge, no studies have linked neuronal population formatting directly to behavior by comparing, in either model or experiment, the effect of causally manipulating brain areas with different formatting.

To generate hypotheses that we could test in further causal experiments, we constructed a recurrent neural network (RNN) that qualitatively reproduced the neuronal population formatting in area MT and dlPFC (Fig. 3a). Our goal was to create a model that could make predictions about the impact of perturbing neuronal activity in area MT and dlPFC on behavior. Our model therefore took as inputs the same information available to the monkey (motion and reward information). We first trained an area MT-like module to receive a noisy motion signal and the locations of increased and decreased reward centers as inputs and to output sharpened estimates of stimulus motion direction and the reward center locations, independently. After training, we froze the recurrent connections and used the motion output of the area MT-like module as the motion input to the dlPFC-like module, alongside a separate reward center location input. This module was trained to output an estimate of the reward associated with each possible choice given the motion direction and reward condition (see Methods: RNN model). We took the model's choice to be the direction associated with the highest expected reward (Fig. 3b).

The goal of the model was not to provide a quantitatively precise account of the neural data, but to capture key qualitative features of population formatting observed in area MT and dlPFC and to generate testable predictions about the effects of perturbing those populations. Consistent with the data, the model represents motion and reward information in a relatively low-dimensional manner and differs across modules in whether these sources of information are represented independently or jointly. Importantly, these differences in formatting lead to distinct predictions about the effects of perturbations, which we then tested experimentally. Like the monkeys, the model appropriately combined information about motion and reward (qualitatively compare Fig. 3b to Fig. 1k). By construction, the recurrent layers in the two modules represent motion and reward in qualitative ways that mimic area MT and dlPFC. Like our population recordings, the model represented task information in a low-dimensional manner (Methods) and demixed-principal component analysis (PCA) revealed that more variance is explained by the interaction of motion and reward information in the dlPFC module (20% of explained variance across conditions) than in the MT module (0% variance explained across conditions; compare to Fig. 2d). The area MT-like module represented motion and reward information separably (Fig. 3c–e), whereas these features were integrated in the dlPFC-like module (Fig. 3f–h). This can be visualized by viewing the stable fixed points that the activity of the hidden layer approaches at the end of a trial (one fixed point for each combination of motion direction and the two reward conditions (Fig. 3c,f), and to more fully visualize the population structure, across a larger range of possible reward conditions (Fig. 3d,g)). The similarities between the model and

https://doi.org/10.1038/s41593-026-02514-7

the data make this model a good tool for generating hypotheses about the relationship between the population formatting and behavior. We noted that the correspondence between model and data is qualitative rather than quantitative and we did not expect all features of the measured population activity to be captured by this simplified model.

Perturbing units in the recurrent layers illuminates the different functional roles of the two modules

We used our models to generate causal hypotheses about how the impact of perturbing units in the recurrent layer of each module[42] affects the way that the module formats information (Methods). Perturbing units in the area MT-like and dlPFC-like modules had qualitatively different effects on model choices. Consistent with results from a previous study that electrically stimulated area MT[31], activating the area MT-like module biased choices toward the preferred direction of the activated unit, provided that the difference between the motion direction and the preference of the activated unit was less than approximately 150° (vector averaging; Fig. 4a–c, column 1). Perturbing the dlPFC-like module was very different. On some ostensibly identical trials, the model chose the motion direction and, on others, the preferred direction of the activated unit and essentially never in between (winner-take-all; Fig. 4a–c, column 3). Furthermore, the behavioral impact of perturbation depended on the reward condition in the dlPFC-like module but not the area MT-like module (compare Fig. 4b and Fig. 4c; Extended Data Figs. 3 and 4).

Manipulating MT and dlPFC activity had qualitatively different effects on choices

We tested the predictions of our model by electrically microstimulating one electrode at a time in area MT or dlPFC on randomly interleaved trials. Consistent with model predictions, microstimulating area MT and dlPFC affected choices in distinct ways. As in previous work[31] and as predicted by our model, microstimulating area MT biased choices midway between the visual stimulus motion direction and the preferred direction measured at the stimulated site (vector averaging; Fig. 4a–c, column 2). In contrast, and consistent with the prediction of our model, microstimulating the dlPFC produced a bimodal choice distribution, where the monkeys sometimes reported the direction of the visual motion and other times reported the preference measured at the stimulated site (winner-take-all; Fig. 4a–c, column 4) with no evidence of vector averaging. We note that the correspondence between model and data was qualitative rather than quantitative and we did not expect all features of the measured population activity to be captured by this simplified model. The fact that the model reproduces the qualitative effects of perturbations despite these differences suggests that the link between population formatting and behavioral impact depends on coarse features of the representation rather than precise quantitative details.

Furthermore, and consistent with the predictions of our model, stimulating area MT and the dlPFC differently interacted with the reward condition. The behavioral impact of MT microstimulation did not depend strongly on whether the increased or decreased reward was aligned with the preference of the stimulated cells (compare Fig. 4b and Fig. 4c, column 2). The intuition is that, because different information sources are encoded approximately independently in area MT, there is no interaction between the behavioral impact of microstimulation and other information that impacts behavior, like reward expectation. In contrast, the behavioral impact of dlPFC microstimulation is much greater when the increased reward center is aligned with the preference of the stimulated cells (compare Fig. 4b and Fig. 4c, column 4). This is consistent with the observation that, in the dlPFC, the representations of motion and reward (and choices) are joint. It would be impossible to stimulate those neurons to impact choices in a way that is independent of reward, because no aspects of the representations are separable. Consistent with these intuitions, the impact of MT microstimulation on

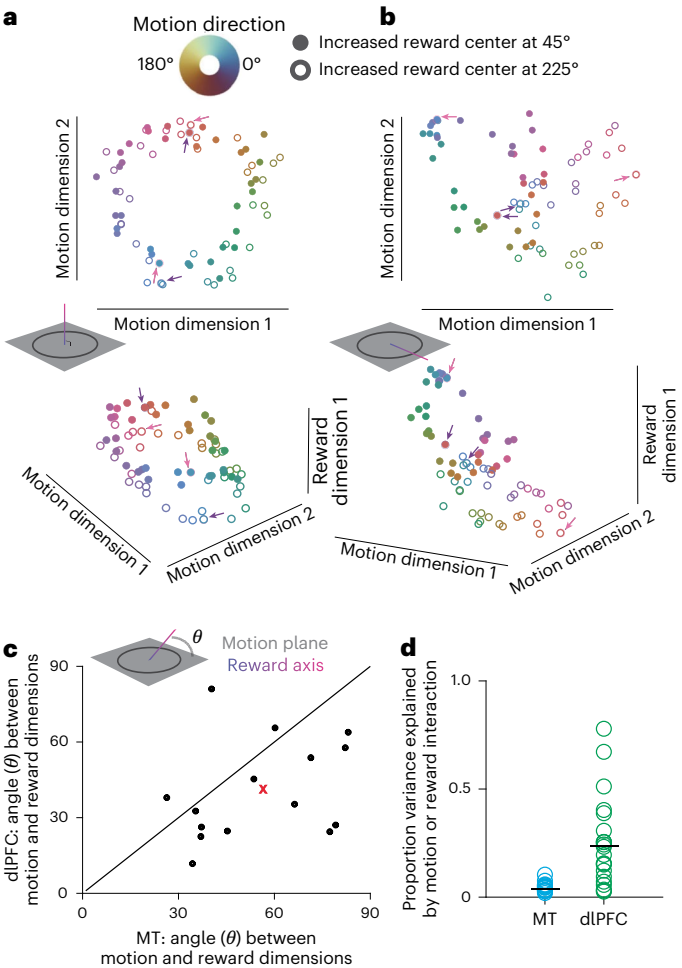


Fig. 2 | Both area MT and the dlPFC encode motion and reward information, but format that information differently. a, MT population responses from an example recording session projected on to the two axes of neuronal population space that account for the most variance in motion direction (top) and the most variance across reward conditions (bottom; *z* axis). Each point is the mean population response to one motion direction (colors) and open and filled circles are from trials in different reward conditions. Pairs of pink and purple arrows of the same orientation indicate responses that align with the reward center in each condition (pink arrow indicates reward increase). **b**, Same as **a**, for dlPFC population responses. **c**, Motion and reward information more separable in MT than dlPFC. In 15 sessions in which MT and dlPFC neurons were recorded simultaneously (10 from monkey 1), the angle between the dimensions accounting for the most variance in motion direction and reward condition (θ in the schematic) was larger in area MT (*x* axis) than in the dlPFC (*y* axis; Wilcoxon's sign-rank test, $P < 0.05$). The red x represents the mean across all 15 sessions. **d**, Demixed-PCA revealing that more variance is explained by the interaction of motion and reward information in the dlPFC than in area MT ($P < 0.001$, two-tailed Wilcoxon's rank-sum test). Data are from 15 sessions in which MT and dlPFC neurons were recorded simultaneously and each data point represents one session of data from each area. The horizontal black lines represent the mean for each area.

choices was indistinguishable in the two conditions (Fig. 5a), whereas dlPFC microstimulation had a larger behavioral impact when the increased reward region was aligned with the preference of the stimulated cells (Fig. 5b). This difference between the areas did not depend on the coherence of the visual motion stimulus (Extended Data Fig. 5).

Interpreting population formatting in relation to task variables

An important issue in interpreting these results is the relationship between what information is represented in a neuronal population

https://doi.org/10.1038/s41593-026-02514-7

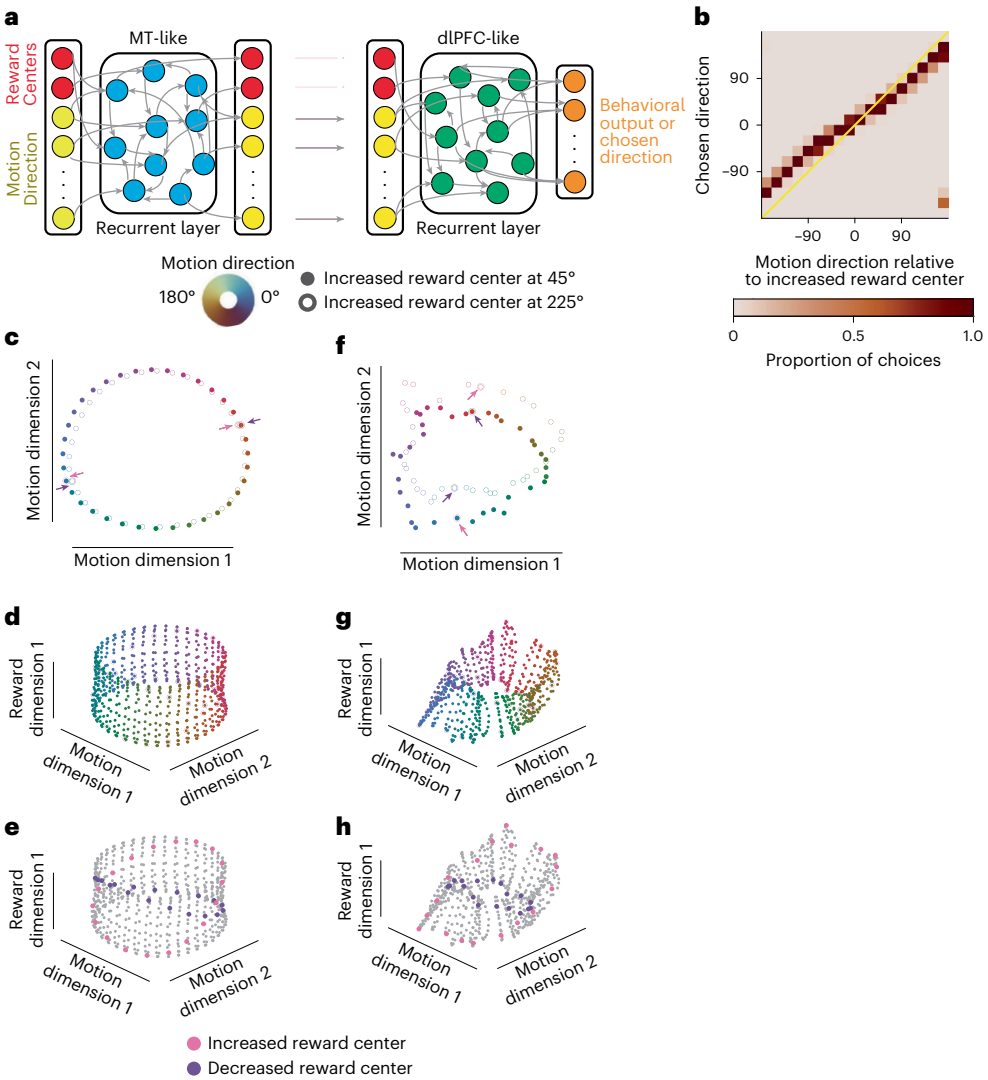


Fig. 3 | A two-stage artificial neural network replicates monkey behavior and population representations in area MT and dlPFC. a, Construction of an RNN that qualitatively reproduced neuronal population formatting in area MT and dlPFC and performance of the same task as the monkeys. An area MT-like module (left) was trained to output the direction of a noisy motion stimulus (yellow) and the increased reward center from two possible locations (red), in separable dimensions. The dlPFC-like module (right) was trained to take the output of the area MT-like module and reward condition on each trial and output an estimate of the reward associated with each possible choice. We took the model's choice to be the direction with the largest predicted reward according to the output units of the dlPFC-like module (orange). **b**, Like the monkeys, the model exhibits behavior that reflects both motion direction and reward information (compare Fig. 1j,k). Conventions are as in Fig. 1k. **c**, Fixed points from the area MT-like model projecting on to the first two axes of population space that account for the most variance in motion direction (compare Fig. 2a). Colors represent different input motion directions and open and filled circles represent each of two possible reward conditions (increased reward center at either 45° or 225°). Pairs of pink and purple arrows of the same orientation point to the responses that aligned with each reward center (pink indicates reward increase). **d**, Same as **c**, but including additional reward conditions projected on to the first two motion dimensions and the first reward dimension (compare Fig. 2a). Reward centers are identified for different reward conditions with pink and purple circles. **e**, Same as **d**, but colored to emphasize the location of the reward centers. **f**–**h**, Same as **c** (**f**), **d** (**g**) and **e** (**h**), for the dlPFC-like module.

and how that information is formatted. In our task, variables such as motion direction, reward condition and choice are inherently related and activity in dlPFC, in particular, is expected to reflect choice-related signals that depend on both motion and reward. Our goal was not to assign exclusive representations of specific variables to each area or to fully dissociate these variables. Instead, we asked whether the formatting of task-relevant information at the population level provides insight into how different areas contribute to behavior. A key is that differences in formatting make distinct predictions about the effects of causal perturbations. When information sources are represented independently, perturbations should affect behavior in a manner that is largely invariant to other task variables. In contrast, when information is represented jointly, the effect of perturbations should depend on the state of those other variables. Consistent with this prediction, the behavioral impact of MT microstimulation was largely independent of reward condition, whereas the impact of dlPFC microstimulation depended strongly on reward alignment. This difference is captured by the population formatting in each area and does not depend on whether dlPFC activity is interpreted as reflecting motion, reward or choice as such. Rather, it reflects how these sources of information are organized within the population response.

As our interpretation does not require a strict dissociation of motion, reward and choice signals, we performed an additional analysis to assess whether the difference in population formatting across areas could be explained by choice-related activity alone (Methods

https://doi.org/10.1038/s41593-026-02514-7

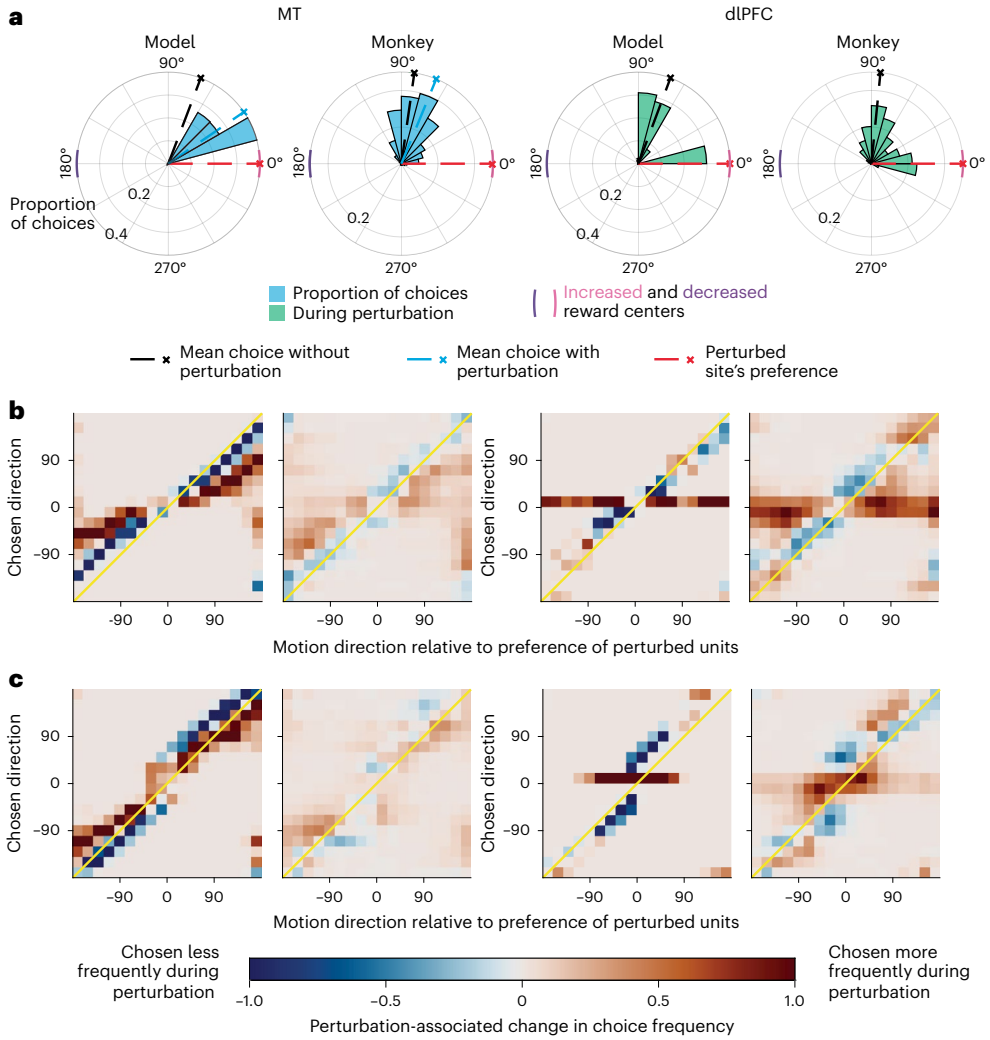


Fig. 4 | Consistent with model predictions, manipulating area MT activity has qualitatively different effects on choices than manipulating dlPFC activity. **a**, Polar histograms of choices from model and monkey in area MT (columns 1 and 2) and dlPFC (columns 3 and 4) on an illustrative subset of trials in which the visual motion direction was 100° counterclockwise relative to the preferred direction of the perturbed units. The pink and purple arcs signify the increased and decreased reward centers, respectively. **b**, Perturbation-related difference in choices when the preference of the perturbed site was aligned with increased reward center for the model (columns 1 and 3) and experimental data (columns 2 and 4). The proportion of choices during nonperturbation trials was subtracted from the proportion of choices during perturbation trials (Extended Data Figs. 3 and 4 display choices from all conditions). In both the area MT-like module (left) and the monkey area MT (right), perturbation biased choices toward the preference of the perturbed site (away from the identity line in columns 1 and 2; monkey data from 29 experimental sessions, comprising 50 sites; some sessions had stimulation on different electrodes on interleaved trials). In both the dlPFC-like module and in monkey dlPFC, perturbation produced bimodal choices that were directed to either the preference of the perturbed site or the mean choice during nonperturbation trials (resulting in choices clustered along the horizontal line representing the preference of the perturbed units in columns 3 and 4; monkey data from 23 experimental sessions, comprising 38 stimulated sites). **c**, Perturbation-related difference in choices when the preference of the perturbed site was aligned with the decreased reward center. Conventions are as in **b**.

and Extended Data Fig. 6). These analyses confirm that dlPFC activity contains strong choice-related components, as expected, but that the key differences in population formatting between areas persist beyond these components.

It is interesting that the requirement of our task that monkeys combine information from visual motion and reward expectation may have heightened the difference between MT and dlPFC stimulation. Unlike previous studies that elicited winner-take-all behavior using MT microstimulation when the angle between the visual motion and the preference of the stimulated site was large[31,43], we typically saw vector averaging whenever we stimulated area MT in our main experiment. We hypothesized that the reward scaling encouraged the monkeys to bias choices toward a larger reward even when the visual motion evidence suggested a choice associated with a smaller reward. Indeed, when we ran a subset of experiments with no reward scaling, we occasionally observed winner-take-all effects during MT microstimulation when the difference between the direction of the visual motion and the preference of the stimulated site was large (Extended Data Fig. 7).

Two controls highlight the robustness of the different impact of MT and dlPFC microstimulation on behavior. First, we investigated the possibility that the impact of stimulation was larger in one brain area than the other, which led to a qualitatively different impact on behavior. We controlled for this by changing the relative strength of the visual and electrical stimuli through a reduction in the visual motion coherence. Rather than qualitatively changing the impact of MT or dlPFC stimulation, this manipulation simply enhanced the impact of microstimulation in both areas, more strongly biasing choices toward the preference of the stimulated site in area MT and increasing the proportion of trials on which the monkey reported the preference of the stimulated site in the dlPFC. The qualitative difference between

https://doi.org/10.1038/s41593-026-02514-7

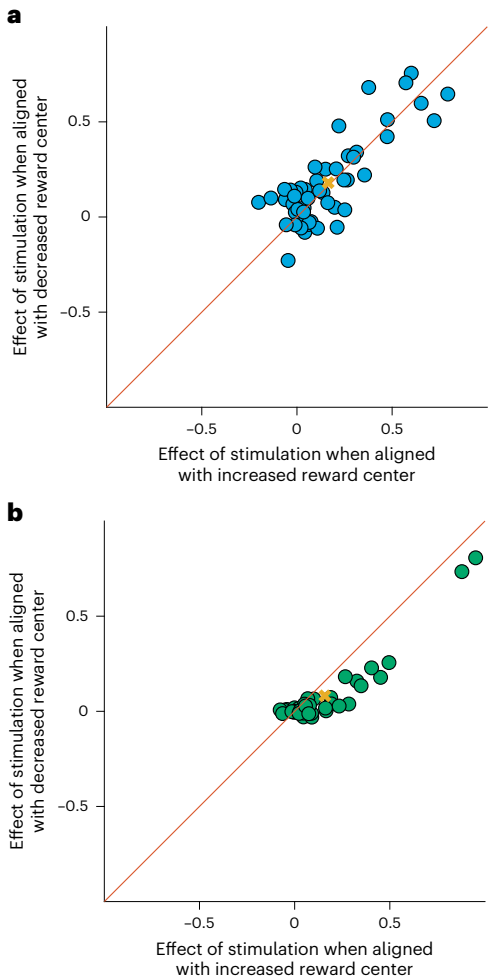


Fig. 5 | The behavioral impact of dlPFC, but not MT, microstimulation depends on reward condition. a, The effect of microstimulation for each stimulated site in area MT summarized as the difference in slope of choices between trials with and without microstimulation (each point summarizes the behavioral impact of one stimulation site during one recording session; $n = 50$ sites from 29 sessions; orange X denotes the mean). Large positive values represent large stimulation-associated biases of choices toward the motion direction preference of the stimulated site. This measure is indistinguishable in the two reward conditions ($P = 0.64$, two-tailed Wilcoxon's rank-sum test). **b**, The effect of dlPFC microstimulation for each stimulated site summarized as the difference in the proportion of choices made toward the response field of the neurons at the stimulated site during trials with and without microstimulation. Large positive values indicate that more choices were directed to the stimulated site's preference. Microstimulation had a larger impact on choices when the reward scaling increase was aligned with the preference of the stimulated site ($P < 0.01$, two-tailed Wilcoxon's rank-sum test). Conventions are as in **a** ($n = 38$ sites from 23 sessions; orange X denotes the mean).

the two areas remained even at low motion coherence (Extended Data Figs. 8 and 9).

Second, we investigated the possibility that the different impact of stimulation in the two areas was attributable to the fact that, for most experiments, the visual motion stimulus overlapped the MT receptive fields and the target ring overlapped the response fields of the dlPFC neurons. We controlled for this difference by conducting a subset of microstimulation experiments when the MT receptive fields overlapped the target ring rather than the visual motion stimulus. In these experiments, we did not systematically bias choices or induce winner-take-all behavior (Extended Data Fig. 10).

Discussion

A causal link between neuronal population formatting and function

Our study made a leap from correlation to causation in the understanding of the relationship between the formatting of neuronal population responses and behavior. As in previous studies[43–47], we found that, when choices are based on combining motion and reward information, both area MT and the dlPFC encode both information sources. However, we demonstrated that those representations are formatted differently. This observation means that this system is a platform for testing predictions about what neuronal population formatting implies about the function of those areas.

Electrical microstimulation has long been used as a tool to understand the contribution of a population of neurons to behavior and there is a well-known relationship between the effects of microstimulation on behavior and the spatial organization of feature tuning in the cortex[48]. Area MT clearly contains more spatial organization for motion direction information than the dlPFC, but the relationship between spatial organization and population formatting is not obvious and is an important avenue for future study.

A growing body of theoretical work has quantified how flexibility or robustness in the way that neuronal population representations could, in principle, guide behavior depends on population formatting[23,49–51]. Indeed, this theoretical framework predicts that the formatting differences between area MT and the dlPFC are consistent with longstanding ideas that visual cortex functions as a close to veridical representation of the world and that the dlPFC flexibly combines information to guide decision-making[33,34,46]. However, we showed, using our model trained only to reproduce the formatting, that there is a causal link between formatting and function: as predicted by our models, electrically stimulating those areas had a qualitatively different impact on behavior. Our results support the exciting idea that the contributions of less well-studied brain areas to less well-studied behaviors can in the future be understood by measuring the way in which they format information from different sources.

We used artificial neural networks in a slightly different way than many recent neuroscience studies[2,49,52–59], because we used them simply as hypothesis generators. The use of RNNs has generated some controversy in neuroscience, mainly out of concern over the fact that the architecture of these networks is vastly different from biological networks[60]. Our work highlights the value of using RNN models that are constrained by neuronal data to generate experimentally testable predictions about the relationships between neurons and behavior rather than as models of a biophysically realistic mechanism.

Finally, our study provides a resolution to the seeming paradox of having multiple brain areas that encode redundant sensory, cognitive and motor information ('everything is everywhere'). Our results suggest that different brain areas divide up neuronal computations in ways that are reflected in the formatting of different information sources. We demonstrated that the ability to decode information from a neuronal population is not sufficient to determine how that population contributes to behavior. Instead, future efforts to understand neuronal computations and to intervene in disorders that affect cognition should focus on how the flexible reformatting of neuronal population responses in different areas enables flexible behavior.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41593-026-02514-7.

References

1. Driscoll, L. N., Shenoy, K. & Sussillo, D. Flexible multitask computation in recurrent networks utilizes shared dynamical motifs. *Nat. Neurosci.* **27**, 1349–1363 (2024).
2. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).
3. Roth, G. & Dicke, U. Evolution of the brain and intelligence. *Trends Cogn. Sci.* **9**, 250–257 (2005).
4. Ungerleider, L. G. & Mishkin, M. *Analysis of Visual Behavior* (eds Ingle, D. J., Goodale, M. A. & Mansfield, R. J. W.) 549–586 (MIT Press, 1982).
5. Tremblay, S., Testard, C., DiTullio, R. W., Inchauspé, J. & Petrides, M. Neural cognitive signals during spontaneous movements in the macaque. *Nat. Neurosci.* **26**, 295–305 (2023).
6. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
7. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
8. Parker, A. J. & Newsome, W. T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* **21**, 227–277 (1998).
9. Born, R. T. & Bradley, D. C. Structure and function of visual area MT. *Annu. Rev. Neurosci.* **28**, 157–189 (2005).
10. Dubner, R. & Zeki, S. M. Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Res.* **35**, 528–532 (1971).
11. Maunsell, J. H. R. & Van Essen, D. C. Functional properties of neurons in middle temporal visual area of the macaque monkey. I. Selectivity for stimulus direction, speed, and orientation. *J. Neurophysiol.* **49**, 1127–1147 (1983).
12. Albright, T. D., Desimone, R. & Gross, C. G. Columnar organization of directionally selective cells in visual area MT of the macaque. *J, Neurophysiol.* **51**, 16–31 (1984).
13. Cook, E. P. & Maunsell, J. H. Attentional modulation of behavioral performance and neuronal responses in middle temporal and ventral intraparietal areas of macaque monkey. *J. Neurosci.* **22**, 1994–2004 (2002).
14. Seidemann, E. & Newsome, W. T. Effect of spatial attention on the responses of area MT neurons. *J. Neurophysiol.* **81**, 1783–1794 (1999).
15. Treue, S. & Maunsell, J. H. R. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* **382**, 539–541 (1996).
16. Treue, S. & Martínez Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).
17. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix high dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016).
18. Kaufman, M. T. et al. The implications of categorical and category-free mixed selectivity on representational geometries. *Curr. Opin. Neurobiol.* **77**, 102644 (2022).
19. Bernardi, S. et al. The geometry of abstraction in the hippocampus and prefrontal cortex. *Cell* **183**, 954–967 (2020).
20. Chung, S. & Abbott, L. F. Neural population geometry: An approach for understanding biological and artificial neural networks. *Curr. Opin. Neurobiol.* **70**, 137–144 (2021).
21. Kriegeskorte, N. & Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17**, 401–412 (2013).
22. Kriegeskorte, N. & Wei, X.-X. Neural tuning and representational geometry. *Nat. Rev. Neurosci.* **22**, 703–718 (2021).
23. Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K. & Kiani, R. Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell* **184**, 3748–3761 (2021).
24. Sorscher, B., Ganguli, S. & Sompolinsky, H. Neural representational geometry underlies few-shot concept learning. *Proc. Natl Acad. Sci. USA* **119**, e2200800119 (2022).
25. Pavuluri, A. & Kohn, A. The representational geometry for naturalistic textures in macaque V1 and V2. Preprint at *bioRxiv* https://doi.org/10.1101/2024.10.18.619102 (2024).
26. Nogueira, R., Rodgers, C. C., Bruno, R. M. & Fusi, S. The geometry of cortical representations of touch in rodents. *Nat. Neurosci.* **26**, 239–250 (2023).
27. Cohen, M. R. & Maunsell, J. Neuronal mechanisms of spatial attention in visual cerebral cortex. in *The Oxford Handbook of Attention* (eds Nobre, K. & Kastner, S.) Ch. 11 (Oxford Univ. Press, 2014).
28. Ruff, D. A., Ni, A. M. & Cohen, M. R. Cognition as a window into neuronal population space. *Annu. Rev. Neurosci.* **41**, 77–97 (2018).
29. Cohen, M. R. & Maunsell, J. H. R. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
30. Srinath, R., Czarnik, M. M. & Cohen, M. R. Coordinated response modulations enable flexible use of visual information. Preprint at *bioRxiv* https://doi.org/10.1101/2024.07.10.602774 (2024).
31. Nichols, M. J. & Newsome, W. T. Middle temporal visual area microstimulation influences veridical judgments of motion direction. *J. Neurosci.* **22**, 9530–9540 (2002).
32. Cloherty, S. L., Yates, J. L., Graf, D., DeAngelis, G. C. & Mitchell, J. F. Motion perception in the common marmoset. *Cereb. Cortex* **30**, 2658–2672 (2020).
33. Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* **2**, 176–185 (1999).
34. Heekeren, H. R., Marrett, S. & Ungerleider, L. G. The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* **9**, 467–479 (2008).
35. Gaymard, B., Ploner, C. J., Rivaud, S., Vermersch, A. I. & Pierrot-Deseilligny, C. Cortical control of saccades. *Exp. Brain Res.* **123**, 159–163 (1998).
36. Charles, F. et al. Brain-wide representations of prior information in mouse decision-making. *Nature* **645**, 192–200 (2005).
37. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
38. Johnston, W. J., Palmer, S. E. & Freedman, D. J. Nonlinear mixed selectivity supports reliable neural computation. *PLoS Comput. Biol.* **16**, e1007544 (2020).
39. Ostojic, S. & Fusi, S. Computational role of structure in neural activity and connectivity. *Trends Cogn. Sci.* **28**, 677–690 (2024).
40. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: high dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016).
41. Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A. & Fiete, I. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nat. Neurosci.* **22**, 1512–1520 (2019).
42. Khona, M. & Fiete, I. R. Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* **23**, 744–766 (2022).
43. Cicmil, N., Cumming, B. G., Parker, A. J. & Krug, K. Reward modulates the effect of visual cortical microstimulation on perceptual decisions. *Elife* **4**, e07832 (2015).
44. Lin, Z., Nie, C., Zhang, Y., Chen, Y. & Yang, T. Evidence accumulation for value computation in the prefrontal cortex during decision making. *Proc. Natl Acad. Sci. USA* **117**, 30728–30737 (2020).

45. Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J. Neurosci.* **31**, 6339–6352 (2011).
46. Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**, 1257–1269 (2003).
47. Rorie, A. E., Gao, J., McClelland, J. L. & Newsome, W. T. Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (LIP) of the macaque monkey. *PLoS ONE* **5**, e9308 (2010).
48. Histed, M. H., Ni, A. M. & Maunsell, J. H. R. Insights into cortical mechanisms of behavior from microstimulation experiments. *Prog. Neurobiol.* **103**, 115–130 (2013).
49. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
50. Sohn, H., Narain, D., Meirhaeghe, N. & Jazayeri, M. Bayesian computation through cortical latent dynamics. *Neuron* **103**, 934–947 (2019).
51. Peixoto, D. et al. Decoding and perturbing decision states in real time. *Nature* **591**, 604–609 (2021).
52. Sussillo, D. Neural circuits as computational dynamical systems. *Curr. Opin. Neurobiol.* **25**, 156–163 (2014).
53. Barak, O. Recurrent neural networks as versatile tools of neuroscience research. *Curr. Opin. Neurobiol.* **46**, 1–6 (2017).
54. Yang, G. R. & Wang, X.-J. Artificial neural networks for neuroscientists: a primer. *Neuron* **107**, 1048–1070 (2020).
55. Ehrlich, D. B., Stone, J. T., Brandfonbrener, D., Atanasov, A. & Murray, J. D. PsychRNN: an accessible and flexible python package for training recurrent neural network models on cognitive tasks. *eNeuro* **8**, ENEURO.0427-20.2020 (2021).
56. Sheahan, H., Luyckx, F., Nelli, S., Teupe, C. & Summerfield, C. Neural state space alignment for magnitude generalization in humans and recurrent networks. *Neuron* **109**, 1214–1226 (2021).
57. Pollock, E. & Jazayeri, M. Engineering recurrent neural networks from task-relevant manifolds and dynamics. *PLoS Comput. Biol.* **16**, e1008128 (2020).
58. Rajalingham, R., Piccato, A. & Jazayeri, M. Recurrent neural networks with explicit representation of dynamic latent variables can mimic behavioral patterns in a physical inference task. *Nat. Commun.* **13**, 5865 (2022).
59. Chaisangmongkon, W., Swaminathan, S. K., Freedman, D. J. & Wang, X.-J. Computing by robust transience: how the fronto-parietal network performs sequential, category-based decisions. *Neuron* **93**, 1504–1517 (2017).
60. Yang, G. R. & Molano-Mazón, M. Towards the next generation of recurrent network models for cognitive neuroscience. *Curr. Opin. Neurobiol.* **70**, 182–192 (2021).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

© The Author(s) 2026

Methods

Population formatting simulation

We conducted a simple simulation to demonstrate that it is possible for populations of neurons with identical visual tuning functions and reward modulation to format information in different ways. We simulated two populations of neurons, each of which consisted of units with the same motion direction tuning functions, which were von Mises' functions with uniformly distributed, preferred directions (preferred directions varied in steps of 10°).

We simulated modulation by reward expectation in ways consistent with the small, heterogeneous, multiplicative gain changes observed in neuronal responses in tasks that vary reward expectation and feature attention[61,62]. We drew gains for each neuron from a uniform distribution between 1 and 1.2 and used these as a multiplier to the response of each neuron. Inspired by single-neuron studies, we scaled this multiplier by the difference between the preferred direction of the neuron and the center of the high reward region. Depending on the extent to which these random gains happened to coordinate across the population, we could choose random draws such that the population formatted motion and reward information differently (Fig. 1a–g).

Participants

We trained two adult male rhesus monkeys, S and O (8 and 10 kg, respectively), to perform a continuous motion estimation task with variable reward expectation. Before training, both monkeys were prepared surgically with a customized titanium head-holding device. To enable electrophysiological recordings, two recording cylinders (Crist instruments Co., Inc.) were placed over (1) the principal sulcus and (2) visual cortex to provide access to the dlPFC and area MT, respectively. All surgical and behavioral procedures conformed to the guidelines established by the National Institutes of Health (NIH) and were approved by the institutional animal care and use committees of the University of Pittsburgh and Carnegie Mellon.

Behavioral task

Procedures. During training and recording sessions, the monkeys sat in a primate chair that provided head restraint. Visual stimuli were presented at a 120-Hz refresh rate on a high-quality LCD monitor (VPixx Technologies Inc.) placed 57 cm from the monkey's eyes. Eye movements were monitored using an infrared optical eye tracker (Eyelink 1000, SR Research). Behavioral control and stimulus presentation were managed by computers using a Unix-based operating system (running MATLAB, MathWorks, Inc. and PsychToolbox[63,64]).

Task and training. A trial began when the monkey fixated a central spot on a black background (inside an invisible 1° diameter window). Then a circular target ring was displayed and remained visible for the remainder of the trial. The target ring was uniform gray except for two segments 180° apart that were colored either pink or purple to signify increased or decreased reward scaling (see below; Extended Data Fig. 1). After a 200-ms to 400-ms delay, a dynamic random dot kinematogram was displayed in a stationary circular aperture inside the target ring. These motion stimuli were identical to those described in ref. 65 and were presented at a single motion strength (coherence), typically between 40% and 50%, for most experiments (exceptions are described in Extended Data Figs. 8 and 9). During multi-coherence experiments, trials using two motion coherences were randomly interleaved (low coherence: 15–20%; high coherence: 35–50%) and during the experiments described in Extended Data Fig. 7, where there was no reward scaling on the target ring and data from only low-coherence trials (15–20%) are shown.

Monkeys were trained to make a saccadic eye movement to a location on the target ring to signify their choice. The monkeys could initiate the eye movement between 150 ms and 1,500 ms after visual stimulus onset. The reward was based on the accuracy of the motion estimate, scaled by the amount indicated by the location of the pink-colored and purple-colored portions of the target ring. The color signified the increased or decreased reward center switched in blocks with randomized, unsignaled changes throughout the experiment (typically two to eight changes per session; Extended Data Fig. 1).

During training, the monkeys first learned to perform the motion continuous estimation task before the reward scaling was introduced. The monkeys learned to saccade to a location on the target ring that matched the direction of motion of the visual stimulus, but not where the vector of motion of the stimulus would cross the target ring. This design ensured that the monkeys' judgments would not be biased by the exact location of either the visual motion stimulus or the target ring, both of which could vary based on the receptive field locations of the recorded neurons.

Electrophysiology experiments

Procedures. Neurophysiological recordings were performed using 24-channel or 32-channel linear probes (V-probes and S-probes; Plexon Inc.) positioned using grids (Crist Instruments Company Inc.) and advanced using a hydraulic microdrive (Kopf instruments). Spiking activity, local field potentials, eye position and task events were recorded at 30,000 samples per second (using Trellis software and Ripple recording hardware). All data were analyzed using customized scripts (MATLAB, MathWorks, Inc.).

Recording locations. Recordings in area MT were made via a chamber positioned over the visual cortex aimed anteriorly toward the MT lobe. This posterior approach provided access to groups of MT units with similar spatial receptive fields, but different direction tuning preferences across the contacts of a multi-channel probe[9,12]. MT units were identified based on stereotactic coordinates, gray–white matter transitions and functional properties. In most experiments, we aimed for MT units with receptive fields that were <12° eccentric, so we could place the motion stimulus within their receptive fields without overlapping the target ring. When we electrically stimulated MT neurons with receptive fields that overlapped the target ring (Extended Data Fig. 10), we used neurons with more eccentric receptive fields. During all recordings, we advanced the probe until we could stably record from as many visually responsive and direction-tuned units as possible. We measured spatial receptive field and direction tuning to verify electrode positioning, guide motion stimulus positioning and, when relevant, identify the tuning preferences of sites recorded on the contacts that we used for electrical stimulation.

Recordings in the dlPFC were made via a chamber positioned over the frontal cortex, anterior to the arcuate sulcus and spanning both sides of the principal sulcus (determined by stereotactic location and visual inspection during surgery). We recorded units along the banks and gyri of the principal sulcus. We used recording locations in which we observed spatial selectivity during any period of a delayed saccadic task and we stopped advancing the probe when we could stably record from as many selective units as possible. We typically recorded from units with response fields at similar eccentricities, with locations that sometimes progressed counterclockwise as penetration depth increased. We preferred units with response fields that were >10° eccentric so that we could position a portion of the target ring within them and avoid overlap with the visual motion stimulus. We placed one colored portion of the ring at a location that was well covered by the response fields of the recorded units and the other colored region was placed 180° away.

Electrical microstimulation

During electrical microstimulation experiments, we selected two contacts on the probe to be stimulated on randomly interleaved trials. Each trial consisted either of no microstimulation or stimulation on one electrode, and there was no statistical relationship between stimulation

and the motion stimulus or reward condition. Microstimulation experiments were only ever performed in one area per session. Stimulation electrodes were randomly selected from suitable candidates based on tuning curves measured at the start of an experiment. The MT sites were required to have a suitable visual receptive field location and strong direction tuning. The dlPFC sites were required to have suitably positioned visual or premotor response fields.

Electrical microstimulation began with the onset of the visual motion stimulus and remained on until the monkey moved its eyes out of the fixation window to initiate a saccade. Microstimulation was biphasic and delivered at 200 Hz with an amplitude that ranged across sessions between 20 μA and 40 μA[66–68].

Quantifying the behavioral impact of microstimulation

To investigate the interaction between reward expectation and the impact of microstimulation on choices, we summarized the impact of microstimulation in each reward and stimulation condition (Fig. 5). As MT stimulation biased choices, we summarized behavior as the slope of the line relating the choices to the direction of the visual motion stimulus (restricted to visual motion stimuli that were within −120° of the center of the region of increased reward scaling). We summarized the stimulation-related change as the change in slope between trials with and without stimulation. As dlPFC microstimulation produced winner-take-all choices, we summarized the stimulation-related change in behavior as the difference in the proportion of choices that were aligned with the preference of the stimulated cells (defined as choices within 10° of the preference of the stimulated site).

Analysis of electrophysiological data

Inclusion criteria. MT units were included for analysis if their average firing rate during a 250-ms epoch after visual motion onset, lagged 50 ms for response latency and were at least 10% higher than during a period of stable fixation on a blank screen. The dlPFC units were included for analysis if their firing rate, during either of the two 250-ms epochs that followed the onset of the target ring or the visual motion stimulus, lagged by 50 ms or the 250-ms epoch that preceded eye movements was at least 10% higher than during a period of stable fixation on a blank screen. Trials were included for analysis if the monkey indicated a choice by looking at the target ring between 150 ms and 1,500 ms after visual motion onset. Spike counts from a 300-ms window, aligned to visual stimulus onset and shifted 50 ms for latency, were used for subsequent analyses.

Mixed selectivity analysis. We analyzed nonlinear mixed selectivity for recordings from 15 sessions in which we recorded the MT area and dlPFC simultaneously. We evaluated the extent to which a linear scaling of motion direction tuning curves from one reward condition could explain the tuning curve in the other reward condition. We normalized the two vectors of mean responses from one unit from each reward condition such that their combined mean equaled 1. Next, we fit a line that scaled the vector from one condition to responses from the other and we evaluated the mean squared error (m.s.e.) in Extended Data Fig. 2.

Dimensionality reduction and angle calculations. To visualize the structure of population activity in monkeys (Fig. 2) and the model (Fig. 3), we performed PCA on trial averaged responses that were either marginalized across visual motion direction or reward conditions. We quantified the extent to which information about the two task features was represented in the same dimensions of neuronal population activity, as the angle between the first PCs for visual motion direction and reward (Fig. 2c). In addition, we investigated neuronal population formatting using demixed-PCA[69] with marginalizations for motion direction, reward condition, time and the combination of motion direction and reward (Fig. 2d).

Regression, residual and decoding analysis. We analyzed the amount of variance explained by different task variables by performing linear regression on responses of individual units from 15 recording sessions in which we recorded area MT and dlPFC simultaneously. A full regression model that included stimulus motion direction, reward condition and choice was compared to a model that only included choice (Extended Data Fig. 6a). Residuals from the choice-only model were then used, alongside the full data, in a decoding analysis to assess the amount of information in MT and dlPFC populations that was separate from choice-related activity. For stimulus motion decoding (Extended Data Fig. 6b,c), the motion direction of the visual stimulus was predicted on individual trials using leave-one-out regression and the full list of predicted directions was correlated with the actual shown directions. For reward condition decoding (Extended Data Fig. 6d,e). the reward condition was predicted on individual trials using leave-one-out regression and the full list of predicted reward conditions was correlated with the actual reward condition. Each decoding analysis was performed separately for the full dataset and the choice-removed residuals.

RNN model

Training and architecture. The full neural network model consisted of two recurrent modules connected in series, such that part of the output of the area MT-like module (module 1 in the equations below) determined the input to the dlPFC-like module (module 2). The modules were trained sequentially.

The dynamics of the modules' recurrent states, x_1 and x_2, are governed by the following equation:

$$\mathbf{x}_{i,t+1} = \alpha \times \left(-\mathbf{x}_{i,t} + A_i \mathbf{r}_{i,t} + B_i \mathbf{u}_{i,t} + \mathbf{b}_{\mathrm{rec}\,i}\right) + \sigma\sqrt{2\alpha}\mathbf{w}_{i,t}, \quad i = 1, 2$$

where the firing rate $\mathbf{r}_{i,t} = \max(0, \mathbf{x}_{i,t})$ is the recurrent state passed through the rectified linear unit activation function, $\mathbf{w}_{i,t} \sim N(0,1)$ is a Gaussian noise process scaled by the recurrent noise factor σ and $\alpha = \frac{\mathrm{d}t}{\tau}$ is the ratio of the step size to the neural time constant. $A_i, B_i, \mathbf{b}_{\mathrm{rec}\,i}$ are the recurrent weight matrices, input weight matrices and recurrent bias vectors of module i, respectively, and $\mathbf{u}_{i,t}$ is the input to module i at time t. The outputs of the modules are given by

$$\mathbf{z}_{1,t} = C_1 \mathbf{x}_{1,t} + \mathbf{b}_{\mathrm{out1}}$$

$$\mathbf{z}_{2,t} = C_2 \mathbf{r}_{2,t} + \mathbf{b}_{\mathrm{out2}}$$

where $C_i, \mathbf{b}_{\mathrm{out}\,i}$ are the output weight matrix and output bias vector, respectively. The behavioral choice is taken to be $\mathrm{argmax}(\mathbf{z}_{2,T})$, where T is the last time point of the trial.

For each timestep while the ring is present but before the onset of the motion stimulus, both the area MT-like and the dlPFC-like modules received a two-dimensional input proportional to $[\sin\theta_R, \cos\theta_R]$ where θ_R is the increased reward center for that reward condition. The area MT-like module outputs a sharpened and amplified version of this information, whereas the dlPFC-like module outputs a nonlinear function of these and its other inputs, as described below. Gaussian noise was added independently to all inputs at each timepoint and no effort was made to match the variability in either neuronal responses or behavior to the experimental data. For each timestep during the motion stimulus, the area MT-like module received a 72-dimensional motion input given by the following von Mises' function:

$$M_1(\theta_i, \theta_M) = -0.3 + e^{c \times k_1 \times \cos(\theta_i - \theta_M)}$$

where $c \in [0,1]$ and θ_M are the coherence and direction of the motion stimulus and the 72 dimensions correspond to evenly spaced directions $\theta_i \in [0, 2\pi]$. The model in Figs. 3 and 4 was trained with $c = 0.6, k_1 = 0.3$. The area MT-like module was first trained separately to take in this noisy

signal and output an amplified, sharper 72-dimensional motion output given by the follwoing von Mises' function:

$$M_2(\theta_i, \theta_M) = a \times (-0.3 + e^{c \times k_2 \times \cos(\theta_i - \theta_M)})$$

with $k_2 = 0.8 > k_1$ and $a = 2$. It was also trained to output $[\sin\theta_R, \cos\theta_R]$ to ensure that the reward condition input is not simply ignored. After the MT area-like module was trained, we froze the weights and trained the dlPFC-like module to take as inputs the motion output from the area MT-like module and combine it with a reward condition input to produce a 72-dimensional reward function output given by

$$R(\theta_i, \theta_M, \theta_R) = \begin{cases} (3.5 + 1.5 \times \cos(1.8 \times d_{\theta_i,\theta_M})) \times (1 + 0.75 \times \cos(d_{\theta_i,\theta_R})), & d_{\theta_i,\theta_M} \le \frac{\pi}{4} \\ 0, & d_{\theta_i,\theta_M} > \frac{\pi}{4} \end{cases}$$

where d_{θ_1,θ_2} is the angular distance between θ_1 and θ_2. This is the same reward scaling as was used in the experiments (Extended Data Fig. 1).

Both modules were trained by supervised learning with m.s.e. loss using the Adam optimizer algorithm[69]. The error signal between the model and desired output was used to define a gradient in the recurrent weights and then updated. The weight matrices were subject to a Dale's law constraint in which all input and output projections are excitatory (≥0) and 20% (30 out of 150) of recurrent units sent solely inhibitory (≤0) projections, whereas the remaining 80% sent solely excitatory projections.

The training and testing of the neural network were implemented with customized code based on open-source packages (PsychRNN[55] and TensorFlow[70]).

In the area MT-like module, the first two dimensions of activity representing direction information accounted for 98.0% of the response variance and the first reward dimension accounted for 55.3% of the variance explained across the possible reward conditions (tested in steps of 20° from 0 to 340). In the dlPFC-like module, activity is higher dimensional; the first two dimensions of activity representing direction information accounted for 52.2% of the response variance across stimulus motion conditions and the first reward dimension accounted for 51.5% of the variance across the possible reward conditions.

Model unit perturbation

To generate predictions about the impact of differences between the information representation in the area MT-like and dlPFC-like module, we performed perturbation experiments in the fully trained models. We perturbed the activity of individual units in the recurrent layers of each module by adding an additional untrained input (that projected only to the stimulated unit) during the motion input period.

Statistics and reproducibility

No statistical method was used to predetermine sample size. No data were excluded from the analyses. The Investigators were not blinded to allocation during experiments and outcome assessment.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request and, for the RNN model, via GitHub at https://github.com/solkm/formattingRNN.

Code availability

Code for analysis and training of the RNN model is available via GitHub at https://github.com/solkm/formattingRNN.

References

61. Martinez-Trujillo, J. C. & Treue, S. Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* **14**, 744–751 (2004).
62. Treue, S. & Martinez-Trujillo, J. C. Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* **399**, 575–579 (1999).
63. Brainard, D. H. The psychophysics toolbox. *Spatial Vis.* **10**, 433–436 (1997).
64. Pelli, D. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vis.* **10**, 437–442 (1997).
65. Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* **12**, 4745–4765 (1992).
66. Moore, T. & Armstrong, K. M. Selective gating of visual signals by microstimulation of frontal cortex. *Nature* **421**, 370–373 (2003).
67. Salzman, C., Britten, K. H. & Newsome, W. T. Cortical microstimulation influences perceptual judgements of motion direction. *Nature* **346**, 174–177 (1990).
68. Salzman, C., Murasugi, C., Britten, K. H. & Newsome, W. T. Microstimulation in visual area MT: effects on direction discrimination performance. *J. Neurosci.* **12**, 2331–2355 (1992).
69. Kingma, D. & Ba, J. Adam: a method for stochastic optimization. Preprint at https://doi.org/10.48550/arXiv.1412.6980 (2017).
70. Abadi, M. et al. Tensorflow: large-scale machine learning on heterogeneous distributed systems. Preprint at https://doi.org/10.48550/arXiv.1603.04467 (2016).

Acknowledgements

We are grateful to K. McKracken for providing technical assistance and to R. Srinath, B. Cowley and members of M.R.C.'s laboratory for comments on earlier versions of this manuscript and other helpful suggestions.

Author contributions

D.A.R. and M.R.C. designed the study. D.A.R. performed the experiments. S.M. and J.K. performed the recurrent network modeling. All the authors analyzed data and wrote the paper.

Funding

This work is supported by the Simons Foundation (Simons Collaboration on the Global Brain, award no. 542961SPI) and the NIH (grant nos. R01EY022930, R01EY034723 and RF1NS121913).

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41593-026-02514-7.

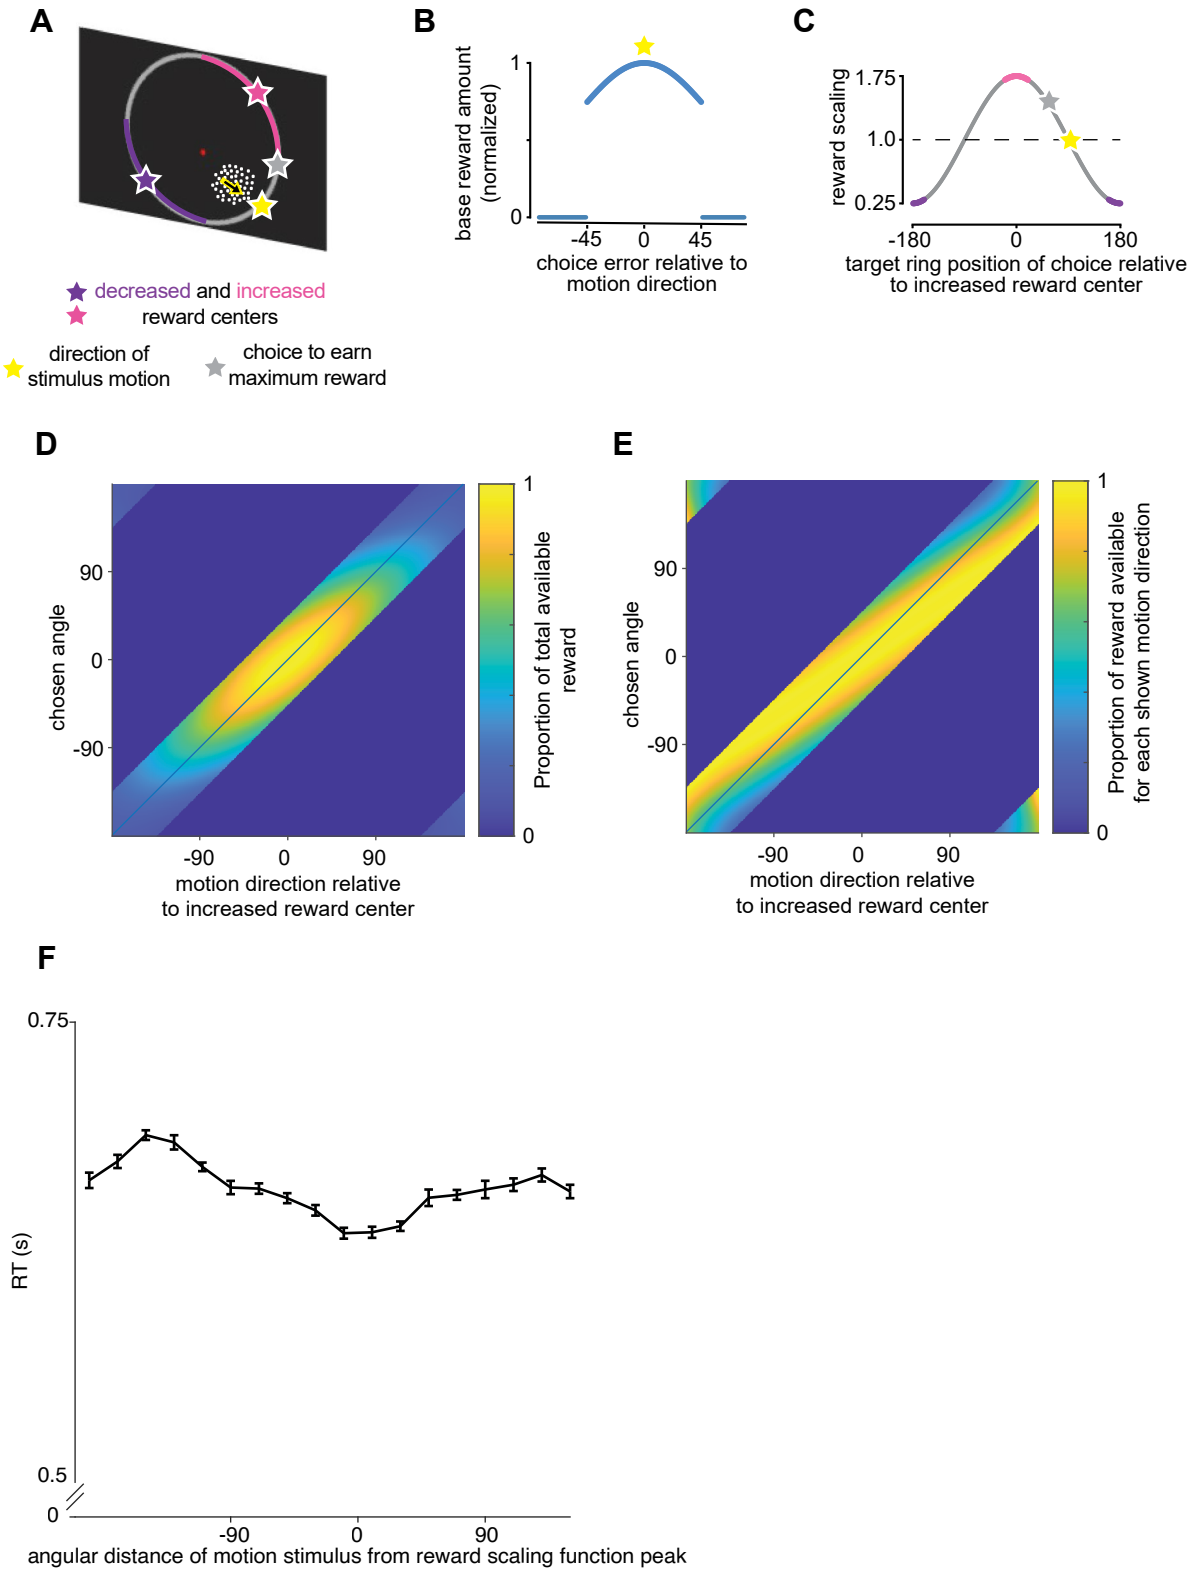
Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41593-026-02514-7.

Correspondence and requests for materials should be addressed to Douglas A. Ruff or Marlene R. Cohen.

Peer review information *Nature Neuroscience* thanks the anonymous reviewers for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

https://doi.org/10.1038/s41593-026-02514-7

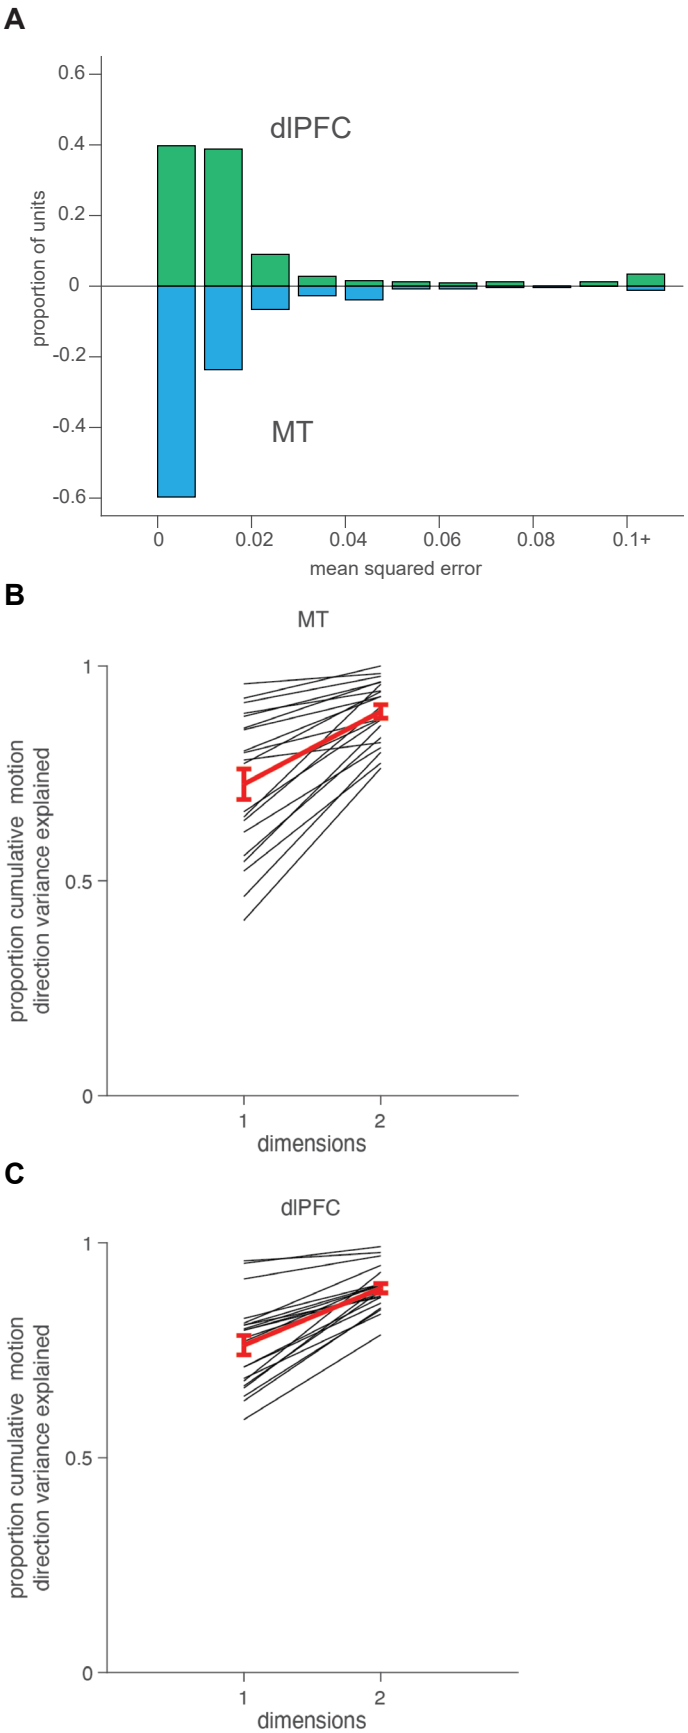


Extended Data Fig. 1 | See next page for caption.

https://doi.org/10.1038/s41593-026-02514-7

Extended Data Fig. 1 | The monkeys were rewarded based on the accuracy of their motion direction judgment, increased or decreased by a reward scaling factor associated with different choices. A) Each trial required a motion direction judgment biased by reward scaling information indicated by the colored portions of the target ring. For the purposes of illustration, the reward scaling centers are noted with stars, the direction of stimulus motion is indicated with a yellow star and the choice associated with the largest reward is marked with a gray star. Both the motion direction and the reward condition were randomly interleaved from trial to trial, and the colors associated with increased or decreased reward varied in blocks with uncued changes. The raw reward amount varied across sessions, but rewards were always based on the (**B**) difference between the monkey's choice and the direction of visual motion (with a hard cut off at +- 45 degrees error) which was scaled by the values in (**C**). (**C**) In this example reward condition, the increased reward center is the middle of the pink arc, and the decreased reward center is the middle of the purple arc. (**D**) Reward amount associated with all choices plotted as a function of the stimulus motion direction relative to the increased reward center (at 0). The largest rewards were available on trials with stimulus motion directions that were close to the increased reward center location. Choices with an error greater than 45 degrees from the motion stimulus were given no reward. (**E**) The same reward data as in D, but normalized within column (motion direction) to show which choices were associated with the most reward for each possible stimulus motion direction. Compare to choices in Fig. 1k. (**F**) Reaction times during the continuous motion estimation task were slightly faster when the motion direction was aligned with the increased reward center. Mean reaction times (+- standard error of the mean) are plotted from 41,106 trials collected from 27 recording sessions (18 from monkey O). Reaction times are defined as the time from the onset of the visual stimulus to the time the monkey initiated the eye movement used to indicate a direction judgment. Reaction times were significantly modulated by the alignment between motion direction and the increased reward center (ANOVA, $p < 0.01$).

https://doi.org/10.1038/s41593-026-02514-7

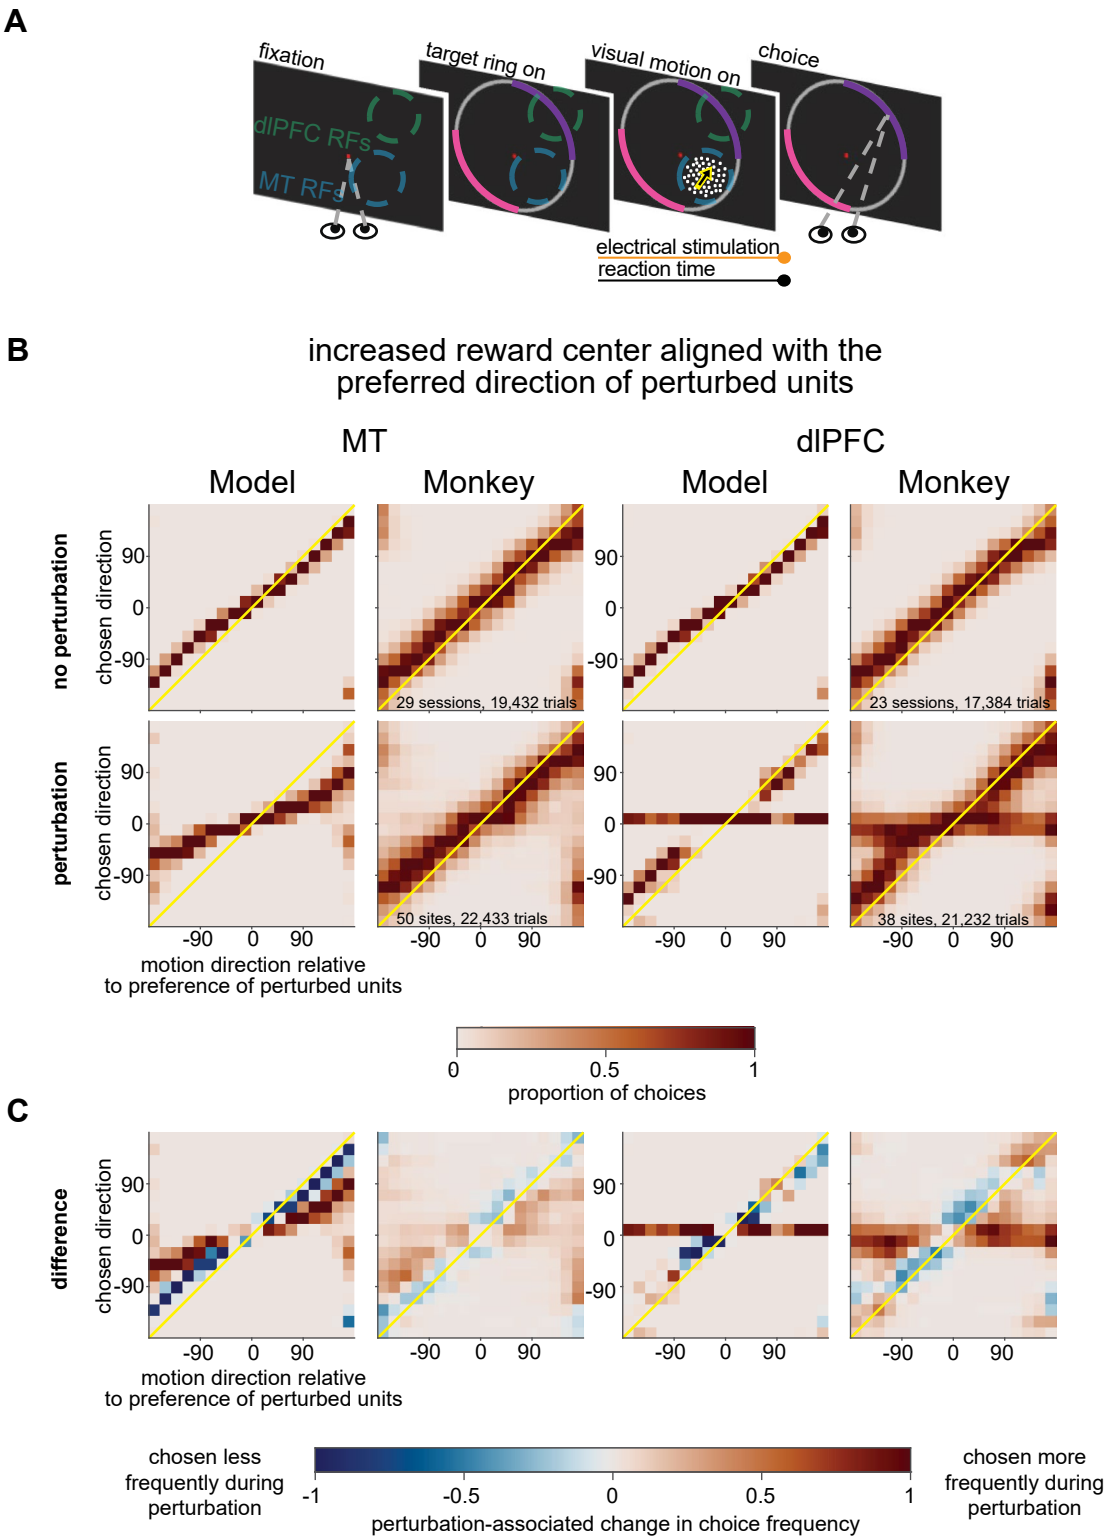


Extended Data Fig. 2 | See next page for caption.

https://doi.org/10.1038/s41593-026-02514-7

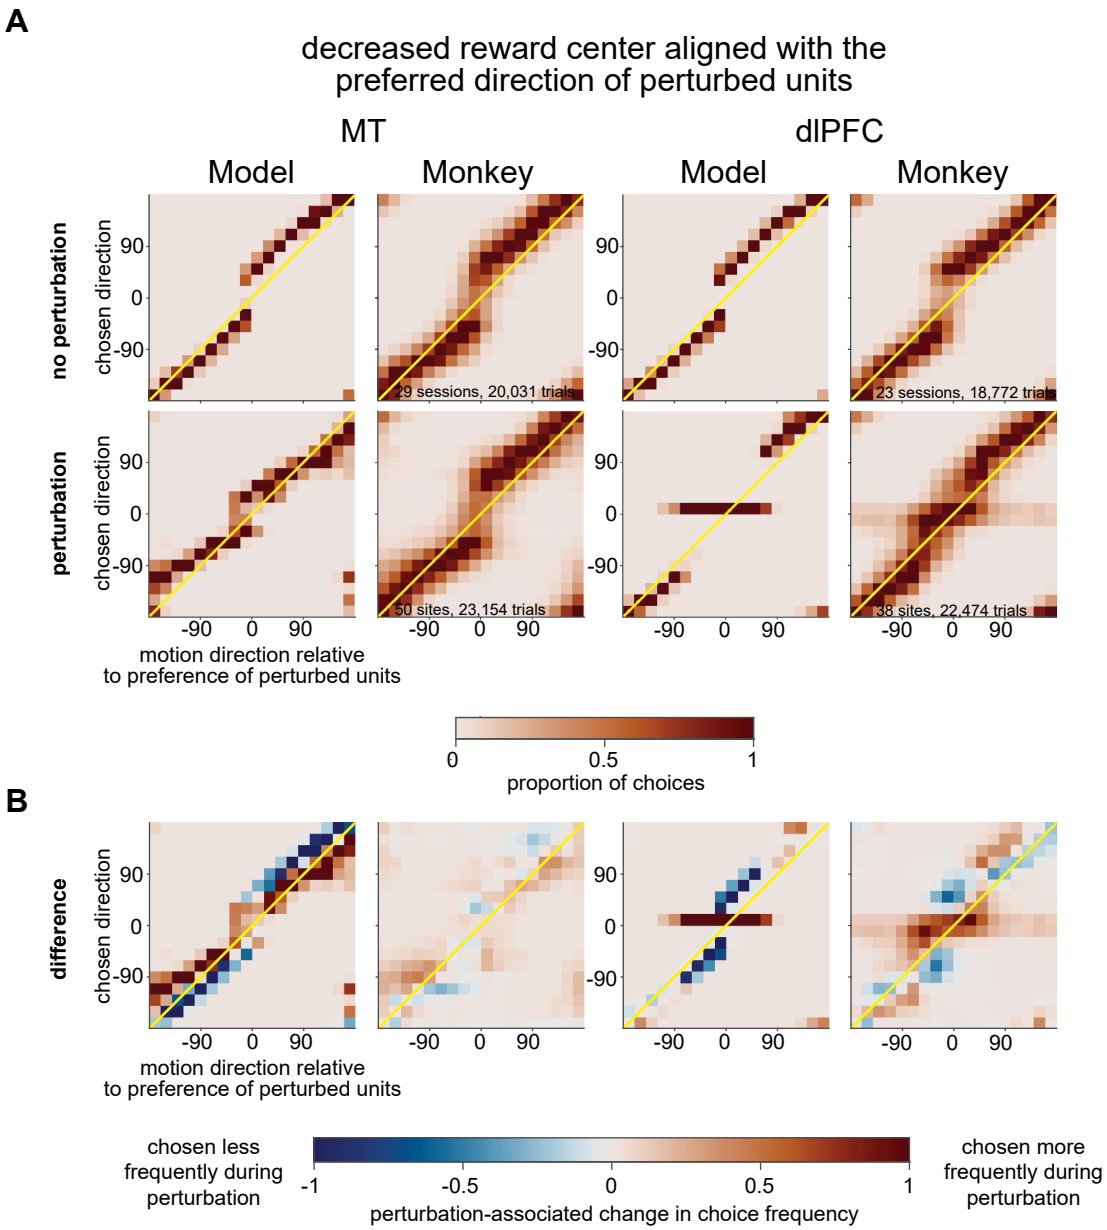
Extended Data Fig. 2 | Information formatting in MT and dlPFC population activity. (**A**) dlPFC units exhibit more non-linear mixed selectivity to visual motion direction and reward information than MT. The histograms of mean squared error of a motion direction tuning curve using a single scaling term to fit the two reward conditions (see Methods section 5.2) show greater error for dlPFC than MT (2-tailed Wilcoxon rank sum test, $p < .01 \times 10^{-8}$). The histograms describe tuning curves for 258 MT units and 322 dlPFC units from 15 recording sessions in which there were simultaneous recordings in the two areas (10 from monkey 1). (**B-C**) Visual motion direction has a low-dimensional representation in both MT and dlPFC. **B**) Cumulative proportion explained variance for first and second dimensions of population averaged MT activity across the same 15 sessions as Fig. 2c. Each black line represents data from a different experimental session, the red line is the mean across sessions. **C**) Same as B, for data from dlPFC during the same recording sessions as in B.

https://doi.org/10.1038/s41593-026-02514-7



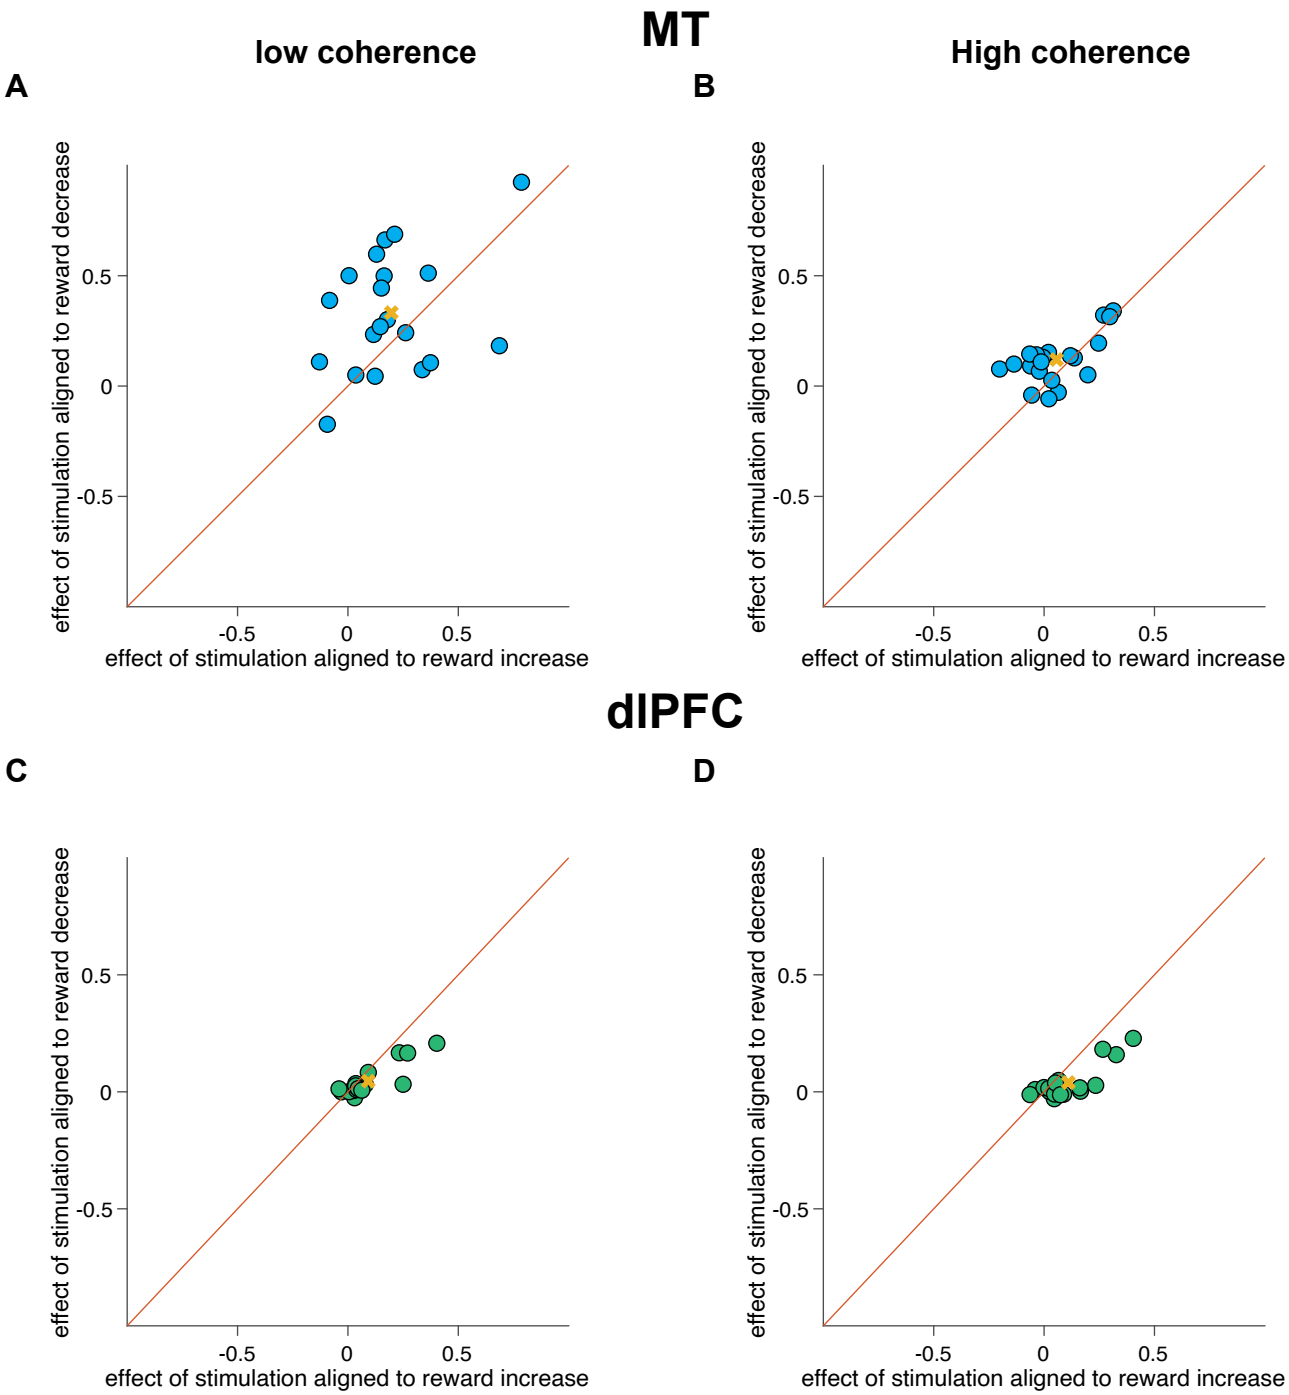
Extended Data Fig. 3 | Behavioral choices from model and monkey during perturbation experiments when the increased reward center was aligned with the preference of the perturbed units. A) During monkey experiments, electrical stimulation coincided with the presentation of the visual stimulus. **B**) Summary of all choices by the model or monkeys when the increased reward center aligned with the preferred direction of the perturbed units. Data from 29 experimental sessions (15 from monkey 1; comprising 50 total stimulated sites across both monkeys; some sessions had stimulation on different electrodes on interleaved trials) of MT microstimulation (column 2), and 23 sessions (8 sessions from monkey 1; comprising 38 total stimulated sites across both monkeys) of dlPFC microstimulation (column 4). Top row are choices during trials with no perturbation, bottom row are choices from trials with perturbation. All trial types, stimulus, and reward conditions were randomly interleaved within a session along with those shown in Extended Data Fig. 4. Data are aligned to the tuning preference of the stimulated site, and sessions are only included in these analyses if the angular difference between that preference and the increased reward center was 30° or less. Data are normalized such that each column sums to 1. **C**) Plots reproduced from Fig. 4b depicting the difference of the plots in (**B**).

https://doi.org/10.1038/s41593-026-02514-7



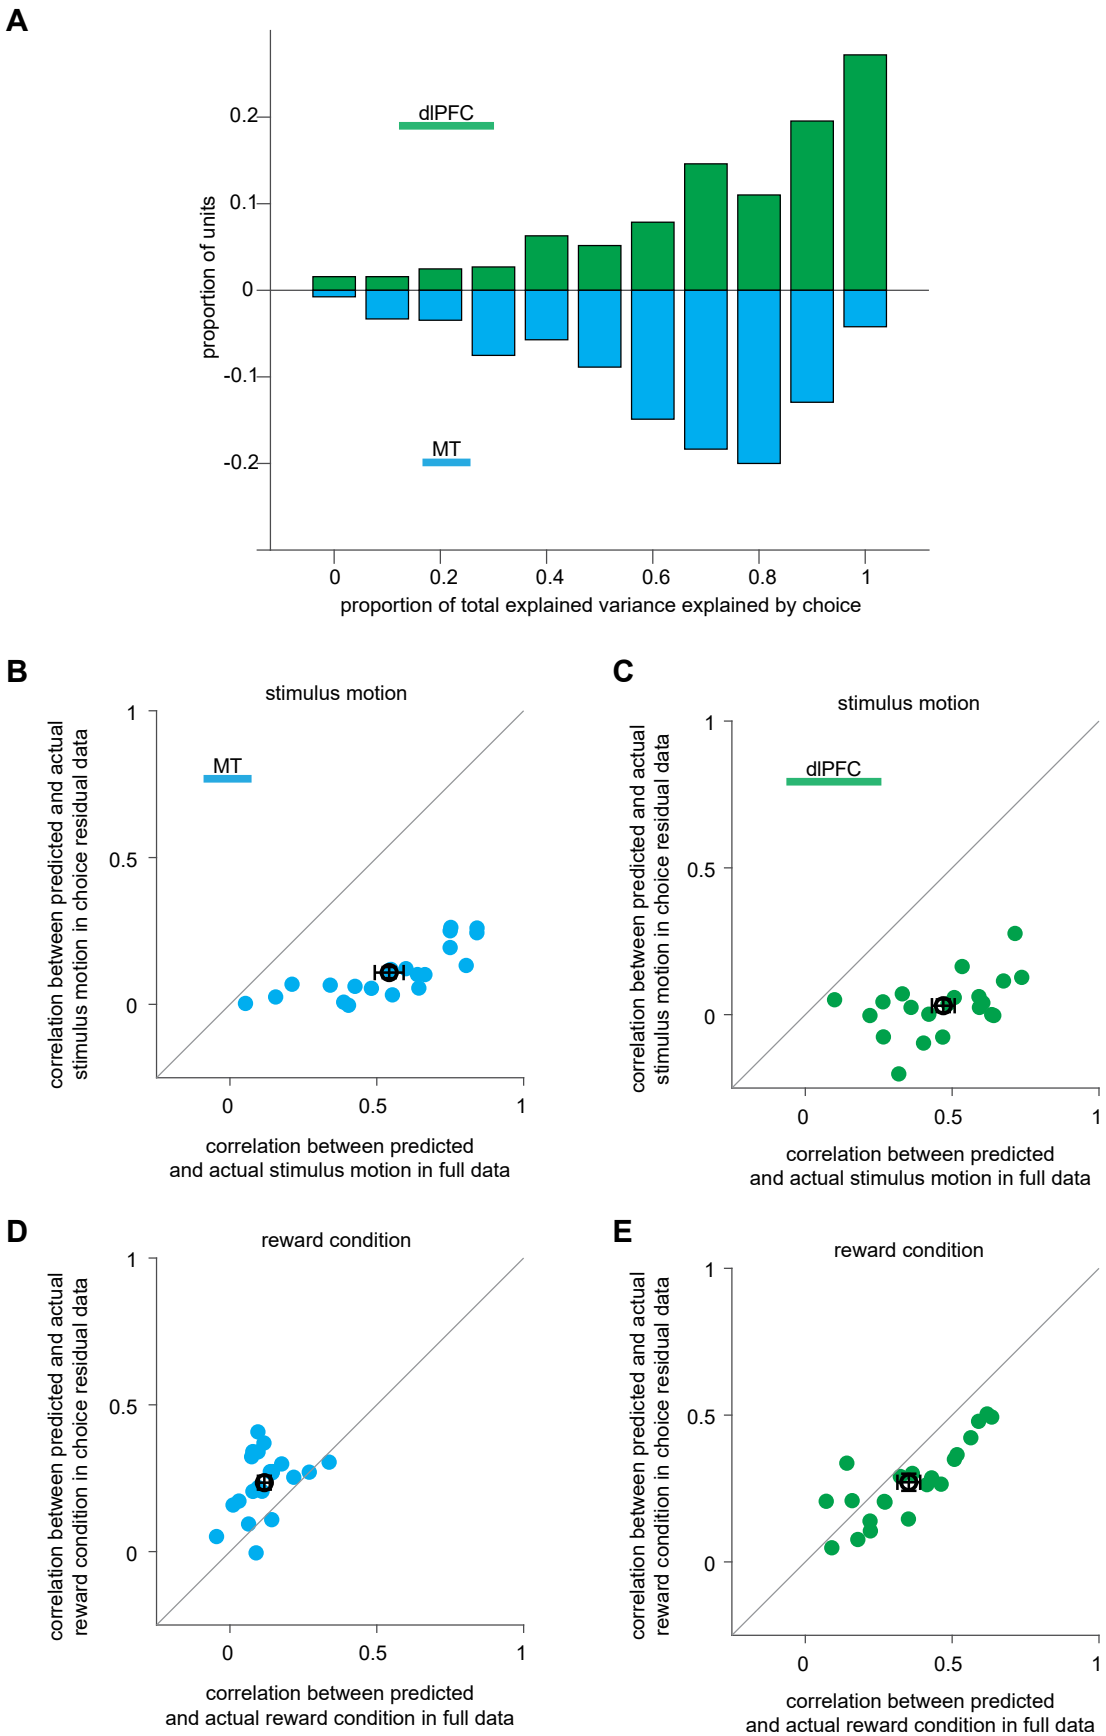
Extended Data Fig. 4 | Behavioral choices from model and monkey during perturbation experiments when the decreased reward center was aligned with the preference of the perturbed units. A) Summary of all choices by the model or monkeys from the same experimental sessions and models as Extended Data Fig. 3, on trials when the decreased reward center aligned with the preferred direction of the perturbed units. Conventions as in Extended Data Fig. 3. **B**) Plots reproduced from Fig. 4c depicting the difference of the plots in (**A**).

https://doi.org/10.1038/s41593-026-02514-7



Extended Data Fig. 5 | The behavioral impact of microstimulation depends on the reward condition in dlPFC but not MT and does not depend on stimulus motion coherence. A) The effect of microstimulation for each stimulated site in MT is summarized as the difference in slope of choices between trials with and without microstimulation (each point summarizes the behavioral impact of one stimulation site on one recording session; orange X denotes the mean) during trials with low coherence visual stimuli. Experiments used two stimulus motion coherences per session (low coherences were chosen between 15% and 20% coherence; high coherences were chosen between 35% and 50% coherence). 12 experimental sessions (7 sessions from monkey 1; comprising 20 total stimulation sites across both monkeys) of MT microstimulation. This measure was indistinguishable in the two reward conditions at low motion coherences ($p = 0.10$, 2-tailed Wilcoxon ranked sum test). **B**) Same as A from the same experimental sessions during interleaved high motion coherence trials ($p = 0.07$, 2-tailed Wilcoxon ranked sum test). **C**) and **D**) The effect of dlPFC microstimulation for each stimulated site is summarized as the difference in the proportion of choices made toward the preferred direction of the neurons recorded at the stimulated site during trials with and without microstimulation. This measure was indistinguishable in the two reward conditions at low coherence (**C**; $p = 0.09$, 2-tailed Wilcoxon ranked sum test) but microstimulation had a larger impact on choices when the reward scaling increase was aligned with the preference of the stimulated site at high coherence (**D**, $p < 0.05$, 2-tailed Wilcoxon ranked sum test). 11 sessions (6 sessions from monkey 1; comprising 17 total stimulation sites across both monkeys) of dlPFC microstimulation.

https://doi.org/10.1038/s41593-026-02514-7

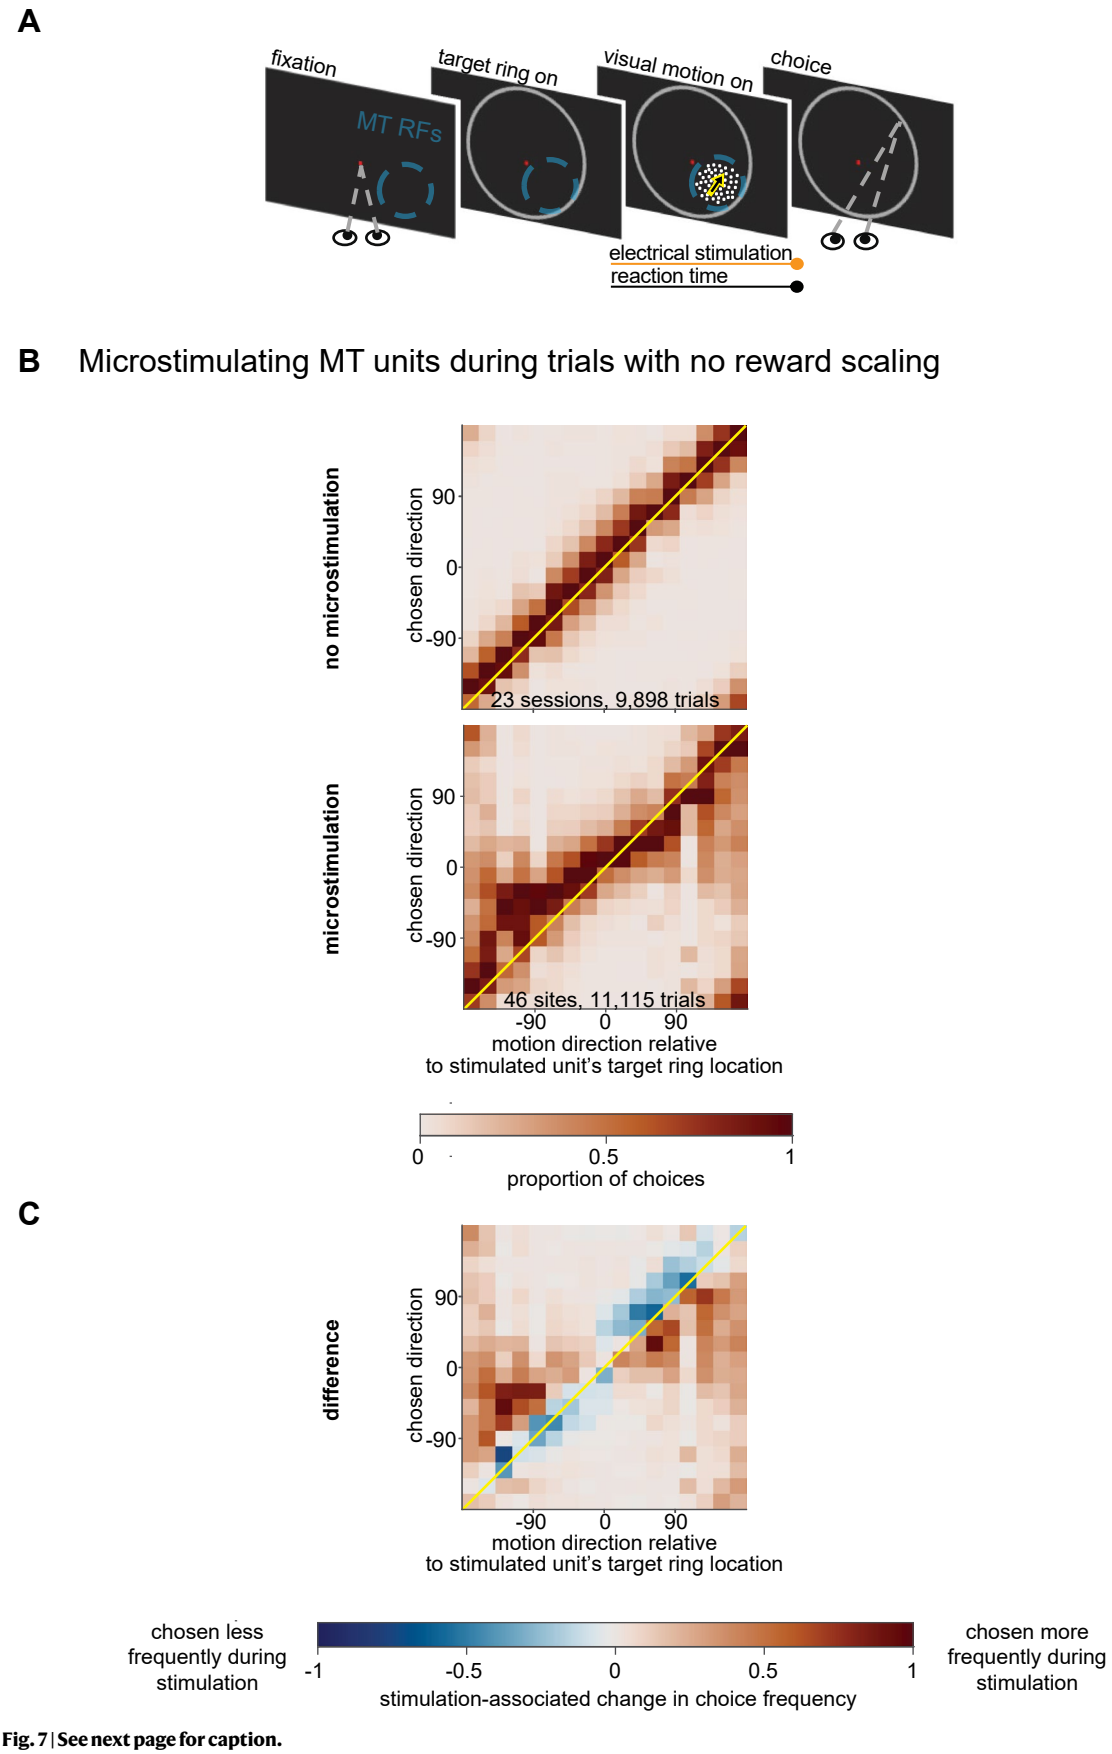


Extended Data Fig. 6 | See next page for caption.

https://doi.org/10.1038/s41593-026-02514-7

Extended Data Fig. 6 | Differences in MT (B, D) and dlPFC (C, E) formatting persist after removing choice-related signals. A) Linear regression performed on individual units recorded in MT and dlPFC reveals that a greater proportion of total explained response variance is explained by choice signals in dlPFC than in MT (2-tailed Wilcoxon ranked sum test, $p < 0.001$). **B**, **C**) A stimulus motion decoder reveals strong overlap of choice and motion related signals in both areas. This is unsurprising given the task design. Removing choice signals from activity (via regression, see Methods) leaves less motion information remaining in the population activity in both areas (2-tailed Wilcoxon ranked sum test, both $p < 0.0001$; means and standard error of the mean are specified by the black point and error bars). **D**)-**E**) A reward condition decoder is less strongly affected by removing choice signals from population activity.

https://doi.org/10.1038/s41593-026-02514-7

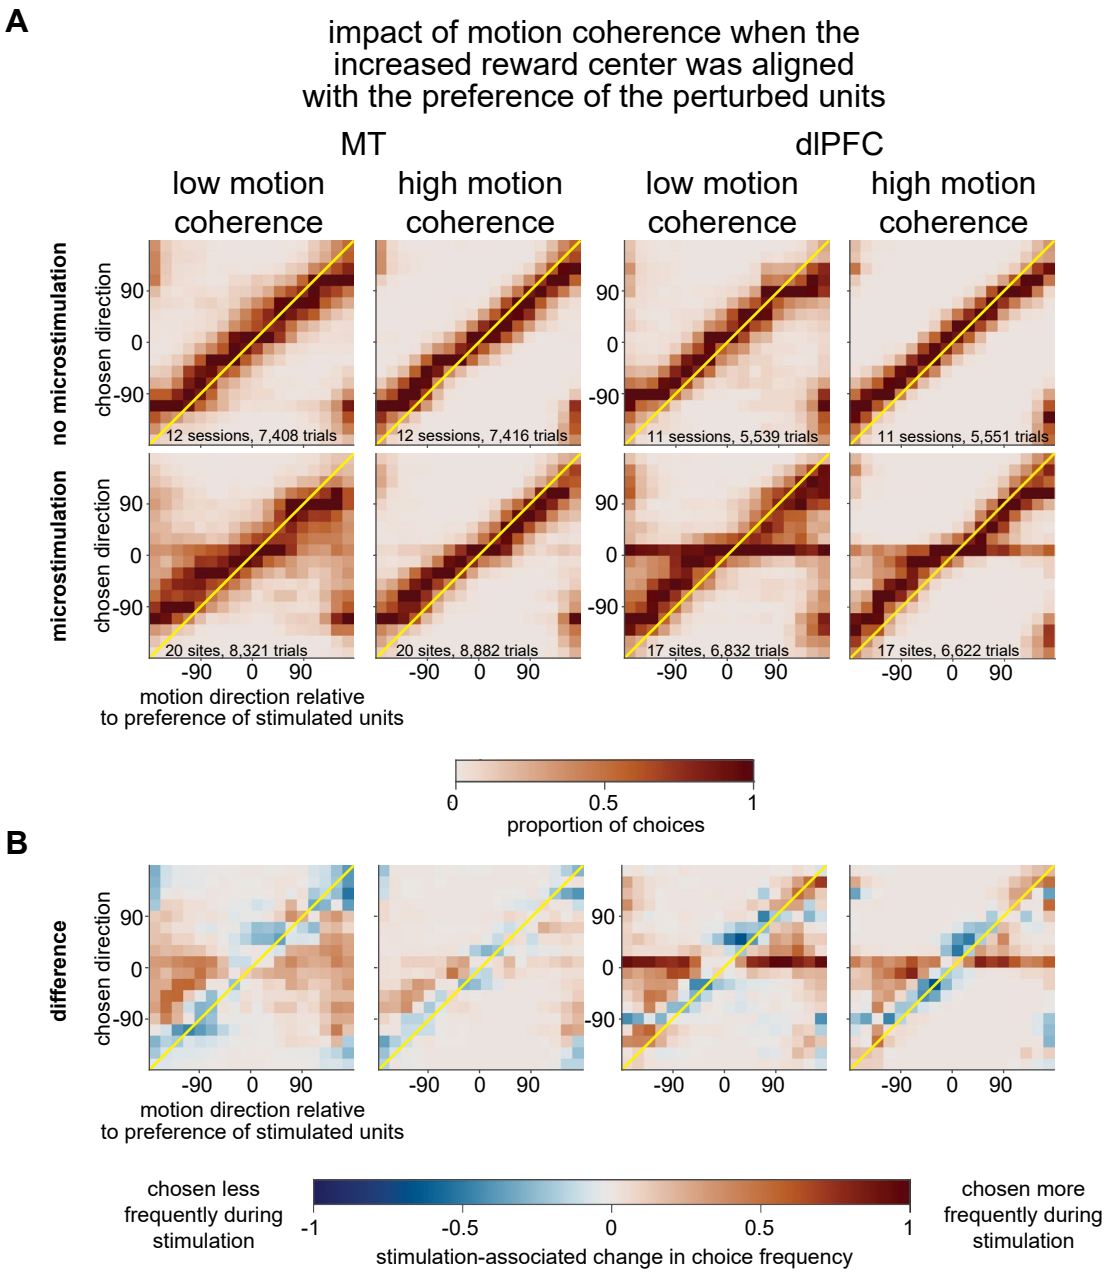


Extended Data Fig. 7 | See next page for caption.

https://doi.org/10.1038/s41593-026-02514-7

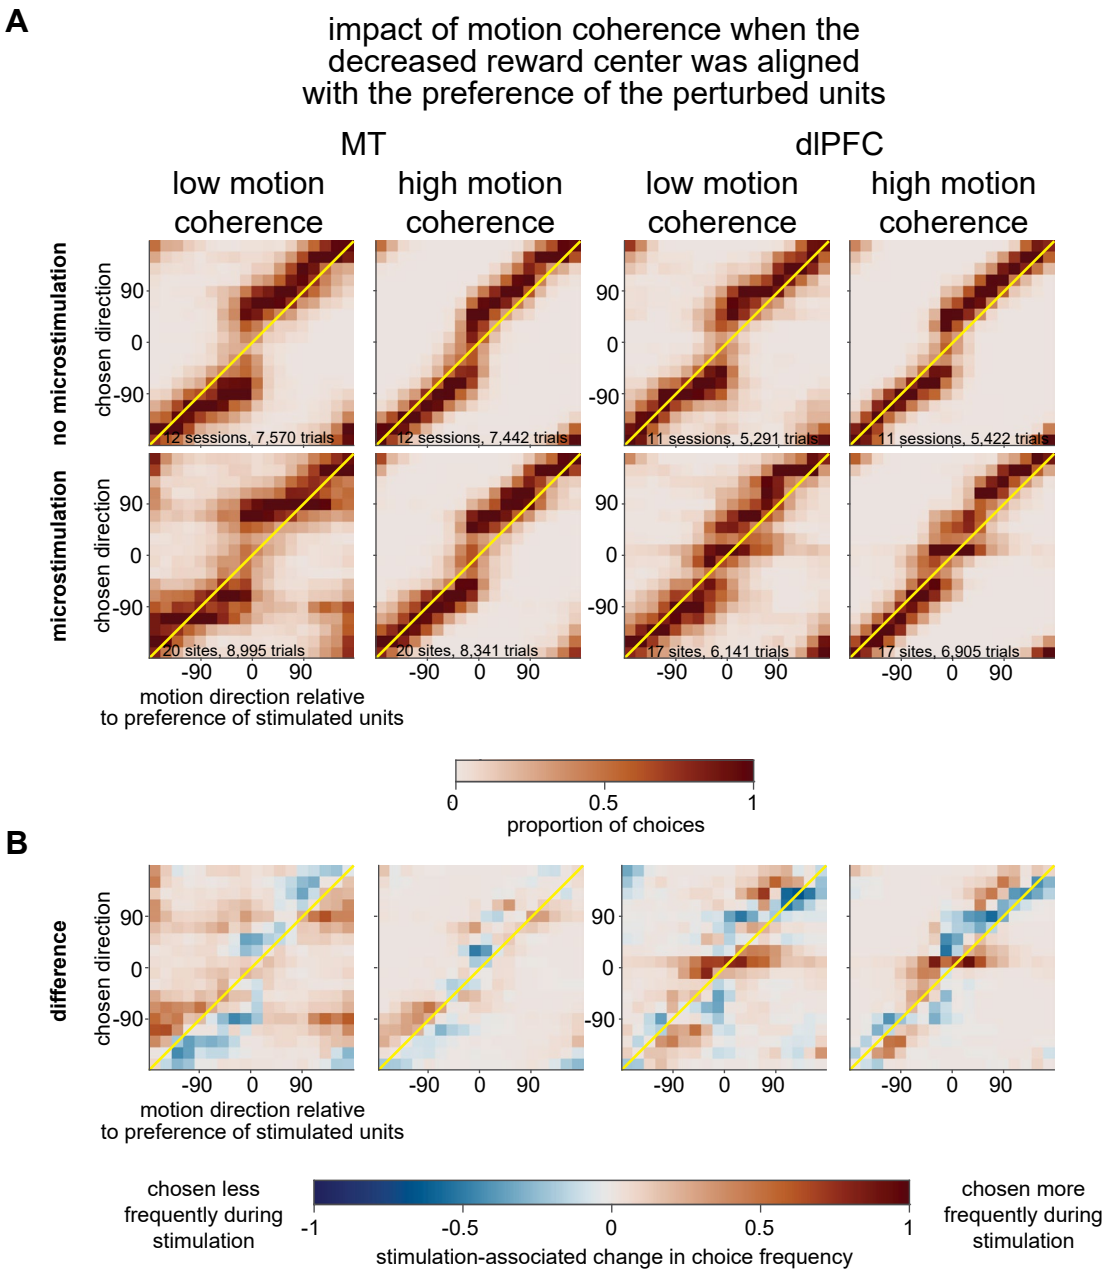
Extended Data Fig. 7 | MT microstimulation led to some winner-take all choices during separate experiments with no reward scaling. A) The task design was identical to the main experiment except that there were no colored portions of the target ring and reward magnitudes were uniform across choices and depended only on motion estimation accuracy. **B**) Summary of choices during 23 experimental sessions (10 sessions from monkey 1; comprising 46 total stimulation sites across both monkeys) of MT microstimulation and low coherence dot motion (coherence less than 20% across experiments). Top row depicts trials with no microstimulation, bottom row depicts trials with microstimulation. **C**) Microstimulation-related difference in choices. The proportion of choices from B during non-microstimulation trials were subtracted from the proportion of choices during microstimulation trials for each reward condition. Choices were more likely to be 'winner-take-all' reflections of the microstimulation (that is, clustered near a horizontal band around of choices around the direction preference of the microstimulated site) than during the main experiment with reward scaling (compare Fig. 4b, c and Extended Data Fig. 8).

https://doi.org/10.1038/s41593-026-02514-7



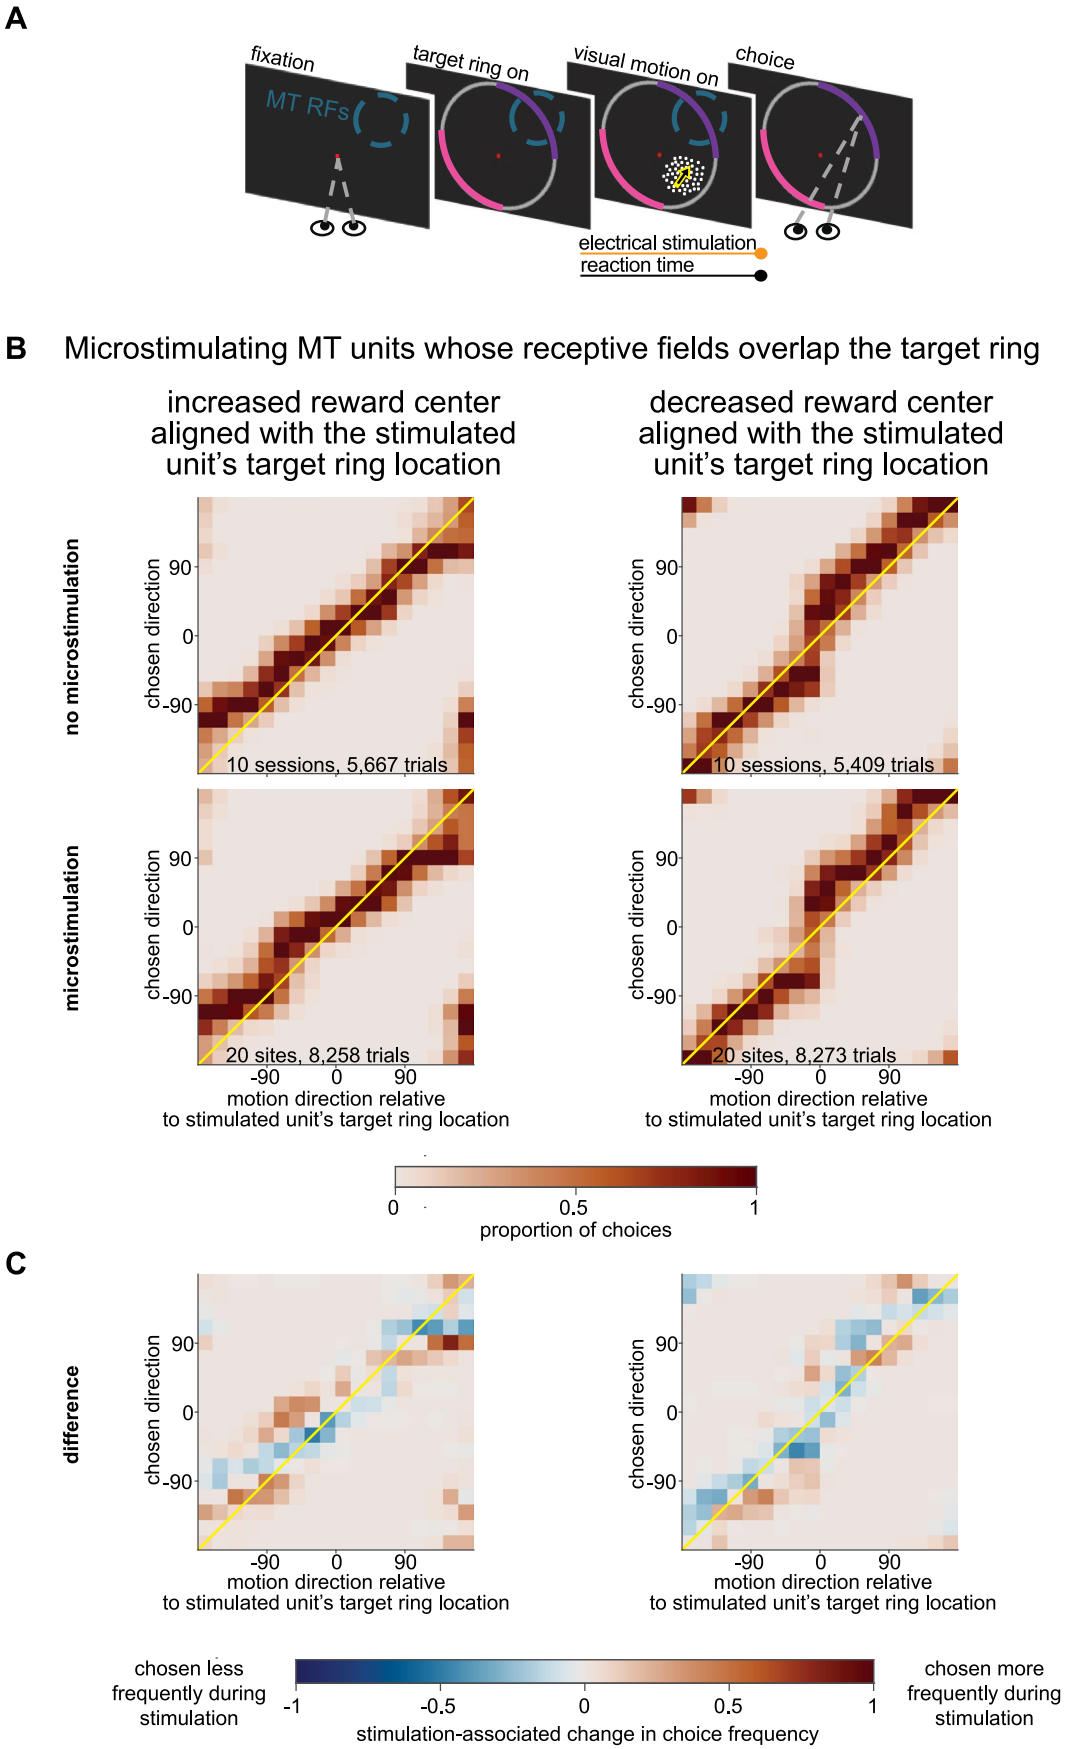
Extended Data Fig. 8 | The magnitude of, but not qualitative difference between, the behavioral impact of MT and dlPFC microstimulation depends on motion strength, on trials when the increased reward center was aligned with the preference of the perturbed units. A) Summary of choices when the increased reward center aligned with the preferred directions of the microstimulated units during 12 experimental sessions (7 sessions from monkey 1; comprising 20 total stimulation sites across both monkeys) of MT microstimulation and 11 sessions (6 sessions from monkey 1; comprising 17 total stimulation sites across both monkeys) of dlPFC microstimulation using two motion coherences per session (low coherences were chosen between 15% and 20% coherence; high coherences were chosen between 35% and 50% coherence). The top row depicts trials with no microstimulation and the bottom row depicts trials with microstimulation. **B**) Microstimulation-related difference in choices in both areas. The proportion of choices from A during non-microstimulation trials were subtracted from the proportion of choices during microstimulation trials at each coherence level. In MT, when the visual motion stimulus was at a lower coherence, microstimulation biased choices more strongly toward the preference of the stimulated site. In dlPFC, when the visual motion stimulus was lower coherence, microstimulation elicited more choices near the preference of the stimulated site.

https://doi.org/10.1038/s41593-026-02514-7



Extended Data Fig. 9 | The magnitude of, but not qualitative difference between, the behavioral impact of MT and dlPFC microstimulation depends on motion strength, on trials when the increased reward center was aligned with the preference of the perturbed units. A) Summary of choices on trials when the decreased reward center aligned with the preferred directions of the microstimulated units. Conventions as in Extended Data Fig. 8A. **B**) Microstimulation-related difference in choices in both areas. Conventions as in Extended Data Fig. 8B.

https://doi.org/10.1038/s41593-026-02514-7



Extended Data Fig. 10 | See next page for caption.

https://doi.org/10.1038/s41593-026-02514-7

Extended Data Fig. 10 | Microstimulation in MT does not bias choices when the target ring, not the visual motion stimulus, overlaps the receptive fields of the microstimulated units. A) As in all stimulation experiments, electrical stimulation coincided with the presentation of the visual stimulus, but unlike in most experiments (e.g. Figure 4), the receptive fields of the stimulated MT units overlapped the target ring, not the visual motion stimulus. **B)** Summary of choices from 10 experimental sessions (5 sessions from monkey 1; comprising 20 total stimulated sites across both monkeys) of MT microstimulation when the receptive fields of the stimulated units overlapped the increased (left) or decreased (right) reward centers. Top row depicts trials with no microstimulation, bottom row depicts trials with microstimulation. **C)** Microstimulation-related difference in choices. The proportion of choices from B during non-microstimulation trials were subtracted from the proportion of choices during microstimulation trials for each reward condition. Choices were not biased by microstimulation in a way that is qualitatively similar to microstimulation in dlPFC (compare Fig. 4b, c), even though in both cases, the receptive fields of the stimulated neurons overlapped part of the target ring.

nature portfolio

Corresponding author(s): Douglas Ruff daruff@uchicago.edu

Last updated by author(s): Aug 4, 2026

Reporting Summary

Nature Portfolio wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Portfolio policies, see our Editorial Policies and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed
☐	☒ The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒ A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒ The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒ A description of all covariates tested
☐	☒ A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒ A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒ For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐ For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐ For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☒	☐ Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Trellis recording software (Ripple, Utah, USA) was used for neuronal data recording.
Data analysis	custom MATLAB and python code was used to analyze data.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Portfolio guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A description of any restrictions on data availability
- For clinical datasets or third party data, please ensure that the statement adheres to our policy

Data and code are available in a repository https://github.com/solkm/formattingRNN

Research involving human participants, their data, or biological material

Policy information about studies with human participants or human data. See also policy information about sex, gender (identity/presentation), and sexual orientation and race, ethnicity and racism.

Reporting on sex and gender	*Use the terms sex (biological attribute) and gender (shaped by social and cultural circumstances) carefully in order to avoid confusing both terms. Indicate if findings apply to only one sex or gender; describe whether sex and gender were considered in study design; whether sex and/or gender was determined based on self-reporting or assigned and methods used.* *Provide in the source data disaggregated sex and gender data, where this information has been collected, and if consent has been obtained for sharing of individual-level data; provide overall numbers in this Reporting Summary. Please state if this information has not been collected.* *Report sex- and gender-based analyses where performed, justify reasons for lack of sex- and gender-based analysis.*
Reporting on race, ethnicity, or other socially relevant groupings	*Please specify the socially constructed or socially relevant categorization variable(s) used in your manuscript and explain why they were used. Please note that such variables should not be used as proxies for other socially constructed/relevant variables (for example, race or ethnicity should not be used as a proxy for socioeconomic status).* *Provide clear definitions of the relevant terms used, how they were provided (by the participants/respondents, the researchers, or third parties), and the method(s) used to classify people into the different categories (e.g. self-report, census or administrative data, social media data, etc.)* *Please provide details about how you controlled for confounding variables in your analyses.*
Population characteristics	*Describe the covariate-relevant population characteristics of the human research participants (e.g. age, genotypic information, past and current diagnosis and treatment categories). If you filled out the behavioural & social sciences study design questions and have nothing to add here, write "See above."*
Recruitment	*Describe how participants were recruited. Outline any potential self-selection bias or other biases that may be present and how these are likely to impact results.*
Ethics oversight	*Identify the organization(s) that approved the study protocol.*

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	Relevant sample sizes are distinct experimental sessions and were estimated using previous literature.
Data exclusions	Sessions were excluded from final inclusion if an insufficient number of behavioral trials were collected or if recording stability was insufficient.
Replication	results were replicated across distinct experimental sessions (days).
Randomization	all relevant experimental conditions were randomly interleaved within daily experimental sessions.
Blinding	blinding is not possible and all experimental conditions are collected within a single session.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology and archaeology
☐	☒ Animals and other organisms
☒	☐ Clinical data
☒	☐ Dual use research of concern
☒	☐ Plants

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other research organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research, and Sex and Gender in Research

Laboratory animals	two male rhesus macaques
Wild animals	lab born macaques.
Reporting on sex	sex is not a relevant variable.
Field-collected samples	NA
Ethics oversight	All surgical and behavioral procedures conformed to the guidelines established by the National Institutes of Health and were approved by the Institutional Animal Care and Use Committees of the University of Pittsburgh and Carnegie Mellon

Note that full information on the approval of the study protocol must also be provided in the manuscript.

Plants

Seed stocks	*Report on the source of all seed stocks or other plant material used. If applicable, state the seed stock centre and catalogue number. If plant specimens were collected from the field, describe the collection location, date and sampling procedures.*
Novel plant genotypes	*Describe the methods by which all novel plant genotypes were produced. This includes those generated by transgenic approaches, gene editing, chemical/radiation-based mutagenesis and hybridization. For transgenic lines, describe the transformation method, the number of independent lines analyzed and the generation upon which experiments were performed. For gene-edited lines, describe the editor used, the endogenous sequence targeted for editing, the targeting guide RNA sequence (if applicable) and how the editor was applied.*
Authentication	*Describe any authentication procedures for each seed stock used or novel genotype generated. Describe any experiments used to assess the effect of a mutation and, where applicable, how potential secondary effects (e.g. second site T-DNA insertions, mosiacism, off-target gene editing) were examined.*