


Representational geometry of perceptual decisions in the monkey parietal cortex

Graphical abstract



Authors

Gouki Okazawa, Christina E. Hatch, Allan Mancoo, Christian K. Machens, Roozbeh Kiani

Correspondence

roozbeh@nyu.edu

In brief

Comparison of neuronal activity during decision-making in face and motion discrimination tasks reveals task-dependent response geometry that shapes optimal readout of the decision variable.

Highlights

- We examine parietal neural activity in motion and face discrimination tasks
- Population neural response dynamics form a curved manifold during decision-making
- Manifolds rotate across tasks, changing optimal readout of the decision variable
- Circuit models of perceptual decisions need revision to include task dependency

Article

Representational geometry of perceptual decisions in the monkey parietal cortex

Gouki Okazawa,¹ Christina E. Hatch,¹ Allan Mancoo,² Christian K. Machens,² and Roozbeh Kiani^{1,3,4,5,*}

¹Center for Neural Science, New York University, New York, NY 10003, USA

²Champalimaud Research, Champalimaud Centre for the Unknown, 1400-038 Lisbon, Portugal

³Neuroscience Institute, NYU Langone Medical Center, New York, NY 10016, USA

⁴Department of Psychology, New York University, New York, NY 10003, USA

⁵Lead contact

*Correspondence: roozbeh@nyu.edu

<https://doi.org/10.1016/j.cell.2021.05.022>

SUMMARY

Lateral intraparietal (LIP) neurons represent formation of perceptual decisions involving eye movements. In circuit models for these decisions, neural ensembles that encode actions compete to form decisions. Consequently, representation and readout of the decision variables (DVs) are implemented similarly for decisions with identical competing actions, irrespective of input and task context differences. Further, DVs are encoded as partially potentiated action plans through balance of activity of action-selective ensembles. Here, we test those core principles. We show that in a novel face-discrimination task, LIP firing rates decrease with supporting evidence, contrary to conventional motion-discrimination tasks. These opposite response patterns arise from similar mechanisms in which decisions form along curved population-response manifolds misaligned with action representations. These manifolds rotate in state space based on context, indicating distinct optimal readouts for different tasks. We show similar manifolds in lateral and medial prefrontal cortices, suggesting similar representational geometry across decision-making circuits.

INTRODUCTION

Perceptual decision-making relies on deliberative processes that evaluate and compare possible alternatives based on sensory information. A commonly suggested mechanism for such decisions is to integrate sensory information into a decision variable (DV), which is compared against a decision criterion or bound to commit to a choice (Green and Swets, 1966; Link, 1992; Ratcliff and McKoon, 2008; Smith and Vickers, 1988). Consistent with these behavioral models, multiple brain structures show activity compatible with representation of the DV (Gold and Shadlen, 2007; Shadlen and Kiani, 2013). For example, neurons in the lateral intraparietal (LIP) area increase their average firing rate in response to sensory evidence supporting their preferred saccade target during a motion direction discrimination task (Bennur and Gold, 2011; Shadlen and Newsome, 2001). The hypothesis that the average LIP firing rate reflects accumulated evidence has successfully accounted for various aspects of decision-making behavior such as choice and reaction time distributions (Churchland et al., 2008; Purcell and Kiani, 2016) as well as choice confidence (Kiani and Shadlen, 2009). Similar encoding of the DV has been found in many other brain structures (Deverett et al., 2018; Ding and Gold, 2010, 2012; Donner et al., 2009; Horwitz and Newsome, 1999; Kiani et al., 2014a; Kim and Basso, 2008; Kim and Shadlen,

1999; de Lafuente et al., 2015; Peixoto et al., 2018; Ratcliff et al., 2003) and in other perceptual tasks (Hanks et al., 2015; Heekeren et al., 2004; Hou et al., 2019; Kumano et al., 2016; O'Connell et al., 2012; Philiastides and Sajda, 2006; Purcell et al., 2010; Thura and Cisek, 2014).

These experimental findings have promoted the proposition that decisions are formed through competition of neuronal modules with specific choice preferences until one module “wins” and dictates the choice. The strongest supporting evidence for this proposition is observed in motor planning regions of the brain, where neurons represent the formation of decisions communicated through their preferred actions. Sensory information biases this competition such that the evidence supporting a choice increases the activity of corresponding neurons. Because the competition happens at the level of action selection, what distinguishes different perceptual tasks is the sensory information being integrated, not the integrator itself. These ideas shape existing theoretical frameworks for implementation of the decision-making process with biophysically realistic neural networks (Beck et al., 2008; Deco et al., 2013; Lo and Wang, 2006; Mazurek et al., 2003; Purcell et al., 2010; Wang, 2002; Wimmer et al., 2015; Wong and Wang, 2006). Such networks have been quite successful in explaining past behavioral and physiological data.

However, several key assumptions at the core of these frameworks have yet to be tested experimentally. First, it is unclear

whether the neural representation of the integration process is shared across tasks and whether potential differences necessitate different readout schemes for downstream circuits. Second, it is unclear whether the DVs are encoded as intermediate response patterns between competing actions: are the pools of neurons competing during decision formation identical to those representing the final actions? Testing these assumptions has been challenging because past studies were largely limited to a single behavioral task. A side-by-side comparison of multiple tasks is necessary for a rigorous test as it creates an opportunity to observe qualitative and quantitative differences of neural activity across tasks.

Here, we examine the geometry of population activity in LIP in seven monkeys and three tasks: a motion direction discrimination task, which has been extensively used to study perceptual decisions, and two variants of a novel face categorization task, where subjects report the species or expression of a face (Okazawa et al., 2018). We find a striking mismatch in patterns of average firing rates among the tasks. Specifically, mean firing rates during the face tasks are inversely correlated with the sensory evidence in favor of the preferred saccade target of neurons—the opposite of the pattern observed in the motion task (Churchland et al., 2008; Kiani et al., 2008; Shadlen and Newsome, 2001). We show that this inconsistency arises because, at each moment, the LIP population encodes the DV on a curved manifold in the state space, and this manifold changes in a task-dependent manner. We establish that the curved manifold provides an explicit code (Misaki et al., 2010) for both the DV and stimulus difficulty but the latter does not have apparent functional contributions to the monkey's confidence about the choice. Further, we show that task dependence of curved manifolds is shaped by task rules and is not explained solely by differences of sensory stimuli. Finally, we demonstrate that similar manifolds emerge in the lateral and medial prefrontal cortices. We propose that the task-dependent curvature of the DV manifold is a fundamental representational property that challenges key assumptions of existing circuit models for perceptual decisions.

RESULTS

Curved population response manifolds in LIP

We examined population response properties in LIP for two variants of a face categorization task (Figure 1A, face task) and a motion direction discrimination task (Figure 1B, motion task). In the motion task, monkeys viewed a dynamic random-dots stimulus and reported the perceived motion direction. We manipulated stimulus difficulty by varying the percentage of coherently moving dots across trials (motion coherence). Within each trial, the stochastic nature of the stimulus caused frame by frame fluctuations of motion coherence around a nominal mean for the trial. In the face task, monkeys classified the species of a set of face stimuli (monkey or human) or their expression (sad or happy) in separate blocks. The discriminated categories in each block were defined by two prototype faces. We manipulated stimulus difficulty by creating a morph continuum between the two prototypes and varying the stimulus morph level across trials. Our customized morphing algorithm allowed independent morphing

of facial features. We chose eyes, nose, and mouth as informative features and morphed them, while fixing other facial features at halfway between the prototypes. The facial features fluctuated randomly around a mean morph level in each trial (SD, 20% morph; updated every 106.7 ms), similar to the within-trial variability of motion in the motion task. These fluctuations, which were interleaved with masks to remain subliminal, enabled us to estimate spatiotemporal weighting of informative features for decision-making (Figure S1). In both tasks, monkeys reported their decision by making a saccadic eye movement to one of the two targets, one placed inside the response field (RF) of the recorded neurons (T_{in}) and the other on the opposite side of the screen, outside the RF (T_{out}).

Three monkeys performed the motion task and two monkeys performed the face task, while we recorded from LIP neurons. In both tasks, the monkey's choice accuracy monotonically improved as a function of stimulus strength (Figures 1C and 1E) and duration (Figures 1D and 1F). The reduction of psychophysical thresholds with longer stimulus durations (face task, log-log regression slope, -0.19 ± 0.02 , $p < 0.001$; motion task, -0.61 ± 0.03 , $p < 0.001$, bootstrap test) indicated that monkeys leveraged multiple stimulus frames for their decisions. These results were quantitatively compatible with drift diffusion models (DDMs) that accumulated sensory evidence toward decision bounds in both tasks (DDM fit to psychometric function, face task, $R^2 = 0.998$; motion task, $R^2 = 0.997$). In both tasks, the model also quantitatively explained the results of psychophysical reverse correlation, including differential effects of sensory evidence on choice over time and differential weighting of informative facial features for species and expression categorizations (Figure S1). Based on the similarity of decision-making behavior between the two tasks, we hypothesized that LIP neurons would exhibit similar decision-related responses in both tasks.

Contrary to our expectation, we found a qualitative difference in the population average LIP responses in the two tasks (Figures 2A and 2B). In the motion task (Figure 2B), neurons had higher firing rates when the stimulus strongly supported the T_{in} choice and lower firing rates when it strongly supported the T_{out} choice. Thus, the firing rates monotonically reflected the amount of evidence for the T_{in} choice. However, LIP neurons recorded in the face task exhibited a reversed order for T_{in} choices (Figure 2A); they showed enhanced activity for weaker sensory evidence. Consequently, the average firing rates no longer monotonically reflected the amount of evidence for the T_{in} choice. This observation is puzzling as it seemingly contradicts the account of LIP responses as a neural correlate of accumulated evidence for a decision. Critically, this pattern also appears to disagree with monkeys' behavioral performance, which monotonically improved as a function of stimulus strength in both tasks (Figure 1).

To quantify the significance of this reversal in firing rates, we plotted the population average firing rates in a window spanning from 250 to 600 ms from stimulus onset as a function of stimulus strength (Figure 2C). For T_{in} choices in the motion task, firing rates increased with the strength of evidence supporting the choice ($\alpha_1 = 0.22 \pm 0.05$ across units, $p = 8.2 \times 10^{-5}$; Equation 1), whereas in the face task, firing rates decreased ($\alpha_1 = -0.19 \pm 0.03$, $p = 3.8 \times 10^{-8}$). This reversal was also seen in many single units and was consistently




Figure 1. Task designs and behavioral results

(A) Face categorization tasks. Monkeys classified face stimuli, reporting either the species (monkey versus human) or expression (sad versus happy) of faces in different blocks. In each task, the stimuli were chosen from a morph continuum between two prototypes. On each trial, following fixation, two targets appeared, one inside (T_{in}) and the other outside (T_{out}) of the neuron RF. Then, a sequence of faces interleaved with masks appeared. During the sequence, morph levels of faces fluctuated around a nominal mean level for the trial, providing noisy sensory evidence. The stimulus was positioned non-foveally and its size was randomly varied across trials by one octave to prevent the monkey from relying on local and low-level visual attributes. After a delay, the fixation point disappeared (Go cue), and the monkey reported its choice by making a saccade to one of the targets.

(B) Motion direction discrimination task. Monkeys viewed a dynamic random-dots stimulus and reported the net motion direction. Stimulus difficulty was controlled by varying the percentage of coherently moving dots.

(C and E) In both the face (C) and the motion (E) tasks, monkeys' accuracy monotonically improved for higher stimulus strengths. Gray lines are the fits of drift diffusion models (see STAR Methods).

(D and F) Psychophysical thresholds (the stimulus strength at 81.6% accuracy) decreased for longer stimulus durations. Error bars are SEM.

See also Figures S1.

present in different monkeys and different task variations (Figures S2A–S2D). In contrast to T_{in} choices, firing rates systematically declined with the strength of evidence in favor of the T_{out} choice during stimulus viewing in both tasks (Figure 2D; motion task, $\alpha_1 = -0.30 \pm 0.05$, $p = 5.0 \times 10^{-9}$; face task, $\alpha_1 = -0.33 \pm 0.03$, $p = 2.9 \times 10^{-20}$). The peri-stimulus time histograms (PSTHs) for the T_{out} choice were better separated for different stimulus strengths (Figure 2B), consistent with past electrophysiological studies (Shadlen and Newsome, 2001) as well as bounded evidence accumulation models for the decision-making process.

Does the reversal suggest different underlying neural computations during face categorization? Because individual units show diverse response profiles (Figure S2B), we sought to better understand the neural code by exploring the patterns of population responses in the two tasks. Figures 2E and 2F depict population activity in a three-dimensional (3D) principal component (PC) space derived from the PSTHs of recorded units (top 3 PCs, explained variance in motion task: 78%, face task: 54%; see Figure 3A for the PCA procedure). The data points in this

state space correspond to the population responses for different stimulus strengths at different times after stimulus onset. The population neural responses gradually changed along a curved line in the state space from the strongest stimulus supporting T_{in} to the strongest stimulus supporting T_{out} choices. The curvature was quite large and unlikely to have been caused by random fluctuation of neural responses (Figure S4A). Further, a bootstrap analysis confirmed that the manifold is curved along a consistent direction in each task ($p < 0.001$). This curvature, which arose from the diverse tuning of neurons for stimulus strength in the population (Figure S3B), was observed in individual monkeys and also found for groups of simultaneously recorded units within single sessions (Figures S4B–S4D). The curved manifold gradually expanded over time after stimulus onset, indicating better separation of population response patterns associated with different stimuli. This expansion is compatible with integration of sensory evidence over time and the representation of the DV by the neural population in both tasks. Later in the trials, immediately before the monkey's saccade, the population responses converged to one of the two states depending on the




Figure 2. Population response patterns form a curved manifold in state space that rotates across tasks giving rise to opposite ordering of population average PSTHs

(A and B) PSTHs in the face and motion tasks showed opposite ordering of firing rates as a function of stimulus strength for T_{in} choices (red lines). Stronger motion toward T_{in} was associated with higher firing rates, matching past studies (Kumano et al., 2016; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001; Shushruth et al., 2018). The face task had a reversed order. Only correct choices were included in the PSTHs to ensure that the order of PSTHs was not shaped by different proportions of choices for different stimulus strengths.

(C) Average firing rates for T_{in} choices (250–600 ms after stimulus onset) increased with the stimulus strength in the motion task but decreased in the face task. Error bars are SEM.

(D) Average firing rates for T_{out} choices decreased with the stimulus strength in both tasks.

(E and F) Population neural responses formed curved manifolds. Principal component analysis (PCA) was performed on trial-averaged PSTHs concatenated across units. In the 3D PC space, each point corresponds to the population neural activity for correct choices to different stimulus strengths at different times after stimulus onset. The gray lines are cubic smoothing splines fit. Error bars are SEM.

(G and H) In the PC space, changes of the population average firing rate correspond to a linear axis (black vector). Depending on the direction of manifold curvature with respect to this axis, changes of population response averages as a function of stimulus strength could become monotonic (motion task, H) or non-monotonic (face task, G). In the 2D illustrations, the projection lines appear non-perpendicular to the averaging axis, but they are perpendicular in the 3D space. See also Figures S2, S3, S4, S5, and S7.

monkey's choice (Figure S5). Correspondingly, the average firing rates reached a common level (Figure S5A), as has been reported previously (Shadlen and Newsome, 2001).

An important implication of the curvature of the population response manifold is that the average firing rate does not necessarily preserve the order of the stimulus strengths. Recall that the average firing rate is calculated as a linear sum of single-neuron firing rates with equal weights across the recorded population. Because projections on the PC space are also linear, changes of the population average firing rate would correspond to a linear axis in the PC space (black vectors in Figures 2G and 2H). By projecting on this axis, one can determine the average firing

rate associated with each point on the manifold. In the face task, the curved manifold was rotated with respect to the average rate axis such that the projections result in a reversed order of average firing rates for different stimulus strengths supporting the T_{in} choice. By contrast, the position of the curved manifold in the motion task did not create such a reversal. Thus, the apparent mismatch of average firing rates between the two tasks was incidental to the position of the manifolds, which changed substantially in the state space for the two tasks. However, the population response manifolds of both tasks shared a common geometry that orderly represented the evidence supporting each choice.

These findings indicate that LIP activity during decision formation is better characterized as a state on a curved manifold than as average firing rate. As we explain in the following sections, this insight leads to important questions about the relationship between neural responses and behavior, as well as the underlying computational mechanisms.

Functional implications of curved manifold for evidence integration

The curved manifold indicates that LIP simultaneously encodes two key variables for the decision-making process: stimulus difficulty and integrated evidence (the DV). Stimulus difficulty, or more generally the experimentally controlled reliability of incoming evidence, is defined as unsigned stimulus strength in both tasks and encoded along the major axis of the curved manifold (Figure 3A). In contrast, the DV is encoded as position along the curved manifold, or alternatively as projections onto a linear axis perpendicular to the major axis of the curved manifold. The explicit encoding of stimulus difficulty contrasts with previous studies, which suggested only an implicit code for stimulus difficulty through the representation of the DV and choice (Kiani and Shadlen, 2009; Pouget et al., 2016). Here, we define explicit as linearly decodable (Misaki et al., 2010) and implicit as readable albeit not linearly. Our results raise an intriguing possibility that the curved manifold may contribute to both the choice and the confidence associated with the choice, because it enables linear readouts of both the DV and stimulus difficulty by downstream areas.

To examine this prediction, we first identified the best linear axes in the state space for reading out the DV and stimulus difficulty, then we quantified the sufficiency of these axes for predicting choice and stimulus difficulty, respectively, and finally, we tested whether projections along the stimulus difficulty axis were predictive of the monkey's confidence. In both the motion and face tasks, the DV for a fixed stimulus duration is proportional to the signed stimulus strength (Okazawa et al., 2018, 2020; Palmer et al., 2005). Further, in both tasks, difficulty is defined by the absolute stimulus strength. We used an orthogonal canonical correlation analysis (CCA) (Cunningham and Ghahramani, 2015) to find a pair of linear axes in the state space, one most correlated with the signed stimulus strength ("DV axis") and another most correlated with the unsigned stimulus strength ("stimulus difficulty axis"). These two axes provided a two-dimensional (2D) perspective of the population neural responses, which reproduced the curved manifold in both tasks (Figure 3B). The weights of neurons in the two axes were unimodally distributed, with no indication that distinct sub-populations were engaged in coding the DV and difficulty. The DV and difficulty encoding axes were stable throughout the decision formation period, reflecting the stability of curved manifolds in the 3D PC space (Figures 2E and 2F).

Critically, the neural responses projected on the DV and difficulty axes were predictive of the choice and stimulus difficulty, respectively (Figures 3C and 3D). To ensure that these predictions were not artificially introduced by our analysis, we adopted a cross-validation approach, where we first derived the DV and difficulty axes from a random half of the trials and then used the other half for the predictions. Projections along the DV axis

gradually diverged for T_{in} and T_{out} choices during stimulus viewing (Figure 3C; $p < 0.001$ in both tasks, bootstrap test of neural responses 350–450 ms after stimulus onset). Further, when the monkey made an error, the neural responses followed a trajectory opposite to that of correct trials, indicating that the DV axis affords a representation of the DV that shapes the choice, not a representation of the objective stimulus. Finally, the neural responses projected on the stimulus difficulty axis distinguished easy and difficult stimuli (Figure 3D; $p < 0.001$), but they did not predict the monkey's choices (Figure 3E; T_{in} versus T_{out} , face task, $p = 0.15$; motion task, $p = 0.46$; 350–450 ms after stimulus onset, bootstrap test). The inability to predict the monkey's performance suggests that the modulation of population responses along the difficulty axis does not reflect potential increase of the monkey's attention or task engagement for difficult stimuli.

Does this stimulus difficulty encoding play a role in confidence judgments? We examined this possibility by analyzing neural responses while monkeys performed the motion task with post-decision wagering (Kiani and Shadlen, 2009; Middlebrooks and Sommer, 2012). In this version of the task, a third target (sure target, T_s) appeared during the delay period on half of the trials (Figure 4A). By choosing this target after the Go cue, the monkey could opt out of direction discrimination for a small but guaranteed reward. It has been previously verified that monkeys choose T_s based on their certainty, selecting it more frequently on difficult trials (Figure 4B) (Kiani and Shadlen, 2009; Komura et al., 2013; Odegaard et al., 2018). Further, previous studies demonstrated that the neural representation of the DV is predictive of the monkey's T_s choices (Kiani and Shadlen, 2009). However, if the neural mechanisms of confidence rely on explicit encoding of stimulus difficulty by the LIP population, projections along the difficulty axis in the state space would also be predictive of T_s choices. In fact, it is even possible that previous reports of the relationship between the DV representation and T_s choices were merely side effects of the joint encoding of the DV and stimulus difficulty in the same curved manifold, because changes of stimulus strengths across trials modify projections along both the DV and difficulty axes.

Although projections along the stimulus difficulty axis distinguished difficult and easy trials (Figure 3D), these projections were not predictive of the monkey's confidence (Figure 4). To illustrate this point, Figure 4C shows the projection of neural responses in the 2D space defined by the DV and difficulty axes at an example intermediate time during decision formation (400 ms after stimulus onset). The DV axis in this figure is defined with respect to the correct choice, with positive and negative values indicating the DVs in favor of the correct and error choices, respectively. Whereas different choices were associated with distinct projections along the DV axis, projections along the stimulus difficulty axis did not distinguish low-confidence T_s choices from the higher-confidence choices where the monkey chose one of the direction targets (correct and error choices).

To quantitatively test these observations, we asked whether, for a fixed stimulus strength, residual projections along the DV and difficulty axes were predictive of confidence over the course of decision formation (Figures 4D and 4E). The residuals were computed separately for each stimulus strength. Then, we combined the residuals of different stimulus strengths for each




Figure 3. LIP population responses jointly encode the DV and stimulus difficulty

(A) Curved manifolds jointly encode the DV and stimulus difficulty in the state space. r_i indicates the activity of unit i .

(B) Orthogonal canonical correlation analysis (CCA) identifies the linear axes that encode the DV and stimulus difficulty. All analyses were cross-validated; we used half of trials to identify the axes and the other half to create the projections. Data points are projected responses for different stimulus strengths. The gray curve is a second-order polynomial fit to the data points. The plot shows responses 400 ms after stimulus onset but similar patterns were also observed at other times (300–600 ms). Error bars are SEM.

(C) The population neural responses projected on the DV axis for each stimulus category (T_{in} or T_{out}) and for each choice outcome (correct or error). The responses were predictive of the monkey's forthcoming choices. Shading indicates SEM.

(D) The population responses projected on the stimulus difficulty axis for easy ($\geq 20\%$ stimulus strength) and difficult ($< 20\%$) stimuli.

(E) The monkey's choices for a particular stimulus strength were not correlated with projections of population responses on the stimulus difficulty axis. Residual responses were calculated for each choice and stimulus strength and then averaged across the stimuli.

choice. This residual analysis addresses the stimulus difficulty confound that causes joint variations along the DV and difficulty axes, because it isolates the effect of variations along one axis while keeping projections on the other axis constant. The residual projections along the DV axis showed a clear separation for the three choices, with the residuals for T_s choices in between those for correct and error choices (Figure 4D; correct versus T_s , $p < 0.001$; error versus T_s , $p < 0.001$, bootstrap test, 350–450 ms after stimulus onset). Thus, the DVs closer to zero were associated with T_s choices and errors occurred because the DV supported the opposite choice (Figures 3C and 4C), consistent with previous reports (Kiani and Shadlen, 2009). In contrast, residual projections along the difficulty axis failed to distinguish different choices (Figure 4E; correct versus T_s , $p = 0.52$; error versus T_s , $p = 0.17$), indicating that this axis did not have noticeable bearing on the monkey's choice and confidence. These results suggest that the functional role of the curved manifold is primarily to encode the DV.

Task dependency of the DV encoding

The observed geometry of the DV encoding has important implications for circuit models of perceptual decision-making. Existing theories assume action-selective neural modules compete to implement evidence integration (Beck et al., 2008; Deco et al., 2013; Lo and Wang, 2006; Mazurek et al., 2003; Purcell et al., 2010; Wang, 2002; Wimmer et al., 2015; Wong and Wang, 2006). In these models, different perceptual tasks are implemented by changing the input to the competing modules, not by changing the competition process itself. However, the position of the DV manifold largely differed between the face and motion tasks (Figures 2G and 2H), suggesting that the DV encoding

is also task-dependent. Revealing the nature of this task dependency—whether it arises due to differences in sensory stimuli or task rules—and its impact on the readout of the DV is important for understanding the properties of decision-making circuits.

We therefore analyzed a subset of neurons ($n = 28$) recorded in two face categorization tasks (species and expression) to quantify task dependency. The face stimuli in both contexts could vary along both species and expression morph axes (Figure 5A). The stimuli around the center of this 2D space were presented in both categorization tasks (gray box in Figure 5A). Thus, the monkeys applied two different categorization rules to the same visual stimuli. Their behavior indeed indicates that they successfully distinguished the two categorization rules without significant interference from the orthogonal rule (Figure 5B).

We found that the change in task rules alone is sufficient to change the neural manifold. In both face tasks, PSTHs showed a reverse order with stimulus strength (Figure S2D), and there was significant curvature in the population response manifolds (Figure 5C; $p < 0.001$). Critically, the manifolds formed by the neural activity within the species and expression categorization tasks were distinct from each other (Figure 5C; different coefficients of polynomial fits to the manifolds, $p < 0.001$, likelihood-ratio test). These manifold differences were apparent in population response patterns for the same stimuli in the two tasks, evident in the projections of population responses on the top three PCs (Figure 5D; $p < 0.001$, permutation test). Projections on the first PC were most striking, because they revealed the distinct category boundaries of the two tasks (top row of Figure 5D; a vertical boundary in the species task and a horizontal boundary in the expression task).




Figure 4. Explicit encoding of stimulus difficulty in LIP is not predictive of confidence judgments

(A) We measured the monkey's confidence about motion direction using a post-decision wagering task. A sure target (T_s) appeared during the delay period on a random half of trials. The monkey could choose T_s to opt out of direction discrimination for a small but guaranteed reward.

(B) Monkeys chose T_s more frequently for weaker motion strengths.

(C) Population neural responses at 400 ms after stimulus onset (100-ms window). The sign of the DV axis is redefined with respect to the correct target such that positive DV indicates neural support for the correct choice. The saturation of the dots indicates stimulus strength. For error and T_s choices, we show only the three weakest stimulus strengths, where these choices were present in all recording sessions. Error bars are SEM.

(D) Residual projected population responses on the DV axis were predictive of the monkey's choices. The residuals were computed by subtracting the mean projection for each stimulus strength and then combining them across different stimulus strengths for each choice. The shading indicates SEM.

(E) Residual projected population responses on the stimulus difficulty axis were not predictive of the monkey's choices.

The differences in the manifolds caused the linear DV axes derived from CCA to be different in the two task contexts. The two DV axes projected onto the 3D state space had a $+19.0^\circ \pm 6.0^\circ$ angle (signed with respect to the axis of average firing rate), which was significantly different from 0° (Figures 5E and 5F; $p < 0.001$, bootstrap test; the angle difference grew larger in higher dimensional state spaces). Consequently, when a DV decoder optimized for one categorization task was used for de-

coding the DV in the other task, a substantial readout bias emerged (Figure 5H; $\Delta\gamma_0 = 10.3 \pm 1.1$ in Equation 9 for decoding with the same- and different-context DV axes at 400 ms, $p = 2.5 \times 10^{-4}$, likelihood-ratio test).

Although the neurons recorded in both face categorization tasks provided the best test for changes of manifolds with task context, the difference could also be inferred from the units recorded only in one categorization task ($n = 104$). For those units, we computed the angle between the DV axis and the average firing rate axis for each task (ω_{spc} and ω_{exp} in Figure 5E). If an area encodes the DV in a task-independent manner, these angles would be identical. However, we found them to be significantly different (Figure 5G; Monkey A, $p < 0.006$; Monkey L, $p < 0.006$, bootstrap test).

Together, the manifolds that encode the DV changed in the state space depending on task contexts, and these changes necessitated the optimal readout axis of the DV to be also task-dependent. In simple decision-making models based on competition of action-selective modules, the balance of activity between the modules encodes the DV, and as such, the same DVs across tasks are associated with similar relative activity between the action-selective units. In contrast, the population neural code in LIP is versatile, reflecting the decision context in a manner that suggests distinct encoding of the DV and action.

Distinct encoding of the DV and action plan

Further support for distinct encoding of the DV and action plan comes from comparing the population responses during decision formation and saccade preparation (Figures 6A–6D). During saccade preparation, responses coalesced into one of two states corresponding to the two choices (Figures 6A and 6C, right, and S5 depicts average firing rates). However, the axis connecting the two action states was misaligned with the curved DV manifold (Figures 6B and 6D). The angle between the two axes (face task, $34.4^\circ \pm 4.0^\circ$; motion task, $11.7^\circ \pm 1.6^\circ$ in 3D space) was significantly greater than those expected from the variability of neural responses ($p < 0.001$, bootstrap test; the expected angles based on response variability were estimated by comparing the choice axes between random halves of trials, face task: $2.9^\circ \pm 1.6^\circ$, motion task: $2.5^\circ \pm 1.3^\circ$).

The misalignment can also be demonstrated by comparing the firing rates between units preferring different choices (T_1 and T_2 , Figure 6E). In the motion task, the differential firing rates were strongly correlated with stimulus strength (slopes of firing rate difference as a function of stimulus strength for T_1 choice: 9.46 ± 1.10), whereas in the face task, they were much less dependent on stimulus strength (slope: 1.61 ± 0.62). $T_{in} - T_{out}$ differential firing rates followed a similar pattern (Figure 6F). In the face task, a significant portion of neurons (27.3%) showed non-positive slopes of $T_{in} - T_{out}$ firing rates as a function of stimulus strength (Figure 6G), larger than the proportion observed in the motion task (9.3%, $p = 0.0002$, Fisher's exact test). Accordingly, ideal observer accuracy for discriminating T_{in} and T_{out} stimuli based on single unit firing rates was much poorer in the face task (Figure 6H; slope of AUC as a function of stimulus strength, face task: 0.047 ± 0.007 , motion task: 0.255 ± 0.019 , $p = 2.1 \times 10^{-21}$).

Congruent with these results, the weights of a large portion of individual neurons in the DV axis were negative for the face task




Figure 5. Task dependency of the population response manifold

(A) The stimuli used in our two face tasks can be visualized in a 2D space. The stimuli at the central region of this space (gray box) appeared in both tasks.
 (B) Performance in both tasks with respect to each categorization rule. Near chance performance for the irrelevant categorization rule (e.g., magenta dots in the species task) indicates that monkeys correctly ignored the irrelevant stimulus fluctuations. Lines are drift diffusion model fits. Error bars are SEM.
 (C) The neural manifolds were distinct in the two tasks. The common PC space was derived from the combined neural responses. The plot was generated using a subset of units ($n = 28$) recorded in both tasks.
 (D) The same stimuli elicited different response patterns in the two tasks. The heatmaps show projections of population responses on the top three PCs for the stimuli in the central region of the morph plane (gray box in A). We divided this region into 16 sub-regions for better visualization.
 (E) Illustration of angles between the DV-encoding axes of the two manifolds (θ), as well as their angles with the average rate axis (ω_{spc} and ω_{exp}). The DV-encoding axes were computed using CCA. Manifolds are identical to those in (C) but from a different perspective.
 (F) The angle of DV encoding axes in the two tasks was significantly greater than zero. Error bars are SEM.
 (G) Angles between the DV encoding and averaging axes were significantly different. Unlike θ , the calculation of ω could be done using the units recorded only in one task. Here, we show results from those units so that (F) and (G) provide independent support for our conclusion. Asterisks indicate significant difference ($p < 0.006$, bootstrap test).
 (H) Using the DV-encoding axis of one task to read out the DVs in the other task would lead to substantial biases. The plots show neural responses of expression categorization task projected on the DV encoding axis of the same task or the species categorization task. The fitting curves are logistic functions (Equation 9).

(24.2%, compared to 10.1% in the motion task, $p = 0.003$, Fisher's exact test). A negative weight indicates a qualitative mismatch between the saccade selectivity of the neuron and its contribution to the encoding of the DV. However, the negative weights were less prevalent than expected if the DV and action codes were uncorrelated (50% incongruence; $p = 1.2 \times 10^{-9}$, binomial test), suggesting partial but imperfect correspondence between the population codes for the DV and action.

Based on these findings, we suggest that there could be a non-monotonic transformation of the neural activity from the DV en-

coding in early phases of decision formation to encoding of a planned action at later times in the trial (Figure 6). In other words, the encoding of the DV in LIP is context-dependent and not tightly bound to the actions associated with the decisions. At later times, however, the DV representation gives rise to the representation of the upcoming saccade and is ultimately supplanted by it.

Encoding of the DV in lateral and medial prefrontal areas

Our analyses thus far focused on LIP neural activity. However, decision-related signals encoding the DV are prevalent




Figure 6. Distinct encoding of the DV and action plan

(A and C) Left: population responses at 400 ms after stimulus onset and the average PSTHs aligned to stimulus onset. Right: population responses at 50 ms before saccade onset and the average PSTHs aligned to the saccade onset.

(B and D) Distinction of the DV and action encoding is visualized in the 3D state space by comparing the two manifolds during decision formation (solid gray line) and saccade preparation (dashed black line). Two different perspectives are shown to illustrate the 3D configuration.

(E) The difference in firing rates of T_1 and T_2 selective neurons created by constructing pseudo-populations with balanced preference for the two choices (see STAR Methods). The difference was larger for stronger stimuli in the motion task but largely similar across stimulus strengths in the face task.

(F) The difference in firing rates of T_{in} and T_{out} trials.

(G) Slope histograms for the difference of T_{in} and T_{out} firing rates as a function of stimulus strength (β_1 in Equation 2).

(H) Ideal observer accuracy for discriminating T_{in} and T_{out} stimuli with equal strength. Accuracy at each time was quantified using the area under curve (AUC) of the receiver operating characteristic analysis applied to individual unit responses. Shading indicates SEM across units.

(I) Schematic of the transformation of neural code. Neural activity during decision formation encodes the DV in a manner partly dissociated from the action selectivity of neurons (black curve in the inset is angled but not orthogonal to the action axis). As the decision-making process progresses, the DV is transformed to a choice representation through changing the activity of neurons according to their action selectivity.

See also Figures S5 and S6.

across multiple brain structures (Deverett et al., 2018; Ding and Gold, 2010, 2012; Horwitz and Newsome, 1999; Kiani et al., 2014a; Kim and Shadlen, 1999; de Lafuente et al., 2015). If the curved manifold is a representational property of decision formation, areas other than LIP are also expected

to have this response geometry. Alternatively, the manifold curvature could be specific to LIP and stem from its circuit and functional properties. To test these possibilities, we examined the responses of two other cortical areas in the frontoparietal network involved in oculomotor decisions




Figure 7. Curved population response manifolds in the lateral and medial prefrontal cortices

(A) PSTHs of prearcuate gyrus (PAG, Utah array recording) and supplementary eye field (SEF, linear probe recording) neurons during a reaction time version of the motion task. Both areas showed monotonically increasing responses for T_{in} choices with larger stimulus strengths, similar to LIP.

(B) Population responses in both areas formed a curved manifold as in LIP. The plots show the manifold at 400 ms after stimulus onset. Error bars are SEM.

during the motion task: the supplementary eye field (SEF) and the prearcuate gyrus (PAG).

Both areas revealed similar response geometries as in LIP (Figure 7). Critically, population responses, visualized using the same PCA procedure as for LIP, formed a curved manifold in both areas (Figure 7B; significance of curvature, $p < 0.001$), with a monotonic arrangement of different stimulus strengths along the manifold. In these recordings, monkeys performed a reaction time variant of the task, where they indicated their decision as soon as they were ready. In the reaction-time task, the monkey did not wait for a Go cue, and the trial terminated as soon as the monkey committed to a choice. Hence the results confirmed that the curvature of the manifold is indeed a property pertaining to decision formation, not a post-commitment representation and not a persistent representation of choice. Together, the results across the frontoparietal network indicate that the curved DV manifold is a ubiquitous property of the brain regions involved in decision formation.

DISCUSSION

We show that the average firing rates of LIP neurons do not change monotonically with the DV in the face task, in contrast to the motion task (Figure 2A). In state space, the neural popula-

tion represents the DV along a curved manifold in both tasks (Figures 2E and 2F). However, the manifolds are rotated and shifted with respect to the axis of average firing rate across tasks, causing the disparate ordering of PSTHs (Figures 2G and 2H), and highlighting the conceptual gap between average firing rates and the DV (Raposo et al., 2014; Schall, 2019). Population level analysis was key to elucidating task-dependent encoding of the DV because a dependency could be obscured in average firing rates (Hou et al., 2019; Kumano et al., 2016). The task-dependent changes of the DV manifold are not simply due to changes in sensory stimuli. Rather, they occur even when the same stimuli were used with different task rules (Figure 5). Critically, the rotation of DV manifolds implies that the readout of the DVs—a necessity for action planning and circuit-circuit coordination in the decision-making network—must take task context into account. In the absence of such “context-aware readouts,” there could be substantial inaccuracies manifesting as choice bias (Figure 5H).

The task-dependent encoding of the DV despite identical actions (saccades) across tasks also suggests partial dissociation of neural codes for the DV and actions. This conclusion is further supported by opposing selectivity of a sizeable fraction of neurons for the DV and actions, as well as misalignment of the DV and action encoding axes in population responses. Our results, however, do not indicate that decision formation takes place in circuits separate from action planning. Rather, the same LIP population that initially encodes the DV on a curved manifold gradually changes its response patterns to encode actions later in the trial (Figures 6 and S5). This sequential use of the same neural population for decision-making and action planning matches prior experimental observations in frontal cortices (Aoi et al., 2020; Kaufman et al., 2015; Peixoto et al., 2021), especially progressive recruitment of choice representation (Peixoto et al., 2018). The partial dissociation between the DV and action planning necessitates network models in which the same neural circuits undergo a transformation of population response patterns over time (Figure 6I).

Our finding that similar response geometries and encoding dynamics are present in LIP, as well as lateral and medial prefrontal cortices (Figure 7), suggests ubiquitous principles that shape the neural code throughout the cortical nodes of the decision-making network. Our conclusions, therefore, transcend present debates about the causal role of LIP in perceptual decisions (Jeurissen et al., 2019; Katz et al., 2016; Shushruth et al., 2018; Zhou and Freedman, 2019). The presence of a similar response geometry in multiple nodes of the saccade planning network supports the possibility that decisions are formed through interactions within a distributed network encompassing multiple brain regions. Pertinently, neural responses in the superior colliculus, caudate, and frontal eye fields have been occasionally reported to vary inversely or non-monotonically with stimulus strength (Ding and Gold, 2010, 2012; Horwitz and Newsome, 2001), hinting at the potential presence of curved DV manifolds in those regions as well.

Why does the curvature exist ubiquitously in all the tasks and brain regions studied here? We suggest that curved manifolds arise automatically due to fundamental constraints on neural computations (Figure S6). The first set of constraints limit the

dynamic range of the firing rates (and hence projections along the encoding axes) that constitute the neural code: the non-negativity of firing rates puts a lower bound on the dynamic range, and various costs associated with spiking (e.g., metabolic costs) (Lennie, 2003) curtail firing rates from above. An energetically efficient code would seek to lower the overall firing rate of neurons, bending the manifold near the non-negativity bound (Keemink and Machens, 2019) (Figure S6C). The second set of constraints arises from the need for a precise code. Encoding a continuous decision variable within a limited dynamic range could become more precise with a curved manifold than a straight line, as the difference of nearby DV values could be encoded with larger, more distinguishable changes in population responses (Figure S6C). However, the exact precision gained by a curved manifold depends on the correlation structure of the population (Bondy et al., 2018; Moreno-Bote et al., 2014) and the feasibility of decoding along the curved manifold instead of a linear readout. Fortunately, decoding along a low-order polynomial in state space, similar to those in our study, is not prohibitively challenging. Further, there is no strong reason to favor a linear readout over other simple readouts (Ritchie et al., 2019). However, the complexity of the manifold cannot increase in an unbounded fashion as the resulting neural code may become susceptible to decoding errors or become too tangled to be easily readable (Figure S6E). We propose that for each task, the sensory and contextual inputs define the sub-region of the state space where the DV can be encoded in, and the constraints above shape a curved manifold that accommodates adequately precise encoding of the DV and an easy readout.

The curved manifold indicates that the LIP population activity explicitly encodes both stimulus difficulty (i.e., reliability of sensory evidence) and the DV (Figure 3A), a departure from past studies that reported only implicit encoding of stimulus difficulty by single neuron responses in LIP (Kiani and Shadlen, 2009; Pouget et al., 2016). This encoding of stimulus difficulty cannot be explained away as an increased attention level or task engagement for more difficult stimuli. Two key observations corroborate our conclusion. First, stimulus difficulty is encoded from the very beginning of the representation of the decision-making process (~200 ms after stimulus onset) (Figure 3D), making it unlikely to be a reactive recruitment of engagement mechanisms due to perceived stimulus difficulty. We remind that stimulus difficulty is not predictable prior to the stimulus onset and inferring the difficulty of a stimulus is time-consuming (Khalvati et al., 2020). Second, and more critically, higher engagement should lead to better behavioral performance, but the activity along the difficulty encoding axis was uncorrelated with the monkey's accuracy (Figure 3E). Moreover, trial-to-trial variations of the encoded difficulty for the same stimulus strength did not correlate with the monkey's confidence in a post-decision wagering task (Figure 4), providing a clear example that decodability does not necessarily amount to functionality (Ritchie et al., 2019). This apparently non-functional encoding of stimulus difficulty lends support to the hypothesis that the manifold curvature is a geometric property of population responses for computing or monitoring changes of the DV, as explained above.

We further rule out that the curvature of the DV manifold and encoding of stimulus difficulty arise from sensory responses or

spatial attention to stimuli. Parietal neurons respond to a variety of stimulus attributes in their RFs (Bisley et al., 2004; Janssen et al., 2008; Lehky and Sereno, 2007; Sarma et al., 2016) and also show strong response modulations depending on the location of spatial attention (Bisley and Goldberg, 2003). However, these response properties do not explain the reversal of firing rates. First, our recording from face-selective regions in the inferior temporal (IT) cortex during the face tasks has revealed that sensory responses are monotonically modulated by the stimulus strength (Figure S7), just like the monotonic tuning to random dots stimuli in the middle temporal (MT) area (Britten et al., 1993). Therefore, the reversal was not inherited from response properties in sensory areas. Second, our sensory stimuli were presented outside the LIP RFs. Increased sensory responses or spatial attention to more difficult stimuli would not affect or rather decrease the activity of the recorded LIP neurons, leading to a response pattern opposite to the reversed order in the face task. To further ascertain that the observed LIP responses were not affected by inadvertent overlaps of the stimuli with RFs, we performed a control experiment in our face task, where we showed the stimuli in the visual hemifield opposite to the LIP RFs. The results replicated the same reversal of firing rates and curved manifolds as in the main task (Figures S2J–S2L). It should also be noted that past studies have reported similar LIP response patterns regardless of whether motion stimuli were presented foveally or parafoveally (Huk and Shadlen, 2005; Shushruth et al., 2018), indicating that differences in stimulus position across tasks do not explain our results.

A hallmark of the decision-making process is its flexibility. To implement flexible decisions, neural circuits must accommodate context-dependent interpretation of information, intricate rules, and a large variety of motor actions to interact with the outside world. The curvature and context-dependency of the DV manifolds and the partial separation of the DV and action encoding, as we report here, are some of the building blocks that the primate brain may have devised for flexible and accurate decisions based on sensory information.

Limitations of the study

The present study uses the PSTHs of neurons recorded over different sessions to infer the geometry of population responses. We therefore do not directly quantify the effect of population noise correlations, which are known to influence the optimal encoding and decoding axes in large neural populations (Moreno-Bote et al., 2014; Rummyantsev et al., 2020). Although we observe the population response geometry and its task dependency robustly across datasets and small pools of neurons recorded simultaneously in the same session (~10 neurons) (Figure S4C), much larger-scale neural recordings from LIP are needed to quantify the true effect of correlations across the population.

One of our key findings is task-dependent changes of the curved DV manifolds. However, we have limited understanding of how this task-dependency is implemented. Future experimental and computational modeling studies are needed to address this gap in our knowledge. It would also be important to examine how widely similar task dependencies exist outside LIP.

STAR★METHODS

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

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **EXPERIMENTAL MODEL AND SUBJECT DETAILS**
 - Monkeys
- **METHOD DETAILS**
 - Behavioral tasks
 - Face task
 - Motion task
 - Neural recording
- **QUANTIFICATION AND STATISTICAL ANALYSIS**
 - Behavioral data analysis
 - Neural data analyses
 - Principal component analysis of neural data
 - Locating the DV and stimulus difficulty axes in the population response state space
 - Testing task dependency of neural responses

ACKNOWLEDGMENTS

The authors would like to thank Alex Pouget, Long Sha, Stanley Kombar, Michael Waskom, Saleh Esteki, Bianca Sieveritz, Mike Shadlen, and Stefano Fusi for helpful discussions and comments on earlier versions of the manuscript. This work was supported by the Simons Collaboration on the Global Brain (542997 and 543009), McKnight Scholar Award, Pew Scholarship in the Biomedical Sciences, and National Institute of Mental Health (R01 MH109180-01). G.O. was supported by post-doctoral fellowships from the Charles H. Revson Foundation and the Japan Society for the Promotion of Science. C.E.H. was supported by National Eye Institute Training Grant in Visual Neuroscience (T32 EY007136).

AUTHOR CONTRIBUTIONS

G.O. and R.K. conceived the study and wrote the manuscript. G.O., C.E.H., and R.K. collected the data. G.O. performed data analysis. A.M. and C.K.M. contributed to interpreting results and developed models. R.K. supervised and guided the study. All authors reviewed and edited the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: June 8, 2020

Revised: December 23, 2020

Accepted: May 17, 2021

Published: June 24, 2021

REFERENCES

Adelson, E.H., and Bergen, J.R. (1985). Spatiotemporal energy models for the perception of motion. *J. Opt. Soc. Am. A* 2, 284–299.

Ahumada, A.J. (2002). Classification image weights and internal noise level estimation. *J. Vis.* 2, 121–131.

Aoi, M.C., Mante, V., and Pillow, J.W. (2020). Prefrontal cortex exhibits multi-dimensional dynamic encoding during decision-making. *Nat. Neurosci.* 23, 1410–1420.

Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., and Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron* 60, 1142–1152.

Bennur, S., and Gold, J.I. (2011). Distinct representations of a perceptual decision and the associated oculomotor plan in the monkey lateral intraparietal area. *J. Neurosci.* 31, 913–921.

Bisley, J.W., and Goldberg, M.E. (2003). Neuronal activity in the lateral intraparietal area and spatial attention. *Science* 299, 81–86.

Bisley, J.W., Krishna, B.S., and Goldberg, M.E. (2004). A rapid and precise on-response in posterior parietal cortex. *J. Neurosci.* 24, 1833–1838.

Bondy, A.G., Haefner, R.M., and Cumming, B.G. (2018). Feedback determines the structure of correlated variability in primary visual cortex. *Nat. Neurosci.* 21, 598–606.

Brainard, D.H. (1997). The psychophysics toolbox. *Spat. Vis.* 10, 433–436.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* 10, 1157–1169.

Churchland, A.K., Kiani, R., and Shadlen, M.N. (2008). Decision-making with multiple alternatives. *Nat. Neurosci.* 11, 693–702.

Cunningham, J.P., and Ghahramani, Z. (2015). Linear dimensionality reduction: Survey, insights, and generalizations. *J. Mach. Learn. Res.* 16, 2859–2900.

de Lafuente, V., Jazayeri, M., and Shadlen, M.N. (2015). Representation of accumulating evidence for a decision in two parietal areas. *J. Neurosci.* 35, 4306–4318.

Deco, G., Rolls, E.T., Albantakis, L., and Romo, R. (2013). Brain mechanisms for perceptual and reward-related decision-making. *Prog. Neurobiol.* 103, 194–213.

Deverett, B., Koay, S.A., Oostland, M., and Wang, S.S. (2018). Cerebellar involvement in an evidence-accumulation decision-making task. *eLife* 7, e36781.

Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. *J. Neurosci.* 30, 15747–15759.

Ding, L., and Gold, J.I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex* 22, 1052–1067.

Donner, T.H., Siegel, M., Fries, P., and Engel, A.K. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Curr. Biol.* 19, 1581–1585.

Fuchs, A.F., and Robinson, D.A. (1966). A method for measuring horizontal and vertical eye movement chronically in the monkey. *J. Appl. Physiol.* 21, 1068–1070.

Gnadt, J.W., and Andersen, R.A. (1988). Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* 70, 216–220.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.

Green, D.M., and Swets, J.A. (1966). *Signal Detection Theory and Psychophysics* (John Wiley).

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520, 220–223.

Heekeren, H.R., Marrett, S., Bandettini, P.A., and Ungerleider, L.G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature* 431, 859–862.

Horowitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161.

- Horwitz, G.D., and Newsome, W.T. (2001). Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *J. Neurophysiol.* *86*, 2543–2558.
- Hou, H., Zheng, Q., Zhao, Y., Pouget, A., and Gu, Y. (2019). Neural correlates of optimal multisensory decision making under time-varying reliabilities with an invariant linear probabilistic population code. *Neuron* *104*, 1010–1021.e10.
- Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* *25*, 10420–10436.
- Janssen, P., Srivastava, S., Ombelet, S., and Orban, G.A. (2008). Coding of shape and position in macaque lateral intraparietal area. *J. Neurosci.* *28*, 6679–6690.
- Jeurissen, D., Shushruth, S., El-Shamayler, Y., Horwitz, G.D., and Shadlen, M.N. (2019). Deficits in decision making after pharmacological and chemogenetic inactivation of area LIP (program no. 061.16). In Society for Neuroscience Meeting.
- Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* *535*, 285–288.
- Kaufman, M.T., Churchland, M.M., Ryu, S.I., and Shenoy, K.V. (2015). Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *eLife* *4*, e04677.
- Keemink, S.W., and Machens, C.K. (2019). Decoding and encoding (de)mixed population responses. *Curr. Opin. Neurobiol.* *58*, 112–121.
- Khalvati, K., Kiani, R., and Rao, R.P.N. (2020). Bayesian inference with incomplete knowledge explains perceptual confidence and its deviations from accuracy. *bioRxiv*. <https://doi.org/10.1101/2020.09.18.304220>.
- Kiani, R., and Shadlen, M.N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* *324*, 759–764.
- Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.* *28*, 3017–3029.
- Kiani, R., Cueva, C.J., Reppas, J.B., and Newsome, W.T. (2014a). Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Curr. Biol.* *24*, 1542–1547.
- Kiani, R., Corthell, L., and Shadlen, M.N. (2014b). Choice certainty is informed by both evidence and decision time. *Neuron* *84*, 1329–1342.
- Kim, B., and Basso, M.A. (2008). Saccade target selection in the superior colliculus: a signal detection theory approach. *J. Neurosci.* *28*, 2991–3007.
- Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* *2*, 176–185.
- Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., and Miyamoto, A. (2013). Responses of pulvinar neurons reflect a subject's confidence in visual categorization. *Nat. Neurosci.* *16*, 749–755.
- Kumano, H., Suda, Y., and Uka, T. (2016). Context-dependent accumulation of sensory evidence in the parietal cortex underlies flexible task switching. *J. Neurosci.* *36*, 12192–12202.
- Lehky, S.R., and Sereno, A.B. (2007). Comparison of shape encoding in primate dorsal and ventral visual pathways. *J. Neurophysiol.* *97*, 307–319.
- Lennie, P. (2003). The cost of cortical computation. *Curr. Biol.* *13*, 493–497.
- Link, S.W. (1992). *The Wave Theory of Difference and Similarity* (Lawrence Erlbaum Assoc).
- Lo, C.C., and Wang, X.J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* *9*, 956–963.
- Mazurek, M.E., Roitman, J.D., Ditterich, J., and Shadlen, M.N. (2003). A role for neural integrators in perceptual decision making. *Cereb. Cortex* *13*, 1257–1269.
- Middlebrooks, P.G., and Sommer, M.A. (2012). Neuronal correlates of meta-cognition in primate frontal cortex. *Neuron* *75*, 517–530.
- Misaki, M., Kim, Y., Bandettini, P.A., and Kriegeskorte, N. (2010). Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *Neuroimage* *53*, 103–118.
- Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., and Pouget, A. (2014). Information-limiting correlations. *Nat. Neurosci.* *17*, 1410–1417.
- O'Connell, R.G., Dockree, P.M., and Kelly, S.P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nat. Neurosci.* *15*, 1729–1735.
- Odegaard, B., Grimaldi, P., Cho, S.H., Peters, M.A.K., Lau, H., and Basso, M.A. (2018). Superior colliculus neuronal ensemble activity signals optimal rather than subjective confidence. *Proc. Natl. Acad. Sci. USA* *115*, E1588–E1597.
- Okazawa, G., Sha, L., Purcell, B.A., and Kiani, R. (2018). Psychophysical reverse correlation reflects both sensory and decision-making processes. *Nat. Commun.* *9*, 3479.
- Okazawa, G., Sha, L., and Kiani, R. (2020). Linear integration of sensory evidence over space and time underlies face categorization. *bioRxiv*. <https://doi.org/10.1101/2020.11.27.396705>.
- Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* *5*, 376–404.
- Peixoto, D., Kiani, R., Chandrasekaran, C., Ryu, S.I., Shenoy, K.V., and Newsome, W.T. (2018). Population dynamics of choice representation in dorsal premotor and primary motor cortex. *bioRxiv*. <https://doi.org/10.1101/283960>.
- Peixoto, D., Verheijn, J.R., Kiani, R., Kao, J.C., Nuyujukian, P., Chandrasekaran, C., Brown, J., Fong, S., Ryu, S.I., Shenoy, K.V., and Newsome, W.T. (2021). Decoding and perturbing decision states in real time. *Nature* *591*, 604–609.
- Philiastides, M.G., and Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cereb. Cortex* *16*, 509–518.
- Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nat. Neurosci.* *19*, 366–374.
- Purcell, B.A., and Kiani, R. (2016). Neural mechanisms of post-error adjustments of decision policy in parietal cortex. *Neuron* *89*, 658–671.
- Purcell, B.A., Heitz, R.P., Cohen, J.Y., Schall, J.D., Logan, G.D., and Palmeri, T.J. (2010). Neurally constrained modeling of perceptual decision making. *Psychol. Rev.* *117*, 1113–1143.
- Raposo, D., Kaufman, M.T., and Churchland, A.K. (2014). A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* *17*, 1784–1792.
- Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* *20*, 873–922.
- Ratcliff, R., Cherian, A., and Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J. Neurophysiol.* *90*, 1392–1407.
- Ritchie, J.B., Kaplan, D.M., and Klein, C. (2019). Decoding the brain: Neural representation and the limits of multivariate pattern analysis in cognitive neuroscience. *Br. J. Philos. Sci.* *70*, 581–607.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* *22*, 9475–9489.
- Rumyantsev, O.I., Lecoq, J.A., Hernandez, O., Zhang, Y., Savall, J., Chrapkiewicz, R., Li, J., Zeng, H., Ganguli, S., and Schnitzer, M.J. (2020). Fundamental bounds on the fidelity of sensory cortical coding. *Nature* *580*, 100–105.
- Sarma, A., Masse, N.Y., Wang, X.J., and Freedman, D.J. (2016). Task-specific versus generalized mnemonic representations in parietal and prefrontal cortices. *Nat. Neurosci.* *19*, 143–149.
- Schall, J.D. (2019). Accumulators, neurons, and response time. *Trends Neurosci.* *42*, 848–860.
- Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. *Neuron* *80*, 791–806.
- Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* *86*, 1916–1936.

- Shushruth, S., Mazurek, M., and Shadlen, M.N. (2018). Comparison of decision-related signals in sensory and motor preparatory responses of neurons in area LIP. *J. Neurosci.* *38*, 6350–6365.
- Smith, P.L., and Vickers, D. (1988). The accumulator model of two-choice discrimination. *J. Math. Psychol.* *32*, 135–168.
- Thura, D., and Cisek, P. (2014). Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* *81*, 1401–1416.
- Tottenham, N., Tanaka, J.W., Leon, A.C., McCarry, T., Nurse, M., Hare, T.A., Marcus, D.J., Westerlund, A., Casey, B.J., and Nelson, C. (2009). The NimStim set of facial expressions: judgments from untrained research participants. *Psychiatry Res.* *168*, 242–249.
- Tsao, D.Y., Freiwald, W.A., Knutsen, T.A., Mandeville, J.B., and Tootell, R.B. (2003). Faces and objects in macaque cerebral cortex. *Nat. Neurosci.* *6*, 989–995.
- Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* *36*, 955–968.
- Wimmer, K., Compte, A., Roxin, A., Peixoto, D., Renart, A., and de la Rocha, J. (2015). Sensory integration dynamics in a hierarchical network explains choice probabilities in cortical area MT. *Nat. Commun.* *6*, 6177.
- Wohrer, A., Humphries, M.D., and Machens, C.K. (2013). Population-wide distributions of neural activity during perceptual decision-making. *Prog. Neurobiol.* *103*, 156–193.
- Wong, K.F., and Wang, X.J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* *26*, 1314–1328.
- Zhou, Y., and Freedman, D.J. (2019). Posterior parietal cortex plays a causal role in perceptual and categorical decisions. *Science* *365*, 180–185.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental models: Organisms/strains		
Rhesus Monkeys (<i>macaca mulatta</i>)	Primate Products	N/A
Software and algorithms		
MATLAB 2018a	Mathworks	https://www.mathworks.com/products/matlab.html ; RRID:SCR_001622
Data acquisition (Central Suite)	Blackrock Microsystems LLC	https://www.blackrockmicro.com/
Plexon Offline Sorter	Plexon	https://plexon.com
Deposited data		
Face and motion datasets and analysis code	This paper	https://github.com/KianiLab/Okazawa_Cell_2021
Other		
Data acquisition system	Blackrock Microsystems LLC	CerePlex Direct
Eye tracker	SR Research	Eyelink 1000
Linear array probe	Plexon	V-probe (16 ch)
Single electrode	FHC	UEWLCHSECN1E
96-ch Utah Array	Blackrock Microsystems LLC	https://www.blackrockmicro.com/electrode-types/utah-array/

RESOURCE AVAILABILITY

Lead contact

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Roozbeh Kiani (roozbeh@nyu.edu).

Materials availability

This study did not generate new unique reagents.

Data and code availability

The datasets and analysis code supporting the current study are available at https://github.com/KianiLab/Okazawa_Cell_2021.

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Monkeys

We recorded responses of parietal and frontal neurons while seven adult macaque monkeys (*Macaca mulatta*; all males, 5–15 years old) performed a variety of perceptual decision-making tasks. Two monkeys (A and L) performed two variants of a novel face discrimination task and five monkeys (D, I, S, O, and N) performed direction discrimination tasks with random dots (Kiani et al., 2008; Roitman and Shadlen, 2002). Monkeys D and I also performed the dots task with post-decision wagering (Kiani and Shadlen, 2009). All experimental procedures conformed to the National Institutes of Health *Guide for the Care and Use of Laboratory Animals* and were approved by the institutional animal care and use committee at New York University. The data from monkeys D, I, and S were previously published (Kiani and Shadlen, 2009; Kiani et al., 2008).

METHOD DETAILS

Behavioral tasks

Monkeys were seated in a semi-dark room in front of a cathode ray tube monitor (frame rate, 75 Hz) with their heads stabilized using a surgically implanted head post. Stimulus presentation was controlled with Psychophysics Toolbox (Brainard, 1997) and MATLAB.

Eye movements were monitored using a high-speed infrared camera (Eyelink, SR-Research, Ontario) or a scleral eye coil (Fuchs and Robinson, 1966). Gaze positions were recorded at 1 kHz.

Throughout the paper, we refer to the face discrimination task as the “face task,” and to the direction discrimination task with random dots as the “motion task.”

Face task

The task required classification of faces into two categories, each defined by a prototype face. Each trial began when the monkey fixated on a small fixation point at the center of the screen (diameter, 0.3°). Shortly afterward, two red targets appeared on opposite sides of the screen, equidistant from the fixation point. After a variable delay (truncated exponential distribution, range, 250–500 ms), a face stimulus appeared 1.6°–2.5°, to one side of the fixation point (contralateral to the recording site in the main experiment or ipsilateral in a control experiment, Figures S2J–S2L). The stimulus size randomly varied across trials by one octave to prevent the monkey from relying on local features (width range, 2.9°–5.8°, mean, 4.2°). The stimulus was presented for a variable duration (truncated exponential distribution, range, 227–1080 ms, mean, 440 ms) followed by a delay period (truncated exponential distribution, range, 300–900 ms, mean, 600 ms). After the delay, the fixation point disappeared (Go cue), and the monkey reported the face category with a saccadic eye movement to one of the targets. Correct responses were rewarded with a drop of juice. To manipulate task difficulty, we created a morph continuum between the two prototypes and presented different intermediate faces (see below). When the face was halfway between the two prototypes on the morph continuum, the reward was given randomly.

We trained monkeys to categorize the same face stimuli in two distinct ways: species and expression categorization (Figure 5A). In the species categorization task, the two prototype stimuli were human and monkey faces, whereas in the expression categorization task, they were faces with “happy” and “sad” expressions. Here, happy and sad are used as designations that facilitate explaining the task. We focus merely on the monkey’s ability to discriminate different expressions and do not imply that they interpret these expressions as humans do. Monkeys performed these two categorizations in blocks that switched either within or across sessions. In the first few trials of each block, we presented only the faces closest to the prototypes of the discriminated categories to cue the task context to the monkey. In both categorization tasks, monkeys easily generalized across different facial identities. We therefore created multiple stimulus sets from different human and monkey identities. Each set enabled testing both categorization tasks (see below). We switched the stimulus set after runs of sessions. In total, three stimulus sets were used during electrophysiological recording sessions.

Each stimulus set can be visualized as a two-dimensional (2D) “face space,” whose axes correspond to the morph level between species and expression category prototypes (Figure 5A). This 2D face space was created with three “seed” faces: photographs of happy and sad expressions of a human face and neutral expression of a monkey face. The seed faces were obtained from the Mac-Brain Face Stimulus Set (Tottenham et al., 2009) and the PrimFace database (<https://visiome.neuroinf.jp/primface>). We manually defined 96–118 anchor points on each seed face and developed a custom algorithm to morph the three seed faces by computing a linear weighted sum of the positions of the anchor points and textures inside the tessellated triangles defined by the anchor points. Our algorithm allowed independent morphing of different stimulus regions. Further, because our morphing changed both the geometry and content of the facial regions, the resulting faces were perceptually seamless.

Natural faces are complex and high-dimensional stimuli that challenge investigation of the decision-making process. To control stimulus complexity and make studying the decision-making process tractable, we limited informative features for face categorization. We selected three features in each face (eyes, nose, and mouth) and limited morphing across faces to those regions, reducing the effective stimulus dimensionality. Regions outside the informative features were always fixed at the midpoint of the morph space between the three seed faces and did not change in the experiment.

Each informative facial feature varied in a 2D “feature space” defined by the corresponding features of the three seed faces. Morphed features were generated as weighted combinations of the three seed features. To change the features along the species axis, the weights for the happy (W_h), sad (W_s), and monkey (W_m) faces varied from [0.5, 0.5, 0] to [0, 0, 1]. We label these two end points as –100% and +100% species morph levels (or species prototypes). For the expression axis, the weights for the prototypes changed from [0.75, –0.25, 0.5] to [–0.25, 0.75, 0.5]. The negative weights indicate linear extrapolation beyond the seed features. We verified that for the weights used in our stimulus sets, all morphed features and resulting faces looked naturalistic and did not show noticeable aliasing. In the 2D space, the species morph level of a feature was defined by $(W_m - W_h - W_s) \times 100\%$, and its expression morph level was defined by $(W_s - W_h) \times 100\%$. Because the three informative features were morphed independently in their 2D feature spaces, our full stimulus space was six dimensional. The 2D face space of Figure 5A shows faces with identical morph levels of eyes, nose, and mouth — a 2D cut of the 6D space.

On each trial, we chose a nominal morph level along the relevant axis for the categorization task (species axis for species categorization or expression axis for expression categorization). This nominal value determined the morph level of the three features in the trial, as well as the response that the monkey was rewarded for. The nominal morph level ranged from –96% to +96%. The exact set of nominal morph levels was adjusted for each stimulus set to keep the monkey’s overall accuracy roughly constant across the sets.

To investigate how monkeys weighted evidence conferred by the three informative features, we allowed the features to fluctuate randomly around the nominal morph level of the trial every 106.7 ms during stimulus presentation. For all nominal morph levels between –12% and +12%, the three features fluctuated independently within their own 2D feature spaces according to a circular Gaussian distribution with a standard deviation of 20% morph level. The features therefore changed both along the relevant and

irrelevant axes for the categorization task, allowing us to gauge which ones influenced the monkey's choices. For strong nominal morph levels (>12%), fluctuations happened only along the relevant axis in order to prevent the monkey from confusing the task context. Sampled values that fell outside the prototype range [−100% +100%] were replaced with new samples inside the range (5.2% of samples).

We used a masking procedure to keep changes of features in a trial subliminal (Okazawa et al., 2018). The masks were created by phase randomization of faces (Heekeren et al., 2004) and interleaved the stimulus fluctuations. In each 106.7 ms sample-mask cycle, a face stimulus was shown for 13.3 ms and then gradually faded out as the mask faded in. In the fading period, the mask and the stimulus were linearly combined, pixel-by-pixel, according to a half-cosine weighting function, such that in the last frame, the weight of the mask was one and the weight of the face was zero. In the next cycle, a new face stimulus with slightly altered informative features was shown, followed by fading with another mask, and so on. Before the first sample-mask cycle in the trial, we showed a mask frame to ensure all stimulus samples were preceded and followed by masks. The masking procedure was quite effective in pilot experiments with humans, preventing the detection of the feature changes. The stimulus in each trial looked like a face appearing and disappearing behind cloud-like patterns. Stimulus viewing duration in each trial consisted of 2–10 stimulus-mask cycles (truncated exponential distribution, mean, 4), which corresponded to 227–1080 ms, including the initial mask.

Two monkeys performed the task in 72 recording sessions (monkey A, 40 sessions; monkey L, 32 sessions), which amounted to 56,582 trials (monkey A, 18,080 and 9,309 trials for species and expression categorizations, respectively; monkey L, 17,077 and 12,116 trials). In the majority of these sessions (55), the face stimulus was contralateral to the recording site. In the control condition where the stimulus was ipsilateral to the recording site (Figures S2J–S2L), we ran 17 sessions (monkey A, 10; monkey L, 7) and collected 12,944 trials (monkey A, 7,584 trials for species categorization; monkey L, 4,041 and 1,319 trials for species and expression categorization, respectively).

Motion task

The experimental settings for the motion task are described elsewhere (Kiani and Shadlen, 2009; Kiani et al., 2008), and summarized here. The monkey began each trial by fixating a fixation point at the center of the screen, followed by the appearance of two red targets on opposite sides of the screen equidistant from the fixation point. After a variable delay (truncated exponential distribution, range, 250–600 ms), the moving random-dots stimulus (Britten et al., 1992) appeared within a 5°, circular aperture centered on the fixation point. The percentage of coherently moving dots determined the strength of motion (coherence). The motion strength varied randomly across trials from the following set: 0%, 1.6%, 3.2%, 6.4%, 12.8%, 25.6%, and 51.2% coherence. The net motion direction was toward one of the two targets and varied randomly across trials. Monkeys I and S performed a variable duration task, where the stimulus was presented for 80–1500 ms (truncated exponential distribution, mean, 311 ms). The Go cue occurred after the stimulus or following a delay period (truncated exponential distribution, range, 500–1000 ms, mean, 847 ms), instructing monkeys to report their perceived motion direction with a saccadic eye movement to one of the targets. Monkey D as well as monkey I performed another variable duration task, where the stimulus duration was 100–900 ms (mean, 286 ms) and always followed by a delay duration of 1200–1800 ms (mean, 1386 ms). Monkeys O and N performed a reaction time task, where they were free to respond any time after the stimulus onset (mean reaction time, monkey O, 729 ms, monkey N, 927 ms). Correct responses were rewarded with water or juice. For 0% coherence motion, the reward was given randomly.

Monkeys D and I also performed a version of the variable duration task with post-decision wagering (Kiani and Shadlen, 2009). During this task, on a random half of the trials, a third target (“sure” target) appeared at a random time during the delay period (500–750 ms after motion offset) and stayed on through the rest of the delay period. After the Go cue, the monkey made a saccade either to one of the two direction targets or to the sure target, if present. Choosing the sure target always yielded a reward, but the reward size was smaller than that for choosing the correct direction target. The reward ratio was adjusted to encourage the monkey to choose the sure target on nearly half of trials.

We collected 67,349 trials in the variable duration task (monkey I, 24,603 trials; monkey S, 21,947 trials; monkey D, 20,799 trials). Among them, 30,991 trials were in the post-decision wagering task (monkey I, 10,192 trials; monkey D, 20,799 trials). We collected 23,756 trials in the reaction time task (monkey O, 6,291; monkey N, 17,465). Psychophysical thresholds were comparable across the tasks.

Neural recording

For recordings from the lateral intra-parietal (LIP) cortex, electrodes were placed in the ventral division of LIP (LIPv), located based on both structural magnetic resonance imaging scans and transition of white and gray matter during recordings. The recording was performed through a plastic cylinder (Crist Instruments, Damascus, MD), implanted on the skull, and a plastic grid (1 mm spacing; Crist Instruments), placed inside the chamber for precise targeting of the electrodes. We used either single tungsten microelectrodes (FHC, Bowdoin, ME) or 16-channel linear array probes (V-Probe; Plexon, Dallas, TX). Action potential waveforms were isolated online using a time-amplitude window discriminator or sorted offline (Plexon offline sorter). In our analyses, we combined both well-isolated single units and multi-units. We confirmed that similar results were obtained using single units alone (Figure S4D).

Units were selected for analysis if they exhibited spatially-selective persistent activity during the delay period of a memory-guided saccade task (Gnadt and Andersen, 1988). While the monkey maintained fixation, a target briefly flashed on the screen and was followed by a delay period (~1000 ms). At the end of the delay period, the fixation point turned off (Go cue), and the monkey made a

saccade to the remembered target location. The target location varied randomly across trials. The response field (RF) of the neurons was identified as the target locations associated with the largest firing rates during the delay period. During the motion or face task, we placed one of the targets (T_{in}) in the RF and the other (T_{out}) outside the RF. For the post-decision wagering results included in this paper, the sure target (T_s) was placed outside the RF.

Our choice to focus on units with persistent delay activity in a memory-guided saccade task was made to remain consistent with previous studies (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). However, the LIP neurons lacking persistent activity also showed similar curved population response manifolds during our main tasks (Figure S2F–S2I). Thus, similar conclusions could be made using all LIP neurons or different subgroups of neurons.

In the motion task, 129 units were recorded from the LIP of three monkeys (monkey D, 46 single units; monkey I, 50 single units; monkey S, 33 single units). In the face task, 132 units were recorded from two monkeys (monkey A, 70 (single: 41); monkey L, 62 (single: 34)) with face stimuli presented contralateral to the recording site. An additional 44 units (monkey A, 31; monkey L, 13) were also recorded during the control condition with ipsilateral stimuli (Figures S2J–S2L).

To examine the decision-related activity of other cortical areas, we also recorded from two prefrontal regions: prearcuate gyrus (PAG) in lateral frontal cortex and supplementary eye field (SEF) in dorsomedial prefrontal cortex. The recordings from PAG were performed from monkey N with a chronically implanted 96-channel microelectrode array (electrode length = 1 mm; spacing = 0.4 mm; Blackrock Microsystems, Salt Lake City, UT). We performed 11 recording sessions, which yielded 84 ± 10 (mean \pm s.d.) simultaneously recorded units per session. The recordings from SEF were performed from monkey O with either single tungsten microelectrodes (FHC) or 16-channel linear array probes (Plexon). In 9 recording sessions, we isolated 34 SEF units. SEF was identified based on stereotactic coordinates, saccade selectivity, and evoked-saccades with low-current microstimulation ($<50\mu A$).

QUANTIFICATION AND STATISTICAL ANALYSIS

Behavioral data analysis

We sorted the stimulus strengths for each task into five levels. For the motion task, we used motion coherence to define the strength levels, compatible with past studies (Kiani and Shadlen, 2009; Kiani et al., 2008; Shadlen and Newsome, 2001): 0%–3.2%, 6.4%, 12.8%, 25.6%, and 51.2% coherence. For the face task, we used the average morph level across stimulus-mask cycles in each trial: 0%–5%, 5%–15%, 15%–25%, 25%–50%, 50%–100% morph. Due to the stochastic nature of the stimuli, this average morph level could differ from the nominal morph level. Specifically, in addition to changing the stimulus along the relevant axis for the categorization task, the fluctuations also shifted the stimulus along the irrelevant axis. Using the average morph level enabled us to quantify the veridical stimulus strength for both axes and explore their effects on the behavior and neural responses. Apart from this benefit, we confirmed that defining the strength levels based on the nominal morph level did not critically change any of our results. Similarly, we confirmed that defining the stimulus strength levels of the motion task based on average motion energy — a measure of veridical stimulus strength (Adelson and Bergen, 1985; Kiani et al., 2014b) — did not change any of our conclusions. The same stimulus strength levels were used for both behavioral and neural data analyses.

Note that the stimulus strength could be in favor of one choice or another. Therefore, for analyses that depended on the exact choice made by the monkey (e.g., T_{in} versus T_{out}), we expanded the definition of stimulus strength levels to add a sign to the numbers above, where positive and negative strengths indicate stimuli supporting the T_{in} and T_{out} choices, respectively. Throughout the paper, we use consistent colors to show different stimulus strength levels in the figures, with red symbols depicting positive strengths and blue symbols depicting negative strengths, while the saturation of symbols depicting absolute stimulus strength.

Psychometric functions (Figures 1C and 1E) were calculated as the probability of making correct choices — choices compatible with the sign of the stimulus strength level — as a function of stimulus strengths. Psychophysical thresholds were defined as the stimulus strength that yielded 81.6% correct choices. Thresholds were estimated by fitting a logistic function to the choice data (similar results were obtained with cumulative Weibull fits). To assess the effect of stimulus duration on behavioral performance, we divided trials into six stimulus duration quantiles, calculating the psychophysical threshold separately for each quantile (Figures 1D and 1F).

To test if the behavioral performance is consistent with an evidence accumulation mechanism, we fit a drift diffusion model (DDM) to the data. In the DDM, signed momentary sensory evidence accumulated over time to create the decision variable (DV). The process continued until the DV reached either an upper or a lower bound, or all the available sensory evidence was integrated. The bound that was reached first dictated the choice. But if the accumulated evidence failed to reach a bound by the end of the trial, the sign of the DV determined the choice. We fit the DDM to individual monkeys' behavior using a maximum-likelihood procedure. Details of the method are described elsewhere (Okazawa et al., 2018). The DDM had two free parameters in the motion task: sensitivity and bound height. Sensitivity, k , determined the linear scaling of the mean momentary evidence in the model with signed stimulus strength, s . Therefore, the drift rate of the diffusion process at each moment, $\mu(t)$, was equal to $ks(t)$. Since stimulus strength was defined based on motion coherence, $s(t)$ was a constant in the model. The bound height, B , determined the amount of evidence that had to be accumulated to reach the upper ($+B$) or lower ($-B$) bound. Because the drift rate and bound height scale with the s.d. of momentary evidence in the model, we set the variance of the momentary evidence to 1 to allow unique solutions for the model fit.

For the face task, we allowed different sensitivity parameters for the three informative features in each categorization task: $\mu(t) = k_e s_e(t) + k_m s_m(t) + k_n s_n(t)$, where $s_e(t)$, $s_m(t)$, and $s_n(t)$ are signed morph levels of eyes, nose, and mouth along the relevant axis for categorization at time t . We did not observe any notable influence of morph levels along the irrelevant axis on the monkey's choices,

within the tested ranges in our tasks (Figure 5B). We therefore excluded them from the equation for drift rate to limit the number of model parameters. In a reaction-time version of the task performed by human subjects, we established that sensitivity parameters were largely constant over time (Okazawa et al., 2020) and that there was minimal loss of information (leak). We do not present these model variants in this paper as they do not directly pertain to our main conclusions. However, we note that we do not find substantial evidence for non-stationary sensitivity or the presence of a significant leak in the monkey data, suggesting a similarity of the decision-making process across species.

The monkeys that performed the same task had consistent behavior and neural responses. We therefore combined data across subjects in the figures. But when appropriate, we show individual monkeys' results too (e.g., Figure S2). The behavioral plots in Figures 1C–1F combine data from all monkeys performing the respective tasks. The model fits (gray lines) are the average of fits for individual monkeys. Figures 1C and 1D combines data and fits across both monkeys and both face categorization tasks. Figure 5B shows the behavior separately for the species and expression categorization tasks. Since the veridical stimulus strengths varied along both task-relevant and irrelevant axes, we plotted two psychometric functions, one for each axis. Both curves in the figure are from the DDM explained above. The model fits do not include any degree of freedom for stimulus strengths along the irrelevant morph axis, as explained above. The non-significant changes of model accuracy along the irrelevant morph axis are due to random stimulus fluctuations along the task-relevant axis that do not totally average out for the limited number of trials in the dataset.

Neural data analyses

Peri-stimulus time histograms (PSTHs) shown in Figures 2A, 2B, 6A, 6C, S2, and S4D were smoothed by convolution with a 100 ms boxcar filter. The PSTHs were aligned to the stimulus onset and cut off at 300 ms after stimulus offset to focus our analysis on the responses pertaining to the decision formation based on sensory inputs. The same convention was used in subsequent analyses. When plotting the PSTHs, we focus on correct trials only to be able to examine changes of firing rate with stimulus strength for the same choice except for Figure S2E.

The firing rates shown in Figures 2C, 2D, and S2K were calculated by counting the number of spikes during decision formation (250–600 ms after stimulus onset) in correct trials. To combine data across units with different ranges of firing rates, we z-scored the firing rates of each unit by subtracting its mean firing rate and dividing by the standard deviation of firing rates across all trials. We then calculated the mean of z-scored firing rates of each unit for each stimulus strength and averaged them across units. To quantify the relationship between the firing rates, $r(s)$, and the stimulus strength, s , we performed linear regressions independently for T_{in} and T_{out} choices:

$$r(s) = \alpha_0 + \alpha_1 |s| \quad (1)$$

where $|s|$ is the absolute (unsigned) stimulus strength. The regression coefficients (α_0 and α_1) were determined for each unit. The lines in Figures 2C, 2D, and S2K were generated from average α_0 and α_1 across units. The significance of slopes across the neural population was determined based on a two-sided t test on the distribution of single unit slopes.

To test if the differences in activity between two pools of neurons preferring two choices implement the DV, we computed the differential firing rates between units preferring two choices (Figure 6E). Since the units were recorded from one hemisphere and most of them had the same preference, we created a pseudo-population, where the preferred choice of a random half of the recorded units is inverted. We then computed the differences in the firing rates between the units with the opposite preferences for each stimulus strength. We repeated this procedure 1,000 times and averaged the PSTHs over these repetitions.

To examine if the difference of T_{in} and T_{out} firing rates increased with stimulus strength, we used the following linear regression:

$$\Delta r(s) = \beta_0 + \beta_1 |s| \quad (2)$$

where $\Delta r(s)$ is the differential firing rate between correct T_{in} and T_{out} choices for the unsigned stimulus strength $|s|$. Figure 6G shows the distribution of the slope coefficient (β_1) across the recorded units.

We implemented a neurometric analysis (Figure 6H) to quantify the accuracy of an ideal observer that uses the responses of LIP neurons to discriminate the stimuli with similar strength on the two sides of the category boundary in each task (Britten et al., 1992; Shadlen and Newsome, 2001). The ideal observer accuracy equals the area under the Receiver Operating Characteristic (ROC) curve of the distribution of spike counts for the two competing stimuli. Figure 6H shows the results for different stimulus strengths using a 100 ms sliding window. To remain conservative, we did not conditionalize spike count distributions on correct choices in this analysis. Because conditioning on correct choices would lead to higher areas under ROC for weaker stimuli, it would further reduce the dependence of the ideal observer's accuracy on stimulus strength, bolstering our conclusion that the difference of the activity of pools of neurons selective for the two choices fails to fully explain the representation of the DV in the face task.

The PSTHs for the reaction-time version of the motion task (Figure 7A) were calculated from 100 ms before stimulus onset to the median reaction time for each stimulus strength. A 50 ms window before saccade onset was excluded from each trial so that the PSTHs were not influenced by the motor burst of the neurons. Using a more conservative exclusion window (100 ms before saccade onset) did not critically change the results. For both the PAG and SEF units, target selectivity was determined using neural responses during a presaccadic period (–300 to 0 ms from saccade onset) and the choice

target that evoked higher activity was defined as the putative T_{in} . Responses in memory- or visually-guided saccade tasks corroborated our designations.

In the following sections, we explain analyses for characterization of population response patterns. We use these population analyses as a vehicle to drive intuition and develop insights, noting that because these methods build on single cell responses, their conclusions can also be derived from careful analysis of single cells across the population.

Principal component analysis of neural data

We performed principal component analysis (PCA) to visualize the neural population response manifolds (Figures 2E–2H, 5C, 6A–6D, 7B, S2, S3, S4, and S5). PCA was performed across units using trial-averaged PSTHs. Each unit contributed multiple PSTHs, one for each combination of choices and stimulus strength levels. The PSTHs focused on the period of decision formation and spanned 250–600 ms after stimulus onset. We detrended the PSTHs of each unit by subtracting the average unit PSTH across all stimulus strengths and choices to focus on modulations of activity associated with task parameters. We confirmed that the detrending did not critically change the observed geometry of neural manifolds. We then concatenated the detrended PSTHs of each unit into a response vector for that unit and combined the vectors of all units into a population response matrix with $T \times C$ rows and N columns, where T is the number of time points in each PSTH (1 ms resolution), C is the number of conditions (5 stimulus strength levels \times 2 choices, T_{in} and T_{out}), and N is the number of units. Only correct T_{in} and T_{out} choices were used for these analyses because error responses were rare for high stimulus strengths. A tiny fraction of elements in the population response matrix were missing ($<0.1\%$) due to low numbers of trials for some units. The missing data were replaced with the average response of the corresponding unit. The first three principal components (PCs) explained 78% of the total variance of the full N dimensional neural space in the motion task and 54% of the total variance in the face task. The analysis included all units recorded either simultaneously or separately. We confirmed that similar results were obtained using simultaneously recorded units only (Figure S4C).

Because we have aligned the PSTHs of LIP neurons with respect to their RFs (Figure S3A; PSTHs were defined with respect to T_{in} and T_{out} , not the right-left location of targets), the PCA effectively captures the structure of population responses within a group of neurons having a shared choice preference. Therefore, the observed manifold geometry arises from the variability of neural responses within that group of neurons (see Figure S3B for intuition). The finding, however, also extends to a neural population with diverse choice preferences. We performed PCA for a pseudo-population created by inverting the preferred choice of a random half of the recorded units. The resulting manifold geometry leads to conclusions similar to our main results (Figures S3C–S3F).

The neural population response for each stimulus strength level at each time can be depicted as a point in the three-dimensional (3D) PC space (Figures 2E–2H, 5C, 6A–6D, 7B, S2, S3, S4, and S5). The data points in the figures were calculated based on spike counts within a 100 ms window centered at the specified times in each figure. To draw a manifold that captures the population responses for different stimulus strength levels (gray lines), we fit the projection of population responses along each PC axis with a cubic smoothing spline as a function of stimulus strength (*csaps* function in MATLAB). The SEM of the population response projections along each PC axis (error bars in Figures 2G, 2H, 7B, and S2L) was computed using a bootstrap procedure. We re-computed the mean firing rate of each unit for each stimulus strength level and time by randomly sampling trials with replacement. Then, we projected the resulting population response patterns onto the 3D PC space. This procedure was repeated 1,000 times, and the standard deviations of the projections were calculated and used as the SEM of the data points in the plots.

To rule out the possibility that the population response manifold is actually a straight line that looks curved due to noise in the measured neural responses, we performed three tests. The first two examined the consistency of the direction of the curvature using a cross-validation procedure, and the third tested whether the observed extent of the curvature could arise solely from neural noise. For the analyses in the paper, the three tests always provided consistent results.

Stochastic variability of neural activity could indeed cause changes in the curvature of the population response manifold, but these noisy changes are in random directions in the state space and would lack consistency. In our cross-validation tests, we divided the trials of individual units into random halves and derived a manifold for each half. In the first test, we directly compared the directions of the major axes of the two manifolds, \vec{v}_1 and \vec{v}_2 . The major axis of each manifold was defined as the vector between the midpoint of the two ends of the manifold and the apex of the manifold. The angles between the two major axes were quite small, $11.9^\circ \pm 6.0^\circ$ in the motion task and $9.3^\circ \pm 4.9^\circ$ in the face task (*mean* \pm *SEM*, calculated over 1,000 iterations of random splitting of data). These angles were significantly smaller than expected from noise (90°), indicating strong consistency in the direction of manifold curvature.

In our second cross-validation procedure, we adopted a more robust approach that averaged the direction of curvature of all locations of the manifold for the second half of the data with respect to the major axis of the manifold from the first half. For the 1D manifold ($\vec{m}(s)$) in the state space at the location corresponding to stimulus s , its curvature is:

$$\vec{T}(s) = \frac{\vec{m}'(s)}{|\vec{m}'(s)|} \quad (3)$$

$$\vec{N}(s) = \frac{\vec{T}'(s)}{|\vec{m}'(s)|} \quad (4)$$

$$\kappa(s) = |\vec{N}(s)| \quad (5)$$

where $\vec{T}(s)$ is the unit tangent vector, $\vec{N}(s)$ is a vector normal to the curve, and the length of $\vec{N}(s)$ equals the curvature, $\kappa(s)$, at the location corresponding to stimulus s . $|\cdot|$ is the L_2 norm. The direction of $\vec{N}(s)$ determines the direction of curvature in the state space. The dot product of $\vec{N}(s)$ with \vec{v}_1 determines how well their directions match each other:

$$\frac{\cos(\phi(s)) = \vec{N}(s) \cdot \vec{v}_1}{|\vec{N}(s)| \cdot |\vec{v}_1|} \quad (6)$$

where $\phi(s)$ is the angle between $\vec{N}(s)$ and \vec{v}_1 at s . $\cos(\phi(s))$ was computed at every point on the manifold and then averaged over the full extent of the manifold. Positive values of $\langle \cos(\phi(s)) \rangle_s$ indicate that the direction of the curved manifold for the second half of data is consistent with the first half. We tested if this value is significantly greater than zero using a bootstrap procedure (iterations, 1,000).

For our third test of the significance of manifold curvature, we compared the curvature of the actual neural data with that of simulated data from a linearized response manifold with similar conditional response variability to the actual data (Figure S4A). The simulated data were generated by straightening the underlying manifold in the actual data. We first redefined the mean firing rates of each unit for different stimulus strengths by linearly interpolating between firing rates for the strongest stimuli supporting T_{in} and T_{out} choices. This interpolation ensured the linearity of the simulated response manifold. For each stimulus strength, we defined the distribution of simulated responses of a unit as a mean-adjusted version of the true distribution of unit responses, such that variance and shape of stimulus-conditioned response distributions remained unchanged. We simulated individual trials by randomly sampling from these adjusted distributions for the same number of trials that each unit was recorded for. Repeating this process for all the units generated a simulated dataset with the same number of units, trials, and conditional response variability as the actual data. We created 1,000 simulated datasets and derived the manifold for each of them using the procedure explained above. Then we calculated the curvature at each point of each manifold using Equation 5 and quantified the overall magnitude of the curvature as the average curvature across the manifold. Finally, the magnitude of the manifold curvature in the data was compared with the distribution of magnitudes obtained from the simulated data.

In the PC space, the projected population average firing rate changed along a linear axis (straight line in Figures 2G, 2H, 5E, and S5B), since PCA is a linear dimensional reduction procedure. A vector that corresponds to this linear projection axis can be readily derived from the PC coefficients:

$$\vec{v}_{pop.avg.} = \left[\sum_{n=1}^N w_{1,n}; \sum_{n=1}^N w_{2,n}; \sum_{n=1}^N w_{3,n} \right] \quad (7)$$

where $w_{i,n}$ is the coefficient of neuron n for the i -th PC and N is the total number of neurons. In Figures 2G and 2H, to illustrate the population average firing rates corresponding to locations on the manifold, we projected a perpendicular line from each point on the manifold to the axis of the average firing rate in the 3D PC space. Due to limitations of 2D illustrations, however, these perpendicular projections in the 3D space appear to have angles different from 90° in the figure.

For the reaction time tasks, PCA was performed as described above but using the detrended PSTHs from 250 ms to 550 ms after stimulus onset. The PAG results in Figure 7 were generated using the simultaneously recorded units in one session ($n = 99$). Similar results were obtained in all recorded sessions.

Locating the DV and stimulus difficulty axes in the population response state space

The curved population response manifold suggests the presence of linear axes in the state space that encode the decision variable (DV) and stimulus difficulty (Figure 3A). To identify the best encoding axes, we performed an orthogonal canonical correlation analysis (CCA) (Cunningham and Ghahramani, 2015) (Figure 3B). CCA finds linear transformations that maximize the correlation of two matrices with each other. One matrix in our analysis was the neural population response, \mathbf{R} , consisting of the trial-averaged, detrended PSTHs of the recorded units. The population response matrix was constructed as we explained for the PCA analysis but was limited to a shorter window (350–450 ms after stimulus onset) in order to allow approximation of the DV and stimulus difficulty based on task parameters. Neural activity around 400 ms after stimulus onset well reflects the decision formation, but similar results could also be obtained at earlier or later times (e.g., 300 ms, 500 ms, or 600 ms). The second matrix in our analysis was the task parameter matrix, \mathbf{P} , which approximated the DV and stimulus difficulty associated with each PSTH as s (signed stimulus strength) and $-|s|$, respectively. For simplicity, we assumed a constant DV and difficulty during the short PSTH snippet included in the response matrix. The average DV across trials of a particular stimulus strength is proportional to the signed stimulus strength in our tasks, especially at early periods of the decision-making process when the decision bound is not yet reached for the majority of trials because the behavior could be successfully modeled with linear integration of stimulus strength (Figures 1C–1F and S1). The proportionality constant, however, changes with time, as the DVs for different stimulus strengths diverge from each other

(Beck et al., 2008; Kiani et al., 2014a; Peixoto et al., 2018). It is therefore important to limit the duration of analysis window to ensure reliable CCA results, while keeping it long enough to allow accurate estimation of mean firing rates in each condition. A 100 ms window met the analysis requirements, but our conclusions were not critically dependent on it. Further, the exact time of the analysis window had minimal impact on our results as long as it reasonably overlapped with decision formation (we tested 350–600 ms after stimulus onset). More complex methods in which we inferred the DV and task difficulty using the exact stimulus fluctuations corresponding to the analysis window and the DDM yielded similar results.

The CCA seeks projection vectors, $\vec{\alpha}_i$ and $\vec{\beta}_i$, that maximize the correlation of canonical variables $\mathbf{R}\vec{\alpha}_i$ and $\mathbf{P}\vec{\beta}_i$. After finding the first pair of canonical variables, the analysis finds a second pair of canonical variables uncorrelated with the first. Orthogonal CCA also ensures that the projection vectors are orthogonal (Cunningham and Ghahramani, 2015). Combining the projection vectors yield 2D transformation matrices $\mathbf{A} = [\vec{\alpha}_1, \vec{\alpha}_2]$ and $\mathbf{B} = [\vec{\beta}_1, \vec{\beta}_2]$, which we used to project the neural responses to the 2D subspace defined by the linear DV and difficulty axes:

$$\vec{x}_{CCA} = \vec{r}^T \mathbf{A} \mathbf{B}^{-1} \quad (8)$$

where \vec{r} is the population response vector and \vec{x}_{CCA} is the position in the 2D subspace. In Figures 3, 4, and 5H, \vec{r} consisted of firing rates of the neurons in 100 ms windows. For the CCA results presented in the paper, we reduced the dimensions of \mathbf{R} to 10 using PCA to attempt denoising the population responses prior to the CCA. However, the denoising or the number of PC dimensions used for denoising was not critical for our conclusions (we tested a wide range of dimensions from 10 to 40; 10 dimensions explained 89% and 76% of the total variance in the motion and face tasks, respectively). The 2D CCA dimensions accounted for 63% (motion) and 36% (face) of the total variance.

All figures and analyses that depended on CCA were cross-validated. The cross validation was implemented by using a random half of trials as a training set to compute the transformation matrices, and then applying them to the other half of trials (test set). In Figures 3B–3E and 4C–4E, the SEM of the data points were estimated using a bootstrap procedure within the test trials (iterations, 1,000). The gray curves in Figures 3B and 4C, which estimated the manifold spanning different stimulus strengths, were computed by fitting the projection of neural responses along each dimension using a second-order polynomial function of stimulus strength.

To plot the time course of neural responses along the two axes defined by CCA (Figures 3C and 3D), we further split the test trials based on the monkey's choice (correct or error) or based on stimulus difficulty (easy or difficult, specified as stimulus strengths greater or less than 20%). When computing the choice-dependent neural responses (Figure 3C), all stimulus strengths were combined except for 0%, for which the correct choice was undefined. When computing the difficulty-dependent neural responses (Figure 3D), we used only correct trials, thus keeping the monkey's choice identical between the easy and difficult conditions.

To examine the relationship between the neural responses and the monkey's confidence, we analyzed a subset of neurons (70/129) recorded during the motion task with post-decision wagering (see the section on Behavioral tasks). Bailing out of the motion direction discrimination by choosing the sure-bet option indicated low confidence. We first performed CCA on a random half of trials with correct choices (training set) to find the best encoding axes for the DV and stimulus difficulty. Then, we projected neural responses of the remaining trials on these axes, independently for correct, error, and T_s choices. In Figure 4C, we adjusted the sign of the DV axis projections such that positive and negative values indicated DVs in favor of the correct and error choices, respectively. To quantitatively test the relationship between the neural responses and the monkey's choices (Figures 3E, 4D, and 4E), we computed the residual activity associated with each choice (correct, error, and T_s) by subtracting the mean projection of all test trials with the same stimulus strengths along each axis. We then averaged the residual responses across different stimulus strengths to achieve better visualization and boost analysis power. For these analyses, we used only trials with low stimulus strengths (3.2%–12.8%), where error and T_s choices were present in all sessions. Because correct and error choices were undefined for 0% coherence trials, they were not included.

Testing task dependency of neural responses

To examine how task contexts influenced LIP activity, we compared the neural responses during the species and expression categorization of the same face stimuli (Figure 5). The analysis used a subset of units ($n = 28$) recorded in both tasks in the same sessions. To create the PC space common to both categorization tasks, we used stimuli present in both tasks (gray region in Figure 5A). The stimuli were grouped into four strength levels ([−15% − 5%], [−5% 0%], [0% 5%], [5% 15%]) for each stimulus axis (species and expression). This yielded 32 conditions (4 species levels \times 4 expression levels \times 2 task contexts). We computed PSTHs of correct trials in these conditions and performed PCA using the same procedure described earlier.

In Figure 5D, we plotted the PC scores of each stimulus strength (4 \times 4 levels) in a window 350–450 ms after stimulus onset. The significance of differences in PC scores between the tasks was determined using a permutation test. For each stimulus level, we shuffled the trial labels between the two tasks (preserving the number of trials) and computed PC scores of the two shuffled groups. We then evaluated the similarity of the two groups of PC scores by calculating the Pearson correlation coefficient across stimuli and PC dimensions. We repeated this procedure 1,000 times to establish the null distribution and compared it against the observed correlation coefficient to report the p -value.

In Figure 5C, we selected the top three PCs derived from the above analysis and projected neural responses calculated for our standard stimulus strength levels in each task as explained in earlier sections. Here, the statistical significance of the difference of

the manifolds between the two tasks was evaluated using a likelihood ratio test. We fit the PC scores of each dimension using a second-order polynomial function of stimulus strength, either simultaneously, using the same coefficients for both tasks, or separately, using different coefficients for each task. The goodness of fits was quantified as the likelihood of data points given the manifolds under a Gaussian noise model. The noise covariance matrix was estimated using a bootstrap procedure (iterations, 1,000). The likelihood ratio test determined whether the likelihood of separate manifolds for the two tasks was significantly larger, taking into account the difference in the degree of freedom between the two fits.

The task dependency of neural responses influenced how the DV was encoded by the population. To examine the extent of this effect, we first determined the DV axis in each task context using the CCA procedure described earlier. Our results did not critically depend on whether we determined the DV axis by performing the 2D CCA (estimating both the DV and difficulty axes) or by performing a 1D CCA (estimating the DV axis only). We then computed the angle between the DV axes of the two face categorization tasks in the 3D PC space (θ in Figure 5E). We used the neural responses in a window 350-450 ms after stimulus onset, but the results were largely invariant to the exact choice of analysis window during decision formation. θ was defined with respect to the axis of average firing rate ($\vec{\Gamma}_{pop.avg.}$ in Equation 7) such that positive values indicated a greater angle between $\vec{\Gamma}_{pop.avg.}$ and the DV axis in expression task, compared to that in species task. Hence, when the DV axes in the two categorization tasks differ, θ should be either significantly greater or less than zero. We used a bootstrap procedure to test the statistical significance of θ (iterations, 1,000).

We quantified the extent to which the angle between the DV axes (θ) affected the linear readout of the encoded DVs. We projected the neural responses recorded in one task context on the DV axis derived for the same or different tasks (Figure 5H; the results of the expression task are plotted). The projected responses were fit with a logistic curve as a function of stimulus strength (s):

$$x_{CCA} = (\gamma_0 + I \cdot \Delta\gamma_0) + \frac{\gamma_1}{1 + \exp(-(\gamma_2 + \gamma_3 s))} \quad (9)$$

where γ_i are regression coefficients and I is an indicator variable (0 for using the DV axis from the same task and 1 for using the DV axis from the different task). $\Delta\gamma_0$ is a regression coefficient that captures the change of bias caused by not adjusting the DV axis based on task context. The significance of $\Delta\gamma_0$ was examined using a likelihood ratio test that compared the likelihood of models with and without $\Delta\gamma_0$. The likelihoods were estimated using a bootstrap procedure (iterations, 1,000).

The analyses above used only the neurons recorded in both task contexts. Because the rest of units ($n = 104$) were recorded for only one task context, their DV axes could not be directly compared. However, the axis of average firing rate ($\vec{\Gamma}_{pop.avg.}$) could be used as a common reference. We reasoned that, if a brain area encodes the DV in a context-independent manner, a population of neurons recorded from this area would have a fixed angle between the DV axis and the axis of average firing rate (i.e., $\omega_{SPC} = \omega_{EXP}$ in Figure 5E). We therefore tested if the angle between the DV axis and $\vec{\Gamma}_{pop.avg.}$ differed between the populations recorded for each face categorization task (Figure 5G). The SEM of the angles and statistical significance of the angle differences were calculated using a bootstrap procedure (iterations, 1,000). Repeating these tests in higher dimensional state spaces provided similar conclusions.

Supplemental figures




Figure S1. Psychophysical kernels of the face tasks support bounded evidence accumulation with task-dependent weighting of facial features, related to Figure 1

(A) In the face task, morph levels of the three informative facial features (eyes, nose, mouth) randomly fluctuated every 106.7 ms around a mean morph level on each trial (nominal morph level). These fluctuations allowed us to perform psychophysical reverse correlation (Ahumada, 2002) to examine spatiotemporal weighting of features for decision-making. Psychophysical kernel, $K_f(t)$, for each facial feature f at time t was calculated as the difference of average fluctuations of morph levels conditional on the monkey's choices:

$$K_f(t) = E[s_f(t)|T_{in\ choice}] - E[s_f(t)|T_{out\ choice}]$$

where $s_f(t)$ is the morph level of feature f at time t . We used trials with nominal morph levels less than 20%, where stimulus fluctuations were potent enough to switch the monkey's decision. To calculate a kernel across different morph levels, we subtracted the nominal morph level of each trial from the actual morph levels presented on the trial and then averaged these residuals across trials. Colored lines show the resulting kernels for species categorization (top row) and expression categorization (bottom row). Kernel amplitudes were different across features and tasks, indicating that (1) the three informative features differentially contributed to the monkey's decisions in each task, and (2) the feature weights varied across tasks. Species categorization relied most heavily on the eyes but expression categorization on the mouth. Further, the kernels decreased over time in our task design, compatible with bounded integration of evidence, although such declines could also arise from other mechanisms (Okazawa et al., 2018). To quantify the weights of the three features for the monkey's choices, we approximated the decision-making process with a drift diffusion model (DDM) that accommodated differential weights for the informative features along the relevant morph axis for each categorization task (see STAR Methods). Fitting the model to the distribution of choices across trials quantitatively explained monkeys' psychometric functions and changes of threshold with stimulus viewing duration (Figures 1C and 1D). Further, it generated predictions for the shape of psychophysical kernels that closely matched the data (gray lines). The predicted kernels were calculated by creating new stimulus sequences that were not used for fitting the model and simulating the DDM for these stimuli with the fitted model parameters. The predicted kernels accurately explained the data (species categorization, $R^2 = 0.97$; expression categorization, $R^2 = 0.98$). (B) Feature weights of the DDM fits in the two face categorization tasks. Error bars are SEM.




Figure S2. Reversed ordering of T_{in} PSTHs with stimulus strength was consistently present across units, monkeys, different categorizations, and conditions in the face task, related to Figure 2

(A) Units recorded in the face task (top row) showed a reversed order of T_{in} PSTHs for different stimulus strengths compared to those recorded in the motion task (bottom row). We show three example units from each task. For clarity, trials were divided into two difficulty levels ($\geq 20\%$, red, and $< 20\%$, black) for T_{in} (solid) and T_{out} (dashed) choices. The PSTHs show only correct responses and are smoothed with a 100 ms boxcar filter. (B) Histograms of the slope of individual unit firing rates (window from 250 to 600 ms) as a function of stimulus strength (α_1 in Equation 1). Overall, units recorded in the face task showed negative slopes for T_{in} . Arrows indicate means. (C) Population average PSTHs for individual monkeys. Conventions are the same as in Figure 2A. (D) Population average PSTHs for species and expression categorizations in the face task. (E) Population average PSTHs including both correct and error trials. Including error trials leads to different proportions of T_{in} and T_{out} choices for different stimulus strengths. A mixture of T_{in} and T_{out} choices for lower stimulus strengths moves their PSTHs toward the center of the firing rate range. Nonetheless, the reversal of the order of the PSTHs was still present in the face task. (F-I) Although our main analyses focused on neurons with significant persistent activity during the delay period of the memory-guided saccade task to remain consistent with past studies, the cells

(legend continued on next page)

that did not elicit significant persistent activity in the saccade task showed similar population response profiles in our main tasks. Including these cells in our analyses for the face task ($n = 111$) did not critically change the reverse ordering of PSTHs (F) or the curved manifold of the population responses (H). In fact, the population of the cells without persistent activity in the memory-guided saccade task replicated both findings on their own (G, I), albeit they had lower firing rates during the stimulus viewing period. (J) Ipsilateral placement of face stimuli eliminated any possible overlap with the RF. Yet, the ordering of the PSTHs remained the same as that in the main task, confirming that the reversal of T_{in} PSTH order with stimulus strength was not caused by inadvertent overlaps between the RFs and the face stimuli. These results also rule out explanations based on attentional mechanisms or visual selectivity of LIP cells that depend on the overlap of the stimulus and the RF. Further support for our conclusions comes from independent analyses of stimuli with different sizes, which changed by an octave across trials. Both large and small stimuli in the main task caused consistent T_{in} PSTH orders and curved manifolds (not shown). Forty-four units were recorded in 17 sessions of the control experiment. (K) The average firing rates in the ipsilateral control experiment decreased with stimulus strength for both T_{in} (Equation 1, $\alpha_1 = -0.15 \pm 0.05$, $p = 0.010$) and T_{out} choices ($\alpha_1 = -0.31 \pm 0.05$, $p = 4.6 \times 10^{-8}$). Error bars are SEM. (L) The population neural responses in the ipsilateral control experiment formed a curved manifold. The plot shows the neural activity at 400 ms after stimulus onset (100 ms window). The curvature of the manifold was significant ($p < 0.001$, based on tests of consistency of curvature direction; see [STAR Methods](#)).




Figure S3. A curved manifold in the PC space arises from the variation of tuning for stimulus strength across the recorded population, related to Figure 2

(A) A schematic description of our principal component analysis (PCA). Our dataset included neurons preferring either of the two targets. T_1 and T_2 indicate the targets on the right and left side of the screen, one of which in the RF. T_{in} and T_{out} indicate the targets inside and outside the neuron RF. Signed stimulus strength in the experiment was originally defined based on the location of the targets (T_1 versus T_2). For each neuron, we redefined the stimulus strength such that negative values matched T_{out} and positive values matched T_{in} . Doing so enabled us to analyze the recorded units as a population of neurons with similar target preferences (i.e., aligned T_{in}) (middle). In our main analyses, PCA was applied to the PSTHs after this alignment (right) to focus on the diversity of stimulus tuning across the population (see panel B). (B) When the preferences of neurons are aligned, a curved manifold arises from the variation of neural tuning across units. To demonstrate this point, we simulated six hypothetical scenarios for a neuron pair with opposite choice preferences (one preferring T_1 and the other T_2), plotting their joint manifolds before and after the T_{in} alignment. If the neurons have symmetric tuning for positive and negative stimulus strengths (slopes with similar magnitudes but opposite signs) and the tuning slopes are identical across the population, the manifold would be a straight line regardless of T_{in} alignment (B1-B2). If the single-unit tuning is asymmetric but similar across the population, the manifold would be curved before alignment (B3) but straight afterward (B4). However, T_{in} alignment does not remove the curvature caused by diverse tuning across neurons (B5-B6). (C) The T_{in} alignment focused our analyses on the diversity of responses across the population. However, all of our main conclusions replicate in the absence of T_{in} alignment too. The left panels show the population response manifolds without T_{in} alignment for a balanced population of neurons preferring T_1 and T_2 (see [STAR Methods](#)). The right panels show the manifolds with T_{in} alignment. (D-F) Context dependency of neural manifolds (see [Figures 5C, 5F, and 5H](#)) tested using the population without T_{in} alignment. The results are qualitatively similar to those with alignment.




Figure S4. The extent of the manifold curvature in LIP population responses exceeded the curvature expected purely from neural response variability, related to Figure 2

(A) To confirm the significance of the manifold curvature, we compared the magnitude of curvature in the actual neural responses with that of simulated data that had the same response variability as the actual data but were constrained to a linearized manifold (see STAR Methods). The linearized manifolds were generated by adjusting mean responses for each stimulus strength to be on an interpolating line between the responses to the strongest stimuli supporting the T_{in} and T_{out} choices. Population responses were then generated using these adjusted means while preserving the response variability across trials at the levels observed in the data. Response variabilities could bend the simulated manifolds due to finite sampling. However, these “noise”-induced curvatures were much smaller than those in the data (right column in the panel, $p < 0.001$; error bars are SEM.). The plots are based on firing rates in the same window as in Figures 2E and 2F (350–450 ms after stimulus onset). Colored lines are second-order polynomial fits to ten bootstrap samples of the observed data (left column) or ten simulated linearized population response manifolds (middle column). (B) Significant manifold curvatures were consistently observed in both monkeys performing the face task (monkey L, $n = 62$; monkey A, $n = 70$). (C) The curvature was also observed in individual sessions of the task when we could record enough units simultaneously to infer the manifold from single sessions (session 1, $n = 8$; session 2, $n = 10$). Conventions in B and C are the same as in Figure 2E. (D) Average PSTHs of well-isolated single units ($n = 75$) and their curved manifolds showed the same patterns obtained from all units.




Figure S5. Population neural responses throughout task epochs, related to Figure 2

(A) Population average PSTHs aligned to multiple task events. Those aligned to the stimulus onset are the same as Figures 2A and 2B. In the face task, the separation of the population average T_{in} and T_{out} PSTHs were smaller, and firing rate dynamics continued into the delay period after stimulus offset. However, in both tasks, T_{in} PSTHs converged to the same level before the saccade for all stimulus strengths, consistent with past studies (Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). (B) Around the time of the Go cue and saccade, the population neural responses in the PC state space converged to one of the two points that represented the monkey's choice. The PC space is identical to those in Figures 2E and 2F. Each panel shows the responses at a particular moment aligned to the Go cue (top) or saccade onset (bottom). (C) Population response manifolds expanded over time after stimulus onset, remaining largely aligned during decision formation but misaligned with the action encoding axis. After the stimulus offset, population response patterns for different stimulus strengths converge onto one of the two action-encoding points in the state space, compatible with the transformation of the DV encoding to choice encoding. Regions not occupied by neural responses on the saccade-aligned manifolds are shown as dashed lines and plotted only to convey misalignment with the manifolds at earlier times. (D) The trajectory of population neural responses for each stimulus strength in the same PC space, aligned to the stimulus onset and saccade onset for each task. Data points represent response patterns for different stimulus strengths at different times. Lines connect data points of the same stimulus strength over time. Later times are shown with larger points and more saturated colors.




Figure S6. Curved manifolds emerge when precise coding of continuous variables must happen with a limited dynamic range of neural responses, related to Figure 6

Various constraints limit the dynamic range of neurons in a biophysically plausible neural circuit (Wohrer et al., 2013), including but not limited to the non-negativity of firing rates and metabolic costs associated with spiking (Keemink and Machens, 2019; Lennie, 2003). These constraints limit the neural responses to a bounded subspace within the population response state space. (A) Encoding a continuous variable (e.g., the DV) along a straight line in the state space can achieve arbitrary high precision in the absence of constraints on the dynamic range of neural responses. (B) Compacting a straight linear manifold within the bounded subspace has two important downsides. First, the diagonal lines that provide the maximum dynamic range go through the center of the bounded subspace, where the whole population has high levels of activity —an energetically expensive code. Second, compacting the straight manifold makes the response patterns associated with different DVs less distinguishable, reducing the encoding precision. (C) Bending the manifold within the bounded subspace will reduce energetic costs (Keemink and Machens, 2019). Further, under a favorable correlation structure in the population responses (Moreno-Bote et al., 2014) and especially if decoding along a curved manifold is feasible, the curved manifold makes the population responses associated with neighboring DV values more distinguishable, increasing the encoding precision. (D) The presence of curvature is not changed by affine transformations of the bounding subspace, e.g., applying different scaling coefficients to the responses of individual units. However, the exact shape and layout in state space could be modified. (E) Although it is possible to further elongate the one-dimensional manifold by twisting and tangling it within the subspace, it may not be advantageous for two reasons. First, although such twisted manifolds allow larger separation of points along the manifold, they may bring distant points closer to each other in the state space, increasing readout errors due to noise. Second, a twisted manifold necessitates a more complex readout (decoder). Whereas a linear or low-order polynomial readout is easily implementable, more complex readouts may not be easy to implement or may lack the robustness offered by simpler decoders.




Figure S7. Neurons in the IT cortex show monotonic tuning for stimulus strength during the face categorization tasks, similar to the monotonic tuning of MT cells for motion strength, related to Figure 2

Firing rates of MT neurons are known to monotonically change as a function of motion coherence of the random dots stimuli. If sensory responses to our face stimuli show non-monotonic, complex tuning, integrating these responses might result in non-monotonic firing rates as observed in LIP. To exclude this possibility, we analyzed the activity of neurons in the IT cortex recorded simultaneously with LIP neurons during the face tasks. The recordings were made from face selective neural clusters in anterior IT (Tsao et al., 2003). The figure shows the population average firing rates of 199 face-selective neurons. Stimulus strengths were sorted according to individual neurons' preferred face category. The firing rates were monotonically modulated by the stimulus strength, providing evidence against the possibility that the reversal of firing rates in LIP is inherited from the activity of sensory neurons.