

Hue tuning curves in V4 change with visual context

Ari S. Benjamin^{1*}, Pavan Ramkumar^{2,3}, Hugo Fernandes², Matthew Smith^{4,5}, Konrad P. Kording^{1,2}

1. Department of Bioengineering, University of Pennsylvania, Philadelphia, PA

2. Department of Physical Medicine and Rehabilitation, Northwestern University and Shirley Ryan Ability Lab,
Chicago, IL

3. System1 Biosciences, San Francisco, CA

4. Department of Biomedical Engineering, Carnegie Mellon University, Pittsburgh, PA

5. Carnegie Mellon Neuroscience Institute, Pittsburgh, PA

6. *Correspondence: aarri@seas.upenn.edu, Richards 401A 3700 Hamilton Walk. Philadelphia PA, 19104

11 **Abstract**

12 Neurons are often probed by presenting a set of stimuli that vary along one dimension (e.g. color) and
13 quantifying how this stimulus property affect neural activity. An open question, in particular where
14 higher-level areas are involved, is how much tuning measured with one stimulus set reveals about tuning
15 to a new set. Here we ask this question by estimating tuning to hue in macaque V4 from a set of natural
16 scenes and a set of simple color stimuli. We found that hue tuning was strong in each dataset but was not
17 correlated across the datasets, a finding expected if neurons have strong mixed selectivity. We also show
18 how such mixed selectivity may be useful for transmitting information about multiple dimensions of the
19 world. Our finding suggest that tuning in higher visual areas measured with simple stimuli may thus not
20 generalize to naturalistic stimuli.

21

22 **New & Noteworthy:** Visual cortex is often investigated by mapping neural tuning to variables selected by the
23 researcher such as color. How much does this approach tell us a neuron’s general ‘role’ in vision? Here we
24 show that for strongly hue-tuned neurons in V4, estimating hue tuning from artificial stimuli does not reveal
25 the hue tuning in the context of natural scenes. We show how models of optimal information processing
26 suggest that such mixed selectivity maximizes information transmission.

27 **Keywords:** *Encoding model, tuning curve, V4, visual cortex, natural stimuli*

28

29 **Introduction**

30 Neurophysiology has long investigated the visual cortex by asking how its various areas encode the visual world. In
31 order to simplify experimental design, the majority of early work constructed stimuli sets in which only a few key
32 visual parameters varied. By then observing which areas show corresponding changes in neural activity, the visual
33 cortex can be described in terms of the variables for which neurons are highly tuned. This program has been
34 successful in characterizing that the response properties of neurons in the ventral stream ascend in complexity. V1 is
35 discussed as responding to “edge-detecting” Gabor filters [1], V2 to variations in local curvature [2], V4 to more
36 complex shapes [3], and IT to specific objects and faces [4], which together have inspired the theory that object
37 recognition proceeds via hierarchical image representations [5-7].

38 In recent decades there has been a greater interest in understanding how much the response to simple, parameterized
39 stimuli is informative about how the cortex encodes naturalistic stimuli of the type that make up everyday
40 experience. In addition to representing a field-wide shift towards ethology and natural paradigms, this question
41 pertains to how much tuning curves estimated from artificial stimuli are good models of neurons' general functional
42 role in visual processing. If tuning varies widely and unpredictably with context, the knowledge gained from
43 simplified stimuli would be unique and particular to the tested stimuli alone. It is therefore critical to test how much
44 tuning inferred from simplified stimuli is informative of tuning for more complicated stimuli.

45 In principle, such a question could be asked by steadily increasing the number of stimulus parameters that are varied
46 until the complexity of stimuli approaches that of natural scenes. The original tuning curve could then be
47 appreciated in the context of all others, and interactions with other variables characterized. However, this approach
48 is prohibitive because the number of stimuli required to be displayed to an animal scales exponentially with the
49 number of parameters separately varied. Any practical experiment of this type would need to leave many potential
50 parameters unvaried.

51 An alternative paradigm, and one that has a long history in early visual areas, is to characterize neurons directly
52 from their responses to natural images by regressing models of the visual encoding [8-13]. To build models without
53 a number of stimuli exponential in their dimension, one must place constraints upon what visual encodings are
54 viewed as possible. For early areas scientists often fit rather simple encoding models, such as linear or linear-
55 nonlinear models. Studies in this vein have found that many aspects, of the V1 response, such as preferred
56 orientation, appear similar between both natural images and drifting gratings [11]. However, other aspects of the V1
57 response are different for natural images [10], which limits how well artificial stimuli responses can inform
58 researchers about natural stimuli responses [12, 14]. In area V1, many tuning properties are shared between
59 responses to artificial and natural stimuli, raising the question if the same is true for higher level areas.

60 For higher areas, however, a regression methodology has not yet allowed such a comparison. In this work we focus
61 on area V4. While the response of V4 to natural scenes been studied [5, 15-18], most knowledge about the tuning of
62 V4 has derived from parameterized stimuli sets and it is not known how much tuning can vary with context. Our
63 specific focus is hue, as V4 was first characterized as a color area [19] before later studies found selectivity for other
64 visual features (such as orientation [20], curvature [3], shape [21-23], depth [24-26], and motion [27]; reviewed in
65 [28]). As revealed by displaying simple colored shapes, color-selective neurons are predominantly located in color

66 ‘globs’ [29], which intersperse ‘interglob’ regions more selective for orientation [30]. Glob cells are further arranged
67 by their color preference [31] and generally in the same hue sequence as is found in perceptual color space [32, 33].
68 It is not known how accurately these tuning curves describe V4’s response to color in natural stimuli.

69 In this work, we tested the limits of how much tuning in macaque V4 generalizes to complex stimuli by estimating
70 the hue tuning curves of color-responsive neurons by varying the hue of simple stimuli, and then again from their
71 responses to natural scenes. That is, we asked how well $P(Y|X, Z=z)$, which is the probability of spike counts Y
72 given hue X and a fixed context z , stands in for $P(Y|X)$, the average hue tuning over natural images. We regressed
73 hue tuning from natural scene responses using a progression of encoding models, including one based on a deep
74 artificial network pretrained to classify images [5], obtaining a similar conclusion from each. Overall we found that,
75 although hue strongly modulates the V4 response, the tuning curves estimated from responses to stimuli of a single
76 hue poorly described how hue affected responses to natural scenes. Lastly, we show how such mixed selectivity may
77 be useful for information transmission.

78

79 **Materials and Methods**

80 **Experimental setup: recordings**

81 We recorded from 96-electrode Utah arrays (1.0 mm electrode length) implanted in visual area V4. At the time of
82 the experiment, Monkey 1 (M1) was aged 5 years, 10 months and Monkey 2 (M2) was aged 9 years, 4 months.
83 Surgical details describing the implantation method can be found in previous publications [34, 35]. The array was
84 located in the left hemisphere for monkey M1 and in the right hemisphere for M2. Spikes were sorted off-line first
85 with an automated clustering procedure [36] and then refined by hand using custom MATLAB software
86 (<https://github.com/smithlabvision/spikesort>) taking into account waveform shape and interspike interval
87 distributions [37].

88 All experimental procedures were approved by the Institutional Animal Care and Use Committee of the University
89 of Pittsburgh.

90 **Artificial stimuli**

91 Both monkeys viewed uniform images of a single hue on a computer screen at 36 cm distance, with a resolution of
92 1024x768 pixels and a refresh rate of 100 Hz on a 21" cathode ray tube display. We found that full-field hues
93 elicited strong and selective responses from a majority of neurons (see Results). The hues were sampled from the
94 hue wheel in CIELUV color space (calculated with a D65 standard illuminant and standard observer) at increments
95 of 1 degree and at a chromaticity ensured to lie in the RGB gamut, and were presented in random sequence. Monkey
96 M1 freely viewed the stimuli, and was rewarded periodically for maintaining eye position on the screen for 4
97 seconds, after which time the static image was refreshed. The trial was ended if the monkey looked beyond the
98 screen during this duration. Monkey M2 was trained to fixate a small dot at the center of the screen for 0.3 seconds,
99 during which three images were flashed for 100ms each. A 0.5 second blank period interspersed each fixation.
100 Monkey 1 viewed 7,173 samples of the uniform hue stimuli over 10 sessions, while Monkey 2 viewed 1,119
101 samples during a single session. The full monitor subtended 55.5 degrees of visual angle horizontally and 43.1
102 degrees vertically. The monitor was calibrated to linearize the relationship between input luminance and output
103 voltage using a lookup table. This calibration was performed for grayscale images, and the color profile of the
104 monitor was not separately calibrated.

105 **Natural images**

106 Both monkeys viewed samples from a dataset of 551 natural images, obtained from a custom-made Google Images
107 web crawler that searched and downloaded images based on keywords such as cities, animals, birds, buildings,
108 sports, etc. Monkey M1 viewed images over 15 separate sessions, for a total of 77961 fixations. Monkey M2 viewed
109 images over two sessions on a single day, for a total of 6713 fixations. We then extracted the features from the
110 image patch centered around each fixation that would serve as model inputs. The image patch around fixation
111 corresponded to the 400 x 400 pixel block surrounding the center of gaze. This corresponds to a region 23.5 visual
112 degrees on a side.

113 **Gaze tracking and fixation segmentation**

114 We employed a free-viewing paradigm for one monkey (M1) and a fixed-gaze paradigm for the other (M2). The
115 location of each monkey's gaze on the screen was tracked with an Eyelink 1000 infrared tracker (SR Research,
116 Ottawa, Ontario, Canada). Visual stimuli were presented and the experimental trials were controlled by custom
117 MATLAB software in conjunction with the Psychophysics Toolbox [38]. For monkey M1, we segmented each

118 fixation as a separate event based on thresholding the position and velocity of the gaze coordinates. We did not
119 analyze activity occurring during eye movements. Once each fixation was separated, the average location of the
120 fixation was recorded and matched to image coordinates. Monkey M2 was trained to fixate on a dot positioned at the
121 center of each image. The gaze was tracked as for M1, but this time only to enforce fixation and terminate the trial if
122 the gaze shifted away from center.

123 **Session concatenation**

124 Although all recordings in M1 were performed with the same implanted Utah array, they were recorded over several
125 sessions. The recordings for M2 were made in a single session. In M1, this introduced the possibility that the array
126 might have drifted, and that a single channel might have recorded separate neurons in different sessions. To address
127 this possibility, we noted that spikes identified in a channel in one session will be less predictive of another session's
128 activity if the neurons are not the same, as we expect tuning to be relatively static across days [39, 40]. We thus
129 filtered out neurons whose uniform hue tuning changed across sessions. We trained a gradient boosting regression
130 model with Poisson targets to predict spike counts in response to the hue of the stimuli. Nuisance parameters, such
131 as duration of stimulus, gaze position, inter-trial interval, etc., were also included as model covariates to increase the
132 predictive power even for neurons that were not hue-tuned. We then labeled a neuron as having static tuning as
133 follows. First, we trained the model on each single session in a 8-fold cross-validation procedure and recorded the
134 mean pseudo- R^2 score. This score reflected how well the model could predict held-out trials on the same session.
135 Then, we re-trained the model on each session and predicted on a different session, for all pairs of sessions. This
136 resulted in a cross-prediction matrix with diagonal terms representing same session predictability (the 8-fold CV
137 score), and off-diagonal terms representing generalization between sessions. We did not concatenate sessions if hue
138 tuning estimated in one session could not predict hue responses in another session (i.e. the CV pseudo- R^2 score was
139 less than 0).

140 The natural image sessions were interspersed with the artificial sessions. If a natural image session occurred between
141 two artificial sessions, and a neuron showed static tuning both artificial sessions as identified in the above manner,
142 then that natural image session was included for the hue tuning comparison and model fitting. The recordings of
143 units from other natural image sessions were not used. This procedure improved our confidence that the neurons
144 recorded in different sessions were the same.

145 **Uniform hue tuning curve estimation**

146 Hue tuning curves were built for each neuron by plotting its spike rate on each fixation against the observed hue.
147 Spike rates were calculated from activity 50ms after fixation onset until 300ms or fixation offset, whichever came
148 first. For the visualizations in the figures, we performed LOWESS smoothing, in which each point of the curve is
149 given by a locally-weighted linear regression model of a fraction of the data. The error envelope of the curve
150 represents the 95% confidence interval given by bootstrapping over individual fixations. To calculate the correlation
151 between tuning curves, we did not correlate the LOWESS-smoothed curves but rather the simple binned averages.
152 We created 16 bins of hues and calculated the average spike rate for all stimulus presentations of those hues, then
153 correlated the 16-dimensional tuning curve vector with the natural image tuning curves.

154 To see how well simple tuning could explain the V4 response, we interpreted these tuning curves as models of the
155 natural image (presented in Results in Figure 2B). This was a linear model whose coefficients are set from the
156 uniform field tuning curve. Prediction was performed such that an image patch that was all a single color would
157 result in a prediction that was the firing rate observed in the uniform field condition, and mixtures of colors would
158 predict linear combinations of the corresponding observed firing rates. More precisely, the predicted firing rate was
159 a dot product of the tuning curve with the (normalized) hue histogram. To extract these hue histograms from image
160 patches, we calculated the hue angle of each pixel in the receptive field during a given stimulus presentation (see
161 Receptive Field estimation below) in CIELUV space. We then binned these hues into histograms with 16 bins of
162 hues, and these histograms served as the representation of hue input to the model. The final predicted response was
163 then added to a constant term to account for the difference in mean firing rate across contexts.

164 **Natural scene models**

165 We fit several models of the V4 response to natural scenes. Each of the models described below differs in the form
166 of the encoding and the manner by which hue tuning curves are reconstructed.

167 *Hue models*

168 Our first model describes neural activity as a function of the hues present in the receptive field on each fixation. We
169 used the same extraction of hues as above: we calculated the hue of pixels in the receptive field in CIELUV and
170 binned these hues into histograms with 16 bins of hues. Since the hue histograms have 16 bins, the base regression
171 problem to describe neural activity from hue is 16-dimensional.

172 As additional controls we included as covariates a small number of features unrelated to the images. To account for
173 possible stimulus adaption, we included the trial number in the session and also the number of times the monkey
174 previously fixated on that image. While all models predict the spike rate, which is already normalized by the fixation
175 duration, we included the fixation duration as an input to control for possible nonlinearities of rate with fixation
176 duration. We also included the duration of the saccade previous to the current fixation, the duration of the saccade
177 after fixation, the location of the fixation, the maximum displacement of the gaze position during the entire duration
178 of the fixation, and whether the pupil tracking was lost (often due to a blink) in the saccade before or after fixation.
179 Including these inputs allowed the nonlinear methods to control for factors which also may affect spike rate.

180 For our nonlinear model, we selected the machine learning method of gradient boosted decision trees as
181 implemented by XGBoost, an open-source Python package [41]. This method allows a Poisson loss function and has
182 previously been shown to be effective in describing neural responses [42]. Briefly, XGBoost trains multiple decision
183 trees in sequence, with each trained on the errors of the previous trees. We chose several regularization parameters
184 using Bayesian optimization for a single neuron. These parameters included the number of trees to train (200), the
185 maximum depth of each decision tree (3), the data subsampling ratio (0.5), the minimum gain (0.3), and the learning
186 rate (0.08). The generalized linear model (GLM) presented in Supp. Fig. 2 (DOI 10.6084/m9.figshare.17957957)
187 was a linear-nonlinear model with an exponential link function and a Poisson loss. We included elastic net
188 regularization, and selected the regularization coefficient for each neuron using cross-validation with k=8 folds in an
189 inner loop in the outer cross-validation for model scoring (see Model Scoring and Cross-Validation). We
190 implemented this with the R package r-glmnet [43].

191 To build tuning curves from hue model we predicted the response to a vector indicating which color was present
192 (that is, a “one-hot” vector with one entry per bin of hues that is all zeros except for the hue that is present). Then, to
193 estimate the measurement error of the tuning curves, we refit the models to the original neural responses resampled
194 with replacement (see Calculation of Error Bound). This resulted in tuning curves from hundreds of bootstrapped
195 model fits. In figures in which we display the tuning curves, the lower and upper error bounds represent the 5th and
196 95th percentiles of the tuning curves observed when refitting the models to the resampled data.

197 ***CNN model***

198 Our convolutional neural network (CNN) encoding model was based on findings that the intermediate layers of
199 pretrained networks are highly predictive of V4 responses [5]. Ours was built from the VGG16 network, which is a
200 large convolutional network trained to classify the images from the ImageNet dataset [44]. It contains 13
201 convolutional layers and 3 fully connected layers. We built an encoding model for each neuron from the activations
202 of layer 14 (the first fully-connected layer), which we found to have the highest predictive power in conjunction
203 with a nonlinear readout. We did not modify or refit this CNN to predict neural responses. Instead, we ran nonlinear
204 Poisson regression (XGBoost) to predict each neuron's response to an image from the values of layer 14 when the
205 VGG network was given the same image. The final model thus takes a fixation image as input, runs the image
206 through 14 layers of the VGG16 CNN, and then through a trained instance of XGBoost to predict the spike rate of a
207 neuron. We call the combination of the CNN model and the trained XGBoost for each neuron the "CNN model".

208 The CNN model could then be used to build tuning curves. We conceptualized this as extracting the average first-
209 order effect of hue upon the responses of this model to natural images. We perform the following cross-validated
210 procedure for each of 8 bins of hues. First, we train the CNN model (i.e. train the XGBoost regressor) on the training
211 set of the natural image dataset. We then modify the test set images by slightly desaturating all pixels whose hue lies
212 within the current hue bin. The bins were chosen to be large (8 in instead of 16) to so as to be less affected by pixel
213 noise and to speed computation. We desaturated by multiplicatively reducing the chroma of colors in LUV color
214 space, the same color space in which we define hue, by a certain factor. For robustness, we modified images at each
215 of many desaturation levels, ranging from 5% to 100% of their original chroma . We then obtained the predictions of
216 the CNN model to the original test set and also for each modified, desaturated test set, and take the average
217 difference of these two predictions across all images. This process is repeated in an 8-fold cross-validation
218 procedure, so that each image serves as the test set once. The resulting series of average differences can be plotted
219 against the desaturation. The slope of this line represents the average first-order contribution of that bin of hues to
220 the images in the dataset. Note that the value of slope reflects the scale the x-axis, which represents the
221 parameterization of the desaturation percentage. It is best to think of the units of slope as arbitrary; the important
222 result is the relative value of the slope between hues. Finally, the process was repeated for each bin of hues, resulting
223 in the tuning curve to hue.

224 We sought to validate this procedure on simulated data (Supp. Fig. 4, DOI 10.6084/m9.figshare.17958104). One
225 important aspect is that predictions are made on images that are as close to the distribution of images in the training

226 set as possible. Since images in which a single bin of hues are desaturated by 5% are visually indistinguishable from
227 the originals, this is not likely to be a concern. Nevertheless, we observed whether this method would be able to
228 reconstruct the hue tuning of simulated neurons. We constructed 20 simulated neurons that responded linearly to the
229 hues present in a receptive field. Each neuron was cosine tuned with a randomly selected hue angle. Linear
230 regression could perfectly reconstruct the hue tuning of these simulated neurons, as expected. The CNN method
231 could also reconstruct the tuning curves, though less well than linear regression. If linear tuning curves do exist,
232 then, the CNN method would be able to reconstruct them.

233 **Receptive field estimation**

234 To estimate hue tuning on natural scenes with the hue models, we needed to know which hues were present within
235 the RF on each fixation. We mapped the RFs by presenting sinusoidal gratings at four orientations, which were
236 flashed sequentially at the vertices of a lattice covering a portion of the visual field suggested by anatomical location
237 of the implant. For monkey 1, this procedure identified an average RF over neurons in the implant of 5.87° in
238 diameter (full-width at half-maximum) centered 8.94° below and 4.99° to the right of fixation, whereas for M2 we
239 found an average RF 7.02° wide centered 7.02° below and 7.02° to the left of fixation. The location of the RFs were
240 confirmed in the natural scene presentations as the pixel block that allowed the best predictions on held-out trials.
241 For the hue model analyses, on each fixation we obtained the model inputs by extracting the hues present in the
242 50×50 pixel block (2.93° of visual angle on a side) surrounding the centroid of the RFs of each monkey.

243 We did not use this RF information in the CNN model, which took as input the entire image region around the
244 fixation. Since information about spatial location preserved in the lower and intermediate layers of the CNN, the RF
245 for any neuron can be learned. This addressed any worry that our conclusions are dependent upon the RF
246 specification in the two hue models, and as well that the RF specification might systematically change for natural
247 images.

248 ***Model scoring and cross validation:***

249 We quantified how well the regression methods described neural responses by calculating the pseudo- R^2 score. This
250 scoring function is applicable to Poisson processes, unlike a standard R^2 score [45]. The pseudo- R^2 was calculated in
251 terms of the log likelihood of the true neural activity $L(y)$, the log likelihood of the predicted output $L(\hat{y})$, and the
252 log likelihood of the data under the mean firing rate $L(\bar{y})$.

$$R^2 = 1 - \frac{\log L(y) - \log L(\hat{y})}{\log L(y) - \log L(\bar{y})} = \frac{\log L(\hat{y}) - \log L(\bar{y})}{\log L(y) - \log L(\bar{y})}$$

253 The pseudo- R^2 is, at left, one minus the ratio of the deviance of the tested model to the deviance of the null model. It
254 can be also be seen, at right, as the fraction of the maximum potential log-likelihood. It takes a value of 0 when the
255 data is as likely under the tested model as the mean rate, and a value of 1 when the tested model perfectly describes
256 the data.

257 We used 8-fold cross-validation (CV) when assigning a final score to the models. The input and spike data were
258 segmented randomly by fixation into eight equal partitions. The methods were trained on seven partitions and tested
259 on the eighth, and this was repeated until all segments served as the test partition once. We report the mean of the
260 eight scores. If the monkey fixated on a single image more than once, all fixations were placed into the same
261 partition. This ensures that the test set contains only images that were not used to train the model.

262 *Calculation of error bounds*

263 Each estimate of a tuning curve represents, in essence, a summary statistic of noisy data. To estimate error bounds
264 on tuning curves, we relied on the nonparametric method of bootstrapping across trials, or for summary statistics of
265 the entire neural population, additionally bootstrapping across neurons. Since the uniform field hue tuning curves
266 used for correlations were simple averages of spike rates, binned over hue, we bootstrapped across trials to compute
267 the confidence intervals. The natural scene tuning curves for the nonlinear hue model represented the predicted
268 response to single hues. For these methods, we computed uncertainty bounds on their predictions to single hues by
269 retraining the methods on resampled datasets (with replacement) and selecting the 5th and 95th percentiles of the
270 predicted output for each bin. For the CNN method, the tuning curves were calculated from linear fits of the
271 difference in test set predictions as a function of hue bin desaturation. The difference in predictions was noisy across
272 images, with large changes predicted for some images but small changes predicted for other images. This noise
273 presented as uncertainty in the linear fit to the data. The error on the CNN tuning curve, then, represented the
274 uncertainty in the linear fit to the test set predictions.

275 The uncertainty on each of the tuning curves was then propagated into the correlation between the natural scene and
276 uniform field tuning curves. This was again done through bootstrapping. For a given natural scene/uniform field
277 correlation, we correlated the natural scene and uniform field tuning curves from hundreds of model fits upon

278 resampled data, yielding a large distribution of correlations. We then reported the mean, 5th, and 95th percentiles of
279 this distribution. The uncertainty of the mean across neurons included a bootstrap across the trials used to build the
280 tuning curves for each neuron, followed by a bootstrap across neurons.

281 *Normative analysis of mixed selectivity*

282 When neurons are nonlinearly selective for mixtures of a feature with others (a situation leading to tuning changing
283 with context) they are said to have nonlinear mixed selectivity. Nonlinear mixed selectivity has previously been
284 argued to be advantageous because it increases the dimensionality of the space of possible neural responses, which
285 allows a greater diversity of linear readouts for downstream tasks [46]. Here we look to the optimal coding literature
286 to find an alternative, more general justification. Our findings are summarized in Results.

287 A well-studied notion of optimality is that of Fisher efficient coding. In this framework the neural code is optimized
288 to increase the Fisher Information it contains about the features of stimuli that are important for behavior. Denoting
289 these features as $\boldsymbol{\theta} = \{\theta_1, \theta_2, \dots, \theta_M\}$, and the population activity of N neurons $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$, the Fisher
290 Information is a matrix defined elementwise as:

$$F(\boldsymbol{\theta})_{i,j} = \left\langle \left(\frac{\partial}{\partial \theta_i} \ln p(\mathbf{x}|\theta_i) \right) \cdot \left(\frac{\partial}{\partial \theta_j} \ln p(\mathbf{x}|\theta_j) \right) \right\rangle_{\mathbf{x}}.$$

291 Here $\langle \cdot \rangle_{\mathbf{x}}$ denotes the expectation over \mathbf{x} given sources of noise. The Fisher Information is intuitively similar to the
292 sensitivity of a representation across all neurons to a given dimension and value of $\boldsymbol{\theta}$. We will ask what
293 representations maximize the Fisher Information about all M encoded features.

294 Maximizing Fisher Information is a good measure of optimality because it describes how well any optimized
295 decoder can read out the features $\boldsymbol{\theta}$ from the response \mathbf{x} . Note that this normative reason is the *potential quality* of
296 the readout, rather than the *overall number* of potential linear readouts as in the work of Fusi et al. [46]. Following
297 the literature on optimal coding of neural populations [47-49], the decoding error can be bounded with the Cramer-
298 Rao inequality [47]:

$$\left\langle (\boldsymbol{\theta} - \hat{\boldsymbol{\theta}})^2 \right\rangle \geq \text{tr}(\mathbf{F}(\boldsymbol{\theta})^{-1}).$$

299 This states that the reconstruction error is lower-bounded by the trace of the inverse of the Fisher Information of the
300 neural population with respect to $\boldsymbol{\theta}$. By this metric, larger Fisher Information matrices allow lower error.

301 Our analysis focuses on the case of neurons that fire with a mean rate $f(\boldsymbol{\theta})$ with additional independent Poisson
 302 noise. In this case the Fisher Information is a sum over neurons [50]:

$$F(\boldsymbol{\theta}) = \sum_i^N f_i(\boldsymbol{\theta})^{-1} \nabla_{\boldsymbol{\theta}} f_i(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} f_i(\boldsymbol{\theta})^T.$$

303 To show that mixed selectivity is advantageous in this setting, we will show that Fisher Information is larger (as
 304 measured by the trace) when neurons respond to mixtures of features rather than code for only one feature. Our
 305 approach rests on the fact that neurons with Poisson noise have lower variance at lower firing rates. Additionally, we
 306 assume that the features that neurons code for are sparse and often not present. We find that mixed selectivity is
 307 better in this setting because many neurons can participate in coding, rather than waiting in silence for their single
 308 feature, and because distributed coding allow lower firing rates for the same sensitivity.

309 Let the spike rates for the aligned response (one feature per neuron) be denoted as $f(\boldsymbol{\theta})$. A random rotation of this
 310 response is $Rf(\boldsymbol{\theta})$, where R is a random rotation matrix. We can additionally restrict R to those rotations that
 311 preserve positivity of the resultant spike rates. If neurons spike with independent Poisson noise, the Fisher
 312 Information of this mixture is

$$F_M(\boldsymbol{\theta}) = \sum_i^N \frac{1}{r_i f(\boldsymbol{\theta})} \nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}))^T.$$

313 Here r_i is the corresponding row of the mixing matrix. In this case the trace of the Fisher Information (our measure
 314 of coding quality) is:

$$\text{tr} F_M(\boldsymbol{\theta}) = \sum_i^N \frac{1}{r_i f(\boldsymbol{\theta})} \text{tr} (\nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}))^T).$$

315 To ask if the rotation improves coding, we can ask if either of the two terms in the summand increase or decrease,
 316 on average. First, we find that the average of the right term does not change. This can be seen via the linearity of the
 317 trace and derivative and the cyclic property of the trace,
 318 $\sum_i \text{tr} (\nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} r_i f(\boldsymbol{\theta}))^T) = \sum_i \text{tr} (r_i \nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}))^T r_i^T) = \text{tr} (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}))^T \sum_i (r_i^T r_i))$. Since R is a
 319 rotation, $RR^T = I$, and the average reduces to its value absent a rotation, $\text{tr} (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) (\nabla_{\boldsymbol{\theta}} f(\boldsymbol{\theta}))^T)$. Thus, only the
 320 average of the scalar term due to Poisson noise $\frac{1}{r_i f(\boldsymbol{\theta})}$ is affected by the rotation.

321 The value of $\frac{1}{r_{if}(\theta)}$ depends on the sparsity of the response. When neurons code for only one feature which is usually
322 not present in a visual scene, most neurons are usually silent. The Fisher Information is controlled by the L active
323 neurons. We can approximate the sum with $\sum_i \frac{1}{f_i(\theta)} \approx \frac{L}{f_{typ}}$ where f_{typ} is the typical firing rate of the active neurons.
324 This will increase for any random rotation with a probability that increases with the number of neurons. In fact, the
325 dot product of a random vector such as r_i with a fixed vector $f(\theta)$ concentrates around 0 in high dimensions, though
326 this is complicated by the positivity constraint on $Rf(\theta)$. Taking a conservative approximation that the average
327 firing rate of the rotated population is the average firing rate of the unrotated population, $\sum_i \frac{1}{m_{if}(\theta)} \approx N \frac{N}{L f_{typ}}$. This is
328 much larger than $\frac{L}{f_{typ}}$ since $L < N$; specifically, a factor of $\frac{N^2}{L^2}$ larger.

329 Thus, the Fisher Information is larger when the response utilizes all Poisson neurons at small firing rates, as in the
330 rotated response, than when the response is concentrated on a sparse subset, as when each neuron codes for one
331 feature. This makes coding for single visual features a disadvantageous strategy.

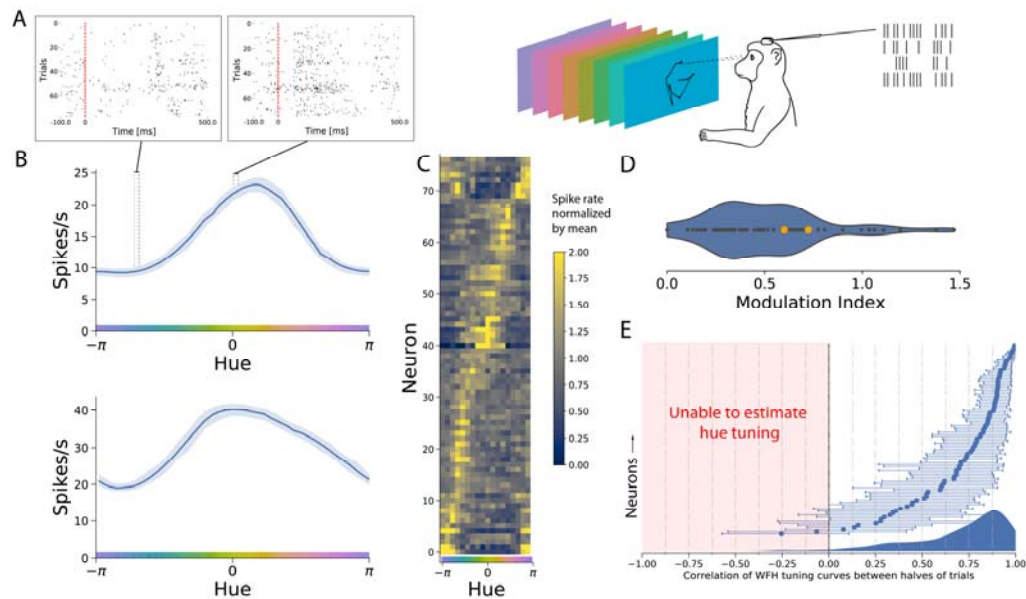
332 **Results**

333 We recorded the spike rates of neurons in area V4 of two macaques as they viewed images on a monitor. One
334 monkey (M1) freely viewed images as we tracked its gaze, while the gaze of the second monkey (M2) was fixed at
335 image center during image presentation. We analyzed the responses of 90 neurons in M1 over several viewing
336 sessions, taking care that the identity of cells on each electrode did not drift across sessions (see Methods: Session
337 Concatenation), and in M2 recorded from 80 neurons in a single session. We then estimated tuning curves from
338 responses to both artificial and naturalistic stimuli in order to ask if and how hue tuning generalizes.

339 **Tuning to hue on uniform screens**

340 We first measured hue tuning by varying the hue of a uniform flat screen (Fig. 1A). We found that most of our
341 neurons were well-tuned to specific hues, consistent with the previous literature on hue tuning in V4 [29, 30, 32].
342 Neurons' strong selectivity for hues evenly tiled the hue circle (Fig. 1C). We characterized the degree of modulation
343 with hue with the Modulation Index, calculated as the peak-to-peak range of the mean-normalized tuning curve (Fig.
344 1D). We also characterized our ability to estimate tuning by correlating the two tuning curves estimated on each half
345 of the trials, selected randomly and bootstrapped for confidence bounds. The confidence interval of this correlation

346 was usually high and excluded zero for 79/90 of neurons in M1 (Fig. 1C), but only for 17/80 neurons in M2 (Fig.
347 5A). A general trend across analyses was that neurons in M2 were more poorly described by hue than the neurons in
348 M1. This difference in monkeys was possibly due to the spatial heterogeneity of color responses in V4 [29, 30]. In
349 later analyses, we compared the hue tuning of neurons only when we could reliably estimate tuning. Thus, the
350 placement of our electrodes in M1 and M2 identified neurons that, as characterized by stimuli of a single hue,
351 appeared to robustly encode the hue of stimuli.



352

353 **Figure 1: Tuning curves estimated from responses to artificial stimuli.** Data from M1; see Fig. 5 for M2. A) We recorded
354 from neurons in area V4 as a monkey viewed fields of a uniform hue and examined the average evoked spike rate 50-300ms after
355 presentation. B) The uniform hue tuning curves for two example neurons, showing strong hue modulation, here displayed with
356 LOWESS smoothing of trial responses. C) Most neurons modulated their activity strongly with hue. Here the unsmoothed tuning
357 curves (mean rate in each bin of hues) are displayed normalized by per-neuron mean firing rate for comparison. D) The degree of
358 hue tuning can be characterized with a Modulation Index, which is the difference in min and max of the tuning curve after
359 normalization. The two example neurons of panel B are marked in orange. E) Our ability to reliably estimate hue tuning was
360 captured by correlating the tuning curve estimated on one half of the trials with the tuning curve estimated on the other half. This
361 correlation would be 1 in the no-noise or infinite-data condition, and if the 95% confidence bounds from bootstrapping include
362 zero we cannot reliably estimate tuning.

363 If these tuning curves capture how these neurons encode hue, these tuning curves should predict responses to other
364 types of stimuli. For example, we might expect that if a neuron preferred uniform fields of orange hue, then that

365 neuron would on average have higher firing rates for scenes containing predominantly orange hues. To test this, we
 366 displayed natural images in alternating sessions to the same monkeys (Fig. 2A). We found that the tuning curves
 367 were not at all informative of the natural image response. Specifically, we asked how well uniform hue tuning
 368 curves could predict natural scene responses by interpreting the curves as the coefficients of a linear response to hue,
 369 and then scoring this model (see Methods). Our scoring method is a pseudo- R^2 score which behaves roughly like an
 370 R^2 but is valid on data with Poisson noise. Scores of 1 indicate perfect prediction, scores of 0 indicate the mean
 371 firing rate is an equally good predictor, and negative scores indicate the mean rate is a better predictor of firing than
 372 the model. Of all but one neuron, the uniform field tuning curves predict natural responses with a negative pseudo-
 373 R^2 (Fig. 2B). Thus, neurons tuned for a certain color presented in isolation did not on average fire more when that
 374 color was present in natural scenes.

375 These observations can be explained if the hue tuning curves themselves are different between two contexts.
 376 Alternatively these neurons could respond more strongly to non-hue features that co-vary within natural images, like
 377 the visual texture of typically green plants. To distinguish these two possibilities, we next estimated tuning to hue
 378 directly from the responses to natural images and compared it with uniform hue tuning.

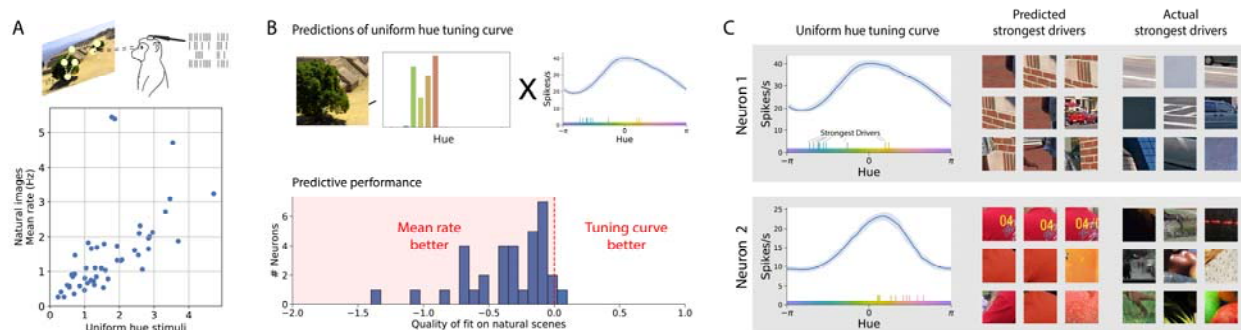


Figure 2. A) We displayed a large set of natural images to the same monkeys in interspersed sessions. Neurons fired at similar rates in these sessions as during presentations of a single hue. B) The hue-tuning on uniform hues (Fig. 1) can be treated as the coefficients of a linear model to predict neural responses to natural scenes, and scored. hue tuning on the artificial hue stimuli could not predict any variance in the natural scene responses. This is Displayed here is the histogram of the Poisson pseudo- R^2 goodness-of-fit scores of the tuning curves' predictions, which is below zero when the predictions underperform the mean firing rate. C) The tuning curves' low predictive utility is exemplified in the disparity between the stimuli that they predicted strong responses for (e.g. the most red), and the stimuli that actually elicited the strongest responses, which were consistently of

different hues.

379

380 Tuning to hue estimated from natural scenes

381 To investigate if hue tuning changes in the context of natural images, we directly regressed the contribution of hue to
 382 the neural response using two separate models. These models are of varying complexity and nonlinearity.
 383 Collectively they control for other visual features that drive V4 neurons, including interactions between hues, and
 384 rule out the possibility that visual confounds could explain the discrepancy between uniform field hue tuning and
 385 natural scene hue tuning.

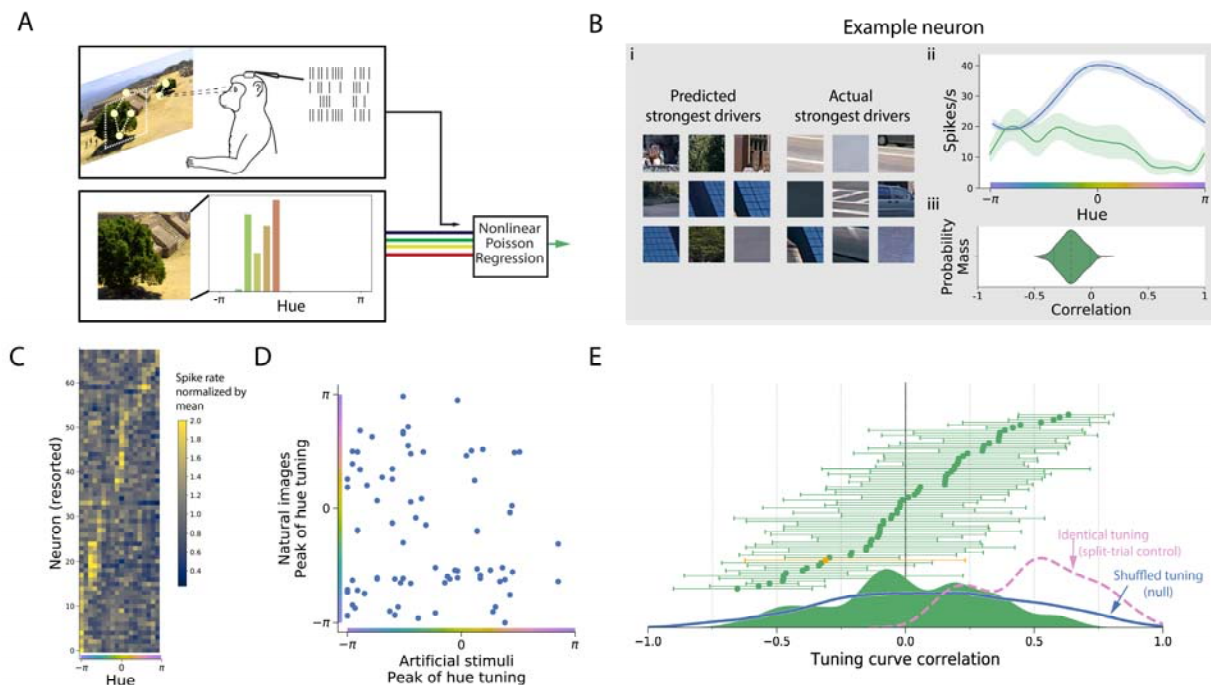


Figure 3. Tuning curves for hue estimated from the responses to natural images. Data from M1; see Fig. 5 for M2. A) We trained a nonlinear model with Poisson output to predict each neuron’s response from the hues present in its receptive field during the natural scene sessions. B) (i) The 9 trials that each model predicted to have highest firing rate looked similar to trials with the actual strongest response, unlike the uniform hue model. (ii) We built tuning curves from the model by observing its response to a single hue. The uncertainty of each curve is given by the 5th and 95th percentiles of hundreds of model fits to the trials resampled with replacement. (iii) This uncertainty is then propagated into the correlation between the uniform hue tuning curves and natural scene tuning curves. C) Tuning curves across all neurons (resorted by hue tuning in this condition). D) The peaks of these tuning curves plotted against the peaks of the tuning curves in the uniform hue condition show no circular

correlation across neurons, $p=0.82$. E) The correlations of the natural scene and the uniform hue tuning curves on each neuron (as shown in B(ii) and (iii)) show this is not an artifact due to multiple peaks in tuning curves. The neurons are sorted by their correlation to show a cumulative distribution. The two example neurons are highlighted in orange. Below: The smoothed density of all neurons' natural scene/uniform hue correlations is similar to what would be expected if neurons randomly shuffled hue tuning between conditions (overlaid, blue). Also overlaid (in pink) is the control distribution of how the correlations might appear if tuning were the same across stimuli, which is limited by neural noise and finite trials. This is estimated conservatively by correlating tuning estimated from one half of the natural scene trials with the tuning estimated on the other half.

386 *Nonlinear model with hue as an input feature*

387 We first modeled natural image responses as a nonlinear function of the hues present in the receptive field of
388 neurons during each fixation (Fig. 3A) plus control features to account for effects such as adaptation. This model,
389 which we refer to as the 'nonlinear hue model', predicted neural activity during natural scenes quite accurately for
390 neurons in both monkeys (Supp. Fig. 2B,C; DOI 10.6084/m9.figshare.17957957). We fit a nonlinear model because
391 nonlinear hue interactions have been previously observed in V4 [51], which would lead to a bias in a generalized
392 linear model (GLM) because hues are correlated in natural scenes (Supp. Fig. 2A). Indeed, the nonlinear model was
393 much more accurate than a GLM fit to predict spike rates using hues of natural images (Supp. Fig. 2B,C). This
394 approach to estimating hue tuning directly regresses the response to bins of hues in natural images.

395 We estimated hue tuning curves for the nonlinear hue by measuring its responses to single hues, in essence
396 reproducing the uniform hue experiment but on the natural scene model. These tuning curves showed clear
397 preference for small ranges of hues (Fig. 3C). We quantified our ability to estimate hue tuning in two ways. First, we
398 repeatedly refit the model on the natural scene trials resampled with replacement, and observed the distribution of
399 coefficients (Fig. 3Bii). This distribution was propagated through to later analyses such as the correlation between a
400 neuron's hue tuning estimated in either stimuli set. Secondly, we visualized how high the correlation of hue tuning
401 across conditions would have appeared if tuning were the same in both contexts, given all sources of noise and
402 measurement error. This we estimated by comparing the hue tuning curves from two non-overlapping halves of
403 natural scene trials (Supp. Fig. 1; DOI 10.6084/m9.figshare.17957408). Note that this split-trial control is a
404 conservative lower bound of our quality of estimation, as the model was fit on only half the number of trials. By this
405 measure, the nonlinear hue model was able to consistently estimate hue tuning for the most neurons in M1 (Supp.

406 Figure 1A, see also Fig. 3E) but for just two neurons in M2 (Fig. 5 D-F), which prevented a statistical analysis in
407 M2. These estimates of uncertainty serve as a baseline limit of how well we can observe changes in hue tuning.
408 We compared the tuning curves across stimulus sets for neurons for which we could consistently estimate hue
409 tuning. The peaks of the tuning curves did not show any correlation across neurons (Fig. 3D; circular correlation of -
410 0.02, unable to reject the uncorrelated hypothesis with $p=0.82$). In addition, the shapes of the tuning curves did not
411 correlate between conditions (Fig. 3E). If hue affected V4 responses in the same way in both contexts, we would
412 have observed the correlations between tuning curves across contexts to be at least as positive as the split-trial
413 control. This was not the case. In M1, the natural scene/uniform field tuning curve correlations were significantly
414 lower than these split-trial correlations ($p=1.0 \times 10^{-14}$, Wilcoxon signed-rank test; Supp. Fig. 1D, DOI
415 10.6084/m9.figshare.17957408), indicating that the observed change in hue tuning across contexts was not a
416 consequence of noise in the estimation of tuning. In fact, the spread of correlations between the two sets of tuning
417 curves was similar to the distribution that would arise if hue tuning shifted randomly between contexts (Fig. 3E
418 inset), which to preserve typical tuning shapes we approximated as the correlations between random neurons'
419 tuning. Thus, the regressed contribution of hue to the neural response had little relationship to the strong tuning
420 observed in response to stimuli of a single hue.

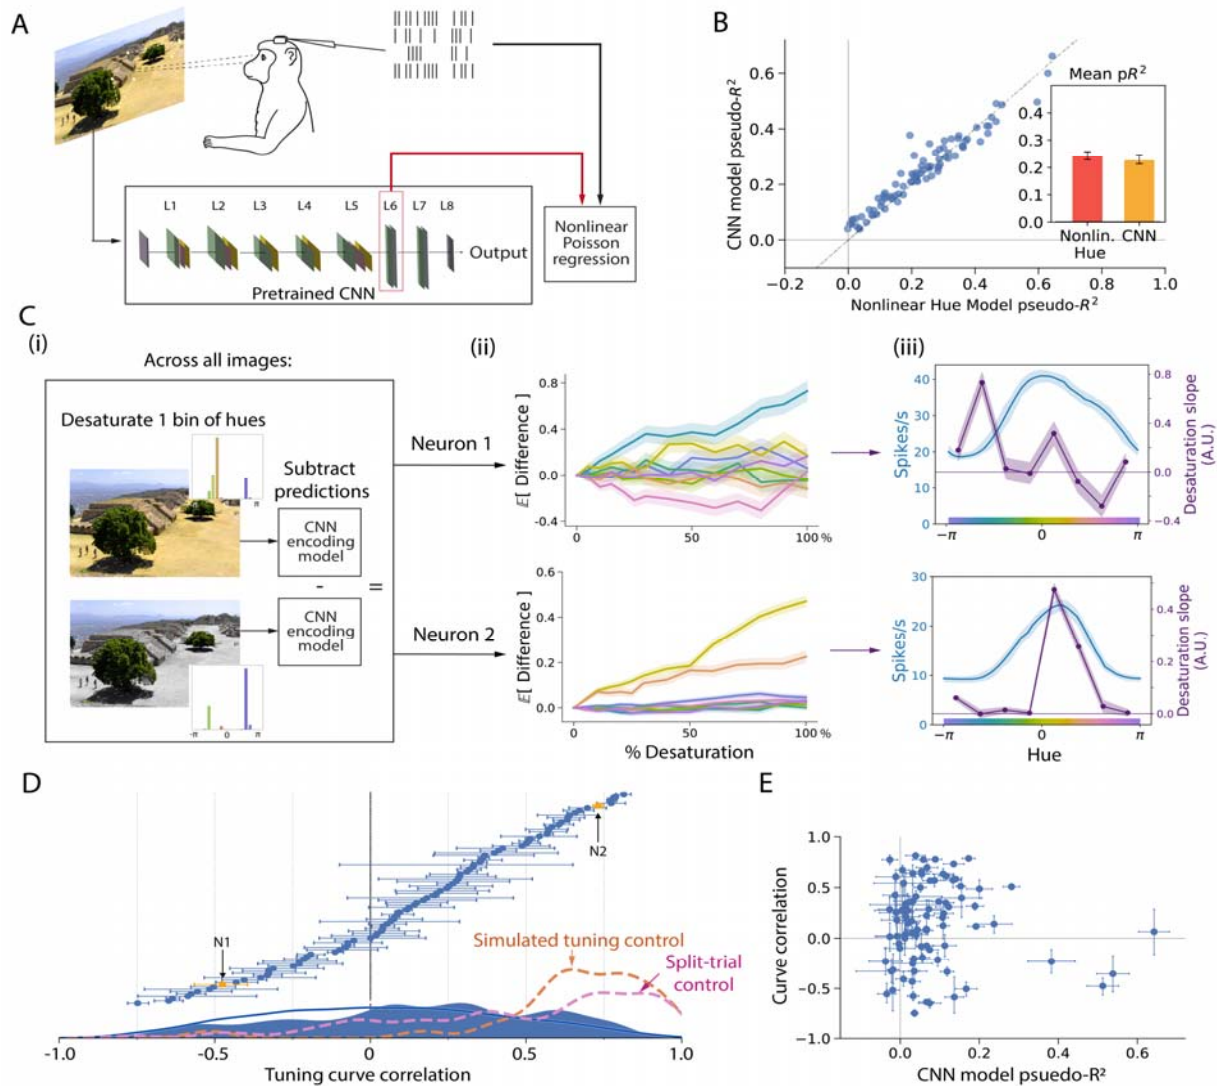


Figure 4. Tuning curves estimated for hue from a model of V4 responses built from a pretrained convolutional neural network (CNN). Data from M1; see Fig. 5 for M2. A) We trained a nonlinear Poisson regression model (gradient boosted trees) to predict the V4 response from the activations of an intermediate layer in the VGG16 network given the visual stimulus. B) The quality of the neural predictions on each neuron, measured by the cross-validated pseudo- R^2 score, were similar between the CNN model and the nonlinear hue model. C) We built hue tuning curves in the following manner: (i) For each image in a test set, we slightly desaturated all pixels in a bin of hues, and subtracted the CNN model’s predictions on the perturbed image from those on the original image. (ii) For each neuron, the average change in the predicted response across all test images was plotted against the percentage by which hues were desaturated. The slope of each line is, to first order, the average effect of that hue on the model response in the test set. The top and bottom plots show the same example neurons as in earlier plots. (iii) The resulting tuning curve (purple) summarizes the average effect of each of the 8 bins of hues – i.e. the slopes of the 8 desaturation

curves. It can be seen that the tuning of neuron 1 was poorly correlated with the uniform hue tuning (blue), while that of neuron 2 was well-correlated, in agreement with the hues of the strongest-driving stimuli shown in Fig. 1B. D) We calculated the correlation between the two tuning curves for all neurons. The distribution of correlations was lower than for the reconstructed hue tuning of simulated neurons (“simulated tuning control”; see also Supp. Fig. 3) as well as the distribution of correlations between tuning curves estimated from two non-overlapping halves of the natural scene trials (“split-trial control”; see also Supp. Fig. 1). E) The quality of the CNN model fit for each neuron did not predict the correlation of the tuning curves.

421 *Neural network model of V4 responses*

422 We next repeated the estimation of hue tuning on natural scenes with a more general model of V4 neurons that does
423 not rely on hand-specified summaries of the features present in a receptive field. This was important to ensure that
424 our results were not sensitive to design decisions in processing the images, as well as to account for the confounds of
425 other, non-hue features contained in the image. The two hue models would provide biased estimates of tuning if
426 neurons also responded to other visual features, and if these features co-varied in the image dataset with hue. If most
427 green objects are plants, for example, the observed dependence on green hues may be partially attributable to a
428 response to the high spatial frequency of greenery. Theoretically, one could include these features as additional
429 covariates, but the list of features that drive the V4 response in a simple manner (e.g. linearly) is not
430 comprehensively known. Good progress has been made with shape and texture [3, 52, 53], but arguably not enough
431 to thoroughly control for all non-hue features in a simple model. Controlling for other drivers of V4 thus requires a
432 model that learns relevant visual features instead of using features chosen by a researcher or parameterized by hand.
433 The model we selected was based on an encoding model of V4 that relates neural responses to the activations to the
434 upper layers of a convolutional neural network (CNN) pretrained to classify images [5]. Such “transfer learning”
435 models have also recently been used to infer optimal stimuli for neurons in V4 [18, 54]. Instead of pre-specifying a
436 receptive field estimated with sparse noise, we allowed the CNN model to learn any location sensitivity itself and
437 thus fed the entire fixation-centered image as input. Predictions of neural activity are obtained by passing the image
438 through the network, obtaining the intermediate network activations, and then passing these to a classifier trained for
439 each neuron (Fig. 4A). The predictions of neural activity given by this model were comparable in accuracy to those
440 of the nonlinear hue model (Fig. 4B) despite the model making many fewer assumptions about how raw pixels
441 related to responses.

442 Our initial, unsuccessful method to estimate hue tuning from this model was to simply observe the model's response
443 to images of a uniform hue, as before. However, this approach failed to reconstruct tuning on simulated data. This
444 interesting parallel to our main finding is likely due to the fact that uniform field test images are far outside the
445 domain of natural scenes on which the CNN was pretrained.

446 Instead, we developed a method to estimate hue tuning from the model that only uses responses to images close to
447 the domain of natural images. By slightly perturbing the hue of input images and observing the change in the learned
448 model's response, we could test the model's sensitivity to hues to in natural images (Fig. 4C). First, for a test set of
449 images not used for training, we desaturated all pixels within a bin of hues by a set percentage (Fig. 4Ci). The
450 percentage of desaturation varied from 0% (i.e. no change) to 100% (in which all pixels of one hue are taken to the
451 isoluminant grey). We took the difference between the model's predictions on the original and perturbed images and
452 examined how severely this difference depended on the level of desaturation (Fig. 3Cii, iii). For each neuron, we
453 averaged over the entire image dataset to yield the average effect of perturbing each hue on natural images. This
454 method established the effect of hue only in the tight neighborhood of each image, and is set up to estimate the
455 average local effect of hue on the natural image response.

456 To ensure that this process could in principle reconstruct correct tuning curves, we built simulated responses (Supp.
457 Fig. 3; DOI 10.6084/m9.figshare.17958104). We generated random cosine tuning curves, then simulated a hue
458 response by applying these as linear filters upon the histograms of the hues present in each image. We then
459 attempted to predict these simulated responses from the activations of the pretrained CNN given the raw images.
460 Using the method of progressively desaturating test images, we found we could reconstruct the original cosine
461 tuning curves with high accuracy (Fig. 4D overlay and Supp. Fig. 3). As a second, more conservative test, we also
462 performed the split-trial control for the actual V4 neurons, which involved repeating the entire analysis separately on
463 two non-overlapping halves of natural scene trials and then correlating the two resulting tuning curves. The split-
464 trial tuning curves showed significantly positive correlations for most neurons in M1 (Fig. 4D overlay) and M2 (Fig.
465 5). This method of querying the effect of hue could thus accurately estimate hue tuning curves from natural scene
466 responses in both monkeys.

467 We next asked if these tuning curves would be similar to the tuning curves to uniform hues. We found that the tuning
468 curves of one context were different from tuning in the other (Fig. 4D for M1 and Fig. 5G for M2). Among those
469 neurons for which we could consistently estimate hue tuning, the natural scene/ uniform hue tuning curve

470 correlations were significantly closer to 0 ($p=1.1 \times 10^{-8}$, Wilcoxon signed-rank test, Supp. Fig. 1D for M1; Fig. 2H for
 471 M2). This difference in tuning curves was not an artifact of our model fit or estimation method, as this would be
 472 measured in the split-trial control, and additionally we observed no correlation between the model's accuracy on
 473 unseen natural images and the natural scene/uniform field correlation (Fig. 4E and Fig. 5I).

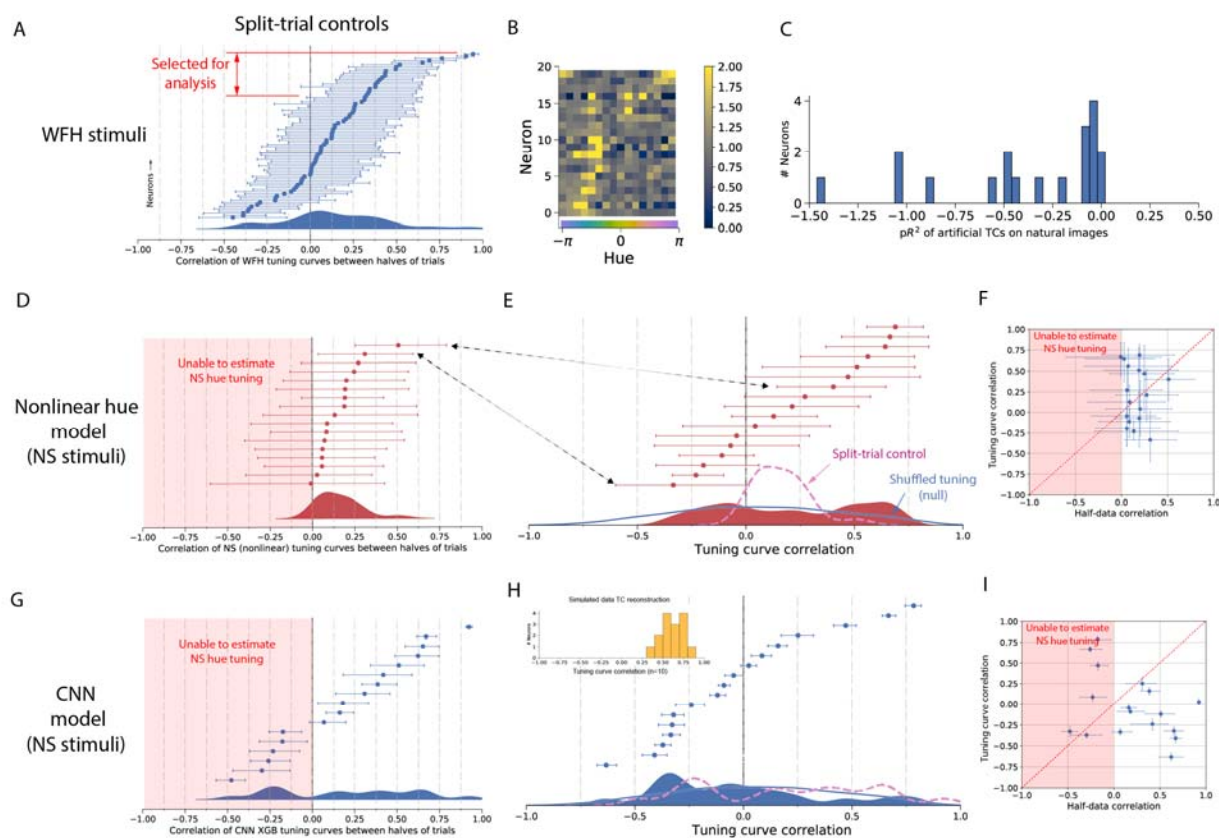


Figure 5: Collected data for M2. A) Most neurons in M2 showed poor hue tuning, and we were not able to consistently estimate uniform hue tuning nearly as well as for M1. B) Binned tuning curves for the neurons selected in A. C) As for M1, the uniform hue tuning curves were worse at predicting natural scene responses than the mean firing rate on natural scenes. D-F) Analysis of the natural scene tuning curves estimated by the nonlinear hue model was inconclusive. D) The natural scene tuning curves could not be estimated as consistently as for M1. E) The natural scene/uniform hue tuning curve correlations as estimated by the nonlinear hue model. Like for M1 (Fig. 2c) we overlay the split-trial distribution and the null distribution expected with random reshuffling of hue tuning. F) By a Wilcoxon signed rank test, we were unable to reject the null hypothesis that natural scene/uniform hue correlations are lower than the split-trial correlations ($p=0.65$) and thus it was not clear from the hue model on M2 neurons whether hue tuning does or does not change. G-I) Analysis of the natural scene tuning curves estimated with the CNN method. G) Distribution of correlations of hue tuning estimated of non-overlapping halves of trials. H) Natural scene/uniform hue correlations. Inserted is the distribution of natural scene/uniform hue

correlations of simulated neurons with cosine hue tuning. Since M2 saw 10x fewer trials than M1, we simulated again on this smaller dataset. I) Among the neurons for which we could consistently estimate hue tuning (i.e. with a positive correlation of tuning curves estimated on split data), all neurons had a higher split-trial natural scene curve correlation than a natural scene/uniform hue correlation. This was significant under a Wilcoxon signed rank test at $p=0.003$.

474

475 In addition to changes in tuning curve shape as captured by correlation, we also examined if the natural scene tuning
476 curves showed changes in the overall degree of hue modulation. We found that hue modulation – the maximum of a
477 tuning curve minus the minimum, normalized by the mean – was related across contexts, but weakly (Supp. Fig. 4;
478 DOI 10.6084/m9.figshare.17958221). Many neurons strongly modulated by hue on uniform fields had weak
479 responses to hue on natural scenes, and vice versa. Overall, the tuning curves estimated with this more advanced
480 method support our previous conclusion that hue tuning on uniform fields does not agree with the effect of hue in
481 natural scenes.




Figure 6: Interactions between features allow neurons to carry more information in their activity. A) In this two-dimensional tuning curve, a hypothetical neuron responds to only hue and carries no information about other variables. B) A hypothetical neuron that additionally responds another non-hue feature is informative about multiple dimensions of stimuli (due to its nonzero derivative). C) We can build a hue tuning curve for this neuron by varying hue with the other feature held fixed. If the average non-hue feature is different between natural images and uniform hues, the tuning curves to hue will differ between contexts.

482

483 Why features interact

484 A straightforward explanation of why hue tuning differs across visual contexts is that these neurons respond to
485 nonlinear combinations between hue and non-hue features, as shown schematically in Figure 6. What computational
486 advantage could explain this coding scheme for visual perception? It is clear that if the role of these neurons were to

487 encode hue alone, then any nonlinear interactions would be detrimental. This is because hue can no longer be
488 unambiguously read out without additional contextual information. Therefore these V4 neurons likely assist in a
489 more general task, like object recognition or segmentation. Other studies have also noted that color vision may be
490 best thought of in terms of task performance [55]; the absorbance spectra of the L and M photoreceptors in primates,
491 for example, are not maximally separated as in birds but rather overlap significantly, possibly because this helps to
492 discriminate and classify fruit and leaves [56]. The question then arises: why would neurons being responsive to
493 multiple features help visual processing?

494 One possible reason is that selectivity to multiple features increases the dimensionality of the space of possible
495 neural responses, which allows a greater diversity of linear readouts for downstream tasks [46]. This justification
496 prioritizes the diversity of possible uses of an area rather than the accuracy of encoding. Here, we look to the
497 optimal coding literature (e.g. [47, 50, 57]) to find an alternative normative justification. Starting with a certain set
498 of behaviorally-relevant visual features, what is the optimal way of representing these M features in a single
499 population of N neurons?

500 Our findings, derived in Methods, show that mixed selectivity allows for better neural codes when neurons fire as
501 Poisson processes and visual features are sparse and often not present. This is because a mixed selectivity strategy
502 allows many neurons to participate in each response, rather than a sparse subset. When each neuron responds to k
503 features, k times more neurons can respond on average to each scene. A distributed response in turn enables lower
504 firing rates for the same sensitivity, which is advantageous because Poisson noise has lower variance at low firing
505 rates. Coding quality, which is related both to the variance of internal noise and the sensitivity of the response when
506 measured by Fisher Information, thus improves when the response is maximally distributed across many neurons.
507 Recent studies using other measures of coding besides Fisher Information support this conclusion that mixed
508 selectivity improves neural encoding [58].

509 It should be noted that an optimal coding argument may not necessarily explain selectivity to all features.
510 Orientation selectivity, for example, is similar across contexts [11] and this may be related to fundamental visual
511 cortical architecture and the importance of visual form for behavior. Where neurons do not show mixed selectivity,
512 anatomical or behavioral constraints may override mixed selectivity's benefits of increased precision and quality of
513 the neural encoding.

514 **DISCUSSION**

515 For populations of V4 neurons in two macaques, we found that varying the hue of simple stimuli produced tuning
516 curves that do not accurately describe hue tuning measured from natural scenes. While some discrepancy may be
517 expected, we found that the two sets of tuning curves correlated not much better than chance. This finding was
518 robust across multiple methods of estimating tuning, which together accounted for the confounds of both hue-hue
519 interactions as well as of non-hue drivers of V4 activity. A hue tuning curve for V4 estimated from any one set of
520 stimuli thus does not universally describe the average response to hue on other stimuli.

521 **Known sources of modulation in visual responses**

522 The V4 response is modulated by a number of factors that change with visual context. These factors are divided in
523 the manner of their relevance to our findings. First are possible reasons why we might have observed low tuning
524 curve correlations even if, in fact, tuning did not change between contexts. The second category of factors are
525 known interactions between hue and other features in the V4 response that may explain why hue tuning in V4
526 changes with visual context. We will review both in turn.

527 Of first concern as a potential confound upon hue tuning estimation is visual attention [59]. A particularly relevant
528 form of attention is feature-based attention, in which neurons tuned for a feature (say, red) increase their firing rate if
529 that feature is attended to (as in the task, “find the red object”) [60, 61]. While the task of M1 was free viewing and
530 involved no instructions, it is likely that the monkey’s attention shifted during the task and that it was influenced by
531 object salience. This effect may bias our results if object salience were correlated with hue. Attention is less likely to
532 have presented a confound in the task of M2, in which gaze was fixed at center and stimuli were presented for
533 100ms. We have not directly controlled for attention, apart from trends in salience that might have been learned by
534 the CNN model, but we believe that the size of the apparent change in hue tuning cannot be attributable to salience-
535 hue correlations.

536 Neurons in V4 are have been shown to preferentially respond to objects near the center of attention, even when
537 attention falls away from fixation [62-64]. This phenomenon of receptive-field remapping is most problematic for
538 our hue models, which required that we extract the hues lying within the receptive field. If the monkeys’ attention
539 frequently strayed away from fixation, we would have extracted hues from an irrelevant image portion. This would
540 introduce some noise in the hue covariates, and therefore some smoothing of hue tuning curves. The CNN model

541 learned any spatial sensitivity directly from the natural scene responses. However, the effect of attention upon
542 receptive fields could not be modeled and it is likely that some smoothing of the hue tuning curve occurred for this
543 technique as well. Smoothing would obscure fine-scale structure in the tuning curves. As the curves were already
544 smooth, however, the natural scene/uniform field correlations should not be much diminished. The smoothing effect
545 is furthermore not consistent with our finding that many neurons have natural scene hue tuning with zero, or even
546 negative correlation with their uniform field tuning while still showing strong hue-dependent modulation.

547 We now turn to potential descriptions of the interactions that might have led to a shift in hue tuning across contexts.
548 One possibility is the behavioral phenomenon of color constancy, which would present as a neural correlate as
549 responses to the inferred surface color of objects rather than their apparent color (which reflects the color of ambient
550 light) [51]. This is a clear example of the nonseparability of the V4 response to hue, and a reason hue tuning might
551 change between any two, single stimuli. It is less obvious, however, that color constancy correlates would cause the
552 average effect of hue over all natural images to be different than on uniform hues. It would be expected that over
553 tens of thousands of images with a broad range of lighting conditions, color constancy would result in some
554 smoothing of the estimated tuning curve due to the difference between the pixels' hue and the inferred hue, and of
555 the same characteristic scale as their typical difference. Additionally we may expect a bias that would result from the
556 discrepancy between pure white and the average lighting condition. We expect this discrepancy to be small, and
557 therefore that natural scene tuning curves would still be strongly (though not perfectly) correlated with the uniform
558 field tuning curves. Though phenomena like color constancy would affect hue tuning on natural scenes, it cannot
559 account for the entire difference we observed, and it is likely that there exists other undocumented sources of
560 nonseparability.

561 A subpopulation of neurons in V4, so-called equiluminance cells, respond to object boundaries defined solely by
562 chromatic boundaries [65]. Such shapes are defined by changes in hue or saturation, and so it is worth asking
563 whether the response function of equiluminance cells includes interactions between hue/saturation and spatial
564 arrangement. However, it was not originally determined if the responses were actually separable in this way, as
565 neurons' hue tuning curves were characterized with a fixed shape. It is possible that equiluminant cells had fixed hue
566 tuning that was then modulated by shape. Thus, it is plausible but undetermined that equiluminance cells would
567 show different hue tuning across shape and explain our results.

568 The apparent shift in hue tuning in natural scenes may be partially explained by a multiplicative or generally
569 nonlinear interaction between shape and color, as was examined in a recent paper that jointly varied the hue and
570 shape of simple stimuli [66]. In a linear model with terms for hue and a multiplicative interaction between hue and
571 shape parameters, the authors observed a significant interaction between shape and color in the majority of cells
572 (44/60). This interaction would cause (linear) hue tuning to appear different for natural images with varying shapes,
573 as we observe. We note that other, undescribed features may also interact with hue, and that conclusively
574 determining which visual features interact would require presenting stimuli tiling many more dimensions of
575 variation.

576 **Implications for V4 and for the tuning curve approach**

577 Color responsivity has long been a defining feature of V4 [19, 67]. Recent studies have shown that localized areas in
578 V4 are strongly responsive to color [29], and furthermore that the anatomical organization of color preference on the
579 neocortex is similar to perceptual color spaces [31-33]. These findings have been taken as evidence that areas within
580 V4 are specialized for the perception of color. However, each of these studies characterized hue tuning by changing
581 the color of simple shapes. Since the color tuning of V4 neurons changes with visual context, as we show here, it is
582 possible that previous conclusions about the functional organization of V4 do not accurately describe how V4
583 processes more naturalistic stimuli.

584 It should be noted that our simple stimuli were not colored shapes chosen by hand to drive V4 neurons strongly, as
585 in many previous studies, but rather uniform screens. We found that these still elicited strong responses and well-
586 identified tuning curves. Nevertheless it may be objected that previous studies' stimuli may still measure hue tuning
587 that generalizes to natural scenes. However, this would require that the factors that modulate hue tuning in V4
588 neurons are only present in uniform screens. It is more parsimonious that hue-tuned neurons are modulated by
589 interactions with a range of spatial features, which collectively will cause tuning on any set of stimuli to not
590 generalize to naturalistic stimuli.

591 Based on the discovery of robust tuning for the color of simple visual stimuli, some studies have concluded that the
592 role of color-responsive areas in V4 is to represent color. Our results do not rule this out; for example these areas
593 might represent color but be modulated by what colors are likely given the surroundings. This would complicate a
594 read-out of color from V4, but may have other advantages like efficiency. It would be interesting to investigate this

595 possibility in future studies. An alternative possibility is that the color-sensitive areas of V4 are not specialized to
596 represent color, *per se*, but rather serve a more complex role within recognition and perception. This is analogous to
597 how V2 appears tuned to orientation but can perhaps be better described as processing naturalistic texture [68].
598 Furthermore, this role aligns with the suggestion that the ventral temporal cortex at large decomposes scenes into
599 neural activity such that object categories are linearly separable [69]. Thus, the color-responsive areas of V4 may
600 represent how color informs an inference of object identity [55]. Whether the color responses of V4 are an end to
601 themselves (i.e. representing color) or intermediate computations in a larger assessment of object identity [70], or
602 both, cannot be decided from this study; both are consistent with the data.

603 Our study joins a longer history of literature observing that, across many brain areas, tuning curves previously
604 characterized with simple stimuli in fact change with context. In V1, for example, researchers found that receptive
605 fields change with certain visual aspects that were not varied within previous stimuli sets, such as the presence of
606 competing orientations [71-74]. Even sound has been shown to modulate V1 receptive fields, at least in mice [75].
607 More recently, it was observed that receptive fields are different in the contexts of dense versus sparse noise for
608 neurons in layer 2/3 of V1 [76]. Spatio-temporal receptive fields of V1 neurons also appear different when estimated
609 on natural movies versus drifting gratings [10, 12] (though note that orientation tuning is similar for static natural
610 scenes versus gratings [11]). In other areas, such as for retinal ganglion cells [77-79] and in macaque M1, S1, and rat
611 hippocampus [42], contextual modulation in the form of nonlinear feature interactions have been identified by
612 showing perturbed natural images instead of white noise or by comparing the performance of a model that assumes
613 separability (such as a GLM) with a nonlinear model that does not. Thus, while tuning curves generalize in some
614 situations (e.g. [11]), it is common that they do not, and any assumption of separability of the neural response should
615 be verified. Furthermore, as derived in Methods, feature interactions are likely optimal for visual processing when
616 the full visual scene is represented in neural activity and should be expected. Unless specifically investigated, it
617 might not be correct to assume that a tuning curve accurately describes the neural response on different stimuli than
618 used to create it.

619 If it cannot be assumed that neural tuning is separable, however, it becomes necessary to test prohibitively many
620 stimuli or else make an alternative simplifying assumption. This is because the stimuli must scatter the entire space
621 of relevant features, rather than be systematically varied along just one feature at a time. Since the number of tested
622 stimuli must follow the number of potential feature combinations, the overall number of stimuli will grow

623 exponentially with the number of features. When there are very many features, even very large recording datasets by
624 today's standards may be insufficient.

625 One possible way forward is to make simplifying assumptions, i.e. to set strong priors of the kinds of tuning curves
626 that could be expected. This is the approach taken, for example, when modeling neurons using the activations of
627 deep neural networks pre-trained on image classification tasks [5, 80] or considering neural responses as
628 implementing a sparse code [9, 81]. To compare with the previous literature, single dimension experiments can then
629 be performed on these complex encoding models, as we demonstrate here, or alternatively performed directly on
630 artificial neural networks to gain intuition about what tuning curves say about information processing [82, 83]. In
631 general, finding suitable priors will require the use of strong theoretical ideas and mechanistic hypotheses. To
632 estimate tuning without assuming separability, then, neurophysiology must embrace and develop theories of neural
633 processing.

634 Acknowledgements

635 We are grateful to Samantha Schmitt for assistance with data collection and spike sorting. K.K., P.R., H.F., and A.B.
636 acknowledge support from NIH grants MH103910 and EY021579. M.S. was supported by NIH EY029250,
637 MH118929, EB026593, and NSF NCS 1954107.

638

639 Author Contributions

640 A.B. prepared the modeling methodology, conducted analyses, prepared figures, and wrote the initial manuscript.
641 P.R. and H.F. conceptualized the research, designed the experimental protocol, curated the data, prepared
642 methodology and initial analyses, and edited the manuscript. M.S. provided funding and resources for the
643 experiment, designed the experimental protocol, performed the data collection, advised the analysis methodology,
644 and edited the manuscript. K.K. conceptualized the project, provided funding, advised the experimental protocol,
645 supervised the data analysis, and edited the manuscript.

646 Competing Interests

647 The authors declare no competing interests.

648 Data Availability

649 All electrophysiological data is available upon request. Data can be provided in various stages of preprocessing to
650 expedite replication: as spike times in each session, or as spike counts paired with raw images, with fixation-
651 centered images, or with extracted hue histograms.

652 Code Availability

653 Our in-house code for spike sorting is available at <https://github.com/smithlabvision/spikesort>, and the code used to
654 process sorted data, run analyses, and create figures is available at <https://github.com/KordingLab/V4py>. Our
655 analyses depended on Python v2.7 and a large number of open-source projects that can be found listed in
656 <https://github.com/KordingLab/V4py/requirements.txt>, including Pandas v0.23, Tensorflow v1.1, Numpy v1.11, and
657 Keras v2.2.

658

659 References

- 660 1. Carandini, M., Demb, J.B., Mante, V., Tolhurst, D.J., Dan, Y., Olshausen, B.A., Gallant, J.L., and Rust, N.C. (2005). Do we
661 know what the early visual system does? *Journal of Neuroscience* 25, 10577-10597.
- 662 2. Hegd , J., and Van Essen, D.C. (2003). Strategies of shape representation in macaque visual area V2. *Visual neuroscience*
663 20, 313-328.
- 664 3. Pasupathy, A., and Connor, C.E. (2001). Shape representation in area V4: position-specific tuning for boundary
665 conformation. *Journal of neurophysiology* 86, 2505-2519.
- 666 4. Hung, C.P., Kreiman, G., Poggio, T., and DiCarlo, J.J. (2005). Fast readout of object identity from macaque inferior
667 temporal cortex. *Science* 310, 863-866.
- 668 5. Yamins, D.L., Hong, H., Cadieu, C.F., Solomon, E.A., Seibert, D., and DiCarlo, J.J. (2014). Performance-optimized
669 hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences* 111,
670 8619-8624.
- 671 6. Connor, C.E., Brincat, S.L., and Pasupathy, A. (2007). Transformation of shape information in the ventral pathway. *Current*
672 *opinion in neurobiology* 17, 140-147.
- 673 7. Logothetis, N.K., and Sheinberg, D.L. (1996). Visual object recognition. *Annual review of neuroscience* 19, 577-621.
- 674 8. Simoncelli, E.P., and Olshausen, B.A. (2001). Natural image statistics and neural representation. *Annual review of*
675 *neuroscience* 24, 1193-1216.
- 676 9. Vinje, W.E., and Gallant, J.L. (2000). Sparse coding and decorrelation in primary visual cortex during natural vision.
677 *Science* 287, 1273-1276.
- 678 10. David, S.V., Vinje, W.E., and Gallant, J.L. (2004). Natural stimulus statistics alter the receptive field structure of v1
679 neurons. *Journal of Neuroscience* 24, 6991-7006.
- 680 11. Touryan, J., Felsen, G., and Dan, Y. (2005). Spatial structure of complex cell receptive fields measured with natural images.
681 *Neuron* 45, 781-791.
- 682 12. David, S.V., and Gallant, J.L. (2005). Predicting neuronal responses during natural vision. *Network: Computation in Neural*
683 *Systems* 16, 239-260.
- 684 13. Felsen, G., and Dan, Y. (2005). A natural approach to studying vision. *Nature neuroscience* 8, 1643.
- 685 14. Olshausen, B.A., and Field, D.J. (2006). What is the other 85 percent of V1 doing. L. van Hemmen, & T. Sejnowski (Eds.)
686 23, 182-211.
- 687 15. Mazer, J.A., and Gallant, J.L. (2003). Goal-related activity in V4 during free viewing visual search: Evidence for a ventral
688 stream visual salience map. *Neuron* 40, 1241-1250.

- 689 16. Sharpee, T.O., Kouh, M., and Reynolds, J.H. (2013). Trade-off between curvature tuning and position invariance in visual
690 area V4. *Proceedings of the National Academy of Sciences* *110*, 11618-11623.
- 691 17. Gallant, J.L., Connor, C.E., and Van Essen, D.C. (1998). Neural activity in areas V1, V2 and V4 during free viewing of
692 natural scenes compared to controlled viewing. *Neuroreport* *9*, 2153-2158.
- 693 18. Cowley, B., Williamson, R., Clemens, K., Smith, M., and Byron, M.Y. (2017). Adaptive stimulus selection for optimizing
694 neural population responses. In *Advances in neural information processing systems*. pp. 1396-1406.
- 695 19. Zeki, S.M. (1973). Colour coding in rhesus monkey prestriate cortex. *Brain research* *53*, 422-427.
- 696 20. Desimone, R., and Schein, S.J. (1987). Visual properties of neurons in area V4 of the macaque: sensitivity to stimulus form.
697 *Journal of neurophysiology* *57*, 835-868.
- 698 21. David, S.V., Hayden, B.Y., and Gallant, J.L. (2006). Spectral receptive field properties explain shape selectivity in area V4.
699 *Journal of neurophysiology* *96*, 3492-3505.
- 700 22. Carlson, E.T., Rasquinha, R.J., Zhang, K., and Connor, C.E. (2011). A sparse object coding scheme in area V4. *Current*
701 *Biology* *21*, 288-293.
- 702 23. Popovkina, D.V., Bair, W., and Pasupathy, A. (2019). Modeling diverse responses to filled and outline shapes in macaque
703 V4. *Journal of neurophysiology* *121*, 1059-1077.
- 704 24. Watanabe, M., Tanaka, H., Uka, T., and Fujita, I. (2002). Disparity-selective neurons in area V4 of macaque monkeys.
705 *Journal of Neurophysiology* *87*, 1960-1973.
- 706 25. Hinkle, D.A., and Connor, C.E. (2001). Disparity tuning in macaque area V4. *Neuroreport* *12*, 365-369.
- 707 26. Hinkle, D.A., and Connor, C.E. (2002). Three-dimensional orientation tuning in macaque area V4. *Nature neuroscience* *5*,
708 665.
- 709 27. Mountcastle, V.B., Motter, B., Steinmetz, M., and Sestokas, A. (1987). Common and differential effects of attentive fixation
710 on the excitability of parietal and prestriate (V4) cortical visual neurons in the macaque monkey. *Journal of Neuroscience* *7*,
711 2239-2255.
- 712 28. Roe, A.W., Chelazzi, L., Connor, C.E., Conway, B.R., Fujita, I., Gallant, J.L., Lu, H., and Vanduffel, W. (2012). Toward a
713 unified theory of visual area V4. *Neuron* *74*, 12-29.
- 714 29. Conway, B.R., Moeller, S., and Tsao, D.Y. (2007). Specialized color modules in macaque extrastriate cortex. *Neuron* *56*,
715 560-573.
- 716 30. Tanigawa, H., Lu, H.D., and Roe, A.W. (2010). Functional organization for color and orientation in macaque V4. *Nature*
717 *neuroscience* *13*, 1542-1548.
- 718 31. Conway, B.R., and Tsao, D.Y. (2009). Color-tuned neurons are spatially clustered according to color preference within alert
719 macaque posterior inferior temporal cortex. *Proceedings of the National Academy of Sciences* *106*, 18034-18039.
- 720 32. Li, M., Liu, F., Juusola, M., and Tang, S. (2014). Perceptual color map in macaque visual area V4. *Journal of Neuroscience*
721 *34*, 202-217.
- 722 33. Bohon, K.S., Hermann, K.L., Hansen, T., and Conway, B.R. (2016). Representation of perceptual color space in macaque
723 posterior inferior temporal cortex (the V4 Complex). *Eneuro* *3*, ENEURO.0039-0016.2016.
- 724 34. Smith, M.A., and Sommer, M.A. (2013). Spatial and temporal scales of neuronal correlation in visual area V4. *Journal of*
725 *Neuroscience* *33*, 5422-5432.
- 726 35. Snyder, A.C., Morais, M.J., and Smith, M.A. (2016). Dynamics of excitatory and inhibitory networks are differentially
727 altered by selective attention. *Journal of neurophysiology* *116*, 1807-1820.
- 728 36. Shoham, S., Fellows, M.R., and Normann, R.A. (2003). Robust, automatic spike sorting using mixtures of multivariate t-
729 distributions. *Journal of neuroscience methods* *127*, 111-122.
- 730 37. Kelly, R.C., Smith, M.A., Samonds, J.M., Kohn, A., Bonds, A., Movshon, J.A., and Lee, T.S. (2007). Comparison of
731 recordings from microelectrode arrays and single electrodes in the visual cortex. *Journal of Neuroscience* *27*, 261-264.
- 732 38. Brainard, D.H., and Vision, S. (1997). The psychophysics toolbox. *Spatial vision* *10*, 433-436.
- 733 39. Bondar, I.V., Leopold, D.A., Richmond, B.J., Victor, J.D., and Logothetis, N.K. (2009). Long-term stability of visual pattern
734 selective responses of monkey temporal lobe neurons. *PLoS one* *4*, e8222.
- 735 40. McMahan, D.B., Jones, A.P., Bondar, I.V., and Leopold, D.A. (2014). Face-selective neurons maintain consistent visual
736 responses across months. *Proceedings of the National Academy of Sciences* *111*, 8251-8256.
- 737 41. Chen, T., and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd*
738 *international conference on knowledge discovery and data mining*. (ACM), pp. 785-794.
- 739 42. Benjamin, A.S., Fernandes, H.L., Tomlinson, T., Ramkumar, P., VerSteeg, C., Chowdhury, R.H., Miller, L.E., and Kording,
740 K.P. (2018). Modern machine learning as a benchmark for fitting neural responses. *Frontiers in computational neuroscience*
741 *12*.
- 742 43. Friedman, J., Hastie, T., and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate
743 descent. *Journal of statistical software* *33*, 1.
- 744 44. Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv*
745 *preprint arXiv:1409.1556*.
- 746 45. Cameron, A.C., and Windmeijer, F.A. (1997). An R-squared measure of goodness of fit for some common nonlinear
747 regression models. *Journal of Econometrics* *77*, 329-342.
- 748 46. Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed
749 selectivity in complex cognitive tasks. *Nature* *497*, 585.

- 750 47. Seung, H.S., and Sompolinsky, H. (1993). Simple models for reading neuronal population codes. *Proceedings of the*
751 *National Academy of Sciences* *90*, 10749-10753.
- 752 48. Wang, Z., Stocker, A.A., and Lee, D.D. (2012). Optimal neural tuning curves for arbitrary stimulus distributions: Discrimax,
753 infomax and minimum Lp loss. *Advances in Neural Information Processing Systems (NIPS)* *3*, 2168-2176.
- 754 49. Brunel, N., and Nadal, J.P. (1998). Mutual Information, Fisher Information, and Population Coding. *Neural Computation* *10*,
755 1731-1757.
- 756 50. Wang, Z., Stocker, A.A., and Lee, D.D. (2013). Optimal neural population codes for high-dimensional stimulus variables. In
757 *Advances in neural information processing systems*. pp. 297-305.
- 758 51. Kusunoki, M., Moutoussis, K., and Zeki, S. (2006). Effect of background colors on the tuning of color-selective cells in
759 monkey area V4. *Journal of Neurophysiology* *95*, 3047-3059.
- 760 52. Portilla, J., and Simoncelli, E.P. (2000). A parametric texture model based on joint statistics of complex wavelet
761 coefficients. *International journal of computer vision* *40*, 49-70.
- 762 53. Okazawa, G., Tajima, S., and Komatsu, H. (2015). Image statistics underlying natural texture selectivity of neurons in
763 macaque V4. *Proceedings of the National Academy of Sciences* *112*, E351-E360.
- 764 54. Bashivan, P., Kar, K., and DiCarlo, J.J. (2019). Neural population control via deep image synthesis. *Science* *364*, eaav9436.
- 765 55. Rosenthal, I., Ratnasingam, S., Haile, T., Eastman, S., Fuller-Deets, J., and Conway, B.R. (2018). Color statistics of objects,
766 and color tuning of object cortex in macaque monkey. *Journal of vision* *18*, 1-1.
- 767 56. Osorio, D., and Vorobyev, M. (2008). A review of the evolution of animal colour vision and visual communication signals.
768 *Vision research* *48*, 2042-2051.
- 769 57. Brunel, N., and Nadal, J.-P. (1998). Mutual information, Fisher information, and population coding. *Neural computation* *10*,
770 1731-1757.
- 771 58. Johnston, W.J., Palmer, S.E., and Freedman, D.J. (2020). Nonlinear mixed selectivity supports reliable neural computation.
772 *PLoS computational biology* *16*, e1007544.
- 773 59. Chelazzi, L., Della Libera, C., Sani, I., and Santandrea, E. (2011). Neural basis of visual selective attention. *Wiley*
774 *Interdisciplinary Reviews: Cognitive Science* *2*, 392-407.
- 775 60. Motter, B.C. (1994). Neural correlates of attentive selection for color or luminance in extrastriate area V4. *Journal of*
776 *Neuroscience* *14*, 2178-2189.
- 777 61. Mirabella, G., Bertini, G., Samengo, I., Kilavik, B.E., Frilli, D., Della Libera, C., and Chelazzi, L. (2007). Neurons in area
778 V4 of the macaque translate attended visual features into behaviorally relevant categories. *Neuron* *54*, 303-318.
- 779 62. Connor, C.E., Preddie, D.C., Gallant, J.L., and Van Essen, D.C. (1997). Spatial attention effects in macaque area V4.
780 *Journal of Neuroscience* *17*, 3201-3214.
- 781 63. Connor, C.E., Gallant, J.L., Preddie, D.C., and Van Essen, D.C. (1996). Responses in area V4 depend on the spatial
782 relationship between stimulus and attention. *Journal of neurophysiology* *75*, 1306-1308.
- 783 64. Gallant, J.L., Connor, C.E., Rakshit, S., Lewis, J.W., and Van Essen, D.C. (1996). Neural responses to polar, hyperbolic,
784 and Cartesian gratings in area V4 of the macaque monkey. *Journal of neurophysiology* *76*, 2718-2739.
- 785 65. Bushnell, B.N., Harding, P.J., Kosai, Y., Bair, W., and Pasupathy, A. (2011). Equiluminance cells in visual cortical area V4.
786 *Journal of Neuroscience* *31*, 12398-12412.
- 787 66. Bushnell, B.N., and Pasupathy, A. (2012). Shape encoding consistency across colors in primate V4. *Journal of*
788 *neurophysiology* *108*, 1299-1308.
- 789 67. Zeki, S. (1980). The representation of colours in the cerebral cortex. *Nature* *284*, 412-418.
- 790 68. Ziemba, C.M., Freeman, J., Movshon, J.A., and Simoncelli, E.P. (2016). Selectivity and tolerance for visual texture in
791 macaque V2. *Proceedings of the National Academy of Sciences* *113*, E3140-E3149.
- 792 69. Grill-Spector, K., and Weiner, K.S. (2014). The functional architecture of the ventral temporal cortex and its role in
793 categorization. *Nature Reviews Neuroscience* *15*, 536.
- 794 70. DiCarlo, J.J., and Cox, D.D. (2007). Untangling invariant object recognition. *Trends in cognitive sciences* *11*, 333-341.
- 795 71. Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. *Visual neuroscience* *9*, 181-197.
- 796 72. Knierim, J.J., and Van Essen, D.C. (1992). Neuronal responses to static texture patterns in area V1 of the alert macaque
797 monkey. *Journal of Neurophysiology* *67*, 961-980.
- 798 73. Sillito, A., and Jones, H. (1996). Context-dependent interactions and visual processing in V1. *Journal of Physiology-Paris*
799 *90*, 205-209.
- 800 74. Fitzpatrick, D. (2000). Seeing beyond the receptive field in primary visual cortex. *Current opinion in neurobiology* *10*, 438-
801 443.
- 802 75. McClure Jr, J.P., and Polack, P.-O. (2019). Pure tones modulate the representation of orientation and direction in the
803 primary visual cortex. *Journal of neurophysiology* *121*, 2202-2214.
- 804 76. Yeh, C.-I., Xing, D., Williams, P.E., and Shapley, R.M. (2009). Stimulus ensemble and cortical layer determine V1 spatial
805 receptive fields. *Proceedings of the National Academy of Sciences* *106*, 14652-14657.
- 806 77. Heitman, A., Brackbill, N., Greschner, M., Sher, A., Litke, A.M., and Chichilnisky, E. (2016). Testing pseudo-linear models
807 of responses to natural scenes in primate retina. *bioRxiv*, 045336.
- 808 78. McIntosh, L., Maheswaranathan, N., Nayebi, A., Ganguli, S., and Baccus, S. (2016). Deep learning models of the retinal
809 response to natural scenes. In *Advances in neural information processing systems*. pp. 1369-1377.
- 810 79. Goldin, M.A., Lefebvre, B., Virgili, S., Ecker, A., Mora, T., Ferrari, U., and Marre, O. (2021). Context-dependent selectivity
811 to natural scenes in the retina. *bioRxiv*.

- 812 80. Ponce, C.R., Xiao, W., Schade, P., Hartmann, T.S., Kreiman, G., and Livingstone, M.S. (2019). Evolving super stimuli for
813 real neurons using deep generative networks. *bioRxiv*, 516484.
- 814 81. Felsen, G., Touryan, J., and Dan, Y. (2005). Contextual modulation of orientation tuning contributes to efficient processing
815 of natural stimuli. *Network: Computation in Neural Systems* 16, 139-149.
- 816 82. Pospisil, D.A., Pasupathy, A., and Bair, W. (2018). 'Artiphysiology'reveals V4-like shape tuning in a deep network trained
817 for image classification. *Elife* 7, e38242.
- 818 83. Morcos, A.S., Barrett, D.G., Rabinowitz, N.C., and Botvinick, M. (2018). On the importance of single directions for
819 generalization. *arXiv preprint arXiv:1803.06959*.
- 820