

High precision coding in mouse visual cortex

Carsen Stringer^{1*}, Michalis Michaelos¹, Marius Pachitariu^{1*}

¹HHMI Janelia Research Campus, Ashburn, VA, USA

* correspondence to ([stringerc](mailto:stringerc@janelia.hhmi.org), [pachitariu](mailto:pachitariu@janelia.hhmi.org)) @ janelia.hhmi.org

ABSTRACT

Single neurons in visual cortex provide unreliable measurements of visual features due to their high trial-to-trial variability. It is not known if this “noise” extends its effects over large neural populations to impair the global encoding of sensory stimuli. We recorded simultaneously from ~20,000 neurons in mouse visual cortex and found that the neural population had discrimination thresholds of 0.3° in an orientation decoding task. These thresholds are ~100 times smaller than those reported behaviorally in mice. The discrepancy between neural and behavioral discrimination could not be explained by the types of stimuli we used, by behavioral states or by the sequential nature of trial-by-trial perceptual learning tasks. These results imply that the limits of sensory perception in mice are not set by neural noise in sensory cortex, but by the limitations of downstream decoders.

INTRODUCTION

Sensory neurons respond with high variability to repeated presentations of the same stimulus [1–8]. This variability is thought to limit the accuracy of perceptual judgements because it is pervasive [9–12], and it is reduced during attentional engagement [13–15] and over the course of perceptual learning [16]. The hypothetical links between neural variability, sensory information and perceptual judgements stand at the foundation of several theoretical frameworks such as the efficient coding hypothesis [17], the ideal observer model [18], the Bayesian coding hypothesis [19–21] and the probabilistic sampling hypothesis [22, 23].

However, it is not clear how variability measured from single neurons or from pairs of neurons scales to local circuits of tens of thousands of neurons [10]. Intuitively, one might expect the noise to be averaged out over large enough numbers of neurons. Theoretical studies have shown that most types of noise are indeed harmless at the population level [24–27]; only a special kind of correlated noise is detrimental to neural coding because it can limit the total amount of information available in the system. This “information-limiting” noise arises when the estimation errors of single neurons are correlated to each other at the level of the population [27].

In geometric terms, noise can only affect the encoding of a stimulus when it aligns to the same neural subspaces which the stimuli drive [28–30]. We now know that at least some of the neural noise is orthogonal to the stimulus subspace, and thus harmless. Specifically, we have shown orthogonality to the stimulus space for the neural variability induced by the animal’s own behavior across its entire brain [31]. However,

this behavior-related neural activity only accounted for ~35% of the total coordinated variability, leaving the possibility that the rest is stimulus-related and thus potentially information-limiting.

Estimating the impact of information-limiting noise on coding is difficult, because even small amounts of information-limiting noise can put absolute bounds on the precision of stimulus encoding [27]. To detect the effects of such small noise, it is thus necessary to record from large numbers of neurons. This introduces a new problem: the information content of such recordings cannot be estimated directly and decoding approaches must be used instead. Previous studies in anesthetized macaques have decoded one-dimensional variables (i.e. the orientation of a grating) from populations of at most 100 simultaneously-recorded neurons and reported decoding errors in the 2–20° range [32–34]. It is not known if these errors represent an absolute lower bound, or may decrease further for larger neural populations. Some studies have suggested that small subsets of neurons are as discriminative as the entire population [33, 35], while others have shown consistent improvements in discrimination with increasing numbers of neurons [36].

Here we aimed to measure the absolute lower bound on decoding error for visual stimulus orientation by recording from populations of ~20,000 neurons in mice. If information-limiting noise exists, we reasoned that the decoding error must asymptote at some non-zero value as we increase the number of neurons we consider [27]. Instead, we found that we could decode the orientation of a grating with mean errors as low as 0.3°, and this decoding error did not asymptote with the number of neurons or stimuli we presented. To achieve this decoding performance, it was important to take correlations into account, but these correlations were not of the information-limiting type. We also show that the low decoding errors can be achieved by a variety of biologically plausible learning strategies based on linear decoders. These results imply that the visual cortex encodes visual features to high precision on a trial-by-trial basis in mice, a species not known for high acuity vision, and which performs poorly in orientation discrimination tasks [37, 38].

RESULTS

We recorded from primary visual cortex in awake, head-fixed mice that were free to run on an air floating ball. Each session lasted for 120–180 minutes during which we presented images to the left eye (Figure 1A). Our main stimuli were static gratings, which lasted for 750ms and were rotated at a random orientation on each trial. We recorded neural activity from visual cortex using multi-plane two-photon calcium imaging,

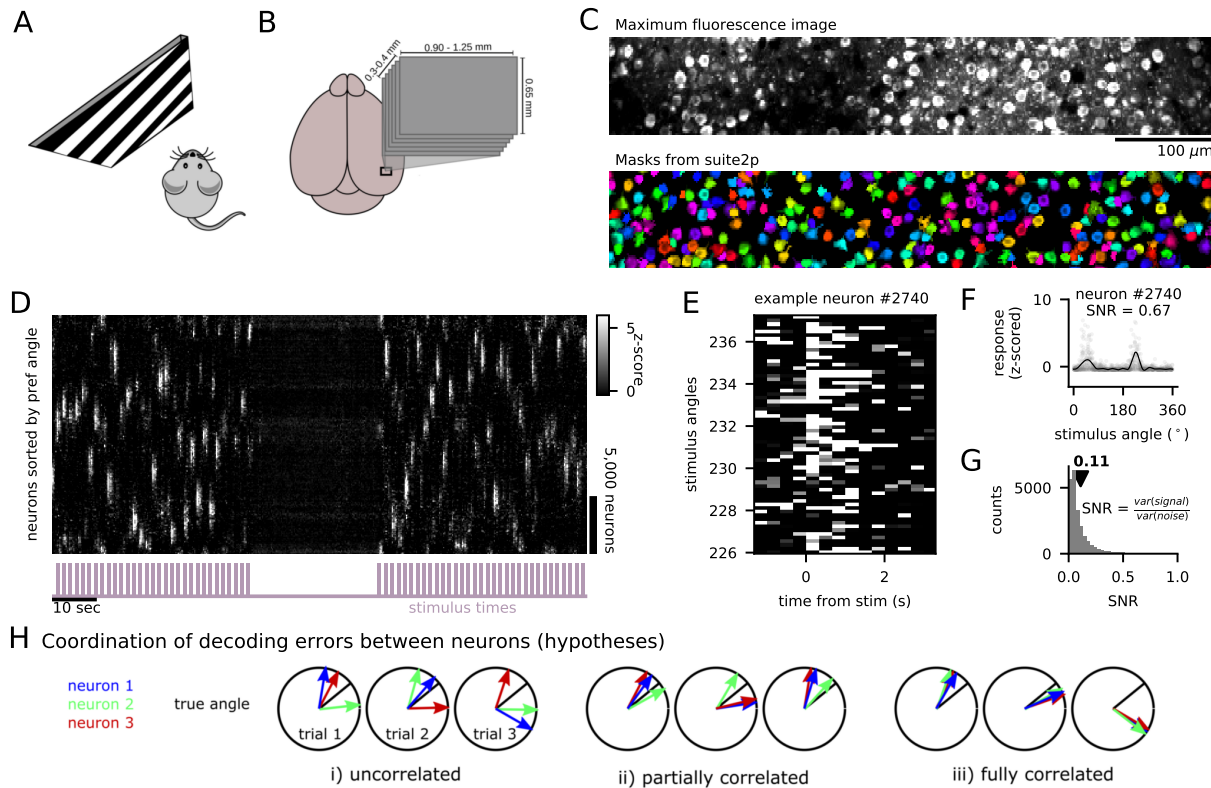


Figure 1: Recording setup and single-neuron variability. (A) Visual presentation setup. (B) Multi-plane imaging setup. (C) Section from a recorded plane. (top plot) maximum fluorescence image. (bottom plot) cell masks detected with Suite2p, randomly colored. (D) Activity raster of all neurons in one recording in response to static gratings of various angles. Neurons are sorted along the vertical axis by their preferred angle, and smoothed with a Gaussian of width 50 neurons. (E) Single trial stimulus responses of an example neuron. (F) Responses from (E) averaged over 3 bins following stimulus onset (gray dots), with average tuning curve (black line). (G) Distribution of the signal-to-noise ratio (SNR) across neurons in this session. (H) Hypotheses for the alignment of decoding errors across the neural population. Each circle represents a different trial. The black line denotes the true angle of the stimulus. Each arrow represents the decoded angle from a single neuron’s response. (i) the errors are uncorrelated across neurons. (ii) some errors are correlated. (iii) all errors are correlated, and therefore information about the stimulus is lost.

with 10-17 planes spaced $25 \mu\text{m}$ apart in depth, scanning the entire stack repeatedly at an average 3 Hz (Figure 1B). We obtained $19,665 \pm 3,062$ (s.d., $n=6$) neurons per recording using the processing pipeline Suite2p [39] (Figure 1C). All analyses were performed on deconvolved data, which localizes in time the extended fluorescence responses of the calcium indicator [40]. We have publicly shared the data and code for this paper [41].

To visualize the patterns of population activity in a raster plot, we sorted neurons by their preferred stimulus and binned their activity along this stimulus-sorted axis (Figure 1D). As previously shown [42], single neurons had high trial-to-trial variability and sometimes failed to respond at all to their preferred stimulus (Figure 1E). We quantified this variability by the signal-to-noise ratio (SNR) (Figure 1F), and found a mean SNR of 0.11 for the example session, and 0.13 ± 0.01 (s.e.m., $n=6$) across recordings, similar to the single-trial SNR previously reported for natural image stimuli [36]. The aligned, population-averaged tuning curves had a mean half-width at half-max of 14.1° (Figure S1). The correlation between neural

response vectors for different orientations decayed smoothly as a function of the difference in orientation (Figure S2A,B). Using a manifold embedding algorithm in three dimensions (ISOMAP, [43]), we found a predominantly one-dimensional representation of the stimulus space, but nonetheless a representation corrupted with noise (Figure S2C).

We distinguish between three types of potential neural variability that can affect coding (Figure 1H). First, the decoding errors of single neurons might be uncorrelated across neurons, in which case averaging over single-neuron predictions would give an unbiased estimate of the true stimulus angle with a small error (Figure 1Hi). A second possibility is that decoding errors are partially-correlated over the population, for example if subsets of neurons have correlated errors (Figure 1Hii). The third and final possibility is that decoding errors are fully correlated over the population, so that averaging over the predictions of even infinitely many neurons would give a biased estimate of the true stimulus orientation (Figure 1Hiii). This situation would indicate the presence of “information-limiting” correlations [27].

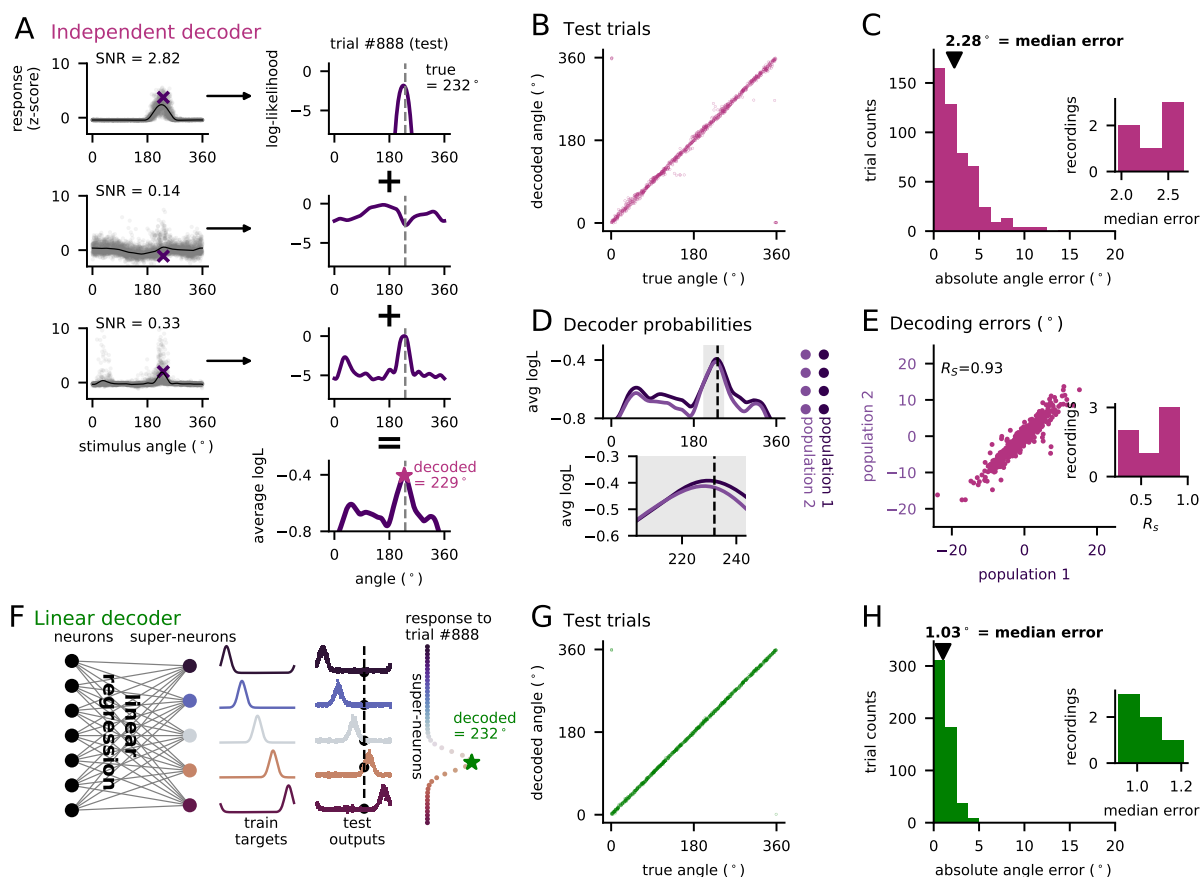


Figure 2: Decoders of stimulus orientation. (A) Single-neuron single-trial stimulus responses (gray dots) with average tuning curves (black line). Right panel: log-likelihood of stimulus orientation based on each neuron’s observed response on trial #888. Actual stimulus is gray dotted line. Right, bottom panel: average log-likelihood over all neurons. (B) True versus decoded stimulus on test trials. (C) Distribution of absolute angle errors across trials. (inset: distribution of median decoding errors across recordings). (D) Neurons were split into two random populations and the decoder was fit separately to each population. (E) Scatter plot of decoding errors from the two populations. The Spearman correlation R_s is reported (inset: distribution of R_s across recordings). (F) Linear decoder schematic, with “super-neuron” responses to test stimuli. (G,H) Same as B,C for the linear decoder.

Independent neuron decoder: 2.4° error

Consider the first possibility, that the stimulus-dependent variability is uncorrelated between neurons. If that was true in the data, we could decode from the neural population using the “Naive Bayes” classifier [44], in which every observed response $R(n, \theta)$ of a neuron n to a stimulus θ determines the probability $P(\theta' | R(n, \theta))$ which that neuron gives to every possible stimulus θ' (Figure 2A). Assuming independence, we can multiply these probabilities across neurons to obtain the likelihood $P(\theta' | R(1, \theta), R(2, \theta), \dots)$ that the entire population gives to every possible stimulus θ' . Equivalently, we can sum the log-likelihoods across neurons (Figure 2A). We then select the stimulus with the highest likelihood as the decoded orientation, using interpolation to predict fractional stimulus values (see Methods). Finally, we can define $P(\theta' | R(n, \theta)) \sim \mathcal{N}(R(n, \theta) | f_n(\theta'), \sigma_n(\theta'))$, where f_n, σ_n are means and variances optimized on the training data (Figure 2A).

This population decoder had a median decoding error of

$2.39^\circ \pm 0.13^\circ$ (s.e.m., $n=6$) (Figure 2B,C). This error may be due to either: 1) single-neuron noise that was not fully averaged out, or 2) correlations in decoding errors between neurons. To test if the decoding errors were correlated, we split the neurons into two populations, decoded from each, and asked if the two decoding errors were in the same direction with respect to the true stimulus (Figure 2D). We found that the errors were highly correlated (Spearman’s $R=0.65 \pm 0.11$ (s.e.m, $n=6$)), which invalidated the independence assumption of the decoder, and suggests that the error may decrease further if correlations are taken into account (Figure 2E).

Linear decoders ignore correlations: 1° error

To account for correlations, a decoder must be able to appropriately weigh neurons, potentially discarding neurons that have too high correlations. This is most easily achieved by simple linear decoders, which we modify here to predict continuous functions of the stimulus angle (Figure 2F),

rather than the stimulus angle itself, which is a circular variable and thus not directly predictable by a linear method. We called these intermediate functions “super-neurons” and chose them to have von Mises tuning to orientation θ (Figure 2F): $F(\theta|\theta_{\text{pref}}) = \exp(\cos(\theta - \theta_{\text{pref}})/\sigma)$, with $\sigma = 0.1$, with different super-neurons having different values of θ_{pref} that tile the full range $[0, 2\pi]$. On test data, we computed the activations of the super-neurons and picked the preferred angle of the super-neuron with the highest activation as the decoded stimulus, using interpolation methods to decode fractional angles (Figure 2F).

The error of the linear decoder was $1.03^\circ \pm 0.04^\circ$ (s.e.m., $n=6$) compared to the 2.39° for the independent decoder (Figure 2G,H). To achieve this decoding error, more than 256 principal components were needed (Figure S3). The best decoding was achieved for horizontal gratings, and the worst for vertical gratings, although the differences were small (Figure S4). The decoder was not affected by the presence of spontaneous activity patterns [31], and performed equally well whether these patterns were subtracted off or not (Figure S5).

To determine if the linear decoder achieved the minimum possible error, we analyzed the asymptotic behavior of the error as a function of the number of neurons and stimuli. As opposed to the independent decoder, which nearly achieved its saturating value of $\sim 2.12^\circ$ (Figure S6A), the linear decoder’s error continued to decrease with increasing numbers of neurons (Figure S6B). We estimated that its error would plateau at 0.71° in the limit of very many neurons. We performed a similar scaling analysis with the number of training set stimuli (Figure S6C). Again we found that the error continued to decrease with increasing numbers of stimuli, presumably due to less overfitting. We estimated that the error would saturate at 0.38° for very many training stimuli.

This asymptotic analysis suggested that the most effective way to further lower decoding errors would be to show more stimuli. To do this we restricted the range of presented angles to $43\text{--}47^\circ$, and presented $\sim 4,000$ stimuli in this range. This increased the stimulus density from 10 trials/deg to 1,000 trials/deg. To avoid boundary effects when decoding orientations in such a limited range, we switched to a decoding task in which we infer if the true stimulus is above or below the middle value of 45° .

Neural discrimination thresholds of 0.3°

For this new decoding task we built structured decoders that predicted a single function f of the presented stimulus θ and the threshold stimulus θ_0 : $f(\theta) = F(\theta|\theta_0 + 15^\circ) - F(\theta|\theta_0 - 15^\circ)$, where F is the von Mises function defined above. We used the sign of the neural prediction on test trials as the predicted class label of the decoder which for convenience we refer to as “left” vs “right”.

The neural discrimination threshold is defined as the angle at which the decoder achieves 75% correct performance. The discrimination threshold was 1.04° (Figure 3B) for the original

stimulus set at 10 trials/deg, and 0.32° for the new stimulus set of 1,000 trials/deg (Figure 3C). An asymptotic analysis shows that the error might have decreased further with more neurons and more trials (Figure 3D). As a sanity check, we performed control recordings with the laser path blocked (shutter on) and found no ability to decode from the very low bleedthrough of the screen into the microscope (see Figure S7).

In contrast to these low neural discrimination thresholds, behavioral discrimination thresholds in mice are $>25^\circ$ [37, 38] (Figure 3B). We wondered if some of the difference between neural and behavioral discrimination could be related to behavioral states, because the mice in our experiments were free to run on an air-floating ball. To quantify running-related differences, we split test trials (but not training trials) in two groups based on locomotion. We found that passive trials had modestly increased discrimination thresholds of 1.25° compared to 0.96° on running trials (Figure 3E).

Thus, behavioral states cannot account for the discrepancy between behavioral and neural discrimination thresholds. We next asked if the discrepancy may be accounted for by stimulus properties, and used neural responses to new stimulus sets (Figure 3F) to investigate this possibility. We varied the size (full field vs 30°) and duration (750ms vs 100ms) of the static grating stimuli. We also presented drifting gratings (2Hz), drifting gratings with low contrast (1 vs 0.05) and drifting gratings with low contrast and large added noise. Finally, we showed a complex stimulus without a well-defined absolute orientation, and rotated it around its center. These manipulations either did not increase discrimination thresholds or did so modestly, up to at most 1.87° for the low-contrast, drifting gratings (Figure 3F).

We also considered the possibility that the neural code might change over the duration of the recording, making decoders obtained from the first half of the recording ineffective on the second half. We therefore split train/test trials chronologically rather than randomly. We found a modest increase in discrimination threshold to 1.14° compared to the original 1.04° (Figure S8). We did not observe a change in the discrimination threshold across cortical layers (L2/3 vs L4, Figure S9).

Thus, neither the stimulus properties nor the behavioral states can account for the discrepancy between behavioral and neural discrimination thresholds. We conclude that mice might not be using their available neural information efficiently in behavioral tasks. This may be a consequence of trial-by-trial learning limitations, which we investigate next in simulations.

Biologically-plausible learning by perceptrons

Our first hypothesis was that the sequential nature of trial-by-trial learning makes it difficult to learn optimal decoders without storing the neural responses to all previously presented stimuli, which might be unfeasible (but see hippocampal replay e.g. [45, 46]). To construct an online decoder, which pro-

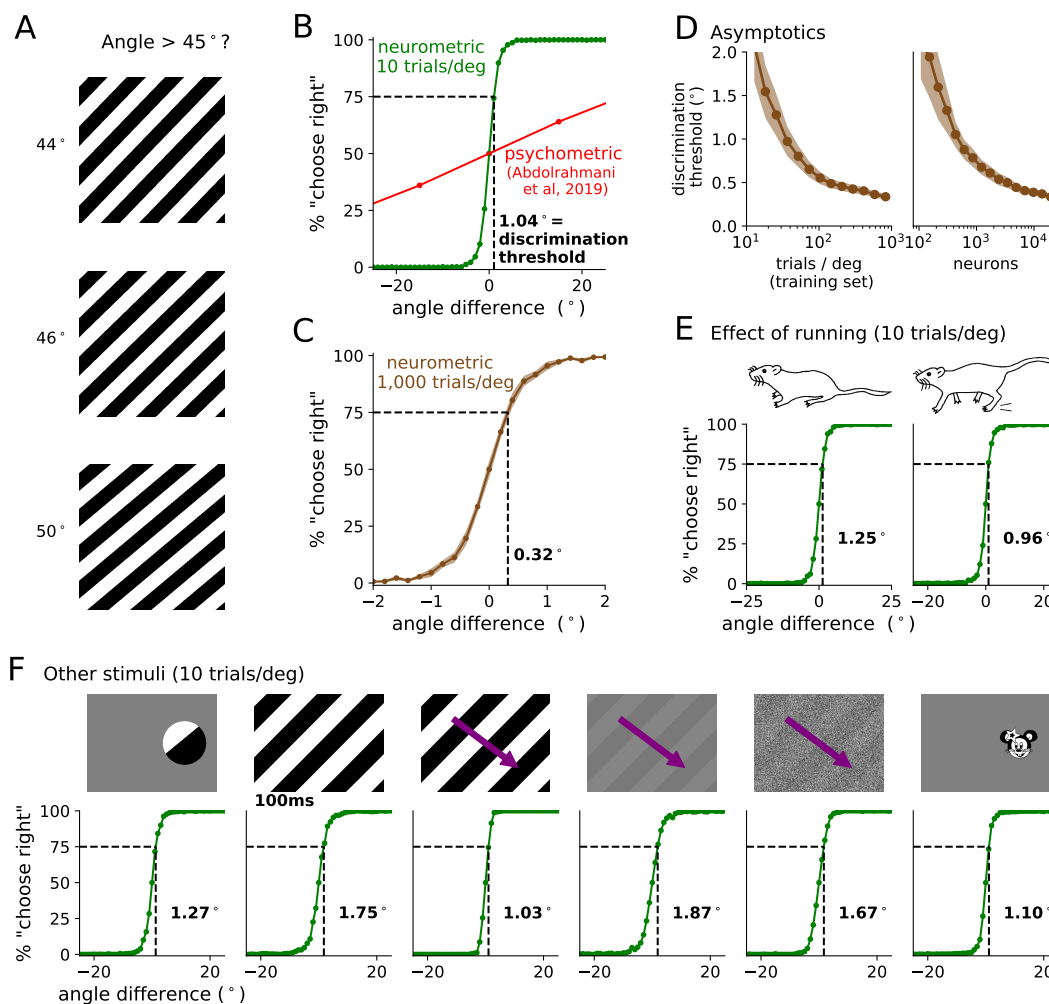


Figure 3: Orientation discrimination. (A) Decoding task set up. (B) Neurometric curves were symmetrized and averaged across sessions. The discrimination threshold was computed as the angle difference for which 75% correct performance is achieved. Psychometric results reproduced from [38]. (C) In another set of recordings, we presented stimuli in the range 43-47°, obtaining a higher density of training trials of 750 trials/degree. The analysis in B was repeated on this data. (D) Asymptotic behavior of the discrimination threshold as a function of the number of trials/deg (left) and the number of neurons (right). (E) Neurometric curves for subsets of trials divided into running and passive trials. (F) Neurometric curves for other stimulus types: localized, short, drifting, drifting + low-contrast, drifting + low-contrast + noise, and complex.

cesses trials sequentially, we used perceptrons [47–51] (Figure 4A). The perceptron sums the activities of its input neurons multiplied by a set of weights and then puts this output (y_{pred}) through a nonlinearity. In our case, the nonlinearity outputs the sign of the perceptron’s output and we use that sign to predict the label (-1,+1 for left and right choices). The objective of learning in a perceptron is to change the weights as to minimize the mean-squared error of the prediction.

Simple forms of online learning in a perceptron can be biologically realistic if they only require global error signals in addition to information available locally at a synapse. We investigate three such versions here. First, we consider a “supervised Hebbian” learner, which changes the weight w_k by $\Delta w_k = y_{\text{label}} \cdot x_k$, with x_k the response of neuron k (Figure 4A). We also consider a gradient descent learner, which

changes the weights by $\Delta w_k = (y_{\text{label}} - y_{\text{pred}}) \cdot x_k$, and a restricted form of gradient descent which uses the reinforcement feedback (correct/wrong) in place of the full prediction error: $\Delta w_k = \text{sign}(y_{\text{label}} - y_{\text{pred}}) \cdot x_k$. We call this last one “reinforcement descent”.

To test these online learning strategies, we designed two tasks, one easy and one hard. In the easy task, the learning agents used the neural responses to our first stimulus set with 10 trials/deg, restricted to angles between -30 and -5 degrees for the “left” choice, and 5 to 30 for the “right” choice (Figure 4B). All three perceptron learners performed this task perfectly, using a small number of training trials (Figure 4B). In the hard task, the learners had to discriminate between positive and negative stimulus angles of up to 2°, using the neural data recorded at 1,000 trials/deg (Figure 4C). The supervised

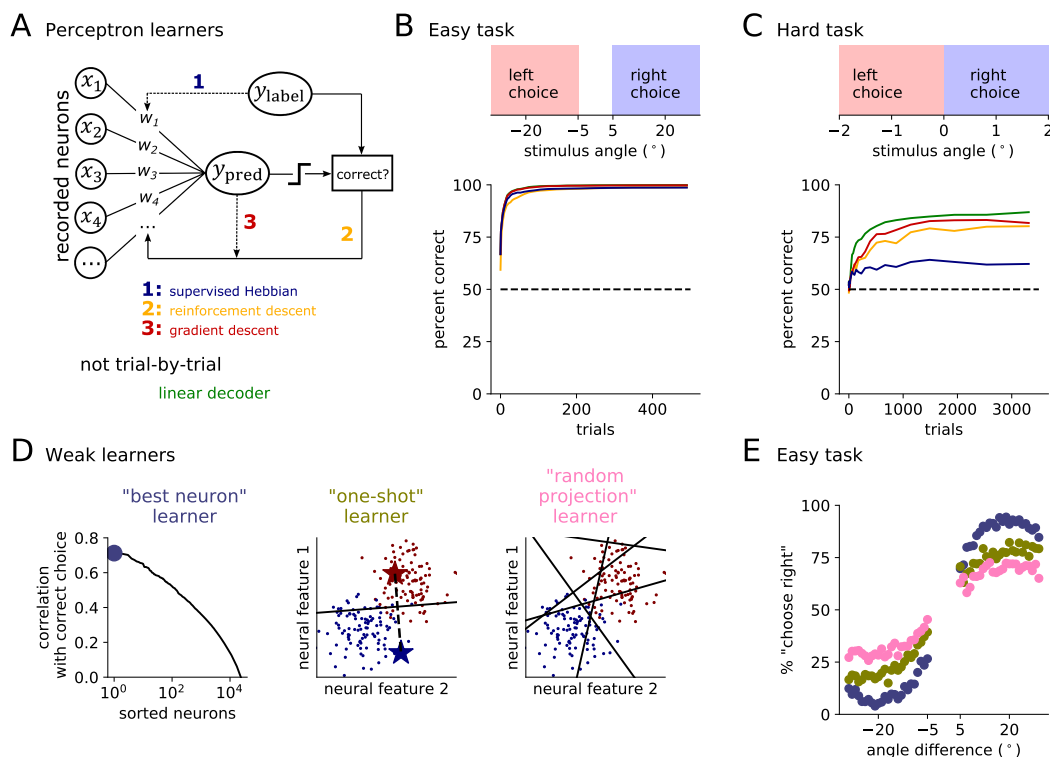


Figure 4: Online learning of discrimination tasks. (A) Trial-by-trial learning for a perceptron can use either 1) only the label information, 2) only the trial reinforcement, or 3) both the trial reinforcement and the continuous valued prediction. Finally, non-incremental learning can use all trials simultaneously, such the linear decoder previously used in Figure 3. (B) “Easy task” configuration and decoder learning strategies, averaged over the neural datasets which provide the inputs x . (C) Same as B for the “hard task”. (D) Alternative learning strategies that are plausible but weak. The “best neuron” learner chooses the neuron most correlated with the correct choice on the training set. The “one shot” learner uses just one training example of each class (left and right) to find a linear discriminant. The “random projection” learner chooses the best performing random projection out of 100. (E) Performance of the weak learners on the easy task.

Hebbian learner was unable to perform well in this task with an asymptotic performance of 66%. However, the learners based on gradient descent and reinforcement descent performed relatively well at 83% and 77% compared to 86% for the optimal linear decoder, which had instantaneous access to all trials in its history (Figure 4C). Therefore, the perceptrons acquired task information nearly as efficiently as the optimal linear decoder did.

We conclude that these simple online decoders can learn the orientation discrimination task from neural data in a sample-efficient way. Therefore, the trial-by-trial limitation of animal perceptual learning experiments does not explain the discrepancy between neural and behavioral discrimination thresholds. It follows that animals use highly suboptimal learning strategies in psychophysical experiments, perhaps because those strategies are favorable in ecological contexts. We end this study by proposing examples of weak but potentially relevant decoders, and leave it to future work to test if any of these weak learners provide insight into animal behavior. The decoders we propose (Figure 4D) are: 1) the “best neuron” learner, which finds the neuron most correlated with the correct choice on the training set and then uses this

neuron to make predictions on the test set; 2) the “one-shot” learner, which uses neural responses from only one trial of each stimulus category and builds a linear decoder along the discriminant direction of these two trials; and 3) the “random projection” learner, which tests 100 different random projections in parallel and retains the one that maximizes training set performance. These decoders had test set performance on the easy task in the range of 65-90% (Figure 4D), which is in line with mouse behavioral performance on similar tasks [37, 38, 52].

DISCUSSION

What sets the limits of sensory discrimination? Our analyses imply it is *not* neural noise. We have shown this by recording from very large populations of neurons in visual cortex and computing neural discrimination thresholds as low as 0.3° in mice, a species thought to have poor visual acuity. Simple linear decoders were sufficient to achieve this performance, although it is possible that nonlinear decoders may achieve even higher performance ([53]; but see Figure S10). The neural discrimination thresholds were $\sim 100\times$ lower than behavioral discrimination thresholds, which we estimated at $>25^\circ$

based on previous studies [38]. We were not able to explain this discrepancy by varying six different stimulus properties, by splitting trials according to behavioral states, or by learning the decoder in a biologically plausible way.

We suggest two potential explanations for the discrepancy between behavior and neurons. First, it might be that downstream decoders in the mouse brain are computationally limited, and cannot find the same discrimination directions in neural space that are found by an ideal external observer. Second, it might be that our methods for probing discrimination in mice are limited and better task designs may result in better training and better discrimination. To better train mice, more naturalistic tasks might be necessary, such as immersive virtual reality navigation [52, 54, 55]. Alternatively, it may be that mice do have latent task knowledge even if their behavior is poor [56], in which case future task improvements should target the expression of this latent knowledge.

Acknowledgments

We thank Dan Flickinger for assistance with the two-photon microscope, Salvatore DiLisio and Sarah Lindo for surgeries, and James Fitzgerald for useful discussions. This research was funded by the Howard Hughes Medical Institute.

All of the processed neural activity is available on figshare [41] (https://figshare.com/articles/Recordings_of_20_000_neurons_from_V1_in_response_to_oriented_stimuli/8279387). The code is available on github (<https://github.com/MouseLand/stringer-et-al-2019a>).

References

- [1] David J Tolhurst, J Anthony Movshon, and Andrew F Dean. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision research*, 23(8):775–785, 1983.
- [2] William R Softky and Christof Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random epsps. *Journal of Neuroscience*, 13(1):334–350, 1993.
- [3] Ehud Zohary, Michael N Shadlen, and William T Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370(6485):140, 1994.
- [4] Zachary F Mainen and Terrence J Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268(5216):1503–1506, 1995.
- [5] Amos Arieli, Alexander Sterkin, Amiram Grinvald, and AD Aertsen. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science*, 273(5283):1868–1871, 1996.
- [6] M N Shadlen, K H Britten, W T Newsome, and J a Movshon. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *Journal of Neuroscience*, 16(4):1486–1510, 1996. ISSN 0270-6474, 1529-2401.
- [7] Rob R de Ruyter van Steveninck, Geoffrey D Lewen, Steven P Strong, Roland Koberle, and William Bialek. Reproducibility and variability in neural spike trains. *Science*, 275(5307):1805–1808, 1997.
- [8] Michael R Deweese and Anthony M Zador. Shared and private variability in the auditory cortex. *Journal of neurophysiology*, 92(3):1840–1855, 2004.
- [9] A Aldo Faisal, Luc PJ Selen, and Daniel M Wolpert. Noise in the nervous system. *Nature reviews neuroscience*, 9(4):292, 2008.
- [10] Bruno B Averbeck, Peter E Latham, and Alexandre Pouget. Neural correlations, population coding and computation. *Nature reviews neuroscience*, 7(5):358, 2006.
- [11] Mark M Churchland, Byron M Yu, John P Cunningham, Leo P Sugrue, Marlene R Cohen, Greg S Corrado, William T Newsome, Andrew M Clark, Paymon Hosseini, Benjamin B Scott, David C Bradley, Matthew a Smith, Adam Kohn, J Anthony Movshon, Katherine M Armstrong, Tirin Moore, Steve W Chang, Lawrence H Snyder, Stephen G Lisberger, Nicholas J Priebe, Ian M Finn, David Ferster, Stephen I Ryu, Gopal Santhanam, Maneesh Sahani, and Krishna V Shenoy. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3):369–378, 2010. ISSN 1097-6256. doi: 10.1038/nn.2501. URL <http://dx.doi.org/10.1038/nn.2501>.
- [12] Marlene R Cohen and Adam Kohn. Measuring and interpreting neuronal correlations. *Nature Neuroscience*, 14(7):811–819, 2011. ISSN 1097-6256. doi: 10.1038/nn.2842.
- [13] Jude F Mitchell, Kristy A Sundberg, and John H Reynolds. Spatial attention decorrelates intrinsic activity fluctuations in macaque area v4. *Neuron*, 63(6):879–888, 2009.
- [14] Marlene R Cohen and John H R Maunsell. Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, 12(12):1594–1600, 2009. ISSN 1097-6256. doi: 10.1038/nn.2439.
- [15] Yong Gu, Sheng Liu, Christopher R Fetsch, Yun Yang, Sam Fok, Adhira Sunkara, Gregory C DeAngelis, and Dora E Angelaki. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron*, 71(4):750–761, 2011.
- [16] Amy M Ni, Douglas A Ruff, Joshua J Alberts, Jen Symmonds, and Marlene R Cohen. Learning and attention reveal a general relationship between population activity and behavior. *Science*, 359(6374):463–465, 2018. URL <https://science.sciencemag.org/content/359/6374/463>.

- [17] Horace B Barlow et al. Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234, 1961.
- [18] Wilson S Geisler. Contributions of ideal observer theory to vision research. *Vision research*, 51(7):771–781, 2011.
- [19] Wei Ji Ma, Jeffrey M Beck, Peter E Latham, and Alexandre Pouget. Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11):1432, 2006.
- [20] Kenji Doya, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao. *Bayesian brain: Probabilistic approaches to neural coding*. MIT press, 2007.
- [21] David C Knill and Alexandre Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12):712–719, 2004.
- [22] Lars Buesing, Johannes Bill, Bernhard Nessler, and Wolfgang Maass. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS computational biology*, 7(11):e1002211, 2011.
- [23] Ralf M Haefner, Pietro Berkes, and József Fiser. Perceptual decision-making as probabilistic inference by neural sampling. *Neuron*, 90(3):649–660, 2016.
- [24] Larry F Abbott and Peter Dayan. The effect of correlated variability on the accuracy of a population code. *Neural computation*, 11(1):91–101, 1999.
- [25] Alexander S Ecker, Philipp Berens, Andreas S Tolias, and Matthias Bethge. The effect of noise correlations in populations of diversely tuned neurons. *Journal of Neuroscience*, 31(40):14272–14283, 2011.
- [26] Haim Sompolinsky, Hyoungsoo Yoon, Kukjin Kang, and Maoz Shamir. Population coding in neuronal systems with correlated noise. *Physical Review E*, 64(5):051904, 2001.
- [27] Rubén Moreno-Bote, Jeffrey Beck, Ingmar Kanitscheider, Xaq Pitkow, Peter Latham, and Alexandre Pouget. Information-limiting correlations. *Nature Neuroscience*, 17(10):1410–1417, 2014.
- [28] Jorrit S Montijn, Guido T Meijer, Carien S Lansink, and Cyriel MA Pennartz. Population-level neural codes are robust to single-neuron variability from a multidimensional coding perspective. *Cell reports*, 16(9):2486–2498, 2016. URL <https://doi.org/10.1016/j.celrep.2016.07.065>.
- [29] Douglas A. Ruff and Marlene R. Cohen. Simultaneous multi-area recordings suggest a novel hypothesis about how attention improves performance. *bioRxiv*, 2019. doi: 10.1101/372888. URL <https://www.biorxiv.org/content/early/2019/05/09/372888>.
- [30] Ranulfo Romo, Adrián Hernández, Antonio Zainos, and Emilio Salinas. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron*, 38(4):649–657, 2003.
- [31] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):255–255, 2019. ISSN 0036-8075. doi: 10.1126/science.aav7893. URL <https://science.sciencemag.org/content/364/6437/255>.
- [32] Rufin Vogels and GA Orban. How well do response changes of striate neurons signal differences in orientation: a study in the discriminating monkey. *Journal of Neuroscience*, 10(11):3543–3558, 1990.
- [33] Arnulf BA Graf, Adam Kohn, Mehrdad Jazayeri, and J Anthony Movshon. Decoding the activity of neuronal populations in macaque primary visual cortex. *Nature neuroscience*, 14(2):239, 2011.
- [34] Philipp Berens, Alexander S Ecker, R James Cotton, Wei Ji Ma, Matthias Bethge, and Andreas S Tolias. A fast and simple population code for orientation in primate v1. *Journal of Neuroscience*, 32(31):10618–10626, 2012.
- [35] William T Newsome, Kenneth H Britten, and J Anthony Movshon. Neuronal correlates of a perceptual decision. *Nature*, 341(6237):52, 1989.
- [36] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris. High-dimensional geometry of population responses in visual cortex. *bioRxiv*, 2018. doi: 10.1101/374090. URL <https://www.biorxiv.org/content/early/2018/08/13/374090>.
- [37] Seung-Hee Lee, Alex C Kwan, Siyu Zhang, Victoria Phoumthipphavong, John G Flannery, Sotiris C Masmanidis, Hiroki Taniguchi, Z Josh Huang, Feng Zhang, Edward S Boyden, Karl Deisseroth, and Yang Dan. Activation of specific interneurons improves v1 feature selectivity and visual perception. *Nature*, 488(7411):379–383, 2012.
- [38] Mohammad Abdolrahmani, Dmitry R. Lyamzin, Ryo Aoki, and Andrea Benucci. Cognitive modulation of interacting corollary discharges in the visual cortex. *bioRxiv*, 2019. doi: 10.1101/615229. URL <https://www.biorxiv.org/content/early/2019/04/23/615229>.
- [39] Marius Pachitariu, Carsen Stringer, Sylvania Schrder, Mario Dipoppa, L Federico Rossi, Matteo Carandini, and Kenneth D Harris. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv*, 2016. URL <https://www.biorxiv.org/content/early/2017/07/20/061507>.
- [40] Marius Pachitariu, Carsen Stringer, and Kenneth D Harris. Robustness of spike deconvolution for neuronal calcium imaging. *Journal of Neuroscience*, 38(37):7976–7985, 2018. URL <http://www.jneurosci.org/content/early/2018/08/06/JNEUROSCI.3339-17.2018>.

- [41] Marius Pachitariu, Michalis Michaelos, and Carsen Stringer. Recordings of 20,000 neurons from V1 in response to oriented stimuli. *figshare*, 6 2019. doi: 10.25378/janelia.8279387.v1. URL https://janelia.figshare.com/articles/Recordings_of_20_000_neurons_from_V1_in_response_to_oriented_stimuli/8279387.
- [42] Saskia E J de Vries, Jerome Lecoq, Michael A Buice, Peter A Groblewski, Gabriel K Ocker, Michael Oliver, David Feng, Nicholas Cain, Peter Ledochowitsch, Daniel Millman, Kate Roll, Marina Garrett, Tom Keenan, Leonard Kuan, Stefan Mihalas, Shawn Olsen, Carol Thompson, Wayne Wakeman, Jack Waters, Derric Williams, Chris Barber, Nathan Berbesque, Brandon Blanchard, Nicholas Bowles, Shiella Caldejon, Linzy Casal, Andrew Cho, Sissy Cross, Chinh Dang, Tim Dolbear, Melise Edwards, John Galbraith, Nathalie Gaudreault, Fiona Griffin, Perry Hargrave, Robert Howard, Lawrence Huang, Sean Jewell, Nika Keller, Ulf Knoblich, Josh Larkin, Rachael Larsen, Chris Lau, Eric Lee, Felix Lee, Arielle Leon, Lu Li, Fuhui Long, Jennifer Luviano, Kyla Mace, Thuyanh Nguyen, Jed Perkins, Miranda Robertson, Sam Seid, Eric Shea-Brown, Jianghong Shi, Nathan Sjoquist, Cliff Slaughterbeck, David Sullivan, Ryan Valenza, Casey White, Ali Williford, Daniela Witten, Jun Zhuang, Hongkui Zeng, Colin Farrell, Lydia Ng, Amy Bernard, John W Phillips, R Clay Reid, and Christof Koch. A large-scale, standardized physiological survey reveals higher order coding throughout the mouse visual cortex. *bioRxiv*, 2018. doi: 10.1101/359513. URL <https://www.biorxiv.org/content/early/2018/06/29/359513>.
- [43] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.
- [44] Jason D Rennie, Lawrence Shih, Jaime Teevan, and David R Karger. Tackling the poor assumptions of naive bayes text classifiers. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pages 616–623, 2003.
- [45] Thomas J Davidson, Fabian Kloosterman, and Matthew A Wilson. Hippocampal replay of extended experience. *Neuron*, 63(4):497–507, 2009.
- [46] Gabrielle Girardeau, Karim Benchenane, Sidney I Wiener, György Buzsáki, and Michaël B Zugaro. Selective suppression of hippocampal ripples impairs spatial memory. *Nature neuroscience*, 12(10):1222, 2009.
- [47] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [48] M Devos and Guy A Orban. Modeling orientation discrimination at multiple reference orientations with a neural network. *Neural Computation*, 2(2):152–161, 1990.
- [49] H S Seung and H Sompolinsky. Simple models for reading neuronal population codes. *Proceedings of the National Academy of Sciences*, 90(22):10749–10753, 1993. ISSN 0027-8424. doi: 10.1073/pnas.90.22.10749.
- [50] Peter Dayan and Laurence F Abbott. Theoretical neuroscience: computational and mathematical modeling of neural systems. 2001.
- [51] Yoram Burak, SueYeon Chung, and Haim Sompolinsky. Quadratic networks for invariant perceptual discrimination. In *Computational and Systems Neuroscience (Cosyne)*, 2012.
- [52] Jasper Poort, Adil G Khan, Marius Pachitariu, Abdellatif Nemri, Ivana Orsolich, Julija Krupic, Marius Bauza, Maneesh Sahani, Georg B Keller, Thomas D Mrsic-Flogel, and Sonja Hofer. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron*, 86(6):1478–1490, 2015.
- [53] Maoz Shamir and Haim Sompolinsky. Nonlinear population codes. *Neural Computation*, 16(6):1105–1136, 2004. doi: 10.1162/089976604773717559. URL <https://doi.org/10.1162/089976604773717559>.
- [54] Laura N Driscoll, Noah L Pettit, Matthias Minderer, Selmaan N Chettih, and Christopher D Harvey. Dynamic reorganization of neuronal activity patterns in parietal cortex. *Cell*, 170(5):986–999, 2017.
- [55] Michael Krumin, Julie J Lee, Kenneth D Harris, and Matteo Carandini. Decision and navigation in mouse parietal cortex. *ELife*, 7:e42583, 2018. URL <https://elifesciences.org/articles/42583>.
- [56] Kishore V Kuchibhotla, Tom Hindmarsh Sten, Eleni S Papadoyannis, Sarah Elnozahy, Kelly A Fogelson, Rupesh Kumar, Yves Boubenec, Peter C Holland, Srdjan Ostojic, and Robert C Froemke. Dissociating task acquisition from expression during learning reveals latent knowledge. *Nature communications*, 10(1):2151, 2019.
- [57] Nicholas James Sofroniew, Daniel Flickinger, Jonathan King, and Karel Svoboda. A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *ELife*, 5:e14472, 2016. URL <https://elifesciences.org/articles/14472>.
- [58] Thomas A Polgruto, Bernardo L Sabatini, and Karel Svoboda. Scanimage: flexible software for operating laser scanning microscopes. *Biomedical engineering online*, 2(1):13, 2003.
- [59] Johannes D Seelig, M Eugenia Chiappe, Gus K Lott, Anirban Dutta, Jason E Osborne, Michael B Reiser, and Vivek Jayaraman. Two-photon calcium imaging from head-fixed drosophila during optomotor walking behavior. *Nature methods*, 7(7):535, 2010.

- [60] Mario Kleiner, David Brainard, Denis Pelli, Allen Ingling, Richard Murray, Christopher Broussard, et al. What's new in psychtoolbox-3. *Perception*, 36(14):1, 2007.
- [61] Stefan Van Der Walt, S Chris Colbert, and Gael Varoquaux. The numpy array: a structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22, 2011.
- [62] Eric Jones, Travis Oliphant, Pearu Peterson, et al. SciPy: Open source scientific tools for Python, 2001. URL <http://www.scipy.org/>.
- [63] Siu Kwan Lam, Antoine Pitrou, and Stanley Seibert. Numba: A llvm-based python jit compiler. In *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, page 7. ACM, 2015.
- [64] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Goullart, and Tony Yu. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014.
- [65] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct): 2825–2830, 2011.
- [66] John D Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9(3):90, 2007.
- [67] Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B Hamrick, Jason Grout, Sylvain Corlay, et al. Jupyter notebooks—a publishing format for reproducible computational workflows. In *ELPUB*, pages 87–90, 2016.
- [68] Johannes Friedrich, Pengcheng Zhou, and Liam Paninski. Fast online deconvolution of calcium imaging data. *PLOS Computational Biology*, 13(3):e1005423, 2017. URL <https://doi.org/10.1371/journal.pcbi.1005423>.
- [69] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.

METHODS

All experimental procedures were conducted according to IACUC.

Animals and surgery

We performed experiments on three mice bred to express GCaMP6s in excitatory neurons: TetO-GCaMP6s x Emx1-IRES-Cre mice (available as JAX 024742 and JAX 005628).

We also performed experiments on three wild-type C57 mice in 17 recordings. In these mice, AAV-GCaMP6s-P2A-nls-dTomato (Addgene plasmid #51084) was expressed virally through injections into visual cortex. These mice were male and female, and ranged from 2 to 8 months of age.

Surgical methods were similar to those described elsewhere [31]. Briefly, surgeries were performed in adult mice (P35-P125) under isoflurane anesthesia (5% for induction, 1–2% during the surgery) in a stereotaxic frame. Before surgery, buprenorphine was administered as a systemic analgesic and lidocaine was administered locally at the surgery site. During the surgery we implanted a head-plate for later head-fixation, and made a craniotomy of 4 mm in diameter with a cranial window implant for optical access. We targeted virus injections (50–200 nl, 1–3 x 10¹² GC/ml) to monocular V1 (2.1–3.3 mm laterally and 3.5–4.0 mm posteriorly from Bregma). To obtain large fields of view for imaging, we typically performed several injections at nearby locations, at multiple depths (~500 μ m and ~200 μ m).

Data acquisition

We used a custom-built 2-photon mesoscope [57] to record neural activity, and ScanImage [58] for data acquisition, obtaining $18,496 \pm 3,441$ (s.d., n=32) neurons in the recordings. Multi-plane acquisition was controlled by a resonance mirror, with planes spaced 25 μ m apart in depth. 10–17 planes were acquired sequentially, scanning the entire stack repeatedly on average at 3 Hz. We synchronized stimulus presentation to the beginning of each frame for the first plane, and computed stimulus responses from the first three frames acquired after stimulus onset for each plane. We used a custom online Z-correction module (now in ScanImage), to correct for Z and XY drift online during the recording.

The mice were free to run on an air-floating ball. For all static image presentations an LED tablet screen at a 45° from the left eye (we recorded in right visual cortex). For drifting grating image presentations a custom circular screen made of LED arrays was placed around the head of the mouse [59]. We also repeated the static grating experiments with this screen and obtained comparable decoding errors.

For each mouse, recordings were made over multiple days, always returning to the same field of view. The field of view was selected on the first recording day such that large numbers of neurons could be observed, with clear calcium transients and a retinotopic location (identified by neuropil fluorescence) that was localized on the screen.

Visual stimuli

We showed various gratings and localized images. To present stimuli, we used PsychToolbox-3 in MATLAB [60]. The stimuli were presented for 750 ms (unless otherwise stated), alternating with a gray-screen inter-stimulus interval lasting on average 650 ms. After every 150 stimuli, the screen was left blank (gray screen) for 32 seconds. The activity during these non-

stimulus periods was used to project out spontaneous dimensions from the neuronal population responses (see below).

All gratings were square and had a spatial frequency of 0.05 cycles / degree; drifting gratings had a temporal frequency of 2 Hz. We showed random orientations/directions of these stimuli on each trial.

The local stimulus was restricted to 30° of visual space and used a lower frequency grating. The complex "minimouse" stimulus was also 30° of visual space and was rotated around its center. Outside of these stimuli, the screen was gray.

Data processing

The calcium imaging processing pipeline and the subsequent analyses use numpy, scipy, numba, scikit-image, and scikit-learn [61–65]. The figures in the paper were made using matplotlib in jupyter [66, 67].

Calcium imaging data was processed using the Suite2p toolbox [39], available at www.github.com/MouseLand/suite2p. Suite2p performs motion correction, ROI detection, cell classification, neuropil correction, and spike deconvolution as described elsewhere [31]. For non-negative deconvolution, we used a timescale of decay of 1.25 seconds [40, 68].

Stimulus responses

We defined the stimulus response as the summed activity of the first three bins (~1 second) following stimulus onset. We split the trials 75/25 into training and testing sets, with every 4th trial assigned to the test set.

Splitting cells into two populations

When looking at correlated decoding errors, we split the neurons into two populations. We first divided the XY plane into 8 non-overlapping strips of width 150 μm, and assigned the neurons in the even strips to one group, and the neurons in the odd strips to the other group, regardless of the neuron's depth. Thus, there did not exist neuron pairs in the two sets that had the same XY position but a different depth. This special division was performed to avoid contamination artifacts between overlapping cells or between consecutive planes.

Tuning curves and SNR

We fit the training trials with $n_{\text{basis}} = 10$ cosine and sine basis functions, where θ is the angle of the stimulus shown in radians:

$$B = \begin{pmatrix} \cos(0) \\ \sin(\theta) \\ \cos(\theta) \\ \vdots \\ \sin(n_{\text{basis}} \cdot \theta) \\ \cos(n_{\text{basis}} \cdot \theta) \end{pmatrix}$$

We performed linear regression from B to the neural responses and used the fitted function $f_n(\theta)$ as the tuning curve. To compute the signal-to-noise ratio, we defined the signal as the variance of the tuning curve and the noise as the variance

of the residual noise after subtracting out the tuning curve value.

Independent decoder

For the independent decoder, we built generative models of the neural data by modelling the mean and standard deviation of each neuron. The mean was obtained as a function of stimulus angle using the tuning curve fits $f_n(\theta)$ above. The standard deviation $\sigma_n(\theta)$ was fit similarly, after subtracting the mean predictions on the training set, by squaring the residuals and fitting them in the same set of basis functions. With the mean and standard deviation defined for each neuron and each stimulus, we computed the probability that a novel neural pattern $R(n)$ was produced by a putative stimulus orientation θ' :

$$P(\theta' | R(n)) \sim \mathcal{N}(R(n) | f_n(\theta'), \sigma_n(\theta')).$$

These probabilities were evaluated in log-space for a discrete set of θ' ($n=48$ orientations) and summed across neurons. To decode orientations more finely, we upsampled the log-probability curves by a factor of 100 using kriging interpolation. The stimulus angle corresponding to the peak of the upsampled curve was used as the decoded orientation.

Linear decoder

To decode circular angles using locally linear decoders, we regressed the neural activity onto "super"-neurons (see Figure 2F). These super-neurons were von Mises tuning curves ($n = 48$) with peaks equally spaced along 360° and with $\sigma = 0.1$:

$$v_k = e^{(\cos(\theta - \theta_k) - 1) / \sigma}$$

$$\theta_k = 360(k - 1) / n$$

We fit the transformation from neurons to super-neurons on training trials and then predicted the super-neuron responses on test trials. As for the independent decoder, we upsampled the super-neuron responses from 48 to 4,800 to decode stimulus angle more finely.

Removal of ongoing activity dimensions

Approximately half the variance of visual cortical population activity represents behavior-related fluctuations [31]. In Figure S5 we projected out the ongoing activity dimensions to show that they had no influence on sensory coding. To do this we computed the top 32 principal components of z-scored and binned (3 frames = ~ 1 second) ongoing activity. The spontaneous, ongoing neural activity was recorded during gray screen presentations, and these were presented after every 150 stimulus presentations for 32 seconds each time. To remove these dimensions from stimulus responses, the stimulus-driven activity was first z-scored (using the mean and variance of each neuron computed from ongoing activity), then the projection onto the 32 top spontaneous dimensions was subtracted.

Asymptotics

To fit the asymptotic error, we modeled the scaling of the median error with the parametrization $\alpha + \frac{\beta}{\sqrt{N}}$, where N is the number of neurons or the number of stimuli (Figure S6). The scaling of $1/\sqrt{N}$ was chosen because it corresponds to the decay of the standard deviation of an average of independent random variables with the same variance. We fit α and β to the last 12 points of each curve in Figure S6 using linear regression.

Linear decoder for discrimination

To perform the discrimination task with a linear decoder, we took a similar approach to the linear decoder described above. Instead of having super-neurons with von Mises tuning, the super-neurons for this task ($n = 32$) are differences of von Mises functions centered above and below the discrimination boundary:

$$d_k = e^{(\cos(\theta - \theta_k - dt) - 1)/\sigma} - e^{(\cos(\theta - \theta_k + dt) - 1)/\sigma}$$
$$\theta_k = 360(k - 1)/n$$
$$dt = 15^\circ$$

Neural networks for decoding

We used PyTorch [69] to train a neural network to perform the discrimination task in Figure 3 on the 10 trials/degree recordings. This network consisted of two rectified-linear layers and an output sigmoid. The input layer consisted of 256 units using the top principal components from the data (this reduced overfitting). The two hidden layers consisted of 100 units each. We trained the network for 50,000 iterations using stochastic gradient descent with momentum of 0.9 and a learning rate of $1e-3$. The cost was a binary cross-entropy loss. We averaged over 5 random initializations of the neural network for each recording.

Random forests for decoding

We used scikit-learn [65] to train a random forest ensemble classifier to perform the discrimination task using the neural activity from the 10 trials/degree recordings. We used 1000 trees and averaged over 5 random initializations of the classifier for each recording.

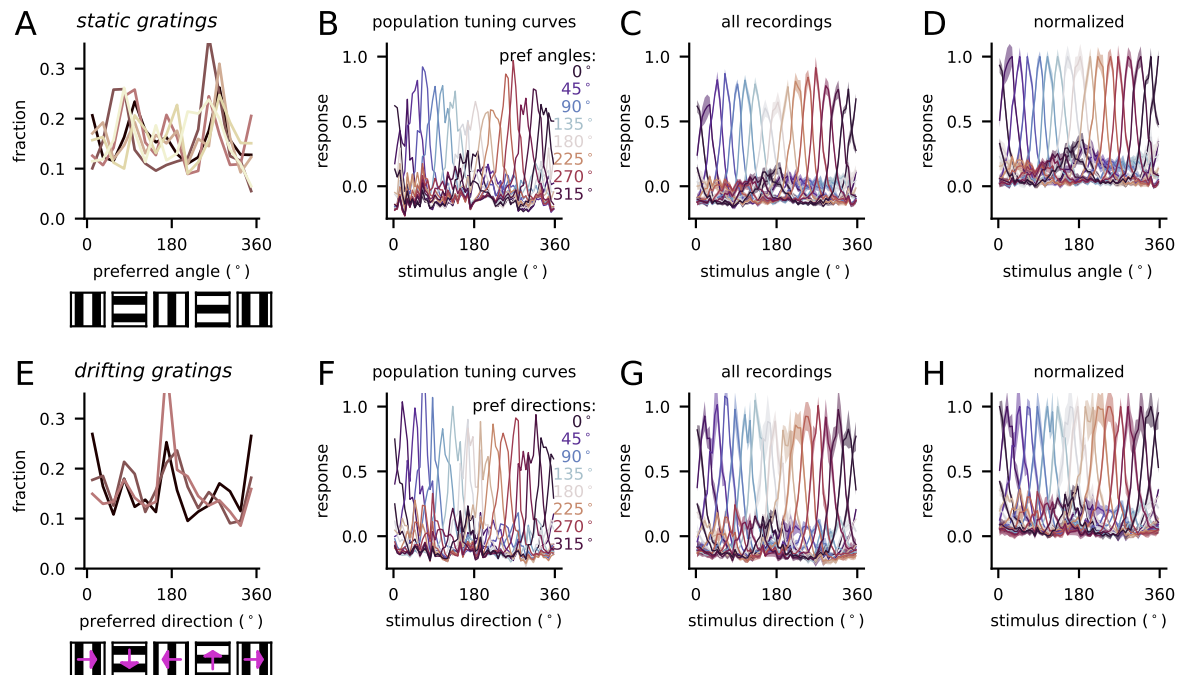
Perceptron learners

The perceptrons used the learning rule

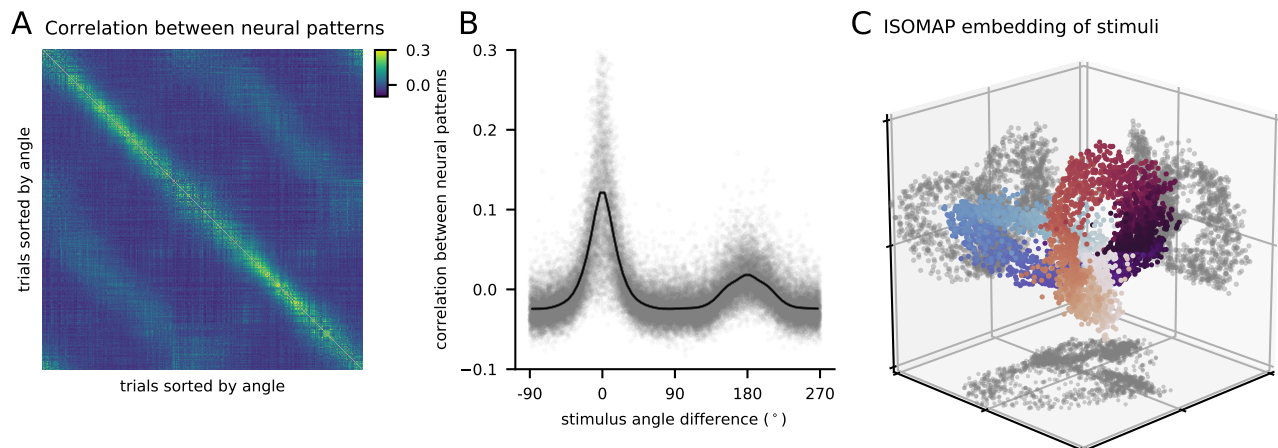
$$w_k = w_k + \eta \cdot \text{err} \cdot x_k,$$
$$\text{err} \in \{y_{\text{label}}, y_{\text{pred}} - y_{\text{label}}, \text{sign}(y_{\text{pred}} - y_{\text{label}})\}$$

where η is the learning rate and the three cases for “err” correspond to the supervised hebbian, the gradient descent and the reinforcement descent decoders. The gradient descent decoder has been derived from the error cost function:

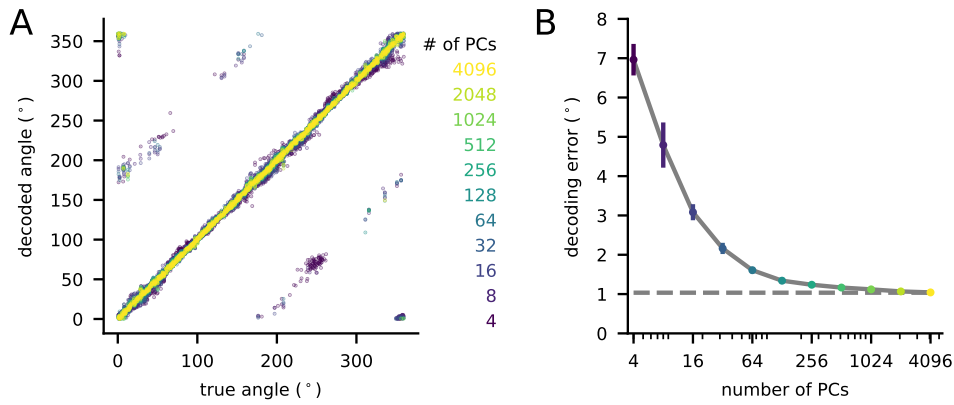
$$\text{Cost}(w) = \sum_{\text{trials}} (\langle \mathbf{w}, \mathbf{x} \rangle - y_{\text{label}})^2$$



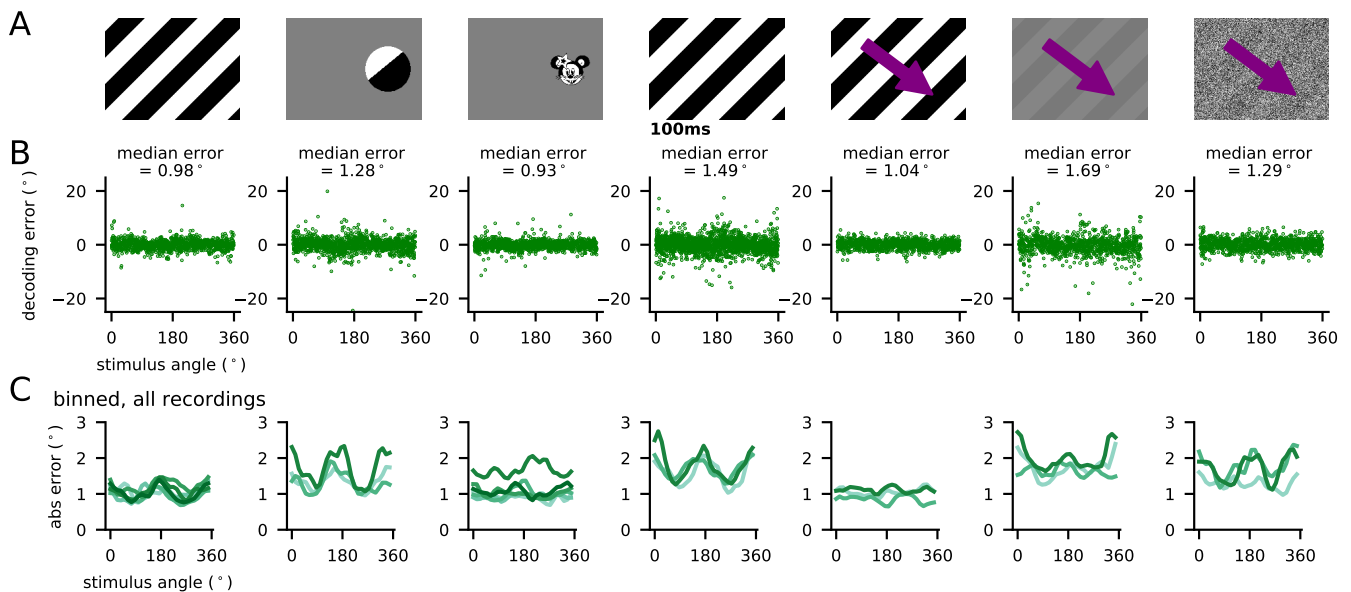
S1: Population tuning curves. (A) Cosine and sine basis functions were fit to all neurons' stimulus responses (train trials) in all recordings (same as in Figure 1F). The preferred angles of those basis function are reported here as a distribution. Each line represents a different recording. (B) On test trials, the responses of neurons with similar preferred angles (determined on train trials) were averaged. (C) Same as B, averaged over all recordings. (D) Same as C, but normalized between 0 and 1. The half-width half-max of these tuning curves are 14.1° . (E-F) Same as A-D for drifting grating responses. The half-width half-max of the tuning curves in H are 15.1° .



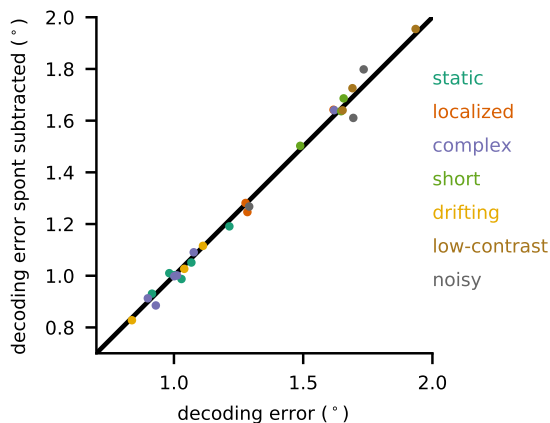
S2: Manifold of stimulus responses. (A) Correlation between neural responses on all pairs of trials of a recording. (B) Same data as A, with correlation plotted as a function of stimulus angle difference. The black line shows the binned and smoothed average. (C) The $\sim 20,000$ -dimensional vectors were embedded into three dimensions using ISOMAP. Points are colored by the angle of the presented stimulus. Diametrically opposite points on the manifold come in closer proximity to each other than expected from a pure circle, due to the correlation of the neural patterns at π angle difference.



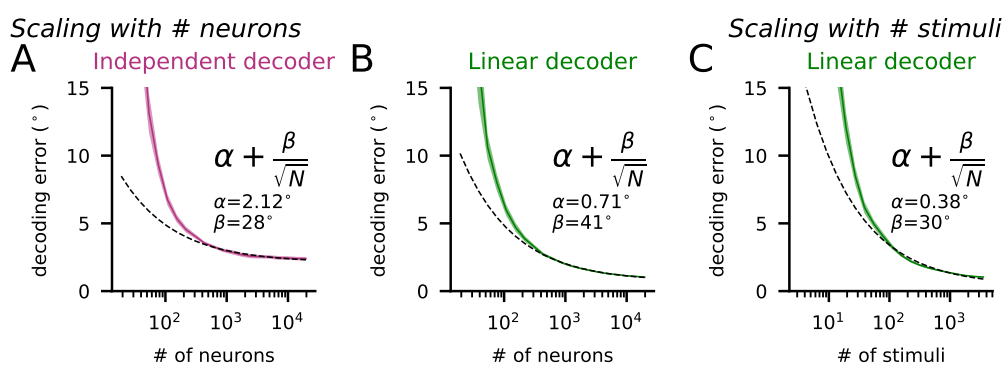
S3: Linear decoding using principal components (A) True angle versus decoded angle using the same linear decoding strategy used for all neurons (Figure 2F), for an example recording. (B) Median decoding error as a function of the number of principal components kept, averaged over recordings. Error bars are standard error. Dashed line is decoding error using full dataset.




S4: Decoding error as a function of stimulus angle. (A) All the stimulus sets. (B) Decoding error vs stimulus angle for each trial of an example recording of each stimulus set. (C) Plot of B after taking the absolute value and binning. Each line is a different recording.



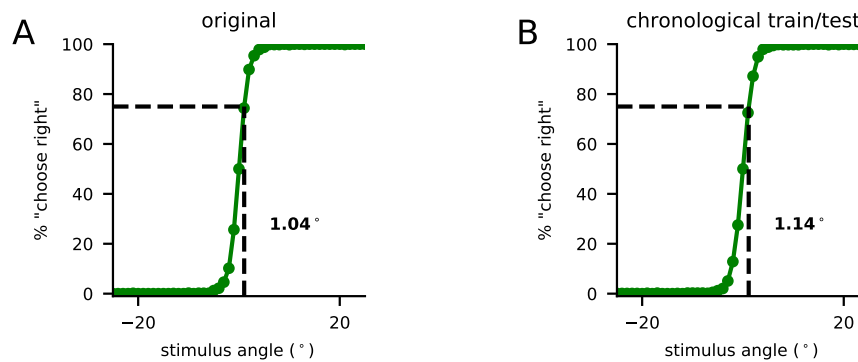
S5: Principal components of spontaneous, ongoing activity do not influence decoding accuracy. 32 principal components of spontaneous, ongoing activity were subtracted from the stimulus responses, and the linear decoder was trained on these responses. The median decoding errors of the subtracted responses are plotted vs the original.




S6: Asymptotics. (A) Independent decoder fit to random subsets of neurons, averaged over all recordings. Error bars are standard error. (B) Same as (A) for the linear decoder. (C) Linear decoder fit using random subsets of stimuli and all neurons.



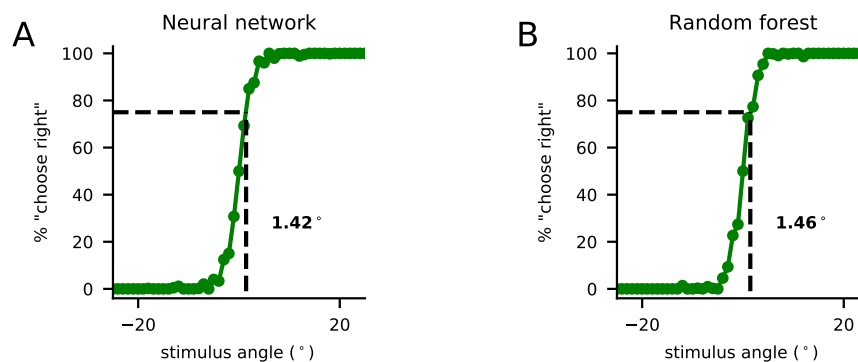
S7: Control (laser off). To ensure that the visual stimulus screen did not contaminate the fluorescence signals collected by the photo-multiplier tube (PMT), we performed recordings in which the shutter on the 2p laser was kept closed, but all other conditions remained the same. With the laser off, the PMT signals were near-zero and reflected mainly auto-fluorescence, photon noise and 60Hz signals, most probably from the monitors. (A) Example frame from one of the recordings used in the paper. (B) Example frame of control recording with laser OFF. (C) Discrimination of stimuli in laser OFF recordings ($n=3$ mice, >4000 total trials/mouse). The discrimination performance appeared to be at chance.



S8: Chronological splitting of train/test trials. (A) Train and test trials interleaved (same as Figure3B). (B) Chronological split across the 120-180 minutes of recording: training trials were first 75% of stimuli presented and test trials were last 25% of stimuli presented. The discrimination threshold was only modestly higher.



S9: L2/3 vs L4. (A) Discrimination of static gratings (10 trials/deg) using only neurons in at depths ~ 125 - $225 \mu\text{m}$. (B) Same as (A) with neurons (~ 375 - $4755 \mu\text{m}$ deep).



S10: Decoding with multilayer neural networks and random forests Same decoding task and data as Figure3B. (A) Two-layer neural network. (B) Random forest classifier.