

Not optimal, just noisy: the geometry of correlated variability leads to highly suboptimal sensory coding

Jesse A. Livezey^{*1,2}, Pratik S. Sachdeva³, Maximilian E. Dougherty⁴, Mathew T. Summers⁵, and Kristofer E. Bouchard^{†‡1,2,6}

¹*Biological Systems and Engineering Division, Lawrence Berkeley National Laboratory, Berkeley, 94720, CA, USA*

²*Redwood Center for Theoretical Neuroscience, University of California, Berkeley, 94720, CA, USA*

³*Department of Physics, University of California, Berkeley, 94720, CA, USA*

⁴*Department of Neurology, University of California, San Francisco, San Francisco, 94143, CA, USA*

⁵*Department of Molecular and Cell Biology, University of California, Berkeley, Berkeley, CA 94720, USA*

⁶*Computational Resources Division, Lawrence Berkeley National Laboratory, Berkeley, 94720, CA, USA*

Abstract

The brain represents the world through the activity of neural populations. Correlated variability across simultaneously recorded neurons (noise correlations) has been observed across cortical areas and experimental paradigms. Many studies have shown that correlated variability improves stimulus coding compared to a null model with no correlations. However, such results do not shed light on whether neural populations' correlated variability achieves *optimal* coding. Here, we assess optimality of noise correlations in diverse datasets by developing two novel null models each with a unique biological interpretation: a uniform correlations null model and a factor analysis null model. We show that across datasets, the correlated variability in neural populations leads to highly suboptimal coding performance according to these null models. We demonstrate that biological constraints prevent many subsets of the neural populations from achieving optimality according to these null models, and that subselecting based on biological criteria leaves coding performance suboptimal. Finally, we show that the optimal subpopulation is exponentially small as a function of neural dimensionality. Together, these results show that the geometry of correlated variability leads to highly suboptimal sensory coding.

Introduction

The brain represents the world through the coordinated firing of neural populations. For instance, neural populations in early sensory areas are thought to transform the features of stimuli and transmit them to downstream cortical areas. Indeed, many studies of sensory areas seek to analyze what sensory features are transmitted in the brain and with what fidelity. Understanding population neural activity necessitates analyzing the joint activity of many neural units, beyond single-neuron analysis. Normative theories, which formalize optimality criteria, are powerful tools in these analyses, as they can establish principles for explaining features of experimentally observed neural activity at the population level. Therefore, it is important to develop methods for quantitatively assessing normative theories based on the features observed in neural data. One prominent feature of neural activity is variability: neural recordings exhibit trial-to-trial fluctuations in response to the same stimulus. From a normative perspective, the geometry of variability in neural

*jesse.livezey@gmail.com

†kebouchard@lbl.gov

‡corresponding author

36 activity impacts how optimally a population of neurons can encode stimuli [1, 2]. However, the optimality
37 of correlated variability has not been assessed.

38 Many studies have found pairwise correlations in the trial-to-trial variability of the firing rates of simul-
39 taneously recorded neurons, often called correlated variability or noise correlations [3–9]. The correlated
40 variability observed in experimental studies typically depends on the tuning and stimuli [10–12]. For ex-
41 ample, **Figure 1a, b** shows the single-trial variability in Ca^{2+} responses ($\Delta F/F$) for two simultaneously
42 recorded mouse retinal ganglion cells (RGCs) in response to drifting bars. The RGCs' correlated vari-
43 ability for a single stimulus is shown in **Figure 1e**. Although correlated variability is typically considered
44 in simultaneous single neuron electrophysiology measurements, it has been observed in calcium imaging
45 recordings [13] and larger scale measurements such as electrocorticography recordings [9]. Correlated vari-
46 ability has many possible biological sources in neural populations (see **Supplementary Fig. 1**) [6, 7, 14–
47 18], which the nervous system may be able to modify. Understanding the impact of correlated variability on
48 population coding is important for revealing the principles governing neural computation [1, 2, 4].

49 Correlated variability impacts the fidelity of a neural code when discriminating stimuli. Theoretical and
50 computational studies have determined how the interplay between correlated variability and tuning proper-
51 ties affect population coding [2, 8, 15, 19–23]. **Figure 1c** shows the mean response curve (black line, defined
52 by the mean firing rate of the neurons in response to various stimuli) from two hypothetical simultaneously
53 recorded neurons across a range of stimulus values (3 neighboring stimuli are demarcated with black dots).
54 From a geometric perspective, if the correlated variability has low variance (**Fig. 1c**, blue ellipse) along the
55 mean stimulus response curve (**Fig. 1c**, black line), the impact on coding will be less detrimental than having
56 high variance (**Fig. 1c**, orange ellipse) along the stimulus response curve. This is because the trial-by-trial
57 fluctuation (blue ellipse) in response to the central stimulus (large black dot) will minimally overlap with the
58 response to the nearby stimuli (small black dots). In early sensory areas, such as retina and primary visual
59 cortex, studies have found that correlated variability enhances population coding [6, 16, 24–28]. Outside of
60 early sensory areas, both the structure of correlated variability and its impact on coding is heterogeneous [12,
61 29, 30]. Brain states can change correlated variability and therefore its effect on population coding [31–33].
62 These studies leave open the possibility that the correlated variability is optimal for coding in sensory areas,
63 which has not been evaluated.

64 The impact of correlated variability on neural coding is typically assessed by comparing the linear Fisher
65 information (LFI) of the experimentally observed correlations to the distribution of LFI under the *shuffle null*
66 *model*, a null distribution with the same per-neuron variability, but no correlations across neurons. LFI quan-
67 tifies how accurately neural population activity can be used to distinguish two stimuli. Many previous studies
68 have shown, by using the shuffle null model, that the geometry of correlated variability can benefit neural
69 coding. However, comparing the experimentally observed correlated variability with the zero correlation
70 version is only one relevant comparison for determining optimality; there are potentially other geometries
71 which are not captured by the shuffle null model. In principle the brain's correlated variability could have
72 produced better (or worse) coding properties. Furthermore, it is unclear whether zero-correlation population
73 activity is the only reasonable null distribution given biological processes such as learning, highlighting the
74 importance of developing tailored null models [34]. Testing normative theories of stimulus coding in neu-
75 ral datasets requires understanding whether the geometry of experimental correlated variability is optimal,
76 however methods for testing the optimality of correlated variability are currently lacking.

77 In order to test the optimality of correlated variability in experimentally observed neural responses, we
78 developed two null models. The *uniform correlation null model* and the *factor analysis null model* each
79 define a null distribution of correlated variability and have a particular biological interpretation. Using these
80 null models, we test the optimality of neural coding in newly acquired data recorded from retinal ganglion
81 cells (RGCs, Retina), previously recorded neurons in primary visual cortex (V1), and newly acquired ECoG
82 electrodes on primary auditory cortex (PAC) (**Fig. 1d-l**). These datasets span neural areas and recording
83 modalities used in many previous studies. Our main finding is that the experimentally observed geometry

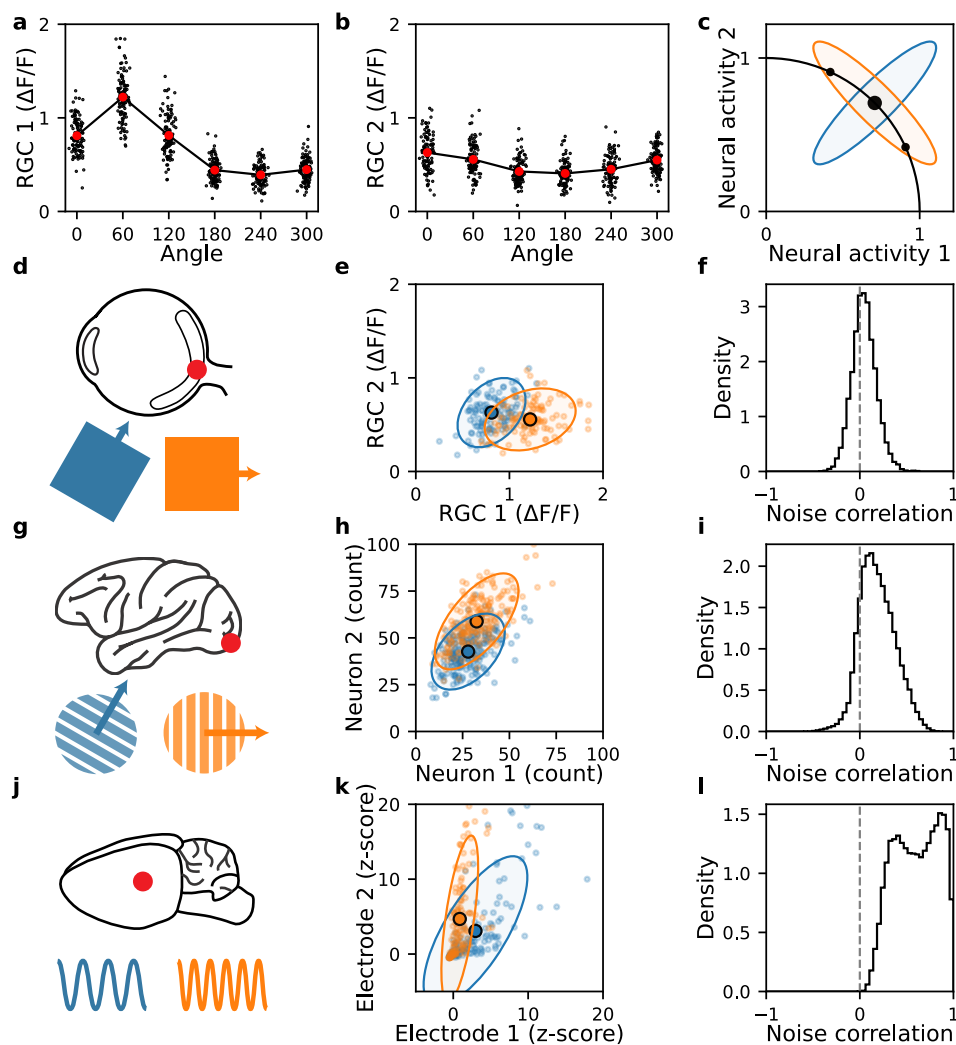


Figure 1: Correlated variability is a pervasive neural phenomenon. **a, b.** Mean activity as a function of bar angle (larger open circles) and trial-to-trial variability (small dots, small angle offsets for visualization) for angle 0 (corresponds to the blue dots for Neuron 1 and Neuron 2, respectively, in **d**). **c.** Illustration of mean stimulus response curve (black line), less detrimental correlated variability (blue ellipse), and more detrimental correlated variability (orange ellipse) for two model neurons. The large black dot is the mean stimulus response corresponding to the covariances. The small black dots are the mean responses for neighboring stimuli. **d-l.** Each row refers to a different experimental dataset, while columns refer to an aspect of the dataset. **d-f.** Calcium imaging recordings from mouse retinal ganglion cells in response to drifting bars. **g-i.** Single-unit spike counts recorded from primary visual cortex of macaque monkey in response to drifting gratings. **j-l.** Micro-electrocorticography recordings (z -scored $H\gamma$ response) from rat primary auditory cortex in response to tone pips at varying frequencies. First column (**d, g, j**) depicts the recording region and stimulus for each dataset. Second column (**e, h, k**) shows the activity of two random RGCs/neurons/electrodes in the population to two neighboring stimuli. Individual points denote the unit activity on individual trials, while covariance ellipses denote the noise covariance ellipse at 2 standard deviations. Third column (**f, i, l**) plots the distribution of pairwise noise correlations, calculated for each pair of units across stimuli.

84 of correlated variability leads to highly suboptimal coding across all datasets and both null models. Fur-
85 thermore, the degree of suboptimality worsens as a function of the number of neural units considered in the
86 neural population. We find that for a large fraction of subsamples of the recorded units, achieving optimality
87 would push the neural responses into regimes that violate biological constraints. However, even when neural
88 units are subsampled to optimize for biological criteria, they remain highly suboptimal. Finally, direct selec-
89 tion of optimal subsamples shows that the optimal population is exponentially small as a function of neural
90 dimensionality. Our results demonstrate that the traditional null model of correlated variability cannot be
91 used to assess the optimality of neural data, and that biological constraints limit the ability of neural activity
92 to achieve optimal correlated variability as defined by our null models. Together, our results show that the
93 geometry of correlated neural variability leads to highly suboptimal sensory coding.

94 Results

95 In order to assess the optimality of correlated variability in neural populations, we used three neural datasets
96 which span animal models, sensory recordings areas, and recording modalities (**Fig. 1**). The newly recorded
97 retina dataset is calcium imaging recordings in mouse retinal ganglion cells (RGCs) (**Fig. 1d-f**). The stimuli
98 are drifting bars at 6 angles with each stimuli being presented 114 times. The previously recorded V1
99 dataset is spike sorted, single unit electrophysiology recordings in macaque V1 (**Fig. 1g-i**) [35]. The stimuli
100 are drifting gratings at 12 angles with each stimuli being presented 200 times. The newly recorded primary
101 auditory cortex (PAC) dataset is high gamma amplitude from μ ECoG recordings in rat primary auditory
102 cortex (**Fig. 1j-l**). The stimuli are tone pips at 30 different frequencies with each stimuli being presented 60
103 times. We will refer to RGCs/neurons/electrodes as neural units. The neural units have various levels of
104 pairwise noise correlations, ρ , across datasets (**Fig. 1f, i, l**), which is a key quantity for analyzing correlated
105 variability. See Methods for more details on dataset recording and preprocessing.

106 Methods for assessing the optimality of neural codes

An abundance of work has aimed to assess whether observed correlated variability is beneficial or detrimen-
tal for neural coding [6, 9–12, 16, 19, 24–33]. These studies often quantify the discriminability or fidelity of
a neural code with the linear Fisher information (LFI, see Section) [36], which is a measure of how well the
neural activity could be used to discriminate between different stimuli. The LFI is a function of the stimulus,
 s , the stimulus-derivative of the mean neural activity, $\frac{d\mathbf{f}(s)}{ds}$, and the variability of the neural activity around
the mean, $\Sigma(s)$, and can be written as:

$$\mathcal{I}(s) = \frac{d\mathbf{f}(s)^T}{ds} \Sigma(s)^{-1} \frac{d\mathbf{f}(s)}{ds}. \quad (1)$$

107 Typically, the impact of correlated variability is assessed by comparing the experimentally observed LFI to a
108 distribution of LFIs generated from the shuffle null model. Trial-shuffling the data will produce a distribution
109 over covariance matrices ($\Sigma(s)$) where the pairwise correlations are all centered near zero (**Fig. 2a**, observed
110 covariance is filled, corresponding shuffle covariance is dashed). However, the shuffle null model does not
111 compare the observed correlations to a broad range of potential non-zero correlations. In principle, neural
112 circuits can support a range of covariance structures with significant nonzero pairwise correlations, many of
113 which can produce higher LFI than having zero correlations. In this case, using the shuffle null model would
114 overestimate the level of optimality in neural data, and therefore cannot be used to assess the optimality of
115 the experimentally observed correlations. To our knowledge, the optimality of correlated variability has not
116 been evaluated on neural data before.

117 In order to assess optimality, the null model should be chosen to adequately span achievable covariance
118 structures. Defining achievable may depend on the experimental context, including the types of neurons

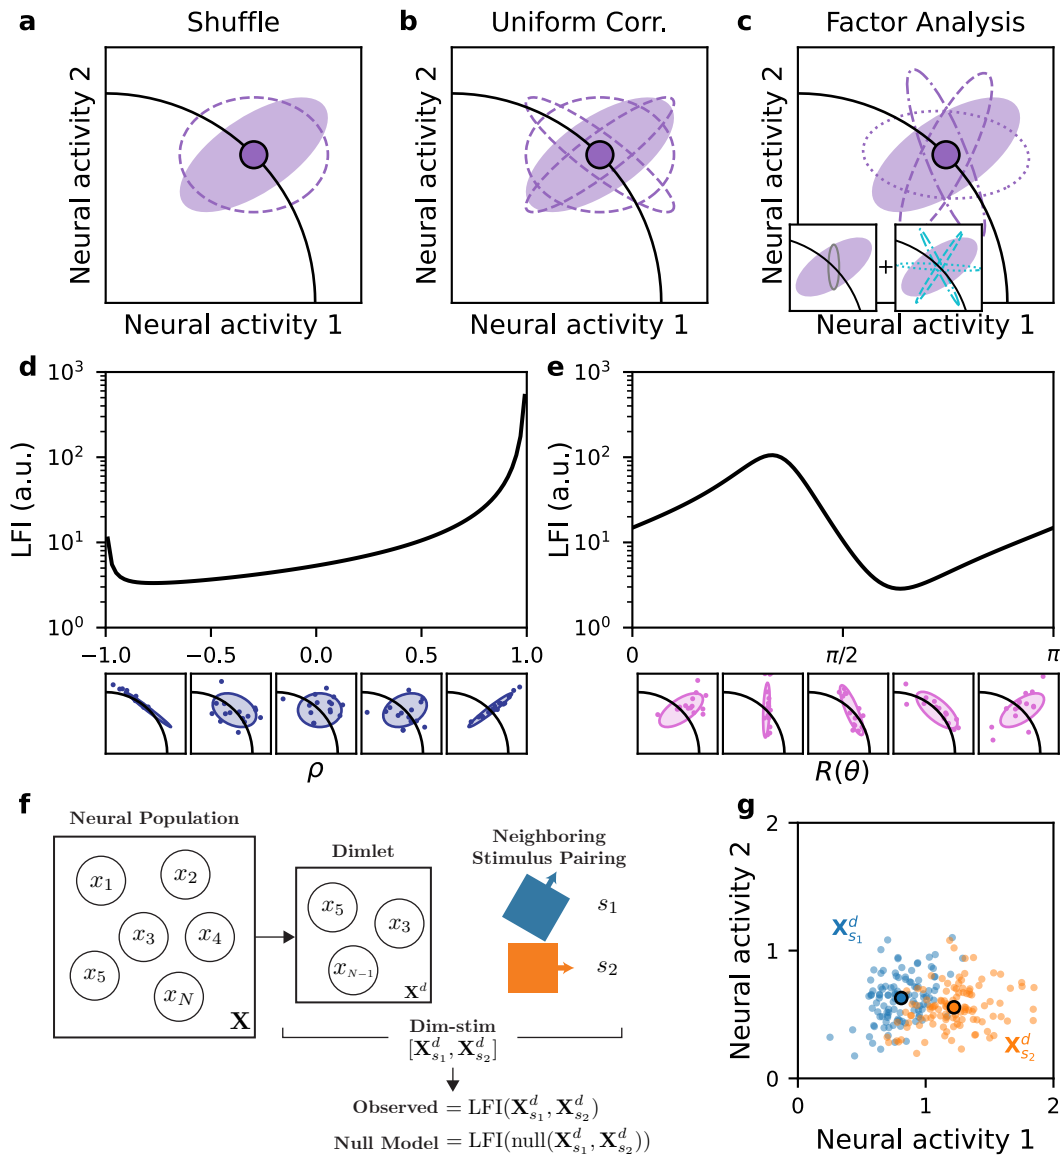


Figure 2: Methods for assessing the optimality of neural codes. **a-c.** Null models of correlated variability. Solid, purple ellipses denote the trial-to-trial variability observed about the mean stimulus activity (solid point). Samples from the null models are depicted by dashed ellipses. **a.** The shuffle null model maintains per-neuron variance and samples correlations near zero. **b.** The uniform correlation null model maintains per-neuron variance and samples uniform correlations. **c.** The factor analysis null model combines a fixed private variability (estimated from the experimental data, left gray inset) with shared variability (right teal inset) that can be rotated to form null samples (dash styles are consistent between the teal shared variabilities in the inset and the purple null samples in the main panel). **d.** For a synthetic 2d dataset, the LFI for the fixed-marginal parameterization as a function of the pairwise correlation, ρ , is shown at the top, the bottom plots are the covariance and samples as a function of ρ . **e.** For a synthetic 2d dataset, the LFI for the factor analysis parameterization as a function of the rotation angle, θ , is shown at the top, the bottom plots are the covariance and samples as a function of θ . **f.** To calculate an observed LFI or percentile under a null model, d units were randomly drawn from the population to form a “dimlet”. Then, two neighboring stimuli, s_1 and s_2 , were chosen. The dimlet and stimulus pairing together constitute a “dim-stim”, or a pair of design matrices $[\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d]$. These dim-stims are the samples inputs into a LFI calculation and form the basis for distributions of calculated quantities. **g.** Dim-stim responses in the retinal data for the depicted stimulus pairing (colors) from **f**.

119 being recorded, their location in the brain, or the recording modality. Thus, it is beneficial if the parameters
120 of the null model have a biological interpretation. We propose two null models that allow us to assess the
121 optimality of experimental neural responses: the uniform correlation (UC) null model and the factor analysis
122 (FA) null model. The uniform correlation null model maintains the per-neural unit distributions of activity,
123 like the shuffle null model. In contrast to the shuffle null model which samples the correlations around
124 zero (**Fig. 2a**), the uniform correlation null model samples the multivariate correlations uniformly (**Fig. 2b**,
125 dashed lines are samples with different correlations) [37]. In the UC null model, neural units maintain their
126 private mean and variance for a particular stimulus, but have the freedom to change their multivariate pair-
127 wise correlations (ρ). Biologically, changing the pairwise correlations could be achieved through recurrent
128 connectivity within the network of neural units. Depending on the correlations, the network could achieve a
129 range of coding fidelities as assessed by the LFI (**Fig. 2d**, covariance structures shown below the plot lead the
130 LFI as a function of the scalar pairwise correlation (ρ). At extreme values of correlation, the LFI can take on
131 the highest values [22]. Mathematically, the UC null model constrains the per-unit variances while sampling
132 the multivariate correlations (ρ) uniformly (see Methods for details). Motivated by experimental findings
133 that the variability in population responses has private and shared components [7, 14], we also developed
134 a factor analysis (FA) null model. The FA null model decomposes the experimentally observed covariance
135 into independent private variances and shared variability [7, 15]. The private variance is fixed (**Fig. 2c**, gray
136 ellipse in the left inset) and the shared variability's weighting on different neural units can change through
137 a rotation (**Fig. 2c**, dashed teal ellipses in the right inset are sampled rotations of the shared variability).
138 Biologically, this models each neuron having fixed private variability and incoming shared variability which
139 could be weighted in different ways. As the shared variability is rotated, the covariance structure varies, and
140 the LFI takes on a smaller range of values than in the UC null model (**Fig. 2e**, covariance structures shown
141 below the plot generate the LFI as a function of rotation angle, $R(\theta)$). Mathematically, the FA null model
142 constrains the factor analysis private variances but applies uniformly sampled rotations to the loading matrix
143 for the shared variability (see Methods for details). Together, these null models define the potential space of
144 covariances based on two different biological motivations and provide suitable tests of optimality.

145 To use the null models, for each neural population (retinal ganglion cells, V1 neurons, electrodes in
146 primary auditory cortex), we randomly sampled “dimlets”, or sub-populations of neural units, of dimension
147 d . We combined dimlets with a variety of neighboring stimulus pairings to obtain a subset of the neural
148 responses which we call a dim-stim (**Fig. 2f**, see Methods). A dim-stim would be the input to the task
149 of constructing a decoder for neighboring stimuli using a neural sub-population's responses across trials
150 (**Fig. 1e, h, k** and **Fig. 2g**). For each dataset, we generated a large number of dim-stims across a set of
151 dimensions $d = 3 - 20$ (see Methods). We calculated the LFI for each dim-stim across dimensions and
152 datasets. We refer to this quantity as the observed LFI. Next, we sampled the null models 1,000 times for
153 each dim-stim, and calculated the LFI for each sample (see Methods). Thus, for each dim-stim, we obtained
154 a single experimental LFI and a corresponding distribution of LFIs for each null model. The 1,000 null LFIs
155 constitute a null distribution to compare the experimentally observed LFI against. In particular, we define
156 the percentile as the fraction of the 1,000 null LFIs which are less or equal to than the observed LFI. Higher
157 percentiles indicate that the observed LFI is larger than more samples from the null model.

158 **The geometry of correlated variability leads to suboptimal neural coding**

159 With the uniform correlation (UC) and factor analysis (FA) null models, we assessed the optimality of
160 the neural code. To characterize the optimality of a wide range of sub-population and stimulus settings,
161 we performed a large scale analysis evaluating the LFI in both the experimentally observed data and null
162 models (see Methods). We compared the experimentally observed LFI to the distribution of LFI from the
163 null models. Specifically, for the experimental data, we compute the median LFI across dim-stims at each
164 dimension (**Fig. 3a-c**, black lines). For the shuffle, uniform correlation (UC), and factor analysis (FA) null

165 models, we first calculated the median LFI from the null distribution for each dim-stim and then report the
166 median across dim-stims (**Fig. 3a-c**, gray, blue, and orchid lines, respectively).

167 As expected, the experimentally observed LFIs across dim-stims grew with dimlet dimension, indicating
168 that increasing the dimension of the neural population improved the stimulus decoding (**Fig. 3a-c**, black
169 lines). Similarly, the median null model LFIs grew with dimlet dimension. The shuffle null model exhibited
170 comparable discriminability relative to the experimental LFI at lower dimensions (**Fig. 3a-c**, gray lines). At
171 higher dimensions, however, the shuffle null model LFIs began to exceed the observed LFIs. In contrast,
172 both the uniform correlation and factor analysis null models exhibited considerably larger median LFIs than
173 the observed data, with the disparity increasing with dimlet dimension. Therefore, on average, the stimuli
174 were more easily discriminable using the covariances sampled from the UC and FA null models than the
175 experimental covariance. We further observed differences across datasets. For example, the factor analysis
176 null model (**Fig. 3a-c**, orchid lines) exhibited similar LFIs to the uniform correlation null model for the PAC
177 dataset. However, in the retina and V1 data, the factor analysis LFIs were more comparable to the observed
178 and shuffle LFIs. Overall, **Figure 3a-c** demonstrates that the uniform correlation and factor analysis null
179 models produce LFIs that generally exceed the LFIs of the observed data, suggesting the neural code is
180 suboptimal.

181 Although the differences between the null model LFIs and observed LFIs were large, the preceding anal-
182 ysis was done at a population level rather than comparing each dim-stim LFI with its own null distribution.
183 Therefore, we quantified the optimality per dim-stim, relative to a null model, with its observed percentile.
184 To calculate the population optimality measure, the median percentile across dim-stims is taken. A higher
185 percentile means that the observed LFIs are greater than a larger fraction of the null LFIs. To operationalize
186 the notion of population optimality, we define three categories for optimality based on the median of the
187 experimental distribution of percentiles. If the median is greater than $2/3$, the population is optimal (Opt),
188 if the median is between $1/3$ and $2/3$ the population is near-chance (NC), and if the median is below $1/3$ the
189 population is suboptimal (Sub). Alternative categorizations could be used, but we chose the even splitting
190 into thirds for simplicity (see Methods for details).

191 We found that each null model exhibits distinct LFI distributions, with further variation depending on
192 the dataset and dim-stim. Example null model distributions for individual $d = 3$ dim-stims are depicted
193 in **Figure 3d-f** (vertical black line indicates the experimental LFI, gray, blue, orchid are the shuffle, UC,
194 and FA null model LFI distributions respectively, note that the uniform correlation null distributions often
195 have long tails and are truncated for visualization). The examples highlight that the percentiles can vary
196 across null models for a dataset (**Fig. 3d-f**, inset text). The heterogeneity in observed percentiles motivated
197 examining their distribution across all dim-stims. Thus, for each dataset, we computed the distribution of
198 observed percentiles across the dim-stims per dimlet dimension ($d = 3$ to $d = 20$). The median observed
199 percentile (calculated across dim-stims) as a function of dimlet dimension is shown in **Figure 3g-i**. Consis-
200 tent with other studies [6, 13, 16], we found that the shuffle null model (gray lines) often had large observed
201 percentiles, indicating that the shuffle null model often showed the benefits of experimentally observed cor-
202 relations versus having no correlations. However, it would be misleading to interpret these results as a test
203 of optimality. Indeed, compared to the uniform correlation (blue lines) and factor analysis (orchid lines)
204 null models, the experimental data exhibited suboptimal observed percentiles (**Fig. 3g-i**, blue and orchid
205 lines). All percentiles decreased with dimlet dimension, implying that the neural representations became
206 less optimal as the number of neurons increases. In theory, this decrease is expected as eventually differ-
207 ential correlations induce information saturation in the populations, however recent work indicates that we
208 should not expect to see the impact of differential correlations at this relatively small scale [38–40]. Indeed,
209 saturation of the LFI was not evident in **Figure 3a-c**. This indicates that the suboptimality observed in
210 **Figure 3g-i** is not due to differential correlation, but from some other biological cause.

211 **Figure 3g-i** also highlights differences across datasets. The shuffle null model had the lowest observed
212 percentiles among the three datasets for the retina data, starting near-chance for small dimlet sizes and

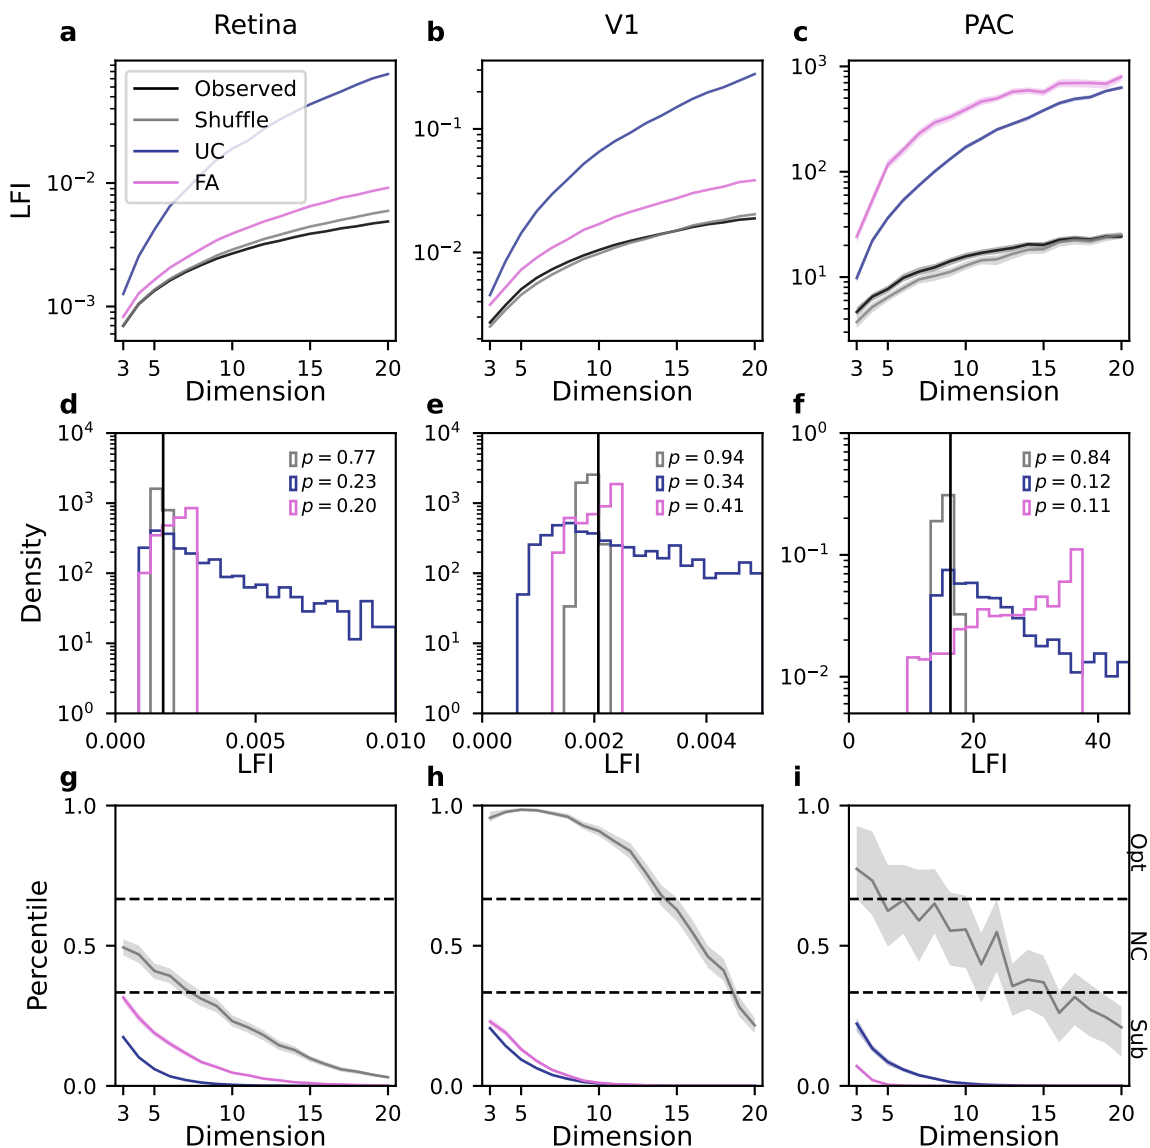


Figure 3: The geometry of correlated variability leads to suboptimal neural coding. Each column corresponds to one of the datasets. Color legend is shared across columns. Color legend is preserved across panels. **a-c.** The median LFI is plotted (solid lines, log-scale y -axis) as a function of the dimlet dimension (x -axis) for the observed correlated variability and null model samples (colors in legend). Shaded regions indicate the 95% CI of the median LFI (note that CIs are often comparable to the median line width). **d-f.** Histograms of null LFIs are shown for the shuffle, uniform correlation, and factor analysis null models for one dim-stims for each dataset. The observed LFI is denoted by the black vertical line in each plot. Percentiles for each null model are reported. **g-i.** Median observed dim-stim percentiles are shown (solid lines) as a function of dimlet dimensions, for each dataset and null model. Shaded regions indicate the 95% CI of the median observed percentile (note that CIs are often comparable to the median line width). Black dashed lines divide optimal (Opt), near-chance (NC), and suboptimal (Sub) regions.

213 dropping below $1/3$ around $d = 7$ (**Fig. 3g**, grey lines). For the V1 data, the shuffle null model clearly
214 exhibited the highest observed percentiles, indicating the coding benefits of correlations compared to zero
215 correlations for small dimlet sizes up to $d = 15$ (**Fig. 3h**, grey lines). In the primary auditory cortex data, the
216 shuffle null model exhibited intermediate observed percentiles, with a larger spread in confidence intervals,
217 indicating a higher heterogeneity in the observed percentiles (**Fig. 3i**, gray shaded region). Meanwhile, the
218 observed percentiles for the uniform correlation and factor analysis null models were more similar across
219 the three datasets, with slightly different magnitudes. In particular, the retinal data exhibited the largest
220 observed percentiles for the factor analysis null model, while the PAC data exhibited the smallest, going to
221 zero around $d = 5$. The uniform correlation null model had the lowest percentiles for the retina dataset and
222 similar percentiles for the V1 and PAC datasets. This behavior roughly tracked the distribution of pairwise
223 correlations amongst the three datasets (**Fig. 1e, h, k**), with the retinal data possessing the lowest average
224 noise correlation, and the PAC data possessing the highest average noise correlation. Critically, across all
225 datasets and dimensions, the percentiles for both the uniform correlation and factor analysis null models
226 were below $1/3$. This indicates that the geometry of correlated variability leads to suboptimal coding, and
227 that the suboptimality becomes more pronounced with increasing neural dimension.

228 **Optimal correlated variability is typically biologically inaccessible**

229 The results of the preceding section indicate that the geometry of correlated variability is highly suboptimal,
230 as opposed to near-chance or optimal. We next sought to understand why this was the case. For the uniform
231 correlation model, we summarize findings about optimal correlations from Hu *et al.* [22]. For the factor
232 analysis model, we compared the structure of the observed covariances to those of the optimal covariances.

233 When the per-neural unit variability is fixed, as in the shuffle and uniform correlation null models, Hu
234 *et al.* [22] showed that the optimal covariance structure will lie on the boundaries of the allowed values of
235 ρ for several measures of coding fidelity, including the LFI (**Fig. 2d**). The authors discussed that points on
236 the boundary may fall outside of biologically allowed regions. Consistent with this, we found that optimal
237 correlation matrices for the uniform correlation null model often had absolute pairwise correlations that
238 are close to 1, which was never observed in the experimental data (see **Supplemental Fig. 3**). Thus, the
239 optimal correlated variability structure suggested by the uniform correlation null model may be biologically
240 inaccessible. Meanwhile, the factor analysis model allows the distribution of highest pairwise correlations
241 to be modified (and generally increased), but does not extend near 1, suggesting that the distribution of noise
242 correlations achieved by the factor analysis null model is more biologically realistic.

243 Both the shuffle and uniform correlation null models will necessarily reproduce the observed single-unit
244 statistics, because they only change the correlations. Therefore, both of these null models will reproduce the
245 Fano factors (FF, $\frac{\text{variance}}{\text{mean}}$) and negative densities (ND, fraction of activity below the smallest responses of the
246 experimental activity) of the observed data. The factor analysis null model, however, can produce covari-
247 ance ellipses that have different single-unit distributions. Thus, some FA-optimal covariances may orient
248 variance in the negative or low-activity regions of the neural space. For the factor analysis null model, we
249 quantified the degree to which the biological inaccessibility of optimal covariances related to the percentiles
250 of the experimental data for each dim-stim. The Fano factor quantifies the variability of neural units relative
251 to their average activity. Typically, Fano factors for single-unit firing rates have been observed to be near
252 1 [41–44], in line with the approximately Poisson nature of firing rates. Thus, a large deviation from the
253 Fano factors observed in the experimental data indicates the single-unit properties of the optimal covariances
254 are biologically implausible (**Supplementary Fig. 2**). First, we examined whether the observed Fano factor
255 diverged from the Fano factors achieved by the FA-optimal covariance on each dim-stim via their absolute
256 log-ratio (see Methods). Large values of this quantity indicate greater difference between optimal and ex-
257 perimental single-unit distributions, suggesting less biological plausibility. Relatedly, a sample-covariance
258 that has negative neural activity can be interpreted as less biologically plausible, because negative activity is

259 either unachievable (for single-unit count variables) or highly unlikely (calcium imaging $\Delta F/F$ or baseline
260 z -scored μECoG) (**Supplementary Fig. 2**). Therefore, the second quantity we examined was the absolute
261 difference in negative density (ND), which captures the degree to which the FA-optimal covariance has neg-
262 ative neural activity (see Methods). Larger values of the negative density imply less biological plausibility.
263 We used these two measures of biological plausibility to assess when the observed neural responses can be
264 optimal according to the FA null model.

265 We determined whether the Fano factor (FF) and negative density (ND) distributions of the optimal co-
266 variances from the FA null model related to the suboptimality of the experimentally observed neural code.
267 To do this, we directly compared the optimal FA null model Fano factors to the experimental Fano factors
268 in **Figure 4a-c**. Across dim-stims, for $d = 3$, **Figure 4a-c** shows 2d-histograms of the absolute log-ratio
269 of Fano factors against the FA percentile, with darker colors corresponding to higher log-density of sam-
270 ples. For each histogram, we additionally plot the median percentile as a function of the log-ratio in blue.
271 We found that when the Fano factors closely matched (i.e., the log-ratio was close to zero), the percentiles
272 spanned a broad range between 0 and 1 (medians percentiles: 0.51, 0.41, 0.15 for the lowest bin across
273 datasets). However, FA-optimal covariances commonly deviated from the observed Fano factors, and when
274 they did, the observed percentiles dropped below 0.5 and were often near 0. Thus, as the biological acces-
275 sibility of the optimal covariance decreased, so did the optimality of the observed neural code. Likewise,
276 for negative density (ND), we directly compared the optimal FA null model NDs to the experimentally ob-
277 served NDs in **Figure 4d-f**. Across dim-stims, for $d = 3$, **Figure 4d-f** shows 2d-histograms of the absolute
278 difference in NDs against the FA percentile with darker colors corresponding to higher log-density. For
279 each histogram, we additionally plot the median percentile as a function of ND difference in red. We found
280 that when the difference was close to zero, the percentiles spanned a broad range between 0 and 1 (medians
281 percentiles: 0.47, 0.60, 0.31 for the lowest bin across datasets). However, the ND of FA-optimal covariances
282 commonly deviated from the observed ND, and when they did, the experimentally observed percentiles were
283 typically closer to 0.

284 We summarized the relationship between biological plausibility and percentile for both FF and ND.
285 At each dimension d , we calculated the Spearman rank correlation between the observed percentile and
286 each measure of biological plausibility (**Fig. 4g-i**). For each dataset, we observed negative correlations
287 that were significantly lower than zero across dimensions ($p < 10^{-5}$, one sample t -test). These negative
288 correlations imply that observed percentiles are smaller (i.e., the neural code is more suboptimal) when
289 optimal correlated variability is biologically inaccessible. Together, these results indicate that the optimal
290 covariances under the FA null model for $d \geq 3$ are not biologically accessible.

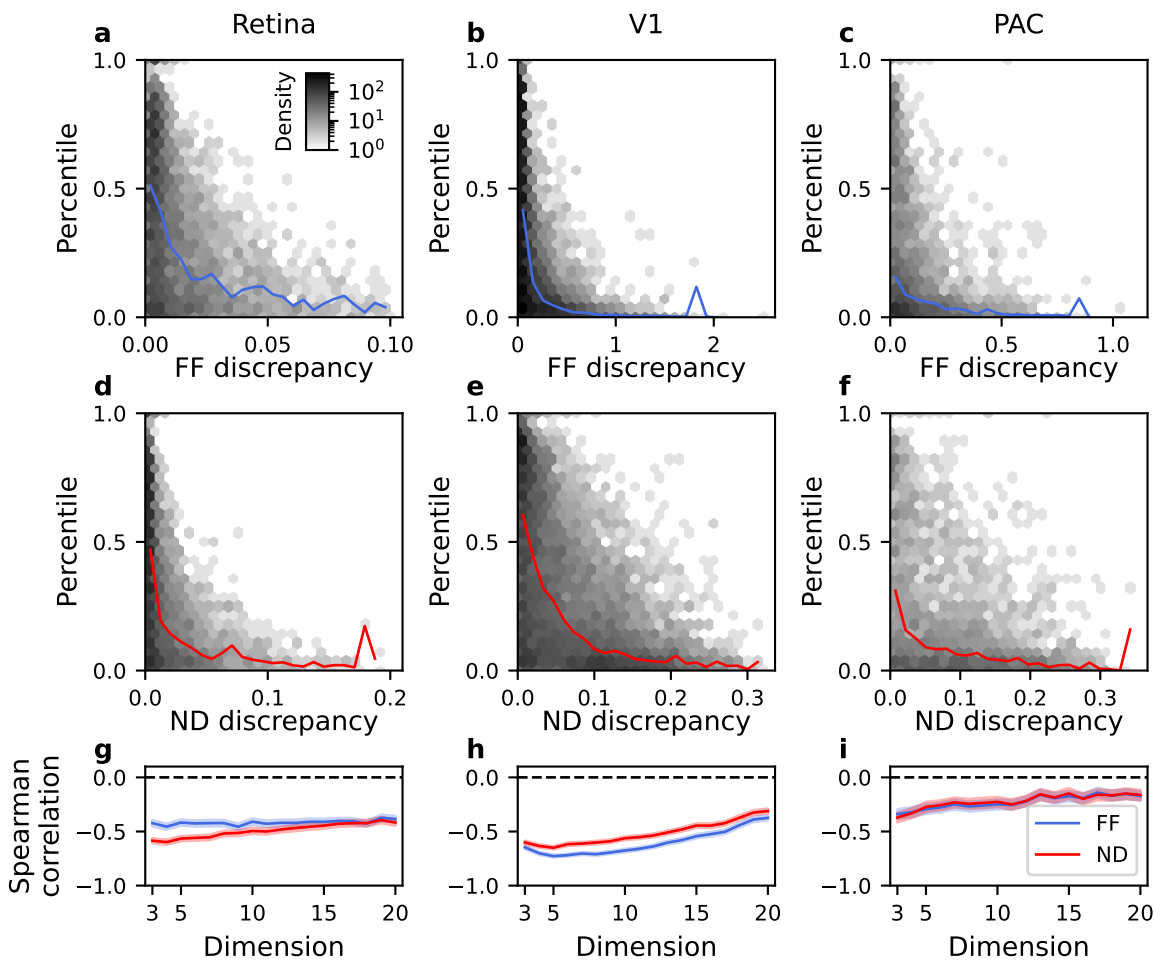


Figure 4: Optimal correlated variability is typically biologically inaccessible. Each column corresponds to a separate dataset. 2d histograms are plotted with a log-density color scale with shared colorbar. Color legend in **i** is shared across panels. **a-c.** 2d-histogram across dim-stims of the observed percentile under the FA null model versus the absolute log-ratio of the observed and FA-optimal covariance Fano factors for $d = 3$. Blue line is the median binned percentile as a function of the absolute log-ratio of observed and FA-optimal covariance Fano factors. **d-f.** 2d-histogram across dim-stims of the percentile under the FA null model versus the absolute difference of negative densities (ND) of the observed and FA-optimal covariance Fano factors for $d = 3$. Red line is the median binned percentile as a function of the absolute difference in NDs. **g-i.** The Spearman correlation coefficient between the observed percentile and absolute log-FF ratio or absolute difference of NDs, respectively is shown as a function of dimlet dimension. Dashed black line indicates zero correlation.

291 **Optimal subpopulations are exponentially small**

292 The results in the preceding section show that a majority of experimental dim-stims could not attain optimal
293 covariances according to the UC and FA null models due to biological constraints. However, it is possible
294 that although a majority of experimental dim-stims are suboptimal, there is a subset that are optimal, and
295 these specific subpopulations are somehow utilized by the nervous system. If this was the case, the uniform
296 sampling strategy over neural units may underestimate optimality as utilized by the nervous system. For
297 example, in the retina, if we are imagining a downstream region like V1 is decoding the stimuli, then a more
298 retinatopic sampling strategy, where retinal ganglion cells are more likely to be considered in a dimlet if they
299 are located spatially near each other in the retina would be preferable. Alternatively, synaptic learning rules
300 in downstream areas may select for neural populations that are tuned for similar stimuli. The responses to
301 the preferred stimuli would be high and therefore we expect less Fano factor and negative density violation.
302 Thus, it is possible that dim-stims subselected by these criteria will be more optimal than dim-stims sampled
303 uniformly.

304 To test if biologically motivated subsampling of dim-stims improved the percentiles, we performed
305 distance- and tuning-based subselection of the neural populations. For the retina and PAC datasets, we had
306 access to the spatial locations of the RGCs/electrodes. We subselected 10% of dim-stims with the small-
307 est average physical distance. Similarly, we subselected the 10% of dim-stims that had the most preferred
308 stimuli (see Methods for details on subselection). We found that distance-based subselection did not re-
309 veal an optimal or near-chance subset of dim-stims (**Fig. 5a, c**, dotted lines and hatched shaded regions).
310 Similarly, for the retina and V1 datasets, the tuning-based subselection did not reveal an optimal subset of
311 dim-stims and the percentiles only improved to near-chance for the PAC dataset at $d = 3$ (**Fig. 5a-c**, solid
312 lines and shaded regions). Furthermore, subselection directly based on the FF and ND criteria also did not
313 find optimal or near-chance percentiles (**Supplementary Fig. 4**).

314 Although these subselection criteria are biologically motivated, the previous results do not address
315 whether any subpopulation of the neural units across stimuli have optimal percentiles, and if so, how small
316 the subpopulation is. Intuitively, given the combination of a large enough neural population, variety of
317 stimuli, and enough dim-stims, one would expect at least a small fraction of the dim-stims to have optimal
318 percentile statistics by chance. To estimate the size of the optimal subpopulation, we calculated the optimal
319 fraction of the neural population, that is, largest fraction of dim-stims that could be retained and still achieve
320 optimal percentile statistics (median $\geq 2/3$) (**Fig. 5d-f**). If the optimal fraction is smaller, optimal subpopu-
321 lations are more rare. As a reference, if the distribution of percentiles was uniform, the largest two-thirds of
322 the percentiles could be retained and their median would be $2/3$, which is optimal (**Fig. 5d-f**, black dashed
323 line). At $d = 3$ for the FA null model (**Fig. 5d-f**, orchid line), across datasets between 14% and 37% of the
324 entire population was optimal if subselected. The optimal fraction according to the FA null model dropped
325 below 10% by $d = 4 - 9$ and below 2% by $d = 13 - 15$ across datasets. At higher dimensions, the optimal
326 subpopulation continued to become exponentially small, although the PAC dataset had a slower decrease.
327 According to the uniform correlation null model, for the retina and V1 datasets, less than approximately
328 0.1% of the population was optimal since almost no subpopulation was found from the finite samples. At
329 $d = 3$ for the PAC dataset, 20% of dim-stims would be considered optimal, but that drops below 1% by
330 $d = 7$ and continued to decrease to the smallest possible estimated value by $d = 12$ since no subpopula-
331 tions were found for higher dimensions. Finally, an alternative analysis of peaks in the percentiles near 1
332 in excess of what would be expected from a uniform distribution confirmed that there were exponentially
333 small optimal populations (**Supplemental Fig. 5**). Together these results show that correlated variability is
334 suboptimal in the neural recordings considered here. Furthermore, biologically motivated selection criteria
335 are not able to find the exponentially small optimal subpopulations.

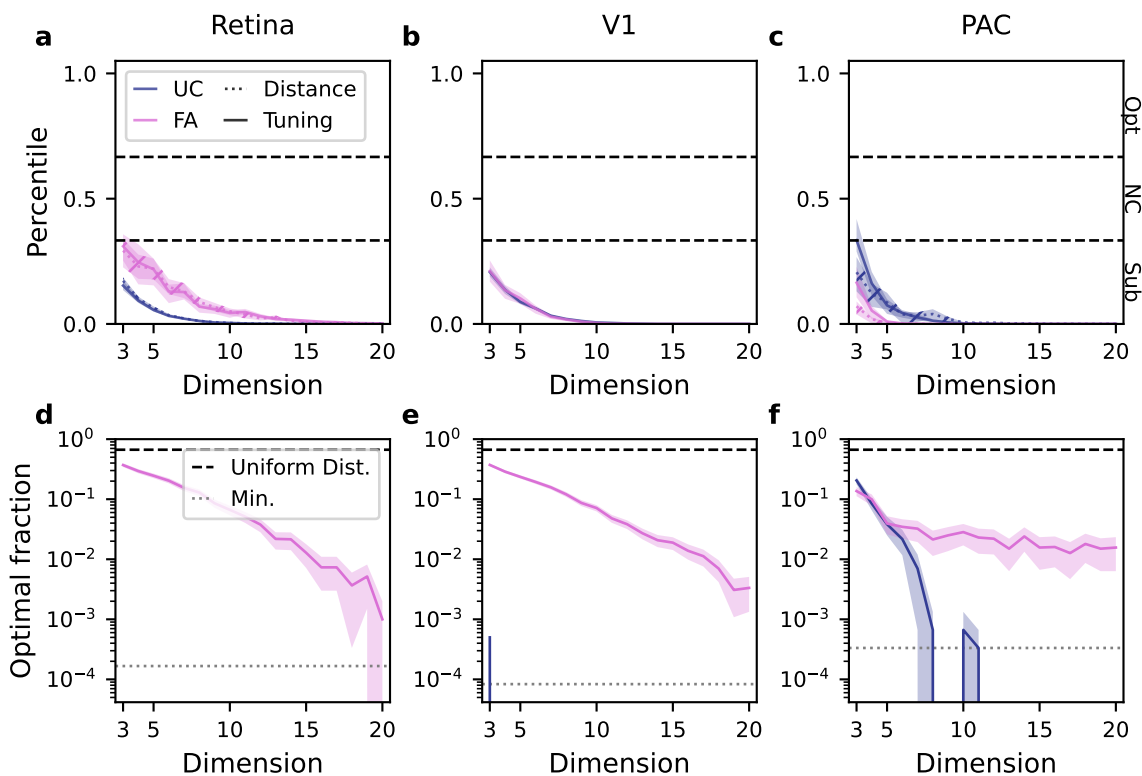


Figure 5: Optimal subpopulations are exponentially small. Color legend in **a** is shared across panels. **a-c.** For the uniform correlation and FA null model, dim-stims were subselected to maximize the units' tuning (solid lines, highest 10% subselected). Additionally, for the retina and PAC datasets, dim-stims were subselected to minimize the average pairwise distance between the RGC RoIs in a dim-stim (dashed lines, lowest 10% subselected). The median percentiles are shown as a function of dimension. Black dashed lines indicate the 1/3 and 2/3 percentile range. Shaded regions indicate the 95% CI of the median percentiles. **d-f.** For each dimension, the largest possible fraction of dim-stim percentiles such that their median is $\geq 2/3$ is plotted. Shaded regions indicate 95% CI. For the uniform correlation null model, dimensions where no samples exceeded the 2/3 threshold are not plotted. Black dashed line indicates the optimal fraction if percentiles were drawn from a uniform distribution. Gray dotted line indicates the minimum non-zero optimal fraction that can be estimated due to finite sampling.

336 Discussion

337 Determining the principles of the neural code is critical for a complete understanding of brain function.
338 Correlated variability is prevalent in neural recordings and has been the subject of numerous studies seeking
339 to understand its mechanistic sources and implication for neural coding. Many previous studies have found
340 that the experimentally observed correlations can be a benefit to neural coding compared to having zero
341 correlations [6, 13, 16, 24, 25, 27]. This suggests that the correlated variability could in fact be optimal.
342 However, the shuffle null model used in these studies is not able to assess optimality. To the best of our
343 knowledge, the optimality of correlated variability in neural data has not previously been assessed.

344 Here, we developed two null models which allow the optimality of observed correlated variability to
345 be directly assessed: the uniform-correlation (UC) and factor analysis (FA) null models. Using these null
346 models, we found that the experimentally observed neural activity across three datasets was consistently
347 suboptimal. As more neural units were included in the neural population, the neural populations became
348 more suboptimal. In order to more fully understand the suboptimality, we evaluated the characteristics of the
349 optimal covariance and found that a consistent picture emerges: for a majority of neural subpopulations, the
350 optimal covariance is biologically inaccessible. We then used biologically motivated subselection criteria
351 to assess whether there were subpopulations with optimal coding statistics. We found that subsampling
352 using criteria based on the tuning of units or the spatial location of the units does not result in increased
353 coding optimality. Finally, we showed that optimal subpopulations based on *post-hoc* selection became
354 exponentially small as the dimensionality of the neural population increased. Thus, we conclude that in the
355 early sensory areas studied here, the geometry of correlated variability leads to highly suboptimal neural
356 coding.

357 We observed suboptimal coding performance as assessed by both the uniform correlation and factor
358 analysis null models. However, the magnitude of the suboptimality, as measured by the observed per-
359 centiles, differed across null models and datasets. The observed percentiles for the uniform correlation null
360 model had a small trend from low to high for the retinal data, the V1 data, and the PAC data, respectively.
361 This trend tracks with the distribution of noise correlations in each dataset (**Fig. 1f, i, l**), with the the retina
362 dataset exhibiting, on average, the smallest magnitude noise correlations, and the PAC datasets exhibiting the
363 largest. The smaller range of noise correlations exhibited by the retina suggests that there may be stronger
364 biological restrictions on its correlated variability compared to V1 and PAC. The observed percentiles for
365 the factor analysis null model trend from just below near-chance to highly suboptimal from retina to PAC.
366 Thus, the larger correlations and more suboptimal coding performance indicates that shared variability in
367 V1 and PAC is more likely to interfere with sensory coding. The retina and V1 recording modalities (cal-
368 cium imaging and single-unit electrophysiology, respectively) measure putative single-unit activity where
369 correlated variability in the recordings corresponds to correlated single neuron activity. Understanding the
370 optimality of the neural code with these two modalities directly addresses decoding as a normative theory in
371 early sensory areas. On the other hand, the correlated variability in the μ ECoG recordings in PAC is likely
372 due to a combination of the correlations between the neural populations under each electrode and local tissue
373 conduction [45, 46]. Due to this, the optimality of the high gamma amplitude correlated variability recorded
374 with μ ECoG is a coarse-grained signal that may not be read-out by any downstream cortical area, but is
375 important for understanding whether limitations in the accuracy of clinical ECoG-based brain-computer
376 interfaces in humans may be due to correlated variability in the input signals.

377 Many studies of correlated variability, including ours, consider the impact of correlated variability from
378 a decoding perspective. However, other normative perspectives exist. In Bayesian models of sensory pro-
379 cessing [47], correlated variability could correspond to sampling from a relevant (posterior) distribution.
380 In this case, correlated variability would be informative for understanding the structure of uncertainty in
381 sensory processing, rather than nuisance variability as in the decoding perspective. Likewise, neural sys-
382 tems likely have other important constraints or ethological goals. Making decisions or generating behavior

383 based on sensory information may be optimized by different correlation structures versus a purely decoding
384 framework [48]. For example, Valente *et al.* [48] find that single-trial responses in posterior parietal cortex
385 which have higher noise correlations also have more correct choices, contrary to expectation. They model
386 this finding with a read-out network that computes an additional nonlinear “consistency” value across the
387 population in addition to the linear sensory information for use in decision making. Huang & Lisberger [49]
388 show that correlated variability in middle temporal visual area could plausibly be the cause of variability
389 in smooth-pursuit eye movements. Even within the normative decoding framework, correlated variability
390 which facilitates decoding as assessed by the LFI may not be the same as the correlated variability which
391 facilitates information propagation or learning in more realistic nonlinear, noisy networks [5, 48, 50]. In
392 these contexts, our formalism for creating null models could be used to test the optimality of neural codes,
393 although as the assumptions on linear decoding are relaxed, it may become difficult to make theoretical
394 predictions that hold generally.

395 The null models we proposed both have parameterizations that are interpreted in a fully Gaussian model.
396 Generalized linear models [51, 52] or correlated multivariate distributions with binary-spike or spike-count
397 distributions [53–55] could potentially better model nonlinearities between the parameters of the model and
398 the non-Gaussian neural responses, which can impact estimates of neural coding optimality. In order to
399 assess optimality in these models when fit to data, a similar formalism for generating null models is needed,
400 where certain parts of the parameterization are fixed and others are given a null distribution. However, the
401 independent parameterization of the mean responses (tuning) and correlated variability is a unique feature of
402 the multivariate Gaussian distribution. Therefore, new analytical results would be needed to directly study
403 the impact of non-Gaussian correlated variability on neural coding. A broader set of null distributions could
404 similarly be used in phenomenological models of correlated variability which combine tuning and various
405 types of (correlated) noise [6, 21, 28, 56] or in mechanistic models, which attempt to simulate some aspects
406 of the neural circuit which lead to correlated variability [5, 15, 16, 57].

407 Correlated variability has been shown to be impacted by behavior and brain states. For example, it has
408 been observed that behavior such as running, whisking, and pupil diameter are encoded in V1 and other
409 brain areas [7]. In these contexts, the behavioral subspaces could be estimated directly (as in [7]) and
410 their optimality could be assessed using the FA null model. In experiments with visual attention, it has
411 been shown that attention can modulate both the within-area and between-area correlated variability [31,
412 58, 59], which can lead to better coding fidelity or better communication of information as assessed by the
413 shuffle null model. Similarly, in an associative task, learning has been shown to modulate the mean response
414 manifold and correlated variability to improve coding in pairs of neurons. The null models developed here
415 could be used to assess whether the modulation due to attention or learning changes the optimality of the
416 correlated variability. Emerging neural recording technologies will allow neuroscientists to simultaneously
417 record from a larger fraction of neurons in a region and more regions, all while the animals are performing
418 naturalistic behaviors. Given these possibilities, the biological origins of correlated variability and how they
419 are modulated by neural circuitry can be further traced and evaluated.

420 In summary, we find that the geometry of correlated variability in sensory areas leads to highly sub-
421 optimal coding for transmission of information about the stimulus. Given the consistency of the findings
422 across datasets, we expect our results would hold true in other organisms, sensory areas, and experimental
423 paradigms. Investigated more broadly, understanding the optimality of correlated variability could lead to
424 a better understanding of the sources of variability in neural circuits and biological constraints that lead to
425 suboptimality. Furthermore, quantitatively evaluating normative theories allows us to adjudicate between
426 competing proposed functions of sensory systems, for example, efficient coding versus predictive informa-
427 tion coding.

428 **Acknowledgements**

429 J.A.L. was supported by the LBNL LDRD “Deep Learning for Science” and the Weill Institute for Neu-
430 roscience at UCSF (Bouchard). P.S.S. was supported by the Department of Defense (DoD) through the
431 National Defense Science & Engineering Graduate (NDSEG) Fellowship Program. M.E.D. was supported
432 by an LBNL LDRD. M.T.S. was supported by the National Science Foundation Graduate Research Fellow-
433 ship (DGE 1752814). K.E.B. was supported by DOE ASCR (FP00009697), NIH (RNS118648A), the Kavli
434 Institute, the Weill Institute for Neuroscience, and the LBNL LDRD “Coordination of Distributed Cortical
435 Circuits.” This research used resources of the National Energy Research Scientific Computing Center, a
436 DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy
437 under Contract No. DE-AC02-05CH11231. We would like to that the Neural Systems and Data Science
438 Lab and Frederic Theunissen for feedback.

439 References

- 440 1. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and neuronal population infor-
441 mation. *Annual review of neuroscience* **39**, 237–256 (2016).
- 442 2. Azeredo da Silveira, R. & Rieke, F. The Geometry of Information Coding in Correlated Neural Pop-
443 ulations. *Annual Review of Neuroscience* **44**, 403–424. ISSN: 1545-4126. [http://dx.doi.org/](http://dx.doi.org/10.1146/annurev-neuro-120320-082744)
444 [10.1146/annurev-neuro-120320-082744](http://dx.doi.org/10.1146/annurev-neuro-120320-082744) (July 2021).
- 445 3. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation.
446 *Nature reviews neuroscience* **7**, 358 (2006).
- 447 4. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nature neuroscience* **14**,
448 811 (2011).
- 449 5. Zylberberg, J., Pouget, A., Latham, P. E. & Shea-Brown, E. Robust information propagation through
450 noisy neural circuits. *PLoS computational biology* **13**, e1005497 (2017).
- 451 6. Franke, F. *et al.* Structures of neural correlation and how they favor coding. *Neuron* **89**, 409–422
452 (2016).
- 453 7. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**
454 (2019).
- 455 8. Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications
456 for psychophysical performance. *Nature* **370**, 140 (1994).
- 457 9. Dichter, B. K., Bouchard, K. E. & Chang, E. F. Dynamic structure of neural variability in the cortical
458 representation of speech sounds. *Journal of Neuroscience* **36**, 7453–7463 (2016).
- 459 10. Kohn, A. & Smith, M. A. Stimulus dependence of neuronal correlation in primary visual cortex of the
460 macaque. *Journal of Neuroscience* **25**, 3661–3673 (2005).
- 461 11. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex.
462 *Journal of Neuroscience* **28**, 12591–12603 (2008).
- 463 12. Ruff, D. A. & Cohen, M. R. Stimulus dependence of correlated variability across cortical areas. *Journal*
464 *of Neuroscience* **36**, 7546–7556 (2016).
- 465 13. Montijn, J. S., Meijer, G. T., Lansink, C. S. & Pennartz, C. M. Population-level neural codes are robust
466 to single-neuron variability from a multidimensional coding perspective. *Cell reports* **16**, 2486–2498
467 (2016).
- 468 14. Deweese, M. R. & Zador, A. M. Shared and private variability in the auditory cortex. *Journal of*
469 *neurophysiology* **92**, 1840–1855 (2004).
- 470 15. Sachdeva, P. S., Livezey, J. A. & DeWeese, M. R. Heterogeneous synaptic weighting improves neural
471 coding in the presence of common noise. *Neural computation* **32**, 1239–1276 (2020).
- 472 16. Zylberberg, J., Cafaro, J., Turner, M. H., Shea-Brown, E. & Rieke, F. Direction-selective circuits shape
473 noise to ensure a precise population code. *Neuron* **89**, 369–383 (2016).
- 474 17. Huang, C. *et al.* Circuit models of low-dimensional shared variability in cortical networks. *Neuron*
475 **101**, 337–348 (2019).
- 476 18. Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E. & Pouget, A. Not noisy, just wrong: the role of
477 suboptimal inference in behavioral variability. *Neuron* **74**, 30–39 (2012).
- 478 19. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code.
479 *Neural computation* **11**, 91–101 (1999).

- 480 20. Yoon, H. & Sompolinsky, H. *The effect of correlations on the Fisher information of population codes*
481 *in Advances in neural information processing systems* (1999), 167–173.
- 482 21. Ecker, A. S., Berens, P., Tolias, A. S. & Bethge, M. The effect of noise correlations in populations of
483 diversely tuned neurons. *Journal of Neuroscience* **31**, 14272–14283 (2011).
- 484 22. Hu, Y., Zylberberg, J. & Shea-Brown, E. The sign rule and beyond: boundary effects, flexibility, and
485 noise correlations in neural population codes. *PLoS computational biology* **10**, e1003469 (2014).
- 486 23. Bujan, A. F., Aertsen, A. & Kumar, A. Role of input correlations in shaping the variability and noise
487 correlations of evoked activity in the neocortex. *Journal of Neuroscience* **35**, 8611–8625 (2015).
- 488 24. Cafaro, J. & Rieke, F. Noise correlations improve response fidelity and stimulus encoding. *Nature* **468**,
489 964 (2010).
- 490 25. Graf, A. B., Kohn, A., Jazayeri, M. & Movshon, J. A. Decoding the activity of neuronal populations
491 in macaque primary visual cortex. *Nature neuroscience* **14**, 239–245 (2011).
- 492 26. Montani, F., Kohn, A., Smith, M. A. & Schultz, S. R. The role of correlations in direction and contrast
493 coding in the primary visual cortex. *Journal of Neuroscience* **27**, 2338–2348 (2007).
- 494 27. Ruda, K., Zylberberg, J. & Field, G. D. Ignoring correlated activity causes a failure of retinal popula-
495 tion codes. *Nature communications* **11**, 1–15 (2020).
- 496 28. Lin, I.-C., Okun, M., Carandini, M. & Harris, K. D. The nature of shared cortical variability. *Neuron*
497 **87**, 644–656 (2015).
- 498 29. Averbach, B. B. & Lee, D. Effects of noise correlations on information encoding and decoding. *Journal*
499 *of neurophysiology* **95**, 3633–3644 (2006).
- 500 30. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nature neuroscience* **14**,
501 811–819 (2011).
- 502 31. Ruff, D. A. & Cohen, M. R. Attention increases spike count correlations between visual cortical areas.
503 *Journal of Neuroscience* **36**, 7523–7534 (2016).
- 504 32. Beaman, C. B., Eagleman, S. L. & Dragoi, V. Sensory coding accuracy and perceptual performance
505 are improved during the desynchronized cortical state. *Nature communications* **8**, 1–14 (2017).
- 506 33. Downer, J. D., Rapone, B., Verhein, J., O’Connor, K. N. & Sutter, M. L. Feature-selective attention
507 adaptively shifts noise correlations in primary auditory cortex. *Journal of Neuroscience* **37**, 5378–5392
508 (2017).
- 509 34. Elsayed, G. F. & Cunningham, J. P. Structure in neural population recordings: an expected byproduct
510 of simpler phenomena? *Nature neuroscience* **20**, 1310–1318 (2017).
- 511 35. Kohn, A. & Smith, M. A. *Utah array extracellular recordings of spontaneous and visually evoked*
512 *activity from anesthetized macaque primary visual cortex (VI)* 2016. [http://dx.doi.org/10.](http://dx.doi.org/10.6080/K0NC5Z4X)
513 [6080/K0NC5Z4X](http://dx.doi.org/10.6080/K0NC5Z4X).
- 514 36. Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. *Pro-*
515 *ceedings of the National Academy of Sciences* **112**, E6973–E6982 (2015).
- 516 37. Joe, H. Generating random correlation matrices based on partial correlations. *Journal of Multivariate*
517 *Analysis* **97**, 2177–2189 (2006).
- 518 38. Kafashan, M. *et al.* Scaling of information in large neural populations reveals signatures of information-
519 limiting correlations. *bioRxiv* (2020).
- 520 39. Rumyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**,
521 100–105 (2020).

- 522 40. Montijn, J. S. *et al.* Strong information-limiting correlations in early visual areas. *bioRxiv* (2019).
- 523 41. Eden, U. T. & Kramer, M. A. Drawing inferences from Fano factor calculations. *Journal of neuro-*
524 *science methods* **190**, 149–152 (2010).
- 525 42. Softky, W. R. & Koch, C. The highly irregular firing of cortical cells is inconsistent with temporal
526 integration of random EPSPs. *Journal of neuroscience* **13**, 334–350 (1993).
- 527 43. Tolhurst, D. J., Movshon, J. A. & Dean, A. F. The statistical reliability of signals in single neurons in
528 cat and monkey visual cortex. *Vision research* **23**, 775–785 (1983).
- 529 44. Van Steveninck, R. R. d. R., Lewen, G. D., Strong, S. P., Koberle, R. & Bialek, W. Reproducibility and
530 variability in neural spike trains. *Science* **275**, 1805–1808 (1997).
- 531 45. Dougherty, M. E., Nguyen, A. P. Q., Baratham, V. L. & Bouchard, K. E. *Laminar origin of evoked*
532 *ECoG high-gamma activity in 2019 41st Annual International Conference of the IEEE Engineering in*
533 *Medicine and Biology Society (EMBC)* (July 2019), 4391–4394.
- 534 46. Baratham, V. L., Dougherty, M. E., Ledochowitsch, P., Maharbiz, M. M. & Bouchard, K. Columnar
535 localization and laminar origin of cortical surface electrical potentials. *bioRxiv* (2021).
- 536 47. Doya, K., Ishii, S., Pouget, A. & Rao, R. P. *Bayesian brain: Probabilistic approaches to neural coding*
537 (MIT press, 2007).
- 538 48. Valente, M. *et al.* Correlations enhance the behavioral readout of neural population activity in associ-
539 ation cortex. *Nature Neuroscience*, 1–12 (2021).
- 540 49. Huang, X. & Lisberger, S. G. Noise correlations in cortical area MT and their potential impact on
541 trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *Journal of Neuro-*
542 *physiology* **101**, 3012–3030 (2009).
- 543 50. Nassar, M. R., Scott, D. & Bhandari, A. Noise correlations for faster and more robust learning. *Journal*
544 *of Neuroscience* **41**, 6740–6752 (2021).
- 545 51. Brown, E. N., Barbieri, R., Eden, U. T. & Frank, L. M. Likelihood methods for neural spike train data
546 analysis. *Computational neuroscience: A comprehensive approach*, 253–286 (2003).
- 547 52. Kass, R. E., Ventura, V. & Brown, E. N. Statistical issues in the analysis of neuronal data. *Journal of*
548 *neurophysiology* **94**, 8–25 (2005).
- 549 53. Inouye, D. I., Yang, E., Allen, G. I. & Ravikumar, P. A review of multivariate distributions for count
550 data derived from the Poisson distribution. *Wiley Interdisciplinary Reviews: Computational Statistics*
551 **9**, e1398 (2017).
- 552 54. Schneidman, E., Berry, M. J., Segev, R. & Bialek, W. Weak pairwise correlations imply strongly cor-
553 related network states in a neural population. *Nature* **440**, 1007–1012 (2006).
- 554 55. Sokoloski, S., Aschner, A. & Coen-Cagli, R. Modelling the neural code in large populations of corre-
555 lated neurons. *Elife* **10**, e64615 (2021).
- 556 56. Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nature neuroscience*
557 **17**, 858 (2014).
- 558 57. Brinkman, B. A., Weber, A. I., Rieke, F. & Shea-Brown, E. How do efficient coding strategies depend
559 on origins of noise in neural circuits? *PLoS computational biology* **12**, e1005150 (2016).
- 560 58. Cohen, M. R. & Maunsell, J. H. Attention improves performance primarily by reducing interneuronal
561 correlations. *Nature neuroscience* **12**, 1594 (2009).
- 562 59. Ruff, D. A. & Cohen, M. R. Attention can either increase or decrease spike count correlations in visual
563 cortex. *Nature neuroscience* **17**, 1591–1597 (2014).

564 **Methods**

565 **Neural Recordings**

566 We examined correlated variability in a diverse set of datasets, spanning distinct brain regions, animal mod-
567 els, and recording modalities. We used calcium imaging recordings from mouse retinal ganglion cells,
568 single-unit recordings from macaque primary visual cortex, and micro-electrocorticography recordings from
569 rat auditory cortex. We briefly describe the experimental and preprocessing steps for each dataset. See **Fig-**
570 **ure 1** and **Table 1** for summaries of the datasets.

Dataset	Animal	Recording	Stimulus	Units	Stimuli	Trials/Stim
Retina	Mouse (Isolated)	Calcium Imaging	Drifting Bars	54	6	114
V1	Macaque	Single-Units	Drifting Gratings	106	12	200
PAC	Rat	μ ECoG	Tone Pips	65	30	60

Table 1: Experimental dataset summary.

571 **Recordings from mouse retina**

572 Mouse retina data was collected via ex vivo 2-photon calcium imaging in an isolated retina preparation [1].
573 The retina was bulk loaded with Cal-520 AM dye using a previously described multicell bolus loading
574 technique [2], and then imaged with ScanImage software [3] at 2.96 Hz in the ganglion cell layer of a 425
575 x 425 μ m area of ventral retina. Visual stimuli were delivered via an ultraviolet LED (375 nm) coupled
576 to a digital micromirror device, and were presented on the flyback of the fast-axis scanning mirror during
577 a scan to interleave the stimuli with imaging [1, 4]. Visual responses were elicited via 600 \times 600 μ m
578 bars drifting for 2.93 s at 750 μ m/s in one of 6 directions (spanning 0° to 300°), with a 5 second intertrial
579 interval. Each direction was presented 114 times, for a total of 684 trials per cell. Fluorescence signals from
580 832 manually selected regions of interest were baseline subtracted and normalized to calculate a $\Delta F/F_0$
581 time series. Of these regions of interest, 54 were used for further analysis after determination of directional
582 tuning via permutation testing and manual screening. Per-trial RGC activity used in the analysis here is the
583 maximum $\Delta F/F_0$ value. Retina data was collected by Summers. Further details on surgical, experimental,
584 and preprocessing steps can be found at [4, 5].

585 **Recordings from macaque primary visual cortex (V1)**

586 Primary visual cortex data (V1) was comprised of spike-sorted units simultaneously recorded in anesthetized
587 macaque monkey. The data was obtained from the Collaborative Research in Computational Neuroscience
588 (CRCNS) data sharing website [6] and was recorded by Kohn and Smith [7]. This dataset contains record-
589 ings from three monkeys, of which the main text presents results from the first one (see Appendix for results
590 on additional two monkeys). Recordings were obtained with a 10 \times 10 grid of silicon microelectrodes spaced
591 400 μ m apart and covering an area of 12.96 mm². The monkey was presented with grayscale sinusoidal
592 drifting gratings, each for 1.28 s. Twelve unique drifting angles (spanning 0° to 330°) were each presented
593 200 times, for a total of 2400 trials per monkey. Spike counts were obtained in a 400 ms bin after stimulus
594 onset. A total of 106 units were isolated in the monkey presented in the main text. These units were chosen
595 by the original authors such that *i*) their signal-to-noise ratio (the ratio of the average waveform amplitude to
596 the standard deviation of the waveform noise) was at least 2.75, *ii*) the best grating stimulus evoked at least
597 2 spikes/s, and *iii*) the variance-to-mean response ratio did not exceed 10. Further details on the surgical,
598 experimental, and preprocessing steps can be found in [8, 9].

599 Recordings from rat primary auditory cortex (PAC)

600 Auditory cortex data (PAC) was comprised of cortical surface electrical potentials (CSEPs) recorded from
601 rats with a custom fabricated micro-electrocorticography (μ ECoG) array. The μ ECoG array consisted of
602 an 8×16 grid of $40 \mu\text{m}$ diameter electrodes. Anesthetized rats were presented with 50 ms tone pips of
603 varying amplitude (8 different levels of attenuation, from 0 dB to -70 db) and frequency (30 frequencies
604 equally spaced on a log-scale from 500 Hz to 32 kHz). We only used samples for the lowest 3 levels of
605 attenuation since these evoked the largest responses. Each frequency-amplitude combination was presented
606 20 times, for a total of $3 \times 30 \times 20 = 1800$ samples. The response for each trial was calculated as the
607 z -scored to baseline, high- γ band amplitude of the CSEP, calculated using a constant-Q wavelet transform.
608 The maximum of the per-trial high- γ activity was used in the analysis here. Of the 128 electrodes, we
609 used 65, selecting those that recorded from primary auditory cortex. Data was recorded by Dougherty &
610 Bouchard. Further details on the surgical, experimental, and preprocessing steps can be found in [10, 11].

611 Linear Fisher information measures coding fidelity

612 A commonly used measure of coding fidelity in the context of decoding is the Fisher information, which
613 provides a limit on how accurately a readout of a neural representation can be used to determine the value of
614 the stimulus [12]. Formally, the Fisher information is a lower bound on the variance of an unbiased estimator
615 for the stimulus. In practice, the Fisher information is analytically intractable. An alternative measure is the
616 linear Fisher information (LFI), defined in Equation 1. The LFI acts as a suitable lower bound to the Fisher
617 information and is the most commonly used measure of coding fidelity in correlated variability analyses
618 [13–20].

Experimental neuroscience datasets only consider discrete sets of stimuli, which are not amenable to the
computation of LFI as posed in Equation 1. In particular, the derivative of the average neural activity must
be estimated by considering the neighboring pairs of stimuli. Thus, in practice, we calculate the coarsened
linear Fisher information [21], which is defined for two stimuli s_1 and s_2 as

$$\mathcal{I}_{\text{coarse}}(\mathbf{f}_1, \mathbf{f}_2, \Sigma_1, \Sigma_2) = \left(\frac{\mathbf{f}_1 - \mathbf{f}_2}{\Delta s} \right)^T \left(\frac{\Sigma_1 + \Sigma_2}{2} \right)^{-1} \left(\frac{\mathbf{f}_1 - \mathbf{f}_2}{\Delta s} \right) \quad (2)$$

619 where $\mathbf{f}_1 = \mathbf{f}(s_1)$, $\mathbf{f}_2 = \mathbf{f}(s_2)$, $\Sigma_1 = \Sigma(s_1)$, $\Sigma_2 = \Sigma(s_2)$, and Δs is the stimulus difference between s_1 and
620 s_2 , whose form may depend on the stimulus structure. In addition, we use the unbiased LFI estimator [20]
621 for the observed LFI values as well as for the sampled from null models. Note that since the corrections to
622 the naïve estimator only depend on the dimensionality of the neural population and number of samples, the
623 corrections only impact the raw LFI values and not percentiles. In this work, we use the terms “coarsened
624 LFI” and “LFI” interchangeably.

625 Assessing the optimality of neural data with null models

626 Information theoretic analyses of neural data often ask whether the observed neural data is “optimal.” In
627 the case of correlated variability, the question can be posed as: are the observed covariances optimal from
628 a decoding perspective? Here, we will quantify the coding fidelity with the linear Fisher Information (LFI,
629 Eq. 1)? In this case, LFI can be infinitely large if $\Sigma \rightarrow 0$ (or at least if the subspace of Σ^{-1} defined by $\frac{d\mathbf{f}(s)}{ds}$
630 diverges). This answer is likely unsatisfying because neural systems have many sources of variability, and so
631 expecting a neural system to become noiseless or exactly remove noise from a subspace seems implausible.
632 Therefore, when assessing the optimality of correlated variability, one must decide which aspects of the
633 correlated variability the neural system could modify and which aspects will remain fixed.

634 In this section, we develop the formalism that will allow us to assess the optimality of observed corre-
635 lated neural variability. The formalism consists of first defining a covariance parameterization for Σ , which

636 is composed of constraints (fixed parameters) and degrees-of-freedom (free parameters). These constraints
637 and degrees-of-freedom define the space of allowed correlated variability. Ideally, these constraints and
638 degrees-of-freedom have some biological interpretation, e.g., fixed private variability or input from other
639 regions of the brain [22, 23]. Then, a null model is defined by combining a covariance parameterization
640 with a null distribution over the degrees-of-freedom. The distribution of some measure, such as the LFI,
641 under the null model can be used to assess the optimality of the observed neural data.

642 We first review the commonly used fixed-marginal constraint for correlated variability using our formal-
643 ism then define the commonly used shuffle and novel uniform correlation null models. Finally, we propose
644 the factor analysis covariance parameterizations and associated null model for assessing optimality which
645 has more biological interpretability. In the following sections we will use the following terminology which
646 we define here:

- 647 • **Covariance Parameterization:** a parameterization of Σ which can combine various constraints
648 (fixed parameters) and degrees-of-freedom (free parameters).
- 649 • **Constraints:** elements of the covariance parameterization which are estimated from data and fixed.
- 650 • **Degrees-of-Freedom:** elements of the covariance parameterization which can potentially be modified
651 or optimized to analyze a null model or optimality.
- 652 • **Optimality:** values for the degrees-of-freedom in a covariance parameterization which maximize a
653 specified objective. Here we assess optimality using the Linear Fisher Information (LFI), although
654 this formalism can be applied to other objectives.
- 655 • **Null Distribution:** distribution of a covariance parameterization's degrees-of-freedom.
- 656 • **Null Model:** combines a covariance parameterization with a baseline or uniform correlation null
657 distribution over the degrees-of-freedom.

658 The standard constraint considered for understanding correlated neural variability is to keep the per-
659 neuron marginal distributions fixed. Since the LFI only depends on the covariance of the correlated variabil-
660 ity, the fix-marginal parameterization is equivalent to constraining the per-neuron variances to be constant
661 (equivalently, the diagonal of Σ is kept constant, $\text{diag}(\Sigma) = \sigma^2$). The corresponding degrees-of-freedom
662 in this parameterization are the positive-definite pairwise correlation matrix, ρ , specifically the symmetric,
663 off-diagonal entries, ρ_{ij} for $i \neq j$, which can vary. Under this parameterization, the observed covariance
664 structure can be compared to other proposed distributions of correlations.

665 When considering the structure that generates Σ , it is desirable that the constraints and degrees-of-
666 freedom be biologically interpretable. This can be achieved by considering the equations that define the
667 mean-centered, single-trial response in terms of the degrees-of-freedom being considered. For the fixed-
668 marginals parameterization, the distribution of the single-trial responses: $\mathbf{f}_t(s)$, can be written in terms of a
669 multivariate normal distribution with the mean response: $\mathbf{f}(s)$, and where the covariance is the element-wise
670 product of the constrained marginal standard deviations: $\sigma\sigma^T$, and the free correlations: ρ ,

$$\begin{aligned} \mathbf{f}_t(s) &= \mathbf{f}(s) + \epsilon \\ \epsilon &\sim \mathcal{N}(0, \sigma\sigma^T \odot \rho). \end{aligned} \tag{3}$$

671 This equation is difficult to directly interpret as a network model, but the correlations could be seen as
672 coming from recurrent activity within the observed neurons.

673 Given a parameterization (fixed-marginal) and a measure of coding fidelity (LFI), it is possible to find
674 optimal covariance structures as a function of the free parameters. In general, the value (or distribution of

675 values) for the degrees-of-freedom that lead to optimality can be derived analytically or optimized numeri-
676 cally. For the fixed-marginal parameterization, this corresponds to finding the points, $\hat{\rho}$, such that

$$\hat{\rho} = \arg \max_{\rho} \text{LFI} \left(\frac{d\mathbf{f}(s)}{ds}, \text{diag}(\Sigma), \rho \right). \quad (4)$$

677 Hu *et al.* [24] characterize the optima of the fixed-marginal parameterization, although they do not provide
678 a constructive way of finding the global optima. We optimize ρ numerically to find optima. We find that the
679 optimization process finds many local maxima for $\hat{\rho}$ in practice.

680 **Novel null models allow the assessment of optimality in neural data**

681 So far, we have laid out a formalism to define the optimal degrees-of-freedom for a specified covari-
682 ance parameterization. However, it is unlikely that observed neural data will precisely match the predicted
683 optimal degrees-of-freedom, even if the biological system is behaving optimally, so the predictions from
684 Eq 4 cannot be used directly to assess optimality in data. In order to assess the optimality of a observed
685 population of neurons, a null model must be constructed for a corresponding parameterization. In this for-
686 malism, constructing a null model corresponds to assuming a null distribution for the degrees-of-freedom
687 of the covariance parameterization. The null distribution should correspond to some notion of “uniform” or
688 “baseline” for the degrees-of-freedom.

689 For example, the shuffle null model, based on the fixed-marginal parameterization, posits that the base-
690 line distribution of correlations is zero correlations. The shuffle null model compares the LFI of the observed
691 response to the distribution of LFIs where the individual neural responses are independently trial shuffled,
692 that is, with fixed-marginal variability, no underlying pairwise correlations, and empirical pairwise corre-
693 lations only arising from finite sampling effects. Under this choice of null model, the observed LFI can
694 be beneficial if it has a high percentile under the null distribution which has no correlations. The shuffle
695 null model provides a limited baseline comparison for the observed LFI. In order to assess optimality, the
696 distribution of parameters should be uniform over the space of allowed covariance matrices, which is the
697 motivation for the uniform correlation null model.

698 Across a population, the median observed percentile across dim-stims can be used to categorize a dataset
699 as optimal: median percentile greater than or equal to 2/3, near-chance: median percentile between 1/3 and
700 2/3, or suboptimal: median percentile less than 1/3. This categorization is motivated by simplicity in having
701 few categories. However, it is also desirable to not have the optimal and suboptimal categories share a
702 boundary. If they do, small changes in percentiles can switch between optimal and suboptimal. In our case,
703 since the null model defines “near-chance”, having 3 categories is natural. The near-chance boundaries could
704 be set in a number of ways besides the choice for an even division into thirds. A Kolmogorov–Smirnov test
705 could compare the distribution of percentiles to a uniform distribution. However, given the large number
706 of dim-stims we use, empirically, no distributions of percentiles in these datasets would be near-chance for
707 p-value thresholds in sensible ranges. Said another way, almost no empirical distributions of percentiles are
708 statistically similar to a uniform distribution (see Supplementary **Fig. 5a-i** for some example distributions).
709 A looser test could be to test whether a binomial distribution with $p = 0.5$ would lead to the observed
710 distribution of percentiles categorically above and below 0.5. We find that with p-values in sensible ranges
711 this gives comparable boundaries to the division into thirds, but the boundaries differ across datasets due to
712 the variation in the number of dim-stims.

713 In some cases, it may also be possible to define a distribution over optimal covariances and categorize
714 whether the observed LFI is likely under the optimal covariance distribution. For instance, if there is a unique
715 optimal covariance, the Wishart distribution could be used to create a sampling distribution of optimal LFIs
716 which the observed LFIs could be compared against. This is not possible in our case since there is not

717 generally a unique optimal covariance. This also suffers from the fragility problem by having a boundary
718 directly between optimal and suboptimal.

719 **Uniform correlation null model**

720 Our first contribution is the uniform correlation null model based on the fixed-marginal parameterization,
721 where the correlations are chosen randomly from a uniform distribution over correlation matrices [25]. This
722 tests whether the observed correlation are optimal with respect to all possible correlations, rather than only
723 comparing against zero correlations. To our knowledge, this null model has not been considered before.
724 Evaluating data under this null model provides a stronger assessment of the optimality of the observed
725 correlated variability than the shuffle null model.

726 At another extreme, we could attribute all trial-to-trial variability to external sources that the network
727 can shape or filter. To prevent trivial solutions, we can restrict the network to only changing the loading of
728 the variability onto the neurons (through a rotation, \mathbf{R}). This model was previously discussed [24], but not
729 analyzed due to its incompatibility with the fixed-marginal constraint.

730 **Factor analysis null model**

731 As a parsimonious combination of the fixed-marginal constraint and pure rotation degrees-of-freedom, we
732 propose using a factor analysis (FA) model to parameterize the correlated variability. Factor analysis de-
733 composes the observed correlated variability into two components: the first is per-neuron private variability,
734 represented as a diagonal matrix $\text{diag}(\sigma_{\text{FA}}^2)$, and the second is a low-rank shared variability component,
735 $\mathbf{L}_{\text{FA}}^T \mathbf{L}_{\text{FA}}$, where $\mathbf{L}_{\text{FA}} \in \mathbb{R}^{k \times d}$, $k < d$. We propose that the FA model has private variability and the spectrum
736 of the shared component as constraints and the rotation of the shared components as the degrees-of-freedom,
737 combining aspects of the fixed-marginal and rotation null models. The single-trial response can be written
738 as a function of the mean response: $\mathbf{f}(s)$, private variances: σ_{FA}^2 , low-rank external sources: \mathbf{z}_{FA} , loading
739 matrix: \mathbf{L}_{FA} , and rotation matrix: \mathbf{R}

$$\begin{aligned} \mathbf{f}_t(s) &= \mathbf{f}(s) + \mathbf{R}^T \mathbf{L}_{\text{FA}}^T \mathbf{z}_{\text{FA}} + \epsilon_{\text{FA}} \\ \mathbf{z}_{\text{FA}} &\sim \mathcal{N}(0, \mathbf{1}) \\ \epsilon_{\text{FA}} &\sim \mathcal{N}(0, \text{diag}(\sigma_{\text{FA}}^2)) \end{aligned} \tag{5}$$

740 To our knowledge, there is no closed-form solution for $\hat{\mathbf{R}}$ in the FA model to maximize LFI. Instead, to
741 optimize the FA model, the rotation can be numerically optimized by gradient ascent. To construct the FA
742 null model, a uniform distribution (Haar distribution) over special orthogonal rotations [26] is applied to the
743 rotations.

744 To estimate the initial σ_{FA}^2 and \mathbf{L}_{FA} , we fit a factor analysis model to the samples [27]. In fitting the
745 model we had two requirements. The first is that we wanted the dimensionality of the shared component,
746 k to be as large as possible so that the observed covariance can be modeled as accurately as possible. In
747 opposition to this, we wanted the factor analysis model parameters to be identifiable, meaning the private
748 variance estimate is unique, which places a limit, which depends on d , on how large k can be [28]. In
749 practice, we find the largest k which is lower than the identifiability bound where different initializations
750 return the same parameters. Note that factor analysis is never identifiable in 2 dimensions, so we do not
751 consider $d = 2$.

752 **Population statistics across dim-stims measure optimality under a null model**

753 Each dataset can be described by a $D \times N$ design matrix \mathbf{X} , where D is the total number of samples
754 and N is the number of units in the population (**Fig. 2f**). We considered distributions of LFI across dim-

755 stims, or sub-components of the design matrix. To create dim-stims, we first selected a dimlet of size d
 756 by subsampling d units from the population at random, resulting in the $D \times d$ design matrix \mathbf{X}^d (**Fig. 2f**).
 757 Next, we created the dim-stim by further subsampling the design matrix according to a specific stimulus
 758 pairing. Specifically, we chose two neighboring stimuli, s_1 and s_2 (**Fig. 2f**), and isolated the samples of \mathbf{X}^d
 759 corresponding to those stimuli, thereby creating a pair of design matrices $[\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d]$. The dim-stim maps to
 760 the task of discriminating between two neighboring stimuli using a sub-population's responses across trials
 761 to those stimuli, which can be visualized in the neural space (**Fig. 2g**).

762 For each dataset, we considered dimlet dimensions $d = 3-20$. As we only allowed neighboring stimulus
 763 pairings, the number of available stimulus pairings for a dimlet was 6 (retinal), 12 (V1) and 29 (PAC). Note
 764 that the retinal and V1 stimulus sets are circular, providing an additional stimulus pairing. In the retinal and
 765 V1 datasets, we drew 1,000 dimlets for each dimension d , and considered all stimulus pairings per dimlet,
 766 resulting in $1,000 \times 6 = 6,000$ dim-stims for the retinal dataset and $1,000 \times 12 = 12,000$ dim-stims for
 767 the V1 dataset. To manage computation time, we considered 3,000 unique dim-stims for the PAC dataset,
 768 selecting both the dimlet and stimulus pairing at random for each dim-stim.

For each dim-stim, we calculate its observed LFI, defined as $\mathcal{I}_{\text{coarse}}(\mathbf{f}_1, \mathbf{f}_2, \boldsymbol{\Sigma}_1, \boldsymbol{\Sigma}_2)$. Specifically, we
 computed

$$\mathcal{I}_{\text{obs}}(\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d) = \mathcal{I}_{\text{coarse}}\left(\text{mean}(\mathbf{X}_{s_1}^d), \text{mean}(\mathbf{X}_{s_2}^d), \text{cov}(\mathbf{X}_{s_1}^d), \text{cov}(\mathbf{X}_{s_2}^d)\right) \quad (6)$$

$$= \left(\frac{\mathbf{f}_{s_1}^d - \mathbf{f}_{s_2}^d}{\Delta s}\right)^T \left(\frac{\boldsymbol{\Sigma}_{s_1}^d + \boldsymbol{\Sigma}_{s_2}^d}{2}\right)^{-1} \left(\frac{\mathbf{f}_{s_1}^d - \mathbf{f}_{s_2}^d}{\Delta s}\right) \quad (7)$$

769 where $[\mathbf{f}_{s_1}^d, \mathbf{f}_{s_2}^d]$ are the dim-stim average responses, $[\boldsymbol{\Sigma}_{s_1}^d, \boldsymbol{\Sigma}_{s_2}^d]$ are the dim-stim covariances, and Δs is the
 770 stimulus difference, or $\Delta s = |s_1 - s_2|$. When necessary, the stimulus difference was taken as a circular
 771 difference (retinal and V1 datasets). Since the LFI is scaled by the units of the stimulus difference, it is only
 772 meaningful to compare observed LFIs within a particular stimulus type. In this work, since all datasets use
 773 a different stimulus the LFIs may not have a meaningful relationship across datasets.

Each null model acts on the design matrices of a dim-stim and outputs a distribution of covariance
 matrices. For example, the fixed-marginal null model shuffles the data within the design matrix, producing
 new design matrices $[\mathbf{X}_{s_1}^{d'}, \mathbf{X}_{s_2}^{d'}]$ and corresponding covariances $[\boldsymbol{\Sigma}_{s_1}^{d'}, \boldsymbol{\Sigma}_{s_2}^{d'}]$. We then calculate the LFI
 using the new covariance matrices. Each null model can be summarized as such: a sampled transformation
 is applied to the observed dim-stim, producing new sampled covariance matrices and therefore a sample of
 LFI from the null. The shuffle null model transformed the data directly, so we write its LFI as

$$\mathcal{I}_{\text{FM}}(\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d) = \mathcal{I}_{\text{obs}}\left(\text{shuffle}(\mathbf{X}_{s_1}^d), \text{shuffle}(\mathbf{X}_{s_2}^d)\right). \quad (8)$$

Meanwhile, the uniform and factor analysis null models transform the covariance parameterization directly,
 so we write their LFIs as:

$$\mathcal{I}_{\text{U}}(\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d) = \mathcal{I}_{\text{coarse}}\left(\mathbf{f}_{s_1}^d, \mathbf{f}_{s_2}^d, \text{sample}_{\text{U}}(\boldsymbol{\Sigma}_{s_1}^d), \text{sample}_{\text{U}}(\boldsymbol{\Sigma}_{s_2}^d)\right) \quad (9)$$

$$\mathcal{I}_{\text{FA}}(\mathbf{X}_{s_1}^d, \mathbf{X}_{s_2}^d) = \mathcal{I}_{\text{coarse}}\left(\mathbf{f}_{s_1}^d, \mathbf{f}_{s_2}^d, \text{rotate}_{\text{FA}}(\boldsymbol{\Sigma}_{s_1}^d), \text{rotate}_{\text{FA}}(\boldsymbol{\Sigma}_{s_2}^d)\right). \quad (10)$$

774 Equations 8 and 10 capture a single application of a null model. Specifically, $\text{shuffle}(\cdot)$ shuffles the
 775 neural data, $\text{sample}_{\text{U}}(\cdot)$ samples a random off-diagonal correlation structure and applies it to the covariance,
 776 and $\text{rotate}_{\text{FA}}(\cdot)$ applies a rotation to the shared component of the covariance. However, we were interested in
 777 characterizing the entire distribution of the null model. Thus, for each dim-stim, we applied 1,000 samples
 778 of the null model to obtain a null model distribution of LFIs. We then calculated observed percentiles as the

779 fraction of samples for which the observed LFI exceeded the null model LFI. Thus, each observed dim-stim
780 has its own corresponding observed percentile, per null model.

781 When summary statistics are reported such as the median LFI, median percentile, or the optimal fraction,
782 95% bootstrap confidence intervals from 1,000 bootstrap resamples are reported [29].

783 **Optimal fraction calculation**

784 The optimal fraction of a population was calculated in the following way. Given a set of dim-stims at a par-
785 ticular dimlet dimension, the observed percentiles were calculated for each dim-stim. Then, the percentiles
786 were sorted from largest to smallest. The optimal fraction of the percentiles is initialized as the largest single
787 percentile. Starting from this initialization, the median percentile of the current optimal fraction is calcu-
788 lated. If the median is greater than or equal to $2/3$, the next smallest percentile is included in the optimal
789 fraction and the process continues to iterate. If the optimal fraction is less than $2/3$, the process terminates.
790 This defines the largest possible fraction of the percentiles that can be retained and have their median be
791 greater than or equal to $2/3$. For reference, the top $2/3$ of a uniform distribution (i.e., $[1/3, 1]$) of percentiles
792 has median equal to $2/3$.

793 **Measures of biological plausibility**

794 We calculated the mean Fano factors (FF) for a dim-stim, based on the per-unit variance and response means

$$FF = \frac{1}{d} \sum_{i=1}^d \frac{\Sigma_{ii}(s)}{f(s)_i}, \quad (11)$$

795 of the observed and optimal covariances matrices directly from the mean response and covariance matrix
796 parameters (Supplemental **Fig. 2**).

797 We calculated the negative density (ND) as follows. For each dim-stim, we calculated $f_i^{1\%}$, the neural
798 activity at the 1st percentile, for each neuron i . We then computed $CDF_i(f_i^{1\%})$, the cumulative density at
799 $f_i^{1\%}$ for a Gaussian obtained from either the observed covariance or the optimal covariance under the null
800 model (Supplemental **Fig. 2**, shaded regions in marginals). The ND, then, was defined as the maximum
801 CDF_i among the neurons in the dimlet (Supplemental **Fig. 2**, dark gray shaded regions).

802 **Distance and tuning ranking dim-stims for subselection**

803 For the retina and PAC datasets, we have access to the spatial locations of the RGC/electrode. For distance-
804 based subselection, we compute the average pairwise distance between neural units for each dim-stim. The
805 dim-stims are ranked by this distance and the 10% of dim-stims with the smallest average distance are
806 subselected.

807 For tuning-based subselection, the stimuli are ranked for each neural unit based on the mean neural
808 activity (tuning). The rank was used because is less sensitive to absolute firing rates compared to using
809 the activity per stimuli, which would biased the subselection towards dim-stims which contain neural units
810 with high firing rates. We then sort the dim-stims by their average tuning rank across dimlets and calculate
811 percentile statistics for the 10% of dim-stims that have the highest tuning ranking.

812 References

- 813 1. Tiriac, A., Smith, B. E. & Feller, M. B. Light prior to eye opening promotes retinal waves and eye-
814 specific segregation. *Neuron* **100**, 1059–1065 (2018).
- 815 2. Stosiek, C., Garaschuk, O., Holthoff, K. & Konnerth, A. In vivo two-photon calcium imaging of neu-
816 ronol networks. *Proceedings of the National Academy of Sciences* **100**, 7319–7324 (2003).
- 817 3. Pologruto, T. A., Sabatini, B. L. & Svoboda, K. ScanImage: flexible software for operating laser scan-
818 ning microscopes. *Biomedical engineering online* **2**, 1–9 (2003).
- 819 4. Caval-Holme, F., Zhang, Y. & Feller, M. B. Gap junction coupling shapes the encoding of light in the
820 developing retina. *Current Biology* **29**, 4024–4035 (2019).
- 821 5. Tiriac, A., Bistrong, K. & Feller, M. Retinal waves but not visual experience are required for develop-
822 ment of retinal direction selectivity maps. *bioRxiv* (2021).
- 823 6. Teeters, J. L., Harris, K. D., Millman, K. J., Olshausen, B. A. & Sommer, F. T. Data Sharing for
824 Computational Neuroscience. *Neuroinformatics* **6**, 47–55 (Mar. 2008).
- 825 7. Kohn, A. & Smith, M. A. *Utah array extracellular recordings of spontaneous and visually evoked*
826 *activity from anesthetized macaque primary visual cortex (V1)* 2016. [http://dx.doi.org/10.](http://dx.doi.org/10.6080/K0NC5Z4X)
827 [6080/K0NC5Z4X](http://dx.doi.org/10.6080/K0NC5Z4X).
- 828 8. Smith, M. A. & Kohn, A. Spatial and temporal scales of neuronal correlation in primary visual cortex.
829 *The Journal of Neuroscience* **28**, 12591–603 (2008).
- 830 9. Kelly, R. C., Smith, M. A., Kass, R. E. & Lee, T. S. Local field potentials indicate network state and
831 account for neuronal response variability. *Journal of computational neuroscience* **29**, 567–579 (2010).
- 832 10. Dougherty, M. E., Nguyen, A. P. Q., Baratham, V. L. & Bouchard, K. E. *Laminar origin of evoked*
833 *ECoG high-gamma activity in 2019 41st Annual International Conference of the IEEE Engineering in*
834 *Medicine and Biology Society (EMBC)* (July 2019), 4391–4394.
- 835 11. Baratham, V. L., Dougherty, M. E., Ledochowitsch, P., Maharbiz, M. M. & Bouchard, K. Columnar
836 localization and laminar origin of cortical surface electrical potentials. *bioRxiv* (2021).
- 837 12. Cover, T. M. & Thomas, J. A. *Elements of information theory* (John Wiley & Sons, 2012).
- 838 13. Abbott, L. F. & Dayan, P. The effect of correlated variability on the accuracy of a population code.
839 *Neural computation* **11**, 91–101 (1999).
- 840 14. Sompolinsky, H., Yoon, H., Kang, K. & Shamir, M. Population coding in neuronal systems with cor-
841 related noise. *Physical Review E* **64**, 051904 (2001).
- 842 15. Yarrow, S., Challis, E. & Seriès, P. Fisher and Shannon information in finite neural populations. *Neural*
843 *computation* **24**, 1740–1780 (2012).
- 844 16. Zylberberg, J., Cafaro, J., Turner, M. H., Shea-Brown, E. & Rieke, F. Direction-selective circuits shape
845 noise to ensure a precise population code. *Neuron* **89**, 369–383 (2016).
- 846 17. Franke, F. *et al.* Structures of neural correlation and how they favor coding. *Neuron* **89**, 409–422
847 (2016).
- 848 18. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and neuronal population infor-
849 mation. *Annual review of neuroscience* **39**, 237–256 (2016).
- 850 19. Sachdeva, P. S., Livezey, J. A. & DeWeese, M. R. Heterogeneous synaptic weighting improves neural
851 coding in the presence of common noise. *Neural computation* **32**, 1239–1276 (2020).

- 852 20. Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences* **112**, E6973–E6982 (2015).
853
- 854 21. Kafashan, M. *et al.* Scaling of information in large neural populations reveals signatures of information-
855 limiting correlations. *bioRxiv* (2020).
- 856 22. Deweese, M. R. & Zador, A. M. Shared and private variability in the auditory cortex. *Journal of*
857 *neurophysiology* **92**, 1840–1855 (2004).
- 858 23. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**
859 (2019).
- 860 24. Hu, Y., Zylberberg, J. & Shea-Brown, E. The sign rule and beyond: boundary effects, flexibility, and
861 noise correlations in neural population codes. *PLoS computational biology* **10**, e1003469 (2014).
- 862 25. Joe, H. Generating random correlation matrices based on partial correlations. *Journal of Multivariate*
863 *Analysis* **97**, 2177–2189 (2006).
- 864 26. Stewart, G. W. The efficient generation of random orthogonal matrices with an application to condition
865 estimators. *SIAM Journal on Numerical Analysis* **17**, 403–409 (1980).
- 866 27. Pedregosa, F. *et al.* Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*
867 **12**, 2825–2830 (2011).
- 868 28. Bekker, P. A. & ten Berge, J. M. Generic global identification in factor analysis. *Linear Algebra and*
869 *its Applications* **264**, 255–263 (1997).
- 870 29. Virtanen, P. *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature*
871 *Methods* **17**, 261–272 (2020).

872 Appendix

873 Geometric contributions to neural correlated variability

874 The geometry of three types of potential contributions to neural variability are shown. Private variability
875 is a zero-correlation contribution (Fig. 1a) [1]. Shared variability can be a low-rank contribution whose
876 orientation depends on the synaptic loading onto the observed neural units (Fig. 1b) [2, 3]. Differential
correlations lie along the $\frac{df(s)}{ds}$ direction (Fig. 1c) [4].




Figure 1: Geometric contributions to neural correlated variability. Each plot depicts the neural space, whose axes correspond to the activities of a specific pair of neurons to a stimulus. Black curves denote the mean responses across different stimuli (i.e., tuning curves). Variability about a specific stimulus mean activity (solid points) may exhibit: **a.** Private, uncorrelated variability in each neural dimension, **b.** Correlated variability, with correlations in the neural space, and **c.** Differential correlations, which lie parallel to the mean activity curve.

877

878 Measures for assessing biological accessibility

879 Consider an example V1 dim-stim for a dimlet of size $d = 3$, with low observed percentiles under both the
 880 null models (e.g., $p_U = 0.001$ and $p_{FA} = 0.0$). We plot the observed covariance structure, projected into
 881 two neural dimensions, in Figure 2a (black covariance denotes average covariance). Next, we compare the
 882 observed structure to that of the optimal structure, both within the factor analysis null model (Fig. 2b) and
 883 the uniform null model (Fig. 2c).

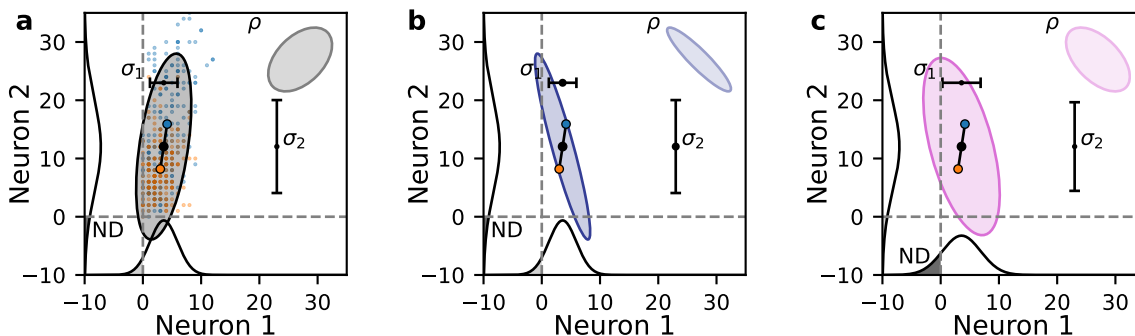


Figure 2: Measures for assessing biological accessibility. Data, fit and optimal covariances (at 2 standard deviations) are from a $d = 3$ dimlet-stim projected into the first 2 neurons. The marginal probabilities of the multivariate Gaussian fits are shown along the axes and the areas with values less than the empirical 1% are shaded grey with the maximum excess negative density in dark grey (annotated with “ND”). The marginal means and standard deviations (for Fano factor calculations) are shown in the black error bars (annotated with “ σ ” and neuron number). For each covariance, the corresponding correlation ellipse (ρ , with an arbitrary uniform scaling) is shown in the top right of the plot. **a:** Observed single-trial neuron responses to stimuli 1 and 2 (orange and blue dots) and the respective means (outlined circles). Their joint mean is the black circle and the observed mean covariance is in gray. **b:** Covariance and marginals from an optimal fixed-marginal correlation. **c:** Covariance and marginals from the optimal Factor Analysis rotation.

884 The observed correlated variability structure (Fig. 2a) exhibits poor discriminability, because a large
 885 amount of variability is oriented parallel to the stimulus manifold (Fig. 2, black lines in the empirical covari-
 886 ance ellipse). We consider several measures of biological plausibility for the optimal covariances. The first
 887 is the median absolute correlation of the optimal covariances (Fig. 2, ellipse labeled ρ in top right shows
 888 optimal correlation), which is most relevant for the uniform correlation null model. The second is the
 889 Fano factors (FF) of the optimal covariance relative to the Fano factors of the observed covariance (Fig. 2,
 890 black mean and standard deviation indicators labeled with σ_1 and σ_2). The third is the cumulative marginal
 891 probability the optimal covariance has below the 1st percentile of the observed data (Fig. 2, gray regions
 892 in marginal distributions labeled ND, negative density). These measures only take on a limited range of
 893 values in measured neural activity, and may impede a neural system from obtaining an optimal correlated
 894 variability structure. The uniform correlation null model preserves the per-RGC/neuron/electrode mean and
 895 variance, and so the FF and ND measures are only relevant for the factor analysis null model.

896 However, the optimal covariance orientations for the factor analysis model may possess different Fano
 897 factors (Fig. 2c). Thus, we aimed to assess whether biologically unachievable Fano factors shared any
 898 relation with the sub-optimality exhibited by the neural codes in our analyses. We summarized each dim-
 899 stim with an aggregate Fano factor, by averaging the Fano factors of that dim-stim’s individual units. We
 900 repeated this process for the optimal noise covariances under each null model, using the variances from the
 901 diagonal of the optimal noise covariance matrix directly when calculating Fano factors.

902 To quantify this phenomenon, we calculated the absolute difference in negative density (ND), which

903 captures the degree to which an optimal covariance puts differing cumulative density in the negative response
904 space (typically higher density). Thus, a larger ND implies that the covariance places an excess of density
905 in the negative or low-activity regions for at least one dimension of the neural space. On the other hand, a
906 lower ND is more biologically plausible, as this implies there is less negative density, although Gaussian fits
907 will always put some non-zero density in the negative.

908 **Optimal correlations for the fixed-marginal parameterization lie on the bound-**
909 **ary of possible correlations**

910 Hu *et al.* [5] show analytically that optimal covariances in the fixed-marginal parameterization lie on the
911 boundary of allowed correlations, which will generally have large absolute pairwise correlations. We repro-
912 duce this results computationally. For each dim-stim, we compare the 90% percentile of the off-diagonal
913 entries absolute correlation matrix for the observed covariance matrix, the optimal uniform correlation (UC)
914 null model matrix, and the optimal factor analysis (FA) null model matrix. The histograms across dim-stims
915 for 4 dimensions is shown in Figure 3. The observed 90% abs. correlations are rarely larger than 0.7. The FA
916 optimal 90% abs. correlations have a larger spread towards higher correlations, but do not have density at 1.
917 However, the UC optimal covariance have 90% abs. correlations that consistently have peaks in probability
918 mass at 1.

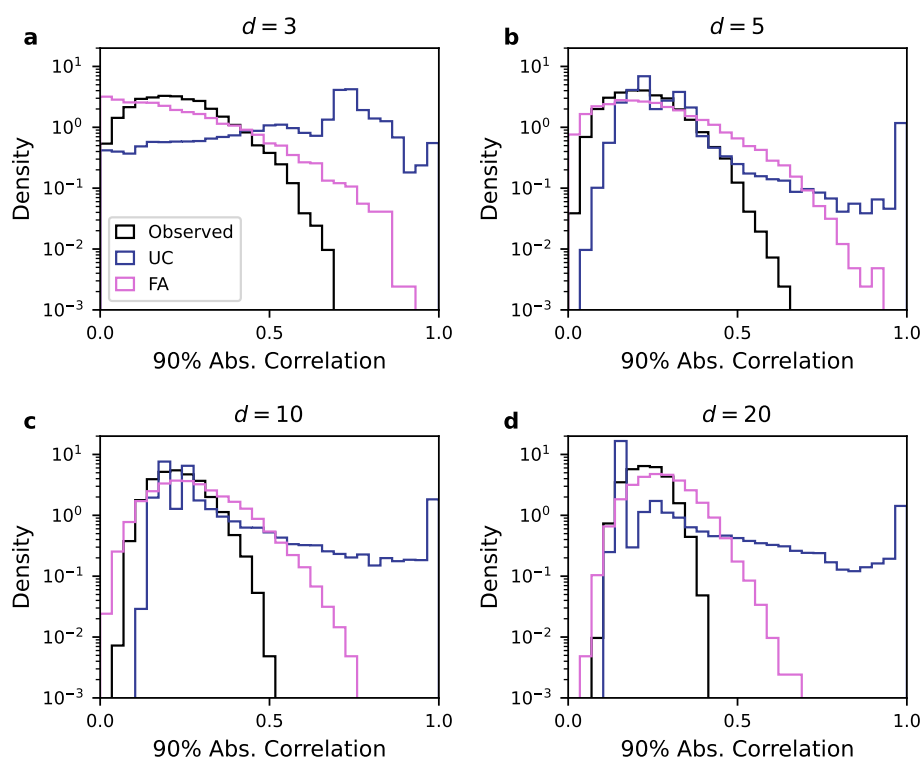


Figure 3: Optimal correlations for the fixed-marginal parameterization lie on the boundary of possible correlations. For each dimension, d , the 90th percentile of the absolute value of the pairwise correlations is histogrammed across dim-stims. Color indicates whether the statistic is from the observed covariance or optimal null model covariance. **a-d.** Dimensions 3, 5, 10, and 20, respectively are shown.

919 **Biologically motivated subselection of dim-stims remains suboptimal**

920 In the main text, biological subselection in Figure 5 was done based on the distance- and tuning-based
921 criteria as they might correspond to biological criteria enforced during development or learning. It is also
922 possible to subselect the dim-stims using the Fano factor (FF) and negative density (ND) criteria directly for
923 the factor analysis null model. Here, we compute the average rank the dim-stims based on their violation of
924 the FF and ND criteria and retain the 10% of dim-stims with the least average violation. This criteria leaves
925 the population percentiles suboptimal (Fig. 4a-c).




Figure 4: Biologically motivated subselection of dim-stims remains suboptimal. The dim-stims are subselected based on biological criteria and their median percentiles are shown as a function of dimension. Black dashed lines indicate the 33rd-66th percentile range. Shaded regions bound the 40th to 60th percentiles of the subselected percentile distributions. **a-c.** For the FA null model, dim-stims were subselected to minimize their average Fano factor and ND deviations (0th-10th percentile). The median and 33-66% of the percentiles for this subpopulation is shown.

926 **There is an exponentially small peak of optimal dim-stims for the factor anal-**
927 **ysis null model**

928 For the factor analysis null model, there is sometimes a peak of percentiles near 1 (Fig. 5a-i). For some
929 dimensions, the peak has higher density than what would be expected from a uniform distribution. To calcu-
930 late the peak width at each dimension, the percentiles are sorted and, starting from the largest percentiles, the
931 observed percentiles are compared with the percentiles that would be expected from a uniform distribution.
932 The peak width is the fraction of percentiles corresponding to the smallest percentile that has a value larger
933 than what is expected from a uniform distribution. Across datasets, the peak width is exponentially small as
934 a function of neural dimension (Fig. 5j-l). The uniform correlation null model does not have peaks near 1
935 for any dimension or dataset.

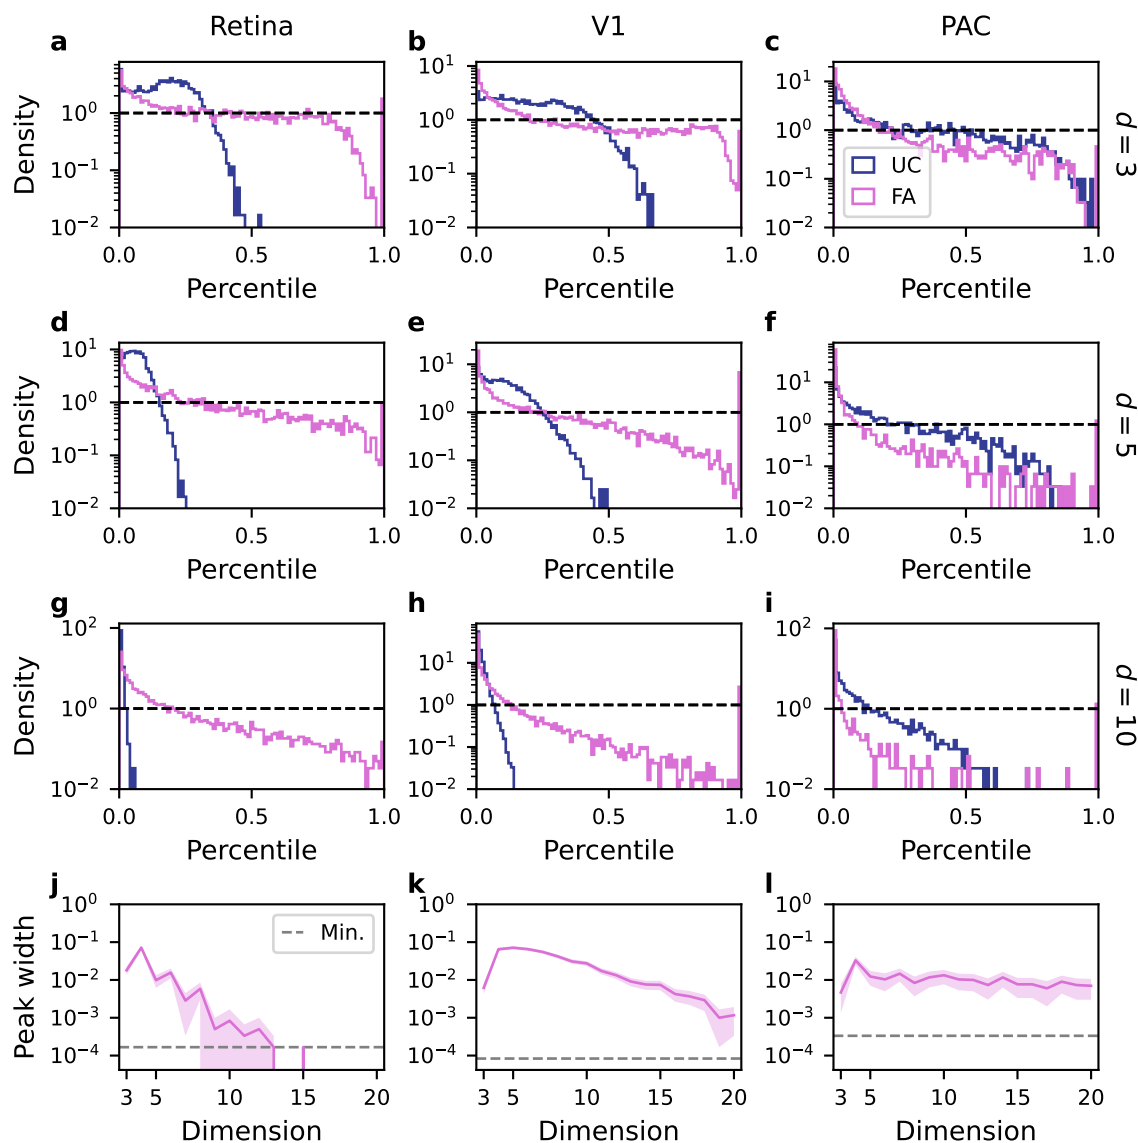


Figure 5: There is an exponentially small peak of optimal dim-stims for the factor analysis null model.

a-i. The histograms of the percentiles distributions are shown across datasets and null models for dimensions 3, 5, and 10 (in rows). Black dashed lines indicate the density of a uniform distribution. Note the y-axis is log-scaled. **j-l.** Across dimensions, the width of the greater-than-uniform peak is shown. Shaded regions are the 95% CI for the peak widths. Gray dashed line indicates the minimum non-zero peak width that can be estimated due to finite sampling.

936 **V1 datasets give similar results across monkeys**

937 The PVC11 dataset (here V1) from CRCNS has data from 3 different monkeys [6]. In the main text, we
938 used monkey 1. Although there are differences in the distribution of pairwise correlations (Fig. 6a), they do
939 not lead to qualitative differences in the results from the main text across animals. Figure 6b-m reproduce
940 the main results for all 3 monkeys.

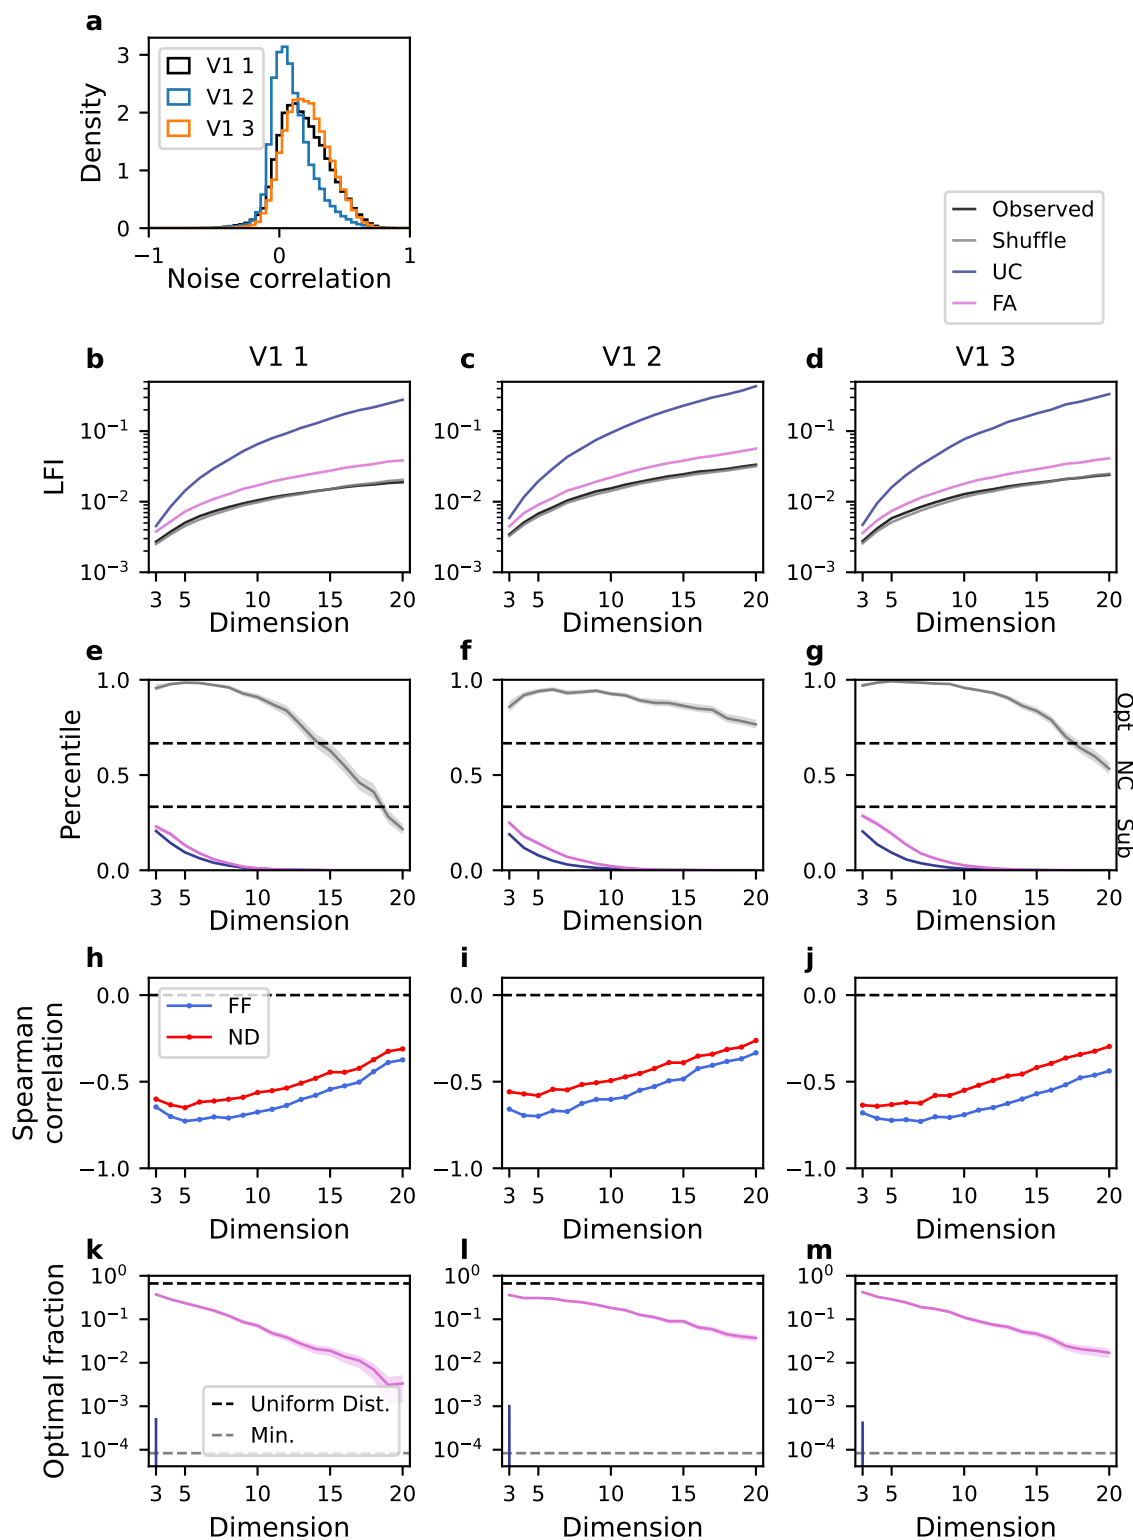


Figure 6: V1 datasets give similar results across monkeys. Main text results are repeated for monkeys 2 and 3 (V1 2 and V1 3, second and third columns) and compared with monkey 1 (V1 1, first column) which is reproduced here. Panel **a** corresponds to main text Figure 1. Panels **b-g** correspond to main text Figure 3. Panels **h-j** correspond to main text Figure 4. Panels **k-m** correspond to main text Figure 5. See main text for panel details.

941 **References**

- 942 1. Deweese, M. R. & Zador, A. M. Shared and private variability in the auditory cortex. *Journal of*
943 *neurophysiology* **92**, 1840–1855 (2004).
- 944 2. Sachdeva, P. S., Livezey, J. A. & DeWeese, M. R. Heterogeneous synaptic weighting improves neural
945 coding in the presence of common noise. *Neural computation* **32**, 1239–1276 (2020).
- 946 3. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**
947 (2019).
- 948 4. Moreno-Bote, R. *et al.* Information-limiting correlations. *Nature neuroscience* **17**, 1410–1417 (2014).
- 949 5. Hu, Y., Zylberberg, J. & Shea-Brown, E. The sign rule and beyond: boundary effects, flexibility, and
950 noise correlations in neural population codes. *PLoS computational biology* **10**, e1003469 (2014).
- 951 6. Kohn, A. & Smith, M. A. *Utah array extracellular recordings of spontaneous and visually evoked*
952 *activity from anesthetized macaque primary visual cortex (V1)* 2016. [http://dx.doi.org/10.](http://dx.doi.org/10.6080/K0NC5Z4X)
953 [6080/K0NC5Z4X](http://dx.doi.org/10.6080/K0NC5Z4X).