

Selective effects of arousal on population coding of natural sounds in auditory cortex

Charles R. Heller^{1,2}, Zachary P. Schwartz², Daniela Saderi², Stephen V. David^{2*}

¹ Neuroscience Graduate Program, Oregon Health and Science University

² Oregon Hearing Research Center, Oregon Health and Science University

*Corresponding author (davids@ohsu.edu)

Abstract

The ability to discriminate between complex natural sounds is critical for survival. Changes in arousal and other aspects of behavioral state can impact the accuracy of sensory coding, affecting both the reliability of single neuron responses and the degree of correlated noise between neurons. However, it is unclear how these effects interact to influence coding of diverse natural stimuli. We recorded the spiking activity of neural populations in primary auditory cortex (A1) evoked by a large library of natural sounds while monitoring changes in pupil size as an index of arousal. Heightened arousal increased response magnitude and reduced noise correlations between neurons, improving coding accuracy on average. Rather than suppressing shared noise along all dimensions of neural activity, the change in noise correlations occurred via coherent, low-dimensional modulation of response variability in A1. The modulation targeted a different group of neurons from those undergoing changes in response magnitude. Thus, changes in response magnitude and correlation are mediated by distinct mechanisms. The degree to which these low-dimensional changes were aligned with the high-dimensional natural sound-evoked activity was variable, resulting in stimulus-dependent improvements in coding accuracy.

Introduction

Humans and other animals are able to discriminate between a multitude of natural sounds. This ability is not static, as the precision of sensory representations by neural activity fluctuates with changes in behavioral state.¹ Arousal, task engagement, and attention have all been

1

2

3

4

reported to modulate the magnitude and selectivity of single neuron auditory responses,²⁻¹² 5
as well as correlated variability across neural populations, often referred to as noise correla- 6
tions.¹³⁻¹⁶ In general, increased arousal and focused attention are associated with increased 7
response magnitude and decreased noise correlations which are believed to enhance the accu- 8
racy of sensory coding.^{1,13,16,17} However, the mechanisms that produce these changes, and 9
the consistency of their effects between different behavioral contexts, are not fully understood. 10

Recent studies have argued that attention-driven changes in both single neuron responses 11
and correlated activity can be modeled as fluctuations in a single, latent signal that coherently 12
modulates the activity of a subset of neurons. These findings suggest that state-dependent 13
neural population activity occurs in a low-dimensional subspace,^{18,19} supporting theoretical 14
models in which a single mechanism accounts for changes in single neuron responses and cor- 15
related variability.^{20,21} Fluctuations in arousal, measured by luminance-independent changes in 16
pupil size, modulate neural activity in similar ways to attention,^{2,13,17} yet these changes occur 17
independent of attention.²² Previous work has not specifically investigated the dimensionality 18
of arousal-dependent signaling and it remains uncertain whether, like other behavioral contexts, 19
it can be explained by a low-dimensional process. 20

Most studies of population coding accuracy rely on relatively small, simple stimulus sets 21
that drive neural activity in stereotyped ways.^{16,23,24} Yet, theoretical work predicts that noise 22
correlations can either enhance or impair coding accuracy, depending on their alignment with 23
the stimulus-evoked activity in the neurons being studied.^{20,25-30} If the effects of arousal 24
are relatively high-dimensional, meaning that they suppress noise along many different di- 25
mensions of neural activity, they should improve coding accuracy of most sensory stimuli 26
equally. Alternatively, if the effects of arousal are confined to a low-dimensional subspace of 27
neural activity, their alignment with sensory-evoked responses should be variable, resulting in 28
stimulus-dependent changes in coding accuracy. 29

In the present study, we investigated the dimensionality of arousal-dependent signaling 30
and its impact on coding accuracy by recording population activity from primary auditory 31
cortex while presenting a large library of natural sounds. We simultaneously monitored arousal 32
level using pupil size.^{2,31} To measure population coding accuracy of natural sound stimuli, we 33
developed a novel dimensionality reduction approach.^{32,33} Overall, arousal improved neural 34
discriminability of natural sounds. However, the degree of improvement varied substantially be- 35
tween stimuli, consistent with the hypothesis that arousal acts on a low-dimensional subspace 36
rather than providing a generalized improvement in coding accuracy. In contrast with attention, 37

modulation of single neuron gain and noise correlations were distinct. These processes operated on different neural populations and timescales. Thus, our results demonstrate that arousal drives robust, selective changes in population coding accuracy across diverse sound stimuli and that these changes act through at least two distinct mechanisms.

Results

We recorded simultaneous single- and multi-unit activity from primary auditory cortex (A1) using single-shank 64-channel, or dual-shank 128-channel linear silicon probes³⁴ ($n = 371$ single-units and $n = 331$ multi-units, Figure 1B). Data were obtained from $n = 25$ recording sites in five awake, head-fixed ferrets. During each recording session, we presented a diverse set of randomly interleaved natural sound excerpts³⁵ (e.g. Figure 2A) in the acoustic field contralateral to the recording hemisphere (Figure 1A). To monitor spontaneous fluctuations in arousal, pupil size was measured continuously during neural recordings using infrared video^{2,31} (Figure 1A, B).

In ferret A1, changes in pupil size are associated with mostly monotonic changes in neural firing rate.² Therefore, to assess overall pupil-related changes in activity, we split the neural data in half based on the median pupil size across each experiment (large pupil/high arousal vs. small pupil/low arousal). Transitions between high and low pupil-indexed arousal were accompanied by changes in the excitability of individual neurons and in the strength of correlations between neurons (Figure 1D). When pupil was large, responses to the same sound were stronger and more reliable than when pupil was small. Across all recorded units, both the baseline firing rate and evoked response gain were positively associated with pupil size (Figure S1). Additionally, during large pupil trials population activity was desynchronized relative to the low arousal state; coordinated, stimulus-independent fluctuations in the population PSTH were primarily observed during small pupil trials only. In support of this observation, we found that pairwise noise correlations were significantly reduced in the high arousal state ($\rho_{small} = 0.047 \pm 0.005$ vs. $\rho_{large} = 0.043 \pm 0.005$, $p = 0.035$, Bootstrap test, $n = 25$ recording sessions, 318 ± 94 unit pairs per session; Figure S1). These results are consistent with previous reports in ferret,² mouse,^{13,17,31} and primate.³⁶

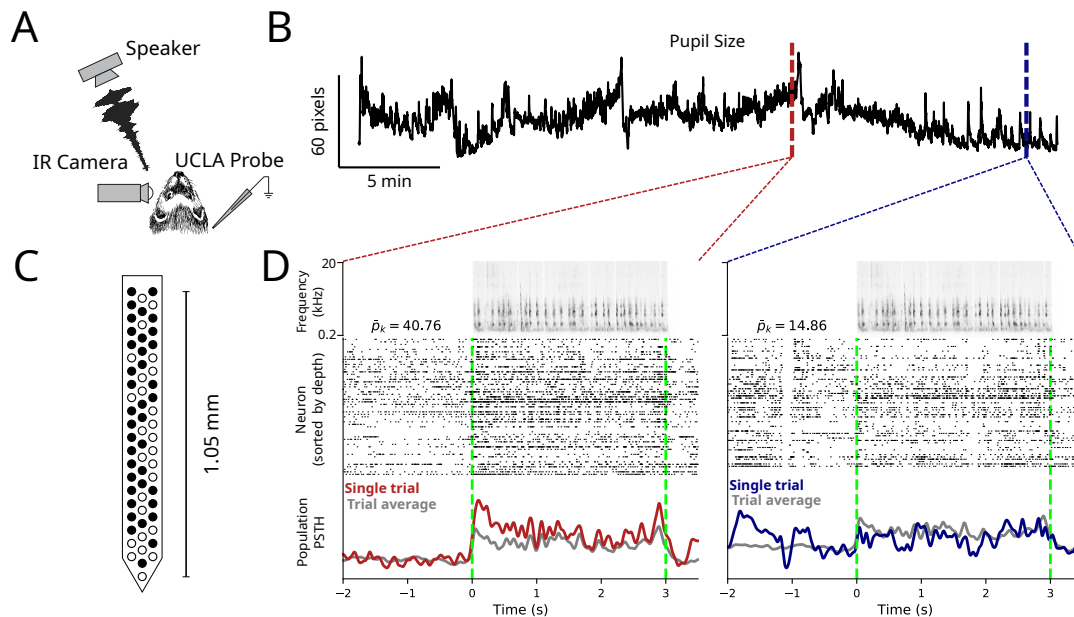


Figure 1. Pupil-indexed arousal modulates neural responses to natural sound stimuli. **A.** Single- and multi-unit activity was recorded from A1 of awake, head-fixed ferrets using laminar electrode arrays during presentation of natural sound stimuli. Pupil size, an index of arousal, was measured simultaneously using infrared video. **B.** Pupil trace from one recording session. Pupil size varied substantially over the course of the recording session, indicating spontaneous transitions between high and low arousal states. **C.** Schematic of 64-channel laminar probe used to record neural activity. Filled circles represent electrode channels on which at least one unit was detected during the same session ($n = 55$ total units). **D.** Top: Spectrogram of one 3 s natural sound excerpt which was presented multiple times during the recording session. Middle: Population raster plot of spiking activity by all simultaneously recorded units during a single stimulus presentation when pupil was large (left, red arrow in B) and when pupil was small (right, blue arrow in B). Bottom: Population peri-stimulus time histogram (PSTH) response, averaged across units during the single trial (red / blue) indicated in B and averaged over all repetitions of this stimulus (gray). \bar{p}_k indicates mean pupil size on each respective trial, k .

The impact of correlated variability on population coding accuracy varies across sensory stimuli

We directly measured the impact of noise correlations on coding of natural sounds using d'^2 , a well-established metric of neural population discriminability.^{26,27,32,37-39} d'^2 describes the ability of population activity to discriminate between two stimuli using an optimal linear decoder. Its value depends both on the mean, stimulus evoked activity and on the variability of responses across trials, or noise correlations, that project along the optimal decoding axis.²⁶⁻²⁸

To prevent overfitting to noise in the high-dimensional population data despite having relatively few repetitions of each stimulus, we performed dimensionality reduction.^{32,33} Before

measuring d'^2 , we projected neural responses onto a two-dimensional plane in neural state space. This projection, which we call decoding-based Dimensionality Reduction (dDR , Figure S3), was defined by two axes: the sensory discrimination axis ($dDR_1 / \Delta\mu$) and the global noise axis (i.e. the first principal component of noise correlations, see Methods). In addition to preventing overfitting to single trial noise in d'^2 estimates (Figure S4), dDR facilitated visualization of the high dimensional neural population data (Figure 2A, C). Increasing the dimensionality of the dDR space beyond two-dimensions by including additional noise components only increased cross-validated estimates of d'^2 marginally and did not impact the effect of arousal on d'^2 , which is the main focus of this work (Figure S6).

Neural discriminability was strongly stimulus-dependent (Figure 2B, C). To determine the source of stimulus dependence, we characterized each stimulus pair by two metrics: discrimination axis magnitude and noise interference (Figure 2B). Discrimination axis magnitude was defined as the vector magnitude of $\Delta\mu$, which described the amount of sensory information contained in the trial-averaged activity. Noise interference was defined as the cosine similarity between $\Delta\mu$ and the correlated variability axis (\mathbf{e}_1 , i.e., the first principal component of stimulus-independent activity in dDR space). Thus, noise interference was a stimulus pair specific metric that described the extent to which single trial variability interfered with the readout of sensory information. These two metrics are used throughout the remainder of this work. To reference their definitions, see Methods: Glossary.

Across the full set of stimulus pairs, we observed substantial variability in both discrimination magnitude and noise interference (Figure S5), which in turn led to changes in discriminability (Figure 2B). For pairs of stimuli with large discrimination magnitude and low noise interference, d'^2 was large (Figure 2C.IV) and for pairs with small discrimination magnitude and high noise interference, d'^2 was small (Fig 2B, C.I). To quantify this, we regressed d'^2 against the per-stimulus noise interference and discrimination axis magnitude (Figure 3E, Eqn. 10). For all sites, noise interference coefficients were negative and discrimination magnitude coefficients were positive ($\beta_{noise} = -0.39 \pm 0.03, p = 0.000071, U = -3.97, n = 11$ recording sessions, Mann-Whitney U Test; $\beta_{discrimination} = 0.80 \pm 0.02, p = 0.000071, U = 3.97, n = 11$ recording sessions, Mann-Whitney U Test), indicating that this pattern was consistent across experiments.

Our results illustrate that baseline neural discriminability of natural sounds varies systematically across the stimulus space. The variation of d'^2 with respect to noise interference emphasizes that correlated variability has very little impact on discrimination for stimulus pairs, but is critically important for others. This dependence of discriminability on noise interference

persisted after controlling for discrimination magnitude, demonstrating that the effect of noise 109
 correlations varies even across pairs of stimuli that differ similarly in their mean response. 110

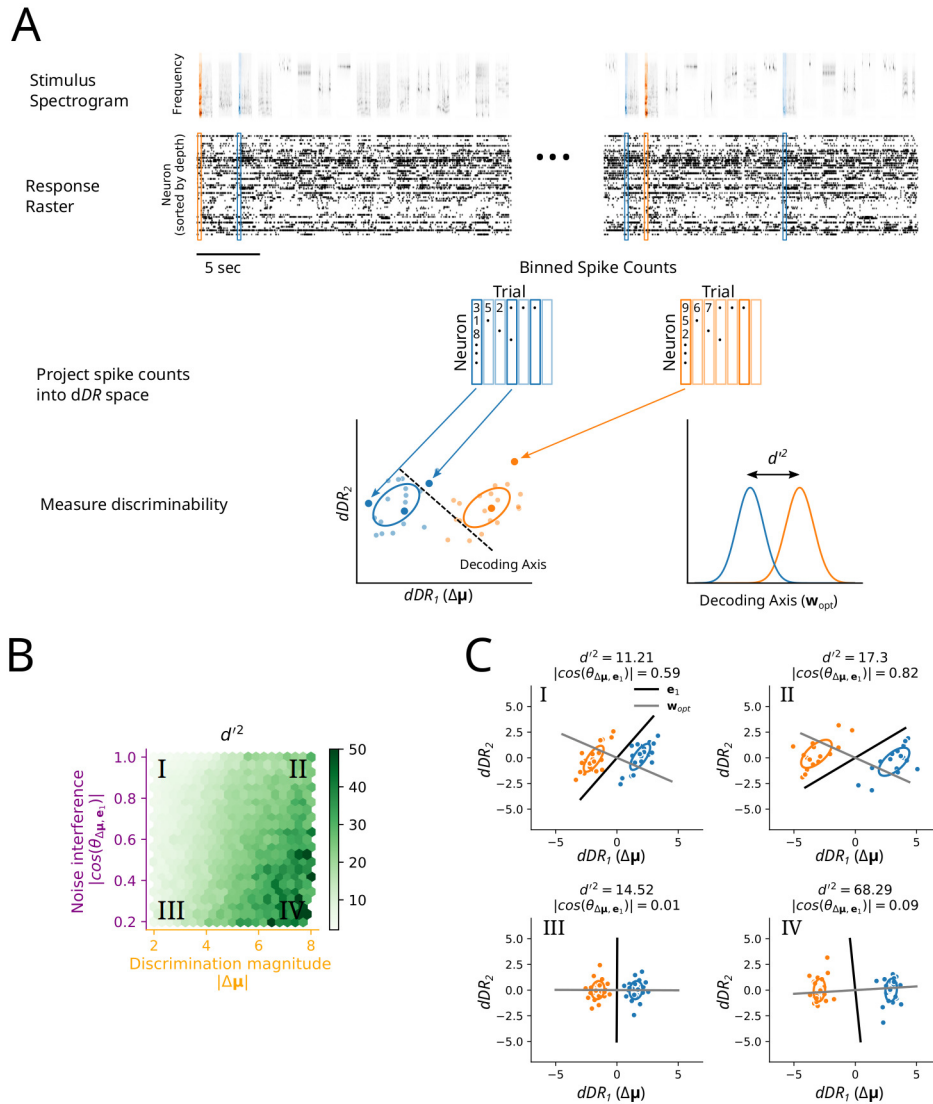


Figure 2. Natural stimulus discriminability in A1 populations varies smoothly across the sensory space. **A.** Procedure for measuring discriminability between natural sound pairs. Spectrogram and response raster from one example recording session are shown for exemplar early (left) and late (right) trials in the experiment. For each 250 ms segment of sound stimuli, spike counts in the corresponding time window were binned and counted as shown. For each pair of unique sound segments (blue vs. orange), the spike counts were then projected into dDR space, where d^2 was measured. **B.** Heat map indicates neural discriminability, d^2 , averaged across stimulus pairs and recording sites. Values were binned according to discrimination magnitude (x-axis) and noise interference (y-axis). **CI-IV.** Single-trial responses to four example sound pairs. Each dot represents the weighted sum of neural population activity on a single trial. Ellipses indicate standard deviation across trials. Example stimulus pairs shown were all collected from the same neural population during a single recording session. Each panel shows a pair of stimuli from a different quadrant (I - IV) in B. e_1 (black line) represents the first eigenvector of the noise in the dDR space and w_{opt} (grey line) represents the optimal decoding axis for separating the two stimulus classes (blue vs. orange).

Arousal selectively improves neural discrimination of natural sound stimuli

We used the *dDR*-based decoder to determine how pupil-indexed arousal modulates population coding accuracy in A1. For each pair of stimuli, we computed a single decoding axis across all pupil states, as above (\mathbf{w}_{opt} , Eqn. 9). We then split the data from each stimulus pair in half based on the median value of pupil across sound repetitions and computed d'^2 separately for small and large pupil conditions. The use of a single decoding axis, combined with the median split of data based on pupil size, prevented biasing d'^2 for one or the other pupil condition. Across stimulus pairs, we found that discriminability of natural sounds was improved on average in large pupil, high arousal states (Figure 3A). This change was significant across the population ($p = 0.02$, $W = 7$, $n = 11$ recording sites, Wilcoxon signed-rank test). Moreover, the magnitude of the effect was correlated with the amount by which pupil varied within an experiment, consistent with it being driven by pupil-indexed fluctuations in arousal (Figure S7, $r = 0.63$, $p = 0.029$, $n = 11$ recording sessions, permutation test).

Our analysis of overall discriminability revealed substantial variation in the impact of noise correlations on d'^2 across stimulus pairs (Figure 2B). Therefore, we next considered the possibility that pupil-dependent changes in coding accuracy also depend on these same features. To test this, we computed the fraction change in discriminability between large and small pupil conditions, $\Delta d'^2$ (Eqn. 12), which ranged from -1 to 1 and could be directly compared across stimulus pairs and recording sites. We binned stimulus pairs according to the discrimination axis magnitude and noise interference (as in Figure 2B) and averaged $\Delta d'^2$ within bins and across recording sites (Figure 3B). Unlike overall discriminability, improvements did not depend on discrimination magnitude (Figure 3D, right). However, $\Delta d'^2$ was positively correlated with noise interference; for stimulus pairs with high noise interference, increased arousal led to a greater relative improvement in discriminability (Figure 3D, left). We quantified these axis-dependent changes in $\Delta d'^2$ using linear regression. $\Delta d'^2$ and noise interference were positively correlated in 8/11 recording sites ($p < 0.05$, t-test), and this correlation was significant across experiments ($\beta_{noise} = 0.102 \pm 0.028$, $p = 0.001$, $U = 3.25$, $n = 11$ recording sessions Mann-Whitney U Test; Figure 3F). No consistent relationship was observed for discrimination magnitude ($\beta_{discrimination} = 0.004 \pm 0.030$, $p = 0.718$, $U = 0.36$, $n = 11$ recording sessions, Mann-Whitney U Test; Figure 3F). These results demonstrate that improvements in stimulus discriminability associated with increased arousal are largest for stimulus pairs where correlated variability

aligns with the discrimination axis. The selectivity of this effect suggests that state-dependent changes are low-dimensional and that they align with the most prominent axis of correlated variability.

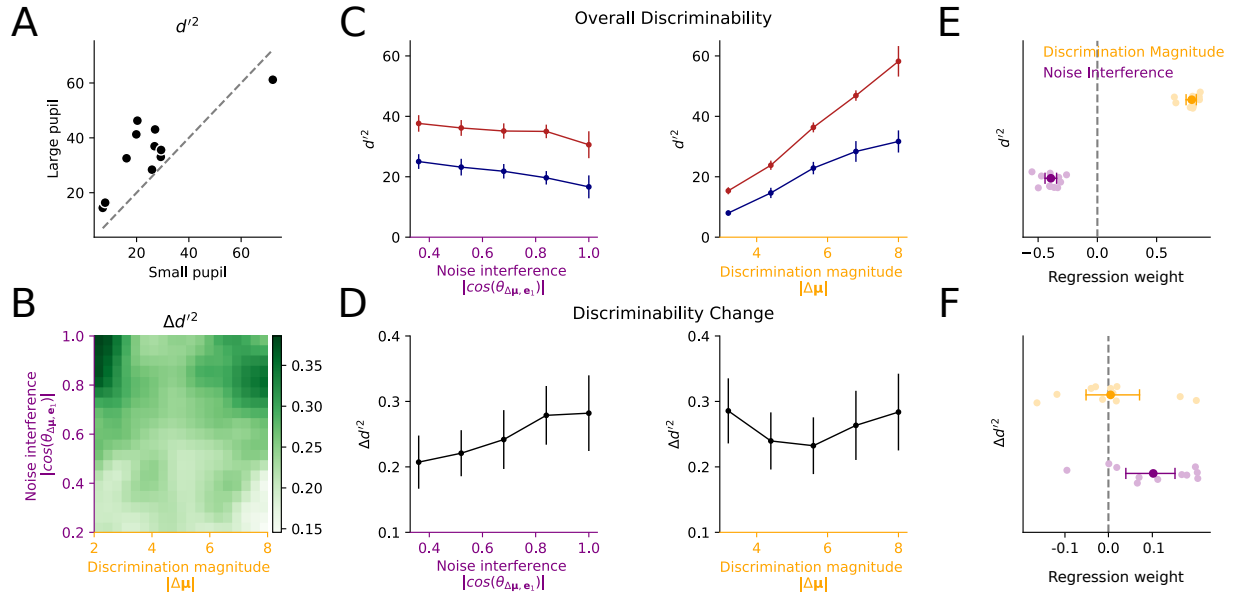


Figure 3. Arousal improves discrimination of natural sounds. **A.** Scatter plot compares mean discriminability (d'^2) across natural sound pairs for small (low arousal) versus large (high arousal) pupil trials. Each point indicates the mean across all stimulus pairs presented during a single recording session. Mean discriminability was nearly always greater during large pupil trials ($p = 0.02$, $W = 7$, $n = 11$ sessions, Wilcoxon signed-rank test). **B.** Heatmap shows the relative change in discriminability for large vs. small pupil ($\Delta d'^2$) as a function of discrimination axis magnitude ($|\Delta\mu|$) and noise interference ($|\cos(\theta_{\Delta\mu, e_1})|$). Results are smoothed by a Gaussian filter of width two bins. **C.** Average d'^2 for large and small pupil trials grouped into five evenly spaced bins by noise interference (left) or discrimination magnitude (right). Points / error bars indicate the mean / standard error across recording sessions. **D.** Arousal-related fraction improvement in discriminability ($\Delta d'^2$) for each session, plotted as in C. **E.** Regression coefficients computed per recording session for dependence of d'^2 on discrimination magnitude and noise interference. Error bars show bootstrapped 95% confidence interval across sessions. $\beta_{Noise} = -0.39 \pm 0.03$, $p = 0.000071$, $U = -3.97$, $\beta_{Discrimination} = 0.80 \pm 0.02$, $p = 0.000071$, $U = 3.97$, $n = 11$ recording sessions, Mann-Whitney U test. **F.** Coefficients for pupil-dependent changes in discriminability, $\Delta d'^2$, plotted as in (E). $\beta_{Noise} = 0.102 \pm 0.029$, $p = 0.001$, $U = 3.25$, $\beta_{Discrimination} = 0.004 \pm 0.030$, $p = 0.718$, $U = 0.36$, $n = 11$ sessions, Mann-Whitney U test.

Arousal improves discriminability through a combination of increased response gain and reduced correlated variability

Changes in either single neuron response magnitude or noise correlations (Figure 1-S1) could lead to enhanced discriminability in the large pupil condition. We reasoned that three

population level factors could capture these effects: discrimination axis (signal) magnitude, 151
shared noise variance, and noise interference (Figure 4A). Discrimination axis magnitude is 152
directly related to d'^2 (Figures 2B and 3C) and is measured using trial-averaged activity. Thus, 153
pupil-dependent changes in this value only reflect changes in the evoked response magnitude 154
of single neurons. Shared noise variance, on the other hand, is independent of trial-averaged 155
responses and reflects the strength of trial to trial variability. Noise interference, as we discuss 156
above, reflects the interaction between these two quantities. 157

Inspection of changes in these factors for individual stimulus pairs showed that discrimi- 158
nation magnitude increased in the large pupil condition and shared noise variance decreased 159
(Figure 4B, left). We used a regression model to determine the extent to which each could 160
explain changes in d'^2 (Figure 4C, D). Consistent with a contribution to the change in discrim- 161
inability, a positive change in discrimination magnitude and a decrease in noise variance both 162
predicted an increase in d'^2 (Figure 4C). Each of these factors contributed significantly to the 163
improved discriminability during high arousal conditions (Figure 4D, discrimination (signal) 164
magnitude: mean cross-validated $R^2 = 0.198$, $p = 0.000071$, $U = 3.97$, shared noise variance: 165
mean $R^2 = 0.177$, $p = 0.000071$, $U = 3.97$, $n = 11$ recording sessions, Mann-Whitney U test). 166
Noise interference, however, changed only marginally between arousal states and did not ex- 167
plain substantial variability in d'^2 (mean cross-validated $R^2 = 0.007$, $p = 0.28$, $U = 1.08$, $n = 11$ 168
recording sessions, Mann-Whitney U test). 169

Both independent, single neuron variance, as well as noise correlations, could impact 170
discriminability. As validation that the our measurement of shared noise variance corresponded 171
primarily to modulation of noise correlations, and not changes in single neuron variance, we 172
simulated population activity. In the simulation, the mean and variance of single neurons was 173
matched to the values measured for the small and large pupil states, but noise correlations were 174
held fixed (Figure S9, Methods). These simulated data were unable to account for the actual 175
observed changes in d'^2 , while a simulation that also incorporated changes in noise correlations 176
accurately reproduced the raw data. Thus, a combination of changes in the trial-averaged, 177
evoked activity of single neurons and in correlated variability between neurons is needed to 178
explain decoding improvements. 179

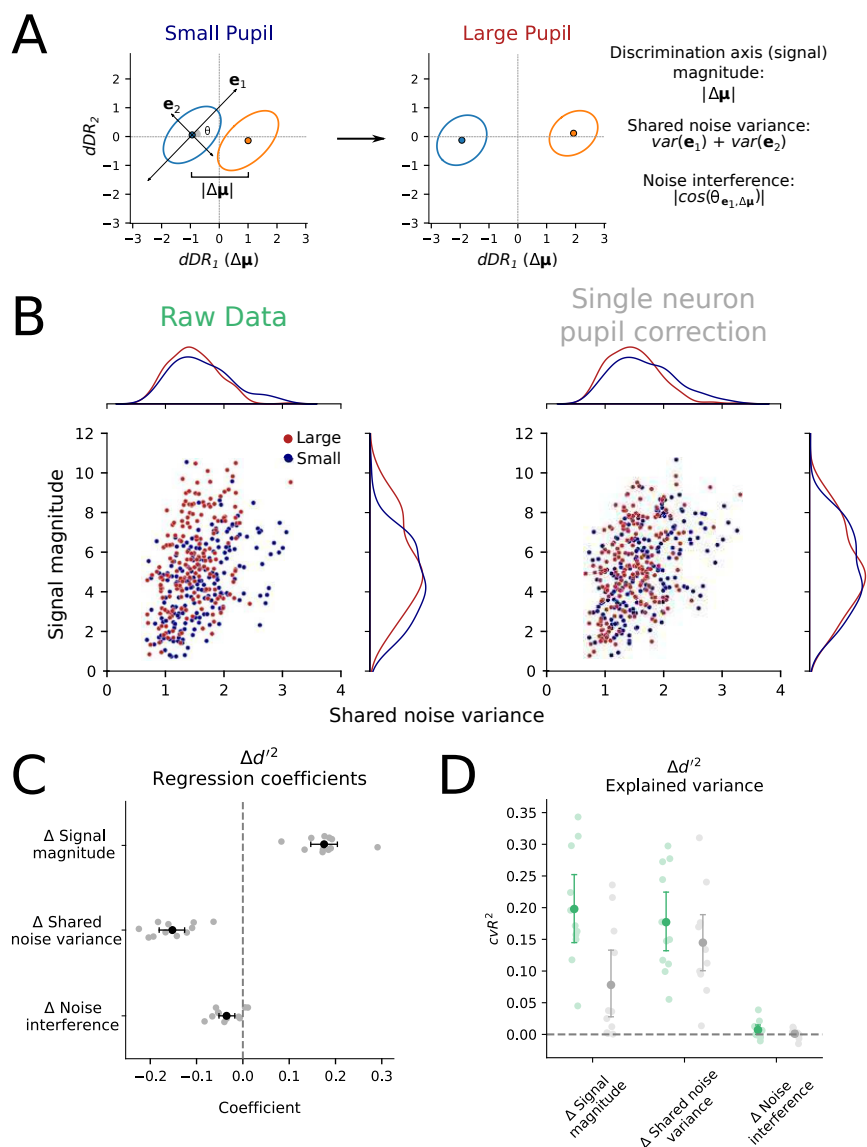


Figure 4. Arousal-dependent changes in both mean activity and trial to trial variability impact stimulus discriminability. **A.** Cartoon schematic of typical arousal-dependent modulation of population responses to two stimuli (blue and orange) in dDR space. A change in mean response affects the distance between ellipsis centers, and a change in correlated variability affects their size. **B.** Scatter plot of signal magnitude (mean difference along discrimination axis) versus shared noise variance for each stimulus pair. Color indicates small versus large pupil conditions, randomly sub-sampled to facilitate visualization ($n = 200/17,778$ stimulus pairs). **C.** Linear regression was used to model changes in discriminability (d'^2) as a function of changes in signal magnitude, shared noise variance, and noise interference (Eqn. 13). Each gray point represents the regression coefficient for data from one experiment ($n = 11$). Error bars indicate standard error of the mean coefficient across experiments. **D.** Cross-validated change in d'^2 explained by each regressor ($\text{cv}R^2$, Methods) before (green / raw) and after (grey / pupil correction) correcting for pupil-dependent changes in single neuron evoked responses. Error bars represent standard error of the mean $\text{cv}R^2$ across experiments.

Arousal-dependent changes in gain and correlated variability operate on distinct timescales and neural populations

Studies of selective attention have shown that attention-related changes in response gain of single neurons and the strength of correlated variability are related. In this framework, changes in correlated variability reflect changes in the variance of coherent modulation of the sensory gain of single neurons across the population.^{18,19} Thus, a single, shared mechanism appears to mediate changes in both gain and correlation. We wondered if a single mechanism also produced effects of arousal or if changes in gain and correlated variability were mediated by separate mechanisms.

If arousal impacts single neuron activity and correlated variability via a single mechanism, we reasoned that removing all pupil-dependent changes in single neuron spiking activity should also remove population-level changes in both discrimination axis magnitude and shared noise variance (Figure 4B). To test this hypothesis, we generated a pupil-corrected data set in which all pupil-explainable response variability for each unit was subtracted from its raw response (Eqn. 8). We then measured the pupil-dependent discrimination axis magnitude, shared noise variance, and stimulus discriminability for the pupil-corrected data. The single neuron correction abolished changes in discrimination axis magnitude, but changes in shared noise variance were largely unaffected (Figure 4B). These results argue that arousal does not modulate correlated variability by reducing shared fluctuations in arousal-dependent gain.^{18,19} Instead, increased arousal suppresses an independent source of correlated variability. These separable mechanisms are distinct from reports for selective attention, where changes in response gain and correlated variability are thought to be produced by a single shared modulator.

To further distinguish mechanisms that impact discrimination magnitude and correlated variability, we used a bandpass filter to partition the data into distinct temporal frequency bands and measured noise correlations separately in each band (Figure 5A, B). Breaking the data into frequency bands allowed us to identify the timescale over which the correlation effects operate. Noise correlations were largest overall in low-frequency bands (< 0.5 Hz), reflecting slow, coordinated fluctuations in the population activity.⁴⁰ Pupil-indexed fluctuations in arousal occur at this slow timescale. Thus, we expected the effects of coherent single neuron modulation to be largest in this band. Indeed, the pupil correction reduced these slow noise correlations significantly ($p = 0.00003$, $W = 6$, $n = 25$ recording sessions, Wilcoxon signed-rank test, Figure 5A) but not in the higher frequency bands, > 2 Hz. In contrast, arousal-dependent changes

in correlation magnitude were restricted to the higher frequency bands, ranging between 0.5 212
and 25 Hz (Figure 5B, $p = 0.019, 0.002, 0.007, 0.00005, W = 76, 51, 63, 11, n = 25$ recording 213
sessions, Wilcoxon signed-rank test). Therefore, we conclude that the arousal-dependent 214
changes in correlated variability are due to modulation of a process that operates on a faster 215
timescale, distinct from the change in arousal itself. 216

If changes in correlated variability are in fact distinct from single neuron effects, the two 217
processes may also operate on different subpopulations of neurons. To determine if this was the 218
case, we compared the magnitude of pupil-dependent modulation observed in each cell with 219
the degree to which arousal state impacted noise correlations for that cell. Single cell effects 220
were quantified by the pupil modulation index (MI) (Eqn. 4). A value of MI= 1 indicated a unit 221
that only responded when pupil was large and MI= 0 indicated no change in responsiveness 222
between large and small pupil. While MI was able to predict noise correlations between two 223
units, it contained no information about the arousal-dependent change in noise correlations 224
(Figure 5C, D). Thus, in addition to acting on distinct timescales, each process also operates on 225
unique populations of neurons in A1, confirming that single neuron changes are distinct from 226
changes in correlated variability. 227

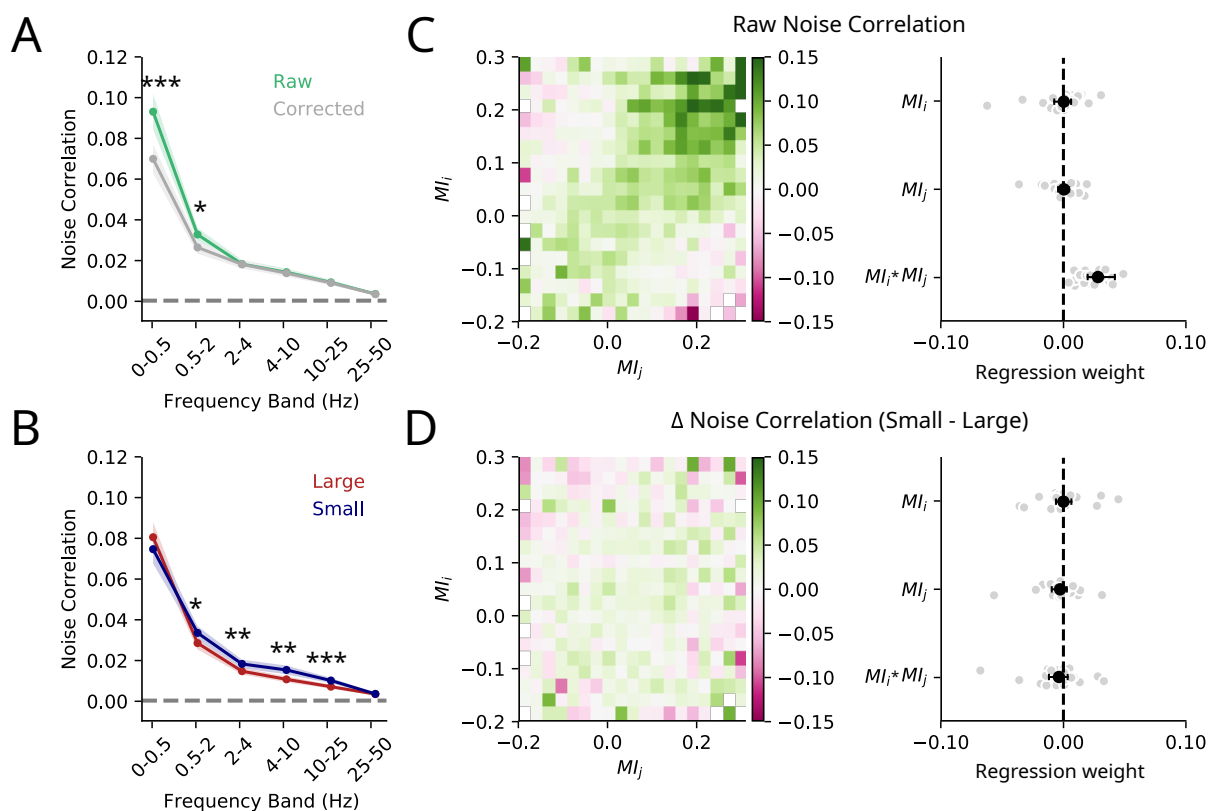


Figure 5. Arousal-dependent changes in evoked response magnitude and correlated variability operate on distinct timescales and neural populations. **A.** Noise correlations, computed with and without pupil correction, split into six non-overlapping temporal frequency bands. To minimize bias from one recording session, noise correlations were averaged across neuron pairs within session before measuring significance of the difference between conditions ($***p < 0.001$, $**p < 0.01$, $*p < 0.05$, Wilcoxon signed-rank test, $n = 25$ recording sessions). Shading indicates standard error across recording sessions. **B.** Pupil-corrected noise correlations measured after splitting data by median pupil size, plotted as in **A**. **C.** Left: Heatmap of mean noise correlations (using 4Hz temporal bins, as in Figure S1), grouped into 20 evenly spaced MI bins for each pair of units (i, j). Right: Regression coefficients, predicting overall noise correlations from MI of individual neurons (MI_i and MI_j) and their interaction ($MI_i * MI_j$). Gray dots represent single recording sessions; error bars represent 95% confidence intervals across sessions. **D.** Same as **C** but for change in noise correlations (Δ noise correlations) between small and large pupil.

Arousal-dependent reduction of correlated variability is low dimensional

Having determined that arousal-dependent changes in correlated variability operate through a distinct mechanisms from changes in individual neuron excitability, we finally sought to determine the dimensionality of the mechanism producing the correlations. Specifically, we determined if changes in noise correlations could be described by a low dimensional latent process, or if they reflected more complicated modulation of high-dimensional neural coupling

across the population. For each session, we computed the axis of maximal change in noise correlations between large and small pupil (Figure 6A-D). In most cases, this pupil modulation space was low-dimensional; in 9/11 recording sessions, we found only one significant dimension along which noise correlations changed between large and small pupil ($p < 0.05$, permutation test, Methods). This dimension was consistently aligned with the first principal component of the pooled noise data (6E, F), suggesting that arousal modulates correlated variability along, low-dimensional, high variance dimensions in the data.

Given that arousal-dependent changes are usually restricted to a single dimension, we hypothesized that their impact on discriminability will depend on how well that axis aligns with the sensory discrimination axis. To test this, we compared the pupil modulation axis to the sensory discrimination axis ($\Delta\mu$) for each stimulus pair. Indeed, this was the case. For pairs of stimuli where the cosine similarity between $\Delta\mu$ and pupil modulation axis was high, we observed that the effect of correlated variability changes on d'^2 were larger (Figure 6G, $p = 0.028$, $W = 4$, $n = 9$, Wilcoxon signed-rank test).

Based on these results, we conclude that arousal-dependent changes in correlated variability are restricted to a relatively low-dimensional subspace. These changes increase sound discriminability when it is aligned with the discrimination axis, a scenario that only occurs for a subset of sensory discriminations.

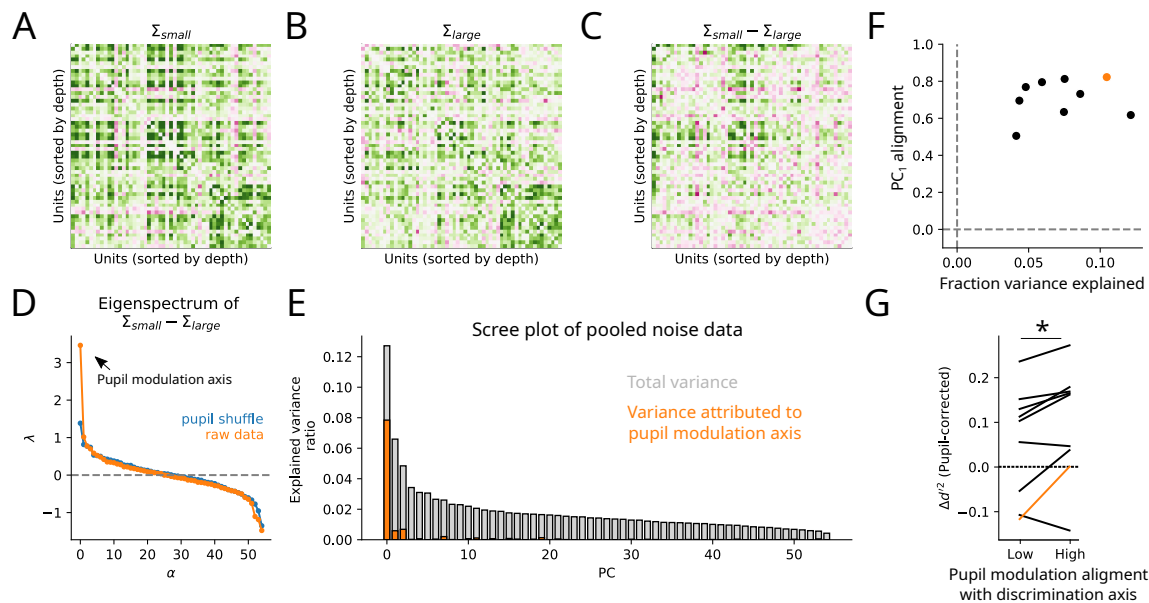


Figure 6. Pupil-related changes in noise correlations are low-dimensional and predict changes in stimulus discriminability. **A.** Matrix of pairwise noise correlations between all units during the small pupil condition in one recording session. **B.** Same as in **A** for large pupil condition. **C.** Change in noise correlations, computed as the difference between matrices in **A** and **B**. Color scale shared between panels **A-C**. **D.** Eigenvalues of the difference matrix in **C** (orange). The first eigenvalue corresponds to the pupil modulation axis. Blue line shows the noise floor, computed by shuffling pupil size before classifying large vs. small pupil conditions. Only sites with at least one significant eigenvalue were included subsequently in panels **F/G** (9/11 sites, $p < 0.05$, permutation test). **E.** Scree plot shows variance explained ratio for each principal component. Orange shading indicates the portion of variance along each PC that can be attributed to the pupil modulation axis. The large orange component in the first position indicates that the pupil modulation axis largely aligns with the first principal component. **F.** Scatter plot compares fraction of total variance in the population activity explained by the modulation axis (x -axis) against the ratio of modulation axis variance to PC_1 variance (y -axis), *i.e.*, the ratio between heights of the orange and grey bars for PC_1 in panel **E**. Example site from **A-E** is shown in orange. **G.** Change in d'^2 between pupil conditions after performing single neuron pupil correction for each site. Mean change is computed across stimulus pairs after first grouping each pair by the alignment of the pupil modulation axis with its discrimination axis (median split between low and high alignment). Stimulus pairs with high alignment (measured with cosine similarity) showed greater discrimination improvements ($\mu_{high} = 0.099 \pm 0.042$, $\mu_{low} = 0.057 \pm 0.041$, $p = 0.028$, $n = 9$, Wilcoxon signed-rank test).

Discussion

Previous studies have suggested that behavior-dependent modulation of neural population coding operates in a low-dimensional space.^{18,19,21,22,40,41} That is, signals reflecting behavioral state are well-described by processes that modulate the activity of many neurons coherently and thus produce correlated variability in sensory responses. However, most previous work has utilized relatively small, focused stimulus sets. This raises questions about

whether the observed low-dimensional processes are a consequence of the stimuli tested, or 259
if they are a general feature of state-dependent modulation. These questions are critical for 260
understanding population coding of sensory stimuli. Theoretical studies have long shown that 261
correlated variability can impact coding accuracy, but only if it aligns with the sensory tuning of 262
neurons in the population.²⁷ Thus, the dimensionality of the mechanisms driving correlated 263
variability and how they interact with sensory selectivity is critical for understanding their 264
impact on sensory processing. 265

In the case of pupil-indexed arousal, we found that correlated activity is modulated in a 266
low-dimensional subspace of primary auditory cortex (A1) which we found to be distinct from 267
the arousal-dependent changes in single neuron responses. These results were consistent 268
across a diverse set of natural sound stimuli. The effect of arousal on neural discrimination of 269
sounds varied substantially with the sound stimulus, as predicted for a low-dimensional signal 270
interacting with high-dimensional stimulus-evoked activity. 271

Effects of shared intrinsic variability on discriminability are stimulus- 272 **dependent** 273

Correlated, intrinsic variability within neural populations is ubiquitous in cortex. Even 274
before this phenomenon was observed experimentally, substantial efforts were made to 275
develop a theoretical understanding of how correlated activity might affect coding by neural 276
populations.^{26,37,42-47} This early work established that correlated variability can interfere 277
with the brain's ability to accurately discriminate sensory stimuli. Therefore, experimental 278
characterization of this phenomenon is critical to fully understand neural population codes. 279

Although evidence for intrinsic correlation is widespread, experimental studies have 280
provided conflicting evidence as to whether or not it does in fact interfere with population 281
coding.^{20,25,28-30,32,38,39} There are at least two reasons why the reported effects of correlated 282
variability might vary across studies. First, in some cases activity along dimensions containing 283
interfering noise could have very low variance. In this case, measuring the noise reliably would 284
require recording large amounts of data, both over many neurons and over many trials, a 285
methodology that has only recently become feasible.^{32,38,39} A second possibility is related to 286
the fact that the impact of correlated noise depends on the tuning of neurons in the population 287
being read out and its relationship with the noise space.^{26,28,30} In this case, discrepancies in 288
previous work might be explained by differences in the neural populations that were sampled 289

or in the stimuli that were tested. Because the effects of intrinsic noise may depend on the stimuli that are presented, it is important to characterize coding accuracy across diverse sets of stimuli. Indeed, our results showed that the effects of correlated variability on coding are highly dependent on the stimuli tested.

Because there is a trade-off between the number of stimuli that can be presented and the number of times that each can be repeated during a single recording session, questions about stimulus-dependent changes in population coding are difficult to completely address in a single study. Unlike recent work,^{32,38,39} we measured neural responses to a large set of stimuli over a relatively small number of repeats and neurons. Thus, we could not measure low-variance dimensions and draw strict conclusions about the presence (or absence) of information-limiting noise.²⁶ Instead, by developing a novel dimensionality reduction approach (*dDR*), we were able to reliably estimate the interaction between the dominant, high-variance noise dimension and sensory discrimination across a large acoustic stimulus space. This approach revealed substantial variability of arousal dependent changes in coding within each recording site. This highlights the practical benefit of dimensionality reduction techniques for studying neural population dynamics across a diversity of stimulus and behavioral contexts.⁴⁸

State-dependent coding in auditory cortex

It is increasingly clear that neural activity in primary sensory regions of the brain is significantly modulated by non-sensory variables, including arousal.^{2,13,17,40,49} Arousal here refers to spontaneous changes in alertness as measured by pupil diameter, even in the absence of a behavioral task.⁴⁹ Similar to previous work, we find that increased arousal is associated with enhanced excitability and reduced noise correlations in A1.^{2,13,31} These effects boost the neural signal to noise ratio in V1¹⁷ and improve population coding accuracy of tonal stimuli in A1.¹³

Building on this previous work, we explored the effects of arousal on population coding accuracy across a large space of natural sounds. Consistent with prior results, increases in pupil-indexed arousal led to improved discriminability between sounds on average. However, the relative magnitude of this improvement varied substantially across stimuli. Improvements were largest when correlated variability interfered with sensory discrimination. Thus, a suppression of correlated variability and increase in evoked response rates during high arousal states may gate behavioral states that require improved perceptual discrimination, like selective attention.

This strong stimulus dependence highlights the importance of a systematic exploration of state-dependent changes in neural coding across the sensory response space. Parametric stimuli might be used to more systematically probe sound feature representations across a range of behaviorally relevant stimuli. For example, one study of auditory processing began to address this question in anesthetized animals.²³ In this work, Kobak et al. measured population coding in A1 of sounds that varied along two dimensions: Inter-aural level (ILD) and absolute binaural level (ABL). By inducing different states of cortical activation with urethane anaesthesia, the authors demonstrated that in the awake (desynchronized) state, noise and signal subspaces shift to become orthogonal, thereby facilitating accurate encoding across both ILD and ABL. Extending this approach to spectro-temporally varying and behaviorally relevant naturalistic stimuli⁵⁰ will be critical for a complete understanding of state-dependent population coding.

Separate mechanisms drive arousal-dependent changes in single neurons and correlated neural variability

Recent studies of selective attention have suggested that correlated variability results from the coherent modulation of many neurons by an intrinsic behavioral state variable.¹⁸⁻²⁰ In a recent study by Denfield et al., macaques were trained on a visual change-detection task in which the stability of spatial selective attention was manipulated between behavioral blocks.¹⁹ Because the gain of evoked responses in visual cortex is known to be modulated by attention,⁵¹ the authors proposed that the magnitude of correlated variability should be highest when attention itself was most variable, as changes in gain are shared across neurons within the receptive field. Indeed, when animals were required to switch attention between multiple locations within a behavioral block, noise correlations were strongest. This idea of correlations produced by a shifting spotlight of attention is consistent with previous characterizations of neural population activity and attention¹⁸ and agrees well with theoretical work.^{20,21} These findings offer a parsimonious explanation for why gain changes are accompanied by a reduction in noise correlations during traditional cued change-detection tasks, where attention is focused stably on a single spatial location.¹⁶

Unlike the case of visual selective attention, we found that arousal-dependent modulation of evoked rates and noise correlations in auditory cortex could be fully dissociated, and thus did not arise from a common signal. Modulation of single neuron activity was slow, on the order of many seconds. These slow changes may be analogous to previously described drift signals

in cortex.^{18,22} Correlated variability between neurons, however, also operated on a faster timescale, from tens to hundreds of milliseconds. The magnitude of these faster correlations was modulated by the slow arousal signal, but we could not predict changes in correlated variability from modulation of single neurons. Instead, neurons undergoing slow changes in excitability and those undergoing modulation of faster noise correlations comprised two independent subpopulations.

Although it was not feasible to directly isolate the circuitry underlying these distinct effects in the current study, we propose that they may arise through a combination of neuromodulation and intracortical feedback. Several studies have shown a strong correlation between slow fluctuations in pupil diameter and brain-wide release of norepinephrine and acetylcholine,⁵² making them good candidates for mediating the slow changes in response baseline and gain across individual neurons. The decrease in correlated activity, on the other hand, may arise due to modulation of feedback from other cortical areas that are themselves targeted by the same neuromodulatory signals.

Intracortical pathways to auditory cortex have been identified from multiple areas, including visual,⁵³ motor⁵⁴ and prefrontal cortex.⁵⁵ These inputs can activate inhibitory networks that desynchronize local network activity, and modulating their strength could produce the correlation effects observed in the current study. Given the diversity of these intracortical signals, it might seem surprising that the arousal-related changes reported here should occur in such a low-dimensional space. Further investigation with selective control of feedback from different cortical areas will determine if, in fact, the impact of signals from these different cortical areas can be dissociated in A1.

Materials and Methods

Surgical procedure:

All procedures were performed in accordance with the Oregon Health and Science University Institutional Animal Care and Use Committee (IACUC) and conform to standards of the Association for Assessment and Accreditation of Laboratory Animal Care (AAALAC). The surgical approach was similar to that described previously.^{2,9,56,57} Five young adult male ferrets were acquired from an animal supplier (Marshall Farms). Head-post implantation surgeries were then performed in order to permit head-fixation during neurophysiology recordings. Two

stainless steel head-posts were fixed to the animal along the midline using UV-cured dental composite (Charisma) or bone cement (Palacos), which bonded to the skull and to stainless steel screws that were inserted into the skull. After a two-week recovery period, animals were habituated to a head-fixed posture and auditory stimulation. At this point, a small (0.5 - 1 mm) craniotomy was opened above primary auditory cortex (A1) for neurophysiological recordings.

Acoustic stimuli:

Digital acoustic signals were transformed to analog (National Instruments), amplified (Crown), and delivered through a free-field speaker (Manger) placed 80 cm from the animal's head and 30° contralateral to the hemisphere in which neural activity was recorded (Figure 1). Stimulation was controlled using custom MATLAB software (<https://bitbucket.org/lbhb/baphy>), and all experiments took place inside a custom double-walled sound-isolating chamber (Professional Model, Gretch-Ken).

Natural sounds stimuli were presented in four different configurations. Set 1 consisted of 93, 3-sec samples (2.5 sec ISI, $n = 3$ sites), set 2 consisted of 306, 4-sec samples (1 sec ISI, $n = 14$ sites), set 3 consisted of 306, 1-sec samples (0.5 sec ISI, $n = 6$ sites), and set 4 consisted of 2, 3-sec samples of ferret vocalizations (2.5 sec ISI, $n = 2$ sites). In sets 1-3, the stimulus sets contained species conspecific and heterospecific vocalizations, speech, music, and environmental sounds chosen to sample diverse spectro-temporal statistics.³⁵ All stimuli were presented at 65 dB SPL. During every experimental session, a subset of samples were repeated at least ten times (set 1: 3 samples, set 2: 18 samples, set 3: 18 samples, set 4: all samples), while the remainder were played only once. In order to study the trial-to-trial variability in neural responses, only the high-repeat sounds were included in this study. The order in which stimuli were presented was generated pseudo-randomly. Stimuli were played continuously until all sound samples in the library had been presented. In the case of set 1, the entire stimulus set was repeated 2-3 times. This meant that experiments lasted approximately 40 minutes. The full sound library can be accessed at <https://bitbucket.org/lbhb/baphy>. Some of data used in this study has been published previously.^{2, 58}

Neurophysiology:

Upon opening a craniotomy, 1 - 4 tungsten micro-electrodes (FHC, 1-5 M Ω) were inserted to characterize the tuning and response latency of the region of cortex. Sites were

identified as A1 by characteristic short latency responses, frequency selectivity, and tonotopic gradients across multiple penetrations.^{59,60} Subsequent penetrations were made with a single (64-channel) or dual shank (128-channel) silicon electrode array.³⁴ Electrode contacts were spaced 20 μm horizontally and 25 μm vertically, collectively spanning 1.05 mm of cortex. On each consecutive recording day, we changed the location of the electrode penetration to access fresh cortical tissue, expanding the craniotomy as necessary. Data were amplified (RHD 128-channel headstage, Intan Technologies), digitized at 30 KHz (Open Ephys⁶¹) and saved to disk for further analysis.

Spikes were sorted offline using Kilosort⁶² or Kilosort2 (<https://github.com/MouseLand/Kilosort2>). Spike sorting results were manually curated in phy (<https://github.com/cortex-lab/phy>). For all sorted and curated spike clusters, a contamination percentage was computed by measuring the cluster isolation in feature space. All sorted units with contamination percentage less than or equal to 5 percent were classified as single-unit activity. All other stable units that did not meet this isolation criterion were labeled as multi-unit activity.

Pupillometry:

During neurophysiological recordings, video of the ipsilateral pupil (relative to the recording hemisphere) was collected using an open source camera (Adafruit TTL Serial Camera) fitted with a lens (M12 Lenses PT-2514BMP 25.0 mm) whose focal length allowed placement of camera 10 cm from the eye. Contrast was increased using infrared illumination. Ambient light levels were fixed for each experiment at roughly 1500 lux to provide maximum dynamic range of pupil size.² Pupil size was measured offline by fitting an ellipse to each video frame using a custom machine learning algorithm (Python and Tensorflow). The minor axis of the fit ellipse was extracted and saved for analysis with neurophysiological data. Blinks were detected and excluded as in² and pupil data was shifted by 750 ms relative to spike times in order to account for the lagged relationship between changes in pupil size and neural activity in auditory cortex and to allow for comparison with previous research.³¹

The pupil tracking algorithm itself utilized a deep learning approach. Our model architecture was based on DenseNet201,⁶³ which is available through Keras (<https://keras.io/>). In order to transform the output of the model to pupil ellipse predictions, we added a single global pooling layer and a final prediction layer in which five pupil ellipse parameters (x-position, y-position, minor axis, major axis, and rotation) were fit to each video frame. In order to initial-

ize model weights, the model was pre-trained on ImageNet,⁶⁴ then fine-tuned using roughly 443
500 previously analyzed, nonconsecutive frames from video of the pupil of multiple different 444
ferrets (data from²). Qualitatively, after this first round of training the model performed well, 445
even on novel video frames of pupil from new animals. However, in cases where the pupil 446
video quality was poor, or differed substantially from the video frames in our training data 447
set, we noticed failures in the model predictions. To further improve the model, we employed 448
an active learning procedure. For each new analyzed video, pupil ellipse fits were analyzed 449
qualitatively by experimenters. If the fit quality was deemed poor, predictions for these frames 450
were manually corrected and added to the training data set. The model was then retrained and 451
the analysis rerun. The network became robust to varying levels of video quality and performed 452
consistently without the need for user intervention. The code for this analysis is available at 453
https://github.com/LBHB/nems_db. 454

Because arousal was not explicitly controlled, identical arousal states were not sampled 455
from day to day, between different animals, or between different stimulus pairs within the 456
same experimental session. To control for this variability, we created a normalized metric of 457
pupil variance for each stimulus, and then computed the mean of this variance metric for each 458
recording session. For each stimulus pair, trials were split in half based on median pupil size 459
across all trials. Next, the difference of mean large pupil size and mean small pupil size was 460
normalized by the standard deviation of pupil across the entire experiment. Therefore, for a 461
stimulus pair that sampled a large range of pupil states (relative to pupil during that particular 462
recording session) this metric was high. Finally, we averaged across all stimulus pairs at a 463
recording site and compared this mean across recording sessions (Figure S7). 464

Pupil-dependent GLM:

 465

In order to characterize the dependence of first-order response statistics on pupil-indexed 466
arousal, we built a state-dependent generalized linear model of sound-evoked activity. For 467
each recorded unit, i , the input to this model was defined as the peri-stimulus time histogram 468
(PSTH) response averaged over all stimulus repetitions ($r_{0,i}(t)$). The predicted firing rate was 469
calculated by scaling the PSTH by a pupil-dependent multiplicative and additive factor to model 470
pupil-dependent changes in gain and baseline firing rate over time (Eqn. 1). To account for a 471
possible nonlinear relationship between pupil size and neuromodulation, the pupil signal was 472
first passed through a static sigmoid nonlinearity, F (double exponential⁶⁵). The parameters of 473

this nonlinearity, as well as the gain (β_0) and baseline (β_1) coefficients were fit independently 474
for each cell using 10-fold jackknifed cross validation using the Neural Encoding Model System 475
(NEMS, <https://github.com/LBHB/NEMS>). 476

$$\hat{r}_i(t) = \beta_0(1 + F(p(t))r_{0,i}(t)) + \beta_1 F(p(t)) \quad (1)$$

The sigmoid transformation applied to the pupil improved model performance, but made 477
it difficult to interpret the gain (β_0) and baseline (β_1) parameters directly. To assess the relative 478
magnitude of pupil effects on baseline and gain, we computed the unique modulation of firing 479
rate due to each, respectively, as the mean difference in prediction between large and small 480
pupil conditions: 481

$$MI_{\beta_0} = \beta_0(\langle F(p_{large}(t)) \rangle - \langle F(p_{small}(t)) \rangle) \quad (2)$$

$$MI_{\beta_1} = \beta_1(\langle F(p_{large}(t)) \rangle - \langle F(p_{small}(t)) \rangle) \quad (3)$$

where $\langle \rangle$ indicates the mean over time and *large* / *small* refer to median splits of pupil 482
size across the experimental session. 483

To quantify overall pupil-dependent modulation without differentiating between baseline 484
and gain, we measured an overall pupil modulation index (*MI*). *MI* was defined by as the mean 485
sound-evoked response when pupil was large minus the mean response when pupil was small, 486
normalized by the sum of these two quantities. Large and small trials were defined based on a 487
median split of pupil size across the entire recording session. 488

$$MI = \frac{\bar{r}_{large} - \bar{r}_{small}}{\bar{r}_{large} + \bar{r}_{small}} \quad (4)$$

Noise correlations: 489

Pairwise noise correlations were measured by grouping spike counts into 250 ms bins, 490
extracting only evoked periods (epochs when sound stimuli were playing), and computing 491
Pearson's correlation between all pairwise combinations of z-scored spike counts. Z-scores 492
were calculated for each stimulus independently, as in Eqn. 5, where $r(t)$ is the single trial 493
response, r_0 is the trial averaged response, and σ is the standard deviation of spike counts 494
across repetitions. Therefore, the z-scored spike counts $Z(t)$ of each neuron i for each stimulus 495

s had mean zero and standard deviation one.

496

$$Z_i(t) = \frac{r_i(t) - r_{0,i}(s)}{\sigma_i(s)} \quad (5)$$

Noise correlations were also measured after first band-pass filtering spike counts. In order to achieve high frequency resolution, data were binned at 100 Hz (10 ms bins) for this analysis. To band-pass filter the data, we first computed the Fast Fourier transform of each unit's spike counts. Next, in the frequency domain we used a Tukey window to extract the frequencies of interest. Finally, we inverted the modified signal back into the time domain, z-scored the data, and computed Pearson's correlation as above. For pupil-corrected noise correlations, we performed the same procedure after first performing the pupil-correction (see below).

497

498

499

500

501

502

503

Pupil-correction:

504

To remove changes in firing rate that were associated with pupil, we regressed out all stimulus-independent variability that could be explained with pupil size. We performed this correction on a per-neuron, per-stimulus basis to completely remove pupil-associated changes in firing.

505

506

507

508

For each neuron (i) and stimulus (s) the mean response across all trials was measured ($\bar{r}_{i,s}$) and subtracted from the true response ($r_{i,s}$, Eqn. 6). A linear regression model was used to predict the residual variability in firing rate with pupil size ($\hat{r}_{resid_{i,s}}$, Eqn. 7). The prediction of this model was then subtracted, and the remaining, pupil-corrected, trial to trial variability was added back to the mean stimulus response (Eqn. 8).

509

510

511

512

513

$$r_{resid_{i,s}} = r_{i,s} - \bar{r}_{i,s} \quad (6)$$

$$\hat{r}_{resid_{i,s}} = \beta_1 p(t) + \beta_0 \quad (7)$$

$$r_{corr_{i,s}} = \bar{r}_{i,s} + (r_{resid_{i,s}} - \hat{r}_{resid_{i,s}}) \quad (8)$$

Pairwise stimulus discrimination:

514

Natural sound samples were broken into non-overlapping 250 ms segments, similar to the procedure followed by Pachitariu et al., 2015.⁶⁶ In order to obtain cross-validated estimates of stimulus discriminability (Figure S2), we required that each stimulus be repeated at least 20 times. This was the case for 11 / 25 recording sites. These 11 recording sites were collected from four different animals. For the remaining sites (which contained only 10 repetitions per stimulus), we could not perform cross-validation. Therefore, for all analyses included in the main text in which we report neural discriminability (d'^2) estimates (Figures 2, 3, 4, and 6) only the high-repetition count recording sessions were included. However, we repeated key analyses using all 25 recording sessions without performing cross-validation and found that results were largely consistent with the more rigorous approach (Figure S8).

515

516

517

518

519

520

521

522

523

524

For each pair of stimulus segments we extracted the N neuron \times k trial response matrices, A and B . Because the number of recorded neurons was greater than the number of stimulus repetitions, we performed dDR to preserve only significant dimensions of the population response (see following section and Figures S3/S4). This allowed us to accurately estimate the population statistics. We quantified encoding accuracy in this reduced-dimensionality space by measuring neural stimulus discriminability, d'^2 , the discrete analog of Fisher information:^{26,27,32,37-39}

525

526

527

528

529

530

$$d'^2 = \Delta\boldsymbol{\mu}^T \boldsymbol{\Sigma}^{-1} \Delta\boldsymbol{\mu} = \Delta\boldsymbol{\mu}^T \mathbf{w}_{opt} \quad (9)$$

where $\Delta\boldsymbol{\mu}$ represents the vector connecting the mean ensemble responses to stimulus A and stimulus B, $\boldsymbol{\Sigma} = \frac{1}{2}(\boldsymbol{\Sigma}_A + \boldsymbol{\Sigma}_B)$ represents the mean noise-covariance matrix, and \mathbf{w}_{opt} is the optimal decoding axis, *i.e.* the vector orthogonal to the optimal linear discrimination hyperplane in state-space. In practice, for our high repetition count data set we estimated \mathbf{w}_{opt} using 50-percent of trials (training data) then projected the held out 50-percent of trials (test data) onto this vector and measured discriminability. For a detailed schematic of this procedure, see Figure S2. Pupil-dependent measurements of d'^2 followed an identical procedure, but before measuring discriminability, the test data was first split in half based on median pupil size.

531

532

533

534

535

536

537

538

Throughout this work, we describe stimulus discriminability as a function of the relationship between the signal and noise subspace. To quantify this relationship, we defined the following two axes which we found to be critical for understanding encoding accuracy:

539

540

541

Discrimination axis magnitude ($|\Delta\boldsymbol{\mu}|$): For a given pair of stimuli A and B, we defined the discrimination axis (*i.e.* the axis containing sensory information) to be the vector $\Delta\boldsymbol{\mu}$. When

542

543

two stimuli drove very different population activity, the magnitude of this vector, $|\Delta\boldsymbol{\mu}|$, was large and discriminability was high.

Noise interference ($|\cos(\theta_{\Delta\boldsymbol{\mu}, \mathbf{e}_1})|$): The structure of trial to trial variability in responses can vary depending on the stimulus.^{25,67,68} Therefore, within the reduced dimensionality dDR space (see below) for each pair of stimuli, we defined the correlated variability axis, \mathbf{e}_1 , as the first eigenvector of the average covariance matrix, $\boldsymbol{\Sigma} = \frac{1}{2}(\boldsymbol{\Sigma}_A + \boldsymbol{\Sigma}_B)$. We then asked how this axis interacted with the signal by measuring the cosine similarity between $\Delta\boldsymbol{\mu}$ and \mathbf{e}_1 . Thus, when the two were perfectly aligned our measure of noise interference was equal to one and when they were perfectly orthogonal, interference was zero.

To quantify the dependence of d'^2 and $\Delta d'^2$ on these two axes, we used the following linear regression models (Eqns. 10, 11). All variables were z-scored prior to model fits.

$$d'^2 = \beta_0 + \beta_1|\cos(\theta_{\Delta\boldsymbol{\mu}, \mathbf{e}_1})| + \beta_2|\Delta\boldsymbol{\mu}| + \beta_3|\cos(\theta_{\Delta\boldsymbol{\mu}, \mathbf{e}_1})||\Delta\boldsymbol{\mu}| \quad (10)$$

$$\Delta d'^2 = \beta_0 + \beta_1|\cos(\theta_{\Delta\boldsymbol{\mu}, \mathbf{e}_1})| + \beta_2|\Delta\boldsymbol{\mu}| + \beta_3|\cos(\theta_{\Delta\boldsymbol{\mu}, \mathbf{e}_1})||\Delta\boldsymbol{\mu}| \quad (11)$$

Decoding-based Dimensionality Reduction:

During our experiments we typically recorded from between 20 and 60 units simultaneously at a single recording site (on one 64-channel probe shank). However, at most, we repeated each individual stimulus only 24 times. Therefore, because the number of neurons was generally greater than the number of trials, estimation of the covariance matrix for a particular stimulus, A ($\boldsymbol{\Sigma}_A$), was unreliable (Figure S4).³² This poses a challenge for estimating neural discriminability as this process depends on accurate estimation of $\boldsymbol{\Sigma}$. One technique that has been proposed to help deal with such limitations is dimensionality reduction.³²

Because we were interested specifically in studying neural discriminability, we developed a decoding-based Dimensionality Reduction (dDR) approach whereby we reduced our N -dimensional space, on a per-stimulus pair basis, to just two dimensions containing the most critical information for estimating stimulus discriminability: The sensory discrimination axis, $\Delta\boldsymbol{\mu}$, and the global noise axis, which we defined as the first principal component of the pooled noise data over all stimuli and neurons (raw spike counts after subtracting the mean stimulus response from each neuron).²³ Thus, the noise axis was fixed between stimulus pairs but $\Delta\boldsymbol{\mu}$ was not. We defined $dDR_1 = \Delta\boldsymbol{\mu}$ and dDR_2 as the axis orthogonal to $\Delta\boldsymbol{\mu}$ and in the plane

spanned by $\Delta\mu$ and the noise axis. We illustrate this process graphically (Figure S3). Analysis of neural discriminability using dDR in simulated data sets (Figure S4) demonstrated the efficacy of this method both in preventing over-fitting and in facilitating accurate estimation of discriminability. Furthermore, though it is not a focus of the current work, we demonstrate that although this method only explicitly preserves the first noise principal component, it is still capable of detecting information limiting correlations²⁶ that recent work has suggested might be restricted to relatively low variance dimensions of the data.^{32,38}

Pupil-dependent changes in stimulus discriminability

d'^2 , measured across pupil states, could vary greatly across the sensory response. Therefore, in order to measure pupil-dependent changes in coding accuracy, we used a normalized metric, $\Delta d'^2$ that allowed us to directly compare relative changes in discriminability for different stimulus pairs and recording sites. To measure this, we computed a pupil-dependent modulation index of d'^2 . For each stimulus pair, $\Delta d'^2$ was defined as the d'^2 measured during large pupil trials minus d'^2 for small pupil trials, normalized by the sum of these two quantities.

$$\Delta d'^2 = \frac{d'^2_{large} - d'^2_{small}}{d'^2_{large} + d'^2_{small}} \quad (12)$$

Contribution of single neuron and correlated variability to changes in discriminability

We observed that pupil modulated both the mean distance between stimuli (discrimination axis / signal magnitude) and the amount of trial to trial variability in the response to the same stimulus. To determine which contributed to pupil-dependent changes in d'^2 , we performed a linear regression between $\Delta d'^2$ and the pupil-dependent changes in both signal magnitude and shared noise variance (Eqn. 13). We also included an interaction term to account for changes in the alignment of the noise and signal axes between arousal states. Each of these statistics was calculated on a per stimulus-pair basis, and the models were fit per experimental session to prevent bias arising from any differences between experiments.

$$\Delta \hat{d}'^2 = \beta_0 + \beta_1 \Delta \text{Signal magnitude} + \beta_2 \Delta \text{Shared noise variance} + \beta_3 \Delta \text{Noise interference} \quad (13)$$

In order to determine if each individual variable explained significant variability in $\Delta d'^2$, we created three different single variable regression models, one for each of the population statistics in the full model (Eqn. 13). We then measured cross-validated estimates of R^2 for each (cvR^2). To do this, models were fit on 90-percent of the data and evaluated on the remaining 10-percent. We used jackknifing so that cross-validated predictions tiled the entire dataset. A single cvR^2 value was measured for each experiment using this prediction. Finally, to determine if the effect of a given predictor was consistent across experiments, we tested whether the distribution of cvR^2 values across experiments was significantly different than zero.

Response simulations:

To dissociate arousal-dependent modulation of single neuron response statistics, such as response gain, from correlated variability effects, we followed a procedure similar to that in Cohen et al., 2009.¹⁶ Simulated data sets were generated independently for both large and small pupil conditions. In the independent variability simulation, the mean and variance of single neurons were allowed to vary between large and small pupil while correlations between neurons were fixed. In the full simulation, correlations between neurons were also allowed to change between large and small pupil.

Independent variability simulation: Pupil-dependent mean and variance of single neurons were measured using the raw data. Next, we built covariance matrices, $\Sigma_{ind.,large}$ and $\Sigma_{ind.,small}$, where the off-diagonal elements in each pupil condition were fixed over all pupil conditions and the diagonal elements of each matrix were set to the measured variance in each respective pupil condition. We simulated population responses, R_{ind} by drawing trials from a multivariate Gaussian distributions as follows:

$$R_{ind.,large} = \mathcal{N}(\mu_{large}, \Sigma_{ind.,large}), \quad R_{ind.,small} = \mathcal{N}(\mu_{small}, \Sigma_{ind.,small}) \quad (14)$$

Full simulations: Pupil-dependent single neuron statistics (mean and variance) and correlated variability between neurons (off-diagonal elements of the covariance matrix) were measured using the raw data. We used these values to define covariance matrices $\Sigma_{full,large}$ and $\Sigma_{full,small}$ where all elements of the matrices were set based on their covariance values in each respective pupil condition. We simulated full population responses, R_{full} as above by

drawing trials from a multivariate Gaussian distributions. All simulations were performed on a per-stimulus basis.

Identification of noise correlation modulation axis

We identified the axis in state space along which noise correlations were most modulated by pupil. First, we measured large and small pupil noise covariance matrices. Noise covariance matrices were defined as the covariance matrix of the z-scored spiking activity (z-scored on a per-stimulus basis). We then set the diagonal of this matrix equal to zero, because for this analysis we were interested only in capturing shared patterns of variability. Finally, we computed the difference covariance matrix (small pupil matrix minus large pupil matrix) and performed an eigen-decomposition. The largest eigenvalue corresponded to the axis along which noise correlations changed most between the two pupil states. Significance of this axis was evaluated with a permutation test.

Statistical analyses

Our data followed a nested structure; multiple cells were recorded from the same animal and many different stimuli were presented during each experimental session. Therefore, it is possible our results could be biased by differences between animals and/or experimental recording session. To account for this, in all of our statistical tests we took one of the following two approaches: (1) Averaged metrics across cells (or pairs of cells) and sound stimuli within a recording session before performing statistical tests or (2) Performed statistical tests using hierarchical bootstrapping.⁶⁹ Although each approach reduces statistical power relative to treating each individual measurement as wholly independent, they provide conservative estimates of p-values and reduce the chance of detecting false positives.⁶⁹ The only exception to this approach was in the estimates of pupil-dependent effects on evoked firing rates in single neurons shown in Figure S1.

For estimating significance of pupil-effects in single neurons, we performed a jackknifed t-test for each individual neuron. Data were split into 20 non-overlapping estimation / validation sets that together tiled the entire experimental session. For each validation set, a prediction correlation was computed for both the full pupil-dependent GLM and the pupil-shuffled GLM. Cells where the mean pupil-dependent model performance was 2 standard errors greater than the mean shuffled model performance ($p \leq 0.05$) were considered to have significant pupil

effects. 654

For all statistical tests measuring large vs. small pupil effects where we first averaged 655
results within recording session, we performed a two-tailed Wilcoxon signed-rank test. For each 656
test, we report the test statistic, W , the p-value, and the exact n number of recording sessions 657
used to perform the test. In cases where we performed a hierarchical bootstrap, we report the 658
direct bootstrap probability of the null hypothesis.⁶⁹ In both cases, we also provide the mean 659
and standard error of the number of measurements per recording session 660

To quantify patterns in d'^2 and $\Delta d'^2$ heatmaps, we performed standard linear regression on 661
a per-recording site basis using the statsmodel package for Python (<https://www.statsmodels.org/stable/index.html>). To determine the significance of regression coefficients at the group 662
level, we performed a Mann-Whitney U test to determine if the parameter distribution over 663
recording sites was significantly different than zero. For each test, we report the test statistic, 664
 U , the p-value, and the exact n number of recording sessions used to perform the test. 665
666

To measure the significance of the correlation between pupil-dependent changes in 667
stimulus discriminability and pupil variance per recording session, we performed a permutation 668
test. We randomly shuffled our distribution of mean $\Delta d'^2$ across recording sites and computed 669
the correlation of these re-sampled values with pupil variance. This was repeated 1000 times 670
in order to calculate a p-value, which is reported in the text along with the true measured 671
correlation coefficient. 672

Finally, to determine if the measured noise correlation modulation axis for a given 673
recording session was significant, we performed a permutation test. For each recording 674
session, pupil size was shuffled in time before classifying trials as large or small pupil. After 675
shuffling, spike covariance matrices were measured for large and small pupil trials and an 676
eigen-decomposition was performed on the difference between these matrices. This was 677
repeated 20 times and the mean / standard error of the resulting eigenvalues was stored. 678
Significant dimensions were those in which the actual measured eigenvalue was 2 standard 679
errors greater than the mean shuffled value. 680

Glossary 681

- *Discrimination/signal axis ($\Delta\mu$):* Vector connecting the mean population response to 682
stimulus A and stimulus B. 683
- *Pooled noise data:* All single trial data after subtracting off mean stimulus responses from 684

each neuron. 685

- *Global noise axis*: First principal component of pooled noise data. 686
- *Correlated variability axis* (\mathbf{e}_1): Stimulus-pair specific estimate of noise axis. First principal component of noise data pooled across the two stimuli within *dDR* space. 687
688
- *Optimal decoding axis* (\mathbf{w}_{opt}): Axis orthogonal to the optimal linear discrimination boundary in *dDR* space. 689
690
- *Discrimination axis magnitude* ($|\Delta\boldsymbol{\mu}|$): Vector magnitude of the discrimination axis. 691
- *Noise interference* ($|\cos(\theta_{\mathbf{e}_1, \Delta\boldsymbol{\mu}})|$): Cosine similarity between the correlated variability axis and discrimination axis. 692
693

Data Availability 694

The datasets analyzed in this study are available from the corresponding author upon reasonable request. 695
696

Author contributions 697

CRH and SVD designed experiments. CRH performed experiments, analyzed data, and developed pupil tracking software. DS assisted with data collection. ZPS assisted with software development for pupil tracking and assisted with data collection. CRH and SVD wrote the manuscript. All authors edited the manuscript. 698
699
700
701

Acknowledgments 702

This work was supported by a National Science Foundation Graduate Research Fellowship (NSF GRFP, GVPRS0015A2) (CRH), the National Institute of Health (NIH, R01 DC0495) (SVD), Achievement Rewards for College Scientists (ARCS) Portland chapter (CRH), and by the Tartar Trust at Oregon Health and Science University (CRH). 703
704
705
706

References

1. Nima Mesgarani and Edward F. Chang. Selective cortical representation of attended speaker in multi-talker speech perception. Nature, 485(7397):233–236, May 2012. Number: 7397 Publisher: Nature Publishing Group.
2. Zachary P. Schwartz, Brad N. Buran, and Stephen V. David. Pupil-associated states modulate excitability but not stimulus selectivity in primary auditory cortex. Journal of Neurophysiology, 123(1):191–208, November 2019. Publisher: American Physiological Society.
3. Daniela Saderi, Zachary P. Schwartz, Charlie R. Heller, Jacob R. Pennington, and Stephen V. David. Dissociation of task engagement and arousal effects in auditory cortex and midbrain. bioRxiv, page 2020.06.16.155432, June 2020. Publisher: Cold Spring Harbor Laboratory Section: New Results.
4. A. Ryan and J. Miller. Effects of behavioral performance on single-unit firing patterns in inferior colliculus of the rhesus monkey. Journal of Neurophysiology, 40(4):943–956, July 1977.
5. Gonzalo H. Otazu, Lung-Hao Tai, Yang Yang, and Anthony M. Zador. Engaging in an auditory task suppresses responses in auditory cortex. Nature Neuroscience, 12(5):646–654, May 2009.
6. Chen-Chung Lee and John C. Middlebrooks. Auditory cortex spatial sensitivity sharpens during task performance. Nature Neuroscience, 14(1):108–114, January 2011.
7. Daniel P. Knudsen and Timothy Q. Gentner. Active recognition enhances the representation of behaviorally relevant information in single auditory forebrain neurons. Journal of Neurophysiology, 109(7):1690–1703, April 2013.
8. Pingbo Yin, Jonathan B. Fritz, and Shihab A. Shamma. Rapid Spectrotemporal Plasticity in Primary Auditory Cortex during Behavior. The Journal of Neuroscience, 34(12):4396–4408, March 2014.
9. S. J. Slee and S. V. David. Rapid Task-Related Plasticity of Spectrotemporal Receptive Fields in the Auditory Midbrain. Journal of Neuroscience, 35(38):13090–13102, September 2015.

10. Kishore V Kuchibhotla, Jonathan V Gill, Grace W Lindsay, Eleni S Papadoyannis, Rachel E Field, Tom A Hindmarsh Sten, Kenneth D Miller, and Robert C Froemke. Parallel processing by cortical inhibition enables context-dependent behavior. Nature Neuroscience, 20(1):62–71, January 2017.
11. Jonathan Fritz, Shihab Shamma, Mounya Elhilali, and David Klein. Rapid task-related plasticity of spectrotemporal receptive fields in primary auditory cortex. Nature Neuroscience, 6(11):1216–1223, November 2003. Number: 11 Publisher: Nature Publishing Group.
12. Zachary P Schwartz and Stephen V David. Focal Suppression of Distractor Sounds by Selective Attention in Auditory Cortex. Cerebral Cortex, 28(1):323–339, January 2018.
13. Pei-Ann Lin, Samuel K. Asinof, Nicholas J. Edwards, and Jeffry S. Isaacson. Arousal regulates frequency tuning in primary auditory cortex. Proceedings of the National Academy of Sciences, November 2019.
14. J. D. Downer, M. Niwa, and M. L. Sutter. Task Engagement Selectively Modulates Neural Correlations in Primary Auditory Cortex. Journal of Neuroscience, 35(19):7565–7574, May 2015.
15. Joshua D. Downer, Brittany Rapone, Jessica Verhein, Kevin N. O’Connor, and Mitchell L. Sutter. Feature-Selective Attention Adaptively Shifts Noise Correlations in Primary Auditory Cortex. The Journal of Neuroscience, 37(21):5378–5392, May 2017.
16. Marlene R Cohen and John H R Maunsell. Attention improves performance primarily by reducing interneuronal correlations. Nature Neuroscience, 12(12):1594–1600, December 2009.
17. Martin Vinck, Renata Batista-Brito, Ulf Knoblich, and Jessica A. Cardin. Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. Neuron, 86(3):740–754, May 2015.
18. Neil C. Rabinowitz, Robbe L. Goris, Marlene Cohen, and Eero P. Simoncelli. Attention stabilizes the shared gain of V4 populations. eLife, 4:e08998, November 2015.
19. George H. Denfield, Alexander S. Ecker, Tori J. Shinn, Matthias Bethge, and Andreas S. Tolias. Attentional fluctuations induce shared variability in macaque primary visual cortex. Nature Communications, 9(1), December 2018.

20. Alexander S. Ecker, George H. Denfield, Matthias Bethge, and Andreas S. Tolias. On the Structure of Neuronal Population Activity under Fluctuations in Attentional State. The Journal of Neuroscience, 36(5):1775–1789, February 2016.
21. Robbe L T Goris, J Anthony Movshon, and Eero P Simoncelli. Partitioning neuronal variability. Nature Neuroscience, 17(6):858–865, June 2014.
22. Benjamin R. Cowley, Adam C. Snyder, Katerina Acar, Ryan C. Williamson, Byron M. Yu, and Matthew A. Smith. Slow Drift of Neural Activity as a Signature of Impulsivity in Macaque Visual and Prefrontal Cortex. Neuron, 0(0), August 2020. Publisher: Elsevier.
23. Dmitry Kobak, Jose L Pardo-Vazquez, Mafalda Valente, Christian K Machens, and Alfonso Renart. State-dependent geometry of population activity in rat auditory cortex. eLife, 8:e44526, April 2019.
24. Jalal K. Baruni, Brian Lau, and C. Daniel Salzman. Reward expectation differentially modulates attentional behavior and activity in visual area V4. Nature Neuroscience, 18(11):1656–1663, November 2015. Number: 11 Publisher: Nature Publishing Group.
25. Joel Zylberberg, Jon Cafaro, Maxwell H. Turner, Eric Shea-Brown, and Fred Rieke. Direction-Selective Circuits Shape Noise to Ensure a Precise Population Code. Neuron, 89(2):369–383, January 2016.
26. Rubén Moreno-Bote, Jeffrey Beck, Ingmar Kanitscheider, Xaq Pitkow, Peter Latham, and Alexandre Pouget. Information-limiting correlations. Nature Neuroscience, 17:1410, September 2014.
27. Bruno B. Averbeck and Daeyeol Lee. Effects of noise correlations on information encoding and decoding. Journal of Neurophysiology, 95(6):3633–3644, June 2006.
28. A. S. Ecker, P. Berens, A. S. Tolias, and M. Bethge. The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. Journal of Neuroscience, 31(40):14272–14283, October 2011.
29. Felix Franke, Michele Fiscella, Maksim Sevelev, Botond Roska, Andreas Hierlemann, and Rava Azeredo da Silveira. Structures of Neural Correlation and How They Favor Coding. Neuron, 89(2):409–422, January 2016.

30. Carsen Stringer, Michalis Michaelos, and Marius Pachitariu. High precision coding in visual cortex. bioRxiv, page 679324, November 2019.
31. Matthew J. McGinley, Stephen V. David, and David A. McCormick. Cortical Membrane Potential Signature of Optimal States for Sensory Signal Detection. Neuron, 87(1):179–192, July 2015.
32. Oleg I. Romyantsev, Jérôme A. Lecoq, Oscar Hernandez, Yanping Zhang, Joan Savall, Radosław Chrapkiewicz, Jane Li, Hongkui Zeng, Surya Ganguli, and Mark J. Schnitzer. Fundamental bounds on the fidelity of sensory cortical coding. Nature, March 2020.
33. Matthew R. Whiteway, Bruno Averbeck, and Daniel A. Butts. A latent variable approach to decoding neural population activity. bioRxiv, page 2020.01.06.896423, January 2020. Publisher: Cold Spring Harbor Laboratory Section: New Results.
34. Jiangang Du, Timothy J. Blanche, Reid R. Harrison, Henry A. Lester, and Sotiris C. Masmanidis. Multiplexed, High Density Electrophysiology with Nanofabricated Neural Probes. PLOS ONE, 6(10):e26204, October 2011.
35. Josh H. McDermott, Michael Schemitsch, and Eero P. Simoncelli. Summary statistics in auditory perception. Nature Neuroscience, 16(4):493–498, April 2013. Number: 4 Publisher: Nature Publishing Group.
36. Tatiana A. Engel, Nicholas A. Steinmetz, Marc A. Gieselmann, Alexander Thiele, Tirin Moore, and Kwabena Boahen. Selective modulation of cortical state during spatial attention. Science, 354(6316):1140–1144, December 2016.
37. L. F. Abbott and Peter Dayan. The Effect of Correlated Variability on the Accuracy of a Population Code. Neural Computation, 11(1):91–101, January 1999.
38. Mohammad Mehdi Kafashan, Anna Jaffe, Selmaan N. Chettih, Ramon Nogueira, Iñigo Arandia-Romero, Christopher D. Harvey, Rubén Moreno-Bote, and Jan Drugowitsch. Scaling of information in large neural populations reveals signatures of information-limiting correlations. bioRxiv, page 2020.01.10.902171, January 2020.
39. Ramon Bartolo, Richard C. Saunders, Andrew R. Mitz, and Bruno B. Averbeck. Information limiting correlations in large neural populations. The Journal of Neuroscience, pages 2072–19, January 2020.

40. Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. Science, 364(6437):eaav7893, April 2019.
41. Alexander S. Ecker, Philipp Berens, R. James Cotton, Manivannan Subramanian, George H. Denfield, Cathryn R. Cadwell, Stelios M. Smirnakis, Matthias Bethge, and Andreas S. Tolias. State Dependence of Noise Correlations in Macaque Primary Visual Cortex. Neuron, 82(1):235–248, April 2014.
42. P. L. E. van Kan, R. P. Scobey, and A. J. Gabor. Response covariance in cat visual cortex. Experimental Brain Research, 60(3):559–563, November 1985.
43. K. H. Britten, M. N. Shadlen, W. T. Newsome, and J. A. Movshon. The analysis of visual motion: a comparison of neuronal and psychophysical performance. Journal of Neuroscience, 12(12):4745–4765, December 1992.
44. Ehud Zohary, Michael N. Shadlen, and William T. Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature, 370(6485):140–143, July 1994.
45. M. N. Shadlen, K. H. Britten, W. T. Newsome, and J. A. Movshon. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. Journal of Neuroscience, 16(4):1486–1510, February 1996.
46. A. J. Parker and W. T. Newsome. SENSE AND THE SINGLE NEURON: Probing the Physiology of Perception. Annual Review of Neuroscience, 21(1):227–277, March 1998.
47. Bruno B. Averbeck, Peter E. Latham, and Alexandre Pouget. Neural correlations, population coding and computation. Nature Reviews Neuroscience, 7(5):358–366, May 2006.
48. John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings. Nature Neuroscience, 17(11):1500–1509, November 2014.
49. Matthew J. McGinley, Martin Vinck, Jacob Reimer, Renata Batista-Brito, Edward Zagher, Cathryn R. Cadwell, Andreas S. Tolias, Jessica A. Cardin, and David A. McCormick. Waking State: Rapid Variations Modulate Neural and Behavioral Responses. Neuron, 87(6):1143–1161, September 2015.

50. Shi Tong Liu, Pilar Montes-Lourido, Xiaoqin Wang, and Srivatsun Sadagopan. Optimal features for auditory categorization. Nature Communications, 10(1):1302, March 2019. Number: 1 Publisher: Nature Publishing Group.
51. Carrie J. McAdams and John H. R. Maunsell. Effects of Attention on Orientation-Tuning Functions of Single Neurons in Macaque Cortical Area V4. Journal of Neuroscience, 19(1):431–441, January 1999. Publisher: Society for Neuroscience Section: ARTICLE.
52. Jacob Reimer, Matthew J. McGinley, Yang Liu, Charles Rodenkirch, Qi Wang, David A. McCormick, and Andreas S. Tolias. Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. Nature Communications, 7:13289, 2016.
53. Jennifer K. Bizley and Andrew J. King. Visual influences on ferret auditory cortex. Hearing Research, 258(1):55–63, December 2009.
54. David M. Schneider, Anders Nelson, and Richard Mooney. A synaptic and circuit basis for corollary discharge in the auditory cortex. Nature, 513(7517):189–194, August 2014.
55. Daniel E. Winkowski, Sharba Bandyopadhyay, Shihab A. Shamma, and Patrick O. Kanold. Frontal Cortex Activation Causes Rapid Plasticity of Auditory Cortical Processing. Journal of Neuroscience, 33(46):18134–18148, November 2013. Publisher: Society for Neuroscience Section: Articles.
56. S. V. David, J. B. Fritz, and S. A. Shamma. Task reward structure shapes rapid receptive field plasticity in auditory cortex. Proceedings of the National Academy of Sciences, 109(6):2144–2149, February 2012.
57. Daniela Saderi, Brad N. Buran, and Stephen V. David. Streaming of Repeated Noise in Primary and Secondary Fields of Auditory Cortex. Journal of Neuroscience, 40(19):3783–3798, May 2020. Publisher: Society for Neuroscience Section: Research Articles.
58. Jacob Pennington and Stephen David. Complementary effects of adaptation and gain control on sound encoding in primary auditory cortex. preprint, Neuroscience, January 2020.
59. Jennifer K. Bizley, Fernando R. Nodal, Israel Nelken, and Andrew J. King. Functional organization of ferret auditory cortex. Cerebral Cortex (New York, N.Y.: 1991), 15(10):1637–1653, October 2005.

60. Diego Elgueda, Daniel Duque, Susanne Radtke-Schuller, Pingbo Yin, Stephen V. David, Shihab A. Shamma, and Jonathan B. Fritz. State-dependent encoding of sound and behavioral meaning in a tertiary region of the ferret auditory cortex. Nature Neuroscience, 22(3):447–459, 2019.
61. Joshua H Siegle, Aarón Cuevas López, Yogi A Patel, Kirill Abramov, Shay Ohayon, and Jakob Voigts. Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. Journal of Neural Engineering, 14(4):045003, August 2017.
62. Marius Pachitariu, Nicholas Steinmetz, Shabnam Kadir, Matteo Carandini, and Harris Kenneth D. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv, page 061481, June 2016.
63. Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. arXiv:1608.06993 [cs], January 2018. arXiv: 1608.06993.
64. Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In 2009 IEEE Conference on Computer Vision and Pattern Recognition, pages 248–255, June 2009. ISSN: 1063-6919.
65. Ivar L. Thorson, Jean Liénard, and Stephen V. David. The Essential Complexity of Auditory Receptive Fields. PLOS Computational Biology, 11(12):e1004628, December 2015.
66. M. Pachitariu, D. R. Lyamzin, M. Sahani, and N. A. Lesica. State-Dependent Population Coding in Primary Auditory Cortex. Journal of Neuroscience, 35(5):2058–2073, February 2015.
67. Adam Kohn and Matthew A. Smith. Stimulus Dependence of Neuronal Correlation in Primary Visual Cortex of the Macaque. Journal of Neuroscience, 25(14):3661–3673, April 2005.
68. Douglas A. Ruff and Marlene R. Cohen. Stimulus Dependence of Correlated Variability across Cortical Areas. Journal of Neuroscience, 36(28):7546–7556, July 2016.
69. Varun Saravanan, Gordon J. Berman, and Samuel J. Sober. Application of the hierarchical bootstrap to multi-level data in neuroscience. arXiv:2007.07797 [q-bio], July 2020. arXiv: 2007.07797.