

1 The Geometry of Masking in Neural Populations

2

3 Dario L. Ringach

4

5 Departments of Neurobiology and Psychology, UCLA, Los Angeles, CA 90095.

6

7

8 **We introduce a geometric approach to study the representation of orientation by**
9 **populations of neurons in primary visual cortex in the presence and absence of an**
10 **additive mask. Despite heterogeneous effects at the single cell level, a simple**
11 **geometric model explains how population responses are transformed by the mask**
12 **and reveals how changes in discriminability and bias relate to each other. We**
13 **propose that studying the geometry of neural populations can yield insights into**
14 **the role of contextual modulation in the processing of sensory signals.**

15

16

17 Introduction

18

19 Individual neurons in primary visual cortex respond to stimulation within restricted areas
20 of the visual field, which define their classical receptive fields¹⁻³. These responses can be
21 modulated by contextual stimuli presented within the classical receptive field or in the
22 surrounding regions⁴⁻⁶. Cross-orientation and surround suppression are two well-known
23 examples of contextual modulation^{5,7-21}.

24

25 The role that contextual modulation plays in cortical function remains an open question.
26 Some have considered such interactions to be directly involved in image processing, such
27 as the detection and enhancement of smooth, spatially extended contours²²⁻³⁷. Others
28 maintain that the fundamental goal of contextual modulation is to generate a sparse,
29 efficient representation of natural images^{6,38-45}. Distinguishing between these theories is
30 not straightforward, as the their goals are not mutually exclusive⁶.

31

32 Here we focus on how contextual modulation transforms the activity of neural populations.
33 Contextual modulation has been studied extensively in single neurons, leading to the
34 development of the influential normalization model^{6,46,47}. We have recently shown,
35 however, that key properties derived from the classic formulation of normalization, such
36 as contrast and subspace invariance, do not strictly hold at a population level⁴⁸. Thus,
37 we and others⁴⁹ see new opportunities in the study of contextual modulation at the level
38 of neural populations. The present study follows up on this line of work by studying how
39 the coding of orientation by neural populations is transformed in the presence of a mask.

40

41 Results

42

43 Measurement of population responses in masked and unmasked conditions

44

45 We measured the responses of neural populations in mouse primary visual cortex using
46 two-photon imaging (**Methods**). Mice were head-restrained but otherwise free to walk on
47 a rotating wheel. The visual stimulus consisted of two conditions (**Fig 1A**). In the
48 *unmasked condition*, a full-field sinusoidal grating was presented while its orientation
49 changed linearly with time $\theta = \pi t/T$ with a period $T = 10s$. This type of continuously
50 rotating stimulus has previously been used to measure orientation maps⁵⁰. In the *masked*
51 *condition*, the same rotating stimulus was presented superimposed on top of a mask
52 consisting of a sinusoidal grating oriented vertically. We estimated the spiking responses
53 of neurons using a standard processing pipeline involving image registration, signal
54 extraction and deconvolution⁵¹. The periodic nature of the stimulus was evident in the
55 temporal responses of cells (**Fig 1B**). This is because neurons tuned to one orientation
56 respond once per cycle. As described in earlier studies⁵², locomotion modulated the
57 overall responses of the population (**Fig 1B**, shaded regions).

58

59 Single neuron responses in masked and unmasked conditions are heterogenous

60

61 We computed the average response of neurons in the unmasked and masked conditions
62 over the cycle of the stimulus (**Fig 1C**). The temporal responses were shifted by the mean
63 stimulus-response delay (see **Methods**). After this correction, the temporal profile of the
64 response can be interpreted as an estimate of the tuning curve of the neuron. In this
65 representation, the mask is present at an orientation of 90 deg (**Fig 1C**, dashed lines;
66 subsequent figures omit the location of the mask to avoid clutter).

67

68 We observed heterogeneity of responses at the single cell level. Some cells were well
69 tuned to orientation in the unmasked condition but were completely suppressed by the
70 addition of the mask (**Fig 1Ca**). Others did show such dramatic suppression, but
71 responded with a scaled down version of their unmasked responses (**Fig 1Cb**) – a
72 behavior consistent with the normalization model^{6,46,53}. Some cells showed little or no
73 difference between the responses in the two conditions (**Fig 1Cc**). Another group saw
74 their unmasked responses scale up by the mask (**Fig 1Cd**). Finally, somewhat
75 surprisingly, a set of neurons showed very weak responses in the unmasked condition
76 but responded vigorously in the presence of the mask (**Fig 1Ce**)⁴⁸.

77

78 We studied the range of behaviors in single cells (**Fig 1C**) by comparing the mean
79 response of the i – th neuron over the stimulation cycle between unmasked and masked
80 conditions, which we denote by μ_u^i and μ_m^i respectively (**Fig 2A**). We found a significant

81 anti-correlation: the stronger a neuron responded in the unmasked condition the weaker
82 its response was in the masked condition and vice-versa ($n = 3920, r = -0.55, p =$
83 5.6×10^{-312}). We refer to neurons at the extremes of the distribution of behavior as
84 *grating* and *plaid* cells (**Fig 2A**, shaded areas). These groups were formally defined as
85 the cells attaining the 10% lowest (grating cells) and highest (plaid cells) ratios of
86 $\log_2(\mu_m^i/\mu_u^i)$ (**Fig 2A**, inset). We note these groups represent behaviors found at the
87 extremes of a unimodal distribution (there is no evidence of discrete classes of neurons).
88 Grating and plaid cells had different preferred orientations. Grating cells were
89 preferentially tuned to the orientation orthogonal to the mask in both conditions (**Fig 2B**,
90 left panels). In grating cells, the introduction of the mask scaled down the responses by
91 about a third but did not affect tuning (note the different y-scales in **Fig 2B**). This is the
92 type of responses one might expect from the classic normalization model^{46,53}. Plaids cells,
93 on the other hand, were preferentially tuned to the orientation of the mask (90 deg) when
94 probed with grating stimuli in the unmasked condition, although their responses were
95 relatively weak (**Fig 2B**, top right). Instead, and somewhat surprisingly, these cells were
96 most responsive to the orthogonal orientation (0 deg) under the presence of the mask
97 (**Fig 2B**, bottom right) – in other words, they responded best when the stimulus was a
98 plaid with orthogonal components.

99

100 **A geometric framework to study contextual modulation in neural populations**

101

102 The data in unmasked and masked conditions can each be represented as a matrix where
103 the columns represent the tuning function of each cell (**Fig 3A**). To ease visualization, we
104 ordered neurons by their preferred orientation. The rows of the matrix represent the
105 population response to a given orientation. We denote the mean population responses
106 as a function of orientation in the unmasked and masked conditions by $r_u(\theta)$ and $r_m(\theta)$,
107 respectively. These vectors can be thought to describe parametric (closed) curves in a
108 high dimensional space as $\theta \in [0, \pi]$ traverses the orientation domain (the dimension
109 being the number of neurons in the population). We aim to understand the shape of these
110 curves, the nature of the transformation $T: r_u(\theta) \rightarrow r_m(\theta)$ introduced by the mask, and
111 how the outcome affects the discriminability of stimuli and biases the estimation of
112 orientation in the masked condition.

113

114 In what follows we denote by $d_u(\theta, \varphi)$ the cosine distance between $r_u(\theta)$ and $r_u(\varphi)$ (**Fig**
115 **3B**, left). The cosine distance is one minus the cosine of the angle between the two
116 vectors. Because these vectors have positive entries representing a spike rate, the
117 distance is bounded between zero and one. Similarly, we define $d_m(\theta, \varphi)$ as the cosine
118 distance between $r_m(\theta)$ and $r_m(\varphi)$ (**Fig 3B**, middle). Under certain assumptions about
119 the uniformity of the noise, the measurements $d_u(\theta, \varphi)$ and $d_m(\theta, \varphi)$ capture the ability
120 of the population to discriminate between two angles in each condition. Finally, $d_{um}(\theta, \varphi)$

121 denotes the cosine distance between the population representation of θ in the unmasked
122 condition and the representation of φ in the masked condition. This measure captures the
123 relative positions of the two curves and will induce biases in the estimation of orientation
124 in the presence of a mask. Namely, non-zero biases result when the structure of $d_{um}(\theta, \varphi)$
125 is not perfectly diagonal (**Fig 3B**, right). In the sequel, we also denote the normalized
126 population vectors by $\hat{r}_u(\theta) = r_u(\theta)/\|r_u(\theta)\|$ and $\hat{r}_m(\theta) = r_m(\theta)/\|r_m(\theta)\|$.

127

128 **Orthogonality of signal and noise subspaces**

129

130 We selected cosine distance as a metric because a substantial component of neural
131 variability in the population response occurs along its direction⁵⁴. To show this, we
132 computed the mean and the covariance of the responses, $r_u(\theta)$ and $\Sigma_u(\theta)$. For each
133 orientation, we compared the direction of the population response with the direction of the
134 largest eigenvector, $v_1(\theta)$, of the covariance matrix (**Fig 3C**). The largest eigenvector
135 accounted for nearly a third of the variability (and more than three times the variance
136 accounted for by the next largest eigenvalue) (**Fig 3D**) and its direction was very similar
137 to that of the largest eigenvector – the correlation coefficient between $r_u(\theta)$ and $v_1(\theta)$
138 was 0.88 ± 0.05 (mean \pm 1SD) (**Fig 3D**, inset). This large component of variability is due
139 to fluctuations in behavioral state which modulates the gain of the response vector^{52,55,56}.
140 An implicit assumption behind the adoption of the cosine metric is that the orientation of
141 the stimulus is coded by the *direction* of the population vector^{46,57}. Thus, the result can be
142 rephrased by stating that for each orientation, the direction of largest variability, $\hat{r}_u(\theta)$, is
143 orthogonal to the direction of the encoding, $\hat{r}'_u(\theta)$, which is tangent to the unit sphere.

144

145 **Multidimensional scaling of population responses.**

146

147 To gain insight about the geometry of the curves and their relative positions we visualized
148 them using multi-dimensional scaling using the cosine distance as a metric (**Fig 4**). The
149 curves represent the embeddings of $\hat{r}_u(\theta)$ (blue) and $\hat{r}_m(\theta)$ (red) in 3D space, while the
150 spheres of matching colors indicate the point where the stimulus cycle begins. We define
151 the mean population response over the entire stimulation cycle as the *white point*, which
152 we denote by μ_u and μ_m . The grey arrows depict the shift of the white points
153 between unmasked and masked conditions, with the stem of the arrow positioned at μ_u
154 and the head at μ_m . These examples are typical of what we observed in our experiments.

155

156 The shapes of $\hat{r}_u(\theta)$ and $\hat{r}_m(\theta)$ are similar, with the masked representation being a scaled
157 down version of the original. The curves are farthest from each other at the beginning of
158 the cycle, when the pattern in the masked condition consists of an orthogonal plaid and
159 the one in the unmasked condition is a horizontal grating. The two curves are closest to
160 each other near the middle of the cycle, when the pattern in the masked condition is a

161 vertical grating with 100% contrast and the one in the unmasked condition is a vertical
162 grating with 50% contrast. The curve $\hat{r}_m(\theta)$ appears to be rotated away from that of $\hat{r}_u(\theta)$,
163 with the axis of rotation passing near the representation the mask. These features were
164 consistent across our experiments suggesting that a scaling and rotation may explain the
165 transformation of $\hat{r}_u(\theta)$ into $\hat{r}_m(\theta)$ induced by the mask. Of course, these visualizations
166 ought to be interpreted with caution, as they are only approximate representations of the
167 geometry of high dimensional objects. Thus, we must check these first impressions of the
168 geometry by doing appropriate calculations in the native space.

169

170 Masking shrinks and rotates population responses

171

172 To verify our perception that curves are shrinking we computed their lengths⁵⁸, $L_u =$
173 $\int_0^\pi \|\hat{r}'_u(\theta)\| d\theta$ and $L_m = \int_0^\pi \|\hat{r}'_m(\theta)\| d\theta$. The arguments represent the angular velocity at
174 which the population changes its orientation and represent a measure of discriminability
175 between nearby angles. The length, therefore, represents local discriminability summed
176 over all orientations^{58,59}. The mask had the effect of reducing the overall length of the
177 curves by a factor of 0.84 ± 0.05 (mean \pm 1SD) (**Fig 5A**). As we will soon demonstrate,
178 this shrinkage is not uniform, but peaks near the orientation of the mask.

179

180 To verify our impression that the mask induces a change in the direction of the mean
181 population activity, we defined the white-point shift as $\Delta = d(\mu_u, \mu_m) / ((\rho_u + \rho_m) / 2)$. Here,
182 ρ_u represents the average radius of the curve in the unmasked condition, calculated as
183 $(1/\pi) \int_0^\pi d(r_u(\theta), \mu_u) d\theta$, and a corresponding definition applies to ρ_m . In other words, we
184 measure the shift of the white point in terms of the average radius of the curves. Across
185 the population we find $\Delta = 0.71 \pm 0.15$ (mean \pm 1SD) – a relatively large fraction (**Fig 5B**),
186 which is consistent with the visualizations from multi-dimensional scaling. We will see this
187 shift is important because it is partly responsible for generating biases in the estimation
188 of orientation in the masked condition.

189

190 Rejection of the linear combination model

191

192 With the geometric formalism in place, we can test a common model of population
193 responses, which postulates that the response to a plaid can be written as a linear
194 combination of the population responses to the individual components^{48,60}. The
195 implication for our experiment is that $r_m(\theta) \in \text{span}\{r_u(\theta), r_u(\pi/2)\}$ (recall the mask has
196 orientation $\pi/2$). One way to test the prediction is to measure the angle formed by the
197 vector $r_m(\theta)$ and the plane $\text{span}\{r_u(\theta), r_u(\pi/2)\}$. The results show a significant departure
198 from the prediction, with angular deviations larger than 30 deg and significantly higher
199 than zero ($p < 0.005$, bootstrap, see **Methods**) (**Fig 5C**). Thus, the present data rule out
200 the linear combination model, thereby confirming and extending a prior result⁴⁸.

201

202 **The mask impairs discriminability and biases the decoding of orientation**

203

204 Next, we analyzed changes in discriminability induced by the mask. Discriminability
205 between any two orientations depends both on the distance between the mean population
206 vectors and the statistics of the noise. As mentioned above, if the statistics of the noise
207 are uniform in the sense that they translate with the direction of the population, we expect
208 discriminability to be proportional to the distances $d_u(\theta, \varphi)$ and $d_m(\theta, \varphi)$. Nevertheless,
209 given we have ~ 100 cycles we were able to compute a proper *d-prime* measure for both
210 masked and unmasked conditions, which we denote by $D_u(\theta, \varphi)$ and $D_m(\theta, \varphi)$ (see
211 **Methods**) (**Fig 6A**). To measure local discriminability (or just noticeable differences) we
212 defined the threshold for detection in the unmasked condition $T_u(\theta)$ as the minimal angle
213 Δ such that $D_u(\theta - \Delta/2, \theta + \Delta/2) \geq 4$ (**Fig 6A**, iso-discriminability contours); we adopted
214 a similar definition for the threshold in the masked condition, $T_m(\theta)$. Comparison of the
215 thresholds in the two conditions revealed that the mask elevated thresholds around the
216 orientation of the mask (at 90 deg) (**Fig 6A**). Interestingly, the thresholds around the
217 orientation orthogonal to the mask (0 deg) were not affected. A similar result is obtained
218 if we perform a similar analysis based on $d_u(\theta, \varphi)$ and $d_m(\theta, \varphi)$ and assume the noise is
219 uniform (data not shown).

220

221 We then analyzed how the presence of the mask can lead to biases in estimates of
222 orientation. We used a decoder based on population voting^{57,61}. The estimated orientation
223 was obtained as $\hat{\theta} = (1/2) \arg \int_{\theta=0}^{\pi} (1 - d_{um}(\theta, \varphi)) \exp(i2\theta) d\theta$. In other words, the
224 population votes for each angle with a weight that depends on the distance to the
225 representation of each angle in the unmasked orientation – the smaller the distance the
226 strongest the vote. The bias is then $b = (\hat{\theta} - \varphi) \bmod \pi$. We observe that except at the
227 orthogonal orientation the estimates are biased towards the orientation of the mask (**Fig**
228 **6B**). These biases arise because $d_{um}(\theta, \varphi)$ does not have a non-diagonal structure -- the
229 local minima of $d_{um}(\theta, \varphi)$ occur slightly off the main diagonal (**Fig 6B**, bottom, white
230 contours). Similar results are obtained using a simpler winner-takes-all decoder, where
231 we pick $\hat{\theta} = \arg \min_{\theta} d_{um}(\theta, \varphi)$.

232

233 **A geometric model for population transformations under masking**

234

235 We tested if a simple geometric model⁶², originally developed to explain the effects of
236 adaptation in psychophysical experiments, could explain our masking data (**Fig 6C**). The
237 model assumes that in the unmasked condition the population response $r_u(\theta)$ describes
238 a trajectory around the unit circle and that the effect of the mask is to translate and scale
239 this response to yield $r_m(\theta)$. Translation is towards the population direction evoked by the
240 mask, and the scaling is a typically a factor smaller than one. The model assumes that

241 orientations are identified by the direction of the population vector, and that the decoder
242 is unaware of the shift in the white point of the population between the two conditions. In
243 other words, estimates of orientation are based on the direction of the population vector
244 measured relative to the origin (which is also μ_u in this case) (**Fig 6C**). The model has
245 only two parameters, the magnitude of the shift of the white point and a scaling factor. Its
246 simplicity allows one to compute an analytical expression for both the threshold and the
247 bias⁶² (see **Methods**). Indeed, this model captures some of the behavior of observed in
248 the data. First, it reproduces the dependence of threshold with orientation in the masked
249 condition, showing a maximum centered around the orientation of the mask. Second, it
250 reproduces the shape of the bias reflecting an attraction towards the orientation of the
251 mask.

252
253 The model, however, fails in three fundamental ways. First, in the model the population
254 responses in both conditions lie within the same plane. As two independent vectors span
255 the entire plane, it has to be the case that $r_m(\theta) \in \text{span}\{r_u(\theta), r_u(\pi/2)\}$ (so long as
256 $r_u(\theta) \neq r_u(\pi/2)$). In other words, this model satisfies linear combination⁶⁰. However, we
257 have already shown this is not the case in the data (**Fig 5D**). Second, both curves make
258 a single revolution around the origin. This means that the lengths of the normalized
259 responses are the same and equal to 2π , predicting a ratio $L_m/L_u = 1$. Another way of
260 stating this is that both $\hat{r}_u(\theta)$ and $\hat{r}_m(\theta)$ are unit circles, and that the two curves are
261 different parametrizations of the same curve. However, the data show the ratios of the
262 lengths to be significantly less than one (**Fig 5A**, tailed sign-test, $p = 9.3 \times 10^{-10}$). Third,
263 the threshold is directly linked to how fast the population response changes its direction
264 with angle, which is given by $\|\hat{r}'_u(\theta)\|$ and $\|\hat{r}'_m(\theta)\|$. The faster the population direction
265 rotates the lower the thresholds for discrimination. Faster rotation speeds However, as
266 we just pointed out the average across all orientations is constant under this model,
267 $\int_0^{2\pi} \|\hat{r}'_u(\theta)\| d\theta = \int_0^{2\pi} \|\hat{r}'_m(\theta)\| d\theta = 2\pi$. This means that if the mask increases
268 discriminability for some orientations it must decrease it for others⁵⁸. This is reflected in
269 the fact that the threshold in the masked condition fluctuates around the mean for the
270 unmasked condition (**Fig 6C**). The data, however, indicates that the effect of the mask
271 was to impair the discriminability around the orientation of the mask, while there is little or
272 no effect at the orthogonal orientation (**Fig 6A**, right column). The data refutes the
273 prediction that increases in threshold at some orientations must be accompanied by
274 decreases in threshold at other orientations (**Fig 6C**).

275
276 Can model be extended to account for our results? We know from the analysis of single
277 cell responses that some neurons are unresponsive in the unmasked condition but
278 respond robustly in the presence of a mask (**Fig 1B** and **2A**). This fact alone indicates the
279 population responses in the masked and unmasked conditions do not lie within the same
280 subspace. Thus, one way to extend the model is to allow the population responses to be

281 displaced relative to each other along a third dimension (**Fig 6D**). Consider the response
282 in the unmasked condition to be the unit circle and the one in the masked condition to be
283 a the result of an affine transformation, $T(r) = \alpha Ar + t$, where A is an orthogonal matrix
284 (representing a rotation), α is a scaling factor, and t a translation. It is then possible to
285 find parameters of the transformation that reproduce the ratio between the lengths of the
286 curves, as well as the dependence of discriminability and bias on orientation (**Fig 6D**,
287 middle and right panels). An affine transformation can be represented in homogenous
288 coordinates as $T(r) = Ar$ where the population vector now has an extra dimension to
289 allow for translation. We can then write the transformation of the *normalized* population
290 responses as $T(\hat{r}) = A\hat{r} / \|A\hat{r}\|$, which we recognize as a projective linear
291 transformation⁶³ (**Fig 6E**). When the model is fit to the data in individual experiments, it
292 nicely accounts for the observed transformations (**Fig 6F**).

293

294 Discussion

295

296 Understanding how populations of neurons encode a physical attribute of a sensory
297 stimulus, and how responses are transformed by contextual modulation is a central
298 question of system neuroscience⁴⁹. Here we considered the simpler question of how the
299 orientation of a sinusoidal grating is transformed by an additive mask.

300

301 At the single cell level, we observed a wide range of responses (**Figs 1C,2**). Interestingly,
302 we found a group of neurons that do not respond to gratings in the unmasked condition
303 but respond strongly to plaids in the masked condition (**Fig 2A**). The maximal response
304 of this type of these *plaid neurons*, occurs when the pattern is an orthogonal grating (**Fig**
305 **2B**). This finding implies that the responses in masked and unmasked conditions do not
306 lie within the same subspace. This explains why the linear model (**Fig 5C**) and the 2D
307 geometric model (**Fig 6C**) fail to account for the data. Grating and plaid cells are
308 reminiscent of pattern and component cells^{64,65}. We use different terms because the
309 definitions are not equivalent. We note, however, that the pattern index used to classify
310 cells as pattern/component correlates with the plaid/grating response we use here⁶⁶ and
311 that mouse primary visual cortex contains a larger proportion of pattern cells than found
312 in non-human primates⁶⁷. Thus, we suspect that the neurons engaged during masking
313 that do not respond strongly to gratings in the masked condition could represent pattern
314 cells.

315

316 We observed that plaid cells, when probed with a single component in the unmasked
317 condition, responded optimally (albeit weakly) to the orientation of the mask (**Fig 2B**).
318 While somewhat puzzling, the behavior in the masked condition might be explained if the
319 addition of a grating orthogonal to a cell preferred orientation (as defined with single

320 gratings) increases its response by releasing it from inhibition from oblique orientations in
321 a ring model of orientation tuning⁶⁸.

322
323 The main contribution of the present study is the introduction of a geometric approach to
324 study contextual modulation of neural populations^{48,62}. The analysis revealed that, despite
325 a substantial heterogeneity in the behavior of individual cells, the map relating population
326 responses in masked and unmasked conditions can be approximated as an affine
327 transformation. When considering normalized responses, the corresponding map is a
328 projective linear transformation⁶³. The finding is so-far limited to masking, but we
329 conjecture it may hold for other types of contextual modulation, such as interactions
330 between the classical receptive field and the surround and sensory adaptation. Indeed,
331 the a 2D model which accounts for psychophysical data on adaptation⁶² (**Fig 6C**) is an
332 instance of an affine transform.

333
334 The geometric approach proved helpful in understanding several important properties of
335 how population responses are modified by the introduction of a mask. First, it offered a
336 straightforward test (and rejection) of a linear combination model⁶⁰. The result could be
337 easily understood as the mask moving the population activity out of its original subspace.
338 Second, the analyses revealed that the transformation cannot be a reparameterization of
339 the same curve, of which the 2D model is a special case⁶² (**Fig 6C**). The reason is that
340 all reparameterizations leave the length of the curve invariant. In contrast, the mask was
341 observed to shrink the length of the representation (**Fig 5A**). Third, we were able to show
342 that the shift in the white point of the population is large compared to the radius of the
343 curve (**Fig 5C**). This explains how a decoder which is unaware of such shift is bound to
344 generate biased estimates. Finally, it clarified how a simple transformation can introduce
345 changes in discriminability and bias in decoding (**Fig 6**).

346
347 Our finding of a white-point shift appears to be at odds with the idea that adaptation keeps
348 the mean population response invariant (population homeostasis)⁶⁹. In our terminology,
349 population homeostasis would have predicted that $\mu_u = \mu_m$, meaning no white-point shift.
350 We suspect the reason for this discrepancy is rooted in the stimuli used. In the referenced
351 study, a sequence of gratings with randomly chosen orientations was presented to the
352 population. In one condition, the orientations were uniformly distributed; in the second
353 condition, one orientation (the adapter) appeared more frequently than the others. In both
354 conditions any one stimulus consists of a single grating. It is possible that such design
355 failed to engage the plaid cells that clearly play an important role in shifting the white
356 point. Similarly, a previous report⁶⁰ selected cells to be analyzed only if their orientation
357 tuning in response to a grating showed good selectivity (circular variance less than 0.85).
358 Perhaps, plaid cells that were either unresponsive or weakly responsive to gratings failed
359 to pass this criterion. The result would be biased towards gratings cells and it is possible

360 that a linear combination model could be satisfactory when applied to this subpopulation
361 of neurons.

362

363 Our findings indicate that analyzing the patterns of activity across large population of
364 neurons we might be able to discover some general principles of sensory representation,
365 including topological⁷⁰ and geometrical structure, that are undetectable at the single cell
366 level. These patterns can allow us to describe the transformations of representations in a
367 simple way, as appeared to be the case for masking. Novel technologies that allow the
368 recording of hundred or thousands of neurons simultaneously provide an ideal test bed
369 for these ideas.

370

371 **Methods**

372

373 **Animals:** All procedures were approved by UCLA's Office of Animal Research Oversight
374 (the Institutional Animal Care and Use Committee) and in accord with guidelines set by
375 the US National Institutes of Health. A total of 5 tetO-GCaMP6s mice (Jackson Labs),
376 both male (3) and female (2), aged P35-56, were used in this study. Mice were housed in
377 groups of 2-3, in reversed light cycle. Animals were naïve subjects with no prior history of
378 participation in research studies. We imaged 30 different fields, and obtained data for
379 3920 cells, for a median of 111 cells per field (range: 50 to 275).

380

381 **Surgery:** Carprofen and buprenorphine analgesia were administered pre-operatively.
382 Mice were then anesthetized with isoflurane (4-5% induction; 1.5-2% surgery). Core body
383 temperature was maintained at 37.5C using a feedback heating system. Eyes were
384 coated with a thin layer of ophthalmic ointment to prevent desiccation. Anesthetized mice
385 were mounted in a stereotaxic apparatus. Blunt ear bars were placed in the external
386 auditory meatus to immobilize the head. A portion of the scalp overlying the two
387 hemispheres of the cortex (approximately 8mm by 6mm) was removed to expose the
388 underlying skull. After the skull was exposed it was dried and covered by a thin layer of
389 Vetbond. After the Vetbond dried (approximately 15 min), it provided a stable and solid
390 surface to affix the aluminum bracket with dental acrylic. The bracket was then affixed to
391 the skull and the margins sealed with Vetbond and dental acrylic to prevent infections.

392

393 **Imaging and signal extraction:** Imaging was performed using a resonant, two-photon
394 microscope (Neurolabware, Los Angeles, CA) controlled by Scanbox acquisition software
395 (Scanbox, Los Angeles, CA). The light source was a Coherent Chameleon Ultra II laser
396 (Coherent Inc, Santa Clara, CA) running at 920nm. The objective was an x16 water
397 immersion lens (Nikon, 0.8NA, 3mm working distance). The microscope frame rate was
398 15.6Hz (512 lines with a resonant mirror at 8kHz). We monitored locomotion using a
399 rotary, optical encoder (US Digital, Vancouver, WA) connected to the rotation axel. The

400 quadrature encoder was read by an Arduino board. We performed motion stabilization of
401 the images, followed by signal extraction and de-convolution to estimate the spiking of
402 neurons. The details of these methods are described elsewhere^{51,55,71}. We used the
403 average delay (387ms) measured in reverse correlation experiments to correct for the
404 stimulus-response delay in the data⁷¹.

405
406 **Visual stimulation:** We measured the responses of neural populations in mouse primary
407 visual cortex using two-photon imaging in tetO-GCaMP6s mice (Jackson Labs #024742).
408 The visual stimulus consisted of two conditions. In the first, unmasked condition, a
409 sinusoidal grating (50% contrast and a spatial frequency in the range 0.04–0.06cpd) was
410 presented with an orientation that changed linearly with time $\theta = \pi t/T$, and a period $T =$
411 10s. The spatial phase of the grating was updated every $T_\phi = 783$ msec by $\phi \leftarrow \phi +$
412 $\pi/2 + n$, where n was a random variable distributed uniformly $n \sim U(-\pi/8, +\pi/8)$. In other
413 words, the grating underwent a “noisy contrast reversal” as its orientation changed
414 continuously with time. This ensured that different spatial phases were present during
415 different cycles of the stimulus. The unmasked condition was displayed for 15 min for a
416 total of 90 cycles around the orientation domain. Immediately after, we added a vertical
417 mask. The vertical mask also underwent a noisy contrast reversal with a period of 717
418 msec. A TTL pulse was generated by an Arduino board at the beginning of each stimulus
419 cycle. The pulse was sampled by the microscope and time-stamped with the frame and
420 line number being scanned at that time.

421
422 The screen was calibrated using a Photo-Research (Chatsworth, CA) PR-650 spectro-
423 radiometer, and the result used to generate the appropriate gamma corrections for the
424 red, green and blue components via an nVidia Quadro K4000 graphics card. The contrast
425 of the stimulus was 99%. The center of the monitor was positioned with the center of the
426 receptive field population for the eye contralateral to the cortical hemisphere under
427 consideration. The locations of the receptive fields were estimated by an automated
428 process where localized, flickering checkerboards patches, appeared at randomized
429 locations within the screen. This measurement was performed at the beginning of each
430 imaging session to ensure the centering of receptive fields on the monitor.

431
432 **Data analysis:** We computed discriminability between two angles θ and φ as follows.
433 Consider the responses in the unmasked condition. Let $r_u^i(\theta)$ be the response of the
434 population in the i -th cycle to a given orientation and let $\mu_u(\theta)$ be the mean population
435 response across all trials. We define $d_u^i(\theta, \varphi) = d(\mu_u(\theta), r_u^i(\varphi))$. We then compute the
436 indices $F_\theta^i = (d_u^i(\theta, \theta) - d_u^i(\theta, \varphi)) / (d_u^i(\theta, \theta) + d_u^i(\theta, \varphi))$ and $F_\varphi^i = -(d_u^i(\varphi, \varphi) -$
437 $d_u^i(\varphi, \theta)) / (d_u^i(\varphi, \varphi) + d_u^i(\varphi, \theta))$. Finally, we compute $D_u(\theta, \varphi)$ as the difference in the
438 means of these distributions normalized by the average standard deviation. The same
439 calculation was applied for the masked condition.

440

441 **Fitting the geometric model to experimental data.** Note that the affine model in d
442 dimensions has a total of $d(d + 1)$ parameters. Our data consists of $s = 155$ equally
443 spaced samples (10 sec period at 15.5 fps) of the continuous curves $r_u(\theta)$ and $r_m(\theta)$.
444 Each sample provides d constraints on the transform. Thus, we must have $d(d + 1) \leq ds$
445 or $(d + 1) \leq s$ to ensure the problem is not under-constrained. We handled this constraint
446 by first projecting the data onto into R^3 using the first three components in the SVD
447 factorization of the data and subsequently fit the lower-dimensional embedding of the
448 curves (**Fig 6F**). Note that projection is a linear operation. Thus, if the data conformed to
449 an affine model in the high-dimensional space, it should also do it its low dimensional
450 embedding (no matter how much distortion we are imposing by the projection). This can
451 be easily shown using a basis set corresponding to the canonical form of the projection.
452 The reverse, of course, it not necessarily true.

453

454 **Analytic computation of threshold and bias:** In the simple geometric model of **Fig 6C**
455 it is possible to compute the threshold and bias. Consider a two-dimensional population
456 code for orientation in the unmasked condition $r_u(\theta) = (\cos \theta, \sin \theta)$, which is transformed
457 by a scaling and translation along the x-axis under the presence of the mask $r_m(\theta) =$
458 $(a + b \cos \theta, b \sin \theta)$. Then, the velocity of $r_m(\theta)$ is

459

$$460 \quad \| r'_m(\theta) \| = \frac{b(b + a \cos \theta)}{(a^2 + b^2 + 2ab \cos \theta)}$$

461

462 The threshold will be inversely proportional to the velocity $T_m(\theta) \propto 1/\| r'_m(\theta) \|$. Given a
463 population direction in the masked condition, which in the plane is simply given by an
464 angle φ , a decoder without knowledge of the white point shift will estimate the orientation
465 by measuring the angle θ formed between the population vector with respect to μ_u (**Fig**
466 **6C**), which a little geometry shows it is given by $\theta = \arctan((a + b \cos \varphi)/\sin \varphi)$. Thus,
467 the bias is given by

468

$$469 \quad bias(\varphi) = [\arctan((a + b \cos \varphi)/\sin \varphi) - \varphi] \bmod 2\pi.$$

470

471

472

473 Figure Legends

474

475 **Fig 1.** Measurement of population responses in masked and unmasked conditions. **(A)**
476 Structure of the visual stimulus. Each of the lines show a single period of the stimulus in
477 unmasked and masked conditions. **(B)** Samples of responses by individual neurons in
478 both conditions. Some cells responded very well in the unmasked conditions (top traces)
479 while others gave a weak response (bottom traces). Periods of locomotion enhanced the
480 overall responsivity of the population (shaded regions). Traces are plotted on a z-scored
481 scale (vertical bar = 10). Horizontal bar represents 1min of stimulation (or 6 periods of the
482 orientation cycle). **(C)** Example of cell responses in unmasked and masked conditions.
483 Each trace shows the response of a neuron over the stimulation cycle after correction for
484 neural delay, so they can be interpreted as a sweep of the orientation tuning curve of the
485 neuron. The dashed line indicates the orientation of the mask. Blue traces represent the
486 responses in the unmasked condition, while red traces represent responses in the
487 masked condition.

488

489 **Fig 2.** Characterization of responses in single neurons. **(A)** Anti-correlation between
490 responses of neurons in masked and unmasked conditions. The mean responses of cells
491 in the unmasked condition, μ_u , are anti-correlated with the responses in the masked
492 condition, μ_m . The inset shows the distribution of $\log_2(\mu_m/\mu_u)$. Cells at the extreme of this
493 distribution are termed grating (shaded green) and plaid (shaded pink) neurons. **(B)**
494 Average tuning of grating and plaid cells in unmasked and masked conditions. The
495 histograms show the distribution of the preferred orientation of the neurons in each case.
496 The red traces show the average tuning of neurons in each condition. The y-axis is
497 labeled by cell count (in black) or by the amplitude of the responses (in red).

498

499 **Fig 3.** Characterization of population responses. **(A)** Responses of a population of
500 neurons in the unmasked and masked conditions. Cells were ordered according to their
501 preferred orientation, thus resulting in a diagonal structure. The rows for these matrices
502 represent the population responses in the unmasked and masked conditions, $r_u(\theta)$ and
503 $r_m(\theta)$. These curves describe a close curve as θ describes one cycle. **(B)**. The intrinsic
504 geometry of the curves is captured by the cosine distances between the representation
505 of two orientations in the unmasked condition, $d_u(\theta, \varphi)$ (left panel), and unmasked
506 condition, $d_m(\theta, \varphi)$ (middle panel). The relative positions of the curves with respect to
507 each other is measured by the cosine distance between $r_u(\theta)$ and $r_m(\varphi)$, denoted by
508 $d_{um}(\theta, \varphi)$ (right panel). **(C)**. Schematic showing a response $r_u(\theta)$ along with the
509 covariance matrix and the direction of the first eigenvector, v_1 . **(D)** The first
510 eigenvector/eigenvalue captured about a third of the total energy of the noise and the
511 direction of the first eigenvector was very close to that of the response itself. The inset

512 shows the distribution of the correlation coefficient between $r_u(\theta)$ and v_1 for all our
513 experiments and orientations tested.

514

515 **Fig 4.** Multidimensional scaling of population responses in unmasked and masked
516 conditions. Each row shows two viewpoints of the result of one experiment. The curves
517 were obtained by doing multidimensional scaling simultaneously on the population
518 responses in unmasked and masked conditions into 3D space using the cosine distance
519 as a metric. The blue curve shows $r_u(\theta)$ and the red curve shows $r_m(\theta)$. The gray sphere
520 represents the origin, and colored spheres represent the beginning of the cycle. The
521 green arrows represent the shift in the white point between conditions. The stimuli
522 represent the patterns at different locations on the curves for the two conditions (blue
523 outline = unmasked condition; red outline = masked condition).

524

525 **Fig 5.** Basic geometric properties of population representations in unmasked and masked
526 conditions. **(A)** Shrinkage of the length of the curves by the introduction of the mask.
527 Scatterplot shows the lengths of the curves in unmasked (L_u) and masked (L_m)
528 conditions. Dashed line represents the unity line. Inset shows the distribution of L_m/L_u
529 across all experiments. **(B)** Distribution of white-point shift (Δ) across all experiments. **(C)**
530 Measurements of the angle between $r_m(\theta)$ and the plane span $\{r_u(\theta), r_u(\pi/2)\}$ across all
531 experiments. Solid line represents the mean, while the shaded area represents ± 2 SEM.

532

533 **Fig 6.** A geometric model of masking. **(A)** Discriminability (d' -prime) between the
534 representation of two orientations in unmasked (left panels) and masked (middle panels)
535 conditions. The top panels show results for one experiment, while the ones at the bottom
536 show the average across all our experiments. Iso-performance contour for the single
537 experiment is shown at $d' = 4$. The iso-performance contours for the average behavior
538 is shown at levels of $d' = 4, 6, 8$. The widening in the iso-performance contours in the
539 masked condition reflect an increase in thresholds near the mask (which has an
540 orientation of 90 deg). This is best shown in the panels on the right, which show the
541 dependence of thresholds in masked (red) and unmasked (blue) conditions as a function
542 of a base angle. In the average data the shaded areas represent ± 2 SEM. **(B)** Mutual
543 distances and bias. Top panels show the mutual distance between orientations across
544 masked and unmasked representations (d_{um}) and the expected bias from a decoder
545 based on the distances. The non-diagonal structure of d_{um} is more evident in the average
546 data (bottom left panel), showing the locations of the minima of the main diagonal (white,
547 dashed line). Bottom right panel shows the average bias across all our experiments.
548 Shaded areas represent 2 SEM. **(C)** Two-dimensional geometric model of population
549 coding⁶². The model assumes $r_u(\theta)$ and $r_m(\theta)$ are two circles in the plane. The
550 displacement of their centers (white points) induce changes in the mutual distances
551 inducing corresponding changes in threshold (middle panel) and bias (right panel). **(D)**

552 The model can be extended by allowing displacement of the curves along a third
553 dimension. (E) Two viewpoints of the same population activity in (D) but now normalized
554 to yield $\hat{r}_u(\theta)$ and $\hat{r}_m(\theta)$. (F) Fits of an affine model to low dimensional representations of
555 $\hat{r}_u(\theta)$ and $\hat{r}_m(\theta)$ in four different experiments. In each case, $\hat{r}_u(\theta)$ represents the
556 population response in the unmasked condition (blue), $\hat{r}_m(\theta)$ represents the population
557 response in the masked condition (red), and $\tilde{r}_m(\theta)$ is the best fit to the response in the
558 masked condition by means of an affine transform.

559
560

561 Acknowledgements

562

563 I thank Elaine Tring for help in surgical preparation and data collection.

564

565 References

- 566 1 Hubel, D. H. & Wiesel, T. N. Receptive fields of single neurones in the cat's striate cortex. *J*
567 *Physiol* **148**, 574-591 (1959).
- 568 2 Hubel, D. H. & Wiesel, T. N. Receptive fields and functional architecture of monkey striate
569 cortex. *J Physiol* **195**, 215-243 (1968).
- 570 3 Spillmann, L. Receptive fields of visual neurons: the early years. *Perception* **43**, 1145-1176,
571 doi:10.1068/p7721 (2014).
- 572 4 Fitzpatrick, D. Seeing beyond the receptive field in primary visual cortex. *Curr Opin Neurobiol* **10**,
573 438-443 (2000).
- 574 5 Allman, J., Miezin, F. & McGuinness, E. Stimulus specific responses from beyond the classical
575 receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons.
576 *Annu Rev Neurosci* **8**, 407-430, doi:10.1146/annurev.ne.08.030185.002203 (1985).
- 577 6 Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat Rev Neurosci*
578 **13**, 51-62, doi:10.1038/nrn3136 (2011).
- 579 7 Morrone, M. C., Burr, D. C. & Speed, H. D. Cross-orientation inhibition in cat is GABA mediated.
580 *Exp Brain Res* **67**, 635-644 (1987).
- 581 8 Morrone, M. C., Burr, D. C. & Maffei, L. Functional implications of cross-orientation inhibition of
582 cortical visual cells. I. Neurophysiological evidence. *Proc R Soc Lond B Biol Sci* **216**, 335-354,
583 doi:10.1098/rspb.1982.0078 (1982).
- 584 9 DeAngelis, G. C., Robson, J. G., Ohzawa, I. & Freeman, R. D. Organization of suppression in
585 receptive fields of neurons in cat visual cortex. *J Neurophysiol* **68**, 144-163,
586 doi:10.1152/jn.1992.68.1.144 (1992).
- 587 10 Sillito, A. M. & Jones, H. E. Context-dependent interactions and visual processing in V1. *J Physiol*
588 *Paris* **90**, 205-209 (1996).
- 589 11 Carandini, M., Movshon, J. A. & Ferster, D. Pattern adaptation and cross-orientation interactions
590 in the primary visual cortex. *Neuropharmacology* **37**, 501-511 (1998).
- 591 12 Sengpiel, F., Baddeley, R. J., Freeman, T. C., Harrad, R. & Blakemore, C. Different mechanisms
592 underlie three inhibitory phenomena in cat area 17. *Vision Res* **38**, 2067-2080 (1998).
- 593 13 Smith, M. A., Bair, W. & Movshon, J. A. Dynamics of suppression in macaque primary visual
594 cortex. *J Neurosci* **26**, 4826-4834, doi:10.1523/JNEUROSCI.5542-06.2006 (2006).

- 595 14 Meese, T. S., Summers, R. J., Holmes, D. J. & Wallis, S. A. Contextual modulation involves
596 suppression and facilitation from the center and the surround. *J Vis* **7**, 7, doi:10.1167/7.4.7
597 (2007).
- 598 15 Walker, G. A., Ohzawa, I. & Freeman, R. D. Asymmetric suppression outside the classical
599 receptive field of the visual cortex. *J Neurosci* **19**, 10536-10553 (1999).
- 600 16 Cavanaugh, J. R., Bair, W. & Movshon, J. A. Selectivity and spatial distribution of signals from the
601 receptive field surround in macaque V1 neurons. *J Neurophysiol* **88**, 2547-2556,
602 doi:10.1152/jn.00693.2001 (2002).
- 603 17 Webb, B. S., Tinsley, C. J., Barraclough, N. E., Parker, A. & Derrington, A. M. Gain control from
604 beyond the classical receptive field in primate primary visual cortex. *Vis Neurosci* **20**, 221-230
605 (2003).
- 606 18 Angelucci, A. & Bullier, J. Reaching beyond the classical receptive field of V1 neurons: horizontal
607 or feedback axons? *J Physiol Paris* **97**, 141-154, doi:10.1016/j.jphysparis.2003.09.001 (2003).
- 608 19 Angelucci, A. & Bressloff, P. C. Contribution of feedforward, lateral and feedback connections to
609 the classical receptive field center and extra-classical receptive field surround of primate V1
610 neurons. *Prog Brain Res* **154**, 93-120, doi:10.1016/S0079-6123(06)54005-1 (2006).
- 611 20 Smith, M. A. Surround suppression in the early visual system. *J Neurosci* **26**, 3624-3625,
612 doi:10.1523/jneurosci.0236-06.2006 (2006).
- 613 21 Samonds, J. M., Feese, B. D., Lee, T. S. & Kuhlman, S. J. Nonuniform surround suppression of
614 visual responses in mouse V1. *J Neurophysiol* **118**, 3282-3292, doi:10.1152/jn.00172.2017
615 (2017).
- 616 22 Gilbert, C. D., Das, A., Ito, M., Kapadia, M. & Westheimer, G. Spatial integration and cortical
617 dynamics. *Proc Natl Acad Sci U S A* **93**, 615-622 (1996).
- 618 23 Kapadia, M. K., Westheimer, G. & Gilbert, C. D. Dynamics of spatial summation in primary visual
619 cortex of alert monkeys. *Proc Natl Acad Sci U S A* **96**, 12073-12078 (1999).
- 620 24 Li, W., Piech, V. & Gilbert, C. D. Learning to link visual contours. *Neuron* **57**, 442-451,
621 doi:10.1016/j.neuron.2007.12.011 (2008).
- 622 25 Li, C. Y. & Li, W. Extensive integration field beyond the classical receptive field of cat's striate
623 cortical neurons--classification and tuning properties. *Vision Res* **34**, 2337-2355 (1994).
- 624 26 Vaiceliunaite, A., Eriskien, S., Franzen, F., Katzner, S. & Busse, L. Spatial integration in mouse
625 primary visual cortex. *Journal of Neurophysiology* **110**, 964-972, doi:10.1152/jn.00138.2013
626 (2013).
- 627 27 Field, D. J., Hayes, A. & Hess, R. F. Contour integration by the human visual system: evidence for
628 a local "association field". *Vision Res* **33**, 173-193 (1993).
- 629 28 Hess, R. & Field, D. Integration of contours: new insights. *Trends Cogn Sci* **3**, 480-486 (1999).
- 630 29 Geisler, W. S., Perry, J. S., Super, B. J. & Gallogly, D. P. Edge co-occurrence in natural images
631 predicts contour grouping performance. *Vision Res* **41**, 711-724 (2001).
- 632 30 Ledgeway, T., Hess, R. F. & Geisler, W. S. Grouping local orientation and direction signals to
633 extract spatial contours: empirical tests of "association field" models of contour integration.
634 *Vision Res* **45**, 2511-2522, doi:10.1016/j.visres.2005.04.002 (2005).
- 635 31 Shaashua, A. & Ullman, S. Structural Saliency: The Detection of Globally Salient Structures Using
636 Locally Connected Network. *ICCV88*, 321-327 (1988).
- 637 32 Alter, T. D. & Basri, R. Extracting Salient Curves from Images: An Analysis of the Saliency.
638 *Network IJCV* **27(1)**, 51-69 (1998).
- 639 33 Guy, G. & Medioni, G. G. Inferring Global Perceptual Contours from Local Features. *IJCV* 113-133
640 (1986).
- 641 34 Lindebaum, M. & Berengolts, A. A Probabilistic Interpretation of the Saliency *ECCV2000* (2000).

- 642 35 Hummel, R. A. & Zucker, S. W. On the foundations of relaxation labeling processes. *IEEE Trans*
643 *Pattern Anal Mach Intell* **5**, 267-287 (1983).
- 644 36 Haralick, R. M. & Shapiro, L. G. The consistent labeling problem: part I. *IEEE Trans Pattern Anal*
645 *Mach Intell* **1**, 173-184 (1979).
- 646 37 Haralick, R. M. & Shapiro, L. G. The Consistent Labeling Problem: Part II. *IEEE Trans Pattern Anal*
647 *Mach Intell* **2**, 193-203 (1980).
- 648 38 Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a
649 sparse code for natural images. *Nature* **381**, 607-609, doi:10.1038/381607a0 (1996).
- 650 39 Olshausen, B. A. & Field, D. J. Natural image statistics and efficient coding. *Network* **7**, 333-339,
651 doi:10.1088/0954-898x/7/2/014 (1996).
- 652 40 Olshausen, B. A. & Field, D. J. Sparse coding with an overcomplete basis set: a strategy employed
653 by V1? *Vision Res* **37**, 3311-3325 (1997).
- 654 41 Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Curr Opin Neurobiol* **14**, 481-487,
655 doi:10.1016/j.conb.2004.07.007 (2004).
- 656 42 Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annu Rev*
657 *Neurosci* **24**, 1193-1216, doi:10.1146/annurev.neuro.24.1.1193 (2001).
- 658 43 Bell, A. J. & Sejnowski, T. J. An information-maximization approach to blind separation and blind
659 deconvolution. *Neural Comput* **7**, 1129-1159 (1995).
- 660 44 Bell, A. J. & Sejnowski, T. J. The "independent components" of natural scenes are edge filters.
661 *Vision Res* **37**, 3327-3338 (1997).
- 662 45 Barlow, H. B. & Foldiak, P. in *The computing neuron* (eds R. Durbin, C. Miall, & G. Mitchison)
663 54-72 (Addison-Wesley Longman Publishing Co, 1989).
- 664 46 Heeger, D. J. Normalization of cell responses in cat striate cortex. *Vis Neurosci* **9**, 181-197 (1992).
- 665 47 Heeger, D. J., Simoncelli, E. P. & Movshon, J. A. Computational models of cortical visual
666 processing. *Proc Natl Acad Sci U S A* **93**, 623-627 (1996).
- 667 48 Tring, E. & Ringach, D. L. On the Subspace Invariance of Population Responses. *Neurons,*
668 *Behavior, Data Analysis, and Theory* (2018).
- 669 49 Snow, M., Coen-Cagli, R. & Schwartz, O. Adaptation in the visual cortex: a case for probing
670 neuronal populations with natural stimuli [version 1; referees: 4 approved]. **6**,
671 doi:10.12688/f1000research.11154.1 (2017).
- 672 50 Kalatsky, V. A. & Stryker, M. P. New paradigm for optical imaging: temporally encoded maps of
673 intrinsic signal. *Neuron* **38**, 529-545 (2003).
- 674 51 Berens, P. *et al.* Community-based benchmarking improves spike rate inference from two-
675 photon calcium imaging data. *PLoS Comput Biol* **14**, e1006157,
676 doi:10.1371/journal.pcbi.1006157 (2018).
- 677 52 Niell, C. M. & Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual
678 cortex. *Neuron* **65**, 472-479 (2010).
- 679 53 Carandini, M., Heeger, D. J. & Movshon, J. A. Linearity and normalization in simple cells of the
680 macaque primary visual cortex. *J Neurosci* **17**, 8621-8644 (1997).
- 681 54 Montijn, J. S., Meijer, G. T., Lansink, C. S. & Pennartz, C. M. Population-Level Neural Codes Are
682 Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. *Cell Rep* **16**,
683 2486-2498, doi:10.1016/j.celrep.2016.07.065 (2016).
- 684 55 Mineault, P. J., Tring, E., Trachtenberg, J. T. & Ringach, D. L. Enhanced Spatial Resolution During
685 Locomotion and Heightened Attention in Mouse Primary Visual Cortex. *J Neurosci* **36**, 6382-
686 6392, doi:10.1523/JNEUROSCI.0430-16.2016 (2016).
- 687 56 Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nat Neurosci* **17**,
688 858-865, doi:10.1038/nn.3711 (2014).

- 689 57 Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. Neuronal population coding of movement
690 direction. *Science* **233**, 1416-1419 (1986).
- 691 58 Ringach, D. L. Population coding under normalization. *Vision Res* **50**, 2223-2232,
692 doi:10.1016/j.visres.2009.12.007 (2010).
- 693 59 Kang, K., Shapley, R. M. & Sompolinsky, H. Information tuning of populations of neurons in
694 primary visual cortex. *J Neurosci* **24**, 3726-3735, doi:10.1523/JNEUROSCI.4272-03.2004 (2004).
- 695 60 Busse, L., Wade, A. R. & Carandini, M. Representation of concurrent stimuli by population
696 activity in visual cortex. *Neuron* **64**, 931-942, doi:10.1016/j.neuron.2009.11.004 (2009).
- 697 61 Seung, H. S. & Sompolinsky, H. Simple models for reading neuronal population codes. *Proc Natl
698 Acad Sci U S A* **90**, 10749-10753 (1993).
- 699 62 Clifford, C. W., Wenderoth, P. & Spehar, B. A functional angle on some after-effects in cortical
700 vision. *Proceedings. Biological sciences* **267**, 1705-1710, doi:10.1098/rspb.2000.1198 (2000).
- 701 63 Wu, W., Klassen, E. & Srivastava, A. Spherical Regression Models Using Projective Linear
702 Transformations AU - Rosenthal, Michael. *J Am Stat Assoc* **109**, 1615-1624,
703 doi:10.1080/01621459.2014.892881 (2014).
- 704 64 Adelson, E. H. & Movshon, J. A. Phenomenal coherence of moving visual patterns. *Nature* **300**,
705 523-525 (1982).
- 706 65 Movshon, J. A., Adelson, E. H., Gizzi, M. S. & Newsome, W. T. in *Scientiarum Scripta* (eds C.
707 Chagas, R. Gattass, & C. Gross) 117-151 (MIT Press, 1985).
- 708 66 Wang, H. X. & Movshon, J. A. Properties of pattern and component direction-selective cells in
709 area MT of the macaque. *J Neurophysiol* **115**, 2705-2720, doi:10.1152/jn.00639.2014 (2016).
- 710 67 Palagina, G., Meyer, J. F. & Smirnakis, S. M. Complex Visual Motion Representation in Mouse
711 Area V1. *J Neurosci* **37**, 164-183, doi:10.1523/JNEUROSCI.0997-16.2017 (2017).
- 712 68 Ben-Yishai, R., Bar-Or, R. L. & Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proc
713 Natl Acad Sci U S A* **92**, 3844-3848 (1995).
- 714 69 Benucci, A., Saleem, A. B. & Carandini, M. Adaptation maintains population homeostasis in
715 primary visual cortex. *Nat Neurosci* **16**, 724-729, doi:10.1038/nn.3382 (2013).
- 716 70 Singh, G. *et al.* Topological analysis of population activity in visual cortex. *J Vision* **8** (2008).
- 717 71 Ringach, D. L. *et al.* Spatial clustering of tuning in mouse primary visual cortex. *Nat Commun* **7**,
718 12270, doi:10.1038/ncomms12270 (2016).

719

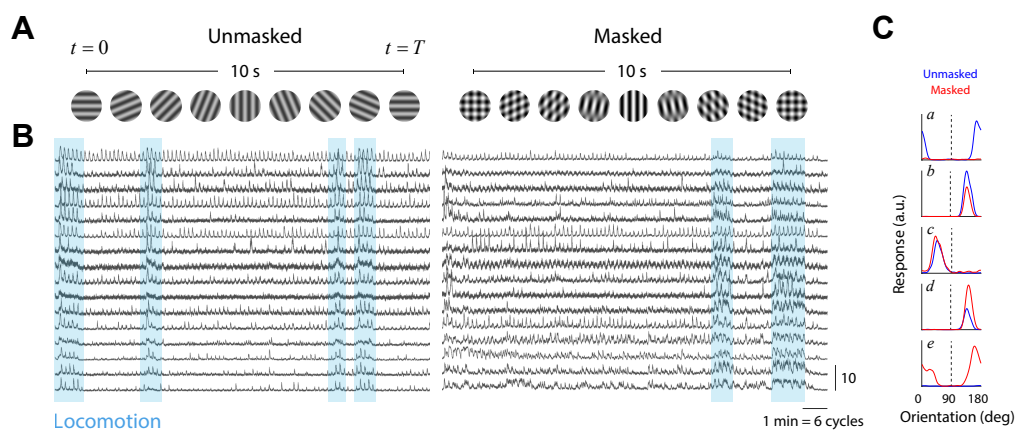


Fig 1

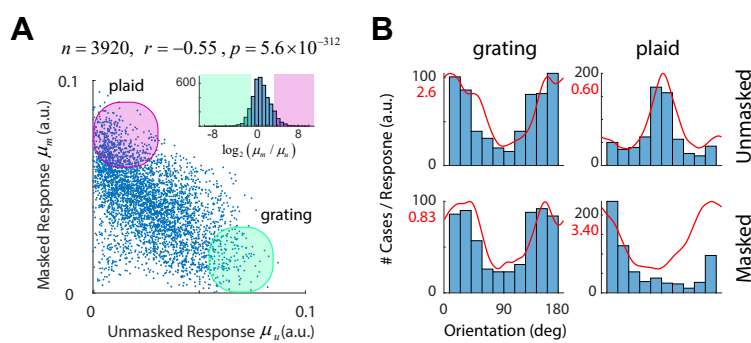


Fig 2

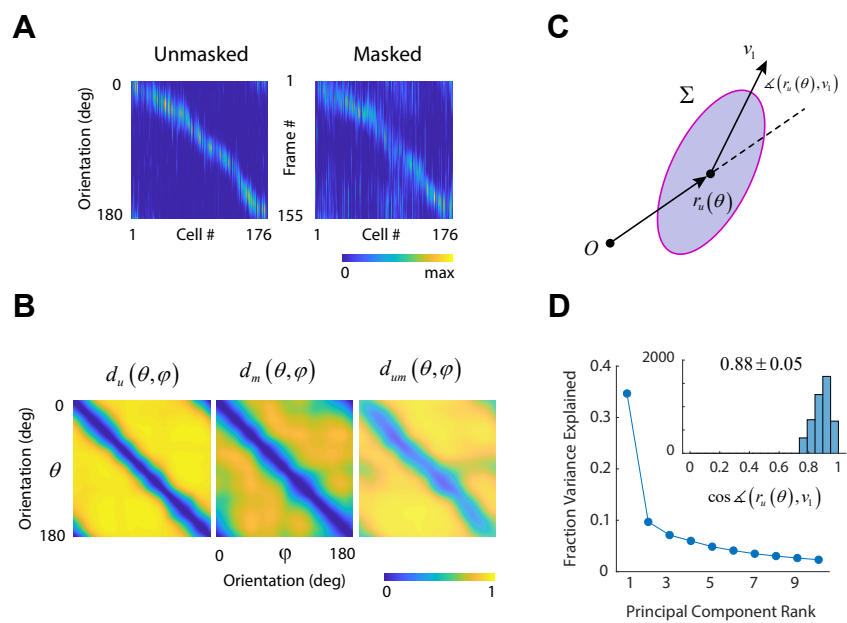


Fig 3




Fig 4




Fig 5

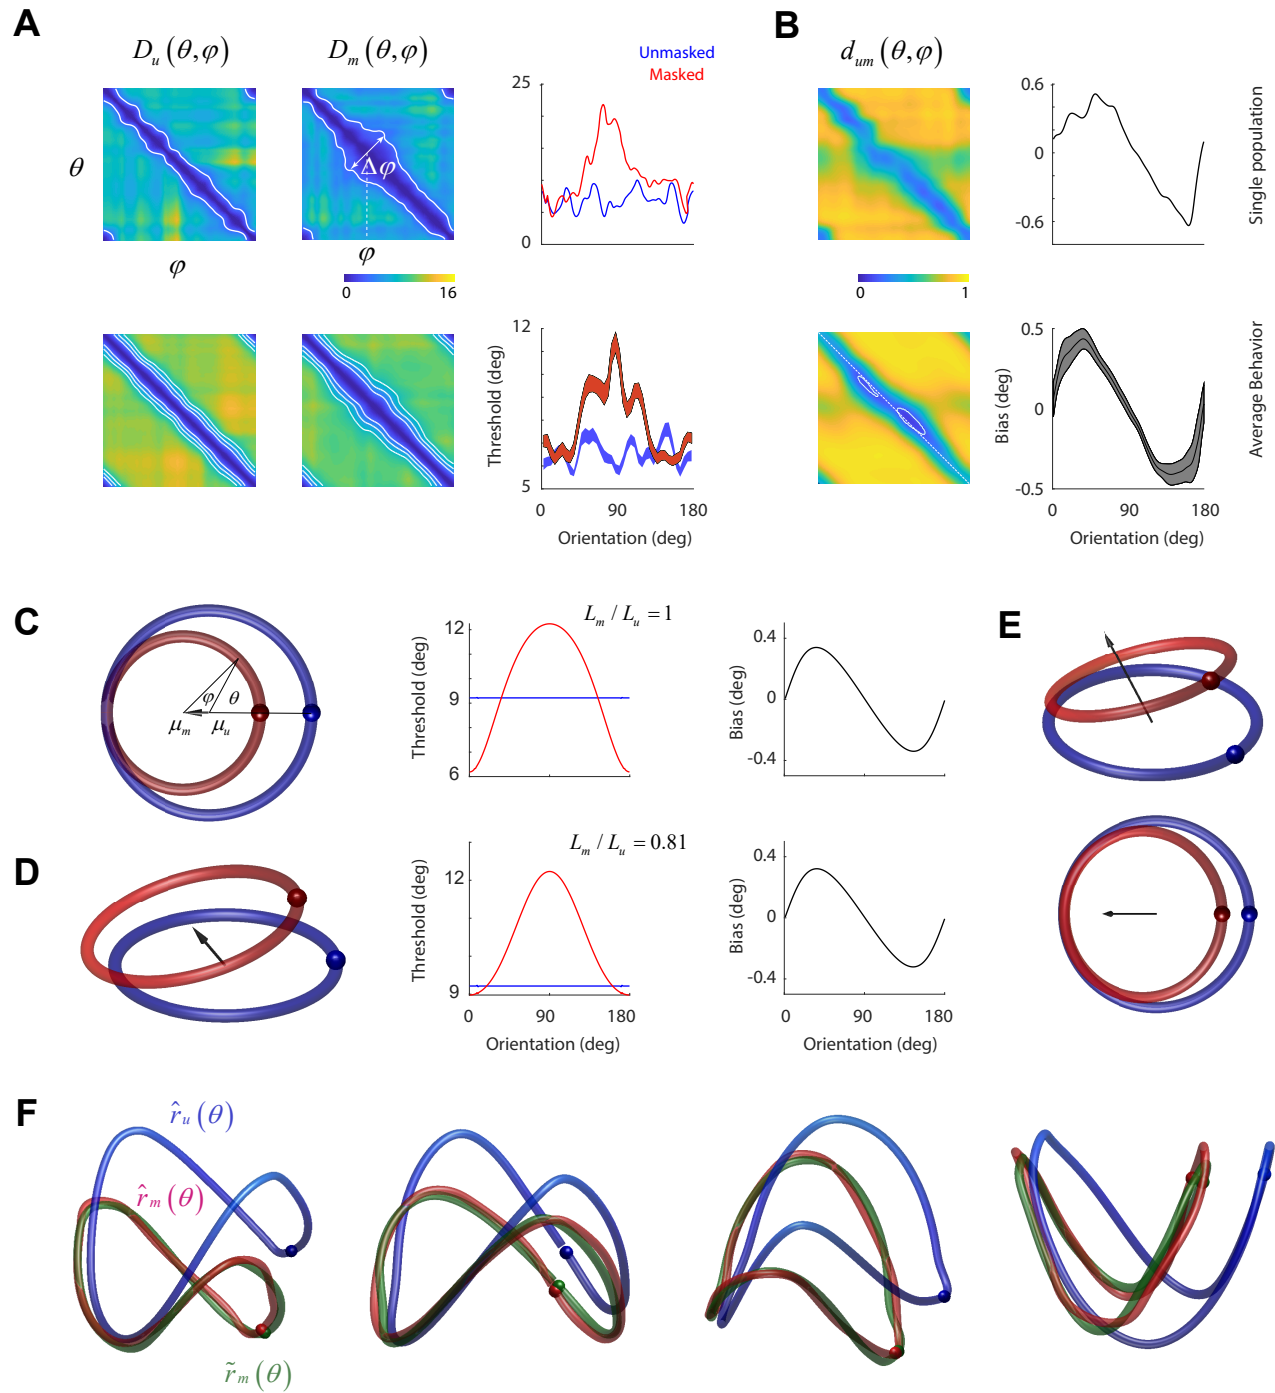


Fig 6