



Correlations enhance the behavioral readout of neural population activity in association cortex

Martina Valente^{1,2,4}, Giuseppe Pica^{1,4}, Giulio Bondanelli^{1,4}, Monica Moroni^{1,4}, Caroline A. Runyan^{1,3}, Ari S. Morcos^{1,3}, Christopher D. Harvey^{1,3}✉ and Stefano Panzeri¹✉

Noise correlations (that is, trial-to-trial covariations in neural activity for a given stimulus) limit the stimulus information encoded by neural populations, leading to the widely held prediction that they impair perceptual discrimination behaviors. However, this prediction neglects the effects of correlations on information readout. We studied how correlations affect both encoding and readout of sensory information. We analyzed calcium imaging data from mouse posterior parietal cortex during two perceptual discrimination tasks. Correlations reduced the encoded stimulus information, but, seemingly paradoxically, were higher when mice made correct rather than incorrect choices. Single-trial behavioral choices depended not only on the stimulus information encoded by the whole population, but unexpectedly also on the consistency of information across neurons and time. Because correlations increased information consistency, they enhanced the conversion of sensory information into behavioral choices, overcoming their detrimental information-limiting effects. Thus, correlations in association cortex can benefit task performance even if they decrease sensory information.

The collective activity of a population of neurons, beyond properties of individual cells, is critical for perceptual discrimination behaviors^{1,2}. A fundamental question is how functional interactions in a population impact both the encoding of sensory information and how this information is read out to guide behavioral choices³. A commonly studied feature of population coding is noise correlations, the correlated trial-to-trial variability over repeated presentations of the same stimulus^{4,5}. Noise correlations can take the form of ‘across-neuron’ correlations between the time-averaged spike rates of different neurons or populations, or ‘across-time’ correlations between the activity of the same neural population at different times.

The impact of correlations has been long debated. Much work has proposed that they limit the information encoding capacity of a neural population^{6–9}. Based on the widely held assumption that perceptual discrimination performance increases proportionally with the amount of sensory information encoded in neural activity¹⁰, this has been taken to imply that information-limiting correlations hinder the ability to discriminate sensory stimuli^{6,10}. Specifically, across-time correlations have been proposed to limit the benefit of integrating noisy information over time for a speed-accuracy trade-off^{11,12}. Further, across-neuron correlations are thought to lessen the benefit of averaging noisy information across neural populations^{6,7,9}.

However, the effect of noise correlations may be more nuanced, as indicated by a separate stream of biophysical studies showing that spatially and temporally correlated spiking can more strongly drive responses in postsynaptic neural populations^{13–17}. It remains poorly tested if and how enhanced signal propagation may have a beneficial impact on behavioral discrimination performance.

We investigated how noise correlations shape behavioral performance in perceptual discrimination by studying together not only how correlations impact the encoding of sensory information, but critically also how they impact the readout of this information by downstream neural circuits to guide behavioral choices.

Correlations of posterior parietal cortex activity limit sensory coding

To examine how noise correlations affect both stimulus coding at the population level and behavioral discrimination performance, we focused on the mouse posterior parietal cortex (PPC). The PPC participates in transforming multisensory signals into behavioral outputs, is essential for perceptual discrimination tasks during virtual navigation¹⁸ and has stimulus information related to an animal’s choices^{19–23}. It is thus a relevant area to study the impact of correlated neural activity on behavior.

We examined across-time and across-neuron correlations in PPC population activity using previously published datasets. To study across-time correlations, we used calcium imaging data from a sound localization task¹⁹ in which mice reported perceptual decisions about the location of an auditory stimulus by navigating through a visual virtual reality T-maze (Fig. 1a). As mice ran down the T-stem, a sound cue was played from one of eight possible locations in head-centered coordinates. Mice reported whether the sound originated from their left or right by turning in that direction at the T-intersection ($78.0\% \pm 0.5\%$ correct). During each session, the activity of ~50 layer 2/3 neurons was imaged simultaneously. Because the same sensory cue (sound location) was presented throughout the trial, this task is well suited to study across-time correlations.

To study across-neuron correlations, we used calcium imaging data from an evidence accumulation task²¹ in which ~350 layer 2/3 neurons were imaged simultaneously for each session. During virtual navigation, mice were presented with six temporally separate visual cues on the left or right walls of a T-maze (Fig. 1f). Mice reported which side had more cues by turning in that direction at the T-intersection ($84.5\% \pm 1.6\%$ correct). We categorized the visual stimuli as having more total left or right cues. Because of the high number of simultaneously imaged neurons, this task is well suited for studying across-neuron correlations.

¹Neural Computation Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy. ²Centro Interdisciplinare Mente e Cervello (CIMeC), University of Trento, Rovereto, Italy. ³Department of Neurobiology, Harvard Medical School, Boston, MA, USA. ⁴These authors contributed equally: Martina Valente, Giuseppe Pica, Giulio Bondanelli, Monica Moroni. ✉e-mail: harvey@hms.harvard.edu; stefano.panzeri@iit.it

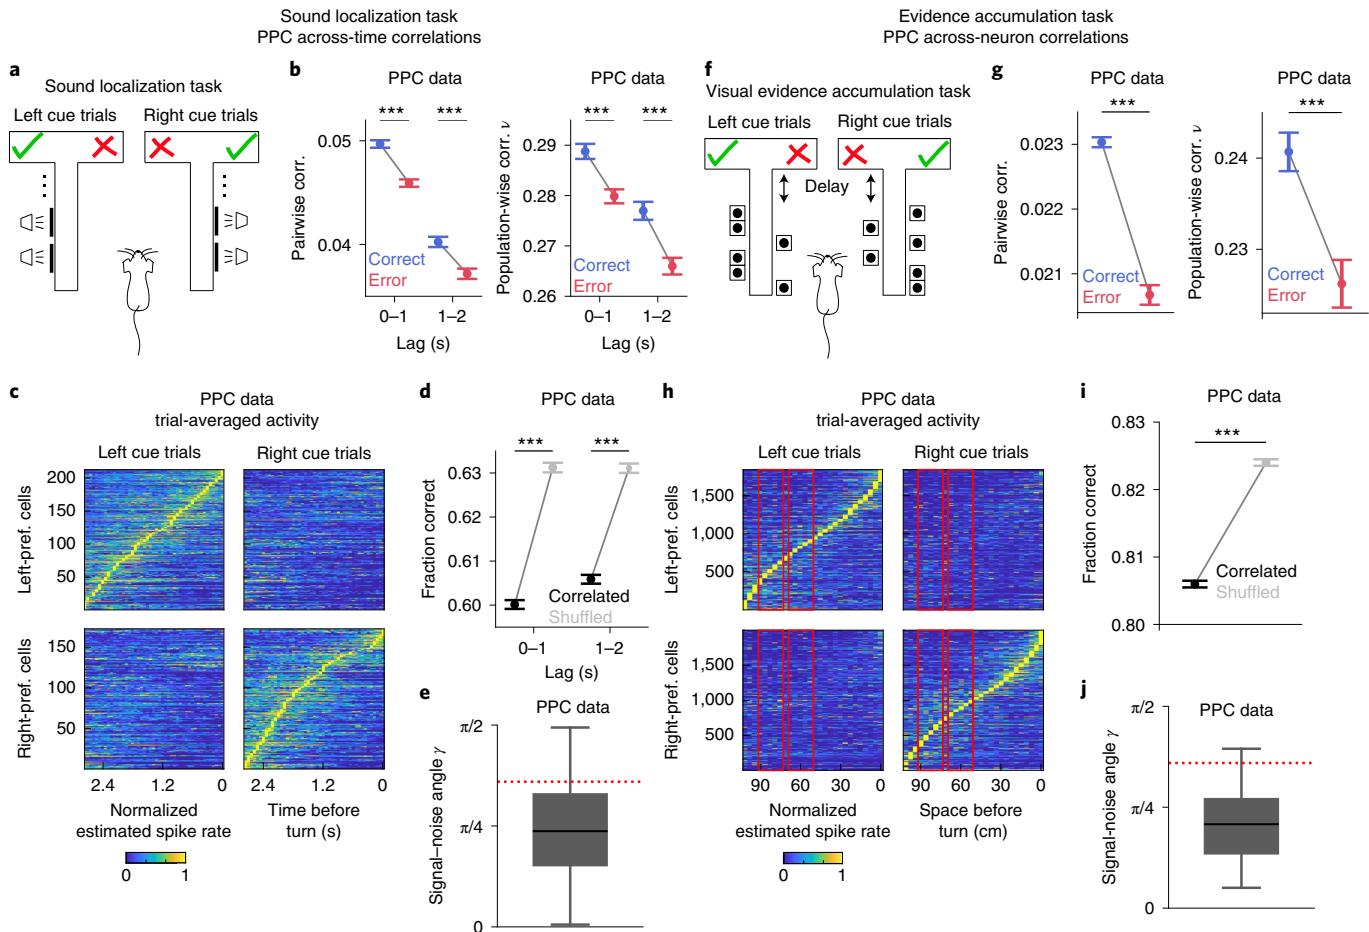


Fig. 1 | Response properties and across-time and across-neuron correlations in mouse PPC during perceptual discrimination tasks. **a–e**, PPC data during the sound localization task. **a**, Schematic of the task. Left and right sound categories (speaker symbols) indicate the rewarded side of the maze (check mark). **b**, Pairwise (left) and population-wise (right) noise correlations in time-lagged activity, for correct (blue) and error (red) trials. **c**, Trial-averaged estimated spike rate traces for PPC cells (left-preferring, $n=212$; right-preferring, $n=172$), aligned to the turn frame, normalized to each cell's peak mean activity and sorted by peak time. **d**, Accuracy of a linear decoder of stimulus category applied to joint population activity, for real recorded (black) or trial-shuffled (gray) data. In **b** and **d**, error bars report the mean \pm s.e.m. across all cell pairs (**b**: left, $n_{\text{pairs}}=133,860$ and $119,562$ for lag 0–1 s and 1–2 s, respectively) and all time-point pairs within the specified lag range from $n=6$ sessions. For all comparisons, $***P=10^{-4}$ (two-sided permutation test). **e**, Distribution of the signal-noise angle, γ (over $n=6$ sessions and all time-point pairs within a 2-s lag). Box plots show the median (line), quartiles (box) and whiskers extending to ± 1.5 times the interquartile range. The red dotted line denotes the analytically computed bound between the information-limiting and information-enhancing regime. **f–j**, PPC data during the evidence accumulation task. **f**, Schematic of the task. The rewarded side of the maze (check mark) is the one identified by the most numerous visual cues (wall segments with black dots patterns). **g**, Pairwise (left) and population-wise (right) noise correlations in time-averaged activity, for correct (blue) and error (red) trials. Error bars report the mean \pm s.e.m. across all cell pairs (left, $n_{\text{pairs}}=1,561,202$) and early delay and late delay epochs from $n=11$ sessions. For all comparisons, $***P=10^{-4}$ (two-sided permutation test). **h**, As in **c**, for all PPC cells (left-preferring, $n=1,840$; right-preferring: $n=2,000$). Activity traces are averaged over spatial bins (about 200 ms). Red rectangles indicate the early delay and late delay epochs. **i**, Accuracy of a linear decoder of stimulus category applied to joint population activity, for real recorded (black) or within-pool trial-shuffled (gray) data. Error bars report the mean \pm s.e.m. across $n=11$ sessions, early delay and late delay epochs and 100 pool splits. $***P=10^{-4}$ (two-sided permutation test). **j**, As in **e**, for across-time correlations, using data from $n=11$ sessions.

For all analyses, we focused on the period toward the end of the T-stem before the mouse had reported its choice and after it had received sensory information. This is a window in which neural activity may carry sensory information used to inform choice.

We tested how noise correlations impacted the encoding of sensory information in population activity. We computed both pairwise correlations (Pearson correlation between activities of two neurons) and population-wise correlations (the fraction of total population activity variance carried by the largest principal component¹) for each stimulus category. Pairwise and population-wise correlations were positively related and varied consistently across conditions in PPC data (Fig. 1b,g and Extended Data Fig. 1g–k,s–w).

and in population encoding models (Supplementary Note 3 and Extended Data Fig. 2a,b). PPC neurons had, on average, positive across-neuron and across-time pairwise and population-wise noise correlations (Fig. 1b,g and Extended Data Fig. 1g–j,s–v). Since many neurons exhibited activity selective for distinct trial types, with different neurons active at distinct time points in the trial (Fig. 1c,h), we could decode the stimulus category significantly above chance from pairs of temporally offset instantaneous population activity vectors (Fig. 1d) and from the population activity vector in one time window (Fig. 1i). Stimulus category decoding performance was higher in the evidence accumulation dataset because of the larger number of recorded neurons. To evaluate how across-time

correlations affected the encoding of stimulus category, we shuffled instantaneous population activity vectors across trials of the same stimulus category, independently at each time point. This shuffle destroyed within-trial temporal relationships while preserving instantaneous population activity. We disrupted across-neuron correlations by randomly splitting the neural population into two non-overlapping pools of neurons of equal size and shuffling the trial labels separately for each pool within the same stimulus category. Importantly, in both datasets, stimulus-decoding performance was higher when across-time or across-neuron correlations were disrupted by shuffling, indicating that both forms of correlations limited stimulus category information in population activity (Fig. 1d,i).

Whether noise correlations limit the information encoded by a neural population depends on how they relate to signal correlations (correlations between trial-averaged responses to individual stimuli^{4,5,24}). We quantified their relationship using the angle^{25,26} between the axis of largest variation at fixed stimulus (noise correlations axis) and the axis of largest stimulus-related variation (signal correlations axis) in the high-dimensional space of population activity (see Fig. 2c for a sketch). Encoding models^{5,24} predict that the smaller the signal–noise angle, the more noise correlations impair stimulus discrimination due to larger overlap between the stimulus-specific response distributions to different stimuli (Figs. 2c and 3a). Supporting the observation above that showed information-limiting correlations in our datasets, most signal–noise angles resided in the information-limiting regime, below the signal–noise angle critical value for the transition to the information-enhancing regime reported in previous work^{27,28} (Fig. 1e,j). Correlations were information-limiting both in ‘easy’ and ‘hard’ trials with high and low levels of sensory evidence, respectively (Extended Data Fig. 1a,b,d,e,m,n,p,q).

If correlations are detrimental to perceptual behaviors, one would expect noise correlations to be lower when animals make correct choices and higher when animals make errors. Contrary to this expectation, both across-time and across-neuron noise correlations were higher in correct trials than in error trials (Fig. 1b,g). Thus, although correlations limit information in population activity, including on correct trials (Extended Data Fig. 1c,f,o,r), they might not impair behavioral performance.

A model of how correlations affect task performance

The above findings lead to the paradoxical suggestion that correlations limit information encoded by a neural population but at the same time may be beneficial for making accurate choices. To reconcile these observations, we developed a simple mathematical model that incorporated both the encoding of stimulus information and the readout of this information to form a choice. We compared two alternative views of how information in population activity may be used to perform a stimulus discrimination task. In the traditional view, choice accuracy is proportional to the amount of information in a neural population, and thus information-limiting correlations constrain task performance. Alternatively, a choice could depend on both stimulus information and features of neural activity that emerge from correlations, in particular the consistency of information across time and neurons in a population.

We simulated a perceptual discrimination task with two possible stimuli that had to be converted into two possible corresponding choices ($c=1$ for $s=1$ and $c=-1$ for $s=-1$; Fig. 2a). We simulated trials of two N -dimensional sets of neural activity features \mathbf{r}_1 and \mathbf{r}_2 , which can alternatively represent neural activity of a pool of N neurons at different points in time (for across-time correlations) or activity of two different pools of N neurons each (for across-neuron correlations). Figure 2 and Extended Data Fig. 3 illustrate the model using a case with unrealistically high noise correlations and one neuron per feature ($N=1$). We display results for $N=20$ (Fig. 3) and $N=10$ (Extended Data Fig. 4) neurons per feature, to show that these are largely independent of N , and to document how they depend on the two parameters that are key for model behavior and comparison with real data: the signal–noise angle and the strength of population-wise noise correlations (Supplementary Notes 1 and 2).

In the encoding model, for each feature, higher trial-averaged activity was associated with one sensory stimulus ($s=1$), while lower mean activity corresponded to the opposite sensory stimulus ($s=-1$), meaning that the two features showed positive signal correlations (Fig. 2c). We simulated noise correlations between \mathbf{r}_1 and \mathbf{r}_2 (noise correlation strength intuitively corresponds to the elongation of the ellipses depicting the distributions of responses to stimuli), and we varied the signal–noise angle across simulations

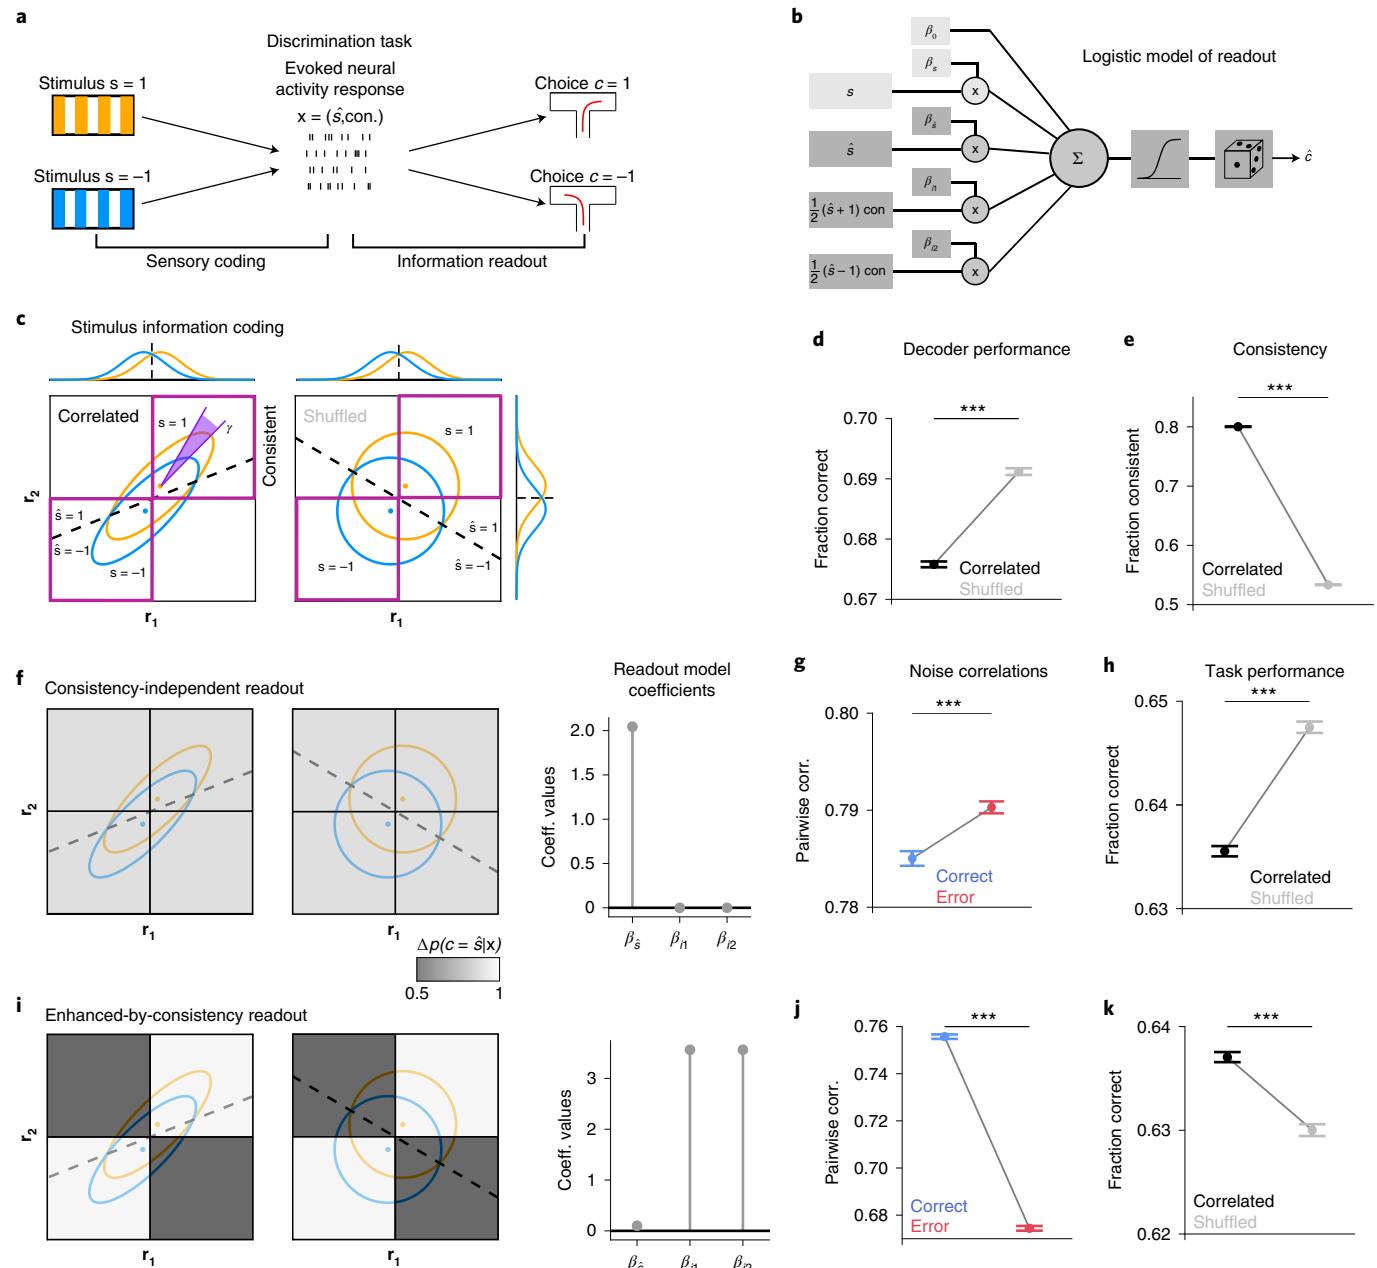
Fig. 2 | A simple encoding-readout model shows how different readouts determine the impact of correlations on task performance. **a**, Schematic conceptualizing two fundamental information processing stages in sensory perception included in the model: sensory coding (mapping from sensory stimuli to neural activity) and information readout (mapping from neural activity to behavioral choice). Task-relevant neural activity is recapitulated by the stimulus predicted from population activity (\hat{s}) and its consistency across features (con.). **b**, Schematic of the readout model used to model choices. Nonneural predictors (bias and real stimulus (light gray boxes on the left); used only in real data analyses to account for the effect of nonrecorded neurons) and neural predictors (predicted stimulus and consistency (dark gray boxes on the left); used both in simulations and data analyses) were weighted, summed and transformed through a sigmoid function that outputs the binomial probability of a binary choice. **c**, Example of simulated response distributions of two one-dimensional neural features (\mathbf{r}_1 and \mathbf{r}_2) to two stimuli ($s=1$ and $s=-1$), modeled as bivariate Gaussian distributions. For the correlated example, noise and signal axes were closely but not perfectly aligned ($\gamma=0.08\pi$). Ellipses denote the 95% confidence intervals of stimulus-specific activity distributions. The dashed black lines represent the optimal stimulus-decoding boundary. Purple squares denote regions in which \mathbf{r}_1 and \mathbf{r}_2 encode consistent stimulus information, that is, the same stimulus is decoded from both features. Marginal response distributions and decoding boundaries are shown next to the graphs. **d**, Accuracy of a linear decoder of stimulus applied to simulated responses was higher for uncorrelated (shuffled) responses (correlations limit the encoded stimulus information). **e**, The fraction of trials in which \mathbf{r}_1 and \mathbf{r}_2 encode consistent stimulus information is higher for correlated responses (correlations increase consistency). Data in **f–h** refer to the consistency-independent readout. **f**, The readout map is superimposed on left and middle plots and represented as a grayscale map in the \mathbf{r}_1 – \mathbf{r}_2 response space. Orange and blue ellipses represent stimulus-specific distributions. Shades of gray represent readout efficacy (probability of transformation from encoded stimulus to choice), which was independent of consistency for this readout. The corresponding readout model regression coefficients are shown on the right (β_s for the predicted stimulus, and β_{11} and β_{12} for the consistency-dependent interaction terms between predicted stimuli and neural consistency). **g**, Pairwise noise correlations were higher in error trials. **h**, Task performance computed using this readout model was higher for uncorrelated (shuffled) responses. Data in **i–k** refer to the enhanced-by-consistency readout (consistency modulation index $\eta=0.9$; Methods). **i**, Same as **f**. The readout map was superimposed on the left and middle plots. The readout efficacy was higher for consistent trials. **j**, Same as **g**. Pairwise noise correlations were higher in correct trials. **k**, Same as **h**. Task performance was higher for correlated responses, indicating that this readout overcame the information-limiting effect of correlations. In **d,e,g,h,j** and **k**, error bars report the mean \pm s.e.m. over 200 simulations with 5,000 trials each with $N=1$, $\gamma=0.08\pi$, $d=\sqrt{0.02}$ and $\sigma=0.2$. Correlated: $\rho=0.8$, equivalent to $\nu=0.9$. Uncorrelated (shuffled): $\rho=0$, equivalent to $\nu=0.5$. For all comparisons, *** $P=10^{-4}$ (two-sided permutation test). Model parameters were purely illustrative and did not match real data.

(for example, in Fig. 2c, this angle is small, as signal and noise were closely but not perfectly aligned). When signal–noise angles were small, noise correlations increased the overlap between the stimulus-specific response distributions (cf. orange and blue ellipses in Fig. 2c) and decreased the stimulus information encoded by the two features jointly (Figs. 2d and 3a). There was a critical signal–noise angle value below which correlations limited information and above which they enhanced information (Fig. 3a). This value depended mildly on the noise correlation strength (Fig. 3a) but was largely independent of the stimulus information level or the population size (Supplementary Note 1 and equation (3)).

We then considered the readout stage of the model. Commonly, the choice during a task trial is expected to follow the decoded stimulus. However, because of the apparent importance of correlations for accurate choices in our experimental data, we hypothesized that the readout of stimulus information might utilize aspects of population activity imposed by correlations. Intuitively, correlations imply that there is greater consistency in the neural population representa-

tions (Fig. 2e). In our model, we defined consistency as a single-trial measure of similarity between the stimuli that are decoded from features r_1 and r_2 separately. In our simulations, a stimulus representation in a trial was classified as consistent when features r_1 and r_2 both signaled the same stimulus (that is, both features were higher than average and thus both signaling $s = 1$ (Fig. 2c) or both features were lower than average and thus both signaling $s = -1$ (Fig. 2c)).

We simulated choices for a binary task discrimination with two alternative readout models, formulated as logistic models of the dependence of choice on several features of stimulus encoding. In the first model, termed the consistency-independent readout, the simulated choice in each trial depended only on the stimulus decoded from the two features jointly (Fig. 2f). This case followed the traditional assumption²⁹ that the choice reflects the stimulus decoded from the full population activity (Fig. 2f). In our experimental data, we did not observe a perfect correspondence between the stimulus decoded from neural activity and the choice of the mouse (the mean \pm s.e.m. over sessions and trials of the fraction



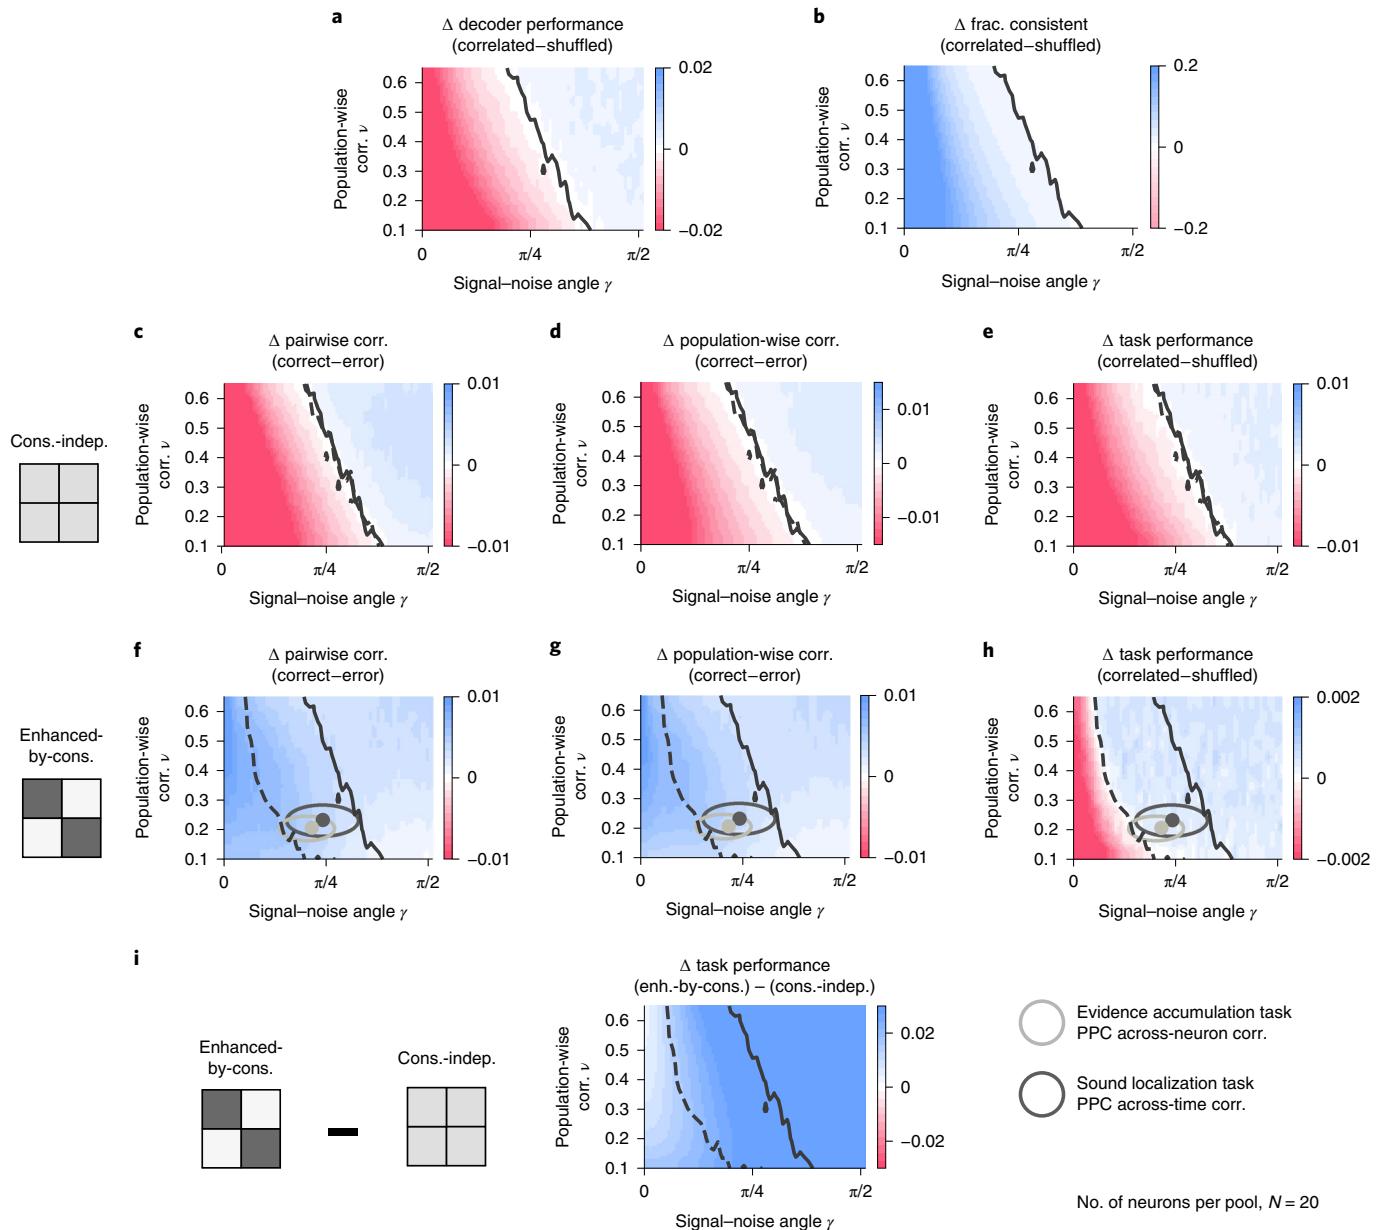


Fig. 3 | Exploration of the parameter space of the encoding-readout model. Data were simulated using the encoding-readout model with $N=20$ neurons in each pool. Data represent the mean over 100 simulations with 300,000 trials each. **a**, Difference in the accuracy of a linear decoder of stimulus applied to correlated and shuffled simulated neural activity for different values of the signal-noise angle (γ) and population-wise correlations (v). The black solid line represents the boundary between a regime with information-limiting correlations and a regime with information-enhancing correlations. **b**, The difference between correlated and shuffled activity in the fraction of trials in which the two neural features encode consistent stimulus information was higher in the information-limiting regime and increases with the population-wise correlations strength. **c–e**, Data refer to consistency-independent readout models. **c,d**, Difference in average pairwise correlations (**c**) and population-wise correlations (**d**) between trials with correct and incorrect predicted task performance for different combinations of model parameters. **e**, Differences in task performance were predicted by applying the consistency-independent readout model to correlated and shuffled simulated neural activity for different combinations of model parameters. **c–i**, Dashed black lines represent the boundary between a regime where task performance was higher for correlated responses and a regime where performance was higher for shuffled responses. The overlap between the continuous and dashed black line indicates that information-limiting correlations are also detrimental for behavior. Data in **f–h** refer to the enhanced-by-consistency readout model (consistency modulation index $\eta=0.85$). **f,g**, Same as in **c** and **d**. The enhanced-by-consistency readout correlations were higher in correct trials. **h**, Same as in **e**. The area between the dashed and the continuous black lines indicates a regime where correlations were information-limiting but task performance was higher for correlated responses. In the parameter range between the two lines, this readout can overcome the negative impact of correlations. Dark and light gray dots and ellipses denote mean values and range between the 25th and the 75th percentiles of the signal-noise angles and population-wise correlations for PPC data from the sound localization task and evidence accumulation task, respectively. **i**, Differences in task performance were predicted by applying the enhanced-by-consistency readout model or the consistency-independent readout model (with matched readout efficacy) for different combinations of model parameters.

of times the mouse's choice matched the decoded stimulus was $61.0\% \pm 0.2\%$ in the sound localization dataset and $91.1\% \pm 0.1\%$ in the evidence accumulation dataset). Therefore, in the model, we set the probability that in a given trial the choice matched the decoded stimulus (termed 'readout efficacy') to a value smaller than 100%.

In the second readout model, termed the 'enhanced-by-consistency' readout, the choice in each trial depended not only on the stimulus decoded from both features jointly, but also on the consistency of the stimulus decoded from the features separately (Fig. 2*i*). If \mathbf{r}_1 and \mathbf{r}_2 reported consistent information about the stimulus, this readout was more likely to use the stimulus encoded in neural activity to inform the choice. This effect was reflected in the positive coefficients assigned to the interaction terms between the decoded stimulus and consistency (Fig. 2*i*). In other words, the readout efficacy was higher when the two features were consistent (Fig. 2*i*). Importantly, the average readout efficacy of this model was matched to the readout efficacy of the consistency-independent model.

For the consistency-independent readout, correlated activity resulted in worse task performance compared to activity in which correlations were absent, across the entire information-limiting regime of the model's parameter space (Figs. 2*h* and 3*e*). This was expected because, with this readout, the task performance directly follows the level of stimulus information, with higher information resulting in higher performance. Further, in the information-limiting regime, noise correlations were higher on simulated error trials than on correct trials (Figs. 2*g* and 3*c,d*), which was notably inconsistent with our PPC data (Fig. 1*b,g*).

For the enhanced-by-consistency readout, larger noise correlations within the information-limiting region increased the fraction of trials with consistent information (Fig. 3*b* and Extended Data Fig. 4*b*). As a result, correlations generated a larger fraction of trials that were better read out by the enhanced-by-consistency readout. Because of this, this readout produced behavioral task performance that was as good as or better than that produced without noise correlations (Figs. 2*k* and 3*h*). When the signal-noise angle was not too small but still in the information-limiting region (so that the information decrease due to correlations was not too large), the enhanced-by-consistency readout compensated and even overcame the information-limiting effects of correlations (Fig. 3*h*). Interestingly, signal-noise angles and noise correlation values estimated from PPC data resided mostly in this specific parameter region (Fig. 3*h*).

Importantly, for the enhanced-by-consistency readout, noise correlations were higher in correct trials than error trials (Figs. 2*j* and 3*f,g*), matching our experimental PPC findings. Thus, the enhanced-by-consistency readout reconciled our experimental observations, providing a mechanism whereby correlations limit information but benefit task performance.

Correlations and consistency contribute to choices

We then used our experimental measurements of PPC neural activity to test for signatures of an enhanced-by-consistency readout. A key prediction of this readout is that the mouse's single-trial choices should depend not only on the correctness of stimulus encoding but also on the consistency of stimulus information. In our experimental data, we defined consistency as the single-trial similarity between the stimuli decoded from population activity at different points in time (across-time consistency) or between the stimuli decoded from separate neuronal pools in the same time window (across-neuron consistency). An example of across-time consistency is a trial in which population activity at time t_1 signaled the same stimulus category as the population activity at time t_2 . We calculated the mouse's performance for four subclasses of trials, defined by the correctness and consistency of the stimulus decoded from neural activity in a given trial. In both datasets, the mouse's task performance was higher for trials with correctly decoded

stimulus information than for incorrectly decoded trials, suggesting that the stimulus information carried by PPC neurons was used to inform behavioral choices (Fig. 4*b,f*). Also, the mouse's task performance was higher for trials with consistent information across time or across neurons, suggesting that the consistency of neural population information was important for accurate choices (Fig. 4*a,e*). Critically, in trials with correctly decoded stimulus information, the mouse's task performance was higher when information was consistent than when it was inconsistent, both across neurons and across time (Fig. 4*b,f*). Further, in trials with incorrectly decoded stimulus information, task performance was lower on consistent trials than on inconsistent trials (Fig. 4*b,f*). These findings indicate that the stimulus information in the PPC was read out in a manner that, in consistent trials, amplified the effect of the decoded stimulus on the mouse's choice, both when the decoded information was correct and incorrect.

To rule out that differences in a mouse's task performance between consistent and inconsistent trials were due to higher stimulus information in consistent trials, we sorted trials according to the stimulus information level and verified that performance in correctly and incorrectly decoded trials was still higher and lower, respectively, when information was consistent across neurons or across time (Extended Data Fig. 5*c-e* and *g-i*).

We further examined the possibility of an enhanced-by-consistency readout by developing an analytical understanding of how single-trial choices were made by the mouse. We used logistic regression to relate PPC activity to the mouse's choices. We expressed a mouse's choice on a given trial as a function of features of the recorded neural activity: the stimulus decoded from the full PPC population activity and its interaction with the across-time or across-neuron consistency (Fig. 2*b*). In addition, we included a predictor for the experimenter-defined stimulus presented to the mouse and a bias term. These two terms captured the stimulus-related and stimulus-unrelated information carried by sources other than the recorded neurons, such as nonrecorded neurons. The inclusion of these terms allowed us to test how much the stimulus information in the recorded neural activity explained the mouse's choice after discounting what could be explained by other sources.

The regression coefficients for the stimulus decoded from neural activity were positive (Fig. 4*d,h*), indicating that the neural stimulus information impacted the mouse's choice. The coefficients for the consistency-dependent terms were also positive, indicating that the readout of PPC activity performed similarly to the enhanced-by-consistency readout model from Fig. 2*j*; that is, the probability that the choice matched the stimulus decoded from neural activity was higher in consistent trials (Fig. 4*d,h*). We tested the specific contribution of the neural-based predictors in explaining the mouse's choices by fitting the logistic regression after shuffling the values of these predictors across trials (Fig. 4*c,g*). Shuffling all neural-based predictors (Fig. 4*c,g*) made it harder to predict a mouse's choice, again demonstrating that neural activity contributed to choices. Moreover, shuffling only the neural consistency values, while leaving the stimulus decoded from neural activity and the experimenter-determined stimulus intact, resulted in worse predictions of a mouse's choices. Using readout models that nonlinearly combined neural population activity to directly decode the stimulus regardless of consistency also failed to explain choices as well as the enhanced-by-consistency readout model (Fig. 4*c,g*), even though they decoded the stimulus well (Extended Data Fig. 1*l,x*). The specific nonlinear interaction between neural consistency and decoded stimulus thus was key to form the mouse's choice.

To rule out that the modulation of the readout by consistency might just reflect differences in overall stimulus information levels between consistent and inconsistent trials, we verified that consistency provided a similar contribution to predicting a mouse's

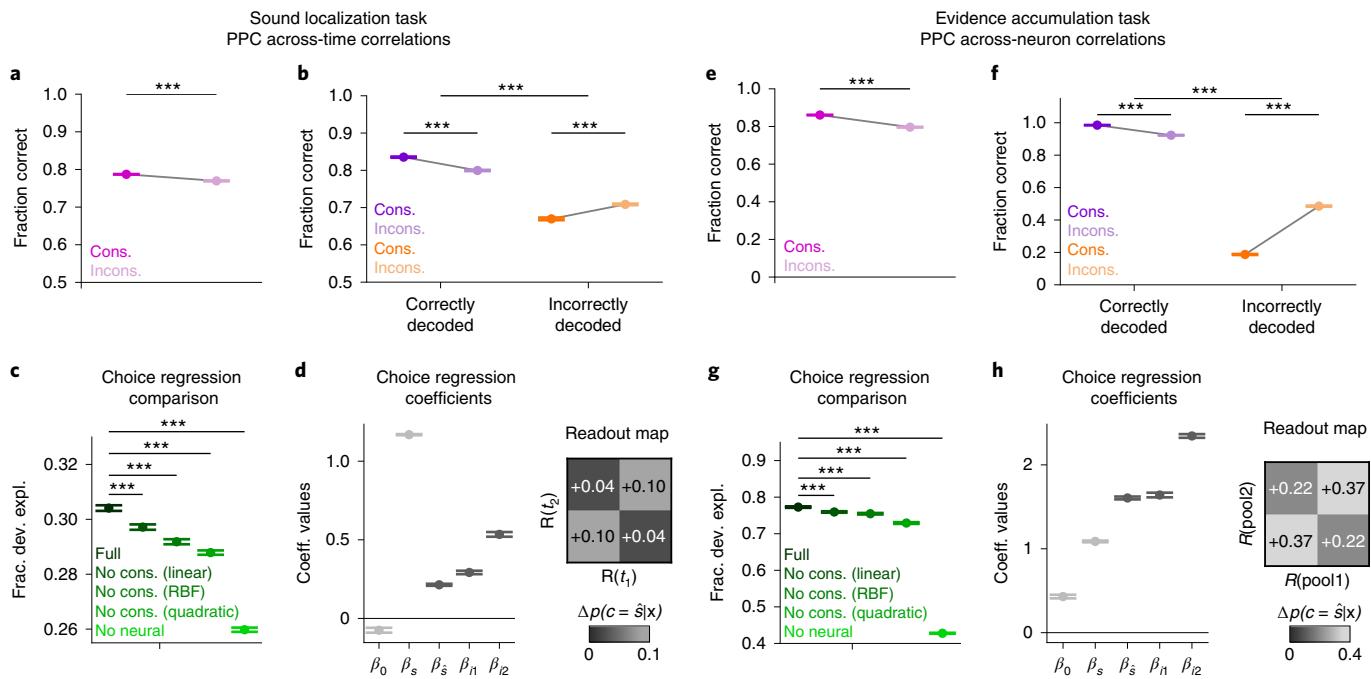


Fig. 4 | Across-time and across-neuron correlations in PPC activity influence mouse choices. Data in **a-d** refer to PPC activity during the sound localization task. **a**, Task performance (fraction correct) was higher when neural population vectors encoded the stimulus consistently across time. **b**, Task performance (fraction correct) was higher in trials with correct stimulus decoding, suggesting that stimulus information is used to inform behavior. Left: task performance in trials with correctly decoded stimulus was higher with consistent information encoding. Right: the opposite happened when information is decoded incorrectly. Thus, stimulus information in neural activity has a larger impact on choices when it is encoded consistently. **c**, Performance (fraction of deviance explained) in explaining single-trial choices of various readout models. Full model uses all predictors (neural and nonneural); other models neglect information from selected predictors as follows. No cons: neglects neural consistency; no neural: neglects stimulus decoded from neural activity and neural consistency. Linear, quadratic or radial basis function (RBF) support vector machines (SVMs) were used to decode stimuli from neural activity as indicated; otherwise, a linear SVM was used. **d**, Left: best-fit coefficients of the full readout model (β_0 and β_s are nonneural coefficients corresponding to bias and stimulus-related drive due to nonrecorded neurons; $\beta_{\hat{s}}$ is the coefficient of the predicted stimulus; and $\beta_n, \beta_{\hat{z}}$ correspond to the consistency-dependent interaction terms between each predicted stimulus and the neural consistency, with positive values amplifying the predicted stimulus effect in consistent trials). Right: readout efficacy estimated from the best-fit coefficients of the full model, above the baseline-level due to nonneural predictors, for consistent and inconsistent population vectors, represented schematically as a readout map in the two-dimensional response space similarly to Fig. 2. In **a-c**, for all comparisons, $***P=10^{-4}$, two-sided permutation test. In **a-c** and in **d** (left), error bars report the mean \pm s.e.m. across $n=6$ sessions and all time-point pairs within a 1-s lag. In **d** (right), the gray levels represent mean values over $n=6$ sessions and all time-point pairs within a 1-s lag. Data in **e-h** refer to PPC activity during the evidence accumulation task. **e-h**, Same as in **a-d**. In **e-h**, consistency and choices were computed from the activity of two neuronal pools. Also in the evidence accumulation task, stimulus information in neural activity had a larger impact on choices when it was encoded consistently, and choices depended critically on neural consistency across pools. In **e-g**, for all comparisons, $***P=10^{-4}$ (two-sided permutation test). In **e-g** and in **h** (left), error bars report the mean \pm s.e.m. across $n=11$ sessions, early delay and late delay epochs and $n=100$ random pool splits. In **h** (right), the gray levels represent mean values over $n=11$ sessions, early delay and late delay epochs and $n=100$ random pool splits.

choices when we used a more sophisticated logistic model that included the magnitude of the stimulus information, instead of only the identity of the decoded stimulus (Extended Data Fig. 5b,f). Further, to control for and discount potential contributions from movement-related neural activity, we verified that neural consistency also contributed to predicting choices when adding to the regression the consistency of the mouse's running speed and direction (Extended Data Fig. 6c-e).

An enhanced-by-consistency readout benefits task performance

Our results show that across-time and across-neuron consistency in the experimental data impact a mouse's choices. We examined the implications of this finding for mouse task performance, in the presence or absence of experimentally measured information-limiting correlations. Because correlations cannot be removed experimentally, we instead created a set of simulated choices using the experimentally fit logistic choice regression from Fig. 4. As input to the experimentally fit choice regression, we used either trials with

simultaneously recorded PPC neural activity or trials with neural activity shuffled to disrupt across-time or across-neuron correlations. We used these simulated choices to estimate how well the mouse would have performed on the task with and without correlations present (Fig. 5a).

We focused only on the contribution of the recorded neural population, by computing task performance from choices simulated using all predictors extracted from experimental data and then subtracting the task performance computed from choices simulated after shuffling across trials the values of neural predictors. This calculation is more precise than the one obtained by simply mapping the session-averaged PPC parameters to the encoding-readout model (Fig. 3h) because it estimates the contribution of correlations to task performance using single trials recorded in each session. In the sound localization dataset, the ~50 recorded neurons were estimated to increase task performance by ~3.5%, and in the evidence accumulation dataset, the ~350 neurons increased task performance by ~25% (Fig. 5b,g). Strikingly, although the stimulus infor-

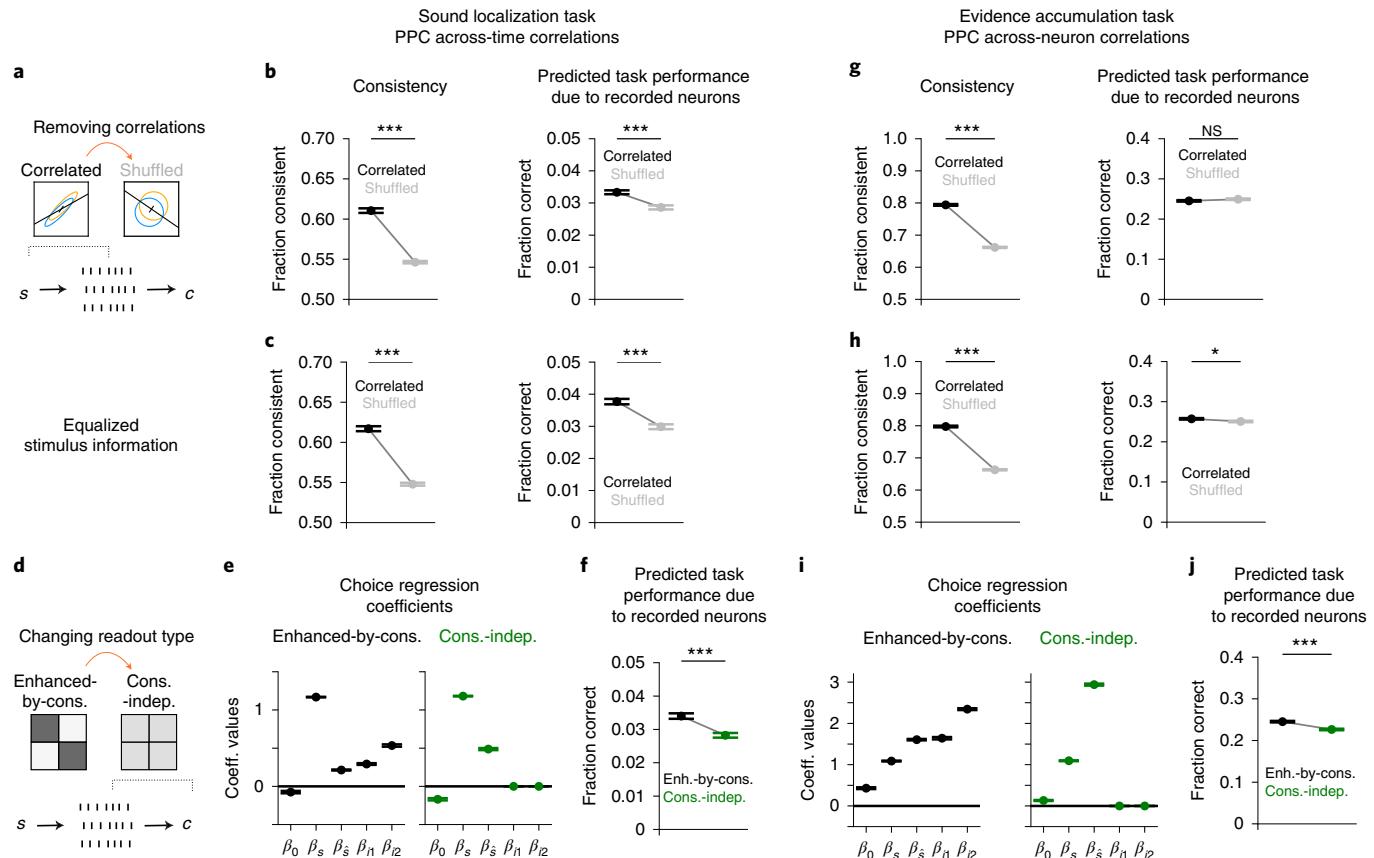


Fig. 5 | Simulated mouse choices show that the best-fit enhanced-by-consistency readout improves task performance in the presence of information-limiting correlations in the PPC. **a**, Schematic of the method used to study the impact of correlations on task performance. The best-fit readout model was applied to the original correlated patterns of neural activity and to artificial patterns where correlations were removed by shuffling. Data in **b** and **c** refer to PPC activity during the sound localization task. **b**, Left: fraction of trials with stimulus information encoded consistently across pairs of time points, for real recorded (black) or trial-shuffled (gray) population vectors. Right: task performance due to the recorded neurons predicted by applying the best-fit readout model to real recorded (black) or trial-shuffled (gray) population vectors. **c**, Same as in **b**, computed after subsampling trials to equalize the encoded stimulus information between correlated and shuffled data. **d**, Schematic of the method used to study the impact of the enhanced-by-consistency readout on task performance. Task performance was predicted by applying the best-fit enhanced-by-consistency model to the original correlated patterns of neural activity and a consistency-independent readout model with readout efficacy matched to the best-fit readout model. Data in **e** and **f** refer to PPC activity during the sound localization task. **e**, Left: best-fit coefficients of the enhanced-by-consistency readout model. Right: coefficients of the consistency-independent model with matched readout efficacy. **f**, Task performance due to the recorded neurons predicted by applying the best-fit enhanced-by-consistency (black) and the consistency-independent (green) readout models to the original correlated patterns of neural activity. In **b,c,e** and **f**, error bars represent the mean \pm s.e.m. across $n=6$ sessions and all time-point pairs within a 1-s lag. For all comparisons, $^{***}P=10^{-4}$ (two-sided permutation test). Data in **g-j** refer to PPC activity during the evidence accumulation task. **g**, Same as in **b**, for across-neuron correlations. **h**, Same as in **c**, for across-neuron correlations. **i**, Same as in **e**, for across-neuron correlations. **j**, Same as in **f**, for across-neuron correlations. $^{***}P=10^{-4}$ and $^{*}P=0.0105$ (two-sided permutation test). In **g-j**, error bars represent the mean \pm s.e.m. across $n=11$ sessions, early delay and late delay epochs and $n=100$ splits in pairs of neuronal pools.

mation in the recorded neurons was lower with correlations intact (Fig. 1d,i), the recorded neurons increased task performance to a greater extent when across-time correlations were intact than when they were removed by shuffling in the sound localization dataset (Fig. 5b). Furthermore, in the evidence accumulation task, the recorded neurons contributed similarly to task performance with and without across-neuron correlations intact, despite lower information with correlations present (Fig. 5g). Thus, the enhanced-by-consistency feature of the experimentally fit readout could overcome, or at least offset, the information-limiting effect of correlations and benefit task performance.

These results incorporate the overall impact of correlations on task performance by combining the effects of the encoding and readout. To quantify the specific contribution of the readout, we again simulated choices from the experimentally fit choice

regression, except we equalized the stimulus information in the correlated and shuffled responses by selecting subsets of trials that had the same fraction of correctly decoded stimuli. With this matching, the correlated and shuffled trials differed only in their neural consistency (Fig. 5c,h), with a proportion of consistent trials equal to those of the full data (Fig. 5b,g). For both datasets, the estimated contribution to task performance of the recorded neurons was higher when correlations were intact than when they were disrupted (Fig. 5c,h). This result shows that the readout of PPC activity was more efficient in extracting information from correlated than from uncorrelated data.

These results also indicate that the readout of stimulus information from PPC activity is suboptimal. From the ~50 neurons recorded in the sound localization task and ~350 neurons in the evidence accumulation task, we were able to decode the stimulus

at ~60% and ~80% correct, respectively (Fig. 1d,i). These populations therefore could have increased task performance by ~10% and ~30% above chance, respectively, if stimulus information was read out optimally. However, these populations only increased task performance by ~3.5% and ~25%, respectively (Fig. 5b,g). Therefore, in both datasets, the recorded neurons apparently increased task performance by a smaller amount than would have been possible if all their stimulus information was converted into choice, indicating that the PPC stimulus information is read out for behavior, but not optimally.

The theoretical analysis of the encoding-readout model (Fig. 3i) predicted that, when population activity is correlated, an enhanced-by-consistency readout leads to higher task performance than a consistency-independent readout with matched readout efficacy. To test this prediction on PPC data, we generated simulated choices by inputting real neural activity into the experimentally fit regression that incorporated across-time and across-neuron consistency. We also generated simulated choices using an alternative choice regression that included only the decoded stimulus predictor, regardless of its consistency (Fig. 5d-f,i,j). For fairness of comparison and to match the experimental data, the coefficients for this second choice regression were selected to yield the same readout efficacy as for the experimentally fit regression. The estimated contribution of the recorded neurons to task performance was higher with the experimentally fit choice regression that used consistency than with the consistency-independent choice regression matched in readout efficacy (Fig. 5f,j). These experimental findings, in agreement with model predictions, suggest that the enhanced-by-consistency readout is well suited for forming behavioral choices in the presence of information-limiting noise correlations, such as those found in the PPC.

In the sound localization experiments, we also had experimental data¹⁹ from the auditory cortex (AC). Relative to the PPC, in the AC we observed similar signal–noise angles, but weaker noise correlations, leading to a smaller information-limiting effect of correlations, and a much lower impact of consistency on the readout (Extended Data Fig. 7). Therefore, an enhanced-by-consistency readout may be more beneficial for PPC activity than for AC activity.

A model of enhanced-by-consistency information transmission

We developed a biophysical model for the downstream transmission of the PPC stimulus information to understand potential mechanisms for the behavioral benefit of information consistency across neurons and time. Our model was based on previous

observations that correlations in the presynaptic inputs to a neuron, either across neurons or time, elicit larger firing rates in postsynaptic neurons with a short integration time constant through a coincidence-detection mechanism^{16,29}. In our model (Fig. 6a), two presynaptic input spike trains, representing the summed inputs from two neuronal pools, were integrated by a postsynaptic ‘readout’ spiking neuron. The neural responses to two different stimuli were simulated (Fig. 6a). We assumed that the average response to the two stimuli was the same across the two input pools, leading to positive signal correlation, and we implemented positive noise correlations, both across input pools and across time (Fig. 6b,c,h). This ensured that noise and signal correlations were aligned and thus information-limiting. Therefore, and in agreement with our encoding models (Fig. 3), higher correlation strengths in the input pools, that is, enhanced across-pool synchrony and/or across-time correlations (Fig. 6b,c,h), more strongly limited the information contained in the inputs (Fig. 6d,i).

We then considered how information-limiting correlations in the inputs to a readout neuron with a short but realistic²⁹ integration time constant (~5–10 ms) affected information transmission. First, input correlations enhanced information transmission by increasing the average firing rate of the readout neuron in response to each stimulus (Fig. 6e,j). However, correlations also limited information transmission by increasing the variance of the readout’s firing activity (Fig. 6e,j). To quantify the trade-off between these factors, we measured the coefficient of variation (CV) of the readout activity and the gain of transmitted information (accuracy of stimulus decoded from the readout neuron’s firing). The coefficient of variation decreased, and the transmitted information increased, with increasing correlations (Fig. 6f,g,k,l). Thus, correlations in inputs to a neuron have advantages in enhancing the neuron’s output rate that outweigh their disadvantages in increasing the neuron’s output noise.

By systematically varying the model parameters, we demonstrated that input correlations enhance readout information from a postsynaptic neuron, even when decreasing input information, when the readout integration time constant is short enough so that the average amount of excitatory postsynaptic potentials received during an integration window is much smaller than the gap between the spiking threshold and resting potential of the readout neuron (Extended Data Fig. 8). In this regime, output firing is driven by input fluctuations. Correlated fluctuations on a short timescale increase the frequency with which the readout neuron reaches the firing threshold, thus enhancing the transmission of both neural activity and information (Supplementary Note 5 and Extended Data Fig. 8).

Fig. 6 | A biophysical model for the enhanced-by-consistency readout model. **a**, Schematic of the model. A leaky integrate-and-fire (LIF) readout neuron receives stimulus-modulated spike trains from two input pools. A linear stimulus classifier of the readout activity generates the transmitted output, *c*. **b**, Cross-correlograms of the two input spike trains for different values of the across-pool (left) and across-time (right) correlation parameters α and τ_c , respectively (mean input rate $R_{in}=2$ Hz). **c**, Schematic illustrating across-pool correlations between the input activity. τ_{obs} is the length of each simulated trial. **d**, Input stimulus-decoding accuracy as a function of the across-pool correlation strength α . **e**, Mean and standard deviation of the readout activity (normalized to their reference value in absence of correlations) as a function of α . **f**, Coefficient of variation (CV) of the readout activity as a function of α , normalized by its value in absence of input correlations. **g**, Gain in the stimulus information transmission from the input to the readout neuron (equation (9) and Supplementary Note 5) as a function of α . In **d-g**, $\tau_m=5$ ms, $R_{in}=2$ Hz, $\tau_c=100$ ms. All the quantities were computed on the input or output spike counts measured on time windows of length τ_{obs} . **h**, Schematic illustrating across-time correlations between the input activity at two different time points t_1 and t_2 . **i-l**, The quantities as in **d-g** showed similar trends when computed as a function of the across-time correlations τ_c . In **i-l** $\tau_m=5$ ms, $R_{in}=2$ Hz, $\alpha=0.9$. In **d-l**, data are presented as mean \pm s.e.m. over $n=20$ simulations. **m**, Time-lagged pairwise noise correlations computed separately on correctly transmitted and incorrectly transmitted simulated trials, across time lags of 500 ms (for 0-ms time lag $P=1.72 \times 10^{-17}$, $t=33.4$ (−1.65, 1.65); for 500-ms time lag $P=1.24 \times 10^{-95}$, $t=57.8$, (−1.65, 1.65); two-sided t-test, $n=200$ independent sets of equalized correct and error trials). **n**, Fraction of deviance explained for the transmitted output *c* (output of the stimulus classifier on readout neuron’s activity) for the enhanced-by-consistency and the consistency-independent readout regressions ($P=4.58 \times 10^{-23}$, $t=21.8$, (−1.69, 1.69), two-sided t-test, $n=20$ independent sets of equalized correct and error trials). **o**, Values of the coefficients of the enhanced-by-consistency readout regressions. The coefficients β_S , β_H and β_{I2} correspond to the stimulus-decoded regressor and the two consistency regressors, respectively (Methods). In **m-o**, box plots show the median (line) and first and third quartiles (box), and whiskers extend to ± 1.5 times the interquartile range; in simulations $\tau_m=5$ ms, $R_{in}=6$ Hz, $\alpha=0.9$, $\tau_c=500$ ms.

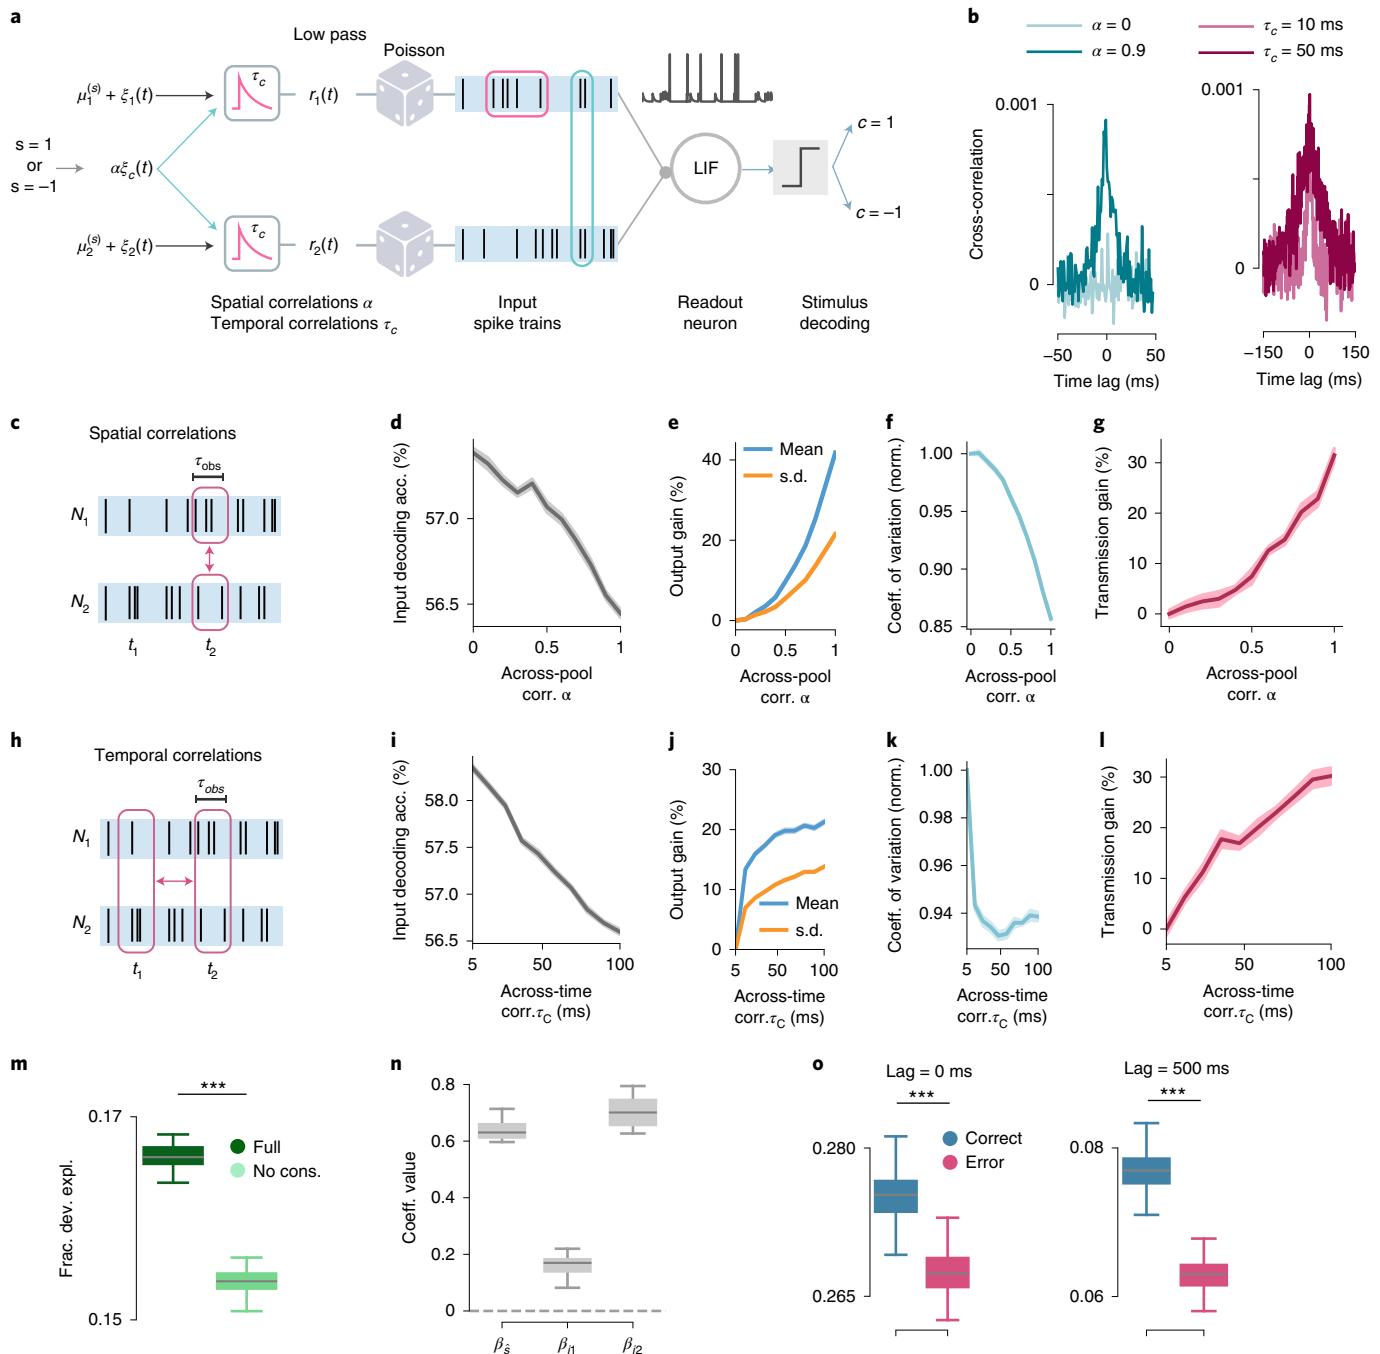
Importantly, for biophysical parameters consistent with the coincidence-detection regime, the model predicted key features of PPC data. We divided the model's simulated trials into correctly transmitted and incorrectly transmitted trials, namely trials in which the stimulus identity was correctly or incorrectly decoded from the activity of the readout neuron. We then analyzed the simulated data with the logistic readout regressions used for the analysis of PPC data. First, as in PPC data, the enhanced-by-consistency readout regression explained a larger fraction of the variance of the model choices (that is, the outcome of the stimulus-decoding algorithm applied on the readout activity) with respect to a consistency-independent readout (Fig. 6m). Also, the enhanced-by-consistency readout fitted on the model choices revealed that the transmission of stimulus information increased when the input activity carried consistent information across pools

(Fig. 6n). Second, correlations in the input spike trains were stronger for correctly transmitted trials than for incorrectly transmitted trials (Fig. 6o), as in the PPC data (Fig. 1b,g).

Thus, a coincidence-detection information transmission model suggests how the enhanced-by-consistency readout may benefit behavior. Correlations in the inputs to a neuron can enhance the transmission of stimulus information from a neuron's inputs to its output, even though these correlations limit the information contained in the inputs.

Discussion

Our results show that noise correlations limit information at the encoding stage, but they also enhance consistency in neural codes, which improves readout. The trade-off of these two effects defines the overall impact of correlations on task performance. Strikingly,



we found that noise correlations can enhance task performance despite limiting the information capacity of a neural population.

Much work has emphasized that the information-limiting effect of correlations in sensory areas may be a bottleneck for behavioral performance^{1,6,7}. A largely separate set of theoretical and biophysical work has alternatively proposed that correlations improve the propagation of neural activity^{13,14,16,17,30}. However, whether the advantages of correlations for signal propagation can overcome their information-limiting effect has not been fully clarified. Theoretical work on signal propagation has seldom specified whether the transmitted activity is informative¹⁷, and its connection to behavior remains unclear. Recent work has proposed that information-limiting across-neuron correlations may benefit information propagation in the presence of output nonlinearities¹⁷. Our models extend these results by identifying the biophysical conditions for which across-time and across-neuron correlations overcome their information-limiting effects by increasing the efficacy of information transmission. Experimental support for a role of correlations in facilitating the readout of population information to aid behavior has also been limited^{31,32}. Although a recent study has suggested a beneficial role of correlations by reporting higher correlation levels during correct behaviors³¹, these effects have not been reported when correlations limit information encoding. Remarkably, in PPC data and in the biophysical model presented here, the advantages of correlations for signal readout were large enough to compensate and overcome their negative encoding effects. Moreover, both our experimental and modeling results revealed a key computation underlying this effect: the amplification of the readout of stimulus information when neural activity is consistent across neurons or time.

Here we developed a formalism to address how the information-limiting effects of correlations on encoding and their benefits for signal readout intersect. Our approach provides a generally applicable framework to dissect the contribution of correlated neural activity to perceptual behaviors. We anticipate that this approach can be applied to different tasks and brain areas. Sensory and association cortices differ in the magnitude of their correlations, with higher correlations in association areas¹⁹. This difference could relate to the potential functions of each area, and our initial observations between the PPC and AC suggest that the best trade-off between the effect of correlations on encoding and readout may also vary across areas. In sensory cortices, a major function may be to encode rapidly changing and high-dimensional sensory features regardless of whether they are used for the immediate behavior at hand. In this case, weaker correlations may be advantageous to lessen information-limiting effects, and a readout sensitive to consistency for propagating the signals may be less critical. This view is compatible with reciprocal relationships between noise correlation levels and behavioral performance in sensory cortices¹. In contrast, because association cortices are closer to behavioral output, they may only need to encode a moderate amount of behaviorally relevant sensory information, but this information should have a strong impact on behavior. In these areas, higher correlations could be beneficial because the consequence of reducing encoded information is small, whereas the ramifications of failing to propagate signals to drive behavior is higher. Thus, in association areas, the best trade-off may involve some redundancy in the neural representation coupled with a readout mechanism that uses this redundancy to enhance signal propagation to inform choice, as we found here. We anticipate that the formalism described here will allow the design of causal tests of the actual readout used in the brain during perceptual discrimination tasks, such as with holographic perturbations³³.

Noise correlations can reflect interactions between cells, shared covariations due to common inputs, general fluctuations in behavioral state or network excitability, or variations with stimuli

within the same category². Previous work and our simulations (Supplementary Note 6 and Extended Data Fig. 9) show that positive information-limiting correlations (such as those observed in PPC) can be created by shared common inputs^{34–37} or stronger excitatory connectivity between neurons with similar stimulus tuning^{35,36}, features which have been reported in cortical circuits^{38,39}. Although our present data cannot disambiguate between these possibilities, this shared variability, regardless of its origin, acts as noise for decoding, because it cannot be reduced by integrating information over more cells or longer times, but it also helps signal propagation by generating more consistent neural representations. Thus, our conclusions are expected to hold regardless of the biophysical origins of the observed noise correlations.

Many studies of neural coding implicitly or explicitly assume that the readout of sensory information is optimal and interpret neural codes with higher sensory information as being more relevant for perception^{6,28,29}. Part of the reason is that the presence and shape of non-optimality are unknown. If the readout is not optimal, then neural codes with higher information are not necessarily the most relevant ones for perception. Our data suggest that stimulus information in population activity is not used optimally to produce accurate behavioral choices. Our work provides a measure of both the nature of readout non-optimality and its implication for the behavioral relevance of a neural code. Previous work has shown that even simple stimulus decoders of population activity trained sub-optimally to decode single-cell activity separately and then joined together can decode stimulus information accurately^{40–42}. Together with our results, this evidence suggests that correlations do not necessarily complicate the decoding of sensory information and may offer advantages for turning sensory information into appropriate behavioral choices.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at <https://doi.org/10.1038/s41593-021-00845-1>.

Received: 11 May 2020; Accepted: 24 March 2021;

Published online: 13 May 2021

References

- Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J. & Cohen, M. R. Learning and attention reveal a general relationship between population activity and behavior. *Science* **359**, 463–465 (2018).
- Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and neuronal population information. *Annu Rev Neurosci.* **39**, 237–256 (2016).
- Panzeri, S., Harvey, C. D., Piasini, E., Latham, P. E. & Fellin, T. Cracking the neural code for sensory perception by combining statistics, intervention and behavior. *Neuron* **93**, 491–507 (2017).
- Gawne, T. J. & Richmond, B. J. How independent are the messages carried by adjacent inferior temporal cortical neurons? *J. Neurosci.* **13**, 2758–2771 (1993).
- Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).
- Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).
- Moreno-Bote, R. et al. Information-limiting correlations. *Nat. Neurosci.* **17**, 1410–1417 (2014).
- Bartolo, R., Saunders, R. C., Mitz, A. R. & Averbeck, B. B. Information-limiting correlations in large neural populations. *J. Neurosci.* **40**, 1668–1678 (2020).
- Rumyantsev, O. I. et al. Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105 (2020).
- Gold, J. I. & Shadlen, M. N. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* **5**, 10–16 (2001).
- Zariwala, H. A., Kepes, A., Uchida, N., Hirokawa, J. & Mainen, Z. F. The limits of deliberation in a perceptual decision task. *Neuron* **78**, 339–351 (2013).

12. Mazurek, M. E. & Shadlen, M. N. Limits to the temporal fidelity of cortical spike rate signals. *Nat. Neurosci.* **5**, 463–471 (2002).
13. Diesmann, M., Gewaltig, M. O. & Aertsen, A. Stable propagation of synchronous spiking in cortical neural networks. *Nature* **402**, 529–533 (1999).
14. Zandvakili, A. & Kohn, A. Coordinated neuronal activity enhances corticocortical communication. *Neuron* **87**, 827–839 (2015).
15. Alonso, J. M., Usrey, W. M. & Reid, R. C. Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* **383**, 815–819 (1996).
16. Salinas, E. & Sejnowski, T. J. Correlated neuronal activity and the flow of neural information. *Nat. Rev. Neurosci.* **2**, 539–550 (2001).
17. Zylberberg, J., Pouget, A., Latham, P. E. & Shea-Brown, E. Robust information propagation through noisy neural circuits. *PLoS Comput. Biol.* **13**, e1005497 (2017).
18. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
19. Runyan, C. A., Piasini, E., Panzeri, S. & Harvey, C. D. Distinct timescales of population coding across cortex. *Nature* **548**, 92–96 (2017).
20. Hanks, T. D. et al. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220–223 (2015).
21. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* **19**, 1672–1681 (2016).
22. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).
23. Pho, G. N., Goard, M. J., Woodson, J., Crawford, B. & Sur, M. Task-dependent representations of stimulus and choice in mouse parietal cortex. *Nat. Commun.* **9**, 2596 (2018).
24. Panzeri, S., Schultz, S. R., Treves, A. & Rolls, E. T. Correlations and the encoding of information in the nervous system. *Proc. Biol. Sci.* **266**, 1001–1012 (1999).
25. Averbeck, B. B. & Lee, D. Effects of noise correlations on information encoding and decoding. *J. Neurophysiol.* **95**, 3633–3644 (2006).
26. Nogueira, R. et al. The effects of population tuning and trial-by-trial variability on information encoding and behavior. *J. Neurosci.* **40**, 1066–1083 (2020).
27. Romo, R., Hernandez, A., Zainos, A. & Salinas, E. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. *Neuron* **38**, 649–657 (2003).
28. Reich, D. S., Mechler, F. & Victor, J. D. Independent and redundant information in nearby cortical neurons. *Science* **294**, 2566–2568 (2001).
29. Koch, C., Rapp, M. & Segev, I. A brief history of time (constants). *Cereb. Cortex* **6**, 93–101 (1996).
30. Reyes, A. D. Synchrony-dependent propagation of firing rate in iteratively constructed networks in vitro. *Nat. Neurosci.* **6**, 593–599 (2003).
31. Shahidi, N., Andrei, A. R., Hu, M. & Dragoi, V. High-order coordination of cortical spiking activity modulates perceptual accuracy. *Nat. Neurosci.* **22**, 1148–1158 (2019).
32. Histed, M. H. & Maunsell, J. H. Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proc. Natl Acad. Sci. USA* **111**, E178–E187 (2014).
33. Emiliani, V., Cohen, A. E., Deisseroth, K. & Häusser, M. All-optical interrogation of neural circuits. *J. Neurosci.* **35**, 13917–13926 (2015).
34. Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
35. Ostojic, S., Brunel, N. & Hakim, V. How connectivity, background activity and synaptic properties shape the cross-correlation between spike trains. *J. Neurosci.* **29**, 10234–10253 (2009).
36. Rosenbaum, R., Smith, M. A., Kohn, A., Rubin, J. E. & Doiron, B. The spatial structure of correlated neuronal variability. *Nat. Neurosci.* **20**, 107–114 (2017).
37. de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K. & Reyes, A. Correlation between neural spike trains increases with firing rate. *Nature* **448**, 802–806 (2007).
38. Cossell, L. et al. Functional organization of excitatory synaptic strength in primary visual cortex. *Nature* **518**, 399–403 (2015).
39. Marshel, J. H. et al. Cortical layer-specific critical dynamics triggering perception. *Science* **365**, eaaw5202 (2019).
40. Pitkow, X., Liu, S., Angelaki, D. E., DeAngelis, G. C. & Pouget, A. How can single sensory neurons predict behavior? *Neuron* **87**, 411–423 (2015).
41. Nirenberg, S., Carcieri, S. M., Jacobs, A. L. & Latham, P. E. Retinal ganglion cells act largely as independent encoders. *Nature* **411**, 698–701 (2001).
42. Karpas, E.D., Maoz, O., Kiani, R. & Schneidman, E. Strongly correlated spatiotemporal encoding and simple decoding in the prefrontal cortex. Preprint at *bioRxiv* <https://doi.org/10.1101/693192> (2019).

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature America, Inc. 2021

Methods

No statistical methods were used to predetermine sample size of imaging experiments in mice. The experiments were not randomized. The investigators were not blinded to allocation during experiments and outcome assessment.

Subjects, behavioral task and two-photon imaging. This study represents an independent analysis of mouse calcium imaging experiments described previously^{19,21} and in publicly available data^{43,44}. A brief summary of the experimental procedures is provided here; full details are described in previous studies^{19,21}. All experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee.

Both experiments used a modified version of a previously described visual virtual reality system¹⁸. Head-restrained mice ran on a spherical treadmill, while images of a virtual maze were projected on a half-cylindrical screen. Forward/backward translation in the maze was controlled by treadmill changes in pitch. Rotation in the virtual environment was controlled by roll of the treadmill. The virtual maze was constructed using the Virtual Reality Mouse Engine (ViRMEn) in MATLAB¹⁵.

Sound localization task dataset. Imaging data were acquired from five male C57BL/6J mice (The Jackson Laboratory), aged 6–8 weeks at the initiation of behavioral task training. Imaging began 4–6 weeks after viral injection and continued for 4–12 weeks.

Mice ran down the stem of the virtual T-maze, while sound stimuli were delivered from eight possible locations (-90° , -60° , -30° , -15° , $+15^\circ$, $+30^\circ$, $+60^\circ$ and $+90^\circ$) using four electrostatic speakers positioned in a semicircular array, centered on the mouse's head. The sound stimulus was activated when the mouse passed an invisible spatial threshold at ~ 10 cm into the T-stem. The stimulus was repeated after a 100-ms gap; repeats continued until the mouse reached the T-stem. Task difficulty was modulated by the direction of the incoming stimulus. To receive a reward (4 μ l water), mice had to judge the location of sound stimuli to be either on the left or right, and to report their decisions by turning left or right at the T-intersection. A 'reward tone' was played as the water reward was delivered on correct trials (when the mouse had reached ~ 10 cm into the correct arm of the T-maze), and a 'no-reward tone' was played when the mouse reached ~ 10 cm into the incorrect arm on error trials. The intertrial interval was 3 s on correct trials and 5 s on error trials. Mice performed ~ 200 trials (range, 125–251) per session.

Imaging was performed on alternating days from the AC and PPC on the left hemisphere (PPC centered at 2 mm posterior and 1.75 mm lateral to bregma; AC centered at 3.0 mm posterior and 4.3 mm lateral to bregma). In each session, ~ 50 neurons (range, 37–69) were simultaneously imaged using a two-photon microscope (Sutter MOM) operating at a 15.6-Hz frame rate and at a resolution of 256×64 pixels ($\sim 250 \mu\text{m} \times 100 \mu\text{m}$). ScanImage (version 3, Vidrio Technologies) was used to control the microscope. Imaging data were acquired at depths of between 150 and 300 μm , corresponding to layers 2/3. Seven AC and seven PPC fields of view from five mice were analyzed.

Evidence accumulation task dataset. Imaging data were acquired from five male C57BL/6J mice (The Jackson Laboratory), aged 8–10 weeks at the initiation of behavioral task training. Imaging began at least 4 weeks after viral injection and was continued for up to 12 weeks.

Mice run down the stem of a virtual T-maze with predominantly gray walls, encountering six visual cues (white wall segments with black dots) at fixed locations. Each cue appeared on either the left or right wall, with only one cue visible at a time. To receive a reward (4 μ l 10% sweetened condensed milk), mice had to determine whether more cues were presented on the left or the right and, after a short stretch of maze without additional cues (90 cm), turn at the T-intersection toward the direction with more cues (left for 6-0, 5-1 and 4-2 trials; right for 2-4, 1-5 and 0-6 trials). Task difficulty was modulated by varying the difference between the number of left and right cues (net evidence). The sequence of cues was determined randomly for each trial of a given net evidence. On trials having zero net evidence (3-3 trials), a random location was rewarded. Intertrial interval duration was 2 s for correct choice and 4 s for incorrect choice. Mice performed ~ 300 trials (range, 231–414) in a typical session.

Imaging data were acquired from the left PPC (PPC centered at 2 mm posterior and 1.75 mm lateral to bregma). In a given session, ~ 350 neurons (range, 188–648) were simultaneously imaged using a custom-built two-photon microscope operating at a ~ 30 -Hz frame rate and at a resolution of 512×512 pixels ($\sim 700 \mu\text{m} \times 700 \mu\text{m}$). The microscope was controlled by ScanImage (version 5; Vidrio Technologies). Imaging data were acquired at depths of between 100 and 200 μm below the dura. Eleven fields of view from five mice were analyzed.

Imaging data processing. After motion correction⁴⁶, correlations in fluorescence time series between pixels within $\sim 60 \mu\text{m}$ were calculated. Fluorescence sources (putative cells) were identified by applying a continuous-valued eigenvector-based approximation of the normalized cuts objective to the correlation matrix, followed by k -means clustering segmentation^{19,21}. To estimate potential neuropil contamination, the cell body fluorescence signal was regressed against the signal from surrounding pixels during the imaging frames when the cell of interest

was not active, and then neuropil contamination was removed during the $\Delta F/F$ calculation by subtracting a scaled version of the neuropil signal from the cell body signal. All fluorescence traces were deconvolved to estimate the relative spike rate in each imaging frame⁴⁷. The deconvolution alleviated the possible artificial lengthening of timescales of across-time correlations due to slow calcium transients. The timescales of single neuron activity of the deconvolved signal were ~ 200 ms^{19,21}, much shorter than the timescales of across-time correlations (~ 1 s), and were shorter in the AC than in the PPC¹⁹, suggesting that the deconvolution was effective at preventing major artificial inflations of across-time correlations timescales.

Data inclusion and task epoch selection for encoding and readout analyses.

Sound localization task dataset. For the analysis of across-time correlations in the PPC, population activity data were temporally aligned to the imaging time frame of the turn, defined as the frame in which the mouse entered the short arm of the maze. Since it is reasonable to assume that the animal computes its choice after the stimulus presentation but before the turn, the analysis focused on the 39 frames preceding the turn frame (this number of frames was chosen because it covered the maximum portion of the pre-turn period that was commonly available across all recording sessions). One of the seven PPC recording sessions used in our previous published work¹⁹ was excluded due to the large unbalance of left/right stimuli that were presented to the mouse across trials in that session, which would result in too few trials available for our analyses.

For the analysis of across-time correlations in the AC, population activity data were temporally aligned to the imaging time frame of the first auditory stimulus presentation, and the analysis focused on the 50 frames after that frame (this number of frames was chosen because it covered the maximum portion of the poststimulus period that was commonly available across all recording sessions). AC neural data aligned to the turn did not encode a sufficient amount of stimulus information for following analyses. One of the seven AC sessions used in our previous published work¹⁷ was excluded due to the large unbalance of left/right stimuli that were presented to the mouse across trials in that session.

Evidence accumulation task dataset. PPC population activity data were first grouped into spatial bins (3.75 cm per bin) covering the whole T-maze (long and short arm) by averaging population activity first in each bin (two or three imaging frames per bin for each trial) and then over epochs of four spatial bins each (about 200 ms). We used the same ten epochs defined in ref.²¹. We plotted results of population activity data recorded in the early delay and late delay epochs. During the delay epochs, the cue presentation was completed but the animal had not yet committed to a final turn (these epochs correspond to the four spatial bins beginning 15 and 37.5 cm, respectively, after the offset of the final cue). Therefore, it is reasonable to assume that the animal's decision is formed in these epochs. All 11 sessions in the original work¹⁹ were used. We did not use trials with zero net evidence (<10% trials in 2/11 sessions).

Selectivity of single cells to stimulus category. For Fig. 1c,h, we computed the selectivity of single cells to stimulus category. For the sound localization task, stimulus category corresponds to the direction of incoming auditory stimuli. Each stimulus category comprises four different sound locations (left: -90° , -60° , -30° and -15° ; right: $+15^\circ$, $+30^\circ$, $+60^\circ$ and $+90^\circ$). For the evidence accumulation task, stimulus category corresponds to the side of the maze where the majority of the visual cues were presented (left: 6-0, 5-1 and 4-2 trials; right: 2-4, 1-5 and 0-6 trials).

The selectivity index¹⁸ (SI) was quantified according to equation (1):

$$\text{SI} = \frac{\text{mean } \Delta F/F_{\text{right trials}} - \text{mean } \Delta F/F_{\text{left trials}}}{\text{mean } \Delta F/F_{\text{right trials}} + \text{mean } \Delta F/F_{\text{left trials}}} \quad (1)$$

Cells with an SI greater or smaller than zero were classified as right-preferring or left-preferring cells, respectively.

Pairwise noise correlations. For each neuron pair and time-points pair in the trial epoch selected for analyses, we quantified across-time pairwise correlations as the Pearson correlation between the activity of neuron 1 at time t_1 and the activity of neuron 2 at time t_2 across trials with the same stimulus category. Results were averaged across all neuron pairs, all time-point pairs with the same time lag, across stimuli and across trials subsamples. We quantified across-neuron pairwise correlations as the Pearson correlation between neuron pairs recorded in a single session, across trials sharing the same stimulus category. Results were averaged across stimuli, across trial subsamples and across early delay and late delay epochs. We quantified noise correlations separately for correct and error trials. To control for differences in trial numerosity, we subsampled trials to equalize the number of correct and error trials in each recorded session. Results were averaged over ten instantiations of random subsampling.

Population-wise noise correlations. We quantified the across-neuron population-wise correlations by performing principal-component analysis (PCA) on the population response to all trials sharing the same stimulus category. The

population-wise noise correlation, ν , was defined as the fraction of variance explained by the first principal component of the whole-population activity (concatenated population activity of the two pools for across-neuron correlations; concatenated activity of the two considered time points for across-time correlations). For across-neuron correlations, results were first averaged across stimuli, then across trials subsamples and eventually pooled across early delay and late delay epochs. For across-time correlations, results were averaged first across all time-point pairs sharing the same lag between each other, then across stimuli and finally across trial subsamples. We quantified across-time and across-neuron population-wise noise correlations¹ separately for correct and error trials. To control for differences in trial numerosity, we randomly subsampled trials to equalize the number of correct and error trials in each recorded session. Results were averaged over ten instantiations of random subsampling. Population-wise noise correlations, when computed on a small number of trials, suffered from finite-sampling bias. Since we equalized trials between correct and error choices, the comparison was not affected. However, to map the experimental data to the model, we corrected for the finite-sampling bias. We estimated the bias by computing the population-wise noise correlation index for randomly selected subsamples of trials with progressively increasing size (from five to the maximum number of available trials) and then using polynomial extrapolation.

Analysis of stimulus encoding and consistency. For encoding and consistency analyses, we considered information about stimulus category. Information about stimulus category carried by population activity was extracted by decoding the most likely stimulus category presented to the animal in each trial using a C-Support Vector Machine (C-SVM) classifier with a linear basis function kernel¹⁸, implemented using the libsvm library¹⁹. For each imaging session, we first subsampled trials randomly such that the left/right stimulus categories were equally represented in the data (sound localization task dataset: no more than 13% of removed trials per session; evidence accumulation task dataset: no more than 15% of removed trials per session). Then, we randomly split the remaining trials ten times into 50/50 training/testing sets, such that left and right stimulus categories were equally represented in both training and testing sets. For each trial split, we trained the C-SVM on the training set and we tested on the test set, which was left out of the fitting procedure. The regularization hyperparameter (C) was selected by maximizing the threefold cross-validated decoding accuracy in the training set. For the analyses that required computing a posterior probability of the decoded stimulus given the observed population activity, we used Platt scaling to calibrate posterior probabilities on the binary outputs of the C-SVM⁵⁰.

Again, using the libsvm library¹⁹, we decoded stimulus category considering also nonlinear classifiers: we decoded the most likely stimulus category using a C-SVM with radial and quadratic basis function kernels.

For the across-time correlations analysis, for any considered pair of time points, we defined the activity in trial as consistent if the stimulus decoded from the population activity at each of the two time points coincided. For the across-neuron correlations analysis, we first split the neuronal population recorded in each session into two randomly selected, equally sized pools of neurons. A total of 100 random splits were performed. For each random split, we defined the activity in trial as consistent if the stimulus decoded from the population activity of each individual pool coincided.

Quantifying the angle between the signal and noise axes. We quantified the angle, γ ($\in [0, \pi/2]$), between the direction of maximum stimulus variation (signal correlations axis) and the direction of maximum noise variation (noise correlations axis)^{25,26} in the neural population response space. The signal correlations axis was defined as the vector connecting the mean responses to the two stimuli. The noise correlations axis was computed as the direction of the first principal component obtained by applying PCA to all single-trial responses at fixed stimulus category. The angle between signal correlation and noise correlation axes was computed separately for each stimulus category ($\gamma_{s=-1}, \gamma_{s=1}$) and then averaged according to equation (2):

$$\gamma = \arccos \left(\sqrt{\cos^2 \gamma_{s=-1} + \cos^2 \gamma_{s=1}} \right) \quad (2)$$

This weighted average of the stimulus-specific angles facilitates comparisons between data and model (Supplementary Note 4).

The computation of γ uses the population's covariance matrix. Since in the accumulation evidence task the dimensionality of the dataset (~350 neurons) was larger than the number of trials per session (~200), we first performed a PCA to keep only those components that explained 95% of the total variance (the dimensionality of the dataset was reduced to 59 ± 18 components (mean \pm s.e.m.) across sessions). This did not substantially change the values of the angles (we obtained a median value of the signal–noise angle of 0.21π in Fig. 1j with the 95% variance cutoff, whereas we would have obtained a value of 0.25π had we used all neurons without variance cutoff, with both values inside the information-limiting region). However, we used the variance cutoff because it led to better stability of individual results when removing random fractions (10%, 20%,...) of data.

For across-time correlations, γ was computed in the space defined by the concatenated population activity at the two time points considered for the analysis.

For across-neuron correlations, γ was computed in the full-dimensional space defined by the population responses, and did not depend on the random split in two pools.

Mathematical model of encoding and readout with two N -dimensional neural features. We developed a simple model of how two N -dimensional neural activity features, r_1 and r_2 (each representing the firing rates of two different pools of N neurons each for across-neuron correlations, or the population activity of the same pool of N neurons at two different times for across-time correlations), encode information about a binary stimulus, and how this information is read out to inform choice in a simulated stimulus discrimination task.

Neural encoding (stimulus–response) model. The encoding (stimulus–response) models describes the neural activity of the two N -dimensional features (r_1 and r_2) in response to two stimuli ($s = -1$ and $s = 1$).

We chose a simple model accounting for the observation that average pairwise correlations in the PPC neural population were positive (Fig. 1b,g). Distributions of stimulus-specific neural responses $r_k^{(s)}$ for each N -dimensional feature were modeled as N -dimensional multivariate Gaussians with mean $\mu_k^{(s)}$ and stimulus-independent covariance Σ_k given by equation (3):

$$\begin{aligned} r_k^{(s)} &\sim \mathcal{N} \left(\mu_k^{(s)}, \Sigma_k \right) \\ \mu_k^{(s)} &= \text{sign}(s) * d * w_{\text{signal},k} \\ \Sigma_k &= \begin{bmatrix} c_{1,1} & c_{1,2} & \cdots \\ c_{2,1} & c_{2,2} & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}, c_{i,j} = \begin{cases} \sigma^2, & \text{if } i = j \\ \rho_{\text{within}} \sigma^2, & \text{if } i \neq j \end{cases} \end{aligned} \quad (3)$$

where s indexes the stimulus category and k the neural feature ($k = 1, 2$); ρ_{within} parametrizes the strength of correlation between neurons within the same feature; and $w_{\text{signal},k}$ represents the signal correlation direction in the N -dimensional space of each feature. For simplicity, we assumed equal variance σ for each neuron i and equal covariances Σ_k for the two pools $k = 1, 2$. The signal–noise angle, γ , was the same across stimuli in this model.

The joint population activity was simulated as a $2N$ -dimensional multivariate Gaussian with mean and covariance given by equation (4):

$$\begin{aligned} \mu^{(s)} &= \text{sign}(s) * d * \begin{bmatrix} w_{\text{signal},1} \\ w_{\text{signal},2} \end{bmatrix} = \text{sign}(s) * \hat{d} * w_{\text{signal}} \\ \Sigma &= \begin{bmatrix} \Sigma_1 & \rho \sigma^2 J_{N \times N} \\ \rho \sigma^2 J_{N \times N} & \Sigma_2 \end{bmatrix} \end{aligned} \quad (4)$$

w_{signal} indicates the normalized signal correlation direction in the $2N$ -dimensional space, $J_{N \times N}$ indicates the matrix with all elements equal to one and ρ determines the correlation between the two neural features. This correlation matrix allowed us to mimic the effect of the real-data shuffling procedure of selectively removing correlations between the two N -dimensional features, while keeping intact single-feature correlation structures, by simply setting ρ to zero. For simplicity we also set $\rho = \rho_{\text{within}}$ for correlated activity. The means of the two response distributions ($\mu^{(s=1)}$ and $\mu^{(s=-1)}$) were symmetrically located around the origin of the $2N$ -dimensional space, at distance \hat{d} . Together, the parameters \hat{d} and σ control the overlap between the two stimulus-specific response distributions.

The first eigenvector of the covariance matrix, representing the direction of the noise correlation axis, is given by $w_{\text{noise}} = (1, 1, \dots, 1) / \sqrt{2N}$, with eigenvalue $\lambda_1 = (N\rho + 1 - \rho)\sigma^2$. The other $2N - 1$ eigenvectors form an orthonormal basis with the vector w_{noise} and the $2N - 1$ corresponding eigenvalues are given by $\lambda_i = (1 - \rho)\sigma^2$. The value of the largest eigenvalue, λ_1 , normalized by the sum of all eigenvalues (total response variance) yields the population-wise noise correlation index ν as a function of the pairwise correlation index ρ , according to equation (5):

$$\nu = \frac{\lambda_1}{\sum_i \lambda_i} = \frac{1 + \rho(2N - 1)}{2N}. \quad (5)$$

We denote with

$$\gamma = \arccos \left(\frac{w_{\text{signal}} \cdot w_{\text{noise}}}{\|w_{\text{signal}}\| \|w_{\text{noise}}\|} \right) \quad (6)$$

the angle between the signal correlations and the noise correlation axes in the $2N$ -dimensional space. The signal correlation axis orientation, w_{signal} , was randomly sampled across all vectors satisfying equation (6) for the fixed γ and w_{noise} .

We quantified the amount of stimulus information carried by the simulated responses (r_1 and r_2) as the accuracy of a linear decoder of stimulus identity applied to the responses. We then applied the same classifier to the responses r_1 and r_2 of each pool separately. For the simulations in Fig. 2 and Extended Data Fig. 3,

we set $N=1$, $\hat{d} = \sqrt{0.02}$, $\sigma=0.2$ and $\rho=0.8$, and we performed 200 simulations with 5,000 trials per stimulus (Fig. 2) or ten simulations with 50,000 trials per stimulus (Extended Data Fig. 3). For the simulations in Fig. 3 and Extended Data Fig. 4, we set $N=20$, 10 , $\hat{d}=0.15$ and $\sigma=0.2$ (consistent with the value found for both experimental datasets; across-time: $\sigma=0.195 \pm 0.005$ (mean \pm s.e.m.); across $n_t=39$ time points and $n=6$ sessions; across-neuron: $\sigma=0.230 \pm 0.012$ (mean \pm s.e.m.) across early delay and late delay epochs and $n=11$ datasets), and we performed 100 simulations with 300,000 trials per stimulus each.

Model of choice generation in a simulated discrimination task. We simulated the process of generating a binary choice in each trial from neural activity through a logistic regression readout model given by equation (7):

$$\text{logit}(p(c=1|x)) = \beta_0 + \beta_s \hat{s} + \frac{\beta_{11}}{2} (\hat{s}+1) \text{con} + \frac{\beta_{12}}{2} (\hat{s}-1) \text{con} \quad (7)$$

where \hat{s} ($\hat{s}=-1$ and $\hat{s}=1$) indicates the stimulus decoded from the concatenated activity of the two N -dimensional neural features; ‘con’ is a ‘consistency’ binary variable that is equal to 1 if the stimuli decoded individually from each neural feature are the same, and 0 otherwise; x indicates the entire set of predictors (s, con).

The model coefficients $\beta_0, \beta_s, \beta_{11}$ and β_{12} control the relative impact of the different model predictors on the simulated choice. The values for the model coefficients were set as follows: We first defined a consistency modulation index η , ranging from 0 to 1, to control the relative strength of neural consistency in the readout. We then derived the readout efficacy, which we defined as the probability of conversion from \hat{s} to c , for each of the four possible combinations of predictors values, from the modulation index η and a reference readout efficacy $\alpha(\hat{s})$ according to equation (8):

$$p(c=\hat{s}|[\hat{s}, \text{con}]) = \begin{cases} \alpha(\hat{s}) + \eta(1 - \alpha(\hat{s})), & \text{con} = 1 \\ \alpha(\hat{s}) - \eta(\alpha(\hat{s}) - 0.5), & \text{con} = 0 \end{cases} \quad (8)$$

where $\alpha(\hat{s})$ takes values between 0.5 and 1. For the simulations in Figs. 2 and 3 and Extended Data Figs. 2–4, we arbitrarily set $\alpha(0)=\alpha(1)=0.75$. Given the readout efficacy values from equation (8), we used equation (7) to compute the model coefficients corresponding to the chosen modulation index η .

Logistic regression of the mouse’s choice. To study how features of recorded neural population activity related to the mouse’s choices, we fitted a logistic regression of the choice (left/right turns) made in each trial to the recorded neural activity.

Our readout model explicitly focuses on the part of the choice signal in neural population activity that relates to the encoded stimulus information (as opposed to the part of the choice signal that is independent of the stimulus information). It differs from other quantifications of choice signals (for example, choice probability^{51,52}) that use either ‘zero-signal’ trials containing no sensory evidence or pooling data across stimulus levels after corrections to remove the stimulus modulation, to infer specifically choice signals in neural activity beyond stimulus-related modulation. Our model also focuses on how the correlation-induced consistency of neural information affects choices, and differs from other models¹⁰ by focusing only on how the total sensory evidence in neural activity influences choices.

Our logistic regression readout model was implemented as follows: For each trial, we considered the choice c made by the mouse ($c=1$: left; $c=-1$: right), the presented stimulus s ($s=1$: left; $s=-1$: right) and the neural population activity for each pair of time points (across-time correlations analysis) or for each pair of randomly selected neuronal pools (across-neuron correlations analysis). For each session, trial split and pair of time points or neuronal pools, we fitted the logistic regression given by equation (9):

$$\text{logit}(p(c=\text{left}|x)) = \beta_0 + \beta_s s + \beta_s \hat{s} + \frac{\beta_{11}}{2} (\hat{s}+1) \text{con} + \frac{\beta_{12}}{2} (\hat{s}-1) \text{con} \quad (9)$$

where \hat{s} ($\hat{s}=1$: left; $\hat{s}=-1$: right) represents the stimulus decoded from the concatenated activity of two time points or neuronal pools; ‘con’ is a binary variable that is equal to 1 if the stimuli decoded individually from each time point or neuronal pool are the same, and 0 otherwise; x indicates the combination of neural (s, con) and nonneural (s) predictors.

Logistic regression fitting was implemented using the statsmodel Python module⁵³. The logistic regression was fit on the testing set using L1-regularized maximum likelihood. The regularization hyperparameter (λ) was selected by maximizing threefold cross-validated fraction of deviance explained.

For control analyses, we fitted the mouse’s choices with more complex choice regressions that included other predictors on top of those described in equation (9).

To discern the genuine role of across-time neural consistency (con) in explaining the mouse’s choices from that of across-time behavioral consistency (con_b), we fitted a logistic regression model including additional behavioral consistency-dependent predictors according to equation (10):

$$\begin{aligned} \text{logit}(p(c=\text{left}|x)) = & \beta_0 + \beta_s s + \beta_s \hat{s} + \frac{\beta_{11}}{2} (\hat{s}+1) \text{con} + \frac{\beta_{12}}{2} (\hat{s}-1) \text{con} \\ & + \frac{\beta_{13}}{2} (\hat{s}+1) \text{con}_b + \frac{\beta_{14}}{2} (\hat{s}-1) \text{con}_b \end{aligned} \quad (10)$$

We performed this control analysis for three behavioral parameters of interest that were measured during the experiments: the lateral running velocity, lateral position and view angle of the mouse in the virtual environment (Extended Data Fig. 6). Two values of lateral running velocity or lateral position at two different time points were defined to be consistent whenever their sign was the same; two values of view angle at two different time points were defined to be consistent whenever they were both higher or both lower than 90°.

We also fitted mouse’s choices with a more sophisticated logistic regression model where the discrete binary variable \hat{s} in equation (9) was replaced with the continuous value of the decoder stimulus posterior probability $p(s=\text{left}|r)$. Because this model also accounts for the magnitude of stimulus information and not only for the identity of the decoded stimulus, we used it to account for possible confounders due to differences in overall stimulus information between consistent and inconsistent trials (Extended Data Fig. 5).

Predictive performance was quantified as the fraction of deviance explained (FDE), evaluated with threefold cross validation. For each fold, we computed the log likelihood l of the test data given the values of the β coefficients of the training data. To calculate a reference null value for the log likelihood, we computed the log likelihood, l_0 , of the test data given the value of the coefficient β_0 of an intercept-only regression fit on the training data. The FDE was then defined according to equation (11):

$$\text{FDE} = 1 - l/l_0 \quad (11)$$

Estimating the impact of across-time and across-neuron correlations on task performance. To estimate the impact of across-time and across-neuron correlations on mouse’s task performance, we generated synthetic choices using the experimentally fit choice regression of equation (9). As input to the regression, we provided predictors extracted from either the real recorded neural data, which included across-time and across-neuron correlations, or hypothetical neural data whose correlations were disrupted by shuffling. We used this analytical approach because current experimental methods cannot remove correlations during task performance, and thus we had to estimate effects with post hoc removal of correlations.

Task performance, $p(c=s)$, the probability that the choice c matches the presented stimulus s , was estimated by computing (using the choice’s logistic regression) the probabilities $p(x)$ of all possible combinations of predictors values x , multiplying them with the corresponding readout probabilities $p(c=s|x)$ obtained from the logistic choice regression, and then summing over x according to equation (12):

$$p(c=s) = \sum_{x \in X} p(c=s|x) * p(x). \quad (12)$$

The same readout probabilities, $p(c=s|x)$, were used for the computation of task performance from both real and shuffled neural data.

We isolated the part of task performance that can be attributed to the (correlated or shuffled) neural activity by subtracting from the total estimated task performance a baseline nonneural estimated task performance. The latter was computed by applying equation (12) after shuffling the values of neural predictors across trials, while keeping the relationship between nonneural predictors and mouse’s choices fixed.

Computation of readout efficacy of the transformation from stimulus information to choice. We termed ‘readout efficacy’ ($p(c=\hat{s})$) the probability that in a given trial the choice c matched the stimulus \hat{s} decoded by neural activity. We computed this probability according to equation (13):

$$p(c=\hat{s}) = \sum_{x \in X} p(c=\hat{s}|x) * p(x), \quad (13)$$

where x represents the set of all possible combinations of predictors values, and $p(c=\hat{s}|x)$ are obtained from the logistic choice regression.

To generate the readout maps in Fig. 4d,h, we computed, separately for consistent and inconsistent trials, readout efficacy as the deviation from the average probability of choice being left or right when the presented stimulus is left or right, according to equation (14):

$$\Delta p(c=\hat{s}|\text{con}) = \sum_{x \in X} \Delta p(c=\hat{s}|[x, \text{con}]) * p(x|\text{con}). \quad (14)$$

In equation (14), x represents the set of all possible combinations of $[s, \hat{s}]$ predictors values and $\Delta p(c=\hat{s}|[x, \text{con}]) = p(c=\hat{s}|[s, \hat{s}, \text{con}]) - p(c=\hat{s}|s)$.

Matching enhanced-by-consistency and consistency-independent readouts in terms of efficacy. To quantify the impact on task performance of the

enhanced-by-consistency experimentally measured readout, we compared the task performance predicted by the experimentally fit choice regression to the one predicted by a consistency-independent choice regression according to equation (15):

$$\text{logit}(p(c = \text{left}|\mathbf{x})) = \beta'_0 + \beta'_s s + \beta'_{\hat{s}} \hat{s}. \quad (15)$$

For a fair comparison, the values of the coefficients β'_0 , β'_s and $\beta'_{\hat{s}}$ were chosen so that the two readouts were matched in terms of readout efficacy (equation (13)). To compute the values of the coefficient of the consistency-independent choice regression in equation (15), we imposed the following conditions according to equations (16a–c):

$$\begin{cases} p_{\text{cons.-indep.}}(c = -1, \hat{s} = -1) = p_{\text{enhanced-by-cons.}}(c = -1, \hat{s} = -1) \\ p_{\text{cons.-indep.}}(c = 1, \hat{s} = 1) = p_{\text{enhanced-by-cons.}}(c = 1, \hat{s} = 1) \\ \beta'_s = \beta_s \end{cases} \quad (16\text{a-c})$$

and then plugged equations (9) and (15) into equation (16a,b), and solved for β'_0 and β'_s .

Shuffling procedure to disrupt across-time or across-neuron noise correlations. Across-time correlations between population vectors at different time points were removed by shuffling trial identities independently for each of the two population vectors within trials with the same stimulus category. With this procedure, across-time signal correlations were maintained, while across-time noise correlations were disrupted. Single-cell autocorrelations were also disrupted. Across-neuron correlations between two neuronal pools were disrupted by shuffling trial identities independently for each pool within trials with the same stimulus category. With this procedure signal correlations were maintained for all pairs of neurons, noise correlations between neurons pairs in two different pools were disrupted, and noise correlations of neuron pairs within the same pool were maintained. Shuffling was performed separately for the training and the testing set.

Biophysical model of consistency-modulated information transmission. We modeled biophysical signal propagation using an integrate-and-fire model neuron that receives inputs from a population of neurons exhibiting correlations across neurons or time.

Our model consisted of an input population, corresponding to PPC population activity in our data, projecting in a feed-forward manner to a single output (or ‘readout’) neuron. We simulated responses of the input population to two different external stimuli (corresponding to the stimulus categories in our data). Stimuli $s = -1$ and $s = +1$ had the lowest and highest mean rates, respectively. We simulated multiple trials for each stimulus. In each simulated trial, the activity of the readout neuron was decoded by an optimal linear classifier (readout neuron activity lower or higher than the optimal decoding boundary was decoded as stimulus -1 or $+1$, respectively), and the outcome c of the decoding algorithm represented the result of the information transmission process through the synapses (which can be compared to the behavioral choice in our neural data) in each trial. We modeled the readout neuron as an input-driven leaky integrate-and-fire neuron with dynamics given by equation (17):

$$\tau_m \frac{dV}{dt} = -V + V_r + w \sum_{n=1}^2 \sum_k \delta(t - t_k^{(n)}), \quad (17)$$

where τ_m represents the membrane time constant of the neuron and V_r is the resting value of its membrane potential V . The rightmost term of equation (17) describes the external drive to the readout neuron coming from two input units (corresponding to two different neural pools). $t_k^{(n)}$ denotes the set of spike times of each input unit. For each input spike received by the output unit, the output membrane voltage V instantaneously increases by a fixed amount w .

To examine the effect of across-pool correlations (analogous to the across-neuron correlations in PPC data) and across-time correlations between the two input units on the activity of the output neuron, we generated correlated input spike trains $t_k^{(1)}$ and $t_k^{(2)}$ as follows: First, we generated stochastic input firing rates for the two pools $r_1^{(s)}(t)$ and $r_2^{(s)}(t)$ with across-pool correlations by modulating the amount of shared noise ξ_c between the two units according to equation (18):

$$\begin{aligned} r_1^{(s)}(t; \alpha) &= \mu_1^{(s)} + \sigma \Delta \mu_1 (\alpha \xi_C(t) + \sqrt{1 - \alpha^2} \xi_1(t)) \\ r_2^{(s)}(t; \alpha) &= \mu_2^{(s)} + \sigma \Delta \mu_2 (\alpha \xi_C(t) + \sqrt{1 - \alpha^2} \xi_2(t)), \end{aligned} \quad (18)$$

where $\mu_i^{(s)}$ is the mean activity of input pool i in response to stimulus s ; $\Delta \mu_i = \mu_i^{(s=1)} - \mu_i^{(s=-1)}$ is proportional to the derivative of the mean activity with respect to the stimulus; and σ equally modulates the variability of the input units. The values of ξ_1 and ξ_2 (private noise) and ξ_C (shared noise) were independently drawn from Gaussian distributions with zero means and unit variance. The parameter α , ranging from 0 (fully uncorrelated activity) to 1 (fully correlated activity) modulates the amount of shared variability between the two input units.

Equation (18) generates correlations aligned with the derivative of the mean activity with respect to the stimulus $\Delta \mu$, which are therefore information-limiting⁷. Additionally, equation (18) ensured that varying α changed only the amount of across-pool correlation but not the variance of the activity of each individual input unit.

We created across-time input correlations by filtering the input activity using a low-pass filter with time constant τ_C given by equation (19):

$$r_i^{(s)}(t; \alpha, \tau_C) = \int_0^\infty e^{-(t-u)/\tau_C} r_i^{(s)}(u; \alpha) du. \quad (19)$$

We then generated the spike trains of the input units as inhomogeneous Poisson processes with firing rates given by $r_1^{(s)}(t; \alpha, \tau_C)$ and $r_2^{(s)}(t; \alpha, \tau_C)$.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The sound localization task data that support the findings of the current study can be downloaded at https://gin.g-node.org/MMoroni/PPC_AC_2p_sound_localization/ (ref. ⁴⁴).

The evidence accumulation task data that support the findings of the current study can be downloaded at https://gin.g-node.org/MMoroni/PPC_2p_evidence_accumulation/ (ref. ⁴³).

Code availability

The code for the biophysical information transmission model (Fig. 6) is available for download at <https://github.com/gbondonelli/BiophysicalReadout/>.

The code for the encoding and readout model (Figs. 2 and 3) is available for download at https://github.com/moni90/encoding_readout_model/.

The code for data analysis is available from the corresponding authors upon reasonable request.

References

43. Morcos, A. S. et al. Dataset of ‘History-dependent variability in population dynamics during evidence accumulation in cortex’. *G-Node* <https://doi.gin.g-node.org/10.12751/g-node.g1xyem/> (2021).
44. Runyan, C. A. et al. Dataset of ‘Distinct timescales of population coding across cortex’ *G-Node* <https://doi.org/10.12751/g-node.tqbad8> (2021)
45. Aronov, D. & Tank, D. W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
46. Greenberg, D. S. & Kerr, J. N. Automated correction of fast motion artifacts for two-photon imaging of awake animals. *J. Neurosci. Methods* **176**, 1–15 (2009).
47. Vogelstein, J. T. et al. Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.* **104**, 3691–3704 (2010).
48. Boser, B. E., Guyon, I. M. & Vapnik, V. N. A training algorithm for optimal margin classifiers. in *Fifth Annual Workshop on Computational Learning Theory* (ed. Haussler, D.) 144–152 (ACM Press, 1992).
49. Chang, C. C. L. & Merz, C. J. LIBSVM: a library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2**, 1–27 (2011).
50. Lin, H.-T. L., Merz, C.-J. & Weng, R. C. A note on Platt’s probabilistic outputs for support vector machines. *Mach. Learn.* **68**, 267–276 (2007).
51. Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis. Neurosci.* **13**, 87–100 (1996).
52. Kang, I. & Maunsell, J. H. Potential confounds in estimating trial-to-trial correlations between neuronal response and behavior using choice probabilities. *J. Neurophysiol.* **108**, 3403–3415 (2012).
53. Seabold, S. P. & Perktold, J. Statsmodels: econometric and statistical modeling with Python. in *Proceedings of the 9th Python in Science Conference* (eds van der Walt, S. & Millman, J.) 92–96 (SciPy, 2010).

Acknowledgements

We thank E. Piasini for early contributions; members of our laboratories for helpful discussions; G. Iurilli, C. Kayser, E. Piasini, C. Becchio and J. Drugowitsch for feedback; and M. Libera for technical support. This work was supported by National Institute of Health grants from the NIMH BRAINS program R01 MH107620 (to C.D.H.), NINDS R01 NS089521 (to C.D.H.), the BRAIN Initiative R01 NS108410 (to C.D.H. and S.P.) and U19 NS107464 (to S.P.) and the Fondation Bertarelli (S.P.).

Author contributions

S.P. and C.D.H. conceived, designed and supervised the study. C.A.R. and A.S.M. acquired the experimental data. M.V., G.P., G.B. and M.M. performed computations. S.P., C.D.H., M.V., G.P., G.B. and M.M. wrote the paper, with feedback from all authors.

Competing interests

The authors declare no competing interests.

Additional information

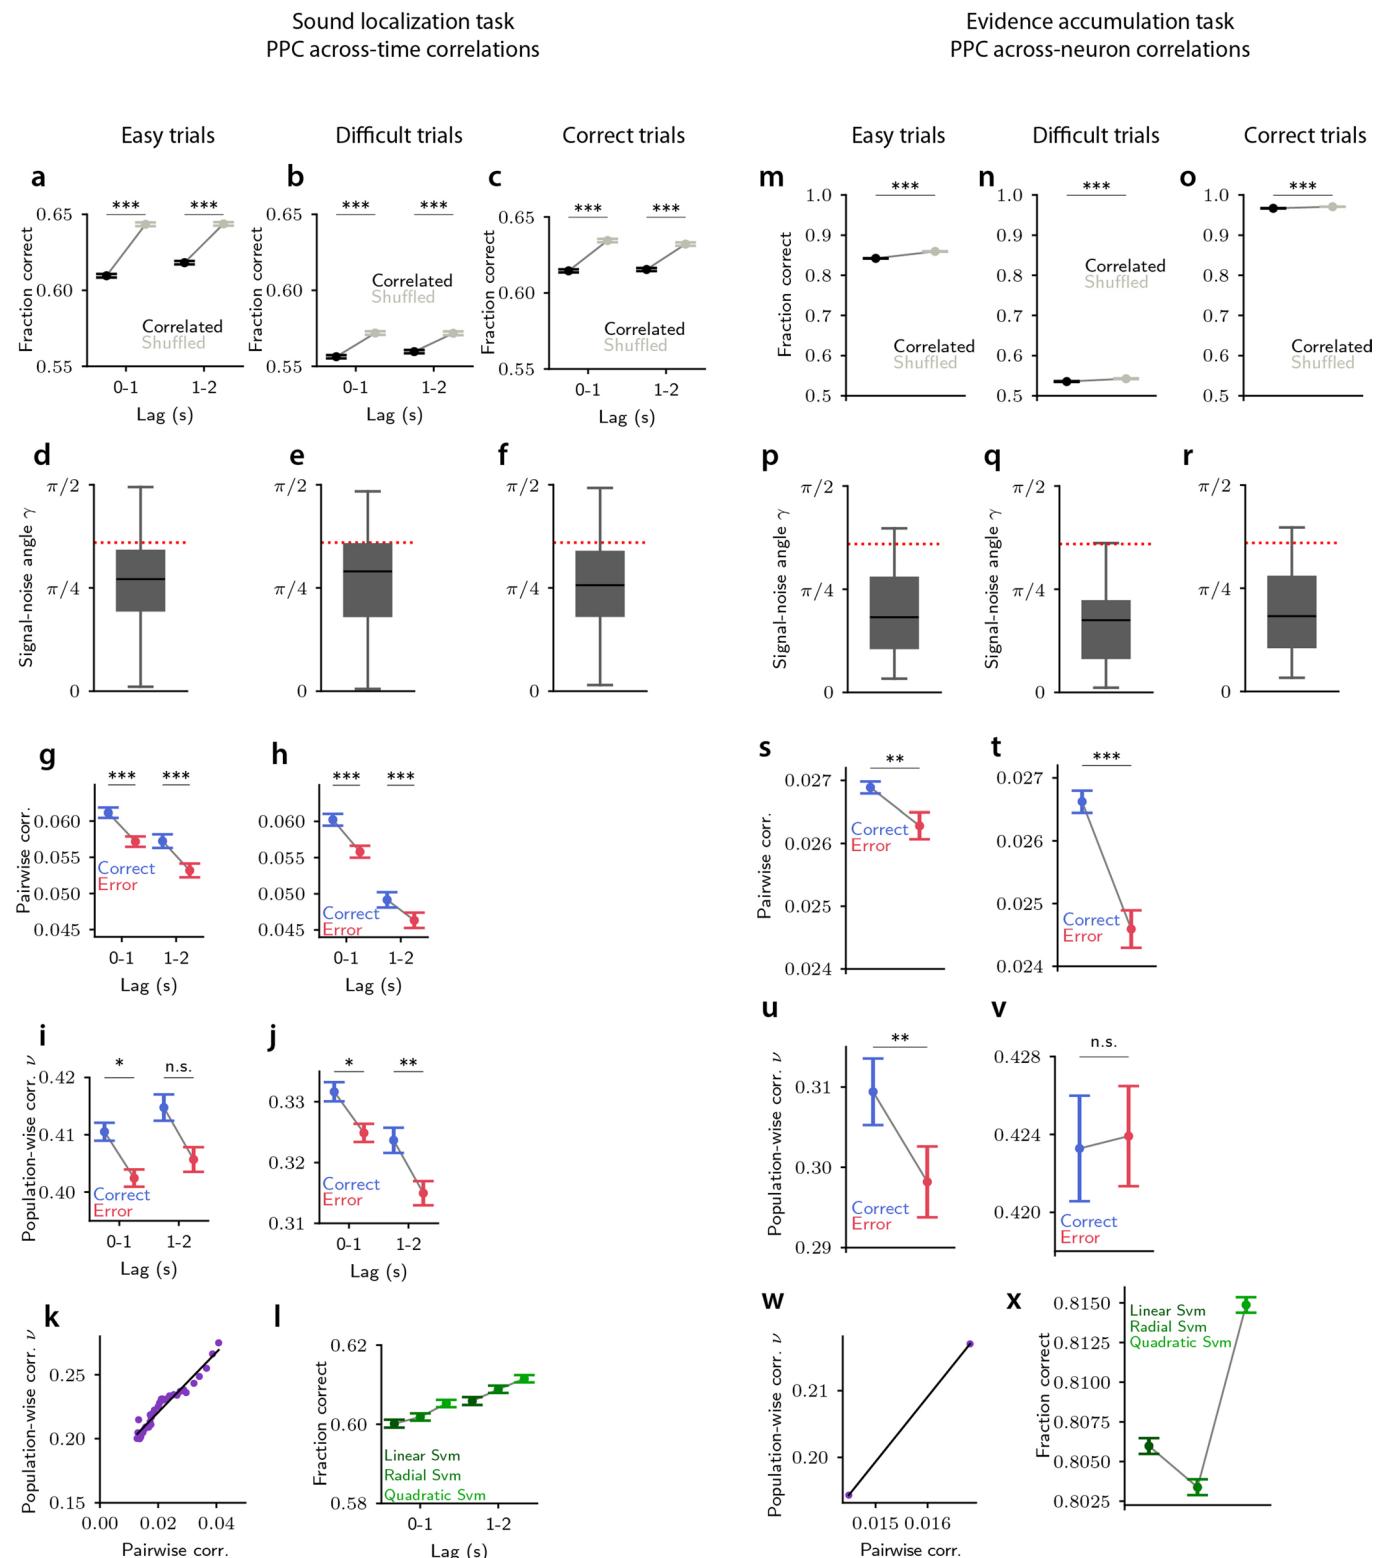
Extended data is available for this paper at <https://doi.org/10.1038/s41593-021-00845-1>.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41593-021-00845-1>.

Correspondence and requests for materials should be addressed to C.D.H. or S.P.

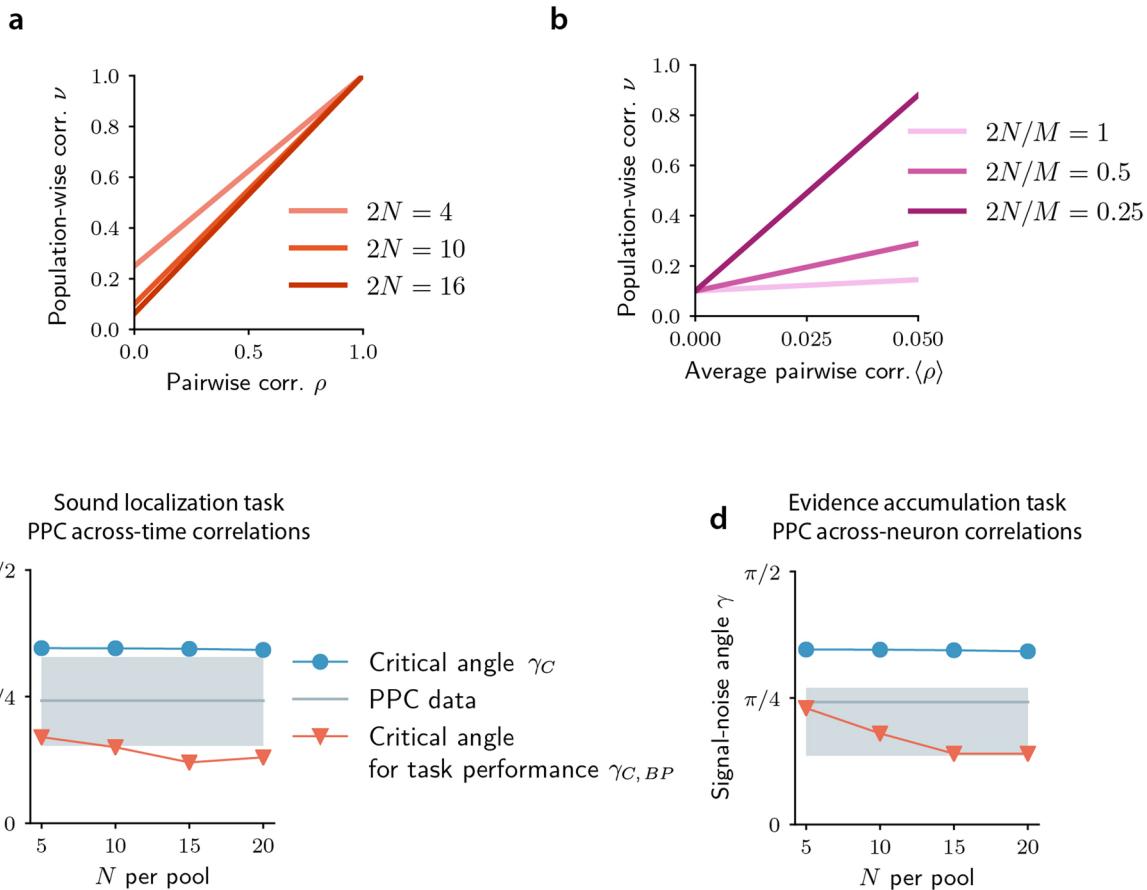
Peer review information *Nature Neuroscience* thanks Yong Gu, Joel Zylberberg, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

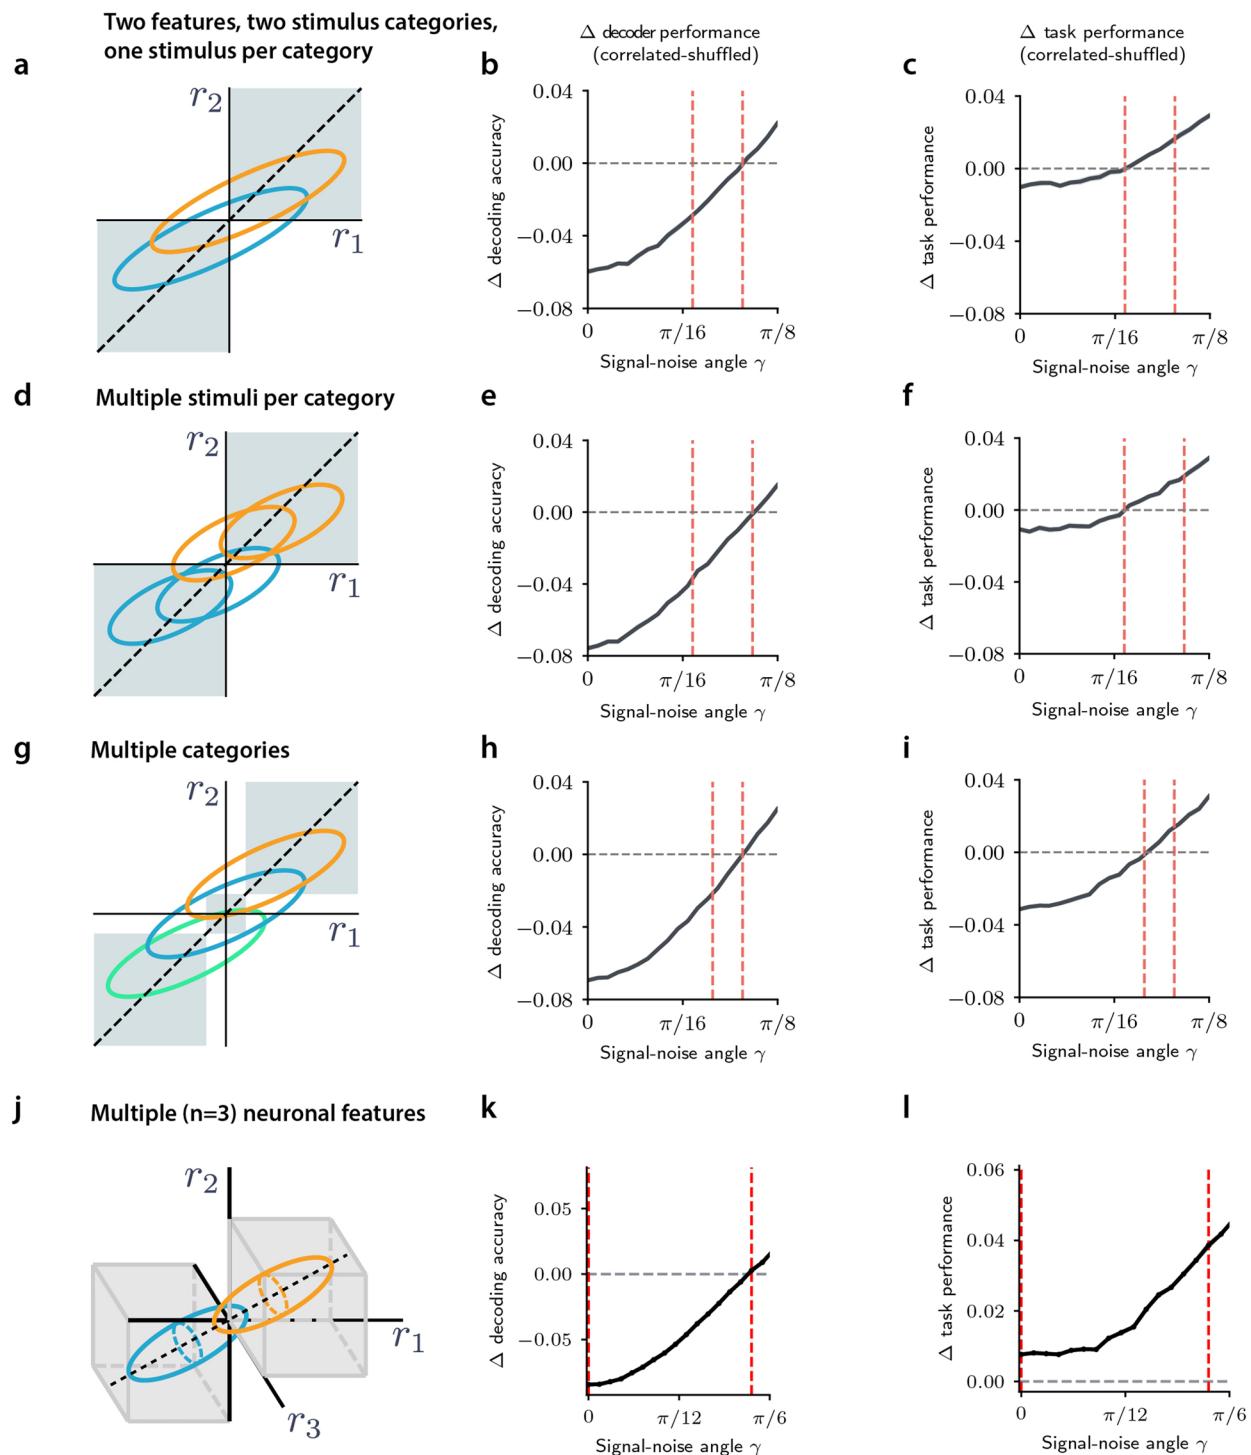


Extended Data Fig. 1 | See next page for caption.

Extended Data Fig. 1 | Response properties and across-time and across-neuron correlations in PPC during perceptual discrimination tasks for different trials categories. Panels **a-l** refer to PPC data during the sound localization task and across-time correlations. **a-c**, Accuracy of a linear decoder of the stimulus applied to the joint population activity at two different time points, for recorded (black) or trial-shuffled (gray) population vectors, for 'easy' trials with high level of sensory evidence (**a**, sound locations further from the midline than 45 deg), 'difficult' trials with low level of sensory evidence (**b**, sound locations closer to the midline than 45 deg) and behaviorally correct trials only (**c**). Errorbars report mean \pm SEM across $n=6$ sessions and all time point pairs within the specified lag range. For all comparisons, $P=10^{-4}$, two-sided permutation test. **d-f**, Distribution of the signal-noise angle γ (over $n=6$ sessions and all time point pairs within a 2 s lag), for 'easy' trials (**d**), 'difficult' trials (**e**) and behaviorally correct trials only (**f**). Boxplots show the median (line), quartiles (box) and whiskers extend to $\pm 1.5 \times$ interquartile range. Red dotted line: theoretical value of the critical angle γ_c between the information-limiting and information-enhancing regime. **g-h**, Pairwise noise correlations in time-lagged activity, for correct and error trials, for 'easy' (**g**) and 'difficult' (**h**) trials. Errorbars report mean \pm SEM across $n=6$ sessions, all time point pairs within the specified lag range and all cell pairs. For all comparisons, $P=10^{-4}$, two-sided permutation test. **i-j**, Population-wise noise correlations in time-lagged activity, for correct and error trials, for 'easy' (**i**) and 'difficult' (**j**) trials. Errorbars report mean \pm SEM across $n=6$ sessions and all time point pairs within the specified lag. In **i**, $P=0.0380$ for Lag 0-1 s, n.s. $P=0.1510$ for Lag 1-2 s, two-sided permutation test. In **j**, $P=0.0480$ for Lag 0-1 s, $P=0.001$ for Lag 1-2 s, two-sided permutation test. **k**, Relation between pairwise and population-wise noise correlations. Each dot represents the average across $n=6$ session and all time points with a given lag. The black line indicates the linear fit. **l**, Accuracy of a linear, quadratic and radial basis function SVM decoder of stimulus identity applied to joint population activity at two different time points for real recorded population vectors. Errorbars report mean \pm SEM across $n=6$ sessions and all time point pairs within the specified lag range. Panels **m-x** refer to PPC data during the evidence accumulation task and across-neuron correlations. **m-o**, Same as in **a-c**. **p-r**, Same as in **d-f**. **s-t**, Same as in **g-h**. **u-v**, Same as in **i-j**. For the evidence accumulation task, 'easy' and 'difficult' trials were defined as trials with net evidence ≥ 4 or < 4 respectively. In panels **m-v** errorbars report mean \pm SEM across $n=11$ sessions, Early and Late Delay epochs and 100 pairs of neuronal pools. In **m-o**, for all comparisons, $P=10^{-4}$. In **s**, $P=0.0120$. In **t**, $P=9 \times 10^{-4}$. In **u**, $P=0.001$. In **v**, $P=0.8641$. For all comparisons, two-sided permutation test. **w**, Same as in **k**. Each dot represents the average across $n=11$ sessions for a given delay epoch. **x**, Same as in **l**, with errorbars reporting mean \pm SEM across all $n=11$ sessions, Early and Late Delay epochs, and 100 pairs of randomly split neuronal pools.

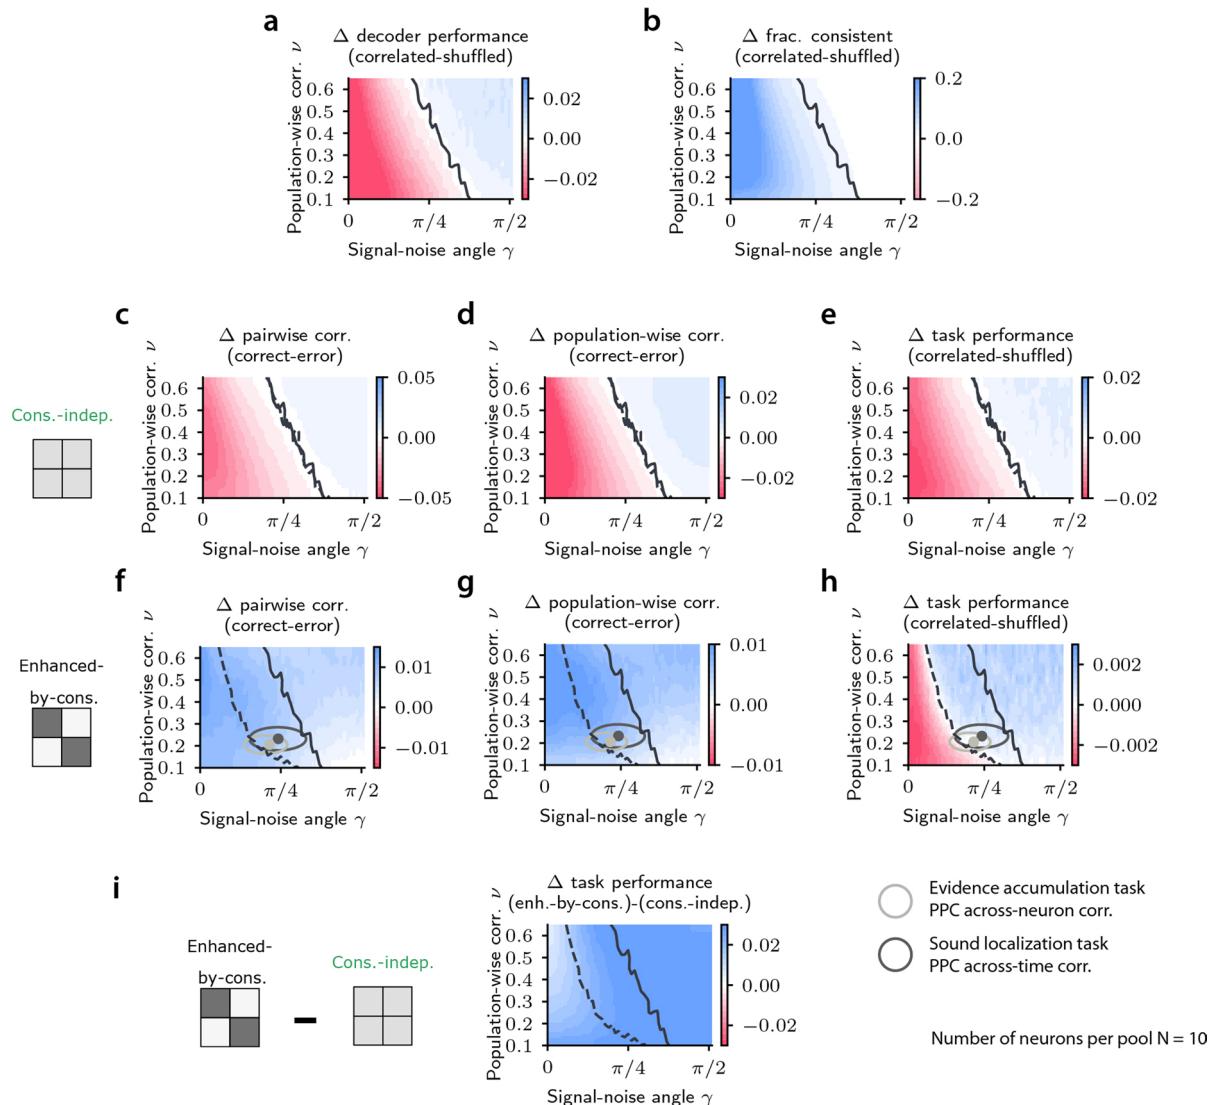


Extended Data Fig. 2 | Parameter exploration of the two-pools encoding-readout model and comparison with PPC data. **a**, Population-wise noise correlations v as a function of the pairwise noise correlations ρ , for different values of the active neurons $2N$ (N neurons per pool). Here we assumed that all neurons were active ($M=2N$). For $\rho=0$, the population-wise correlation is equal to $v=1/(2N)$. **b**, Population-wise noise correlations v as a function of the average over all pairs of pairwise noise correlations ρ (where \cdot denotes the average over neuron pairs), for different fraction of active neurons $2N/M$ (total number of neurons given by $M=2N+K$). By decreasing the fraction of active neurons, the constant of proportionality between v and ρ increases. **c**, Blue line: critical signal-noise angle γ_c below which correlations are information-limiting in the model, as a function of the number of neurons per pool N , computed using the experimental value of the PPC across-time population-wise correlation for the sound localization task. Red line: critical value $\gamma_{c,BP}$ for the angle above which the task performance in correlated data is higher than that in shuffled data. The experimental distribution of PPC signal-noise angles is reported for comparison ($n=6$ sessions and all time point pairs within a 2 s lag). Horizontal gray line indicates the median. Box edges indicate the first and third quartile. **d**, Same as **c** for the evidence accumulation task and across-neuron PPC correlations.

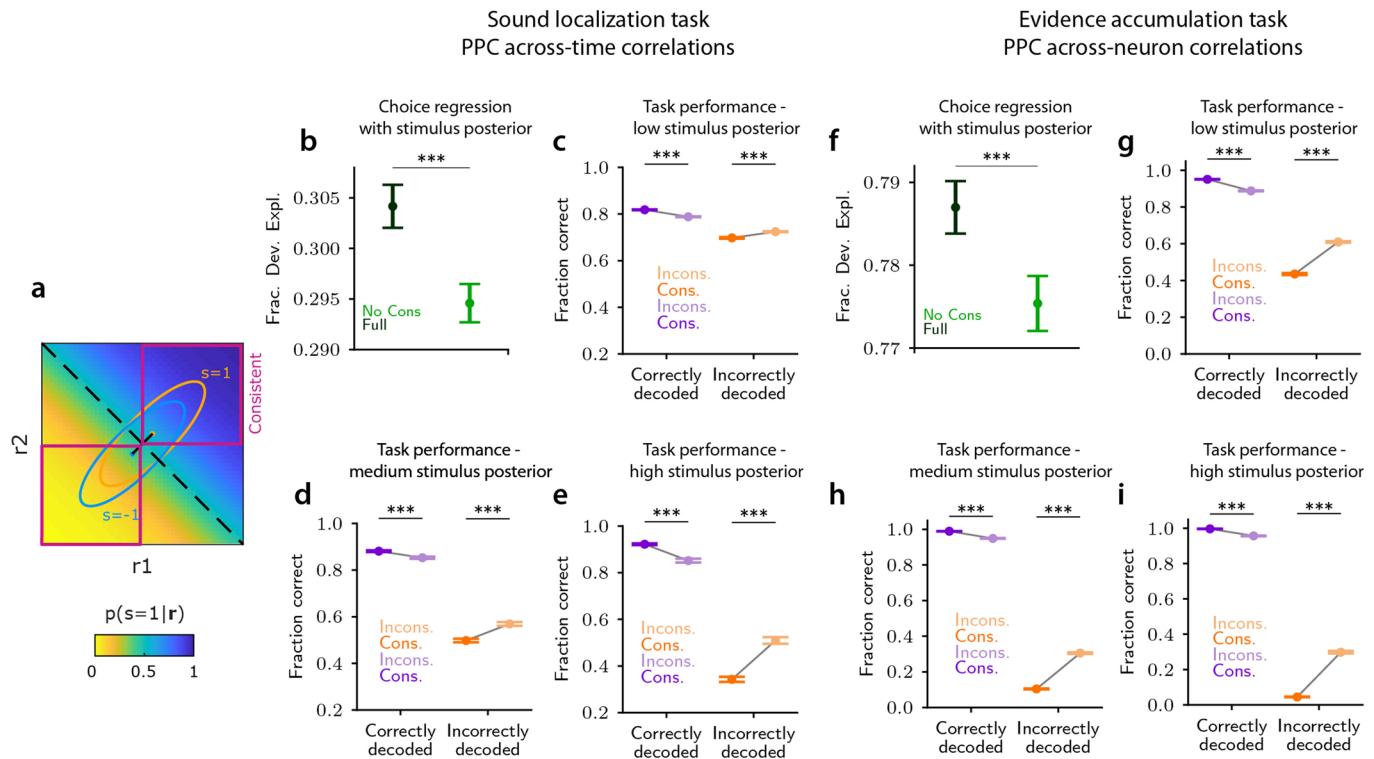


Extended Data Fig. 3 | See next page for caption.

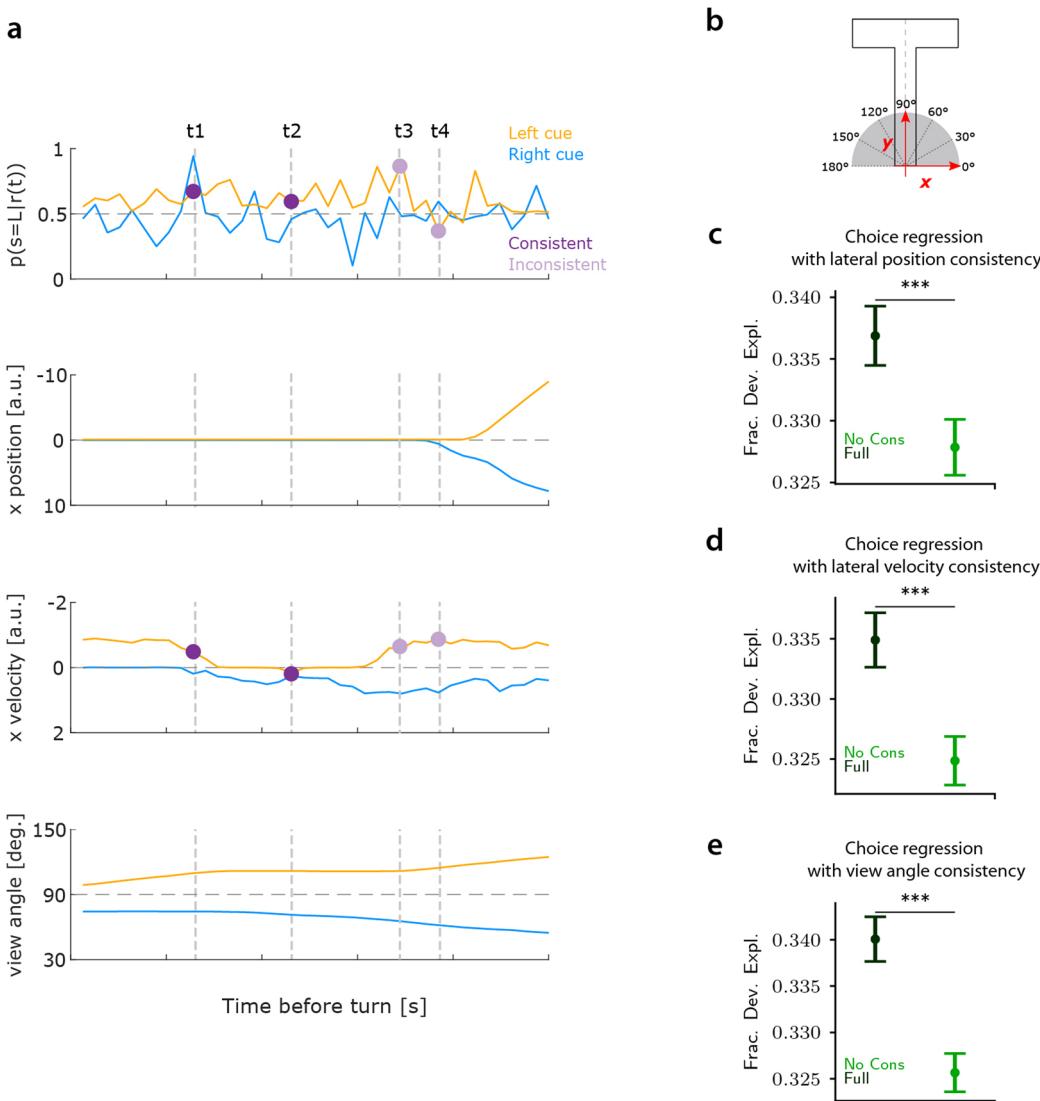
Extended Data Fig. 3 | Extension of the encoding readout model to multiple features, multiple stimuli per category and multiple categories. **a**, Left: schematic of the encoding readout models with two neural features, two categories, two stimuli per category. Each axis represents the activity of a single feature. Colored ellipses: 95% confidence intervals for the simulated neural responses to two different stimuli ($s=1, s=-1$). Dashed black line: stimulus axis. Gray shaded areas: regions of the response space in which stimulus information is encoded consistently across pools and the behavioral readout efficacy is enhanced. **b**, Difference in stimulus classification accuracy between correlated and shuffled responses computed using a linear decoder applied to the joint population activity across the two features, as a function of the signal-noise angle γ . **c**, Difference in task performance between correlated and shuffled responses predicted by an enhanced-by-consistency readout of simulated neural activity, as a function of the signal-noise angle γ . In panel **b-c**, the red dashed lines delimit the parameter range where correlations are information-limiting but task performance is enhanced for correlated data. Data are mean \pm SEM over $n=10$ simulations with 50,000 trials each, with $d = \sqrt{0.02}$, $\rho = 0.8$, $\sigma = 0.2$, $\eta = 0.7$. **d-f**, Same as in **a-c**, but for an encoding model with two neural features, two categories, and multiple ($n=2$) stimuli per category. Within each category, stimulus-specific distributions are symmetrically displaced on either side of the between-category signal axis. Within each category, the noise axes of the individual distributions are aligned to each other and aligned to the vector of differences of mean activity. Data are mean \pm SEM over $n=10$ simulations with 50,000 trials each. We set half the distance between the centers of the distributions of the two categories to $d = \sqrt{0.02}$, and the distance between the centers of the distributions of individual stimuli within each category to $d_2 = 0.3$. In simulations, we set $\rho = 0.8$, $\sigma = 0.13$ (for the distributions of individual stimuli within each category), $\eta = 0.7$. **g-i**, Same as in **a-c**, but for an encoding readout model with two pools and multiple ($n=3$) stimulus categories. Mean responses to the three stimulus categories are aligned along a unique signal axis, and the noise axes of individual distributions form an angle γ with the stimulus axis. Data are mean \pm SEM over $n=10$ simulations with 50,000 trials each, with $d = \sqrt{0.02}$ (distance between the distributions across individual categories), $\rho = 0.8$, $\sigma = 0.2$ (for individual distributions), $\eta = 0.7$. **j-l**, Same as in **a-c**, for a model with three one-dimensional neural features. Neural activity is considered consistent if the same stimulus is decoded from all three features. Data are mean \pm SEM over $n=10$ simulations with 50,000 trials each, with $d = \sqrt{0.02}$, $\rho = 0.8$, $\sigma = 0.2$, $\eta = 0.7$.



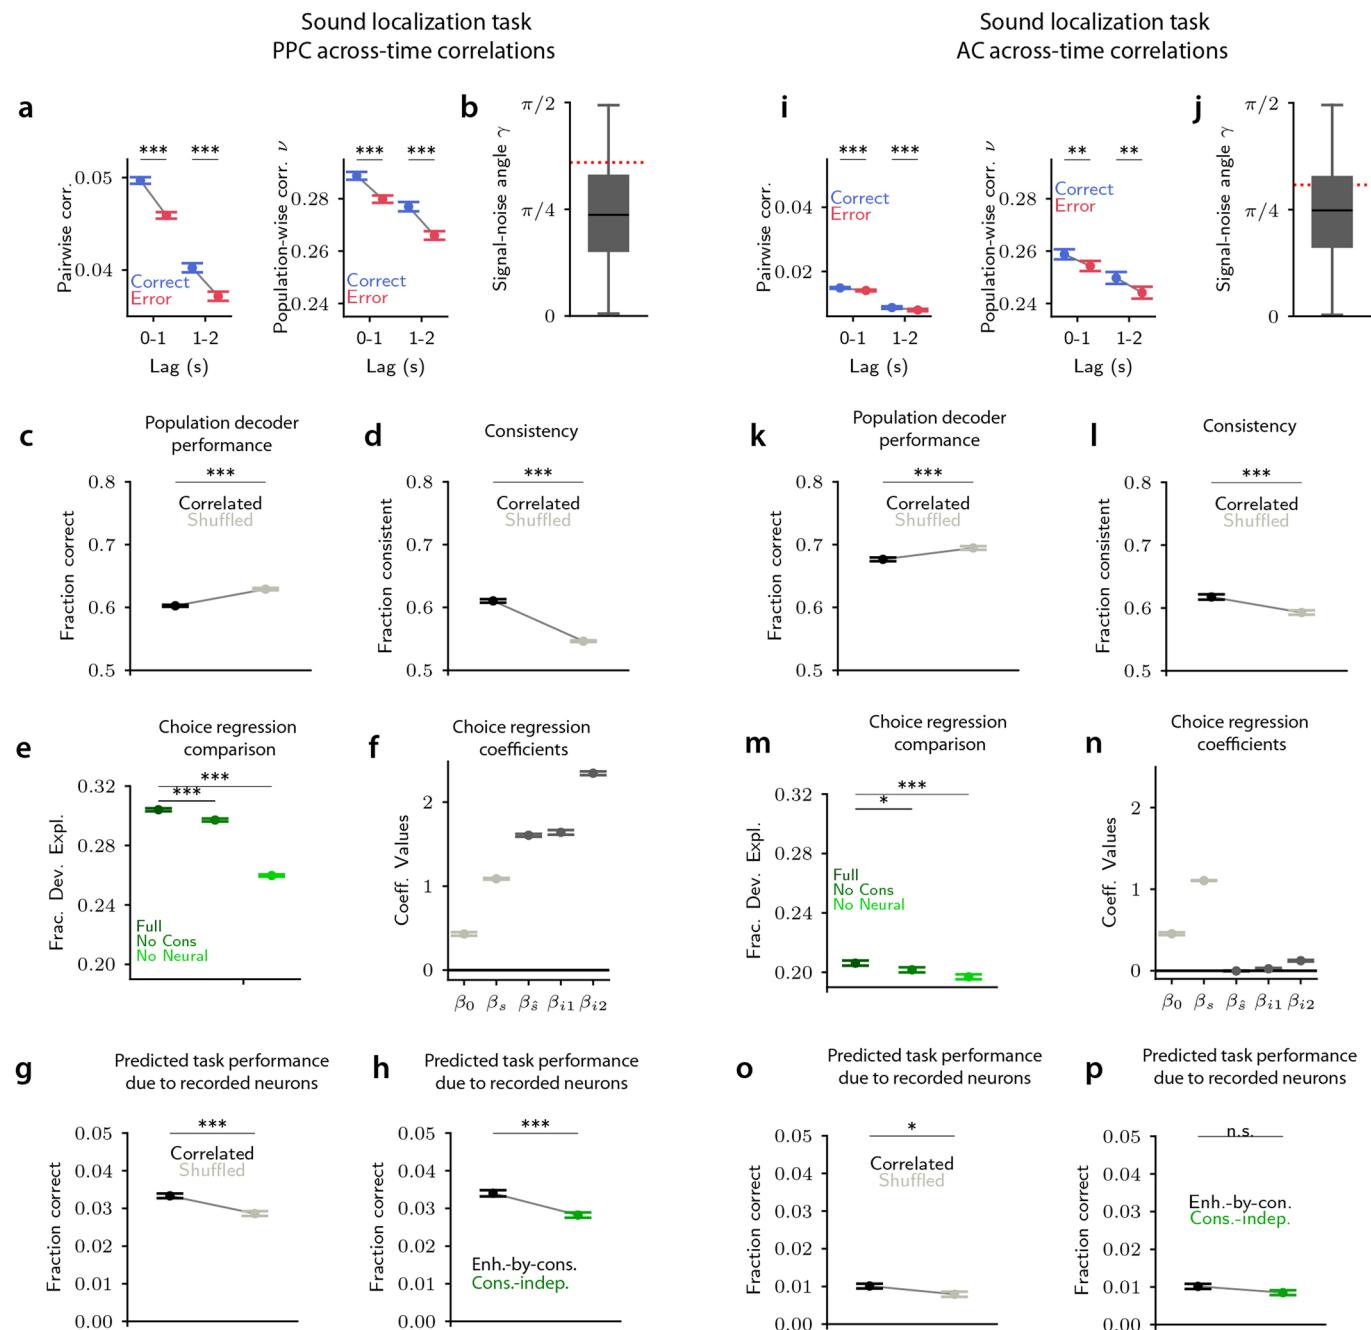
Extended Data Fig. 4 | Exploration of the parameter space of the encoding readout model. Same as in Fig. 3, for an encoding readout model with $N=10$ neurons in each pool. **a**, Difference in the accuracy of a linear decoder of stimulus applied to correlated and shuffled simulated neural activity for different values of the signal-noise angle (γ) and population-wise correlations (ν). For all panels, black solid line: boundary between a regime with information-limiting correlations and information-enhancing correlations. **b**, The difference between correlated and shuffled activity in the fraction of trials in which the two neural features encode consistent stimulus information is higher in the information-limiting regime and increases with population-wise correlations strength. Panels **c-e** refer to the consistency-independent readout. **c-d**, Difference in average pairwise correlations (**c**) and population-wise correlations (**d**) between trials with correct and incorrect predicted task performance for different combinations of model parameters. **e**, Difference in task performance predicted by applying the consistency-independent readout to correlated and shuffled simulated neural activity for different combinations of model parameters. For panels **c-i**, dashed black line: boundary between a regime where task performance is higher for correlated responses and a regime where performance is higher for shuffled responses. The overlap between the continuous and dashed black line indicates that correlations that limits information are also detrimental for behavior. Panel **f-h** refer to the enhanced-by-consistency readout (consistency modulation index $\eta=0.85$). **f-g**, Same as in **c-d**. With the enhanced-by-consistency readout correlations are higher in correct trials. **h**, Same as in **e**. The area between the dashed and the continuous black line indicates a regime where correlations are information-limiting but task performance is higher for correlated responses. Thus, in the parameter range between the two lines, the readout is able overcoming the negative impact of correlations. Dark and light gray dots and ellipses: mean values and range between the 25th and the 75th percentile of the signal-noise angles and population-wise correlations for PPC data from the sound localization task and evidence accumulation task, respectively. **i**, Difference in task performance predicted by applying the enhanced-by-consistency readout or the consistency-independent readout with matched readout efficacy for different combinations of model parameters. The enhanced-by-consistency readout yields increased task performance with respect to the consistency-independent readout. Panels represent the mean over $n=100$ simulations with 300,000 trials each.



Extended Data Fig. 5 | The effect of neural correlations on the mouse's single trial choices cannot be explained by higher stimulus information associated to consistent neural representations. **a**, Schematic example showing response distributions along two neural features (r_1, r_2) to two stimuli ($s = -1$: blue, $s = 1$: orange). Black dashed line: optimal decoding boundary of a linear decoder trained on the simulated neural responses. The background color represents the linear decoder posterior probability that stimulus $s=1$ has occurred given the observation of the neural response $\mathbf{r} = (r_1, r_2)$. Intuitively, the farther neural response \mathbf{r} is from the decoding boundary, the farther $p(s=1|\mathbf{r})$ is from 0.5, and the more ‘informative’ \mathbf{r} is about the stimulus. Note that, in the example shown, consistent trials have on average higher posterior probability than inconsistent trials, which might represent a confounder for the effect of consistency on mouse's choices. To control for potential confounders due to differences in the levels of stimulus information between trials with consistent and inconsistent stimulus information, we fitted to the data a readout model that predicted choice using the posterior probability of the stimulus and posterior probabilities consistency given the neural responses, rather than just the decoded stimulus identity (**b, f**). We further repeated the analyses of Fig. 4 on trials partitioned into those with low ($|p(s=1|\mathbf{r}) - 0.5| < 0.16$), medium ($|p(s=1|\mathbf{r}) - 0.5| > 0.16 \wedge |p(s=1|\mathbf{r}) - 0.5| < 0.32$), or high ($|p(s=1|\mathbf{r}) - 0.5| > 0.32$) ‘stimulus information’ (**c-e, g-i**). Panels **b-e** refer to PPC data during the sound localization task. **b**, Performance (fraction of deviance explained) in explaining single-trial choice of models using neural predictors based on posterior probabilities. Full model includes all predictors values, comprising stimulus posterior probability and posterior probability consistency. No Cons model neglects neural consistency by shuffling consistency values across trials. **c-e**, Left (purple dots). Task performance in trials with correctly decoded stimulus is higher when information is encoded consistently than inconsistently. Right (orange dots). The opposite happens for trials with incorrectly decoded stimulus. Thus, stimulus information in neural activity has a larger impact on choices when it is encoded consistently across time, even when subsets of trials having approximately the same posterior are used. For all comparisons in **b-e**, $P = 10^{-4}$, two-sided permutation test. Errorbars represent mean \pm SEM across $n = 6$ sessions and all time point pairs within a 1 s lag. Panels **f-i** refer to PPC data during the evidence accumulation task. **f**, Same as in **b**. $P = 6 \times 10^{-4}$, two-sided permutation test **g**, Same as in **c**. **h**, Same as in **d**. **i**, Same as in **j**. In panels **f-i**, consistency and mouse choices are computed from the activity of two pools of neurons. For all comparisons in **g-i**, $P = 10^{-4}$, two-sided permutation test. In **f-i**, errorbars represent mean \pm SEM across $n = 11$ sessions, Early and Late Delay epochs and 100 pairs of neuronal pools. From **b-i**, the fact that information in neural activity informs choice more effectively when it is consistent cannot be explained by differences in overall stimulus information level. Rather, for a given amount of sensory information, more information can be extracted to guide behavioral choices if it is distributed redundantly across neurons or across time.

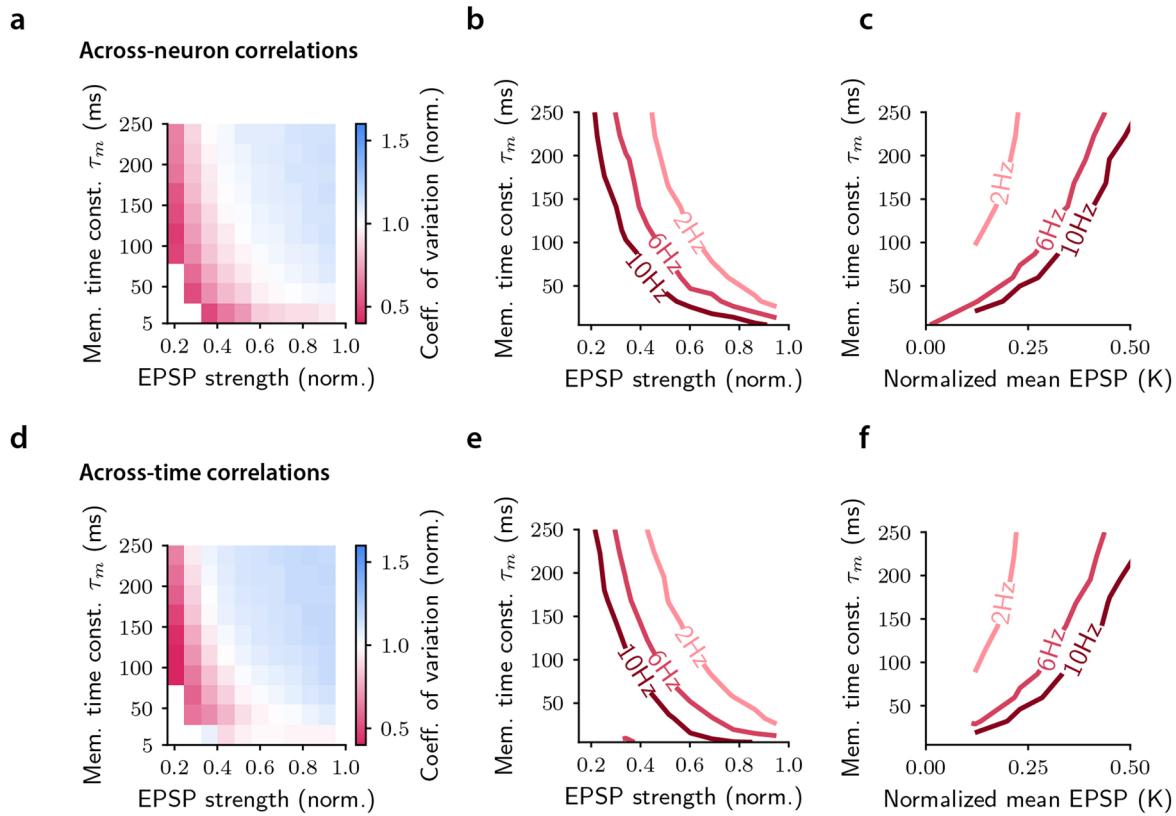

Extended Data Fig. 6 | The role of neural consistency in the readout of PPC activity is not due to the consistency of measured behavioral parameters.

Panels **a-e** refer to PPC neural activity during the sound localization task. To rule out the concern that the impact of across-time consistency of PPC activity on the mouse's choice does not only reflect the effect of running related parameters (whose temporal consistency may correlate with both the mouse's choice and the temporal consistency of neural activity), we developed and fit to PPC data a more sophisticated readout model that explicitly includes such contributions in predicting choices. **a**, The temporal evolution of the decoder posterior probability of left stimulus presentation given the recorded PPC population activity is shown along with the corresponding temporal evolution of a selection of three concurrently-measured behavioral parameters (lateral position, lateral velocity, view angle), for an example left (orange) and right (blue) cue trial. Colored dots indicate two example time point pairs with consistent (t_1-t_2 , dark purple) or inconsistent (t_3-t_4 , light purple) neural information. Colored dots in the first and third row show that neural consistency is not necessarily associated to behavioral consistency (when considering lateral running speed, t_1-t_2 are behaviorally inconsistent while t_3-t_4 are behaviorally consistent). **b**, Schematic representation of the virtual T-maze with corresponding x-y coordinates labelling and mouse's view angle (for a mouse oriented along the y axis). **c-e**, Performance (fraction of deviance explained) in explaining (using two population vectors at different points) single-trial mouse choice of models that use both neural and behavioral consistency (**c**: lateral position, **d**: lateral velocity, **e**: view angle). Full model includes all predictors values, comprising neural and behavioral consistency. No Cons model neglects neural consistency by shuffling consistency values across trials. **c**, $P=10^{-4}$. **d**, $P=2 \times 10^{-4}$. **e**, $P=10^{-4}$, two-sided permutation test. Errorbars report mean \pm SEM across $n=6$ sessions and all pairs of time point within a 1 s lag. Results in **c-e** show that neural consistency still contributed to predicting choices when we added the consistency of running-related variables to the choice regression. This suggests that consistency of the instantaneous PPC population activity across time genuinely influences the behavioral readout of the stimulus information, above and beyond what can be predicted about choice from the consistency of measured behavioral variables.

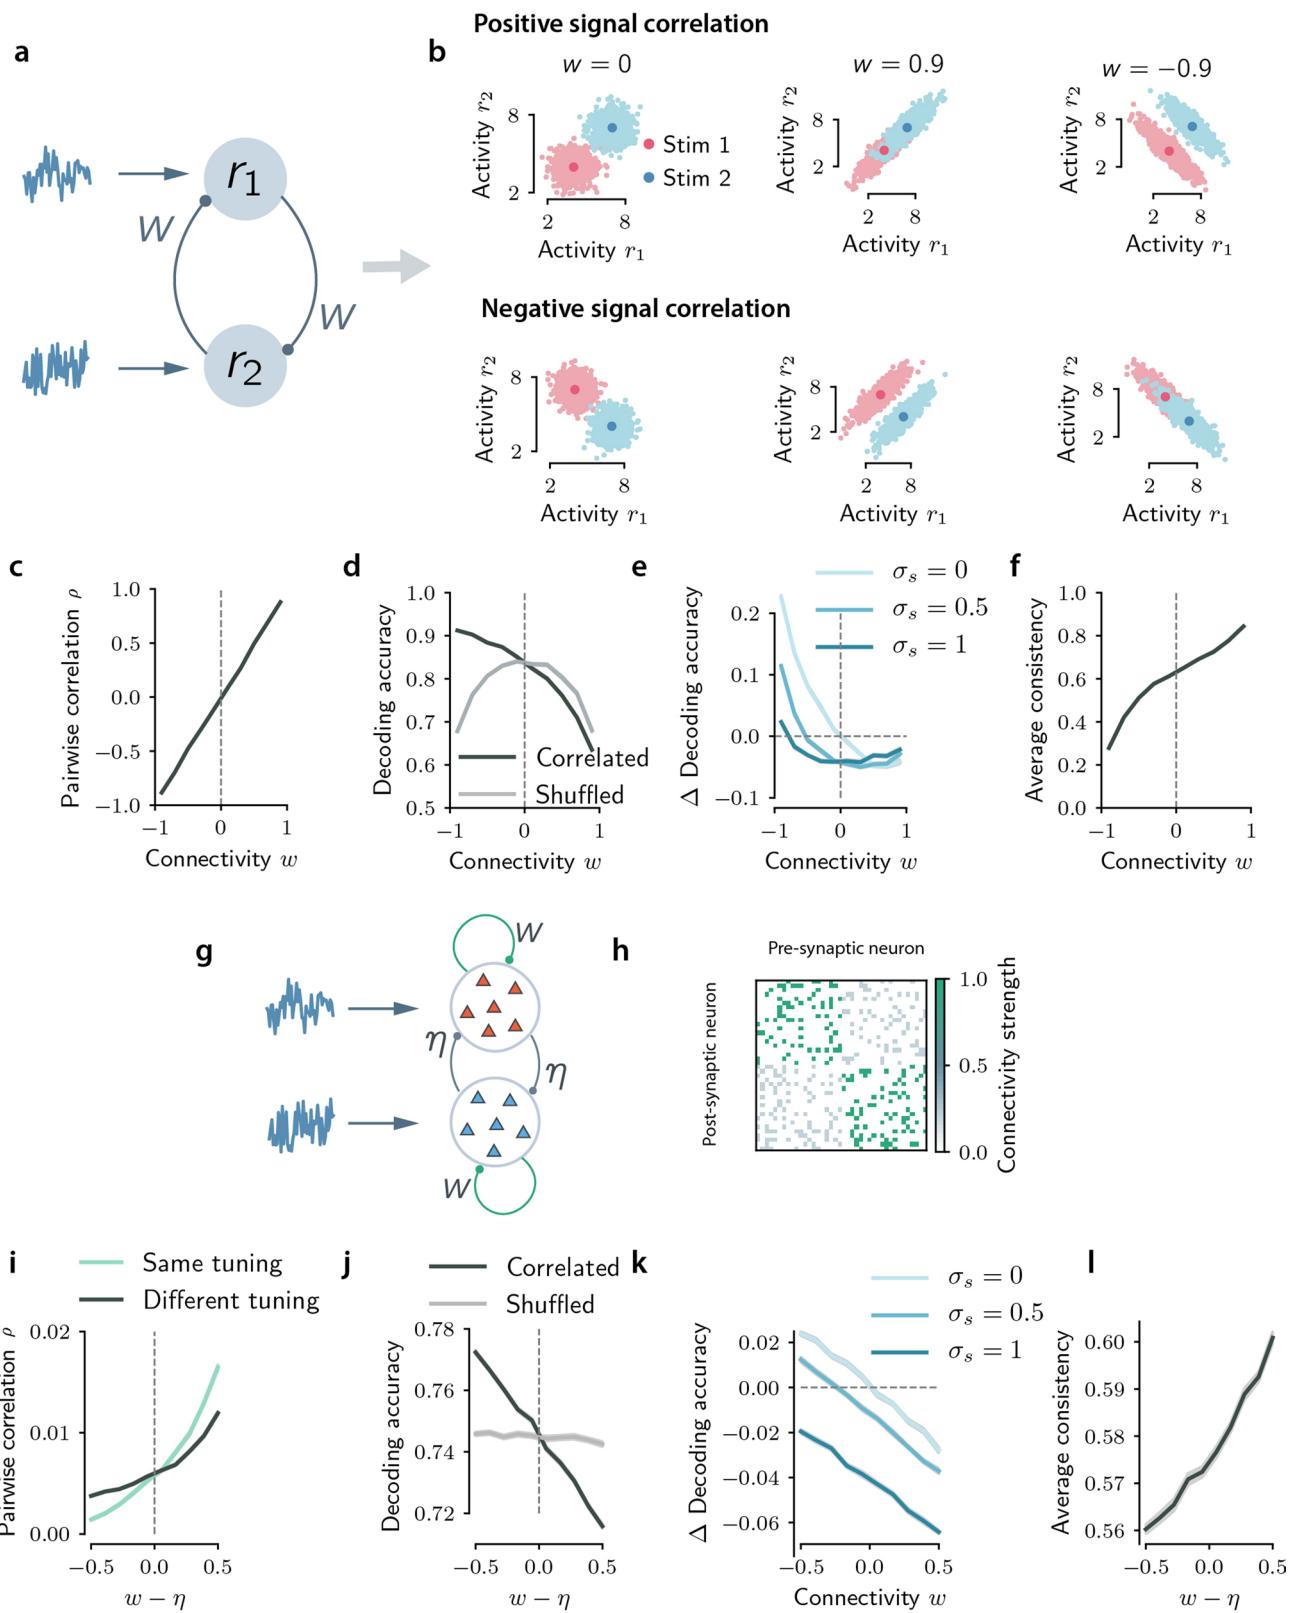


Extended Data Fig. 7 | See next page for caption.

Extended Data Fig. 7 | Across-time correlations in AC do not benefit task performance as they do in PPC. Panels **a-h** refer to PPC neural activity during the sound localization task. **a-h**, Summary of the main results of the analysis of across-time correlations in PPC activity (from Fig. 1, Fig. 4 and Fig. 5), useful for the comparison with AC data. Panels **i-p** refer to AC neural activity during the sound localization task. **i**, Pairwise (left) and population-wise (right) noise correlations in time-lagged activity, for correct and error trials. Overall, noise correlations strength is lower in AC than in PPC. **j**, Distribution of the signal-noise angle γ (over $n=6$ sessions and all time point pairs within a 2 s lag). Boxplots show the median (line), quartiles (box) and whiskers extend to $\pm 1.5 \times$ interquartile range. Red dotted line: analytically computed bound between the information-limiting and information-enhancing regime. **k**, Accuracy of a linear decoder of the stimulus applied to joint population activity at two different time points, for real recorded (black) or trial-shuffled (gray) data. The decoder accuracy is higher in AC than in PPC (fraction correct: 0.676 ± 0.003 in AC, 0.602 ± 0.001 in PPC, $P=10^{-4}$, two-sided permutation test), compatible with the view that AC is involved in the encoding of sound information. Across-time correlations limit the encoding of stimulus information also in AC, but with a smaller effect than in PPC (average increase in decoder accuracy by shuffling: 0.018 ± 0.001 in AC, 0.026 ± 0.001 in PPC, $P=10^{-4}$, two-sided permutation test). Equivalent percentage increase of above-chance (that is above 50%) decoding performance: 10.5% in AC, 25.5% in PPC. **l**, Fraction of trials in which stimulus information is encoded consistently across time, for real recorded (black) or trial-shuffled (gray) data. The increase in consistency due to across-time correlations is smaller in AC than in PPC (-0.025 ± 0.001 in AC, -0.064 ± 0.01 in PPC, $P < 10^{-4}$, two-sided permutation test). **m**, Performance (fraction of deviance explained) in explaining single-trial choices of several readout models (see Methods). Full model uses all predictors (neural and non-neural). ‘No Cons’ model neglects neural consistency. ‘No Neural’ model neglects stimulus decoded from neural activity and neural consistency. A linear SVM is used to decode the stimulus from neural activity. Across-time consistency in AC provides negligible improvements in behavioral choice predictions when compared to PPC (increase in fraction of deviance explained when comparing the Full with the ‘No Cons’ model: 0.0045 ± 0.0001 in AC, 0.0083 ± 0.0012 in PPC, $P=10^{-4}$, two-sided permutation test). **n**, Best-fit coefficients of the Full readout model. AC neural predictors are characterized by low weights. **o**, Task performance predicted by applying the best-fit readout model to real recorded (black) or trial-shuffled (gray) data. Task performance attributable to recorded neurons is much lower in AC than in PPC (~1% in AC, ~3.5% in PPC). Correlations in AC activity enhance task performance, but the effect is small. **p**, Task performance predicted by applying to real recorded population vectors the best-fit enhanced-by-consistency (black) and the consistency-independent readout model (green). Task performance attributable to the recorded AC neural activity would not be substantially different if the behavioral readout was consistency-independent. In **i**, **k-p**, errorbars report mean \pm SEM across all cell pairs (only **b-left**) and all time point pairs within the specified lag range or within a 1 s lag from $n=6$ sessions. For **i**, left, $P=10^{-4}$ for all comparisons, right, $P=0.0016$ for lag 0-1 s, $P=0.001$ for lag 1-2 s. For **k**, **l**, $P=10^{-4}$. For **m**, $***P=10^{-4}$, $*P=0.0324$. For **o**, $P=0.0191$. For **p**, $P=0.0690$. All comparisons, two-sided permutation test.



Extended Data Fig. 8 | Exploration of the parameters of the biophysical model for the enhanced-by-consistency readout. **a**, Normalized coefficient of variation (CV) computed for different values of the membrane time constant τ_m of the readout neuron and EPSP strength w (connection strength from the input to the readout neuron). The mean input rate was set to $R_{in}=6$ Hz. The red parameter region corresponds to the region where the standard deviation of the readout firing rate increases less than the readout mean firing rate with the value of spatial correlations. **b**, Contour lines corresponding to the parameter values (τ_m, w) where the normalized CV is equal to 1, for different values of the input firing rate. **c**, Contour lines for which the normalized CV is equal to unity, in the parameter space defined by the membrane time constant τ_m and the mean EPSP input in a window τ_m normalized by the voltage gap $\Delta V = V_{threshold} - V_r$, that is $K = wR_{in}\tau_m/\Delta V$. Regions of parameters on the left of the contour lines correspond to the parameter values where the standard deviation of the readout neuron increases less than its mean with spatial correlations. **d-f**, Same as **a-c** for temporal correlations.



Extended Data Fig. 9 | See next page for caption.

Extended Data Fig. 9 | Encoding model internally generating correlated activity through recurrent dynamics. **a**, Schematic illustrating the basic setup of the encoding recurrent model. Two neurons receive stimulus-dependent feedforward input (which determines the signal correlations) and input noise, and are connected through recurrent synapses with strength w . **b**, Noise correlations are generated through recurrent connectivity, and depend on the sign of w (for $w=0$ responses are uncorrelated). Top: for positive signal correlations, positive (resp. negative) values of the connectivity generate information-limiting (resp. information-enhancing) noise correlations. Bottom: for negative signal correlations, positive (resp. negative) values of the connectivity generate information-enhancing (resp. information-limiting) noise correlations. **c-f**, Average pairwise noise correlation (over $n=10000$ random pairs of neurons) (**c**), decoding accuracy for correlated and shuffled responses (**d**), difference in decoding accuracy between correlated and shuffled responses for different values of shared noise (**e**) and average consistency (**f**) as a function of connectivity strength w . In **c,d,f** the external input noise is uncorrelated across the two neurons. **g**, Schematic illustrating the $2N$ -dimensional encoding recurrent model. Two N -dimensional neuronal groups with opposite stimulus selectivity receive stimulus-dependent feedforward input and input noise. The connectivity strength is excitatory, sparse and takes the value $w > 0$ between neurons belonging to the same group, and $\eta > 0$ between neurons belonging to different groups. **h**, Example of a connectivity matrix adopted in these analyses. All matrix entries are positive (excitatory synapses) and sparse with connection probability p . **i-l** Same quantities computed in **c-f** as a function of the difference between the within-group connectivity and between-groups connectivity strength, $w - \eta$. We set $N=50$, $p=0.5$, $\eta=0.5$. In **c-f**, **i-l** data are presented as mean \pm SEM over $n=50$ simulations.

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection	Calcium imaging data were acquired using ScanImage from Vidrio Technologies (versions 3 and 5), and the virtual reality environments were designed and controlled using ViRMEs
Data analysis	All analyses were performed in Matlab (R2019b) and Python 3.6 used custom software developed by our labs. For logistic regression, a toolbox was used (statsmodel, cited in reference list as ref 53). For SVM analysis, a toolbox was used (LIBSVM, cited in reference list as ref 49). Code availability: The code for the biophysical readout model (Fig. 6) is available for download at: https://github.com/gbondanelli/BiophysicalReadout . The code for the encoding and readout model is available for download at: https://github.com/moni90/encoding_readout_model The code for data analysis is available from the corresponding authors upon reasonable request.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The sound localization task data that support the findings of the current study can be downloaded at <https://gin.g-node.org/MMoroni/>

PPC_AC_2p_sound_localization (doi 10.12751/g-node.tqbad8).

The evidence accumulation task data that support the findings of the current study can be downloaded at <https://gin.g-node.org/MMoroni/>

PPC_2p_evidence_accumulation. (doi 10.12751/g-node.g1xyem).

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

- Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	All data collected that met the inclusion criteria were included in the results.
Data exclusions	All data exclusions are described in the Methods section, and were determined in the previous publications describing the two datasets used here (Runyan et al Nature 2017; Morcos and Harvey Nature Neuroscience 2016). Sound localization dataset: one of the seven PPC and one of the seven AC recording sessions used in our previous published work was excluded due to the large unbalance of left/right stimuli that were presented to the mouse across trials in that session, which would result in too few trials available for our analyses. Evidence accumulation dataset: we did not use trials with zero net evidence (<10% trials in 2/11 sessions).
Replication	The major results of the paper were present in each individual dataset.
Randomization	There are no applicable 'experimental groups' in this study.
Blinding	There are no applicable 'experimental groups' in this study to blind. All data were collected, processed, and assessed for quality prior to analysis.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
<input checked="" type="checkbox"/>	Antibodies
<input checked="" type="checkbox"/>	Eukaryotic cell lines
<input checked="" type="checkbox"/>	Palaeontology and archaeology
<input type="checkbox"/>	Animals and other organisms
<input checked="" type="checkbox"/>	Human research participants
<input checked="" type="checkbox"/>	Clinical data
<input checked="" type="checkbox"/>	Dual use research of concern

Methods

n/a	Involved in the study
<input checked="" type="checkbox"/>	ChIP-seq
<input checked="" type="checkbox"/>	Flow cytometry
<input checked="" type="checkbox"/>	MRI-based neuroimaging

Animals and other organisms

Policy information about [studies involving animals](#); [ARRIVE guidelines](#) recommended for reporting animal research

Laboratory animals

Dataset 1 (Runyan et al Nature 2017): Imaging data were acquired from five male C57BL/6J mice (The Jackson Laboratory), aged 6-8 weeks at the initiation of behavioral task training. Imaging began 4-6 weeks after viral injection and continued for 4-12 weeks.
Dataset 2 (Morcos and Harvey Nature Neuroscience 2016): Imaging data were acquired from five male C57BL/6J mice (The Jackson Laboratory), aged 8-10 weeks at the initiation of behavioral task training. Imaging began at least 4 weeks after viral injection and was continued for up to 12 weeks.

Wild animals

The study did not involve wild animals.

Field-collected samples

This study did not involve samples collected from the field.

Ethics oversight

All experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee.

Note that full information on the approval of the study protocol must also be provided in the manuscript.