Time-dependent processing of correlated signals between visual cortical areas improves neuronal population encoding in a sensory discrimination task.

Jackson E. T. Smith* and Andrew J. Parker

Department of Physiology, Anatomy and Genetics
University of Oxford
Oxford
Oxfordshire
OX1 3PT
United Kingdom

*Correspondence: jackson.smith@dpag.ox.ac.uk

Character count (with spaces): Actual/Maximum

- Title: 148/150
- Summary: 142/150 words
- Introduction: 3,690
- Results: 21,811
- Discussion: 12,135
- Main figure legends: 7,344
- Total (Intro, Results, Discussion, Fig legends): 44,980/45,000

SUMMARY

Human psychophysical studies demonstrate that visual detection thresholds are sometimes close to the limits imposed by the physics of the stimulus. Another potential limit is the variable spiking of cortical neurons, which theoretically reduces information about the stimulus. The correlation structure of noise in the cortical population is critical in setting this limit. We investigated these correlations by recording simultaneously from visual cortical areas V1 and V4 in macaque monkeys during performance of a stereo depth detection task. We found evidence of rapid, information-limiting, noise correlations within each area, at a temporal scale of tens of ms. However, the correlation structure between the two areas had a different pattern at temporal scales of 100+ms that attenuated within-area correlations, potentially supporting recovery of information. We suggest that processing in multiple, distinct cortical areas may aid the estimation and attenuation of information-limiting noise correlations.

Key words: V1, V4, Cortex, Vision, Rhesus, Noise correlation, Stereo depth, Encoding, Behaviour, Utah array

INTRODUCTION

The weakest detectable visual stimuli are close in intensity to the limits imposed by the physics of the stimulus (Hecht et al., 1941, Parker and Newsome, 1998, Levick et al., 1983). Processing in central neocortical areas may also place limits on sensory signals. Earlier work questioned how the firing induced by stimuli could be detected against the variability of cortical neuronal activity (Tolhurst et al., 1983, Parker and Hawken, 1985, Barlow et al., 1987, Newsome et al., 1989), elucidating how detection relates to the statistics of firing in single neurons. This question has been explored with different tasks and sensory systems (Parker and Newsome, 1998), and remains a topic of intensive research (Pitkow et al., 2015, Krug et al., 2016, Krause and Ghose, 2018). A good model is the processing of stereoscopic information, which clearly requires the cortex to combine separate images from the two eyes to see stereo depth.

Early work on the somatosensory periphery showed that the covariation of residual noise in a pool of sensory neurons is also an important factor (Johnson et al., 1973). Without covariance, additive pooling results in a steady improvement of sensory thresholds with increasing pool size. Covariance limits the improvement. Recent work expounds that insight, showing how the structure of correlations is critical in setting limits on the encoding and decoding of information within sensory cortex (Abbott and Dayan, 1999, Pouget et al., 2003, Averbeck et al., 2006, Moreno-Bote et al., 2014, Kanitscheider et al., 2015b). This theoretical work has been supported by the development of techniques that record simultaneously from multiple cortical neurons.

Yet, few studies make the critical measurement of correlation structure in the population responses of animals performing a sensory discrimination task (Pitkow et al., 2015, Yu and Gu, 2018, Bondy et al., 2018, Sanayei et al., 2018, Cohen and Maunsell, 2009, Ruff and Cohen, 2014, Ruff and Cohen, 2016, Ni et al., 2018, Verhoef and Maunsell, 2017). Most employ tasks that involve a binary decision, potentially limiting which cortical neurons are relevant to the task. We addressed this by training monkeys to discriminate stereo depth, while we probed a range of neuronal encoding by recording multiple neurons simultaneously in two cortical sites: the primary visual cortex (V1) and the extrastriate area V4.

An open question is whether correlations limit the passage of signals through stages of cortical processing (Zylberberg et al., 2017). Since human psychophysical performance is near the physical limits of sensory detection (Hecht et al., 1941), communication through the cortex must be almost noise-free under some circumstances. We tackle this issue by examining the correlation structure and information encoding of neuronal activity within and between cortical areas V1 and V4. Specifically, we test how information may pass between cortical areas, with respect to neuronal correlations during task performance.

We couple our approach with a time-dependent measure of correlations (Bair et al., 2001), and extend this approach to develop time-dependent measures of information encoding. Like previous studies, we find evidence of information-limiting correlations for the pattern of cortical activity within single cortical areas. Our results suggest a simple mode, which

demonstrates that the potentially deleterious effects of information-limiting correlations do not build up as sensory signals pass from one cortical area to another. We suggest that this is an important and previously unrecognized principle of neural encoding across multiple areas of the cerebral neocortex.

RESULTS

We trained two Rhesus macaques (*Macaca mulatta*) to observe 4 patches of random dot stereograms (RDS) and perform an odd-one-out task (Figure 1A&B), in which they detected the one patch that presented a change in binocular disparity. In the 'present RDS' phase (Fig.1Aiii), all 4 targets had the same disparity, which was chosen from a set of 13 different values. The 'popout' phase (iv) presented the change in depth.

The monkeys' behaviour indicated attention to the RDS stimuli in the 'Present RDS' phase, as they were better at reporting the location of the disparity step change as it grew in magnitude (Fig.1B, % correct), while their reaction times decreased (RT). On catch trials with no change in any RDS, the animals had to guess, so their accuracy was at chance levels (Fig.1B, dotted line) and their RTs were longest.

V1 and V4 neurons were measured using Utah array electrodes implanted chronically in each area. The electrode placement meant that the recorded neurons were stimulated by the lower-right RDS. The RDS was located to drive significant co-activation of the V1 and V4 units (Fig.1C&D, Fig.S1). It was impossible to optimally stimulate every unit recorded by the arrays. Hence, we applied entry criteria, such that each unit had statistically significant tuning to disparity (one-way ANOVA at the 1% level) and fired at least 5 spikes per second to its preferred disparity. The recordings were made as the animals attended to the RDS stereo depth, to perform the task. We focus on the spiking from the 'Hold fixation' and 'Present RDS' phases of the trial (Fig.1Aii&iii).

V1 and V4 units were selective for binocular disparity

The V1 and V4 neural signals that we recorded could have supported task performance. Disparity tuning curves are shown for example V1 (circles) and V4 (plus sign) units in Figure 2A, along with the best-fitting Gabor functions (solid lines; V1, gold; V4, blue; Prince et al., 2002), which describe the response profiles as a function of disparity. To quantify a unit's level of disparity selectivity, we computed the rate of mutual information between the RDS disparity and the spike counts, called the disparity mutual information (DMI). This measure was well correlated with the established measure of selectivity (Fig.2B & Fig.S2A) – the disparity discrimination index (DDI; Prince et al., 2002) – for the V1 (dots; Spearman correlation, Initial, $r = 0.976$, $p = 9.559e{-}130$) and V4 (plus sign; $r = 0.953$, $p = 4.705e{-}111$) neurons.

We tested when the units became selective for disparity by convolving the spike trains with a causal exponential kernel (20ms time constant) to obtain the firing rate and DMI time series for each unit. The average firing rates (Fig.2C & Fig.S2B; V1, gold; V4, blue) had a transient burst of activity in response to the onset (black dotted) of the RDS, which then settled down to a steady rate. As expected, V1 began responding to the RDS before V4.

Units in both areas were selective for binocular disparity almost as soon as they began responding to the RDS, with the average DMI rising and falling in a similar way to the firing rate (Fig.2D & Fig.S2C). We verified this similarity for each unit by finding the Spearman

correlation of the firing rate and DMI, matched by time bin, in an 800ms window starting 44ms after the RDS onset. The result was significantly positive for V1 (mean r = 0.192, right-tailed t-test p = 3.339e−22) and V4 (r = 0.245, p = 7.819e−31), showing that units were more selective for disparity when their firing rate was high.

Thus, we partitioned each trial into three consecutive windows for further analysis (Fig.2C & S2B, top). The Spontaneous window (leftward arrow) captured neural responses during the 'Hold fixation' phase of the trial (Fig.1Aii), in which the subjects held their gaze prior to the arrival of RDS signals in V1. The Initial window (thick bar and black solid lines) captured the transient burst of firing rate and DMI in response to the onset of the RDS. Lastly, the Steady state window (rightward arrow) captured responses to the RDS for the remainder of the 'Present RDS' phase of the trial (Fig.1Aiii). Again, within-unit comparisons showed that disparity selectivity was stronger in the Initial than in the Steady state window, for V1 (mean paired difference = −0.577 bits/s, left-tailed t-test p = 1.294e−12) and V4 (−0.220 bits/s, p = 6.827e−11).

Information limiting correlations

If the visual cortex performs linear computations, then a special class of noise correlation – differential correlation (Moreno-Bote et al., 2014) – has the unique ability to limit the information in a sensory pool. It could build up in one layer of neurons (*e.g.* V4) after pooling common noisy inputs from another (V1; Kanitscheider et al., 2015c). We give an example pair of V4 neurons in Figure 3A–C to show how differential correlations could limit information.

The pair had similar tuning for binocular disparity (*i.e.* $0 < r_{signal}$, Fig.3A). Both units had a common negative slope in their tuning curves around 0.1° (black dashed), where they were sensitive to the same small changes in disparity. Joint responses to 0.1° trials were obtained by counting spikes in the Initial window (Fig.3B, grey). At 0.1°, the pair had a positive noise correlation (Fig.3C). Critically, the principal component (arrow) of the covariance was nearly parallel to the tangent line (grey) of the tuning curves at 0.1°.

To see the consequence of this, consider a linear discriminator that classifies the joint responses as one of two disparities on either side of 0.1°. The optimal decision boundary is perpendicular to the tangent line (grey), but the noise correlation (arrow) makes the joint responses vary across the decision boundary, causing classification errors (Averbeck et al., 2006). Put simply, correlated noise would mimic true changes in the stimulus, limiting the sensitivity of this pair.

This outcome contrasts a second, inter-cortical pair (Fig.3D–F). As before, both units had steep, overlapping slopes in their tuning curves (Fig.3D). Unlike the V4/V4 pair, the V1 and V4 units were selective for opposite disparities ($r_{signal} < 0$). Spike counts from both units in response to 0° (Fig.3D, black dashed) were obtained as before (Fig.3E), and the joint responses were examined (Fig.3F). Although this pair has a negative signal correlation, they have a positive noise correlation: the noise correlation does not align to the tangent. In this case, correlated noise would be less likely to leak across the decision boundary.

r_{CCG} measures noise correlation at different time scales

Noise correlations encompass the synchronous firing of spikes up to long-term co-fluctuations in spike numbers (Smith and Kohn, 2008, Smith and Sommer, 2013, Cohen and Kohn, 2011). But measurement of spike count correlation (r_{SC}) often confounds these slow and rapid correlations. Another, time-based measure – r_{CCG} (Bair et al., 2001) – was introduced to disentangle noise correlations at different temporal scales. r_{CCG} is an integral of the spike-train cross-correlation, and converges upon r_{SC} as the integral width (τ) reaches that of the analysis window. We asked if information-limiting correlations occupied a limited set of temporal scales by using r_{CCG} to examine noise correlations within and between V1 and V4 (Figures 4 & S3).

Within V1 (Fig.4A), the time scale of noise correlations changed as the trial progressed. Before the RDS appeared, Spontaneous correlations were slower, on the order of hundreds of ms (dashed). The correlations got faster after RDS onset, peaking at τ = 20ms; r_{CCG}(20ms) was greater in the Initial window (solid black) than the Spontaneous window (mean paired difference = 2.141e−2, paired t-test p = 1.592e−124). Even in steady-state responses to the RDS, the slower correlations were supressed, with a drop in r_{CCG} (100ms) from the Spontaneous window to the Steady state (magenta, −1.092e−2, p = 1.005e−21).

These changes in r_{CCG} over time reflect the V1 to V1 cross-correlograms (CCG, Fig.4B). The Spontaneous CCG was positive and had a broad peak, coinciding with the steady increase of Spontaneous r_{CCG}. In contrast, the Initial and Steady state CCGs had narrow positive peaks flanked by negative side lobes, showing that the majority of correlated spikes occurred within tens of ms of each other. This matched the peak in Initial and Steady state r_{CCG} at τ = 20ms.

Unlike V1, correlations within V4 (Figs.4D & S3A) began saturating at τ = 20ms in all three analysis windows. V4 correlations were also stronger in the Spontaneous window, with r_{CCG}(20ms) dropping in the Initial window (−2.188e−2, p = 9.053e−105) and then changing little from the Initial to Steady state windows (9.464e−4, p = 0.041). These differences appeared in the V4 CCGs (Figs.4E & S3B) in which the Initial and Steady state CCGs were mainly attenuated versions of the Spontaneous CCG.

In marked contrast, the correlations between V1 and V4 (Fig.4G) saturated at τ = 100ms. Though weak, they were significantly positive (mean r_{CCG}(100ms); Initial = 8.075e−3, right-tailed t-test p = 1.512e−43; Steady state = 7.398e−3, p = 4.088e−61), but only in response to the RDS (Spontaneous = −1.286e−3, p = 1). There was no change in the V1/V4 correlations from the Initial to the Steady state window (mean difference = −6.767e−4, paired t-test p = 0.159). The V1/V4 CCGs (Fig.4H) reveal a consistent delay whereby V1 spikes tended to occur 50ms earlier than correlated V4 spikes.

Importantly, r_{CCG} provides a low-variance estimate of short-term correlations, compared to r_{SC} (Bair et al., 2001). τ can be chosen so that $r_{CCG}(\tau)$ integrates the central peak of the CCG and discards the slow components, which r_{SC} captures. This point is made by a direct

comparison of Steady state r_{CCG}(100ms) and r_{SC} in V1 (Fig.4C), V4 (Fig.4F & Fig.S3C), and between areas (Fig.4I). There was a strong, positive Spearman correlation of r_{CCG} and r_{SC} in all cases (V1 r = 0.750, V4 r = 0.772, V1/V4 r = 0.807, all $p \rightarrow 0$), showing that they measured a common source of correlation. But the standard deviation (σ) of r_{SC} (V1 σ_{SC} = 0.135; V4 σ_{SC} = 0.112; V1/V4 σ_{SC} = 0.112) was greater than that of r_{CCG}(100ms) (V1 σ_{CCG} = 0.044, two-sample F-test $F_{2418,2418}$ = 9.569; V4 σ_{CCG} = 0.056, $F_{2225,2225}$ = 3.955; V1/V4 σ_{CCG} = 0.030, $F_{4597,4597}$ = 13.949; all $p \rightarrow 0$). Hence, r_{CCG} is the more accurate measure of short-term correlations.

Noise correlations varied with disparity tuning, temporal scale, and time

A property of differential correlations is that they limit discrimination thresholds in the neighbourhood of a specific stimulus value (Moreno-Bote et al., 2014). Therefore, we asked if the noise correlations had any consistent relationship to disparity selectivity (Figure 5 & S4). If so, then embedded differential correlations may have been present. Thus, we plotted r_{CCG} as a function of the signal correlation between pairs (r_{signal}), grouped by temporal scale (τ, Fig.5A,C,E & S4A), or as a function of τ, grouped by r_{signal} (Fig.5B,D,F & S4B).

For pairs in V4, there was a clear positive relationship of r_{CCG} and r_{signal} (Fig.5C & S4A) at short (τ = 20ms, dashed) and long (τ = 150ms, solid) time scales in the Initial (black) and Steady state (magenta) windows: Spearman correlations were significant in all four cases (Initial τ = 20ms, r = 0.172, p = 4.133e−16; Initial τ = 150ms, r = 0.141, p = 2.448e−11; Steady state τ = 20ms, r = 0.170, p = 5.674e−16; Steady state τ = 150ms, r = 0.131, p = 5.906e−10). This coincided with an increase in r_{CCG} at a range of τ for pairs with similar disparity preferences (Fig.5D & S4B, $r_{signal} > 0$, dashed) compared to those with opposite preferences ($r_{signal} < 0$, solid). For pairs in V1, a similar but weaker positive relationship was observed (Fig.5A&B) that was significant in the Initial window (τ = 20ms, r = 0.042, p = 0.038; τ = 150ms, r = 0.050, p =0.013) and Steady state window for long time scales (τ = 150ms, r = 0.054, p = 0.009) but not short (τ = 20ms, r = 0.030, p = 0.144).

The correlation patterns between V1 and V4 (Fig.5E&F) were therefore unexpected in two ways. First, there was no measurable relationship between r_{CCG}(20ms) and r_{signal} (Initial r = −0.013, p = 0.392; Steady state r = −0.004, p = 0.779). second, at longer time-scales, a weak relationship emerged, but it was negative (τ = 150ms; Initial r = −0.032, p = 0.032; Steady state r = −0.025, p = 0.091). That is, V1/V4 pairs tended to have opposite noise and signal correlations (like in Fig 3D–F), unlike the pairs within each area (*e.g.* Fig.3A–C).

Empirical evidence of differential noise correlations

Theory predicts that measured noise correlations are the sum of two sources, one that limits the sensory threshold with differential correlations and one that does not (Moreno-Bote et al., 2014). The differential correlations will scale with the product of the first derivatives of the tuning curve (Fig.3C). The second source does not limit information, but it does mask the differential correlations. We reasoned that the low variance measure of noise correlation, r_{CCG}, would be more likely to unmask weak and rapid differential

correlations. Thus, we examined the relationship of r_{CCG} and the product of derivatives over a range of time scales.

In V1 (Fig 6A–C) and within V4 (Fig 6D–F & Fig.S5E–G), we found that noise correlations had a positive relationship with the product of derivatives (f'×f'), shown by plotting r_{CCG} against f'×f' (V1 Fig6A &B 6B; V4 Fig 6D & 6E, Fig.S5E & S5F) using the time scales (τ, see arrows in Fig 6C & 6F) that maximized regression line slopes (orange). In these examples, the Spearman correlation was significantly positive for V1 (Fig 6A, Initial r = 0.033, p = 1.137e−6; Fig 6B, Steady state r = 0.049, p = 3.423e−13) and V4 (Fig 6D, Initial r = 0.072, p = 1.437e−24; Fig 6E, Steady state r = 0.079, p = 2.726e−29). To see how this relationship changed with the temporal scale of the noise correlations, we plotted the Spearman correlation as a function of τ for V1 (Fig 6C) and V4 (Fig 6F & S5G). This showed that r_{CCG} was the most related to f'×f' in a narrow range of time scales ($10 \leq \tau \leq 100$ms).

Again, the result between V1 and V4 was unexpected (Fig 6G–I). Instead of a positive relationship, we observed a significantly negative relationship of r_{CCG} and f'×f' in the Initial (Fig 6G, r = −0.026, p = 1.023e−7) and Steady state (H, r = −0.019, p = 9.298e−5) windows that peaked for a narrow range of τ close to 100ms (I).

Differential noise correlations are attenuated at long time scales

Differential correlations will result in a saturation of the information in a neuronal pool of a certain size. Therefore, we measured the depth information in pools of increasing size (Figure 7 & S6). Using Initial responses, we found the normalized and averaged bias-corrected linear Fisher information (Kanitscheider et al., 2015b) for pools of V1 units (Fig.7A), V4 units(Fig 7B & Fig.S6A), or V1 and V4 units (Fig 7C). Fisher information increases when the population response can better distinguish between different RDS disparity values, and results from the addition of informative neurons or from changes in the pool's correlations. Since we added units to each pool from the most selective for disparity down to the least selective (Sanayei et al., 2018), part of the apparent saturation in Fig 7A – C (& Fig.S6A) is not due to the correlations.

We then compared the Fisher information in the empirical pool against the information in an equivalent pool of decorrelated units (Fig 7A–C, thick dashed), for the cases of rapid (τ = 20ms, dotted) and slow (τ = 150ms, solid) noise correlations. For all pool types, we found that the empirical information at τ = 20ms began to diverge from the decorrelated information for pools of 10 units. In pools with just V1 or V4 units, there was a similar but weaker divergence when τ = 150ms. However, for the mixed V1 and V4 pool, the empirical information at τ = 150ms was closely aligned with the decorrelated information (Fig 7C).

To get more detail about the effect of time scale, we plotted the average raw Fisher information of the Initial responses as a function of τ (V1, Fig 7D; V4, Fig 7E & Fig S6B; mixed, Fig 7F). The empirical information (solid) was compared against the decorrelated information (dashed). For all pool types, the empirical information tended to increase with τ. But only the empirical information only converged on the decorrelated information in the

mixed V1 and V4 pool, around τ = 100ms. The normalized, paired difference in decorrelated and empirical information yielded similar results (Fig 7L, left; Fig.S6D, top).

We confirmed this result by plotting the decorrelated information against the empirical information at τ = 150ms (V1 ,Fig 7G; V4 ,Fig 7H & Fig.S6C; mixed Fig 7I). The decorrelated information was greater than empirical for the V1 (mean paired difference 'PD' = 66.117, paired-sample t-test p = 4.700e−6) and V4 (PD = 32.316, p = 1.745e−7) pools. But there was no significant difference for the mixed pool (PD= −20.558, p = 0.183). Similar results were observed using Steady state responses (Fig.7L, right; Fig.S6D, bottom, & E–G), except that the empirical information never entirely converged on the decorrelated information.

It may be that the increase in the mixed V1/V4 information was due to its larger pool size, compared to the V1 or V4 pools, and not due to any change in its noise correlations. This was tested by plotting the information of the mixed pool with 12 units against the sum of information from separate V1 and V4 pools of size 6, each (Initial, Fig 7J; Steady state, Fig.S6H). The value of 6 was chosen as this was the V1 pool size limit in one of the experiments. The average paired difference in information (Fig 7K) revealed that the mixed pool had extra information, over a range of τ. This difference increased with τ; hence information increased in the mixed pool as its covariance structure changed.

A model V4 output neuron can attenuate differential noise

Could information-limiting correlations be removed between V1 and V4? Some pairs of V1 and V4 units had opposite disparity preferences and positive noise correlations (*e.g.* Fig.3D–F), satisfying one possible criterion for cancelling the correlated noise without reducing the sensory signal (Abbott and Dayan, 1999, Averbeck et al., 2006). Thus, we built a simple model of a V4 output neuron (Figure 8), of a type that might project to downstream areas. The model had two afferent neurons (Fig 8A), a V1 projection neuron and a local V4 unit; this detail could be varied without a loss of generality.

Correlated noise was removed by subtracting the weighted V1 responses from the input V4 responses, both integrated in a Δt ms window, to produce the V4 output responses. Each V1 and V4 pair in our dataset was used to predict the responses of a model V4 output neuron, the disparity selectivity of which was then measured. We hypothesized that the V1/V4 input pairs with positive noise correlation ($r_{noise} > 0$) and negative signal correlation ($r_{signal} < 0$) would yield V4 output neurons with better disparity selectivity (Fig 8B, V4 input variance, dashed, output V4 variance, grey).

We tested our hypothesis by plotting the average difference of the V4 output DMI and V4 input DMI (ΔDMI) against the signal correlation (r_{signal}) of the input pair, using responses from the Initial phase of firing and a 160ms integration window (Fig 8C). Values above zero indicate model V4 output neurons with more selectivity for disparity. When empirical data were used (solid), the V4 output neuron could improve its selectivity by cancelling noise or by pooling signals. To see the improvement due to cancellation, we recalculated the models' responses after shuffling the trials so that the noise correlation of the afferent inputs was broken but their disparity tuning was preserved (thick dotted). Although the ΔDMI of both

data sets was positive, there was an increase owing to cancelled noise when signal correlation of the input pairs was negative (r_{signal} < −0.5, mean difference = 1.246e−2bits/s, paired t-test p = 3.298e−5; r_{signal} > +0.5, 7.107e−3, p = 0.070).

A limitation of DMI is that the analysis procedure assigns the data into discrete bins, potentially underestimating the disparity selectivity. We verified the DMI results by measuring the disparity selectivity of each V4 unit, using the area under the receiver operating characteristic curve (AROC), which does not require binning. The difference in AROC between the V4 model output and input (Fig 8D) confirmed the result above (r_{signal} < −0.5, 2.961e−3, p = 4.974e−112; r_{signal} > +0.5, −7.894e−4, p = 3.269e−10).

As differential correlations seemed to be reduced by slower, V1 to V4 correlations (Fig 7), we asked if the V4 output neuron would improve with a wider integration window. Thus, we evaluated the Initial V4 output responses over a set of Δt values. For each Δt, the empirical V4 output DMI (Fig 8E) or AROC (Fig 8F) was subtracted by the V4 output selectivity measure after shuffling the inputs to break correlations, but not the tuning. When the input pair had r_{signal} < −0.5, there was a non-significant tendency for ΔDMI to increase with Δt (Δt = 20ms vs 160ms, mean paired difference = 6.676e−3, p = 0.072); however, the ΔAROC measure demonstrated a significant increase in selectivity (1.540e−3 , p = 9.490e−25).

When repeated for Steady state responses (Figure S7), we found that V4 output neuron DMI did not improve. The AROC measure, however, detected a weak improvement (Fig S7B) when the input pair had opposite disparity preferences (*none* vs *keep tuning*, mean paired difference = 4.9181e−4, paired t-test p = 3.1995e-26). No improvement in selectivity was found with increasing Δt (Δt = 100ms vs 800ms, mean paired difference = −3.902e−4, p = 7.986e−16).

Discussion

Cortical areas V1 and V4 contain noise correlations that could limit the information they convey about stereo depth. This noise limits information chiefly at short time scales as the animal is discriminating depth. Unexpectedly, noise correlations between the two areas were found to have a different structure at longer time scales, which could reflect the brain's capacity to limit harmful correlations.

Our results lead us to propose a general framework for thinking about the effect of noise correlations on neuronal processing. First, within-area noise correlations that potentially limit sensory discrimination are a by-product of local information-processing (Harris and Mrsic-Flogel, 2013, Kanitscheider et al., 2015c). Second, we suggest that the signalling between areas acts to remove the harmful consequences of these local, intrinsic correlations. The removal process may use at least two strategies: one is the connectivity between cortical areas (*e.g.* Figure 8) and the other is dynamic processing that adapts to the ongoing statistics of the neural signals. An implication of this framework is that one function of keeping diverse neocortical areas may be to contain the impact of differential correlations. Thus, the perception of stereo depth – and possibly other stimuli – could be driven by signals that are reliably passed between areas to guide accurate motor behaviour.

Detecting differential correlations

Empirical evidence of differential correlations has rarely been reported, due in part to the prediction that they are so small (Moreno-Bote et al., 2014) that thousands of neurons and thousands of trials are needed to detect them (Kanitscheider et al., 2015c, Kohn et al., 2016).

But the magnitude of noise correlations is dynamic, and changes with a number of factors (Cohen and Kohn, 2011), including state of consciousness (Ecker et al., 2014), stimulus presentation (Churchland et al., 2010), perceptual grouping (Cohen and Newsome, 2008, Poort and Roelfsema, 2009, Bondy et al., 2018, Wasmuht et al., 2019), attention (Cohen and Maunsell, 2009, Mitchell et al., 2009, Ruff and Cohen, 2014, Verhoef and Maunsell, 2017), and perceptual learning (Ni et al., 2018, Sanayei et al., 2018). Our experiment brought many of these global factors under control by using awake subjects that performed a demanding perceptual task.

Another advance we have made in detecting differential correlations is in how they are measured. Spike count correlation (r_{SC}) is often used despite the fact that correlations vary over temporal scales (Bair et al., 2001, Kohn and Smith, 2005, Smith and Kohn, 2008, Smith and Sommer, 2013, Denfield et al., 2018). If differential correlations are rapid, then they could be masked by the slower correlations that are captured by r_{SC}. By using the r_{CCG} measure introduced by Bair et al. (2001), we found evidence of differential correlations operating at tens of milliseconds. This is much shorter than the width of the analysis windows that have typically been used to measure r_{SC} (Cohen and Kohn, 2011).

Another complication is that the time scale of noise correlation is dynamic, and can change within an area due to stimulation or behavioural state. This can be seen by comparing our results with previous studies of V1 (Kohn and Smith, 2005, Smith and Kohn, 2008) and V4 (Smith and Sommer, 2013) that used r_{CCG}. They found noise correlations that saturated around 100ms during stimulation, in both areas; but the subjects were either anaesthetized (Kohn and Smith, 2005, Smith and Kohn, 2008) or passively fixating (Smith and Sommer, 2013).

We saw a similar r_{CCG} profile in the spontaneous responses of V1; but with stimulation, r_{CCG} saturated at 20ms. In V4, saturation was also near 20 ms – before and during stimulation. r_{CCG} also saturates at tens of ms in MT when monkeys perform a perceptual task (Bair et al., 2001, Wasmuht et al., 2019). Thus, the rapid, differential correlations that we observed may have only been detectible during a perceptual task.

Recording depth and topography

We used 1mm electrodes. So our recordings were likely from the supragranular layers near the layer 3-4 border, from a mixture of neurons that projected nearby or to other cortical areas (Harris and Mrsic-Flogel, 2013). The correlations may have carried a stronger differential component if many neurons projected locally. This stems from the theory that differential correlations build up between neurons with intersecting noisy afferents (Kanitscheider et al., 2015c). Given the functional heterogeneity across layers in V1 and V4 (Buffalo et al., 2011, Nandy et al., 2017, Pettine et al., 2019, Denfield et al., 2018), cross-laminar recordings will be required to resolve this issue.

Some of the signal from V1 will have been relayed to V4 by other areas; but there are direct axon projections from V1 to V4 and from V4 to V1 (Ungerleider et al., 2008). Although we did co-activate the V1 and V4 neurons with a single stimulus, the measured V1 RFs were not perfectly coincident with those of V4. Hence, some of the excitatory drive for the V4 neurons presumably came from other, unmeasured V1 neurons. Importantly, we saw qualitatively similar V4 correlations in a set of control experiments with the RDS centred on the aggregate V4 RF (see Figs.S1–S6).

In V1 and V4, noise correlations span a range of cortical (*i.e.* retinotopic) distances when neurons respond to the same stimulus (Smith and Kohn, 2008, Poort and Roelfsema, 2009, Smith and Sommer, 2013, Cohen and Maunsell, 2009) – which we observed (mean r_{CCG}(20ms); V1 pairs < 1.44mm apart = 4.987e−2; V1 > 1.44mm = 1.805e−2; V4 < 1.44mm = 3.416e−2; V4 > 1.44mm = 2.279e−2; all t-test $p \rightarrow 0$; median distance = 1.44mm). Moreover, noise correlations in both areas increase when neurons respond to a common stimulus amongst an array of stimuli (Poort and Roelfsema, 2009, Ruff and Cohen, 2014). Thus, it is likely that the correlations of perceptually grouped neurons vary systematically across the cortex. We suggest that by stimulating V1 and V4 with the same object that we increased the likelihood of the visual system having grouped them into a perceptual pool.

Mechanisms of noise removal

In a sense, our results are puzzling. Neuronal correlations in V1 and V4 show positive relationships between the neuron pairs with similar stimulus tuning and tuning curve slope (differential correlations). This kind of pattern is reported for motion in V5/MT and orientation in V1. With simple feedforward connections, one ought to observe the same relationship between V1 and V4. However, the general pattern of V1/V4 correlations is inverted (Fig.5E&F, Fig.6I).

Noise correlations like that can be removed linearly without losing information (Abbott and Dayan, 1999, Averbeck et al., 2006). Thus, our simple model (Fig.8) suggests that the combined action of V1 and V4 attenuates the local correlations through linear operations. However, differential correlations cannot be linearly removed by later processing (Moreno-Bote et al., 2014, Pitkow et al., 2015, Kanitscheider et al., 2015c). One way of resolving this discrepancy rests on our result that the covariance structure between areas is not static, but evolves across temporal scales. Notably, the V1/V4 pattern emerged at slower time scales than the rapid differential correlations within each area.

Related results come from MT (Cohen and Newsome, 2008) and V4 (Ruff and Cohen, 2014) in which positive noise correlations strengthened between neurons with opposite tuning (but see Bondy et al., 2018). In V1, positive correlations that increase the information in linearly-pooled responses can appear between neurons that respond to different stimuli (Poort and Roelfsema, 2009). Lastly, to change the V1 to V4 correlation structure in a trial would require the kind of fast and highly specific adaptations that are observed between them in attention tasks (Grothe et al., 2012, Rohenkohl et al., 2018).

Adaptable covariance may imply an underlying, non-linear computation. But what could be biologically plausible? A spike-count variance code could overcome correlations if correlated inputs were high-pass filtered in a spatial domain (*e.g.* across neurons) followed by a quadratic non-linearity to compute variance (Shamir and Sompolinsky, 2004). Another potential choice is divisive normalisation (Carandini and Heeger, 2012), which can model the structure of correlations between areas V1 and MT (Ruff and Cohen, 2016). Divisive normalisation might decorrelate responses in two ways (see Eq. 3.1 in Tripp, 2012). First, spike count variance can be saturated by the correlated variance of other neurons, attenuating correlations in the normalised response. Second, the normalised responses of correlated units can move in opposite directions to their raw values, decorrelating them. One mechanism need not exclude others, and the brain might profit from the parallel action of several computations.

Sources of differential correlations

Present analysis suggests both bottom-up (Kanitscheider et al., 2015c) and top-down (Ecker et al., 2016) origins of differential correlations. In the former, stimulus noise is sent through diverging projections to downstream areas; differential correlations then build up between neurons with common noisy inputs (Kanitscheider et al., 2015c), especially if sub-optimal computations are used (Beck et al., 2012). In the latter, noise in the selection of an attended

feature value can lead to a gain profile that shifts position relative to the stimulus preferences of a population of neurons – inducing differential correlations (Ecker et al., 2016).

Our results invite a compromise. Both sources are present, but take effect at different moments. The initial burst of firing is a bottom-up event in which information is most reduced by rapid noise correlations, which are unlikely to have a top-down origin (Denfield et al., 2018). In our data, only the Initial responses could reach the decorrelated information level, at a temporal scale of 100ms. The Steady state information did increase at wider temporal scales, yet failed to hit the decorrelated level. One interpretation is that rapid differential correlations in the Initial response had a bottom-up source, followed by slow differential correlations with a top-down source that could not be locally attenuated. This fits well with the common reports of top-down or contextual effects that follow the initial response of visual neurons to a stimulus (Roelfsema et al., 1998, Pack and Born, 2001, Cohen and Maunsell, 2009, Nienborg and Cumming, 2009, Poort et al., 2012, Herrero et al., 2013, Smith et al., 2015).

Neural pooling and perception

The statistical relationship between a pool of sensory neurons and perceptual behaviour is connected to the noise correlations within the pool (Shadlen et al., 1996, Nienborg et al., 2012, Parker, 2013, Haefner et al., 2013, Pitkow et al., 2015). That relationship may evolve over time, with a sensory pool driving a brief, bottom-up effect on behavioural choice, followed afterwards by a top-down reflection of choice bias (Smith et al., 2012). These effects are difficult to separate in many experiments, which employ long stimulus viewing times (Britten et al., 1996, Dodd et al., 2001, Nienborg and Cumming, 2009, Poort et al., 2012, Pitkow et al., 2015, Yu and Gu, 2018).

However, a subject may base perceptual decisions on only a few hundred ms of viewing in natural conditions (Cohen and Newsome, 2009). As such, studies using mixtures of reaction time tasks and short-duration stimuli often observe that the transient response is accounted for by a small pool of neurons (Ghose and Harrison, 2009, Harrison et al., 2013, Weiner and Ghose, 2015, Krause and Ghose, 2018) with a bottom-up link to behaviour (Price and Born, 2010, Smith et al., 2011, Weiner and Ghose, 2014). Our results suggest that the transient response is a high-fidelity signal that could be decoded from a small pool of neurons. Later, top-down signals may arrive that impose slow and widespread differential correlations in a larger pool, thus injecting choice correlations into neurons that are never read out (Cohen and Newsome, 2009, Nienborg and Cumming, 2009, Nienborg and Cumming, 2010, Pitkow et al., 2015, Yu and Gu, 2018).

ACKNOWLEDGMENTS

The authors wish to thank Andy Emberton and the staff at The University of Oxford's Biomedical Services, plus Daniel Heaton and the DPAG workshop for expert technical support. Thanks go to Ruben Coen-Cagli for help in understanding the computation of bias-corrected linear Fisher information, and also to Alexandre Pouget, Bruce Cumming and Robert Dayan for their insightful conversations. This work was funded by MRC grant MR/K014382/1.

AUTHOR CONTRIBUTIONS

AP was the principal investigator and designed the study. JS collected and analysed the data. JS and AP wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

METHODS

Animals

We recorded from two adult male Rhesus macaque monkeys (*Macacca mulatta*) acquired from the UK Public Health England breeding colony at Porton Down through the Centre for Macaques (M135 and M138). Each animal was implanted with a titanium headpost (Gray Matter Research) and a CerePort pedestal (Blackrock Microsystems) in separate, aspectic surgical procedures under general anaesthesia. Each CerePort was connected to a pair of 64-channel Utah arrays that were implanted subdurally in the upper layers of the exposed cerebral cortex. The exact locations for surgical implantation were determined by a preliminary magnetic resonance imaging (MRI) scan of the animal's head to reveal bone and soft tissue structures.

Animals were trained to fixate on a target and then to perform an odd-one-out task using saccadic eye movement responses. All procedures were done in compliance with Home Office (UK) regulations. The key results were qualitatively similar in both animals, and so we have pooled the data across animals.

Visual stimulus

Visual stimuli were presented with a Wheatstone stereoscope comprising of a pair of CRT monitors (Eizo Flexscan F78 or ViewSonic P225f) at 84cm from the subject's eyes. Cold mirrors were used to reflect images into the eyes so that the infrared light from an eye tracking camera (Hi Speed Primate, SensorMotoric Instruments GmBH), set behind the mirrors, could pass through. Monitors had a resolution of 1600 x 1200 pixels, a refresh rate of 85Hz, and subtended 26.7° × 20.1° of the subject's visual field. Each pixel was 1.6e−2° wide. The VGA signal was split and duplicated, so that the left-eye monitor received three copies of the red channel for its RGB input, while the right-eye monitor received three copies of the green channel. When viewed on a conventional single colour monitor by human observers, the stereo images were visible as red/green anaglyph images; but viewed through the stereoscope, the images appeared as grayscale images presented to each eye of the subjects.

Quadro K2200 (NVIDIA) video cards in multi-core Intel processor computers running Ubuntu 14.04 (Linux kernel 4.4.0 lowlatency) were used to drive the monitors. Stimuli were programmed using PsychToolbox (3.0.14) running in Matlab R2015b (The Mathworks). All stimuli were presented on a mid-gray background, using appropriate OpenGL alpha-blending (GL_SRC_ALPHA + GL_ONE_MINUS_SRC_ALPHA) to obtain sub-pixel resolution. A 4 × 4cm square was shown in the upper-left corner of each monitor during trials. This was recorded with a photodiode (PIN-25DP, OSI Optoelectronics) by the electrophysiology system, in order to synchronise high-precision frame time stamps from PsychToolbox with the neural activity. The square was blocked from the subject's view. A full white, 0.3° diameter circular dot was used as the gaze fixation target.

Random dot stereograms (RDS) were used to present binocular disparities for the behavioural, odd-one-out task and to stimulate recorded neurons. Each RDS comprised of a circular patch of 0.16° diameter circular dots. Dot positions were randomly sampled on each frame; the frame rate matched the refresh rate of 85Hz. Four RDS were presented, but dot

positions were sampled independently for each RDS. Half of the dots were black, and half were white. Dots occluded each other in a random order. Each RDS had two regions, a 6° diameter circular centre and a 1° wide annular surround. The dot density in both regions was 25% (area occupied by dots / total area), assuming no dot overlap; in other words, the number of dots per RDS was 25% of the RDS area divided by the area of a single dot, rounded up. The RDS centre varied in binocular disparity such that a flat circular plane of dots was shifted convergently (negative disparities, toward subject) or divergently (positive, away). Surround dots were always on the fixation plane at zero disparity *i.e.* the surface of the monitors' screens. The surround functioned to mask monocular cues when non-zero disparities were presented in the centre, and to help the subject maintain vergence. When the RDS centre was at a non-zero disparity, each of its monocular images had a horizontal shift that overlapped the RDS surround on one side while leaving a gap on the other. Surround dots in the overlap region were discarded, while uncorrelated dots were used to fill the gaps. Thus, the resulting monocular images carried no trace of the disparity of the RDS centre, so that the subjects could only perform the task using the stereoscopic information.

Task

Subjects performed an odd-one-out task (Figure 1A) that was divided into four epochs. At the start of each trial, a fixation point was presented (Fig.1Ai). The subject then had to maintain its gaze at the fixation point for 0.5 second to initiate the trial (Fig.1Aii). Following this, four RDS were presented for 1 second (M135) or 2 seconds (M138), one in each quadrant of the visual field (Fig.1Aiii). Subjects had to maintain their binocular gaze within a 0.75° radius of the centre of the fixation dot to both initiate the trial and progress through to the end of the presentation phase. During the presentation phase, all four RDS had the same baseline disparity in their central region. Baseline disparities were randomly selected on each trial with equal probability from a set of 13 values: 0°, ±0.01°, ±0.02°, ±0.05°, ±0.10°, ±0.20°, ±0.50° (similar to Shiozaki et al., 2012). At the end of the presentation period, one of the four RDS had a disparity step change in its central region, gaining an additional 0°, ±0.05°, or ±0.10° (Fig.1Aiv). Both the location and value of the step change were randomly selected on each trial with equal probability. Thus, the RDS that changed became the odd-one-out (*i.e.* the popout target), which the subject was required to detect by making a saccade towards its location. With four possible target locations, the probability of correctly selecting the target by chance was 0.25 (Figure 1B). A juice reward was given for correct performance. No reward was given for incorrect saccades, or if the subject responded prior to the step change. After the training or testing session, animals had ad libitum access to water in the home enclosure for a limited period of time every day. Animals were weighed each day and assessed.

Recording

The animals' binocular gaze positions were sampled at 500Hz and recorded alongside the neural responses. In cortex, each subject was implanted with two Utah arrays (Blackrock Microsystems). Both arrays had 64 platinum electrodes, each 1mm in length, arranged in an 8 × 8 grid with 400µm spacing. It is unclear exactly which layer of the cortex that each electrode tip had entered, due to the curvature of the cortical surface relative to the flat array, and from slight deviations of the arrays from the plane tangential to the cortex. Nonetheless, given the electrode lengths, the presumption is that the recordings

reported here are primarily from the upper, supragranular layers of the cortex. One array was implanted in cortical area V1 and the other in V4. A CerePlex E headstage (Blackrock Microsystems) was attached to the CerePort implant. The headstage applied 0.3Hz 1st order high-pass and 7.5kHz 3rd order low-pass anti-aliasing Butterworth filters to the analogue signals before 16-bit digitizing all 128 channels at 30kHz. Digitized signals were transmitted to the Cerebus (128 channel electrophysiology system, Blackrock Microsystems) using fibre optics, thus minimizing noise at the analogue stage. All computers, monitors, and the reward pump were located in a separate control room while the subject, headstage, stimulus monitors, digital hub, and CRTV camera were within a test room that had Faraday shielding. A 2nd order 250 to 5000Hz bandpass digital Butterworth filter was applied online to each data channel before applying a threshold of −6.5 times the root-mean-squared (RMS) of the estimated noise level to detect action potentials (Rizk and Wolf, 2009); the thresholds were newly estimated on each experiment for each electrode. The −6.5×RMS threshold was the default setting for the Cerebus, which worked well in practice to detect multi-unit activity without saturating the acquisition system and losing data.

The simultaneous V1/V4 data in this study came from 6 recording sessions with M135 and 4 sessions with M138, yielding 8279 trials from M135 (1380 per-session mean [60.7 standard deviation]) and 6073 from M138 (1518 [58.0]). Approximately similar number of trials with each baseline disparity were obtained. We made no assumptions about the identity of each unit, treating each day's recording separately. Here, the term 'unit' simply means a cluster of spike waveforms from one electrode (see below). This yielded a total of 63 V1 (10.5 [3.2]) and 111 V4 (18.5 [3.5]) units from M135, and 132 V1 (33.0 [3.5]) and 102 V4 (25.5 [3.7]) units from M138. Thus, we retrieved 315 V1 (52.5 [32.6]), 975 V4 (162.5 [63.5]), and 1206 inter-cortical (201.0 [91.6]) neuron pairs from M135 and 2104 V1 (526.0 [108.4]), 1251 V4 (312.8 [90.6]), and 3392 inter-cortical (848.0 [183.7]) pairs from M138. All pairwise analyses compared units from separate electrodes (Figures 3 – 6 & 8).

In order to stimulate both V1 and V4 receptive fields (RF) with the central region of the lower-right RDS, the array of all four RDS stimuli was sometimes shifted upwards and to the left, in which case they had an asymmetrical arrangement relative to the fixation point (Figure 1C). A second set of V4 control data was obtained in which the lower-right RDS was centred on part of the aggregate receptive field of the V4 Utah array (Figure S1); in this case, the four RDS locations were symmetrically placed around the fixation point. In these sessions, 11 baseline disparity values were used, including: 0°, ±0.05°, ±0.1°, ±0.2°, ±0.3°, ±0.4°. The task was slightly different in that the odd-one-out did not have a disparity step change. Rather, it reversed the binocular correlation of its central dots from +100% correlated to -100%. That is, dots had the same contrast in the left and right-eye images during the presentation period. After the presentation period, odd-one-out dots in the left image were paired with dots of opposite contrast in the right image. However, we only analysed responses to the +100% correlated dots, prior to the popout phase of the trial. Thus, despite this change of task, the control data were obtained under similar sensory stimulation conditions to the test data. For M135, the diameter of each RDS centre was 4°. For M138, it was 1°. 7 control sessions were recorded from M135, yielding a total of 5257 trials (751.0 [135.7]), 413 V4 control units (59.0 [8.4]) and 1536 pairs (219.4 [88.4]). 3

control sessions were recorded from M138, yielding 2082 trials (694.0 [240.8]), 266 V4 control units (88. 7 [15.3]) and 1101 pairs (367.0 [181.7]).

Analysis

The data were analysed using programs that were written and executed in Matlab R2015b (The Mathworks).

Pre-processing

In both animals, we found incidences of cross-talk between a sub-set of electrodes. This was apparent as a fraction of tightly synchronous waveform times between certain pairs of electrodes, defined as waveforms occurring within ±1.3e−4 second *i.e.* ±4 samples at 30KHz of each other. Synchronous waveforms often had highly similar, stereotypical shapes. Testing with a Blackrock digital neural signal simulator indicated that all components downstream of the CerePort functioned properly. Electrodes were discarded from the analysis if they had at least 6% synchronous spikes with any other electrode and the median Pearson correlation of the synchronous waveforms was 0.5 or more. An electrode was also discarded from analysis, if there was no significant change in the multiunit spiking rate following onset of the RDS stimulus. This was defined by counting spikes in two 500ms wide analysis windows, one from −500ms to the appearance of the RDS and the other from 50ms to 550ms. A paired-sample t-test was used to compare the firing rates before and after RDS appearance at a significance level of 5%. When a significant difference was found then the electrode's spikes were sorted for further analysis. On average, 9 V1 electrodes (minimum 5, maximum 14) and 14 V4 (min 12, max 17) electrodes from M135 were used, while 27 V1 (min 22, max 31) and 20 V4 (min 19, max 23) electrodes were used from M138 per session.

Spikes detected online with the Cerebus were sorted using our implementation of the clustering algorithm by Fee et al. (1996) and the heuristic choices of UltraMegaSort2000 (Hill et al., 2011; https://neurophysics.ucsd.edu). An additional step was taken after aligning the spike waveform peaks but prior to the principal component analysis. Each waveform was weighted by a Gaussian shaped window with a standard deviation of 1ms that was centred on the waveform peak; area under the Gaussian that overlapped the waveform was normalized to preserve the overall magnitude of the spike. This emphasised the peak region of the waveforms during sorting. Automated spike sorting was then verified and corrected manually. We based our analyses on both single units and multiunit clusters and use the term 'unit' to refer to either, as in other work with Utah array electrodes (Ruff and Cohen, 2014).

The disparity tuning curve was found for each unit, using spike rates from an analysis window of 44 to 1044ms (M135) or 44 to 2044ms (M138) following onset of the RDS stimuli. Only units with significant disparity tuning (one-way ANOVA, $p < 0.01$) were kept for further analysis. Additionally, any unit with a tuning curve peak of less than 5 spikes/second was also discarded.

Single unit analysis

The selectivity of a unit for binocular disparity was quantified in two ways. Disparity discrimination index (DDI; Prince et al., 2002) was used to validate disparity mutual information (DMI). DDI was defined as:

$$\frac{N_{max} - N_{min}}{(N_{max} - N_{min}) + 2R_{error}} \quad (1)$$

where N_{max} and N_{min} were the largest and smallest responses of the unit's disparity tuning curve, and R_{error} was the square root of the residual variance around the means over the whole tuning curve. DDI was computed using the square root of the firing rates (as in Prince et al., 2002).

DMI was computed non-parametrically (Wallisch et al., 2014) by first estimating the joint probability distribution of the disparity and spike count using 13 evenly spaced bins in both marginal dimensions to get a two-dimensional histogram (one marginal bin per disparity value, 11 marginal bins used for V4 control data). Disparities were ranked ordinally before binning. Joint probabilities were empirically estimated by:

$$P_{xy}(i,j) = \frac{1}{M}\sum_{m=1}^{M} J_m(i,j) \quad (2)$$

where M was the total number of trials and $J_m(i,j)$ had a binary value:

$$J_m(i,j) = \begin{cases} 1, \text{ordinal disparity was in bin } i \text{ and spike count was in bin } j \text{ on trial } m \\ 0, \text{otherwise} \end{cases} \quad (3)$$

The marginal probabilities were computed in a similar manner. For ordinal disparity values (P_x) and spike counts (P_y), they were:

$$P_x(i) = \frac{1}{M}\sum_{m=1}^{M} x_m(i) \quad (4)$$

$$P_y(j) = \frac{1}{M}\sum_{m=1}^{M} y_m(j) \quad (5)$$

where $x_m(i)$ and $y_m(j)$ had a value of 1 if the data fell into the respective bin on trial *m*, and 0 otherwise. The raw empirical mutual information was estimated by:

$$I_{raw} = \sum_{i=1}^{B}\sum_{j=1}^{B} P_{xy}(i,j) log_2 \frac{P_{xy}(i,j)}{P_x(i)P_y(j)} \quad (6)$$

where B was the number of bins along the marginal. However, this information estimate will be biased – without further correction. The correction was found by applying equation (6) to randomly shuffled copies of the data. Using 30 repeats (Hatsopoulos et al., 1998), a shuffle-correction term was estimated as:

$$I_{shuffle} = \frac{1}{30}\sum_{s=1}^{30} i_{shuffle}(s) \tag{7}$$

where $i_{shuffle}(s)$ was the output of equation (6) following the s^{th} random shuffling of the data. The final, corrected mutual information between the disparity and spike count was taken as:

$$I = \frac{I_{raw} - I_{shuffle}}{\Delta t} \tag{8}$$

where Δt was the width of the spike-counting analysis window, in seconds, providing a value in bits/second.

A time series analysis of firing rates and DMI was conducted on spike trains that were convolved with a causal exponential kernel (20ms time constant). DMI was measured independently for each time point using Δt = 0.02s (Eq.8) to estimate the instantaneous rate of information. The firing rate and DMI time series were used to define three analysis windows for the pairwise and population analyses. The Spontaneous window started 500ms before RDS onset, ended 44ms after RDS onset, and covered the period of gaze fixation prior to the units' first responses to the RDS. The Initial window was located to cover the first transient response to the RDS, which was measured from 44ms to 200ms after the RDS appeared. A final Steady state window ran from 200ms to 1044ms for M135 and from 200ms to 2044ms in M138, covering responses to the RDS up to but not including those to the disparity step-change.

Paired signal and noise correlations

In the analysis of pairs of neurons, we focused on three key metrics, as defined by Bair et al. (2001). First was the degree to which the firing rate of two units was driven by the same stimulus values – the pair's signal correlation. Signal correlation was defined as the Pearson correlation of the pair's disparity tuning curve values:

$$r_{signal} = \frac{E\left[N_j^d N_k^d\right] - E\left[N_j^d\right]E\left[N_k^d\right]}{\sigma_{N_j^d}\sigma_{N_k^d}} \tag{9}$$

where N_j^d and N_k^d were the average spike counts of units j and k in response to a disparity d, E was the expected value of the average spike counts across disparity values, and σ was the standard deviation of the average spike counts across disparity values.

The second metric was the degree to which the residual variances in the tuning curves were shared by the two units, the pair's noise correlation. If the stimulus value is fixed over a set of trials, then noise correlation can be quantified as spike-count correlation (r_{SC}) using:

$$r_{SC} = \frac{E[n_j n_k] - E[n_j]E[n_k]}{\sigma_{n_j}\sigma_{n_k}} \tag{10}$$

where n_j and n_k were the spike counts on a single trial from units j and k, E was the expected value across trials, and σ was the standard deviation across trials. In order to pool across trials with different disparity values we first z-score normalized the spike counts of each unit for each sub-set of trials, grouped by disparity; r_{SC} was then computed as:

$$r_{SC} = E[z_j z_k] \tag{11}$$

where z_j and z_k were the z-scored firing rates on a single trial for units j and k.

A limitation of r_{SC} is that it only measures noise correlation at the time scale of the spike-counting window, which is typically on the order of seconds. However, noise correlations may occur at a range of time scales, from near synchrony at the scale of milliseconds to slower modulations over seconds. Noise correlation can be estimated at different time scales in the same analysis window (Spontaneous, Initial, or Steady state) using r_{CCG}, which is the normalized integral of the cross-correlation between paired spike-trains. By integrating from 0 to ±Δt lags under the cross-correlation, r_{CCG} measures rapid noise correlations when Δt is small and includes slower noise correlations when Δt is large. r_{CCG} converges upon r_{SC} when Δt approaches the width of the analysis window. To compute r_{CCG}, the spike train of unit k on trial m was first binned into a series of ms-wide bins; the spike train become a binary series of zeros and ones such that:

$$x_k^m(t) = \begin{cases} 1, \text{if unit } k \text{ fired an action potential during ms } t \text{ on trial } m \\ 0, \text{otherwise} \end{cases} \tag{12}$$

If there were M trials and the analysis window was T ms long, then $1 \leq m \leq M$ and $1 \leq t \leq T$. The post-stimulus time histogram (PSTH) of unit k during time t was:

$$P_k(t) = \frac{1}{M}\sum_{m=1}^{M} x_k^m(t) \tag{13}$$

The spike-train correlation function at τ ms of lag is:

$$\mathrm{C_{jk}}(\tau) = \frac{1}{M}\sum_{m=1}^{M}\sum_{t=1}^{T} x_j^m(t)x_k^m(t+\tau) \tag{14}$$

If $j = k$ then $C_{jk}(\tau)$ is the auto-correlation, otherwise if $j \neq k$ then it is the cross-correlation between units j and k. Likewise, the auto- or cross-correlation of the PSTH was:

(15)

$$S_{jk}(\tau) = \sum_{t=1}^{T} P_j(t) P_k(t+\tau)$$

In fact, $S_{jk}(\tau)$ serves as a shift-predictor for $C_{jk}(\tau)$. Thus, the integral of the shift-corrected spike-train auto- or cross-correlation was:

$$A_{jk}(\theta) = \sum_{\tau=-\theta}^{+\theta} \left[C_{jk}(\tau) - S_{jk}(\tau) \right] \tag{16}$$

where the area $A_{jk}(\theta)$ was summed symmetrically from $-\theta$ to $+\theta$ ms of lag.

Following this, r_{CCG} was defined as:

$$r_{CCG}(\tau) = \frac{A_{jk}(\tau)}{\sqrt{A_{jj}(\tau) A_{kk}(\tau)}} \tag{17}$$

Using (17), equation (10) for r_{SC} can be recast as:

$$r_{SC} = \frac{A_{jk}(T)}{\sqrt{A_{jj}(T) A_{kk}(T)}} \tag{18}$$

where $1 \leq \tau \leq T$. When τ is small, then noise correlations on a short time scale are measured; whereas, when τ is large, slower noise correlations are estimated. We found r_{CCG} separately for each disparity value. The mean r_{CCG} across disparities was found by taking the weighted average, weighing by the number of trials per disparity condition.

A third metric of paired activity is the normalized spike-train cross-correlogram (*CCG*), which shows the relative timing of correlated spikes between the pair of units, defined here as:

$$CCG(\tau) = \frac{C_{jk}(\tau)}{\Theta(\tau)\sqrt{\lambda_j \lambda_k}} \tag{19}$$

$$\Theta(\tau) = T - |\tau| \tag{20}$$

where λ_k was the average firing rate (spikes per second) of unit k, and $-T \leq \tau \leq +T$. *CCG* was also found separately for each disparity value, and the mean across disparities was weighted by the number of trials. For presentation purposes, *CCGs* were convolved with the symmetrical 5ms kernel [0.05, 0.25, 0.40, 0.25, 0.05] (Kohn and Smith, 2005).

First derivatives of the tuning curves

In order to estimate the first derivative of the disparity tuning curves, we found the best-fitting, one-dimensional Gabor function for the tuning curves of each unit, in both the

Initial and Steady state analysis windows (Prince et al., 2002, Tanabe et al., 2004). The Gabor was expressed as:

$$G(d) = c_0 + c_1 e^{-\left(\frac{(d-c_2)^2}{2c_3^2}\right)} cos(2\pi c_4(d - c_2) + c_5) \quad (21)$$

where $G(d)$ was the value of the Gabor function at a disparity value of d, and c_i were the parameters that were optimized for each tuning curve. The baseline offset (c_0) was constrained between zero and the maximum observed response on any trial. The amplitude of the Gaussian envelope (c_1) was constrained between zero and twice the difference between the maximum observed response on any trial and the minimum. The horizontal offset of the Gaussian envelope (c_2) was constrained between the minimum and maximum disparity values that were tested. The width of the Gaussian envelope (c_3) was constrained between 0.1 and the total range of tested disparity values. The disparity frequency of the cosine (c_4) was estimated from the power spectrum of the tuning curve by first linearly interpolating (function *interp1*, MATLAB, The Mathworks) the z-scored tuning curve from the minimum to maximum tested disparity, at 0.001° steps. The discrete Fourier transform at frequency f ($y(f)$; function *fft*, MATLAB, The Mathworks) was estimated from the interpolated tuning curve, and the power spectrum at f was computed using:

$$p(f) = \frac{|y(f)|^2}{D} \quad (22)$$

where D was the number of interpolated points. We found the frequency between 0 and 50 cycles per degree that maximized the power spectrum (f_{max}) and then constrained the disparity frequency to be $c_4 = f_{max} \pm 0.1 f_{max}$. The phase of the cosine (c_5) was constrained to be within ±π. The parameters were fit using the trust-region-reflective algorithm to minimize the square of the residual error (function *lsqcurvefit*, MATLAB, The Mathworks). After optimizing the parameters for each tuning curve, we estimated the first derivative of the tuning curves at each tested disparity with:

$$\mathrm{G}' = \frac{\mathrm{G(d + \Delta d) - G(d - \Delta d)}}{2\Delta \mathrm{d}} \quad (23)$$

using a value of 0.5e−3° for Δd. The Gabors better captured the responses to smaller disparities near zero, hence the 9 smallest disparities were used to estimate f'×f'.

The results in Figure 6 were verified using a second method to estimate the first derivative of the tuning curves, in which the raw tuning curves were linearly interpolated at $d \pm \Delta d/2$ and the difference was divided by Δd, for each tested disparity value of d.

Population response analysis

In order to estimate the ability of simultaneous neuronal responses to distinguish one disparity value from another, we calculated the bias-corrected linear Fisher information as defined by Kanitscheider et al. (2015b); their code formed the basis of our own Matlab implementation (Kanitscheider et al., 2015a). Linear Fisher information quantifies the ability

of a linear decoder to distinguish two sets of simultaneous population spike counts that arose in response to two different stimulus values. Hence, this was calculated by:

$$\hat{I}_{bc}(\tau) = \frac{\Delta N^{tr}}{\Delta d}\ \Sigma(\tau)^{-1}\frac{\Delta N}{\Delta d}\left(\frac{2\mathrm{M_{min}} - \mathrm{U} - 3}{2\mathrm{M_{min}} - 2}\right) - \frac{2\mathrm{U}}{\mathrm{M_{min}}\Delta \mathrm{d}^2} \tag{24}$$

$$\Delta d = d_a - d_b \tag{25}$$

$$\Delta N = N^{d_a} - N^{d_b} \tag{27}$$

$$\Sigma(\tau) = \left(\Sigma^{d_a}(\tau) + \Sigma^{d_b}(\tau)\right)/2 \tag{26}$$

where d_a and d_b were two of the tested disparity values such that $d_a \neq d_b$, N^d was a column vector of the average spike counts for U units in response to disparity value d, $\frac{\Delta N}{\Delta d}$ was also a column vector in which each element of the column vector ΔN was divided by Δd, tr was the transposition operator that turned $\frac{\Delta N}{\Delta d}$ into a row vector, $\Sigma^d(\tau)$ was the $U \times U$ covariance matrix for disparity value d at a timescale of τ ms, X^{-1} was the inverse of matrix X (Moore-Penrose psudoinverse, function *pinv*, Matlab, The Mathworks), and $\hat{I}_{bc}(\tau)$ was a scalar value with a unit of degrees^{-2}. We randomly discarded trials so that each disparity value had the same number of trials as that which had the minimum (M_{min}), as assumed in the derivation by Kanitscheider et al. (2015b). In practice, few trials were discarded, with the remainder being well above the minimum required for the covariance matrix to be invertible, which is $(U + 2)/2$. For units j and k, we defined elements of $\Sigma^d(\tau)$ to be:

$$\Sigma^d_{j,k}(\tau) = r^d_{CCG}(\tau)\sigma^d_j\sigma^d_k \tag{28}$$

where $r^d_{CCG}(\tau)$, σ^d_j and σ^d_k was the r_{CCG} and spike-count standard deviations of units j and k on trials with disparity d. Thus, we could estimate the effect of noise correlations at different time scales on the linear Fisher information of the population response. It was possible to estimate the variance of the linear Fisher information for different time scales of noise correlation as:

$$\mathrm{Var}\hat{I}_{bc}(\tau) = \frac{2\hat{I}_{bc}(\tau)^2}{2M_{min} - U - 5}\left(1 + \frac{4(2\mathrm{M_{min}} - 3)}{\mathrm{M_{min}}\hat{\mathrm{I}}_{\mathrm{bc}}(\tau)\Delta \mathrm{d}^2} + \frac{4\mathrm{U}(2\mathrm{M_{min}} - 3)}{\left(\mathrm{M_{min}}\hat{\mathrm{I}}_{\mathrm{bc}}(\tau)\Delta \mathrm{d}^2\right)^2}\right) \tag{29}$$

A theoretical measure of linear Fisher information is obtained when all units in the population are decorrelated with each other such that the variance of their spike counts is statistically independent. This is often achieved numerically by shuffling the data. However, it is possible to estimate this value analytically without shuffling – using only the diagonal of the covariance matrix (called Ibc,shuffle in Kanitscheider et al., 2015b):

$$\hat{I}_{bc,decorr} = \hat{I}_{bc,decorr}(\tau) = \frac{2U}{M_{min}\Delta d^2}\sum_{i=1}^{U}\frac{(\Delta N_i/\Delta d)^2}{\Sigma_{i,i}(\tau)} - \frac{M_{min} - 2}{M_{min} - 1} \tag{30}$$

Eq. 30 emphasizes the fact that with this calculation the decorrelated linear Fisher information is the same for all time scales of noise correlation, because the auto-correlated value of r_{CCG} is always 1 (see eq. 18 when $j = k$) and σ_j^d does not vary with τ (eqs. 27 & 28).

Any difference between the empirical and decorrelated information could signal the presence of information-limiting noise correlations in the population. In order to make a direct comparison between all combinations of disparity values, and for different compositions of neuronal population, we computed the normalized difference as:

$$\Delta\hat{I}_{bc}(\tau) = \frac{\hat{I}_{bc,decorr}(\tau) - \hat{I}_{bc}(\tau)}{\sqrt{Var\hat{I}_{bc}(\tau)}} \tag{31}$$

in a unit of standard deviations of $\hat{I}_{bc}(\tau)$ (see Figs.7L & S6D). Outliers were discarded when the estimated variance resulted in a normalized difference of 40 or greater. With 13 tested disparities, there were 78 unique pairs of disparity values and therefore 78 values of $\hat{I}_{bc}(\tau)$, $Var\hat{I}_{bc}(\tau)$, $\hat{I}_{bc,decorr}$, and $\Delta\hat{I}_{bc}(\tau)$ were each computed for all neuronal pool sizes, compositions, and correlation time scales. Neuronal pools comprised of V1 units only, V4 units only, or a mixture of V1 and V4 units. Over 10 sessions, this yielded 780 information values in each case.

To see if linear Fisher information saturated with increasing population size, we first ordered a given set of units by the DMI of their Initial responses, from largest to smallest. Units were added to the neuronal pool one at a time in this order, and the Fisher information was computed at each step; that is, the most disparity-selective unit was added first, and the least selective unit was added last. The result was an information curve that was a function of pool size. Initial DMI was always used to order the units because their Initial responses were the most selective for disparity (see Figs.2 and S2). This ordering imposed a saturating shape on even the decorrelated information curves. However, we could then examine whether the saturation worsened as a result of noise correlations. Information curves were found for the Initial and Steady state responses at time scales of 20 and 150ms (Initial) or 20, 150, 400, and 800ms (Steady state) and for decorrelated units (both). In order to average across experimental sessions and combinations of disparity values, it was necessary to normalize the information values in a way that would preserve the relative shape of information curves at different time scales. To do this, the curves were first grouped by experiment and disparity values. Normalization was then achieved within a group by:

$$I_{norm}(\tau, s_{pop}) = \frac{I_{raw}(\tau, s_{pop}) - \beta_{grp}}{\alpha_{grp}} \tag{32}$$

where $I_{raw}(\tau, s_{pop})$ was the value of $\hat{I}_{bc}(\tau)$ or $\hat{I}_{bc,decorr}$ at population size s_{pop} units. β_{grp} was the minimum value of $I_{raw}(\tau, 1)$ across values of τ. α_{grp} was then taken to be the maximum value of $I_{raw}(\tau, s_{pop}) - \beta_{grp}$ across values of τ and s_{pop}. Results obtained using the unit order above were confirmed in a control analysis that randomly permuted the unit

order before computing the information curves; the control was repeated 100 times for a total N = 78,000.

A model V4 output neuron can attenuate differential noise

We modelled a theoretical V4 output neuron in order to test whether a simple mechanism could attenuate differential noise that was present upstream in V1 and V4. Thus, the model unit had only two afferent inputs, a V1 unit and a V4 unit (see Fig.8A). There were two chief parameters of the model, the signal correlation of the V1/V4 input pair and the width of a window used to integrate their responses. V4 output responses were predicted by subtracting a weighted version of the V1 input from the V4 input. V4 output neurons were modelled for all recorded V1/V4 pairs as afferent inputs, and for window widths of 20, 40, 80, and 160ms for Initial responses or 100, 200, 400, and 800ms for Steady state responses. The V1 input weight (*w* in Fig.8A) should reflect the magnitude of noise shared by the afferent pair that was not explained by the random sampling of disparity values from trial to trial. To measure it, we used a simple generalized linear model (GLM) to predict the spike count of the V4 input from its disparity tuning curve and the spiking of the V1 input, using the ln link function:

$$\ln n_{V4}^{m} = c_0^{glm} + c_1^{glm} G_{V4}(d_m) + c_2^{glm} n_{V1}^{m} \quad (33)$$

n_{V4}^{m} was the spike count of the V4 input unit on trial m, G_{V4} was the Gabor function that best fit the V4 input's tuning curve, G_{V4} was evaluated at the disparity value d_m used on trial m, and n_{V1}^{m} was the spike count of the V1 input unit on trial m. The GLM was used to fit (function *glmfit*, MATLAB, The Mathworks) three parameters: the baseline spike count (c_0^{glm}), the amount of disparity signal (c_1^{glm}), and the amount of shared variance between the afferent pair (c_2^{glm}). The Gabor term was important, as it helped to push c_2^{glm} towards accounting for shared residual variance rather than shared signal.

Once obtained, we predicted the V4 output responses with:

$$n_{V4,OUT}^{m} = n_{V4}^{m} - exp\left(c_2^{glm} n_{V1}^{m}\right) + \beta \quad (34)$$

where $n_{V4,OUT}^{m}$ was the V4 output spike count on trial m following. Thus, the V4 output was the V4 input minus the weighted V1 input. β was a constant term that restored the baseline value of the modelled V4 output disparity tuning curve to match that of the V4 input. Any negative values were capped at zero. This produced V4 output spike counts with a similar dynamic range to the V4 input. The exponential function was used to transform the scaled V1 response because the logarithmic link function was used in the GLM.

To avoid over-fitting, we applied equations (33) and (34) to different halves of the trials in a two-fold, cross-validation procedure. Trials were divided into even and odd-numbered groups. Thus, the training and test sets had a similar size and composition, balancing the need to train the GLM against the need to accurately compute DMI and AROC (see below). One group of trials was used to train the model with Eq. (33) and then the other group was used to evaluate it with Eq. (34). The process was repeated a second time after reversing

the roles of each group. DMI and AROC were used to evaluate the disparity selectivity of the V4 output responses for each test set, yielding two estimates of each metric per condition that were averaged together.

The training, testing and evaluation procedure was applied to the empirical data, in which the inter-cortical noise correlations were intact (Fig.8C&D, *none*). However, it was necessary to determine the degree to which any change in disparity selectivity of the V4 output responses was due to the combined disparity information signalled by the afferent pair. Therefore, the procedure was applied a second time to data that were shuffled. This had the effect of preserving the disparity selectivity of both the V1 and V4 input units, while destroying their noise correlations (*keep tuning*). A third evaluation was done using shuffled data in which the V4 input preserved its disparity selectivity, but both the noise correlations and V1 input disparity selectivity were abolished by shuffling across all available trials rather than those with the same disparity (*all trials*).

We evaluated summation of the V1/V4 inputs with integration windows of different widths (Figure 8E&F). This was accomplished by sub-dividing each analysis window (Initial or Steady state) into a set of smaller, contiguous windows that were each Δt ms wide. Spike counts were obtained and equations (33) and (34) were applied separately for each sub-division; that is, an independent set of GLM parameters was trained and tested for each sub-division of the analysis window. The V4 Gabor function was scaled to produce the spike count expected in a window of size Δt ms. Before computing DMI or AROC, the V4 output spike counts were summed back together across sub-divisions:

$$n^{m}_{V4,OUT} = \sum_{\omega=1}^{W} n^{m}_{V4,OUT}(\omega) \tag{35}$$

where $n^{m}_{V4,OUT}(\omega)$ was the V4 output spike count obtained in the ω^{th} sub-division of the analysis window, out of W sub-divisions with a width of Δt ms.

In computing DMI non-parametrically, it is possible that the true information was underestimated as a result of discretizing the data through binning. Therefore, we verified the DMI results using a second, continuous method based on signal detection theory (Green and Swets, 1966) that has previously been used to quantify the sensitivity of single neurons to different stimulus values (Newsome et al., 1989). We used a similar approach to estimate disparity selectivity. In this approach, one builds a receiver operating characteristic curve (ROC) by labelling one set of spike counts as true positives and another as false positives. If the true-positive spike counts were from trials with one disparity value and the false positives were from trials with another, then the area under the ROC (AROC) becomes the probability that an ideal observer could correctly classify the true-positive spike counts. This leads directly to a measure of how well the neuron can discriminate the two tested disparities. A chance value of AROC = 0.5 means that the two spike-count distributions were indistinguishable. Values of 0 or 1, however, indicate perfect separation of the distributions, with intermediate values of AROC indicating different amounts of overlap. We calculated AROC using trapezoidal integration of the true- and false-positive cumulative distribution functions (F_T and F_F, respectively):

$$AROC = \frac{1}{2}\sum_{i=1}^{V}[F_F(x_i) - F_F(x_{i-1})][F_T(x_i) + F_T(x_{i-1})] \quad (36)$$

where $x_i \in \{x_1, x_2, x_3, \ldots x_V\}$ is the ordered (ascending) set of V unique spike counts from either the true- or false-positive distribution, and $F_T(x_0) = F_F(x_0) = 0$. The absolute difference of AROC from chance is a measure of separation between two spike count distributions regardless of which was labelled as the true- or false-positives:

$$|AROC| = |AROC - 0.5| \quad (37)$$

$|AROC|$ was computed for each of the 78 unique pairs of disparity values. With 4598 V1/V4 afferent pairs, we obtained a total of 4598 DMI and 358,644 AROC predicted values for the model V4 output neurons. The difference in V4 output and V4 input disparity selectivity is shown in Figures 8C–D and S7.

Bootstrap confidence intervals

95% bootstrap confidence intervals were found using the bias corrected and accelerated percentile method (function *bootci*, MATLAB, The Mathworks) with at least 1000 bootstrap samples.

FIGURE LEGENDS

Figure 1. *Task and stimulus.* A) (i) Trials started with flashing fixation point (FP, white circle). (ii) Once fixated (dashed circle), FP flash stops. (iii) 500ms later, four random dot stereograms (RDS, white dots) appeared. (iv) After fixed duration (1s M135, 2s M138), central dots of one RDS change disparity (illustrated by red/green dots) and FP disappears. Animals made saccade onto 'odd-one-out' target RDS (arrow) for reward. Neural responses measured from phases (ii) and (iii). B) Percent correct trials (chance=25%, dashed line) and reaction time (RT) over popout disparity (M135 N = 8279, M138 N = 6073): error bars 95% binomial confidence intervals (CI) or standard error of the mean (SEM). C) Positions of V1 (yellow circle) and V4 (blue cross) receptive field (RF) centres. RDS stimuli positioned so that lower-right RDS stimulated V1 and V4 neurons (dashed lines). D) Aggregate RF of all V4 units (average firing rate, heat map). C&D) Left panels show M135, right panels show M138.

Figure 2. *V1 and V4 units were selective for binocular disparity.* A) Example disparity tuning curves from V1 (circles, yellow) and V4 (crosses, blue); symbols show average firing rates; lines show best-fitting Gabors: SEM below symbol size. B) Disparity mutual information (DMI) for V1 (yellow dot, N = 195) and V4 (blue cross, N = 213) units compared with disparity discriminability index (DDI). C) Average firing rate (mean line, SEM ribbon) across all V1 (yellow) and V4 (blue) units, over trial sequence: analysis windows cover spontaneous responses (left arrow), Initial phasic response (solid bar and lines), later Steady state responses (right arrow); dashed line marks the onset of RDS. D) As in C but for DMI.

Figure 3. *Alignments of noise correlations and disparity tuning.* A) Disparity tuning curves of pair V4/1 (black) & V4/2 (grey; M138.336.94/147), signal correlation (r_{signal} = +0.851); crosses and SEM show average firing, lines show best-fitting Gabors (V4/1 r^2 = 89%, V4/2 r^2 = 93%). Vertical dashed line at +0.1° shows trials (N = 108) examined in panels B&C. B) Raster plots and average firing rates for V4/1 (top) and V4/2 (bottom); RDS onset at t=0; for 108 trials at baseline disparity +0.1°; shading shows time window for firing rates in C. C) Dots show joint firing rates; slight jitter applied for visualization; noise correlation, r_{SC} = 0.241, p = 0.012. Tangent line (grey) shows joint tuning curves at +0.1°; arrow shows first PCA component (×40, for visibility) of joint responses. D – F) the same as A – C but for an example V1/V4 pair (M138.336.57/87; V1 black in D, r^2 = 98%; V4 grey in D, r^2 = 94%) with negative signal correlation (r_{signal} = −0.982). Trials with baseline disparity of 0° (D, black dashed, N = 116) yield significant positive noise correlation (F, r_{SC} = 0.203, p = 0.029).

Figure 4. r_{CCG} *measures noise correlation at different time scales.* A) Average r_{CCG} as function of temporal scale τ for all V1/V1 pairs' (N = 2419) spontaneous (dashed black), Initial (solid black), and Steady state (magenta) responses (see Fig 2C); ribbons show SEM. B) As A, but has average V1/V1 cross-correlograms (CCG). C) Scatter plot of Steady state r_{CCG} at τ =100ms versus spike-count correlation (r_{SC}). Marginal distributions show each metric separately. D – F) the same as A – C but for V4/V4 pairs (N = 2226). G – I) the same as A – C but for V1/V4 pairs (N = 4598).

Figure 5. *Noise correlations varied with disparity tuning, time scale, and time.* A) Noise correlation (r_{CCG}) for V1/V1 as function of signal correlation (r_{signal}), grouped by time scale (τ) of r_{CCG} (20ms, dashed; 150ms, solid) and analysis window (Initial, black; Steady state,

magenta); each bin has 10% of the data; error bars are SEM. B) Average V1/V1 r_{CCG} over time scale τ, grouped by signal correlation (positive r_{signal}, dashed; negative r_{signal}, solid) and analysis window, as A; ribbon shows SEM. C&D) same as A&B but for V4/V4 pairs. E&F) same as A&B but for V1/V4 pairs.

Figure 6. *Empirical evidence of differential noise correlations.* A) $r_{CCG}(\tau)$ from V1 at τ = 41ms plotted over product of first derivatives of disparity tuning curves (f'×f'), for Initial response phase, N = 21,771 cases (2419 V1/V1 pairs at 9 disparities −0.1° to +0.1°). Best-fitting, least-squares regression line (orange) $r_{CCG}(\tau)$ = k + m had positive slope (m): k = 3.452e−2; m = 1.244e−7. B) As A but for Steady state: τ = 95ms; k = 1.594e−2; m = 2.812e−7. C) Spearman correlation coefficient of r_{CCG} and f'×f' for Initial (black) and Steady state (magenta) responses, at all $\tau \leq$ 400ms. Filled circles show $p < 0.05$. Arrows show τ used in A (black) and B (magenta). D – F) Same as A – C but for V4/V4 pairs and using τ = 21ms for D and τ = 32ms for E, with N = 20,034. Positive regression slopes in D: k = 3.020e−2; m = 3.039e−6, and E: k = 3.330e−2; m = 3.891e−6. G – I) Same as A – C but for V1/V4 pairs and using τ = 88ms for G and τ = 109ms for H, with N = 41,382. Negative regression slope in G: k = 7.782e−3; m = −4.711e−7, and H: k = 6.354e−3; m = −2.804e−7.

Figure 7. *Differential noise correlations are attenuated at long time scales.* A – C) Average normalized bias-corrected linear Fisher information, over neuronal pool size. Information shown for rapid (dotted), slow (solid), and no (decorrelated, dashed) noise correlations, N = 780 (78 disparity pairs × 10 experiments). D – F) Average Fisher information r_{CCG} τ. Information is empirical (thick solid) or decorrelated (dashed). G – I) Decorrelated vs. empirical information at τ = 150ms. A – I, pools had V1 (A,D,G), V4 (B,E,H), or V1 and V4 (C,F,I) units. J) Information of mixed V1/V4 pool (size 12) vs. summation over V1 and V4 pools (size 6) for rapid (left) and slow (right) correlations. K) Average paired differences of data in J (left) over τ. Right, same but for Steady state. * significant paired t-tests for adjacent data (Initial α = 0.05, Steady state α = 0.05/3). L) Average normalised difference of decorrelated and empirical information over τ (Initial, left; Steady state, right) for V1 (yellow), V4 (blue), or V1 and V4 (green). Units added to pools by decreasing order of Initial DMI. A – J, Initial data. Max. pool sizes in D – I & L. 95% bootstrap CI.

Figure 8. *A model V4 output neuron attenuates differential noise.* A) Diagram of model V4 output neuron. Diamonds, afferent neurons. Rasters, afferent spikes in Δt ms window. Weighted V1 count ($w\Sigma_{V1,IN}$) subtracted from V4 count ($\Sigma_{V4,IN}$) gets V4 output count ($\Sigma_{V4,OUT}$). B) Disparity tuning prediction, afferent V1/V4 $r_{signal} < 0$ and $r_{noise} > 0$. V4 input variance ($\Sigma_{V4,IN}$, dotted). Model V4 output variance ($\Sigma_{V4,OUT}$, grey). C) Average change in DMI over r_{signal} of afferent V1/V4 pairs, Δt = 160ms. Trials shuffled to break V1/V4 correlations but not disparity selectivity (*keep tuning*, thick dotted), to break both (*all trials*, thin dashed), or not shuffled (*none*, solid). D) As C but using AROC disparity selectivity measure, N = 358,644 (78 disparity pairs × V1/V4 pairs). C&D, each bin has 10% of data. E) Average difference in DMI over Δt, V4 model output DMI with no shuffling minus shuffled DMI with preserved tuning. Afferent V1/V4 $r_{signal} < -0.5$ (black, N = 794) and $r_{signal} > +0.5$ (grey, N = 835). F) As E, but for AROC. N = 4598 model V4 output neurons, per recorded V1/V4 pair. Initial data shown. SEM error bars.

REFERENCES

ABBOTT, L. F. & DAYAN, P. 1999. The effect of correlated variability on the accuracy of a population code. *Neural Comput,* 11**,** 91-101.

AVERBECK, B. B., LATHAM, P. E. & POUGET, A. 2006. Neural correlations, population coding and computation. *Nat Rev Neurosci,* 7**,** 358-66.

BAIR, W., ZOHARY, E. & NEWSOME, W. T. 2001. Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci,* 21**,** 1676-97.

BARLOW, H. B., KAUSHAL, T. P., HAWKEN, M. J. & PARKER, A. J. 1987. Human contrast discrimination and the threshold of cortical neurons. *Journal of the Optical Society of America a-Optics Image Science and Vision,* 4**,** 2366-2371.

BECK, JEFFREY M., MA, WEI J., PITKOW, X., LATHAM, PETER E. & POUGET, A. 2012. Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron,* 74**,** 30-39.

BONDY, A. G., HAEFNER, R. M. & CUMMING, B. G. 2018. Feedback determines the structure of correlated variability in primary visual cortex. *Nat Neurosci,* 21**,** 598-606.

BRITTEN, K. H., NEWSOME, W. T., SHADLEN, M. N., CELEBRINI, S. & MOVSHON, J. A. 1996. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience,* 13**,** 87-100.

BUFFALO, E. A., FRIES, P., LANDMAN, R., BUSCHMAN, T. J. & DESIMONE, R. 2011. Laminar differences in gamma and alpha coherence in the ventral stream. *Proceedings of the National Academy of Sciences,* 108**,** 11262.

CARANDINI, M. & HEEGER, D. J. 2012. Normalization as a canonical neural computation. *Nature Reviews Neuroscience,* 13**,** 51.

CHURCHLAND, M. M., YU, B. M., CUNNINGHAM, J. P., SUGRUE, L. P., COHEN, M. R., CORRADO, G. S., NEWSOME, W. T., CLARK, A. M., HOSSEINI, P., SCOTT, B. B., BRADLEY, D. C., SMITH, M. A., KOHN, A., MOVSHON, J. A., ARMSTRONG, K. M., MOORE, T., CHANG, S. W., SNYDER, L. H., LISBERGER, S. G., PRIEBE, N. J., FINN, I. M., FERSTER, D., RYU, S. I., SANTHANAM, G., SAHANI, M. & SHENOY, K. V. 2010. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci,* 13**,** 369-78.

COHEN, M. R. & KOHN, A. 2011. Measuring and interpreting neuronal correlations. *Nature Neuroscience,* 14**,** 811.

COHEN, M. R. & MAUNSELL, J. H. 2009. Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci,* 12**,** 1594-600.

COHEN, M. R. & NEWSOME, W. T. 2008. Context-Dependent Changes in Functional Circuitry in Visual Area MT. *Neuron,* 60**,** 162-173.

COHEN, M. R. & NEWSOME, W. T. 2009. Estimates of the Contribution of Single Neurons to Perception Depend on Timescale and Noise Correlation. *The Journal of Neuroscience,* 29**,** 6635.

DENFIELD, G. H., ECKER, A. S., SHINN, T. J., BETHGE, M. & TOLIAS, A. S. 2018. Attentional fluctuations induce shared variability in macaque primary visual cortex. *Nat Commun,* 9**,** 2654.

DODD, J. V., KRUG, K., CUMMING, B. G. & PARKER, A. J. 2001. Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. *J Neurosci,* 21**,** 4809-21.

ECKER, A. S., BERENS, P., COTTON, R. J., SUBRAMANIYAN, M., DENFIELD, G. H., CADWELL, C. R., SMIRNAKIS, S. M., BETHGE, M. & TOLIAS, A. S. 2014. State dependence of noise correlations in macaque primary visual cortex. *Neuron,* 82**,** 235-48.

ECKER, A. S., DENFIELD, G. H., BETHGE, M. & TOLIAS, A. S. 2016. On the Structure of Neuronal Population Activity under Fluctuations in Attentional State. *J Neurosci,* 36**,** 1775-89.

FEE, M. S., MITRA, P. P. & KLEINFELD, D. 1996. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J Neurosci Methods,* 69**,** 175-88.

GHOSE, G. M. & HARRISON, I. T. 2009. Temporal Precision of Neuronal Information in a Rapid Perceptual Judgment. *Journal of Neurophysiology,* 101**,** 1480-1493.

GREEN, D. M. & SWETS, J. A. 1966. *Signal detection theory and psychophysics,* New York London, Wiley.

GROTHE, I., NEITZEL, S. D., MANDON, S. & KREITER, A. K. 2012. Switching Neuronal Inputs by Differential Modulations of Gamma-Band Phase-Coherence. *The Journal of Neuroscience,* 32**,** 16172.

HAEFNER, R. M., GERWINN, S., MACKE, J. H. & BETHGE, M. 2013. Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat Neurosci,* 16**,** 235-42.

HARRIS, K. D. & MRSIC-FLOGEL, T. D. 2013. Cortical connectivity and sensory coding. *Nature,* 503**,** 51.

HARRISON, I. T., WEINER, K. F. & GHOSE, G. M. 2013. Inattention blindness to motion in middle temporal area. *J Neurosci,* 33**,** 8396-410.

HATSOPOULOS, N. G., OJAKANGAS, C. L., PANINSKI, L. & DONOGHUE, J. P. 1998. Information about movement direction obtained from synchronous activity of motor cortical neurons. *Proceedings of the National Academy of Sciences of the United States of America,* 95**,** 15706-15711.

HECHT, S., SHLAER, S. & PIRENNE*, M. H. 1941. ENERGY AT THE THRESHOLD OF VISION. *Science,* 93**,** 585.

HERRERO, JOSE L., GIESELMANN, MARC A., SANAYEI, M. & THIELE, A. 2013. Attention-Induced Variance and Noise Correlation Reduction in Macaque V1 Is Mediated by NMDA Receptors. *Neuron,* 78**,** 729-739.

HILL, D. N., MEHTA, S. B. & KLEINFELD, D. 2011. Quality Metrics to Accompany Spike Sorting of Extracellular Signals. *The Journal of Neuroscience,* 31**,** 8699.

JOHNSON, K. O., DARIAN-SMITH, I. & LAMOTTE, C. 1973. Peripheral neural determinants of temperature discrimination in man: a correlative study of responses to cooling skin. *Journal of Neurophysiology,* 36**,** 347-370.

KANITSCHEIDER, I., COEN-CAGLI, R., KOHN, A. & POUGET, A. 2015a. MatLab tools for estimating linear Fisher information from population data along with synthetic data and recorded spike count responses from neurons in macaque primary visual cortex to grating images with different orientations and white noise. *CRCNS.org. http://dx.doi.org/10.6080/K0PK0D3B*.

KANITSCHEIDER, I., COEN-CAGLI, R., KOHN, A. & POUGET, A. 2015b. Measuring Fisher information accurately in correlated neural populations. *PLoS Comput Biol,* 11**,** e1004218.

KANITSCHEIDER, I., COEN-CAGLI, R. & POUGET, A. 2015c. Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences of the United States of America,* 112**,** E6973-E6982.

KOHN, A., COEN-CAGLI, R., KANITSCHEIDER, I. & POUGET, A. 2016. Correlations and Neuronal Population Information. *Annu Rev Neurosci,* 39**,** 237-56.

KOHN, A. & SMITH, M. A. 2005. Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci,* 25**,** 3661-73.

KRAUSE, B. M. & GHOSE, G. M. 2018. Micro-pools of reliable area MT neurons explain rapid motion detection. *J Neurophysiol,* 120**,** 2396-2409.

KRUG, K., CURNOW, T. L. & PARKER, A. J. 2016. Defining the V5/MT neuronal pool for perceptual decisions in a visual stereo-motion task. *Philos Trans R Soc Lond B Biol Sci,* 371**,** 20150260.

LEVICK, W. R., THIBOS, L. N., COHN, T. E., CATANZARO, D. & BARLOW, H. B. 1983. Performance of cat retinal ganglion cells at low light levels. *The Journal of General Physiology,* 82**,** 405.

MITCHELL, J. F., SUNDBERG, K. A. & REYNOLDS, J. H. 2009. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron,* 63**,** 879-88.

MORENO-BOTE, R., BECK, J., KANITSCHEIDER, I., PITKOW, X., LATHAM, P. & POUGET, A. 2014. Information-limiting correlations. *Nat Neurosci,* 17**,** 1410-7.

NANDY, A. S., NASSI, J. J. & REYNOLDS, J. H. 2017. Laminar Organization of Attentional Modulation in Macaque Visual Area V4. *Neuron,* 93**,** 235-246.

NEWSOME, W. T., BRITTEN, K. H. & MOVSHON, J. A. 1989. Neuronal correlates of a perceptual decision. *Nature,* 341**,** 52-4.

NI, A. M., RUFF, D. A., ALBERTS, J. J., SYMMONDS, J. & COHEN, M. R. 2018. Learning and attention reveal a general relationship between population activity and behavior. *Science,* 359**,** 463-465.

NIENBORG, H., COHEN, M. R. & CUMMING, B. G. 2012. Decision-Related Activity in Sensory Neurons: Correlations Among Neurons and with Behavior. *Annual Review of Neuroscience, Vol 35,* 35**,** 463-483.

NIENBORG, H. & CUMMING, B. 2010. Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality? *Current Opinion in Neurobiology,* 20**,** 376-381.

NIENBORG, H. & CUMMING, B. G. 2009. Decision-related activity in sensory neurons reflects more than a neuron’s causal effect. *Nature,* 459**,** 89.

PACK, C. C. & BORN, R. T. 2001. Temporal dynamics of a neural solution to the aperture problem in visual area MT of macaque brain. *Nature,* 409**,** 1040-1042.

PARKER, A. & HAWKEN, M. 1985. Capabilities of monkey cortical cells in spatial-resolution tasks. *Journal of the Optical Society of America A,* 2**,** 1101-1114.

PARKER, A. J. 2013. A micro-pool model for decision-related signals in visual cortical areas. *Front Comput Neurosci,* 7**,** 115.

PARKER, A. J. & NEWSOME, W. T. 1998. SENSE AND THE SINGLE NEURON: Probing the Physiology of Perception. *Annual Review of Neuroscience,* 21**,** 227-277.

PETTINE, W. W., STEINMETZ, N. A. & MOORE, T. 2019. Laminar segregation of sensory coding and behavioral readout in macaque V4. *Proceedings of the National Academy of Sciences,* 116**,** 14749.

PITKOW, X., LIU, S., ANGELAKI, D. E., DEANGELIS, G. C. & POUGET, A. 2015. How Can Single Sensory Neurons Predict Behavior? *Neuron,* 87**,** 411-23.

POORT, J., RAUDIES, F., WANNIG, A., LAMME, VICTOR A. F., NEUMANN, H. & ROELFSEMA, PIETER R. 2012. The Role of Attention in Figure-Ground Segregation in Areas V1 and V4 of the Visual Cortex. *Neuron,* 75**,** 143-156.

POORT, J. & ROELFSEMA, P. R. 2009. Noise Correlations Have Little Influence on the Coding of Selective Attention in Area V1. *Cerebral Cortex,* 19**,** 543-553.

POUGET, A., DAYAN, P. & ZEMEL, R. S. 2003. INFERENCE AND COMPUTATION WITH POPULATION CODES. *Annual Review of Neuroscience,* 26**,** 381-410.

PRICE, N. S. C. & BORN, R. T. 2010. Timescales of Sensory- and Decision-Related Activity in the Middle Temporal and Medial Superior Temporal Areas. *The Journal of Neuroscience,* 30**,** 14036.

PRINCE, S. J. D., POINTON, A. D., CUMMING, B. G. & PARKER, A. J. 2002. Quantitative Analysis of the Responses of V1 Neurons to Horizontal Disparity in Dynamic Random-Dot Stereograms. *Journal of Neurophysiology,* 87**,** 191-208.

RIZK, M. & WOLF, P. D. 2009. Optimizing the automatic selection of spike detection thresholds using a multiple of the noise level. *Medical & biological engineering & computing,* 47**,** 955-966.

ROELFSEMA, P. R., LAMME, V. A. F. & SPEKREIJSE, H. 1998. Object-based attention in the primary visual cortex of the macaque monkey. *Nature,* 395**,** 376-381.

ROHENKOHL, G., BOSMAN, C. A. & FRIES, P. 2018. Gamma Synchronization between V1 and V4 Improves Behavioral Performance. *Neuron,* 100**,** 953-963 e3.

RUFF, D. A. & COHEN, M. R. 2014. Attention can either increase or decrease spike count correlations in visual cortex. *Nature Neuroscience,* 17**,** 1591-1597.

RUFF, D. A. & COHEN, M. R. 2016. Stimulus Dependence of Correlated Variability across Cortical Areas. *The Journal of Neuroscience,* 36**,** 7546.

SANAYEI, M., CHEN, X., CHICHARRO, D., DISTLER, C., PANZERI, S. & THIELE, A. 2018. Perceptual learning of fine contrast discrimination changes neuronal tuning and population coding in macaque V4. *Nat Commun,* 9**,** 4238.

SHADLEN, M. N., BRITTEN, K. H., NEWSOME, W. T. & MOVSHON, J. A. 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J Neurosci,* 16**,** 1486-510.

SHAMIR, M. & SOMPOLINSKY, H. 2004. Nonlinear Population Codes. *Neural Computation,* 16**,** 1105-1136.

SHIOZAKI, H. M., TANABE, S., DOI, T. & FUJITA, I. 2012. Neural activity in cortical area V4 underlies fine disparity discrimination. *J Neurosci,* 32**,** 3830-41.

SMITH, J. E. T., BELIVEAU, V., SCHOEN, A., REMZ, J., ZHAN, C. A. A. & COOK, E. P. 2015. Dynamics of the functional link between area MT LFPs and motion detection. *Journal of Neurophysiology,* 114**,** 80-98.

SMITH, J. E. T., MASSE, N. Y., ZHAN, C. A. & COOK, E. 2012. Linking Neural Activity to Visual Perception: Separating Sensory and Attentional Contributions. *In:* MOLOTCHNIKOFF, S. & ROUAT, J. (eds.) *Visual Cortex: Current Status and Perspectives.* IntechOpen.

SMITH, J. E. T., ZHAN, C., AN, A. & COOK, E. P. 2011. The Functional Link between Area MT Neural Fluctuations and Detection of a Brief Motion Stimulus. *The Journal of Neuroscience,* 31**,** 13458.

SMITH, M. A. & KOHN, A. 2008. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci,* 28**,** 12591-603.

SMITH, M. A. & SOMMER, M. A. 2013. Spatial and temporal scales of neuronal correlation in visual area V4. *J Neurosci,* 33**,** 5422-32.

TANABE, S., UMEDA, K. & FUJITA, I. 2004. Rejection of false matches for binocular correspondence in macaque visual cortical area V4. *J Neurosci,* 24**,** 8170-80.

TOLHURST, D. J., MOVSHON, J. A. & DEAN, A. F. 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Research,* 23**,** 775-785.

TRIPP, B. P. 2012. Decorrelation of Spiking Variability and Improved Information Transfer Through Feedforward Divisive Normalization. *Neural Computation,* 24**,** 867-894.

UNGERLEIDER, L. G., GALKIN, T. W., DESIMONE, R. & GATTASS, R. 2008. Cortical connections of area V4 in the macaque. *Cereb Cortex,* 18**,** 477-99.

VERHOEF, B.-E. & MAUNSELL, J. H. R. 2017. Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nature Neuroscience,* 20**,** 969.

WALLISCH, P., LUSIGNAN, M. E., BENAYOUN, M. D., BAKER, T. I., SETH DICKEY, A. & HATSOPOULOS, N. G. 2014. Information Theory. *MATLAB for Neuroscientists : An Introduction to Scientific Computing in MATLAB.* 2 ed. London: Elsevier Science & Technology.

WASMUHT, D. F., PARKER, A. J. & KRUG, K. 2019. Interneuronal correlations at longer time scales predict decision signals for bistable structure-from-motion perception. *Scientific Reports,* 9**,** 11449.

WEINER, K. F. & GHOSE, G. M. 2014. Rapid shape detection signals in area V4. *Frontiers in Neuroscience,* 8**,** 294.

WEINER, K. F. & GHOSE, G. M. 2015. Population coding in area V4 during rapid shape detections. *J Neurophysiol,* 113**,** 3021-34.

YU, X. & GU, Y. 2018. Probing Sensory Readout via Combined Choice-Correlation Measures and Microstimulation Perturbation. *Neuron,* 100**,** 715-727 e5.

ZYLBERBERG, J., POUGET, A., LATHAM, P. E. & SHEA-BROWN, E. 2017. Robust information propagation through noisy neural circuits. *PLoS Comput Biol,* 13**,** e1005497.

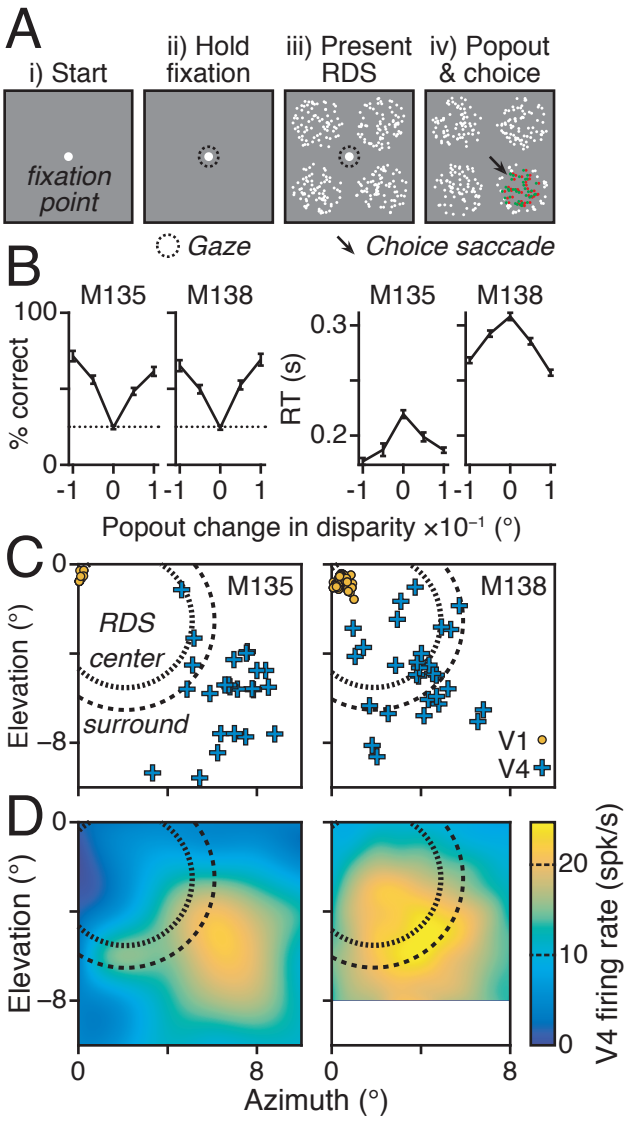

Figure 1

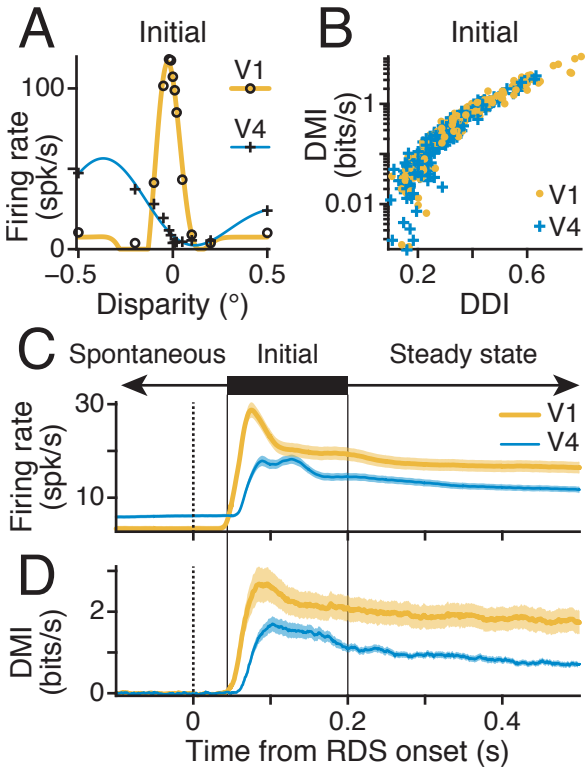

Figure 2

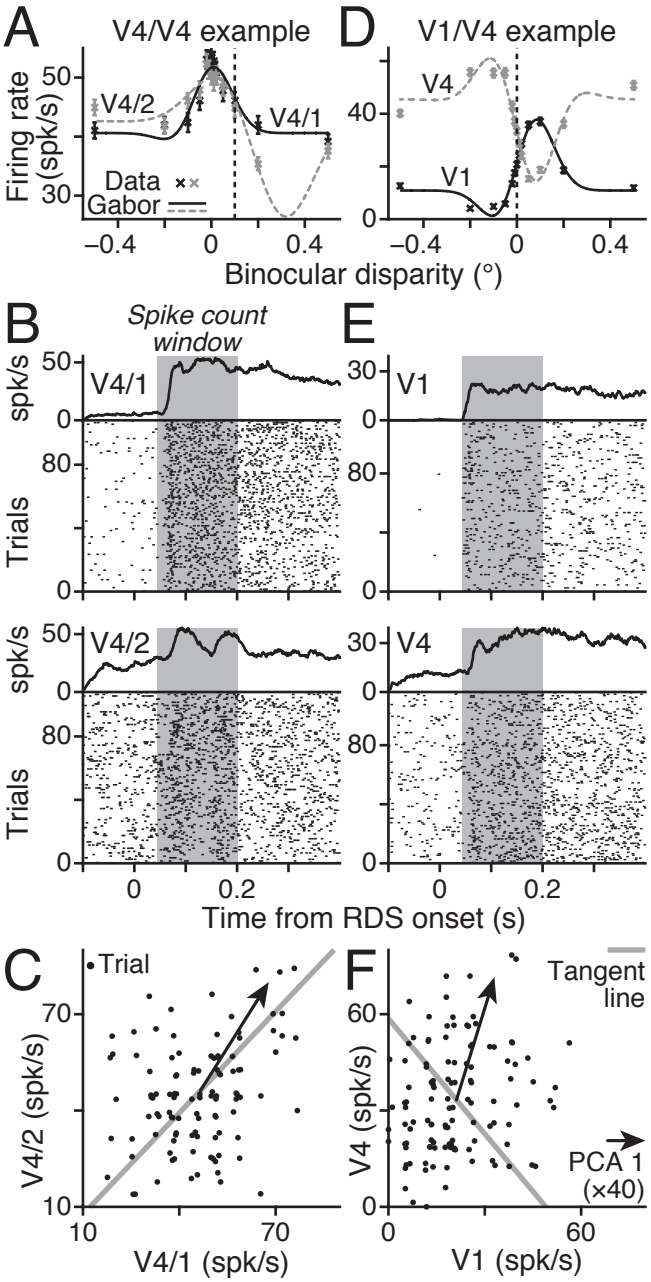

Figure 3

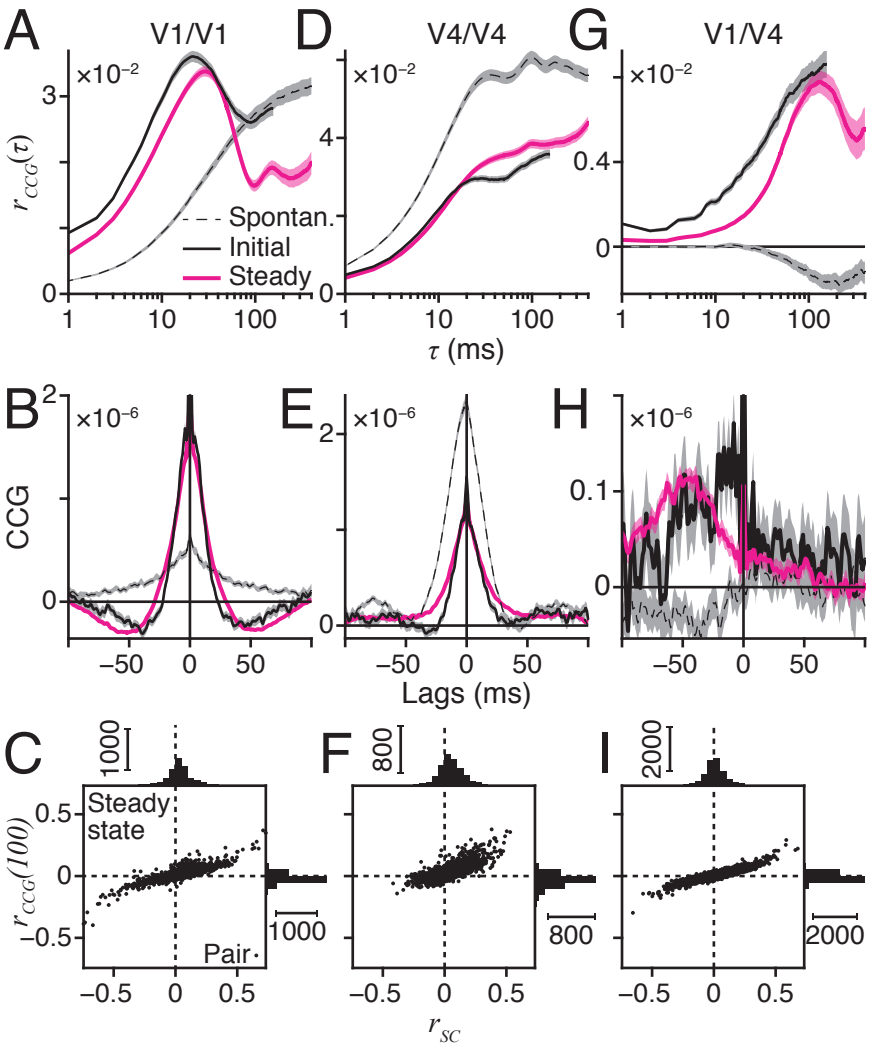

Figure 4

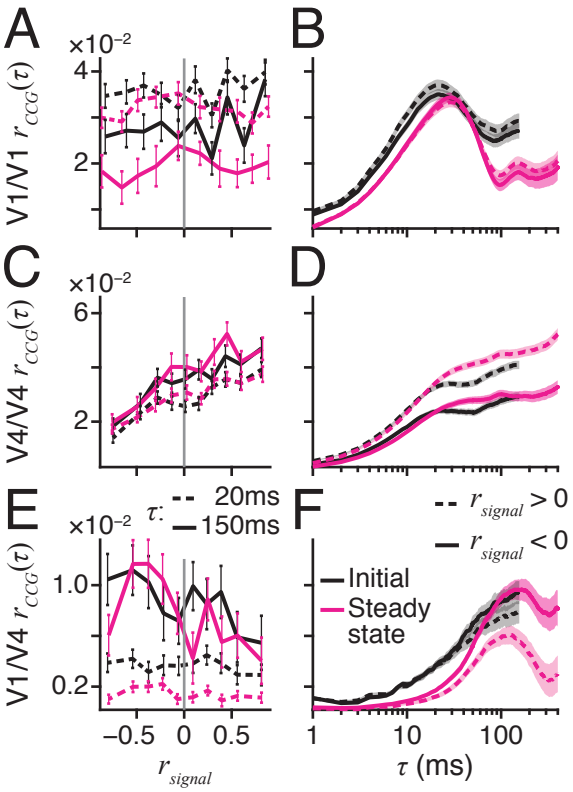

Figure 5

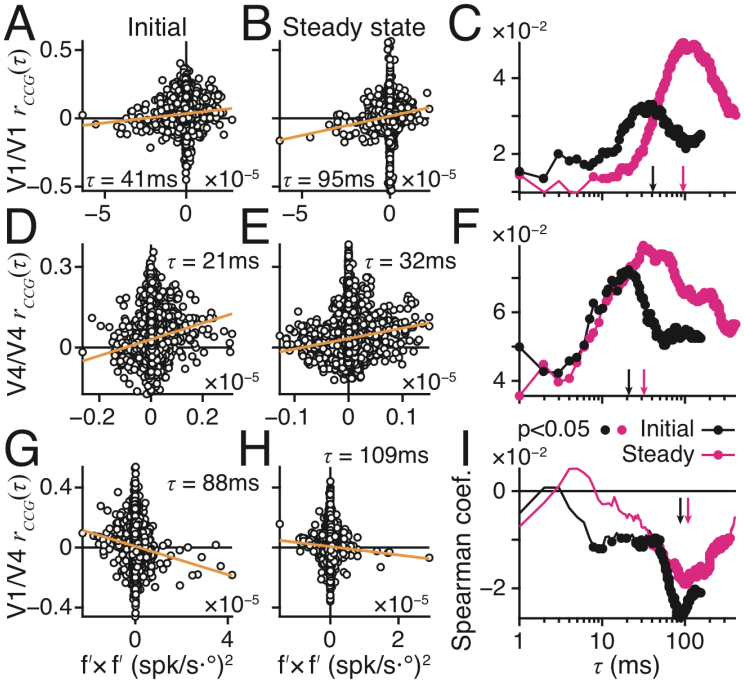

Figure 6

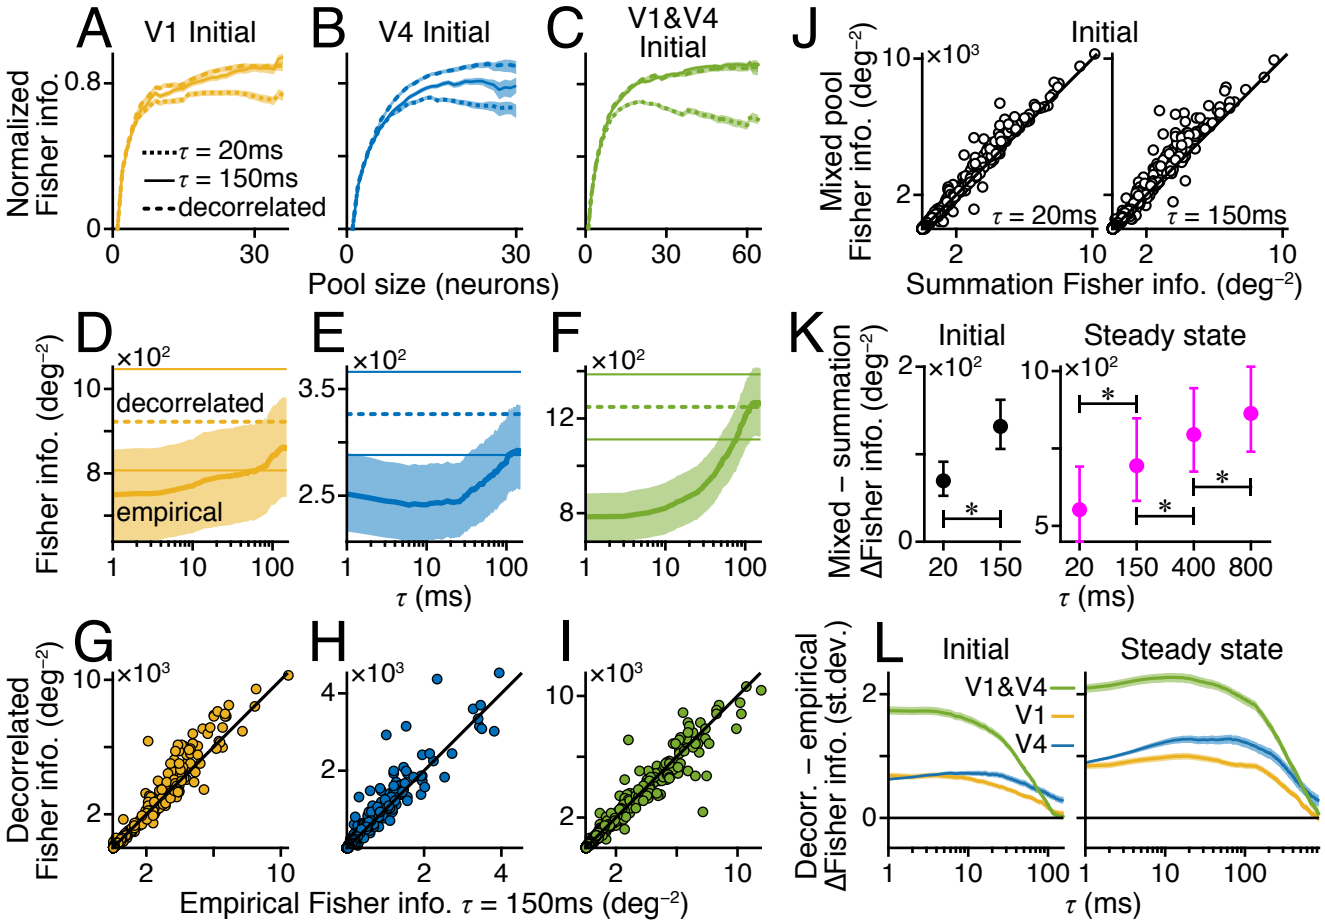

Figure 7

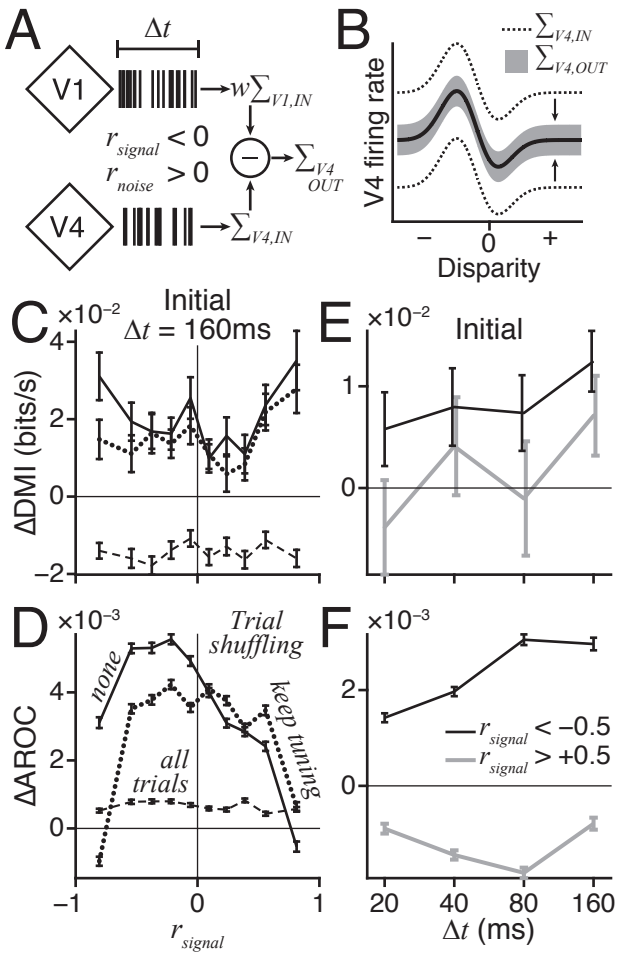

Figure 8