Article

Current Biology

Inactivation of primate cortex reveals inductive biases in visual learning

Highlights

- V4's role in form perception depends on recent training history
- Training with preferred stimuli led to a simple, scalar population readout
- Training with non-preferred stimuli led to a more distributed readout
- Reversible inactivation reveals the shift between scalar and distributed strategies

Authors

Pooya Laamerad, Matthew R. Krause, Daniel Guitton, Christopher C. Pack

Correspondence

pooya.laamerad@pennmedicine.upenn.edu

In brief

Laamerad et al. use reversible inactivation of primate area V4 to examine how the brain learns new perceptual tasks. The results reveal an inductive bias for sensory readouts that are easy to learn, even when they ignore task-relevant information that could theoretically be decoded from the neural population.

Laamerad et al., 2025, Current Biology *35*, 4699–4713
October 6, 2025
https://doi.org/10.1016/j.cub.2025.08.027

CellPress

Current Biology

CellPress

Article

Inactivation of primate cortex reveals inductive biases in visual learning

Pooya Laamerad,[1,2,3,*] Matthew R. Krause,[1] Daniel Guitton,[1] and Christopher C. Pack[1]
[1]Department of Neurology and Neurosurgery, Montreal Neurological Institute, McGill University, Montreal, QC H3A 2B4, Canada
[2]Department of Neuroscience, University of Pennsylvania, Philadelphia, PA 19104, USA
[3]Lead contact
*Correspondence: pooya.laamerad@pennmedicine.upenn.edu
https://doi.org/10.1016/j.cub.2025.08.027

SUMMARY

Humans and other primates are capable of learning to recognize new visual stimuli throughout their lifetimes. Most theoretical models assume that such learning occurs through the adjustment of the large number of synaptic weights connecting the visual cortex to downstream decision-making areas. While this approach to learning can optimize performance on behavioral tasks, it can also be costly in terms of time and energy. An alternative hypothesis is that the brain favors simpler learning rules that do not necessarily optimize the readout of information from visual cortical neurons. Here, we have examined these hypotheses by reversibly inactivating visual area V4 in non-human primates at different stages of training on form discrimination tasks. We find that V4 inactivation generally has a behavioral effect for only a subset of the stimuli that are encoded in the V4 population activity, specifically those that can be represented efficiently in the population firing rate. As a result, neural measures of discriminability do not necessarily predict the causal contribution of V4 neurons to task performance. This pattern of results can be explained by incorporating a strong inductive bias for simpler perceptual readouts into existing theoretical frameworks. Such a simplicity bias is suboptimal in the sense that it ignores information that could theoretically be extracted from the neural population, but it has the likely advantage of facilitating efficient learning on ecologically relevant timescales.

INTRODUCTION

Neurons along the ventral visual pathway of the primate cortex acquire exquisite stimulus selectivity during early development.[1] Although neural plasticity declines with age, stimulus representations in the visual cortex can still change well into adulthood.[2,3] Indeed, when adult subjects are trained on a new form recognition task, neurons in ventral cortical areas like V4[4,5] and the inferotemporal cortex[6,7] adjust their selectivity in such a way as to improve task performance.

Such changes are often modest relative to behavioral improvements,[5] and so it has been hypothesized that learning requires adjustments in the synaptic strength of projections from the ventral visual cortex to other brain areas. This is often called "reweighting,"[8] because the readout of sensory information from visual cortex changes, while the representation of the relevant stimuli in the visual cortex remains relatively fixed.[9] Models of perceptual learning, as well as artificial neural networks, generally rely on reweighting to explain changes in behavioral performance with training.[10] While these models often achieve impressive performance on visual tasks, a common criticism is that they rely on complex learning that requires the adjustment of many millions of parameters.

If learning involves a reweighting of visual cortical signals, it might also be shaped by the stimulus representations in each cortical area. Specifically, visual cortical areas generally exhibit *representational biases*, which are a tendency to devote more neurons to the encoding of stimuli that are commonly encountered in the environment. Examples include biases for cardinal orientations,[11] faces,[12] and expanding optic flow patterns,[13,14] to name a few. Theoretical models have suggested that these representational biases can serve as *inductive biases* that can simplify learning.[15–18]

Here, we have examined this possibility experimentally. We trained two non-human primates to discriminate between different classes of stimuli and used reversible inactivation to assess how neural activity was read out during perceptual decision-making. Our focus was on visual cortical area V4, which has a representational bias for curved or circular stimuli.[19,20]

We found that training on different behavioral tasks powerfully shaped the readout of V4 activity. In particular, training on a task that involved the stimulus preferred by a majority of neurons led to a readout that seemed to rely on total spiking activity in the local V4 population. This is a *simple* readout, because it does not require reweighting of the outputs of different neurons and indeed can be achieved with the adjustment of a single parameter, namely a threshold applied to the population response.

In contrast, training on a task that involved non-preferred stimuli led to a more conventional readout that made use of different V4 subpopulations, weighted by reliability. This readout is more complex, because it requires a reweighting of many different neuronal outputs according to their selectivity for the relevant

stimuli. Although the complex readout extracts more information from the neuronal population, we found that the brain generally preferred to rely on the simpler readout when it was relevant to the task.

We conclude that the representational biases in the visual cortex lead to strong inductive biases during the learning of new tasks. These inductive biases favor readout strategies that are simpler to learn, even though this approach potentially ignores much of the stimulus information that is encoded in the corresponding neural populations.[21] We speculate that a preference for simpler readout strategies renders perception susceptible to the kinds of biases that are widely reported for other cognitive operations.[22–25]

RESULTS

Cortical area V4 is organized into domains or clusters of neurons that are specialized for representing particular stimulus classes.[26–31] To thoroughly examine how these representational biases relate to behavior, we restricted our recordings to a single small patch of V4, roughly 2 × 2 mm, in each of two non-human primates (see STAR Methods and Figure S1A). We first characterized the representation of visual stimuli in these patches, in two experimentally naive animals. We then asked how the readout from these patches depends on stimulus representations and on learning.

V4 contains domains with a representational bias toward circular stimuli

Prior to training on any discrimination task, we assessed local neural selectivity by presenting various stimuli to the animals during passive fixation. The size and position of the stimuli were chosen to cover the receptive fields (RFs) of the neurons under study (Figure S1B). During this preliminary phase of the experiment (Figure 1A), we recorded from a total of 208 neurons (74 neurons in monkey 1 and 134 in monkey 2).

To assess selectivity, we presented a set of 152 diverse images spanning 19 categories, from synthetic shapes to natural images (see STAR Methods and Figure S1C). We then generated representational dissimilarity matrices (RDMs[32]) that characterize the pairwise differences in population firing patterns elicited by different stimuli (see STAR Methods for details). As shown in Figure 1C, the RDM revealed strong pairwise discrimination between circular patterns and all other stimuli (see STAR Methods, hierarchical clustering analysis).

This latter finding is consistent with previous work,[19,27,30] which has also found that many local V4 domains exhibit a preference for curved or circular patterns. In our data (Figure S1D), this representational bias was most pronounced in response to low spatial-frequency grating patterns: circular grating stimuli elicited higher firing rates than the orthogonal stimulus, a radial pattern (Figure 1B). Overall, 74% (55/74) of the neurons in monkey 1 and 66% (88/134) of the neurons in monkey 2 preferred circular gratings over radial gratings, leading to a strong circular preference at the population level in both monkeys (monkey 1, $F(3,70) = 9.13$, $p < 0.01$; monkey 2, $F(3,132) = 6.21$, $p < 0.01$; ANOVA, false discovery rate [FDR]-corrected).

Subjects exploit biased V4 representational biases in a perceptual decision-making task

To determine how representational biases in V4 shaped learning, we trained both animals to perform a delayed match-to-sample (DMS) task (Figure 1D) that involved discriminating between circular and radial gratings. On each trial, the animals maintained fixation while a noisy sample stimulus was presented in the RFs of the neurons being recorded. After a brief delay, the animals had to saccade to the unpredictable location of the stimulus that matched the sample. In phase 1 of the experiment, we used gratings of low spatial frequency, specifically 0.75 cycles per degree (cpd) for monkey 1 and 1.0 cpd for monkey 2, as these generated the largest responses in V4 (Figure 1B). Average neurometric discriminability of these stimuli for individual V4 neurons is shown in Figure S2B.

After several weeks of training (70 sessions for monkey 1 and 39 sessions for monkey 2), both animals learned to perform this task accurately and without a behavioral bias for either stimulus (see STAR Methods and training period inset in Figure 2A). We then probed the readout of neuronal information by injecting muscimol into the targeted V4 domains. Muscimol is a powerful GABA receptor agonist, and previous work has shown that injecting this volume into visual areas abolishes spiking activity over a radius of 1–2 mm,[33–35] sufficient to cover the stimulus-selective domains in V4.[30] In each session (10 sessions in monkey 1 and 4 sessions in monkey 2), we verified that neural activity was silenced near the site of injection (mean firing rate before injection, 14.34 ± 3.15; 45 min after injection, 0.21 ± 0.02; $p < 0.001$, Wilcoxon rank sum [WRS] test).

Figure 2A (black lines) shows the behavioral results for each animal, with the percentage of radial choices represented on the *y* axis of each plot. Following inactivation with muscimol, a bias toward radial choices appeared, starting 45 min after injection (orange lines), peaking 18 h later (purple lines) and returning to baseline after 2 days (gray lines). A similar result was observed consistently across individual inactivation sessions (see Figure S2A). Considering the two grating types separately (see STAR Methods), muscimol increased the behavioral threshold for circular gratings by an average of 51% compared with the pre-injection baseline (45 min, 46% ± 14% for monkey 1 and 37% ± 18% for monkey 2; 18 h, 66% ± 12% for monkey 1 and 57% ± 20% for monkey 2; $p < 0.05$ for both monkeys at 45 min and 18 h, WRS test Figure 2B). In contrast, for radial grating stimuli, the behavioral threshold changed by an average of only 2% ± 3% (45 min, 2% ± 3% for monkey 1 and 7% ± 5% for monkey 2; 18 h, 0% ± 2% for monkey 1 and −1% ± 2% for monkey 2; $p > 0.05$ for both monkeys at 45 min and 18 h, WRS test Figure 2B). Thus, both animals developed a significant bias toward radial grating stimuli after inactivation (Figure 2C; monkey 1, mean bias difference = 0.34 ± 0.06, $p = 0.02$; monkey 2, mean bias difference = 0.44 ± 0.1, $p = 0.026$, WRS test).

Interestingly, inactivation of V4 had little effect on performance for noisy stimuli (0%–8% SNR; Figure 2C). Because the least noisy stimuli elicited the strongest firing rate preference (Figure S2C), the results are consistent with a simple readout strategy that specifically exploits the preference for circular stimuli in V4, as explored further in the next section. Note that the

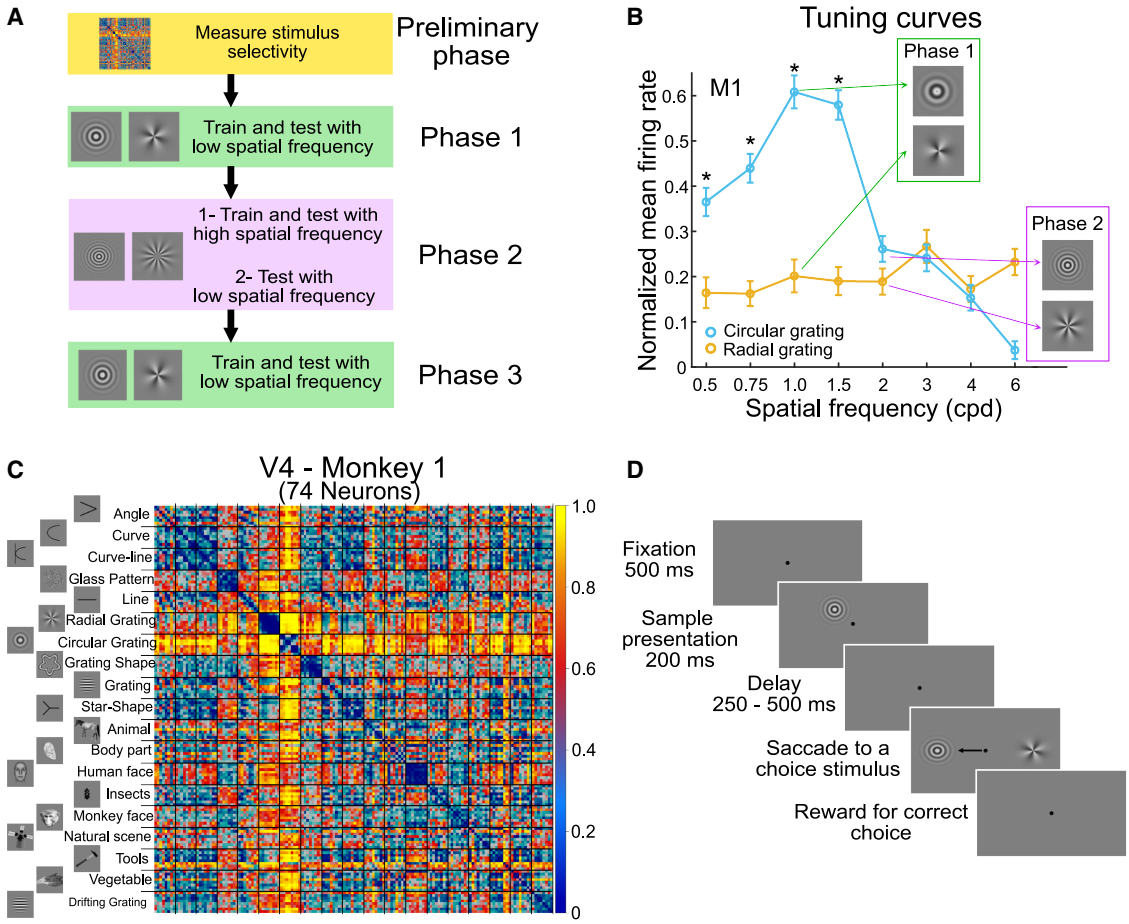

Figure 1. Experimental phases and V4 neural coding

(A) Training phases: following the preliminary phase identifying V4 neuronal selectivity, in phase 1, animals trained and were tested with low-frequency gratings. In phase 2, they trained with high-frequency gratings and were tested first with high and then low frequencies during inactivation. In phase 3 (monkey 1 only), the animal retrained with low-frequency gratings and was tested during inactivation.

(B) Tuning curves across spatial frequencies for noiseless circular and radial gratings in monkey 1. Circular gratings consistently elicited stronger responses at low frequencies, with the largest difference at 1 cpd. Low-frequency gratings were used during phase 1 training (green arrows) and high-frequency gratings during phase 2 (purple arrows).

(C) Representational dissimilarity matrices (RDMs) for V4. Each element color codes the dissimilarity between two neuronal response patterns, with yellow indicating high dissimilarity and blue high similarity.

(D) Delayed match-to-sample task: animals maintained fixation for 500 ms before stimulus onset. After a 200 ms sample (noisy circular or radial grating), a randomized delay (250–500 ms) followed. The task concluded with circular and radial gratings appearing randomly on either side of the fixation point. Error bars indicate the standard error of the mean (SEM). Asterisks denote statistically significant differences. Cpd: cycles per degree.

See also Figure S1.

apparently uniform upward shift of the logistic functions for monkey 2 (Figure 2A, right) is a consequence of asymmetry in the baseline behavioral data (Figure 2C, right).

To control for non-specific effects of the inactivation procedure, we performed a series of control experiments (Figure S3). We first verified that behavioral impairments were not due to the injection per se, since injecting only saline into the same V4 region had no effect on the animals' performance ($p > 0.05$, Wilcoxon signed-rank [WSR] test, Figure S3A). Behavioral thresholds were also not affected by muscimol injections when the visual stimuli were displaced to a different retinal location, indicating that the effects were specific to the RFs of the inactivated V4 neurons ($p > 0.05$, WSR test, Figures S3B and S3C).

A scalar readout strategy can explain perceptual biases

The results shown in Figure 2 are reminiscent of the predictions of *scalar readout* models,[17,36,37] in which the presence or absence of a stimulus is inferred from the total neural population response. For our task, a large V4 population response is associated with circular stimuli and a small population response with radial stimuli (Figure 1B).[17] Relying on this asymmetry can simplify the process of learning the DMS task; indeed, in the limit the only parameter that has to be optimized is a threshold applied to the total activity in the V4 population[17] (Figure 3A). This kind of simple model is distinct from the *distributed readout* models typically used in perceptual learning and decision-making[38] (Figure 3B), which adjust the readout weights of

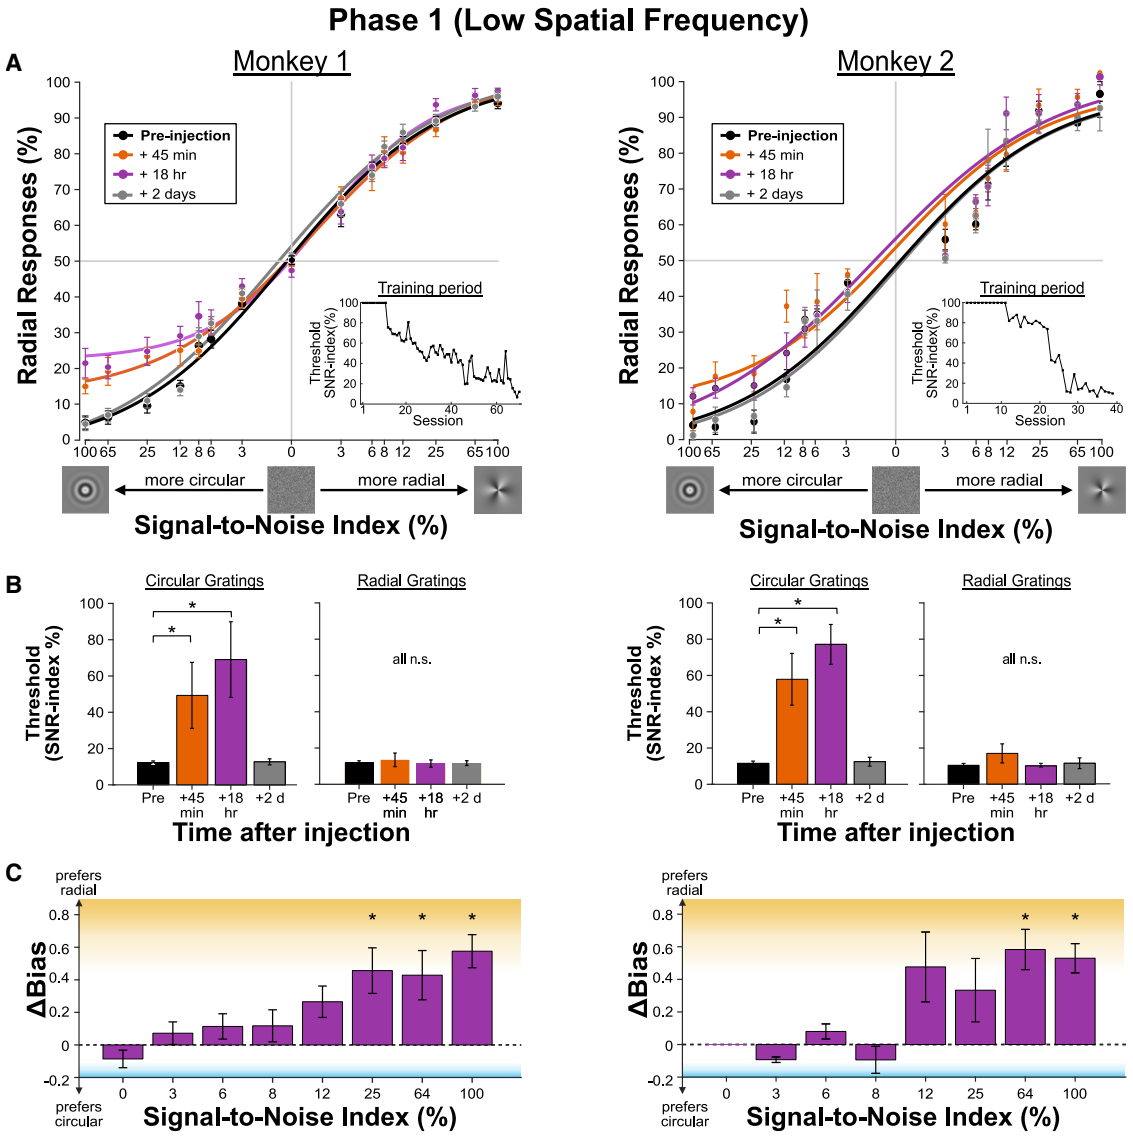

Figure 2. Behavioral impact of V4 inactivation in phase 1

(A) Psychometric functions in phase 1. Psychometric functions fitted to behavioral data at various time points after inactivation show that inactivation of V4 neurons in both animals substantially impaired detection of low spatial frequency circular gratings. No behavioral deficits were observed at any subsequent time points after inactivation for radial gratings. For visualization purposes, solid lines show fits of a logistic function to the behavioral data at different time points relative to muscimol inactivation.

(B) Behavioral thresholds increased for circular gratings following inactivation. The left images depict changes in psychometric thresholds for circular gratings; the right images depict changes for radial gratings. Inset: the psychophysical behavioral thresholds across training sessions, showing the gradual improvement in performance leading up to stable task acquisition prior to muscimol inactivation.

(C) Behavioral bias changed after injection of muscimol. We calculated the difference in behavioral bias between 18 h post-injection and pre-injection sessions across all noise levels. For monkey 1, significant bias differences were observed at noise levels of 100%, 65%, and 25% ($p < 0.05$, WSR test). For monkey 2, significant differences were observed at noise levels of 100% and 65% ($p < 0.05$, WSR test). Positive values in this plot represent a bias toward radial gratings and negative values represent a bias toward circular gratings. Error bars indicate the SEM. Asterisks denote statistically significant differences ($p < 0.05$).

See also Figure S2.

many neurons across the population in such a way as to eliminate biases and to optimize the readout of sensory information.[17,39–41]

One prediction of the scalar readout model is that reducing the overall population activity in V4 should lead to a bias toward radial stimuli.[17] This can be accomplished by lowering the contrast of the stimuli, which has the effect of reducing firing rates across the visual cortex,[42] albeit not as powerfully as muscimol. For a scalar readout of circular stimuli, reducing contrast should lead to a perceptual bias toward radial stimuli (Figure 3A). In contrast, for a more distributed readout (Figure 3B), lower contrast should lead to decreased performance for both classes of stimuli. We therefore tested this hypothesis by performing additional experiments, in which

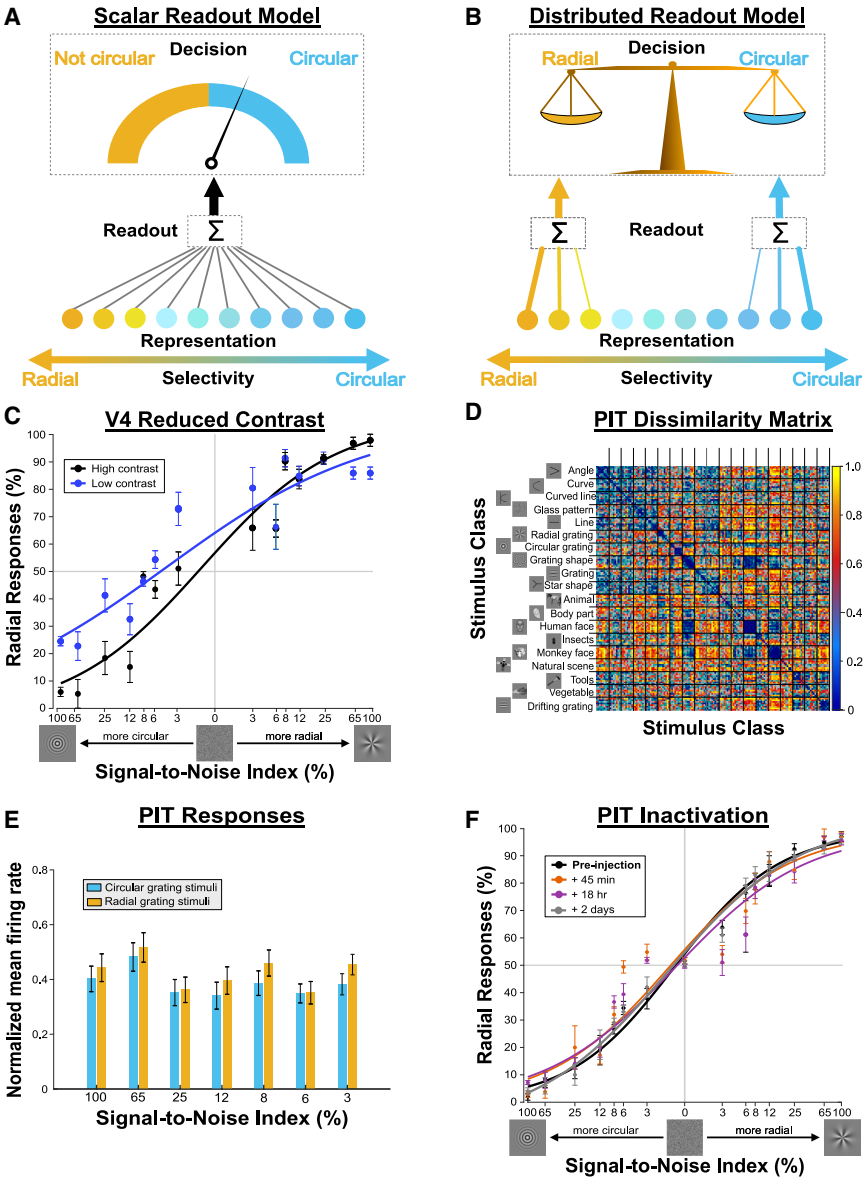

Figure 3. Readout models, impact of contrast reduction and PIT inactivation

(A) In the scalar readout model, perceptual decisions are based on the total population response of V4. Because more V4 neurons prefer circular gratings than radial ones, the brain can set a decision criterion based on the overall firing rate: when total neural activity exceeds this criterion, the stimulus is perceived as a circular grating; when activity falls below the threshold, it is perceived as a radial grating.

(B) In the distributed readout model, perceptual decisions are based on the selective pooling of V4 neurons tuned to different stimuli (circular-selective neurons shown in blue and radial-selective neurons in orange). To optimize discrimination, the brain learns to assign greater weights to neurons with stronger selectivity, with line thickness representing the strength of these weights.

(C) Effect of reduced contrast on grating discrimination. This image shows the average behavioral performance across three sessions from a contrast manipulation experiment, in which high-contrast (100%) and low-contrast (50%) low spatial frequency gratings were interleaved in each session.

(D) RDM for PIT. Unlike V4, PIT exhibited slightly greater selectivity for natural images.

(E) PIT showed no significant differences in neuronal responses between the two grating types (circular vs. radial) at any noise level ($p > 0.05$, t test).

(F) Behavioral impact of PIT inactivation. Muscimol inactivation of PIT neurons (6 sessions) resulted in a small but consistent behavioral deficit for both types of gratings. However, the impact was more pronounced for noisier stimuli. Error bars indicate the SEM.

See also Figure S3.

we interleaved high-contrast (100%) and low-contrast (50%) stimuli during three behavioral sessions.

As shown in Figure 3C, reducing stimulus contrast produced a significant bias toward radial choices for low-contrast images (mean bias difference = 0.18 ± 0.07, $p = 0.041$, WRS test), similar to what we observed during V4 inactivation. This is again consistent with a readout that relies on total firing rates in V4 and highlights the vulnerability of this strategy to changes in the stimulus statistics.[43] As shown below, the brain is actually capable of learning a distributed readout from the same V4 population but apparently has an inductive bias for the simpler strategy afforded by the population response.

Different readouts for different brain regions in the same task

Our results thus far suggest that the firing rate preference in V4 shapes the readout of sensory information by decision-related brain regions. However, the opposite direction of causation is also possible: perhaps the particular behavioral salience of circular stimuli[44] causes them to be read out preferentially.[45] In that case, one would expect to find a similar readout strategy elsewhere in the brain.

To examine this issue, we repeated the inactivation experiment in a second area, the posterior inferotemporal (PIT) cortex. We first recorded neural activity and confirmed that this area has a general capacity to discriminate natural images[46] (Figure 3D) but no firing rate preference for circular or radial gratings (Figure 3E). As shown in Figure 3F, inactivation of PIT (6 sessions) caused a modest reduction in performance 18 h after injection. Thresholds increased by 8.7% ± 3.2% for circular gratings ($p > 0.05$, WRS test) and 10.2% ± 4.0% for radial gratings ($p > 0.05$, WRS test). However, in contrast to the effects of V4 inactivation, the effects were similar for the two stimulus classes, and there was no significant perceptual bias introduced by inactivation ($p > 0.05$, WRS test). Interestingly, muscimol inactivation of PIT led to the strongest impairments for the noisiest stimuli (SNR 3%, mean deficit = 13.25% ± 3.1%, $p < 0.01$; SNR 6%,

mean deficit = 13.8% ± 3.5%, $p = 0.02$, sign test; Figure 3F), in contrast to V4 inactivation, which most strongly affected performance for the least noisy stimuli (Figure 2C).

These results suggest that the readout makes use of different readout strategies, depending on the cortical representation of the trained stimuli.[47–49] We next examined this idea in detail.

V4's contribution to the discrimination task is much weaker for other kinds of stimuli

To the extent that representational biases shape visual learning, we should be able to alter the readout of sensory information by changing the stimuli used during training. To examine this possibility directly, we retrained the animals with gratings of higher spatial frequencies (2 cpd), across 54 sessions for monkey 1 and 48 sessions for monkey 2 (see STAR Methods and training period inset in Figure 4A). For these stimuli, the average responses of V4 neurons were similar for circular and radial patterns (Figure 1B; $p > 0.05$, t test), and the neuronal population exhibited a balanced preference (monkey 1, 55% preferred high spatial frequency circular gratings, 45% radial; monkey 2, 52% circular, 48% radial). Nevertheless, neural discriminability was largely independent of spatial frequency—a support vector machine (SVM) classifier trained on the V4 population achieved a mean accuracy of 70.24% ± 0.54% for low spatial frequencies and 68.47% ± 0.69% for high spatial frequencies, with no significant difference between conditions ($p > 0.05$, t test; see STAR Methods). Corresponding neurometric data are shown in Figure S4B.

After a few weeks of training on the high spatial frequency task (phase 2; Figure 1A), the readout strategies were again probed with muscimol inactivation and neuronal recordings at the same sites as in phase 1 (4 sessions in monkey 1 and 3 sessions in monkey 2). RF eccentricities were not significantly different between the two phases (1.6° ± 0.7° in phase 1 and 1.5° ± 0.6° in phase 2 for monkey 1, $p < 0.05$, WRS test; 2.4° ± 0.9° in phase 1 and 2.6° ± 0.8° in phase 2 for monkey 2, $p < 0.05$, WRS test).

Surprisingly, despite the robust discriminability of the relevant stimuli in V4 (Figure 1C), muscimol inactivation caused negligible impairments in performance on the high spatial frequency grating discrimination task (Figures 4A and S4A; mean change in threshold values relative to pre-injection were as follows: 45 min, 2.31% ± 1.32% for monkey 1 and 9.23% ± 5.19% for monkey 2; 18 h, 2.47% ± 2.16% for monkey 1 and 13.81% ± 6.28% for monkey 2). There was no discernible threshold change after inactivation in monkey 1 and a weak effect in monkey 2 that did not reach significance ($p > 0.05$, WSR test, monkey 1 and monkey 2).

These results were similar for the two stimulus classes. For circular gratings, the threshold increase at 45 min and 18 h post-injection was 5% ± 10% and 7% ± 8% (45 min, 4% ± 5% for monkey 1 and 9% ± 4% for monkey 2; 18 h, −1% ± 3% for monkey 1 and 14% ± 6% for monkey 2; Figure 4B). A similar pattern was observed for radial gratings (45 min, 1% ± 3% for monkey 1 and 10% ± 6% for monkey 2; 18 h, 6% ± 5% for monkey 1 and 13% ± 4% for monkey 2; Figure 4B). As a result, the animals did not develop a bias toward either grating type with inactivation (Figure 4C). These results suggest that the readout had largely switched to some another brain region[33,50] or was distributed across many brain regions,[51,52] in such a way that the contribution of V4 was unimportant. The latter would be likely if other brain regions exhibited a similarly unbiased encoding of high-frequency grating stimuli. In any case, the role of V4 in decision-making appeared to be defined by the representational bias in the local population response.

To further examine this possibility, we retested the animals with the preferred (low spatial frequency) stimuli, immediately after training and data collection with high spatial frequency stimuli. Single-unit recordings revealed that the V4 population maintained a strong preference for low spatial frequencies and circular gratings in this phase, as indicated by the difference in mean normalized activity for low spatial frequency circular and radial gratings (preliminary phase, 0.17 ± 0.037; phase 1, 0.18 ± 0.041; phase 2, 0.16 ± 0.052; $p < 0.05$ for all phases, t test).

Although these stimuli were identical to those used in phase 1, muscimol inactivation during phase 2 led to significant deficits in performance for *both* grating types (5 sessions for monkey 1 and 5 sessions for monkey 2) (Figures 5A and S5B). For circular gratings, the average threshold increase 45 min after injection was 37% ± 10%, rising to 59% ± 8% at 18 h post-injection (45 min, 49% ± 9% for monkey 1 and 25% ± 7% for monkey 2; 18 h, 65% ± 6% for monkey 1 and 52% ± 12% for monkey 2; Figure 5B). Similarly, for radial gratings, inactivation caused a significant decrease in performance, with the average threshold increase 45 min post-injection being 45% ± 17%, increasing to 57% ± 7% at 18 h (45 min, 39% ± 16% for monkey 1 and 54% ± 14% for monkey 2; 18 h, 64% ± 5% for monkey 1 and 48% ± 11% for monkey 2; Figure 5B). Corresponding neurometric data are shown in Figure S5A.

Thus, while there were large increases in behavioral threshold for both monkeys (mean increase 18 h after injection: 58% ± 5% for monkey 1 and 55% ± 9% for monkey 2, $p < 0.05$, WSR test), there was no bias toward either grating type at any time point after injection (monkey 1, $p > 0.05$; monkey 2, $p > 0.05$, WRS test; Figure 5C). This suggests that training with non-preferred stimuli led to the development of a distributed readout strategy (Figure 3B) more consistent with typical discrimination models.[53] Because this readout was shaped by training with non-preferred stimuli, it was likely suboptimal for the low spatial frequency gratings, and indeed behavioral performance on these stimuli was worse in phase 2 than in phase 1 (mean behavioral d' across all noise levels [monkey 1], phase 1 = 1.54 ± 0.81, phase 2 = 1.21 ± 0.98, $p = 0.024$; mean d' [monkey 2], phase 1 = 1.31 ± 0.83, phase 2 = 1.16 ± 0.94, $p = 0.011$, t test). Nevertheless, the causal impact of V4 remained apparent only for the preferred (low spatial frequency) stimuli.

Overall, these results indicate that the effects of causally manipulating neural activity are better predicted by representational biases than by standard measures of discriminability. To examine this issue at the neuronal level, we next attempted to estimate the readout weights of the V4 population across experimental phases.

Single-neuron readout weights reveal a change in strategy with training

The actual contributions of individual V4 neurons to perceptual decisions cannot be measured directly in our experiments, but under reasonable assumptions they can be inferred from

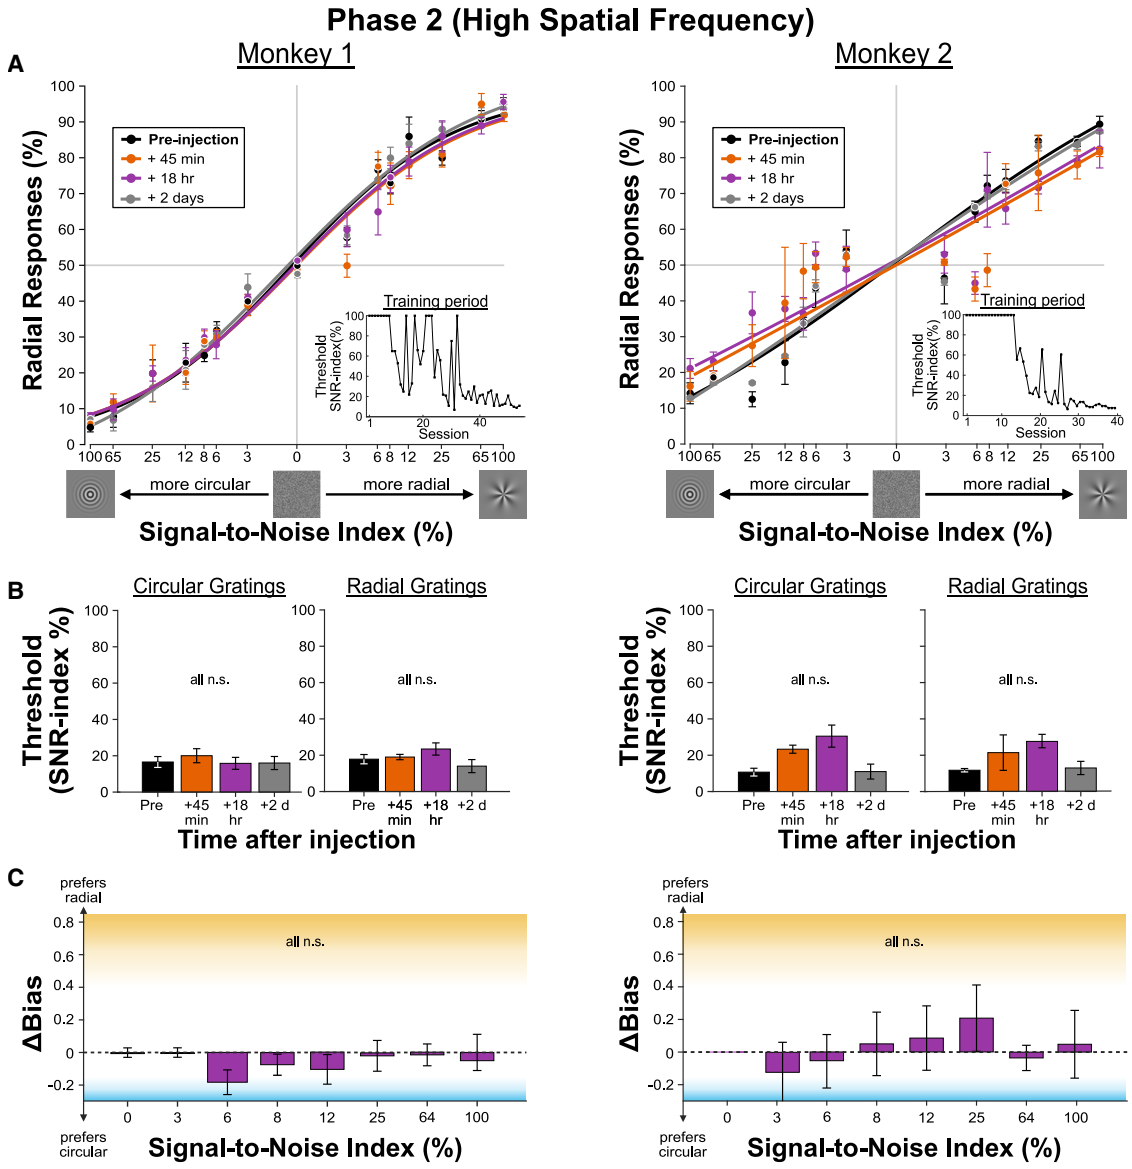

Figure 4. Behavioral impact of V4 inactivation in phase 2 for high spatial frequency stimuli

(A) Psychometric functions in phase 2. Following training with high spatial frequency gratings, V4 inactivation had little effect on behavioral performance. Inset: the psychophysical behavioral threshold across training sessions, showing the gradual improvement prior to muscimol inactivation.

(B) Behavioral thresholds did not change significantly for either grating type in either animal ($p > 0.05$, WSR test).

(C) Behavioral biases were similar between 18 h post-injection and pre-injection sessions across all noise levels for both animals ($p > 0.05$, WSR test).

Error bars indiate the SEM.

See also Figure S4.

correlations between neural firing and behavioral responses.[53] Specifically, the readout weight for each neuron can be estimated from its choice probability (CP) and its noise correlations with other neurons.[54–56] CP measures the extent to which the variability in a neuron's responses predicts an animal's behavioral choices,[53] while noise correlations are shared variability across neurons. For these analyses, we combined data across animals, although the results were similar at the individual level (Figure S6A).

Figure 6A shows the readout weights recovered for neurons with different stimulus preferences across different experimental phases, using the method devised by Haefner and colleagues.[55] In phase 1, readout weights were significantly positive for circular-preferring neurons and significantly negative for radial-preferring neurons (Table S1; F [5,269] = 4.77, $p < 0.01$; ANOVA, FDR-corrected). The negative weights mean that higher firing rates in this population were associated with a *lower* likelihood that the animals would report seeing its preferred (radial) stimulus. This is consistent with the idea that higher firing rates in the population response were used to infer the presence of a circular grating,[17] irrespective of the preferences of any individual neuron (Figure S6B).

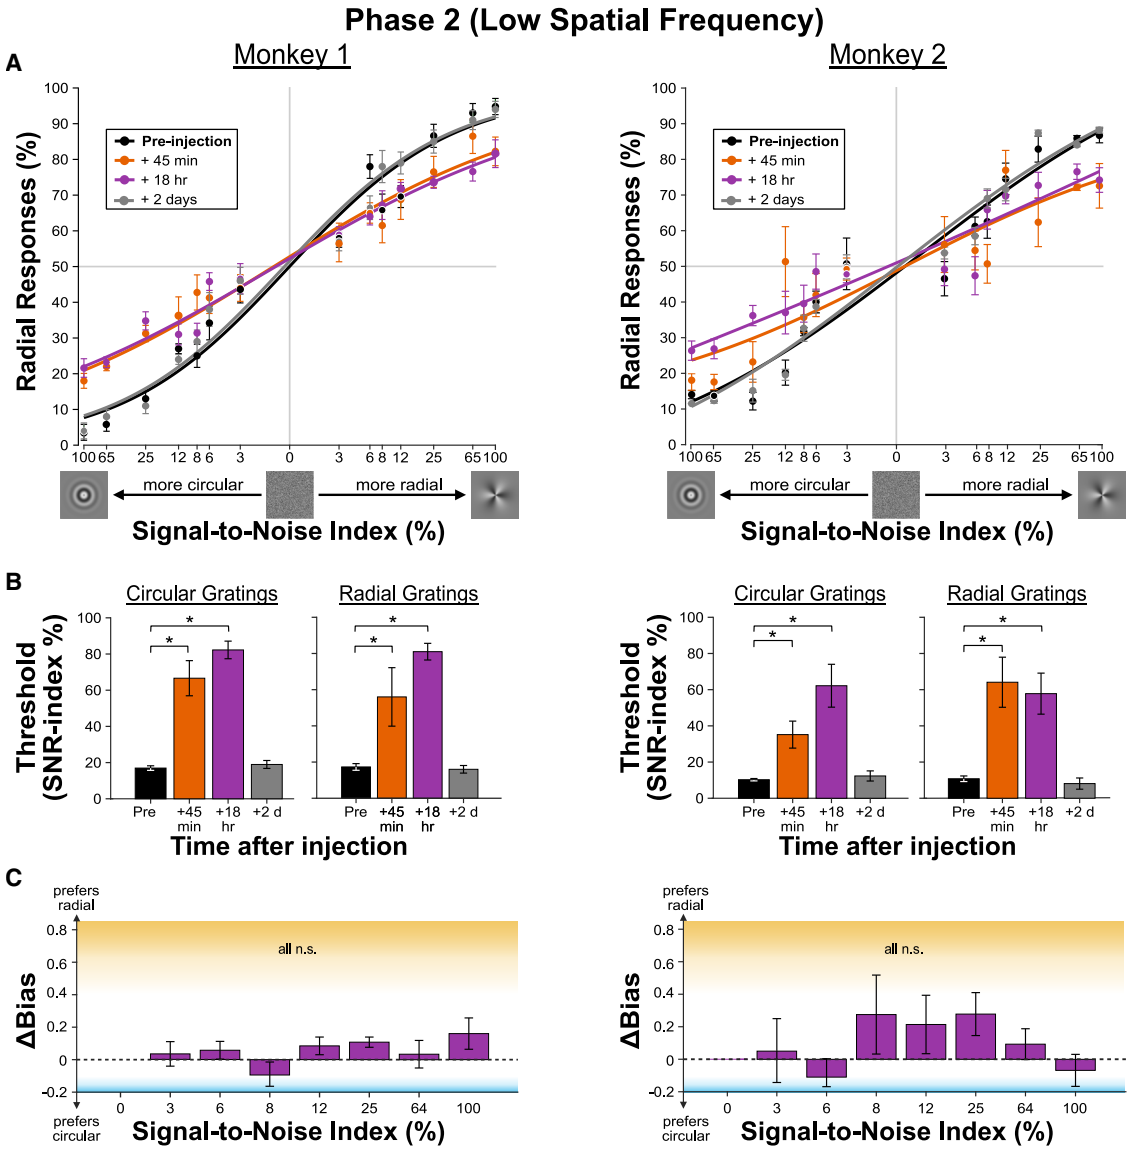

Figure 5. Behavioral impact of V4 inactivation in phase 2 for low spatial frequency stimuli

(A) Psychometric functions for low spatial frequency stimuli in phase 2. Following training with high spatial frequency gratings, when the animals were tested with low spatial frequency gratings, inactivation impaired the detection of both circular and radial grating images in both animals.

(B) Behavioral thresholds increased significantly for both types of gratings after inactivation ($p < 0.05$, WSR test, monkey 1 and monkey 2).

(C) There was no difference in behavioral bias between 18 h post-injection and pre-injection sessions across all noise levels in both animals ($p > 0.05$, WSR test). Error bars indicate the SEM. Asterisks denote statistically significant differences.

See also Figure S5.

In phase 2, however, the readout weights for both circular- and radial-preferring neurons became significantly positive for both low and high spatial frequencies (Table S1), indicating a shift toward a more distributed discrimination strategy (Figure S6B). The increase in readout weights for radial-preferring neurons from phase 1 to phase 2 was statistically significant ($F[5,269] = 4.77$, $p < 0.05$; ANOVA, FDR-corrected).

The changes in readout weights were mostly attributable to changes in CP, although there were changes in noise correlations as well (Table S1; Figure S6C). The CP patterns mirrored the readout weights (Figure 6B). In phase 1, radial-selective neurons had CPs below chance level (0.5), while in phase 2, their CPs significantly increased above 0.5 ($F[5,269] = 5.26$, $p < 0.05$; ANOVA, FDR-corrected). In contrast, circular-selective neurons maintained CPs above 0.5 across all phases.

Of particular relevance to the nature of the V4 readout was the correlation between each neuron's CP and its ability to discriminate between the two stimuli[9,55,56] (d'). In phase 1, this was positive for the neurons with circular preferences ($r = 0.47$, $p < 0.01$, Pearson correlation; Figure 6C) but negative for the radial-preferring neurons ($r = -0.34$, $p = 0.02$, Pearson correlation; Figure 6C), again suggesting a scalar readout strategy based only on

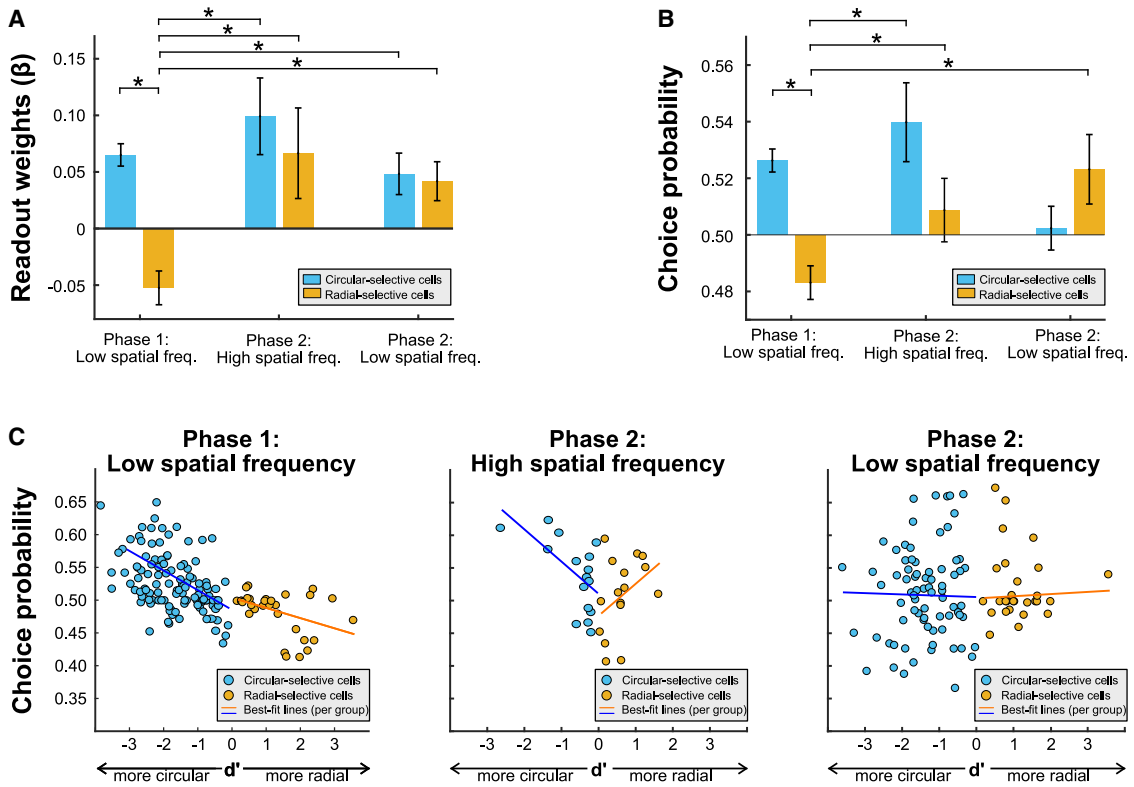

Figure 6. Phase-dependent changes in V4 neuronal readout weights, choice probabilities, and noise correlations

(A) Mean readout weights of V4 neurons: readout weights were calculated for neurons with different stimulus preferences across experimental phases. In phase 1, radial-selective neurons (orange) had negative weights, which were significantly different from those of circular-selective neurons (blue) ($p < 0.05$; ANOVA, FDR-corrected). In phase 2, there was no significant difference between the two groups for both low and high spatial frequency stimuli.

(B) Mean CPs of V4 neurons: changes in CPs mirrored the changes seen in the readout weights.

(C) Relationship between CPs and d′: in phase 1, radial-selective neurons (orange) showed a negative relationship between CPs and d′, while circular-selective neurons (blue) exhibited a positive relationship. In phase 2, this relationship was positive for both groups, for both low and high spatial frequency stimuli. Error bars indicate the SEM. Asterisks denote statistically significant differences.

See also Figure S6; Table S1.

population firing rate. By contrast, in phase 2, for high spatial frequencies, the correlation between CP and d' was positive for both circular- and radial-selective neurons, although it reached statistical significance only for circular-selective neurons (r = 0.6, p = 0.015, Pearson correlation; Figure 6C) and not for radial-selective neurons (r = 0.39, $p > 0.05$, Pearson correlation; Figure 6C), possibly due to the small sample size.

Thus, between phase 1 and phase 2, the readout from V4 began to weight neurons according to their selectivity for both circular and radial gratings, as found in previous studies.[9,57] As suggested by the inactivation results (Figure 5), this distributed weighting also influenced the perceptual readout for low spatial frequency gratings, even though there was no relationship between CP and d' for these stimuli (r = 0.08, $p > 0.05$ for circular-selective neurons; r = 0.07, $p > 0.05$ for radial-selective neurons, Pearson correlation; Figure 6C).

In contrast to the large changes in readout weights, we found little evidence for a change in single-neuron stimulus representations with training. For low spatial frequency gratings, the average neurometric sensitivity was typically lower than behavioral sensitivity (Figure S5B; monkey 1, neurometric threshold = 28.63 ± 11.68, behavioral threshold = 5.32 ± 2.11, p = 0.012; monkey 2, neurometric threshold = 30.47 ± 14.21, behavioral threshold = 6.69 ± 3.57, p = 0.009; WSR test), and neurometric thresholds did not change across experimental phases ($p > 0.05$ for both monkeys, WSR test).

Decision readouts prefer to rely on representational biases when possible

Finally, we considered the possibility that the readout formed in phase 2 was a cumulative result of training in phases 1 and 2, rather than a reaction to the specific stimuli used in phase 2. Indeed, some artificial neural networks learn low-dimensional solutions early in training, which is refined after further exposure to the relevant inputs.[16,43,58] We therefore retrained one animal for 10 sessions (see training period inset in Figure 7B) in a third phase, which was identical to phase 1, being comprised only of exposure to low spatial frequency gratings. We then inactivated the same population of V4 neurons.

As shown in Figure 7A, the effect of inactivation on radial grating trials decreased rapidly across sessions, while it remained consistent for circular grating trials. By the final session (Figures 7A and 7B), threshold increases were significant for circular gratings at both 45 min and 18 h post-injection compared

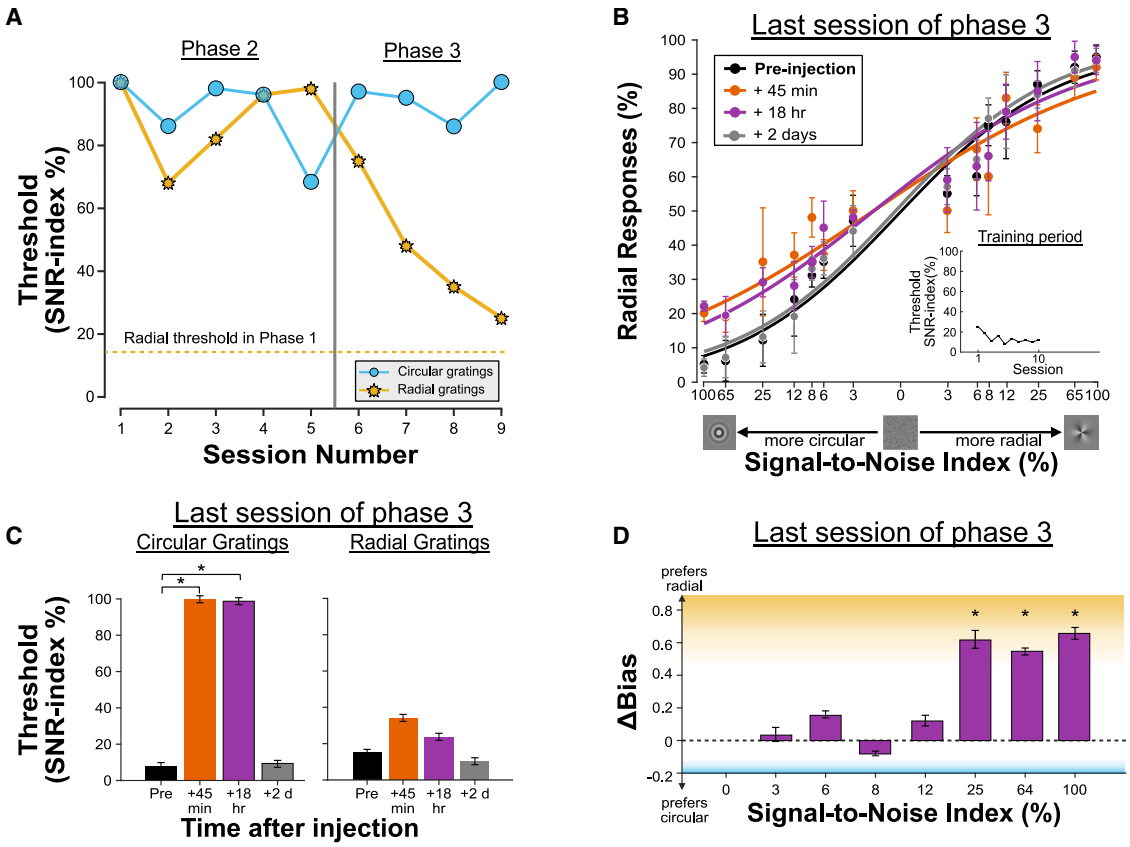

Figure 7. Behavioral impact of V4 inactivation in phase 3

(A) For monkey 1, psychophysical thresholds are plotted against the number of inactivation sessions, comparing phase 2 and phase 3. In phase 3, for four sessions, we inactivated V4 neurons while testing the monkey with low spatial frequency gratings (sessions after the vertical line). The effect of inactivation remained constant for circular gratings, but decreased session by session for radial gratings, with thresholds approaching those seen in phase 1 (dashed line).

(B) Fitted psychometric functions of the last session of inactivation in phase 3. Inset: the psychophysical behavioral threshold across training sessions, prior to muscimol inactivation.

(C) Behavioral thresholds increased significantly for circular gratings following inactivation, consistent with the pattern observed in phase 1, while thresholds for radial gratings remained largely unchanged.

(D) A similar pattern of behavioral bias was observed between pre-injection and 18-h post-injection. Error bars represent standard errors computed from bootstrapped data (see STAR Methods). Asterisks indicate statistically significant differences ($p < 0.05$, permutation test).

with pre-injection levels ($p < 0.05$, permutation test), with no significant change observed for radial grating stimuli at any time point after injection ($p > 0.05$, permutation test). As a result, during phase 3, the animal developed a bias toward radial grating stimuli following muscimol injection. By the final session of this phase, both the behavioral threshold pattern and the magnitude of the bias closely resembled those observed in phase 1 (Figures 7C and 7D). As in phase 1, the radial bias at 18 h post-injection was strongest for the stimuli with the lowest noise levels (100%, 65%, and 25% SNR) ($p = 0.032$, permutation test). Thus, the readout strategy, as assessed with muscimol inactivation, had reverted to the one used in phase 1, suggesting that the cortex exploits population representational biases when they are available, regardless of the total amount of time spent on training.

DISCUSSION

In this work, we have probed the readout of visual information from area V4 of the primate visual cortex. V4 has a representational bias for certain classes of stimuli, and our results demonstrate that this specialization powerfully affects the readout used in behavioral tasks. In particular, the brain learns to exploit representational biases (Figure 1B) to generate a scalar readout of preferred stimuli. As a result, reducing neural activity in V4 leads to predictable behavioral biases (Figures 2 and 3C), which are not seen with inactivation of nearby area PIT (Figure 3F). For non-preferred stimuli, a more conventional, distributed readout of V4 activity forms, but it has little impact on behavior (Figure 4). Thus, the contribution of V4 to behavior and perception appears to be determined in large part by representational biases (Figures 5 and 7), as confirmed by single-neuron correlations with behavioral choices (Figure 6).

Comparison with previous studies

A number of studies have created permanent lesions in V4, revealing deficits in the perception of color,[59,60] three-dimensional objects,[61] and textures.[62] Effects on the perception of two-dimensional features are somewhat variable,[60,63–65] but

this likely reflects the extent of the lesion relative to the scale of functional domains in V4.[66]

Other studies have used temporary methods to manipulate neural activity in V4. In one case, cryogenic inactivation of larger (centimeter-scale) domains caused reliable form discrimination deficits,[67] though this was not related to the underlying neural selectivity. Another study used electrical stimulation to bias the perception of depth from retinal disparity.[68] This latter result is similar to what is often observed in microstimulation studies of other areas,[69] and it bears a superficial resemblance to the behavioral biases we have observed in V4 (Figure 2). However, these previous studies have microstimulated much smaller (100 micron-scale), feature-selective columns[69] than the millimeter-scale domains that we have inactivated. Microstimulation of larger (millimeter-scale) domains generally causes balanced deficits, rather than biases toward particular stimuli.[70]

As shown in Figure 3F, we also inactivated a second region, area PIT. Although we did not attempt to tailor the stimuli to the preferences of the local stimulus domain, we observed small but reliable effects on behavior. These were strongest for the noisiest stimuli, which suggests that this area could play a role in spatial integration, as found in other areas.[71] In the mouse cortex, high-level areas perform an analogous function in tasks that require temporal integration.[72]

Rajalingham and DiCarlo[34] performed a more thorough study of local IT domains and concluded that the behavioral effects of muscimol inactivation were best predicted by the discriminability of the relevant stimuli in the local neuronal population,[34] as found previously.[38,57,73,74] In contrast to our findings (Figure 3), overall response levels in IT did not predict the behavioral effects of inactivation. This suggests that there might be differences in the ways that information is read out from V4 and IT. One possibility is that object discrimination is generally performed by IT, while detection of specific stimuli is initiated by lower-level areas like V4, as has been suggested for face processing.[20,75,76]

However, in some experiments, Rajalingham and DiCarlo[34] did observe a significant "choice bias" away from the locally preferred stimulus after IT inactivation (their Figure S2), similar to our findings in V4 (Figure 2A). Thus, it would be interesting to compare V4 and IT readouts more directly, by inactivating IT domains with strong firing rate biases for specific stimuli. It would also be interesting to examine how the IT readout changes with experience in different behavioral tasks.[77]

Our finding that the V4 readout can change with training (Figures 4 and 5) is reminiscent of previous work showing a flexible readout of stimulus information from the middle temporal (MT) area.[9,33,50,78] We have suggested previously that some of this plasticity reflects a preference for simpler readout strategies,[17] as has been found in mouse V1,[37] in artificial neural networks,[43] and in psychophysics.[17] These results highlight the fact that the stimulus representations in neuronal populations are often insufficient to predict the causal impact of that population on behavior.[33,47,50,79,80]

Although this work has focused on changes in readout strategies, our neural recordings indicate that the V4 representation maintained its preference for low spatial frequencies and circular stimuli across all phases of the experiment, and that there was no obvious change in measures of neural discriminability across phases (Figures S2B, S4B, and S5B). Thus, while stimulus representations in V4 often change in more subtle ways with training,[4] the fundamental stimulus preferences of cortical domains appeared fixed throughout our experiments.

Theoretical implications

Many modern artificial neural networks exhibit a "simplicity bias," which is a tendency to find learning shortcuts that exploit the lower-dimensional structure of the input stimuli or the task.[43,81] These strategies emerge early in training and are gradually replaced by more complex ones as training continues.[16,58] Our work shows a neural correlate of this inductive bias, since the readout of V4 during phase 1 relied on a simple (low-dimensional) readout. However, our data from phase 3 suggest that this bias can reemerge even after a more complex one has been learned (Figure 7), consistent with a more general preference for simpler readouts.[82] Given the low-dimensional nature of this kind of readout, one might expect abrupt changes in task performance during learning, although this was not something that our behavioral training methods were designed to detect.

Previous studies related to area V1,[37] area MT,[78,83] area MST,[17] and the parietal cortex[84] have revealed evidence for similar low-dimensional or scalar readouts. In our experiments, the scalar readout from V4 seemed to co-exist with a more distributed readout from PIT (Figure 3). Indeed, although V4 provides a robust representation of the relevant visual stimuli, none of our experiments indicate that the subjects relied exclusively on V4 for task performance. The inactivation results from phase 1 of the experiment suggest that the brain used V4 primarily for detection of low-noise circular patterns (Figure 2), while noisier stimuli were read out in part from PIT (Figure 3). In phase 2, the role of V4 was greatly reduced in favor of some unknown brain region or regions (Figure 4).

At present, we do not know how different cortical areas are recruited for different tasks. One possibility[85] is that biased cortical domains generate population responses that are more amenable to learning by downstream areas,[16,18,58] by virtue of their shared stimulus preferences and noise correlations.[55,86,87] This would be consistent with the idea that task learning involves a competition among cortical domains for access to downstream decision-making areas.[88–91] Such a competition could lead to a distributed readout for stimuli for which representational biases are weak or absent, as suggested in phase 2 of our experiment (Figure 4A). Given that many perceptual learning tasks rely on (non-cardinal) orientation discrimination, a similarly distributed readout may emerge in these cases.

Implications for perception

The V4 readout that emerged in phase 1 of our experiments (Figure 2) appears to ignore much of the information in the neuronal population and actually assigns weights with the wrong sign to many neurons (Figure 6). This highlights an important point that is often raised in the psychology literature[23]: the brain, having evolved to solve real-world problems, favors fast solutions (heuristics) that can be reached with noisy or incomplete data, rather than an energetically demanding[92] optimization process.[24] From this perspective, seemingly suboptimal strategies

can provide important information about the specific inductive biases used by the brain.[93]

Our results suggest that biased cortical domains are a way of implementing these inductive biases. Indeed, neurons in these domains typically prefer environmentally common stimuli, such as faces (IT cortex),[12] cardinal orientations (V1 and V2),[11] and expanding optic flow patterns (area MST).[13,14] Given that subjects readily learn readouts that rely on this kind of bias (Figure 2), it seems likely that similar readouts have formed in previous experiments on perceptual decision-making. That is, psychophysical subjects might perform discrimination tasks for some classes of stimuli by simply detecting the presence of a preferred stimulus on each trial.[36]

RESOURCE AVAILABILITY

Lead contact

Further information and requests should be directed to and will be fulfilled by the lead contact, Pooya Laamerad (pooya.laamerad@pennmedicine.upenn.edu).

Materials availability

This study did not generate new, unique reagents.

Data and code availability

- Raw neural recordings from area V4 have been deposited in the Open Science Framework (OSF) at https://osf.io/bh4wz/ and are publicly available as of the date of publication.
- All code used for behavioral analyses has been deposited in the OSF at https://osf.io/53vz2/ and is publicly available as of the date of publication.
- Any additional data, code, or resources related to this paper are available from the lead contact upon reasonable request.

ACKNOWLEDGMENTS

This work was supported by grants from the International Development Research Centre (108876-001) and the Canadian Institutes of Health Research (PJT178071) to C.C.P. We thank Josh Gold and Shahab Bakhtiari for helpful discussions and Julie Coursol for outstanding technical support.

AUTHOR CONTRIBUTIONS

P.L., D.G., and C.C.P. conceived and designed the study. P.L., M.R.K., and C.C.P. supervised the animal surgery and the implantation of recording chambers. P.L. collected and analyzed data. P.L., M.R.K., and C.C.P. created the visualization. P.L. and C.C.P. drafted the article. P.L., M.R.K., D.G., and C.C.P. revised the article. C.C.P. supervised the study and acquired the funding.

DECLARATION OF INTERESTS

The authors declare no competing interests.

STAR★METHODS

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

- KEY RESOURCES TABLE
- EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS
- METHOD DETAILS
  - Electrophysiological recordings and pharmacological injections
  - Behavioral tasks and visual stimuli
  - Receptive field mapping
  - Passive fixation task
  - Delayed match-to-sample task
  - Neural data
  - Muscimol injection
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Psychometric curve fitting
  - Bias calculation
  - Neural discriminability for high and low spatial frequencies
  - Neural sensitivity (d')
  - Choice probability
  - Noise correlations
  - Readout weights
  - Hierarchical clustering analysis of neural selectivity
  - Statistical comparisons

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.cub.2025.08.027.

Received: May 13, 2025
Revised: July 18, 2025
Accepted: August 13, 2025
Published: September 5, 2025

REFERENCES

1. Arcaro, M.J., and Livingstone, M.S. (2021). On the relationship between maps and domains in inferotemporal cortex. Nat. Rev. Neurosci. *22*, 573–583. https://doi.org/10.1038/s41583-021-00490-4.
2. Li, N., and DiCarlo, J.J. (2010). Unsupervised natural visual experience rapidly reshapes size-invariant object representation in inferior temporal cortex. Neuron *67*, 1062–1075. https://doi.org/10.1016/j.neuron.2010.08.029.
3. Schoups, A., Vogels, R., Qian, N., and Orban, G. (2001). Practising orientation identification improves orientation coding in V1 neurons. Nature *412*, 549–553. https://doi.org/10.1038/35087601.
4. Adab, H.Z., and Vogels, R. (2011). Practicing coarse orientation discrimination improves orientation signals in macaque cortical area v4. Curr. Biol. *21*, 1661–1666. https://doi.org/10.1016/j.cub.2011.08.037.
5. Yang, T., and Maunsell, J.H.R. (2004). The effect of perceptual learning on neuronal responses in monkey visual area V4. J. Neurosci. *24*, 1617–1626. https://doi.org/10.1523/JNEUROSCI.4442-03.2004.
6. Adab, H.Z., Popivanov, I.D., Vanduffel, W., and Vogels, R. (2014). Perceptual learning of simple stimuli modifies stimulus representations in posterior inferior temporal cortex. J. Cogn. Neurosci. *26*, 2187–2200. https://doi.org/10.1162/jocn_a_00641.
7. Sigala, N., and Logothetis, N.K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. Nature *415*, 318–320. https://doi.org/10.1038/415318a.
8. Dosher, B.A., Jeter, P., Liu, J., and Lu, Z.L. (2013). An integrated reweighting theory of perceptual learning. Proc. Natl. Acad. Sci. USA *110*, 13678–13683. https://doi.org/10.1073/pnas.1312552110.
9. Law, C.T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat. Neurosci. *11*, 505–513. https://doi.org/10.1038/nn2070.
10. Dosher, B., and Lu, Z.L. (2017). Visual Perceptual Learning and Models. Annu. Rev. Vis. Sci. *3*, 343–363. https://doi.org/10.1146/annurev-vision-102016-061249.
11. Li, B., Peterson, M.R., and Freeman, R.D. (2003). Oblique effect: a neural basis in the visual cortex. J. Neurophysiol. *90*, 204–217. https://doi.org/10.1152/jn.00954.2002.
12. Tsao, D.Y., Freiwald, W.A., Tootell, R.B.H., and Livingstone, M.S.A. (2006). A cortical region consisting entirely of face-selective cells. Science *311*, 670–674. https://doi.org/10.1126/science.1119983.

13. Graziano, M.S., Andersen, R.A., and Snowden, R.J. (1994). Tuning of MST neurons to spiral motions. J. Neurosci. *14*, 54–67. https://doi.org/10.1523/JNEUROSCI.14-01-00054.1994.

14. Mineault, P.J., Khawaja, F.A., Butts, D.A., and Pack, C.C. (2012). Hierarchical processing of complex motion along the primate dorsal visual pathway. Proc. Natl. Acad. Sci. USA *109*, E972–E980. https://doi.org/10.1073/pnas.1115685109.

15. Bakhtiari, S., Awada, A., and Pack, C.C. (2020). Influence of stimulus complexity on the specificity of visual perceptual learning. J. Vision *20*, 13.

16. Bordelon, B., and Pehlevan, C. (2022). Population codes enable learning from few examples by shaping inductive bias. eLife *11*, e78606. https://doi.org/10.7554/eLife.78606.

17. Laamerad, P., Awada, A., Pack, C.C., and Bakhtiari, S. (2024). Asymmetric stimulus representations bias visual perceptual learning. J. Vis. *24*, 10. https://doi.org/10.1167/jov.24.1.10.

18. Nassar, M.R., Scott, D., and Bhandari, A. (2021). Noise Correlations for Faster and More Robust Learning. J. Neurosci. *41*, 6740–6752. https://doi.org/10.1523/JNEUROSCI.3045-20.2021.

19. Jiang, R., Andolina, I.M., Li, M., and Tang, S. (2021). Clustered functional domains for curves and corners in cortical area V4. eLife *10*, e63798. https://doi.org/10.7554/eLife.63798.

20. Wilkinson, F., James, T.W., Wilson, H.R., Gati, J.S., Menon, R.S., and Goodale, M.A. (2000). An fMRI study of the selective activation of human extrastriate form vision areas by radial and concentric gratings. Curr. Biol. *10*, 1455–1458. https://doi.org/10.1016/s0960-9822(00)00800-9.

21. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806. https://doi.org/10.1016/j.neuron.2013.10.047.

22. Gardner, J.L. (2019). Optimality and heuristics in perceptual neuroscience. Nat. Neurosci. *22*, 514–523. https://doi.org/10.1038/s41593-019-0340-4.

23. Gigerenzer, G., and Gaissmaier, W. (2011). Heuristic decision making. Annu. Rev. Psychol. *62*, 451–482. https://doi.org/10.1146/annurev-psych-120709-145346.

24. Summerfield, C., and Tsetsos, K. (2012). Building Bridges between Perceptual and Economic Decision-Making: Neural and Computational Mechanisms. Front. Neurosci. *6*, 70. https://doi.org/10.3389/fnins.2012.00070.

25. Tversky, A., and Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. Science *185*, 1124–1131. https://doi.org/10.1126/science.185.4157.1124.

26. Ghose, G.M., and Ts'o, D.Y. (1997). Form processing modules in primate area V4. J. Neurophysiol. *77*, 2191–2196. https://doi.org/10.1152/jn.1997.77.4.2191.

27. Hu, J.M., Song, X.M., Wang, Q., and Roe, A.W. (2020). Curvature domains in v4 of macaque monkey. eLife *9*, e57261. https://doi.org/10.7554/eLife.57261.

28. Li, P., Zhu, S., Chen, M., Han, C., Xu, H., Hu, J., Fang, Y., and Lu, H.D. (2013). A motion direction preference map in monkey V4. Neuron *78*, 376–388. https://doi.org/10.1016/j.neuron.2013.02.024.

29. Namima, T., Kempkes, E., Zamarashkina, P., Owen, N., and Pasupathy, A. (2025). High-Density Recording Reveals Sparse Clusters (But Not Columns) for Shape and Texture Encoding in Macaque V4. J. Neurosci. *45*, e1893232024. https://doi.org/10.1523/JNEUROSCI.1893-23.2024.

30. Tang, R., Song, Q., Li, Y., Zhang, R., Cai, X., and Lu, H.D. (2020). Curvature-processing domains in primate V4. eLife *9*, e57502. https://doi.org/10.7554/eLife.57502.

31. Tanigawa, H., Lu, H.D., and Roe, A.W. (2010). Functional organization for color and orientation in macaque V4. Nat. Neurosci. *13*, 1542–1548. https://doi.org/10.1038/nn.2676.

32. Kriegeskorte, N., Mur, M., Ruff, D.A., Kiani, R., Bodurka, J., Esteky, H., Tanaka, K., and Bandettini, P.A. (2008). Matching categorical object representations in inferior temporal cortex of man and monkey. Neuron *60*, 1126–1141. https://doi.org/10.1016/j.neuron.2008.10.043.

33. Chowdhury, S.A., and DeAngelis, G.C. (2008). Fine discrimination training alters the causal contribution of macaque area MT to depth perception. Neuron *60*, 367–377. https://doi.org/10.1016/j.neuron.2008.08.023.

34. Rajalingham, R., and DiCarlo, J.J. (2019). Reversible Inactivation of Different Millimeter-Scale Regions of Primate IT Results in Different Patterns of Core Object Recognition Deficits. Neuron *102*, 493–505.e5. https://doi.org/10.1016/j.neuron.2019.02.001.

35. Edeline, J.-M., Hars, B., Hennevin, E., and Cotillon, N. (2002). Muscimol diffusion after intracerebral microinjections: a reevaluation based on electrophysiological and autoradiographic quantifications. Neurobiol. Learn. Mem. *78*, 100–124. https://doi.org/10.1006/nlme.2001.4035.

36. Herzog, M.H., and Fahle, M. (1998). Modeling perceptual learning: difficulties and how they can be overcome. Biol. Cybern. *78*, 107–117. https://doi.org/10.1007/s004220050418.

37. Jin, M., Beck, J.M., and Glickfeld, L.L. (2019). Neuronal Adaptation Reveals a Suboptimal Decoding of Orientation Tuned Populations in the Mouse Visual Cortex. J. Neurosci. *39*, 3867–3881. https://doi.org/10.1523/JNEUROSCI.3172-18.2019.

38. Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. Science *293*, 2425–2430. https://doi.org/10.1126/science.1063736.

39. Petrov, A.A., Dosher, B.A., and Lu, Z.L. (2005). The dynamics of perceptual learning: an incremental reweighting model. Psychol. Rev. *112*, 715–743. https://doi.org/10.1037/0033-295X.112.4.715.

40. Putzeys, T., Bethge, M., Wichmann, F., Wagemans, J., and Goris, R. (2012). A new perceptual bias reveals suboptimal population decoding of sensory responses. PLOS Comput. Biol. *8*, e1002453. https://doi.org/10.1371/journal.pcbi.1002453.

41. Gold, J.I., and Ding, L. (2013). How mechanisms of perceptual decision-making affect the psychometric function. Prog. Neurobiol. *103*, 98–114. https://doi.org/10.1016/j.pneurobio.2012.05.008.

42. Sclar, G., Maunsell, J.H., and Lennie, P. (1990). Coding of image contrast in central visual pathways of the macaque monkey. Vision Res. *30*, 1–10. https://doi.org/10.1016/0042-6989(90)90123-3.

43. Shah, H., Tamuly, K., Raghunathan, A., Jain, P., and Netrapalli, P. (2020). The pitfalls of simplicity bias in neural networks. Preprint at ArXiv, 2006.07710.

44. Wilkinson, F., Wilson, H.R., and Habak, C. (1998). Detection and recognition of radial frequency patterns. Vision Res. *38*, 3555–3568. https://doi.org/10.1016/s0042-6989(98)00039-x.

45. Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. Science *308*, 529–534. https://doi.org/10.1126/science.1109676.

46. Ponce, C.R., Lomber, S.G., and Livingstone, M.S. (2017). Posterior Inferotemporal Cortex Cells Use Multiple Input Pathways for Shape Encoding. J. Neurosci. *37*, 5019–5034. https://doi.org/10.1523/JNEUROSCI.2674-16.2017.

47. Pinto, L., Rajan, K., DePasquale, B., Thiberge, S.Y., Tank, D.W., and Brody, C.D. (2019). Task-dependent changes in the large-scale dynamics and necessity of cortical regions. Neuron *104*, 810–824.e9. https://doi.org/10.1016/j.neuron.2019.08.025.

48. Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. Science *348*, 1352–1355. https://doi.org/10.1126/science.aab0551.

49. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273. https://doi.org/10.1038/s41586-019-1787-x.

50. Liu, L.D., and Pack, C.C. (2017). The Contribution of Area MT to Visual Motion Perception Depends on Training. Neuron *95*, 436–446.e3. https://doi.org/10.1016/j.neuron.2017.06.024.

51. Jing, R., Yang, C., Huang, X., and Li, W. (2021). Perceptual learning as a result of concerted changes in prefrontal and visual cortex. Curr. Biol. *31*, 4521–4533.e3. https://doi.org/10.1016/j.cub.2021.08.007.

52. Mayhew, S.D., Li, S., and Kourtzi, Z. (2012). Learning acts on distinct processes for visual form perception in the human brain. J. Neurosci. *32*, 775–786. https://doi.org/10.1523/JNEUROSCI.2033-11.2012.

53. Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., and Movshon, J.A. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100. https://doi.org/10.1017/s095252380000715x.

54. Crapse, T.B., and Basso, M.A. (2015). Insights into decision making using choice probability. J. Neurophysiol. *114*, 3039–3049. https://doi.org/10.1152/jn.00335.2015.

55. Haefner, R.M., Gerwinn, S., Macke, J.H., and Bethge, M. (2013). Inferring decoding strategies from choice probabilities in the presence of correlated variability. Nat. Neurosci. *16*, 235–242. https://doi.org/10.1038/nn.3309.

56. Nienborg, H., Cohen, M.R., and Cumming, B.G. (2012). Decision-related activity in sensory neurons: correlations among neurons and with behavior. Annu. Rev. Neurosci. *35*, 463–483. https://doi.org/10.1146/annurev-neuro-062111-150403.

57. Law, C.T., and Gold, J.I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. Nat. Neurosci. *12*, 655–663. https://doi.org/10.1038/nn.2304.

58. Benjamin, A.S., Zhang, L.Q., Qiu, C., Stocker, A.A., and Kording, K.P. (2022). Efficient neural codes naturally emerge through gradient descent learning. Nat. Commun. *13*, 7972. https://doi.org/10.1038/s41467-022-35659-7.

59. Heywood, C.A., Gadotti, A., and Cowey, A. (1992). Cortical area V4 and its role in the perception of color. J. Neurosci. *12*, 4056–4065. https://doi.org/10.1523/JNEUROSCI.12-10-04056.1992.

60. Schiller, P.H. (1993). The effects of V4 and middle temporal (MT) area lesions on visual performance in the rhesus monkey. Vis. Neurosci. *10*, 717–746. https://doi.org/10.1017/s0952523800005423.

61. Merigan, W.H., and Pham, H.A. (1998). V4 lesions in macaques affect both single- and multiple-viewpoint shape discriminations. Vis. Neurosci. *15*, 359–367. https://doi.org/10.1017/s0952523898152112.

62. De Weerd, P., Desimone, R., and Ungerleider, L.G. (1996). Cue-dependent deficits in grating orientation discrimination after V4 lesions in macaques. Vis. Neurosci. *13*, 529–538. https://doi.org/10.1017/s0952523800008208.

63. Heywood, C.A., and Cowey, A. (1987). On the role of cortical area V4 in the discrimination of hue and pattern in macaque monkeys. J. Neurosci. *7*, 2601–2617. https://doi.org/10.1523/JNEUROSCI.07-09-02601.1987.

64. Walsh, V., and Butler, S.R. (1996). The effects of visual cortex lesions on the perception of rotated shapes. Behav. Brain Res. *76*, 127–142. https://doi.org/10.1016/0166-4328(95)00191-3.

65. Walsh, V., Le Mare, C., Blaimire, A., and Cowey, A. (2000). Normal discrimination performance accompanied by priming deficits in monkeys with V4 or TEO lesions. NeuroReport *11*, 1459–1462.

66. Roe, A.W., Chelazzi, L., Connor, C.E., Conway, B.R., Fujita, I., Gallant, J. L., Lu, H., and Vanduffel, W. (2012). Toward a unified theory of visual area V4. Neuron *74*, 12–29. https://doi.org/10.1016/j.neuron.2012.03.011.

67. Girard, P., Lomber, S.G., and Bullier, J. (2002). Shape discrimination deficits during reversible deactivation of area V4 in the macaque monkey. Cereb. Cortex *12*, 1146–1156. https://doi.org/10.1093/cercor/12.11.1146.

68. Shiozaki, H.M., Tanabe, S., Doi, T., and Fujita, I. (2012). Neural activity in cortical area V4 underlies fine disparity discrimination. J. Neurosci. *32*, 3830–3841. https://doi.org/10.1523/JNEUROSCI.5083-11.2012.

69. Salzman, C.D., Murasugi, C.M., Britten, K.H., and Newsome, W.T. (1992). Microstimulation in visual area MT: effects on direction discrimination performance. J. Neurosci. *12*, 2331–2355. https://doi.org/10.1523/JNEUROSCI.12-06-02331.1992.

70. Murasugi, C.M., Salzman, C.D., and Newsome, W.T. (1993). Microstimulation in visual area MT: effects of varying pulse amplitude and frequency. J. Neurosci. *13*, 1719–1729. https://doi.org/10.1523/JNEUROSCI.13-04-01719.1993.

71. 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. https://doi.org/10.1017/s0952523800010269.

72. Pinto, L., Tank, D.W., and Brody, C.D. (2022). Multiple timescales of sensory-evidence accumulation across the dorsal cortex. eLife *11*, e70263. https://doi.org/10.7554/eLife.70263.

73. Zatka-Haas, P., Steinmetz, N.A., Carandini, M., and Harris, K.D. (2021). Sensory coding and the causal impact of mouse cortex in a visual decision. eLife *10*, e63163. https://doi.org/10.7554/eLife.63163.

74. Williams, M.A., Dang, S., and Kanwisher, N.G. (2007). Only some spatial patterns of fMRI response are read out in task performance. Nat. Neurosci. *10*, 685–686. https://doi.org/10.1038/nn1900.

75. Afraz, A., Boyden, E.S., and DiCarlo, J.J. (2015). Optogenetic and pharmacological suppression of spatial clusters of face neurons reveal their causal role in face gender discrimination. Proc. Natl. Acad. Sci. USA *112*, 6730–6735. https://doi.org/10.1073/pnas.1423328112.

76. Tsao, D.Y., and Livingstone, M.S. (2008). Mechanisms of face perception. Annu. Rev. Neurosci. *31*, 411–437. https://doi.org/10.1146/annurev.neuro.30.051606.094238.

77. Csorba, B.A., Krause, M.R., Zanos, T.P., and Pack, C.C. (2022). Long-range cortical synchronization supports abrupt visual learning. Curr. Biol. *32*, 2467–2479.e4. https://doi.org/10.1016/j.cub.2022.04.029.

78. Laamerad, P., Liu, L.D., and Pack, C.C. (2024). Decision-related activity and movement selection in primate visual cortex. Sci. Adv. *10*, eadk7214. https://doi.org/10.1126/sciadv.adk7214.

79. 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. https://doi.org/10.1038/nature18617.

80. Zhou, Y., and Freedman, D.J. (2019). Posterior parietal cortex plays a causal role in perceptual and categorical decisions. Science *365*, 180–185. https://doi.org/10.1126/science.aaw8347.

81. Huh, M., Mobahi, H., Zhang, R., Cheung, B., Agrawal, P., and Isola, P. (2021). The low-rank simplicity bias in deep networks. Preprint at arXiv. https://arxiv.org/abs/2103.10427.

82. Lampinen, A.K., Chan, S.C., and Hermann, K. (2024). Learned feature representations are biased by complexity, learning order, position, and more. Preprint at arXiv. https://arxiv.org/abs/2405.05847.

83. Tsui, J.M.G., Hunter, J.N., Born, R.T., and Pack, C.C. (2010). The role of V1 surround suppression in MT motion integration. J. Neurophysiol. *103*, 3123–3138. https://doi.org/10.1152/jn.00654.2009.

84. Fitzgerald, J.K., Freedman, D.J., Fanini, A., Bennur, S., Gold, J.I., and Assad, J.A. (2013). Biased associative representations in parietal cortex. Neuron *77*, 180–191. https://doi.org/10.1016/j.neuron.2012.11.014.

85. Nienborg, H., and Cumming, B.G. (2014). Decision-related activity in sensory neurons may depend on the columnar architecture of cerebral cortex. J. Neurosci. *34*, 3579–3585. https://doi.org/10.1523/JNEUROSCI.2340-13.2014.

86. Nienborg, H., and Cumming, B.G. (2006). Macaque V2 neurons, but not V1 neurons, show choice-related activity. J. Neurosci. *26*, 9567–9578. https://doi.org/10.1523/JNEUROSCI.2256-06.2006.

87. Zohary, E., Shadlen, M.N., and Newsome, W.T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance2006. Nature *370*, 140–143. https://doi.org/10.1038/370140a0.

88. Bock, A.S., and Fine, I. (2014). Anatomical and functional plasticity in early blind individuals and the mixture of experts architecture. Front. Hum. Neurosci. *8*, 971. https://doi.org/10.3389/fnhum.2014.00971.

89. Jacobs, R.A., Jordan, M.I., Nowlan, S.J., and Hinton, G.E. (1991). Adaptive mixtures of local experts. Neural Comput. *3*, 79–87. https://doi.org/10.1162/neco.1991.3.1.79.

90. O'Doherty, J.P., Lee, S.W., Tadayonnejad, R., Cockburn, J., Iigaya, K., and Charpentier, C.J. (2021). Why and how the brain weights contributions from a mixture of experts. Neurosci. Biobehav. Rev. *123*, 14–23. https://doi.org/10.1016/j.neubiorev.2020.10.022.

91. Sotiropoulos, G., Seitz, A.R., and Seriès, P. (2018). Performance-monitoring integrated reweighting model of perceptual learning. Vision Res. *152*, 17–39. https://doi.org/10.1016/j.visres.2018.01.010.

92. Pache, A., and van Rossum, M.C.W. (2023). Energetically efficient learning in neuronal networks. Curr. Opin. Neurobiol. *83*, 102779. https://doi.org/10.1016/j.conb.2023.102779.

93. Sinz, F.H., Pitkow, X., Reimer, J., Bethge, M., and Tolias, A.S. (2019). Engineering a less artificial intelligence. Neuron *103*, 967–979. https://doi.org/10.1016/j.neuron.2019.08.034.

94. Desimone, R., and Schein, S.J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form. J. Neurophysiol. *57*, 835–868. https://doi.org/10.1152/jn.1987.57.3.835.

95. Hill, D.N., Mehta, S.B., and Kleinfeld, D. (2011). Quality metrics to accompany spike sorting of extracellular signals. J. Neurosci. *31*, 8699–8705. https://doi.org/10.1523/JNEUROSCI.0971-11.2011.

96. Brainard, D.H., and Vision, S. (1997). The psychophysics toolbox. Spat. Vis *10*, 433–436. https://doi.org/10.1163/156856897X00357.

97. Allen, T.A., Narayanan, N.S., Kholodar-Smith, D.B., Zhao, Y., Laubach, M., and Brown, T.H. (2008). Imaging the spread of reversible brain inactivations using fluorescent muscimol. J. Neurosci. Methods *171*, 30–38. https://doi.org/10.1016/j.jneumeth.2008.01.033.

98. Wichmann, F.A., and Hill, N.J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. Percept. Psychophys. *63*, 1293–1313. https://doi.org/10.3758/bf03194544.

99. Prins, N., and Kingdom, F.A.A. (2018). Applying the model-comparison approach to test specific research hypotheses in psychophysical research using the Palamedes toolbox. Front. Psychol. *9*, 1250. https://doi.org/10.3389/fpsyg.2018.01250.

100. Wickens, T.D. (2001). Elementary Signal Detection Theory (Oxford University Press). https://doi.org/10.1093/acprof:oso/9780195092509.001.0001.

101. Schwartz, A.B., Kettner, R.E., and Georgopoulos, A.P. (1988). Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement. J. Neurosci. *8*, 2913–2927. https://doi.org/10.1523/JNEUROSCI.08-08-02913.1988.

102. Lubba, C.H., Sethi, S.S., Knaute, P., Schultz, S.R., Fulcher, B.D., and Jones, N.S. (2019). catch22: CAnonical Time-series CHaracteristics: Selected through highly comparative time-series analysis. Data Min. Knowl. Discov. *33*, 1821–1852. https://doi.org/10.1007/s10618-019-00647-x.

103. Caliński, T., and Harabasz, J. (1974). A dendrite method for cluster analysis. Commun. Stat. Theor. Methods *3*, 1–27. https://doi.org/10.1080/03610927408827101.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Chemicals, peptides, and recombinant proteins		
Muscimol	Sigma	M1523 Sigma; CAS # 2763-96-4
Experimental models: Organisms/strains		
Rhesus macaque (Mucaca mulatta)	N/A	N/A
Software and algorithms		
MATLAB	MathWorks	https://www.mathworks.com/products/matlab.html
NIMH MonkeyLogic	NIH	https://monkeylogic.nimh.nih.gov/
Intan RHX Data Acquisition Software	Intan Technologies system	https://intantech.com/downloads.html?tabSelect=Software&yPos=100
UltraMegaSort2000 – Spike sorting	MathWorks	https://github.com/danamics/UMS2K
Other		
Plexon V probes	Plexon	PLX-VP-16-15ED-100-SE-140-65(460)-CT-300-F5/6
RHD Recording System	Intan Technologies system	https://intantech.com/RHD_system.html

EXPERIMENTAL MODEL AND STUDY PARTICIPANT DETAILS

Two adult female rhesus monkeys (5 – 7 kg) participated in this study. Initially, under general anesthesia, MRI-compatible titanium head posts were attached to each monkey's skull to stabilize their heads during training and experimental sessions. Target brain areas were accessed through sterile plastic recording chambers (Crist Instruments) that were permanently implanted on the skulls. For both animals, these chambers provided access to the ventral region of area V4 in the right hemisphere, via a dorsal-posterior approach (Figure S1A). All experimental procedures followed the guidelines of the Canadian Council on Animal Care and were approved by the Institutional Animal Care Committee at the Montreal Neurological Institute.

METHOD DETAILS

Electrophysiological recordings and pharmacological injections

Area V4 was identified using anatomical MRI scans, the relationship between receptive field size and eccentricity,[94] and transitions from white to gray matter through different brain regions (Figure S1A). Each recording session began with the installation of a grid system within the recording chamber, to ensure precise electrode placement. The grid was aligned to predefined coordinates that were identical across all phases of the experiment in each animal.

This was followed by the penetration of the dura using a 23-gauge stainless steel guide tube. Relative to the cortical surface, the electrodes were advanced to an average depth of 22.1 ± 0.7 mm for Monkey 1 and 20.7 ± 0.8 mm for Monkey 2 (Figure S1A), using an Oil Hydraulic Micromanipulator (Narishige International USA, INC.).

We recorded single-unit activity using linear microelectrode arrays (V-Probe; Plexon). For the passive fixation tasks, 32-channel V-probes were employed, while 16-channel electrodes were used during sessions that involved muscimol injections. Neuronal signals were recorded using an Intan Technologies system and filtered between 0.5 and 7 kHz. Initial spike detection involved identifying crossings that exceeded a threshold of ±3 standard deviations, robustly estimated for each channel. Short segments surrounding each threshold crossing were then extracted and clustered using UltraMegaSort 2000, a *k*-means-based clustering algorithm.[95]

Behavioral tasks and visual stimuli

Animals were seated in a standard primate chair (Crist Instruments). Visual stimuli were back-projected onto a semi-transparent screen using an LED video projector (VPixx Technologies, PROPixx) with a refresh rate of 120 Hz. The screen spanned an area of 80° × 50° of visual angle at a viewing distance of 81 cm. A neutral gray (54 cd/m^2) served as the background color for all tasks. Eye movements for both animals were tracked using an infrared eye tracking system (EyeLink1000, SR Research) with a sampling rate of 1,000 Hz.

Receptive field mapping

Before the behavioral sessions began, we conducted initial measurements to assess stimulus selectivity (Figure 1) and receptive field properties for each animal. During these sessions, animals maintained fixation while sparse noise stimuli were presented at different

spatial locations. The resulting data were fit with a 2D Gaussian function, to recover the receptive field centers and sizes (Figure S1B). These measurements were performed at the same grid position used for recordings and injections.

Passive fixation task

For the pre-training assessment of stimulus selectivity, we used 152 diverse images, which we divide into three broad categories. The first category included angles, curves, curve-line combinations, lines, and star shapes, with each sub-category featuring 8 variations across different orientations at 45-degree increments. The second category included Glass patterns, circular gratings, radial gratings, polygonal gratings, linear gratings, and drifting gratings, each with 8 variations in shape, spatial frequency, or motion direction. Drifting grating stimuli were only presented to Monkey 1, as Monkey 2 consistently broke fixation when presented with moving stimuli. The third category included realistic images of animals, body parts, human faces, insects, monkey faces, natural scenes, tools, and vegetables, each with 8 different exemplars. Our selection of stimulus categories was influenced by those used in previous studies.[46] A sample image from each sub-category is shown in Figure S1C. All images were generated in MATLAB using extensions from the Psychophysics Toolbox.[96]

Animals were required to fixate a 0.5° green square at the center of the screen and maintain their gaze within ± 1.0° of the fixation point. After a 500 ms baseline period, the first stimulus—randomly selected from the set of 152 images—was displayed on the screen for 200 ms. This was followed by a 200 ms delay during which only the fixation point remained visible, before the next image appeared. This on-off cycle was repeated with up to 10 different images per trial. One reward was dispensed after a random number of images (ranging from 3 to 7); if the animal maintained fixation until the end of the cycle, a double reward was dispensed. Stimuli were centered on the measured RFs of the neurons under study (Monkey 1: 2.1 ± 0.5° radius at 1.5 ± 0.4° eccentricity; Monkey 2: 1.5 ± 0.7° radius at 2.4 ± 0.9° eccentricity; Figure S1B). Stimulus sizes were 3.5° × 3.5° for Monkey 1 and 4° × 4° for Monkey 2.

Delayed match-to-sample task

Monkeys were trained on a delayed match-to-sample (DMS) task that required them to identify which of two choice stimuli (a circular grating or a radial grating) matched a previously presented sample image. Each trial began with an initial fixation period of 500 ms, followed by the presentation of a sample stimulus for 200 ms. This sample stimulus, either a circular or a radial grating, had Gaussian noise added at one of eight levels (from 100% SNR, a noiseless image, to 0% SNR, pure noise) using Psychophysics Toolbox and MATLAB's image processing toolbox. The size and position of the images were identical to those used during the passive fixation task for each animal (see Receptive Field Mapping and Visual Stimuli section). Luminance and contrast were balanced for all images across all phases of the experiment. For both low and high spatial frequency gratings, the average contrast of the images, measured using the RMS contrast method, was 0.086 ± 0.008. There was no significant difference in image contrast compared to the images used during the passive fixation task ($p = 0.92$, t-test).

After the sample was removed, a randomized delay period of 250 - 500 ms followed. Subsequently, noiseless (100% SNR) circular and radial grating stimuli appeared on either side of the fixation point. These response cues were positioned ±7° from the fixation point along the horizontal meridian. To avoid the development of a fixed sensorimotor mapping,[78] the left/right position of each grating was randomized from trial to trial. The monkey had to make a saccade towards the stimulus that matched the sample and maintain fixation on it for 800 ms to receive a fluid reward. A new trial began after a 500 ms interval during which no stimulus was presented. If the monkey chose incorrectly, we imposed a 1.5-second time-out before the next trial.

To familiarize the animals with the concept of the DMS task, we first trained each animal using a white circular image as the sample stimulus. During this phase, the choice stimuli consisted of black and white circles, and the animals were required to select the white image to receive a reward. After approximately two weeks, we switched the sample stimulus to a black image and repeated the same procedure. The animals typically learned the black sample image within a few days.

Next, we began alternating between white and black sample images on a day-by-day basis. Following this, we interleaved trials in blocks of 100 (e.g., 100 trials with white samples followed by 100 trials with black samples). Gradually, we reduced the block size until the animals were able to match the color of the choice stimulus to the sample stimulus in a randomized, trial-by-trial manner. By the final five sessions of this phase, their performance consistently exceeded 90%. This familiarization phase took approximately two months to complete.

After the initial familiarization phase, the monkeys were introduced to low spatial frequency circular and radial gratings. We did not observe any systematic bias toward circular or radial gratings when the two stimulus types were first introduced after the initial familiarization phase. At the beginning, animals occasionally showed mild and inconsistent preferences, sometimes favoring radial gratings and sometimes circular, but these tendencies were random and not sustained across sessions. In fact, the more common behavioral bias observed during early training was spatial (i.e., toward one side of the screen), rather than stimulus-type specific.

To prevent the reinforcement of any such biases, we actively monitored behavior throughout training and dynamically adjusted stimulus contingencies, for instance by increasing the frequency of correct stimuli on the side or of the type that the animal was avoiding. This approach ensured balanced exposure and discouraged the development of fixed biases.

To prevent loss of motivation, we avoided presenting highly noisy stimuli early in training. Animals often disengaged when the task was too difficult, responding randomly to obtain fluid rewards 50% of the time. After two sessions with the 12% SNR stimulus, additional noise levels were introduced incrementally at the beginning of each session. All noise levels were fully integrated by session 16 for Monkey 1 and session 18 for Monkey 2 in Phase 1.

Monkey 1 was trained five days per week and completed 70 sessions (~14 weeks), achieving an average threshold of 12.4% ± 3.1% during the final five sessions of Phase 1. Monkey 2 was trained daily and completed 39 sessions (~5.5 weeks), reaching an average threshold of 9.3% ± 1.2% during its final five sessions.

In Phase 2, the animals were trained on the same task structure but with high spatial frequency gratings (2.0 cpd). Despite familiarity with the task structure, both monkeys exhibited poor initial performance with the new stimuli. Even for noiseless trials, their accuracy dropped below 65%. This increased lapse rate is due to an aversion to changes in routine, rather than to a lack of task understanding.

As in Phase 1, training began with noiseless stimuli, and noisier stimuli were introduced gradually. Monkey 1 met the performance criterion to add the 12% SNR stimulus after 8 sessions; Monkey 2 required 16 sessions. In both cases, the full range of noise levels was introduced over the next two sessions. Monkey 1 completed 54 sessions (~11 weeks) in Phase 2, with an average threshold of 12.4% ± 3.1% in the final five sessions. Monkey 2 completed 48 sessions (~7 weeks), achieving a threshold of 9.3% ± 1.2%.

In Phase 3 (Monkey 1 only), all noise levels were introduced from the first session. This phase lasted two weeks (10 sessions), during which the performance threshold remained stable, with an average of 10.8% ± 2.6%.

In all phases, muscimol injections began only after the completion of training and stabilization of performance. All behavioral data presented in Figures 2, 3, 4, 5, and 7 were collected after training was completed in each respective phase.

Note that we did not present the 0% SNR stimulus to Monkey 2, as she consistently refused to do trials containing this stimulus. Also in Phase 2, for low spatial frequency stimuli, Monkey 1 consistently closed her eyes whenever the 0% SNR stimulus appeared on the screen. As a result, we did not obtain a sufficient number of trials at this noise level to include in the analysis.

Neural data

In Phase 1, we recorded from 93 neurons in Monkey 1 (of which 73 preferred circular gratings) and 53 neurons in Monkey 2 (32 preferred circular gratings). In Phase 2, we recorded from 80 neurons in Monkey 1 (56 preferred circular gratings) and 35 neurons in Monkey 2 (22 preferred circular gratings). In Phase 3, we recorded from 50 neurons in Monkey 1, of which 31 preferred circular gratings. Stimulus preference was determined by comparing the mean firing rates evoked by circular versus radial gratings, using only the noiseless (100% SNR) stimuli across trials.

Muscimol injection

At the start of each behavioral session, we returned to the grid positions used during the initial mapping of stimulus selectivity and manually estimated multi-channel receptive fields by positioning moving grating stimuli within the approximate visual field of the neurons. This process ensured that the inactivated neurons had receptive fields consistent with those identified during the initial mapping sessions. We then performed muscimol inactivation.

The linear array included a glass capillary with an inner diameter of 40 μm, positioned between contacts 5 and 6 of the array (with contact 1 being the most dorsal). The other end of the capillary was connected to a Hamilton syringe through plastic tubing. Muscimol was injected using a mini-pump, typically at a volume of 2 μL and a rate of 0.05 μL/min, with a concentration of 10 mg/mL.[33] We confirmed cessation of neural activity before starting behavioral experiments, 45 minutes post-injection. Muscimol injections and initial behavioral tests were conducted on the afternoon of the first day, with further testing at 18 hours and 2 days later. As shown in Figure S3, behavioral impairments were localized to the RF area of the neurons near the injection site. Given this specificity and the fact that the spread of muscimol is typically less than 2 mm,[97] it is unlikely that the drug spread into nearby retinotopic regions of cortex. Muscimol experiments were performed no more than once per week.

QUANTIFICATION AND STATISTICAL ANALYSIS

Psychometric curve fitting

To quantify perceptual performance, we fit psychometric functions to behavioral data using a four-parameter logistic model implemented with the Palamedes Toolbox[96] in MATLAB. The logistic function is mathematically well-suited for binary decision and perceptual discrimination tasks as it captures the characteristic S-shaped progression of performance as a function of stimulus strength, with bounds between 0% and 100%.[98] Model parameters included the threshold (α), slope (β), guess rate (γ), and lapse rate (λ), which were optimized via maximum likelihood estimation. The best-fitting function was selected based on the highest log-likelihood and lowest deviance.

The logistic function used is defined as:

$$\psi(x) = \gamma + (1 - \gamma - \lambda) \times \frac{1}{1+e^{-\beta(x-\alpha)}}$$

where $\psi(x)$ is predicted proportion of correct responses at signal level x (SNR %).

While the logistic function provided good fits for visualization, particularly for illustrating performance across all SNR levels in Figures 2A, 3F, 4A, 5A and 7B, we used a Weibull function to derive a more accurate and reliable estimate of behavioral thresholds for each type of grating. The Weibull function is well-suited for estimating perceptual thresholds in discrimination paradigms,

especially when applied to single stimulus types. For each grating condition (circular and radial), we fit a Weibull function to the behavioral data, defining threshold as the SNR level at which performance reached 82% correct, a standard in 2AFC tasks corresponding to 63.2% of the rise from chance to perfect accuracy.[98]

Behavioral performance as a function of signal level (SNR %) was modeled using the cumulative distribution function (CDF) of the Weibull function:

$$\psi(x) = \gamma + (1 - \gamma - \lambda) \times \left[1 - \exp\left(-\left(\frac{x}{\alpha}\right)^{\beta}\right)\right]$$

where α determines the threshold, β corresponds to the slope of the function, γ is the guess rate and λ is the lapse rate. This function was fit using the maximum likelihood algorithm from MATLAB Palamedes toolbox for analyzing psychophysical data[99] (Figures 2B, 4B, 5B and 7D).

Bias calculation

We used the equation below to calculate behavioral biases[100]:

$$Bias = \left(Norminv^2(Circular\ grating\ hit\ rate) - Norminv^2(Radial\ grating\ hit\ rate)\right) / 2$$

where Norminv is the inverse cumulative distribution function of the standard normal distribution. A log-transform was used to remove a nonlinear effect of bias. Negative bias values indicate a bias toward radial stimuli, and positive values indicate a bias toward circular stimuli (Figures 2C, 4C, and 5C).

Neural discriminability for high and low spatial frequencies

To quantify neural discriminability between stimulus conditions, we trained a support vector machine (SVM) classifier on neural responses recorded from each monkey during the Preliminary Phase. Neural activity was represented as a feature matrix, where each row corresponded to a trial and each column represented the firing rates of recorded neurons. Pairs of circular and radial gratings with low and high spatial frequency images were analyzed separately, with trials labeled according to their corresponding stimulus condition. We applied a linear SVM using MATLAB's (fitcsvm) function and evaluated classification performance using 5-fold cross-validation to ensure robustness. Discriminability was assessed as the mean classification accuracy across folds, reflecting how well neural activity patterns differentiated between the two stimulus conditions. This procedure was repeated 20 times to ensure reliability. To statistically compare discriminability across image pairs, we performed a t-test on classification accuracies, testing whether SVM performance differed significantly between the two spatial frequency conditions.

Neural sensitivity (d')

The d' measure of neural sensitivity quantifies the strength of selectivity across all noise levels, as follows:

$$d' = \frac{R_{circular} - R_{radial}}{\sqrt{\frac{\sigma^2_{circular} + \sigma^2_{radial}}{2}}}$$

where $R_{circular}$, R_{radial}, $\sigma_{circular}$ and σ_{radial} represent the mean responses and standard deviations in response to circular and radial grating across all noise levels (Figure 6D).

Choice probability

Choice probability (CP) was used to quantify the relationship between behavioral choice and response variability.[53] For an identical stimulus, the responses can be grouped into two distributions based on whether the monkeys made the choice that corresponds to the neuron's preferred type of grating or the non-preferred grating. As long as the monkeys made at least 6 choices for each type of grating with a minimum of 3 trials involving a wrong choice, ROC values were calculated from these response distributions. The area underneath the ROC curve provides the CP value (Figure 6C).

To analyze the dynamics of choice-related activity, after calculating the CP for each neuron across stimulus conditions, spanning from 300 ms before to 500 ms after stimulus onset, using a window size of 75 ms and a step size of 5 ms. The emergence of statistically significant CP was detected with the Cumulative Sum (CUSUM) algorithm.[101]

Noise correlations

Noise correlations were calculated as the Pearson correlation coefficient representing the trial-by-trial covariation of responses from pairs of neurons recorded simultaneously within a 200-ms window, spanning from 50 ms to 250 ms after stimulus onset.[87] Each neuron's responses were z-scored by subtracting the average response and dividing by the standard deviation across multiple stimulus presentations. This procedure eliminated the impact of stimulus strength and direction on the average response, allowing noise correlation to solely capture correlated trial-to-trial fluctuations around the mean response. To avoid correlations influenced by extreme values, only trials with responses within ±3 standard deviations of the mean were considered.[87] Additionally, varying the window size between 100 ms and 400 ms did not affect the pattern of noise correlation results compared to the 200-ms window size.

Readout weights

We estimated the readout weight (β) for each neuron, using the procedure outlined by Haefner and colleagues.[55] According to this measure, readout weights depend on CP and the noise correlation matrix *C*:

$$CP \approx \frac{1}{2} + \frac{\sqrt{2}}{\pi} \frac{(C\beta)_k}{\sqrt{C_{kk}\beta^T C\beta}} \quad (1)$$

Since in the denominator ($\sqrt{\beta^T C\beta}$) is constant across all neurons in each session, we neglected this term and simplified the equation to:

$$\beta = C^{-1}\left(\frac{\pi \times \sqrt{C_{kk}} \times \left(CP - \frac{1}{2}\right)}{\sqrt{2}}\right)$$

We computed *C* in each session with more than three neurons, and used this information, along with each neuron's CP, to calculate the β values.

Hierarchical clustering analysis of neural selectivity

We employed hierarchical clustering to categorize image groups based on their neuronal response patterns, applying complete linkage with a threshold value of 0.2.[102] The primary objective of this method was to identify whether circular and radial grating stimuli emerged as the most distinct clusters among all image categories. Initially, we analyzed the mean normalized neuronal activity for each image category, comprising eight images per category, to assess the selectivity of the V4 neuronal population. Subsequently, we calculated the Calinski-Harabasz criterion[103] to determine the optimal and most reliable number of clusters for each monkey's dataset. After establishing the ideal cluster count, we evaluated linkage distances, which measure the dissimilarity between clusters, to quantify the separation between clusters.[102] Linkage distance provides a numerical representation of how far apart clusters are in the hierarchical tree, with larger distances indicating more distinct differences in neuronal response patterns between the clusters. For both monkeys, our analysis demonstrated that circular and radial grating categories consistently formed the most distinct clusters, with the greatest separation distance observed between them. These results highlight that these two image types elicited significantly dissimilar neuronal responses within the V4 population.

Permutation testing was then performed to validate whether the observed clustering and the large distance between the circular and radial grating categories were statistically meaningful. To achieve this, a null distribution was generated by shuffling the data labels and recalculating the clusters 1,000 times. For each permutation, hierarchical clustering was repeated, and the linkage distance between the radial and circular grating clusters was calculated. This process resulted in 1,000 linkage distance values, representing what would be expected under random conditions. Comparing the actual linkage distance to this null distribution revealed that the observed separation between the circular and radial gratings was significantly larger than expected by chance ($p < 0.001$, t-test). These results confirmed the statistical robustness of the clustering and highlighted the distinctiveness of these two image categories.

Neurometric Function Analysis

To quantify the neural discriminability of circular vs. radial grating stimuli across different noise levels, we computed neurometric functions based on the spike activity of individual V4 neurons. For each neuron, we first identified trials corresponding to circular and radial stimuli at each SNR level. Spike counts were measured in a 200 ms window starting at 50 ms after stimulus onset, and we computed the mean firing rate across trials for each condition.

For each neuron and SNR level, we then calculated the area under the receiver operating characteristic (ROC) curve (AUC), quantifying how well the neuron's spike rate distinguished between the two stimulus classes. Neurons were classified as circular- or radial-preferring based on the sign of their overall d-prime value across all conditions.

To compare population-level neural performance with behavior, we averaged AUC values across neurons within each preference group and computed the standard error of the mean (SEM). The resulting average neurometric data were then fit with a logistic function, using the same fitting procedure as for the behavioral psychometric data.

Statistical comparisons

Statistical comparisons of behavioral thresholds and bias values were based on the Wilcoxon signed-rank (WSR) test when the sizes of the data sets were equal; otherwise, the Wilcoxon rank-sum test (WRS) was used. For CPs, since we reported means, we used a t-test (MATLAB function) to evaluate their significance. Generally, using rank-based statistical methods produced similar patterns of p-values, sometimes even yielding lower p-values. Whenever multiple comparisons across conditions were required, we used Tukey's Honestly Significant Difference method to adjust for multiple comparisons, maintaining the false discovery rate (FDR) at 5% for all tests.

To estimate variability in the animals' performance, including psychometric thresholds and biases for Figure 7, we employed a nonparametric bootstrap procedure. Trial-level data were resampled with replacement within each stimulus (noise) level to generate 1,000 surrogate datasets. To estimate the standard error of the animals' performance at each noise level, we computed the standard

deviation of the resampled distribution, and the 2.5th and 97.5th percentiles were used to construct 95% confidence intervals. For each bootstrap sample, the psychometric function was re-fitted using the same method as applied to the original data. Thresholds—defined as the stimulus level corresponding to 75% correct performance—were extracted from each fit. The standard deviation of the resulting threshold distribution provided the standard error, and confidence intervals were computed as above. To estimate behavioral bias, we applied the same bootstrap procedure separately to data from pre-injection and 18-hour post-injection sessions. Bias was computed for each resampled dataset at each noise level, and the standard error was estimated from the distribution of bias values.

Current Biology, Volume *35*

Supplemental Information

Inactivation of primate cortex reveals inductive biases in visual learning

Pooya Laamerad, Matthew R. Krause, Daniel Guitton, and Christopher C. Pack

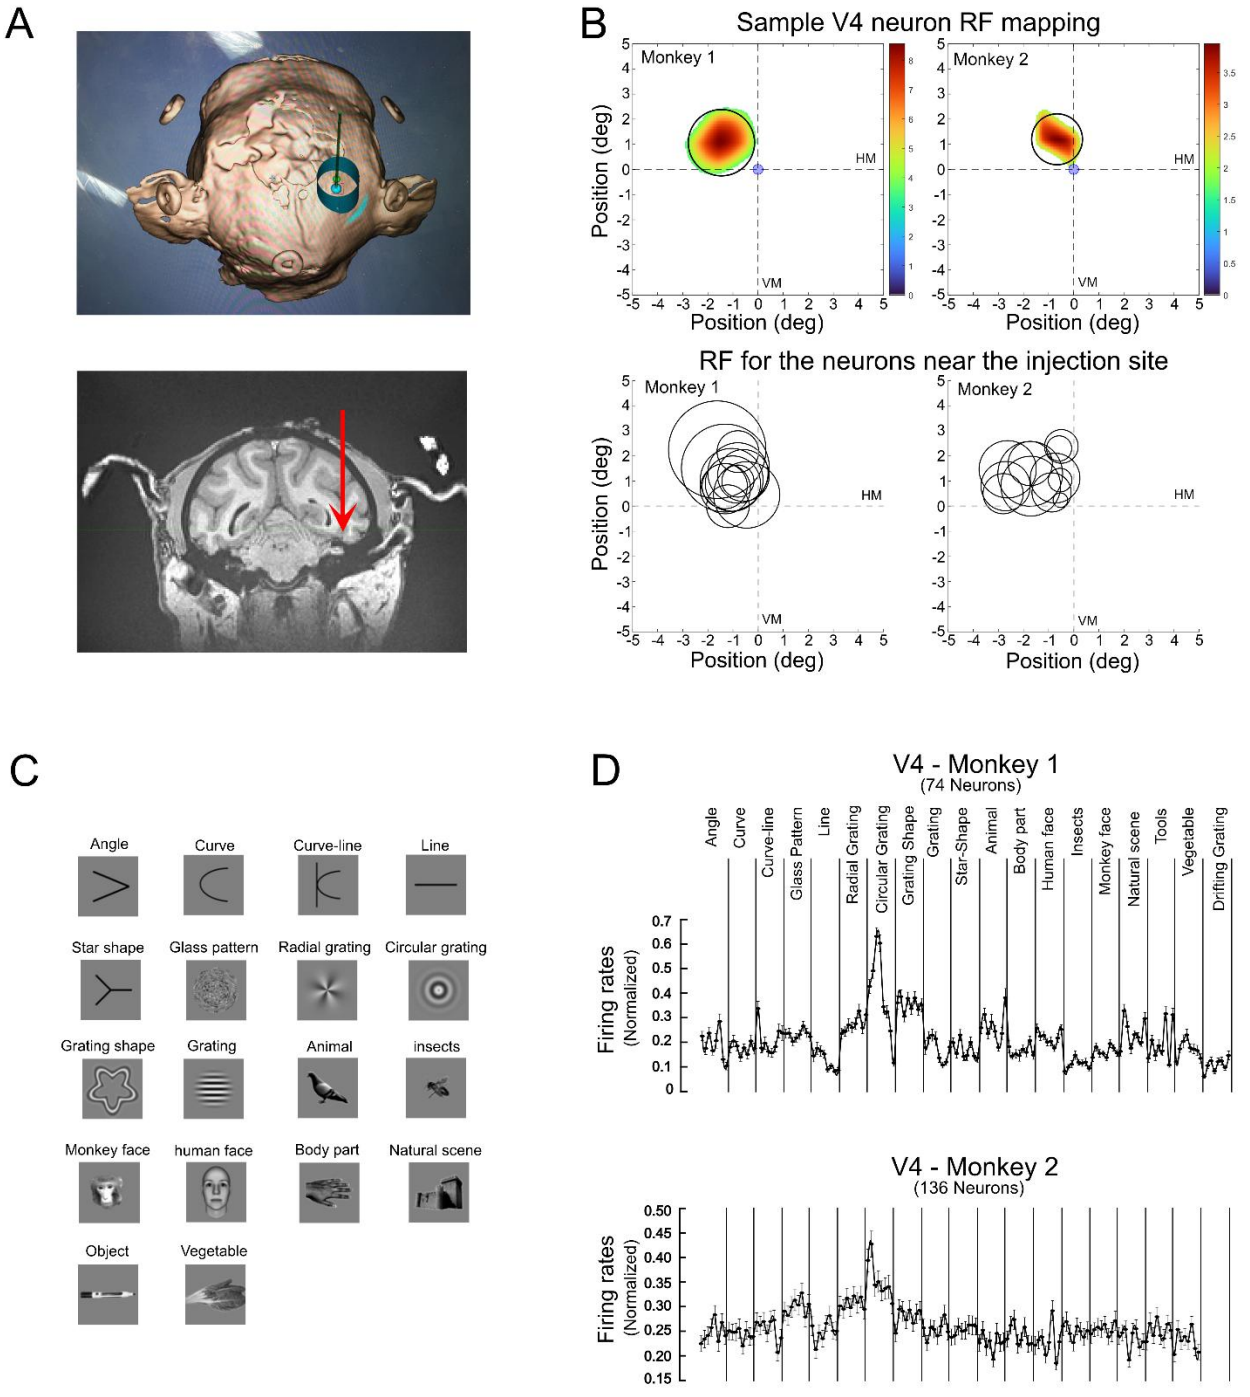

Figure S1: Electrode localization, receptive field mapping, and V4 visual responses. Related to Figure 1. A) Recording chamber and electrode array locations. The left panel displays the location of the electrophysiological recording chamber. It was positioned to access the brain region between the lunate sulcus and the superior temporal sulcus (STS). According to a standard macaque brain atlas and previous literature, dorsal and ventral V4 are positioned beneath this area. Each day, we lowered the electrode approximately 20-23mm to reach ventral V4. The right

panel shows a coronal MRI view from one of the monkeys. The red arrow indicates the approximate location of the electrode tip post-penetration. The injection sites were identical to the recording sites and remained consistent across all experimental phases for each monkey. **B)** Receptive field mapping for the neurons near the injection sites Top panel: Example RFs from each monkey based on average neuronal responses to sparse-noise stimuli. Colors within the circles indicate response strength, and each circle represents the best-fit Gaussian of the RF. Bottom panel: mean RF eccentricity was 1.4 ± 0.7º before training and 1.6 ± 0.7º after training in Monkey 1 ($p > 0.05$, WRS test), and 2.1 ± 0.6º and 2.4 ± 0.9º in Monkey 2 ($p > 0.05$, WRS test). The stimulus placement was based on the RF mapping. HM and VM indicate the horizontal and vertical meridian. **C)** Visual stimuli set. Sample images from each category used during the passive fixation task in the Preliminary Phase. Each category consists of 8 similar stimuli, varying by orientation, direction, or shape, depending on the category. For Monkey 1, drifting grating motion stimuli were also presented. This set of images was used in both the V4 and PIT areas. **D)** Response Activity of V4 Neurons to the set of visual stimuli. Normalized mean firing rates of V4 neurons in response to all stimuli presented during the passive fixation task in the Preliminary Phase. Points within each category represent variations of the stimulus class (e.g., different line orientations or faces). The top panel represents Monkey 1, and the bottom panel represents Monkey 2. For both monkeys, low spatial frequency circular gratings elicited the highest response activity in the V4 area.

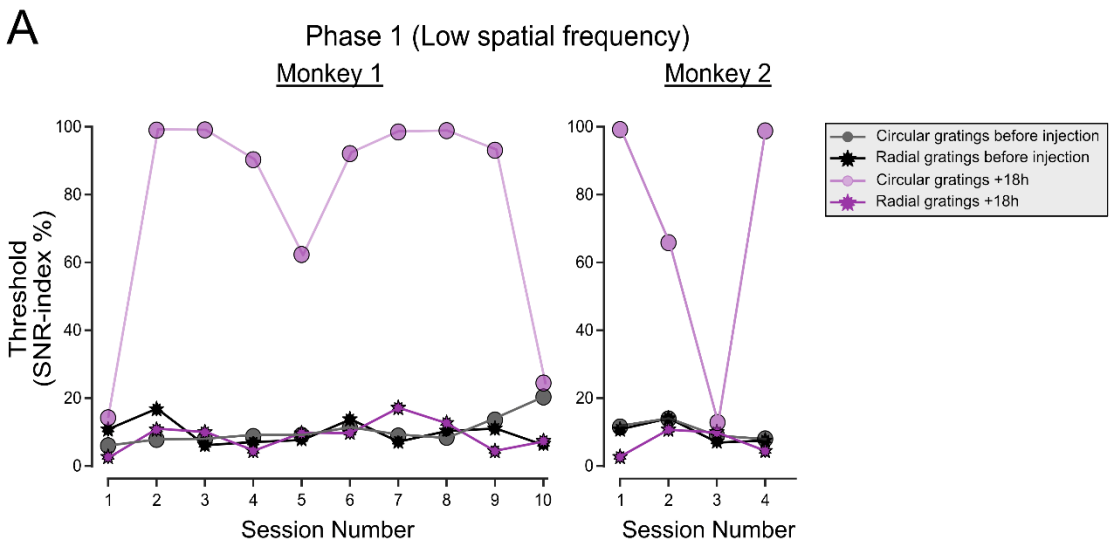
A
Phase 1 (Low spatial frequency)
Monkey 1
Monkey 2
Threshold (SNR-index %)
Session Number
Session Number
Circular gratings before injection
Radial gratings before injection
Circular gratings +18h
Radial gratings +18h

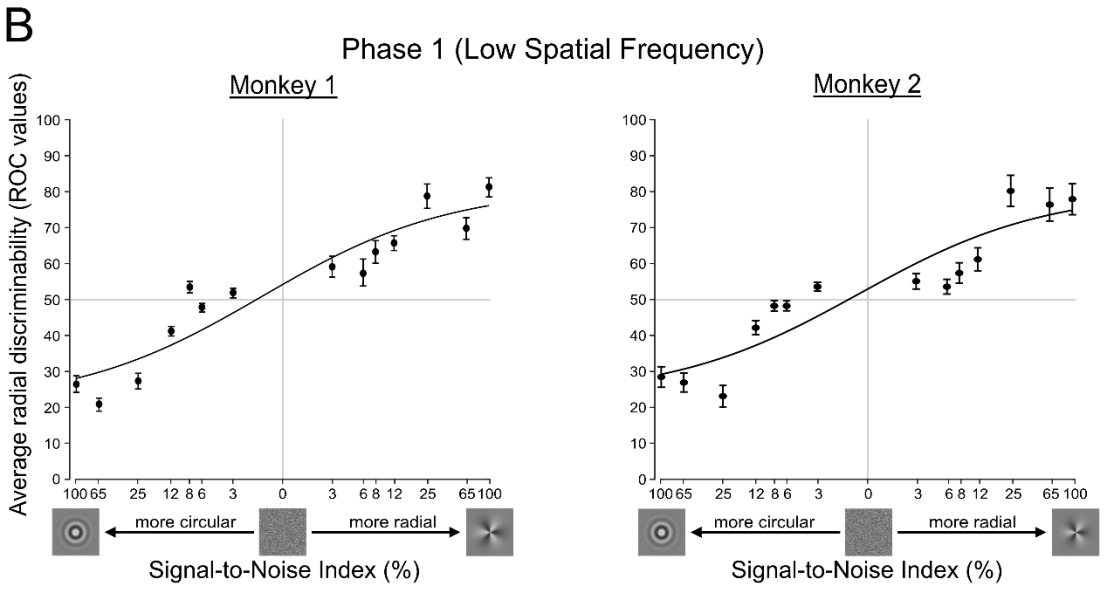
B
Phase 1 (Low Spatial Frequency)
Monkey 1
Monkey 2
Average radial discriminability (ROC values)
more circular
more radial
Signal-to-Noise Index (%)
Signal-to-Noise Index (%)

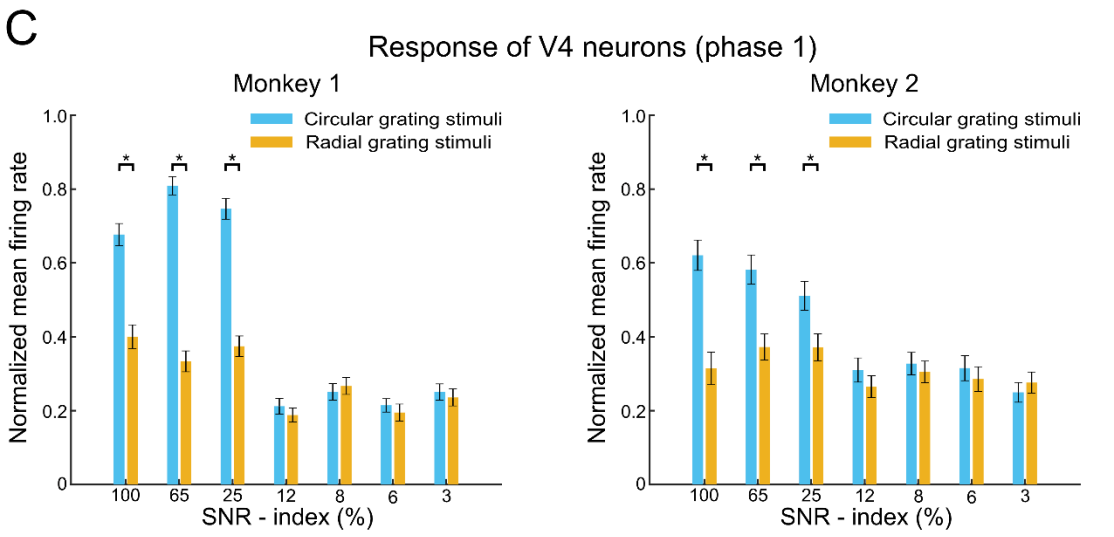
C
Response of V4 neurons (phase 1)
Monkey 1
Monkey 2
Circular grating stimuli
Radial grating stimuli
Normalized mean firing rate
SNR - index (%)
SNR - index (%)

Figure S2: Behavioral and neural responses to circular and radial gratings during Phase 1. Related to Figure 2. A) Psychophysical thresholds for both monkeys in Phase 1 (low spatial frequency stimuli). **B)** Average neuronal discriminability (AUC) as a function of SNR for circular and radial grating discrimination. Each panel corresponds to Phase 1 training for one monkey, matching the layout of the behavioral psychometric function figures. Neurons were grouped by stimulus preference (circular or radial) based on d-prime values. AUCs were averaged within each group, and logistic fits illustrate population-level neurometric performance. **C)** Normalized mean firing rates of V4 neurons across noise levels during Phase 1. Population responses differed significantly between the two grating types for the three least noisy stimuli ($p < 0.05$, t-test). Error bars indicate SEM; asterisks mark significant differences.

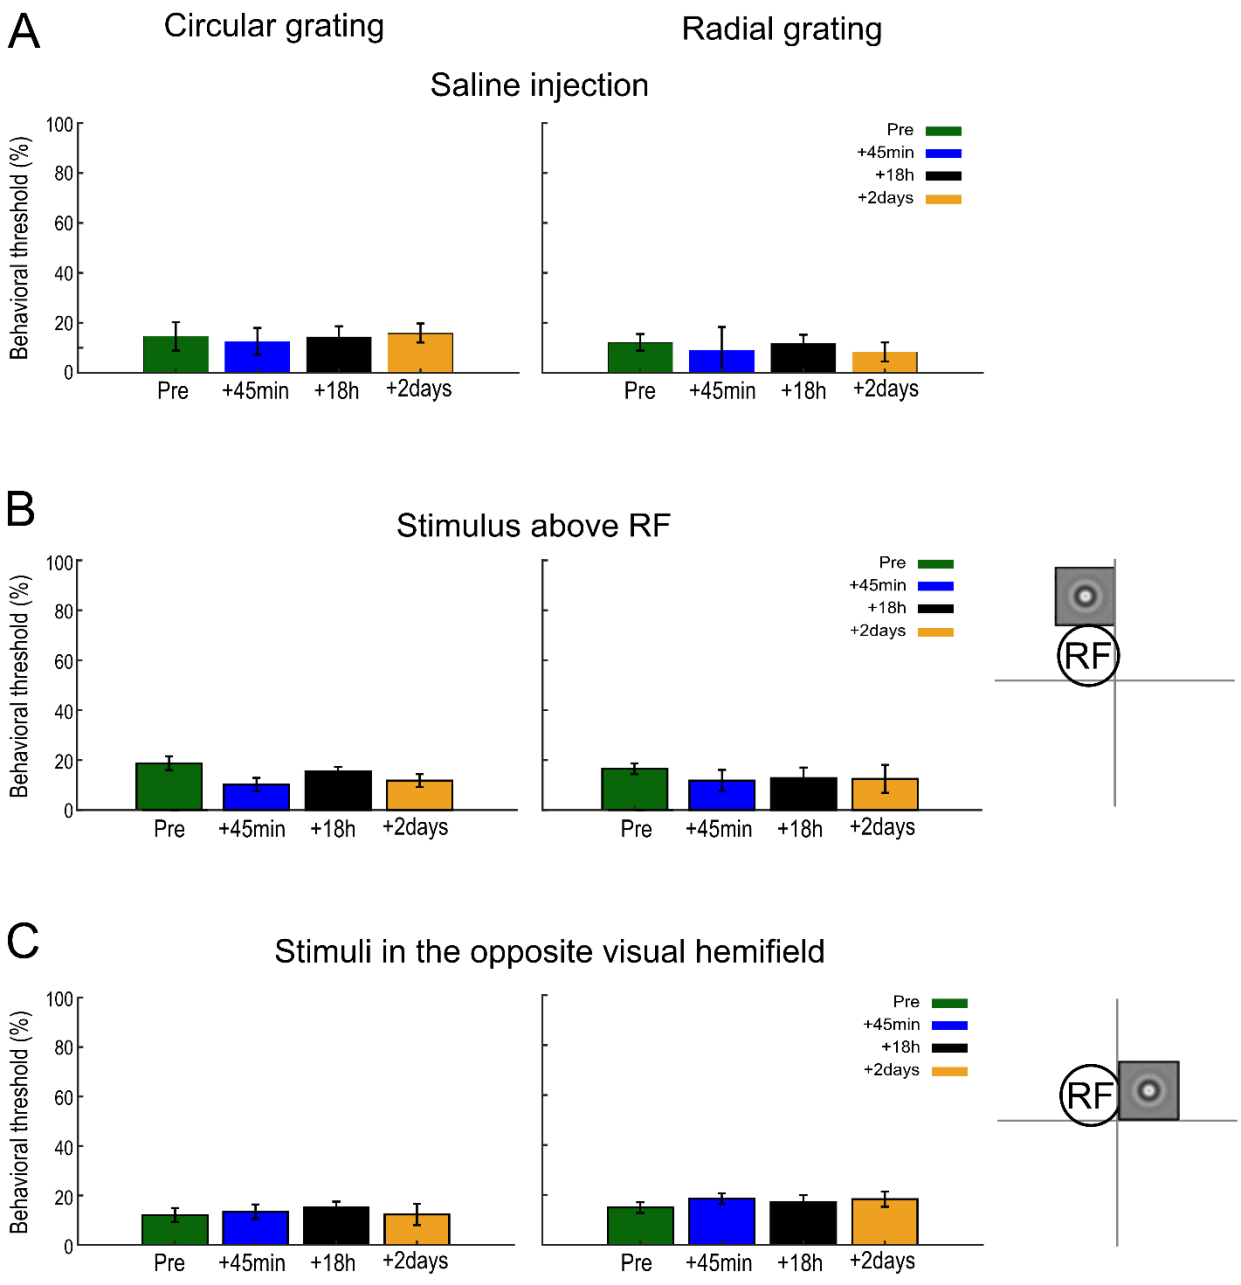

Figure S3: The effect of muscimol on behavioral performance is localized and specific. Related to Figure 3. A) As a control for potential non-specific effects of our inactivation procedure, we injected saline into the same sites and in the same amount as in the muscimol sessions. Saline injections had no effect on the animals' performance (top panel), confirming that the behavioral deficits were specifically due to muscimol. **B** and **C)** We also verified that the behavioral deficit was localized to the vicinity of the RFs of the neurons at the injection site. In Monkey 2, we placed gratings of the same size around the injection site, within the vicinity of the RFs of V4 neurons recorded. When the stimulus was positioned either above the vicinity of the RFs (B) or in the other hemisphere on the right visual field (C), muscimol injection did not change the behavioral

thresholds. This demonstrates that the effect of injection is spatially localized and specific to the type of drug used for inactivation.

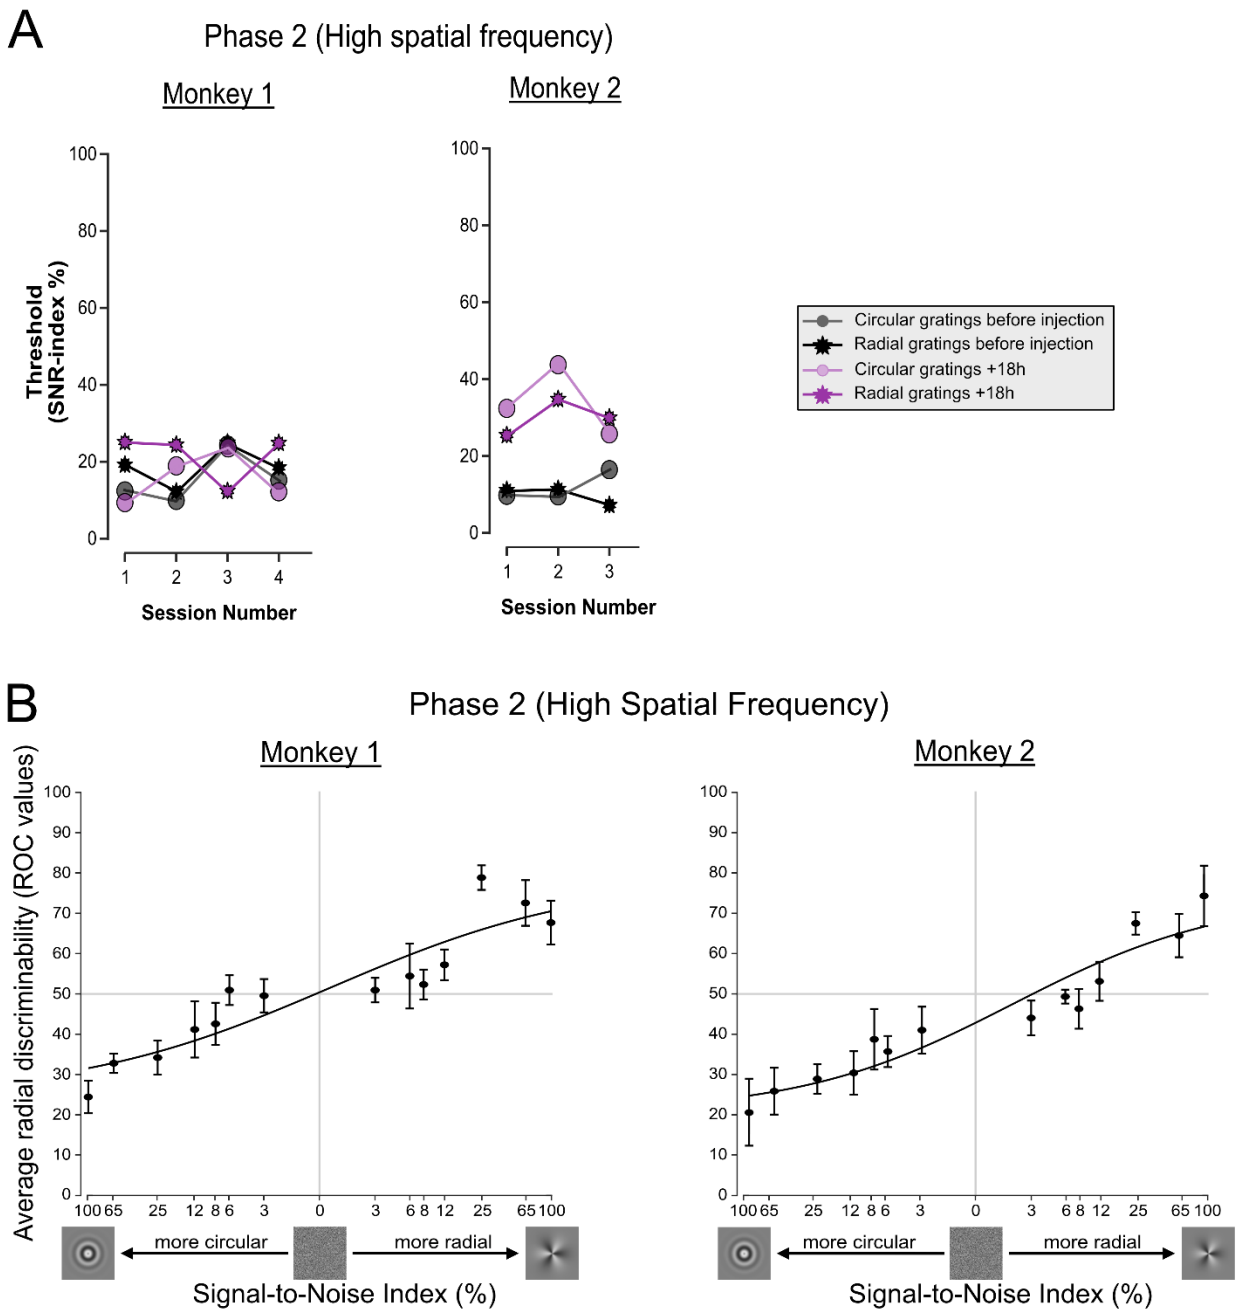

Figure S4: Behavioral and neural responses to circular and radial gratings during Phase 2 (high spatial frequency stimuli). Related to Figure 4. A) Psychophysical thresholds for both monkeys in Phase 2 (high spatial frequency stimuli). **B)** Average neuronal discriminability (AUC) as a function of SNR for circular and radial grating discrimination. Each panel corresponds to Phase 2 training for one monkey, matching the layout of the behavioral psychometric function figures. Neurons were grouped by stimulus preference (circular or radial) based on d-prime values. AUCs were averaged within each group, and logistic fits illustrate population-level neurometric performance. Error bars indicate SEM.

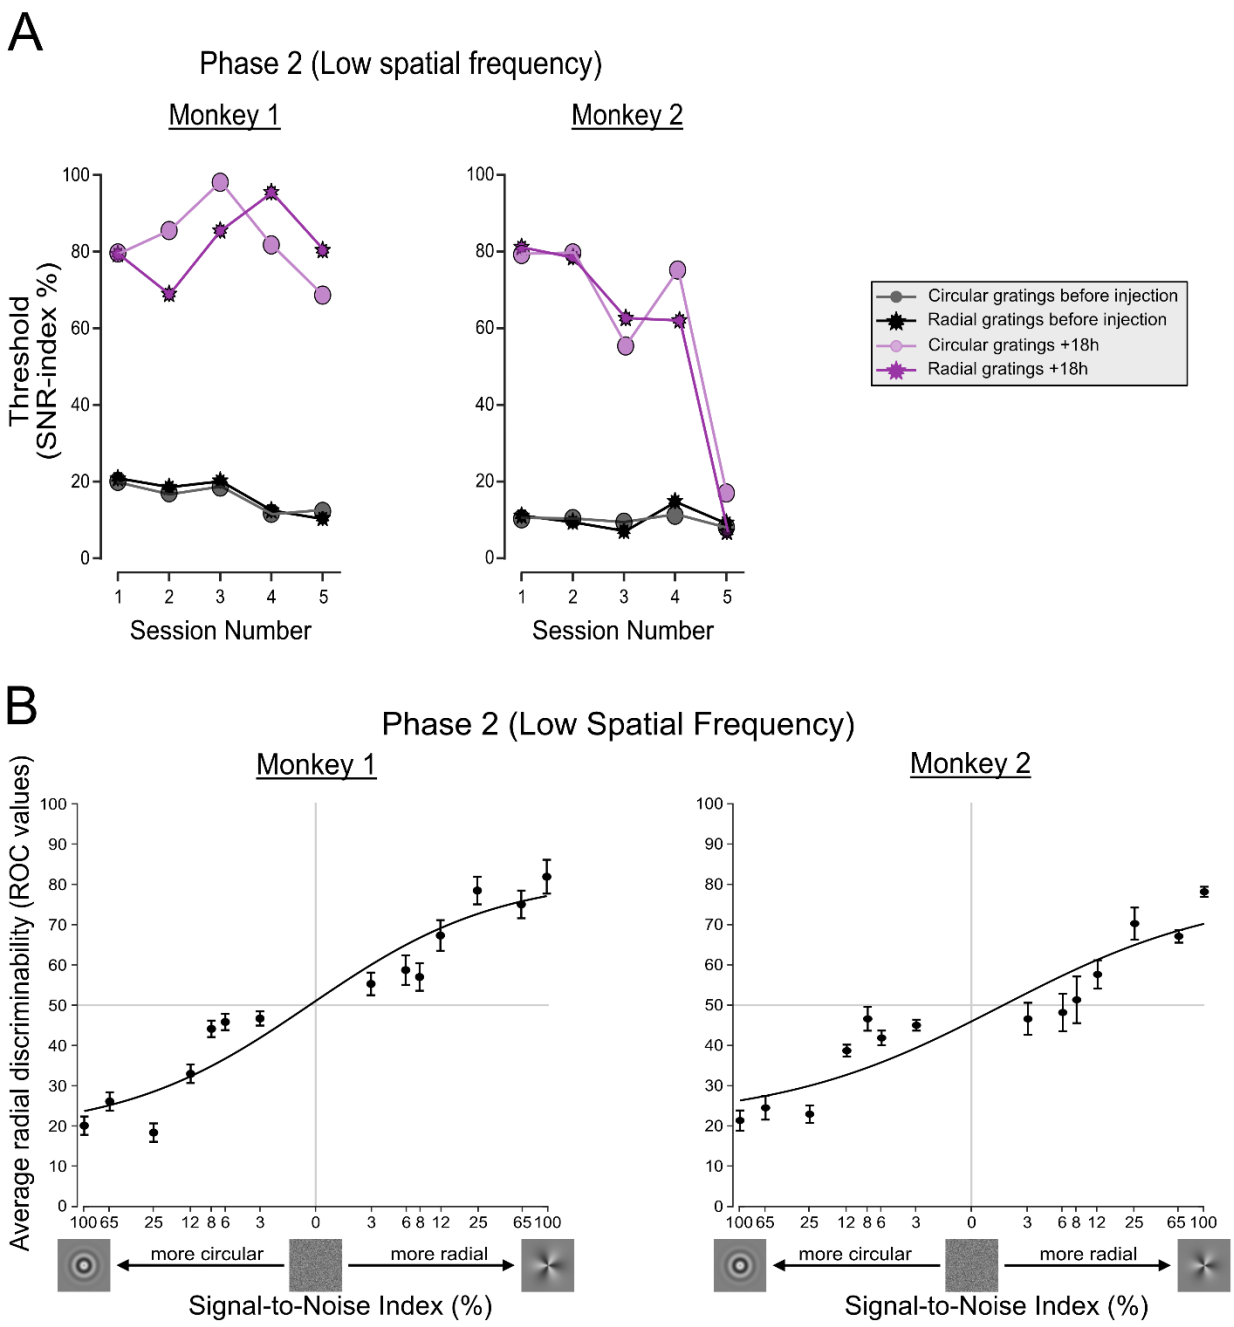

Figure S5: Behavioral and neural responses to circular and radial gratings during Phase 2 (low spatial frequency stimuli). Related to Figure 5. A) Psychophysical thresholds for both monkeys in Phase 2 (low spatial frequency stimuli). **B)** Average neuronal discriminability (AUC) as a function of SNR for circular and radial grating discrimination. Each panel corresponds to Phase 2 training for one monkey, matching the layout of the behavioral psychometric function figures. Neurons were grouped by stimulus preference (circular or radial) based on d-prime values. AUCs were averaged within each group, and logistic fits illustrate population-level neurometric performance. Error bars indicate SEM.

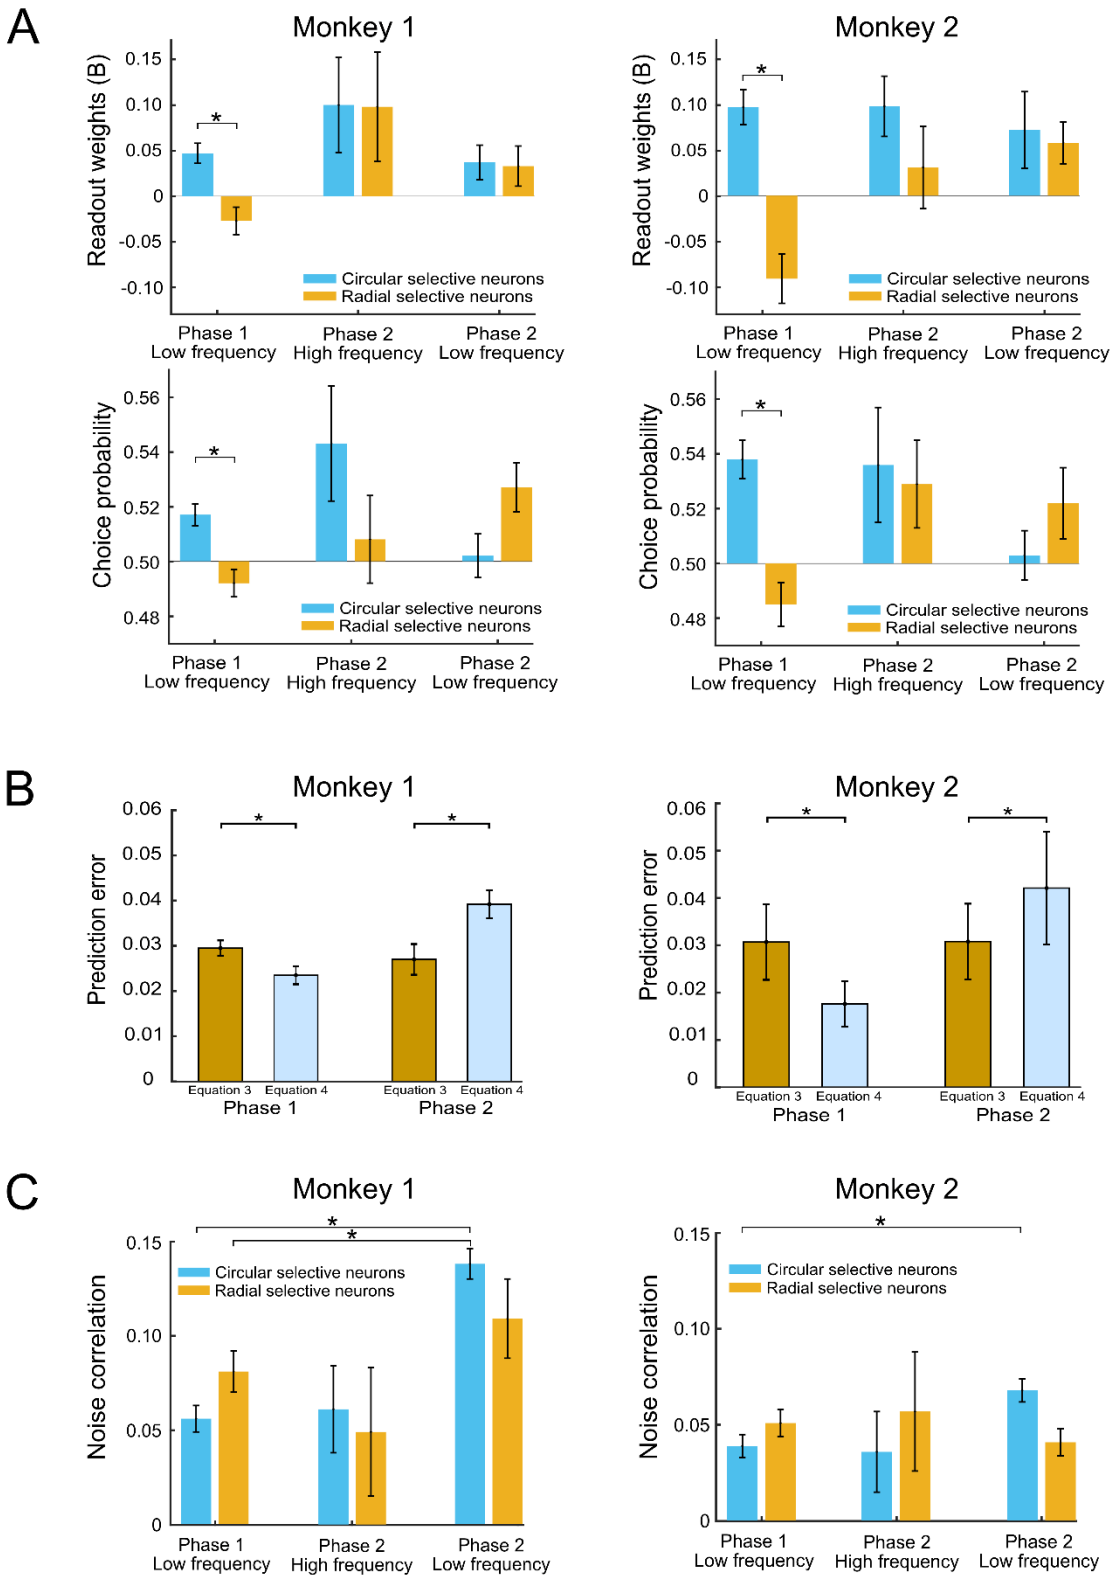
A
Monkey 1
Monkey 2
Readout weights (B)
Choice probability
Circular selective neurons
Radial selective neurons
Phase 1
Low frequency
Phase 2
High frequency
Phase 2
Low frequency
B
Monkey 1
Monkey 2
Prediction error
Equation 3
Equation 4
Phase 1
Phase 2
C
Monkey 1
Monkey 2
Noise correlation

Figure S6: Phase-dependent modulation of readout weights, choice probabilities, and noise correlations in V4. Related to Figure 6. A) Phase-dependent changes in V4 neuronal readout weights and choice probabilities for each monkey. Mean readout weights of V4 neurons. Readout weights were calculated for neurons with different stimulus preferences across experimental phases and spatial frequencies. In Phase 1, for both monkeys, radial-selective neurons (orange) had negative weights, which were significantly different from circular-selective neurons (blue) ($p < 0.05$; ANOVA, FDR-corrected). In Phase 2, no significant difference was observed between the two groups for both low- and high-spatial-frequency gratings. Mean choice probabilities (CPs) of V4 neurons. A similar pattern was observed in CPs, reflecting the trends seen in the readout weights. **B)** CP prediction error based on relative CP. This analysis characterizes noise correlations within and between subpopulations of neurons with different stimulus preferences (Equation 3; see Methods) or across all neurons as a single unified population (Equation 4; see Methods). In Phase 1, Equation 4 (single unified population) predicted CP values significantly better than Equation 3 ($p < 0.05$, t-test). However, in Phase 2, Equation 3 (two subpopulations) outperformed Equation 4 ($p < 0.05$, t-test). Prediction errors were calculated by comparing the predicted CP values to the real (observed) values using the Mean Squared Error. **C**) Noise correlations of V4 neurons across phases. We observed a significant increase in pairwise noise correlations, from Phase 1 to Phase 2 in both monkeys for low spatial frequency stimuli. This increase was primarily driven by circular-selective neurons, which exhibited a significant rise in noise correlations from Phase 1 to Phase 2 ($F(5,602) = 5.41$, $p < 0.001$ for Monkey 1, and $F(5,334) = 4.62$, $p = 0.002$ for Monkey 2; ANOVA, FDR-corrected). In contrast, radial grating-selective neurons showed no significant change in noise correlations between phases ($p > 0.05$ for Monkey 1 and 2, ANOVA, FDR-corrected). Noise correlations computed for high spatial frequency stimuli did not significantly differ from those for low spatial frequency stimuli across both phases ($p > 0.05$ for Monkey 1 and 2, ANOVA, FDR-corrected). Error bars indicate the standard error of the mean (SEM). Asterisks denote statistically significant differences.

Phase	Spatial frequency	Stimulus preference	Choice probability (CP)	Significance of CP vs. 0.5	Noise correlation	Readout weights (β)	Significance of weights vs. 0
1	Low	Circular	0.526 ± 0.004	p < 0.01	0.051 ± 0.006	0.065 ± 0.099	p < 0.01
1	Low	Radial	0.483 ± 0.006	P < 0.01	0.071 ± 0.009	-0.052 ± 0.014	P < 0.01
2	High	Circular	0.539 ± 0.013	P = 0.012	0.050 ± 0.009	0.099 ± 0.034	P < 0.01
2	High	Radial	0.509 ± 0.014	P = 0.32	0.051 ± 0.008	0.066 ± 0.039	P = 0.047
2	Low	Circular	0.502 ± 0.008	P = 0.76	0.124 ± 0.007	0.048 ± 0.018	P = 0.013
2	Low	Radial	0.523 ± 0.011	p = 0.011	0.089 ± 0.011	0.041 ± 0.017	p = 0.032

Table S1 - Choice probability (CP), noise correlation, and readout weights (β) across phases for both monkeys, analyzed for each subpopulation of V4 neurons. Related to Figure 6.