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

Martina Valente[1*], Giuseppe Pica[1*], Caroline A. Runyan[2], Ari S. Morcos[2], Christopher D. Harvey[2,3] & Stefano Panzeri[1,3]

[1] Neural Computation Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy
[2] Department of Neurobiology, Harvard Medical School, Boston, MA, USA
[3] Correspondence should be addressed to: harvey@hms.harvard.edu and stefano.panzeri@iit.it
[*]Equal first author contribution

The spatiotemporal structure of activity in populations of neurons is critical for accurate perception and behavior. Experimental and theoretical studies have focused on "noise" correlations – trial-to-trial covariations in neural activity for a given stimulus – as a key feature of population activity structure. Much work has shown that these correlations limit the stimulus information encoded by a population of neurons, leading to the widely-held prediction that correlations are detrimental for perceptual discrimination behaviors. However, this prediction relies on an untested assumption: that the neural mechanisms that read out sensory information to inform behavior depend only on a population's total stimulus information independently of how correlations constrain this information across neurons or time. Here we make the critical advance of simultaneously studying how correlations affect both the encoding and the readout of sensory information. We analyzed calcium imaging data from mouse posterior parietal cortex during two perceptual discrimination tasks. Correlations limited the ability to encode stimulus information, but (seemingly paradoxically) correlations were higher when mice made correct choices than when they made errors. On a single-trial basis, a mouse's behavioral choice depended not only on the stimulus information in the activity of the population as a whole, but unexpectedly also on the consistency of information across neurons and time. Because correlations increased information consistency, sensory information was more efficiently converted into a behavioral choice in the presence of correlations. Given this enhanced-by-consistency readout, we estimated that correlations produced a behavioral benefit that compensated or overcame their detrimental information-limiting effects. These results call for a re-evaluation of the role of correlated neural activity, and suggest that correlations in association cortex can benefit task performance even if they decrease sensory information.

Introduction

The collective activity of a population of neurons, beyond properties of individual cells, is thought to be 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 downstream to guide behavioral choices[3]. A commonly studied feature of coding in neural populations is noise correlations, that is correlated trial-to-trial variability during repeated presentations of the same stimulus[4,5]. These correlations can take the form of correlations between the spike rates of two individual cells, which we call across-neuron noise correlations. Similarly, neural population activity at a given time in response to a stimulus is often correlated with the activity of the same population at other times, which we refer to as across-time noise correlations.

The impact of across-neuron and across-time correlations has been long debated. Much experimental and theoretical work has proposed that correlations limit the information capacity of a neural population[6–9]. Because these correlations reflect trial-to-trial variability that is shared across neurons or time, the detrimental effect of variability on stimulus decoding cannot be eliminated by averaging the activity of neurons or time points. Noise correlations are therefore proposed to fundamentally constrain what neural networks can compute, and limit the accuracy by which subjects can judge differences between stimuli[6,7]. A reason to suspect that the effect of noise correlations may be more nuanced, however, comes from a separate stream of biophysical and theoretical studies. This line of work has shown that spatially and temporally correlated spiking can more strongly and more reliably propagate signals and drive responses in postsynaptic neural populations[10–14]. However, it remains poorly tested if and how enhanced signal propagation may have a beneficial impact on behavioral discrimination performance.

We reasoned that we could investigate how noise correlations shape behavioral performance in perceptual discrimination by studying at the same time not only how correlations impact the encoding of sensory information, as has been emphasized frequently, but critically also how they impact the reading out of this information by downstream neural circuits to guide behavioral choices.

Correlations of PPC activity limit sensory coding during perceptual discrimination

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). PPC is thought to participate in transforming multisensory signals into behavioral outputs and is thus likely at the interface of encoding and reading out stimulus information. Furthermore, PPC is essential for perceptual discrimination tasks during virtual-navigation[15,16], and its activity has been shown to contain stimulus information that relates to an animal's choices[16–22]. 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 calcium imaging datasets. To study across-time correlations, we used calcium imaging data from a sound localization task[17] 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, real-world 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

86   2/3 neurons was imaged simultaneously. This task had the same stimulus category (sound location)
87   presented throughout the trial and thus was well suited to probe across-time correlations in neural
88   activity. To study across-neuron correlations, we used calcium imaging data from an evidence
89   accumulation task[19] in which ~350 layer 2/3 neurons were imaged simultaneously per session, thus
90   providing a large population size to test across-neuron correlations. In the evidence accumulation
91   task, during virtual navigation, mice were presented with six temporally separate visual cues on the
92   left or right walls of a T-maze (Fig. 1f). Mice reported which side had more cues by turning in that
93   direction at the T-intersection (84.5 ± 1.6% correct). We categorized the visual stimuli as having more
94   total left or right cues. For all analyses, we focused on the period toward the end of the T-stem before
95   the mouse had reported its choice and after it had received sensory information; this is a window in
96   which neural activity may carry sensory information that is used to inform the choice (Methods).

97

98   We tested how noise correlations impacted the encoding of sensory information in population
99   activity. For the sound localization task, we quantified across-time correlations by computing the
100   Pearson correlation between the activity of one neuron at one time point and the activity of another
101   neuron at another time point, for each pair of neurons during trials with the same sound location
102   category. Consistent with previous reports[17,19,23], noise correlations were present even at lags longer
103   than 1 s (Fig. 1b). Since many neurons exhibited activity selective for distinct trial types, with
104   different neurons active at distinct time points in the trial (Fig. 1c), we were able to decode the
105   stimulus category (left vs. right sound location) significantly above chance from pairs of temporally
106   offset instantaneous population activity vectors (Methods; Fig. 1d). To evaluate how across-time
107   correlations affected the encoding of stimulus category information, we shuffled instantaneous
108   population activity vectors across trials of the same stimulus category, independently at each time
109   point. This shuffle destroyed within-trial temporal relationships between population vectors while
110   preserving instantaneous population activity (Methods). Importantly, stimulus decoding performance
111   was higher when correlations were disrupted by shuffling, indicating that across-time correlations
112   limited sound category information in neural activity (Fig. 1d). Similarly, for the evidence
113   accumulation task, we computed across-neuron Pearson correlations between the activity of a pair of
114   neurons in a given time window for trials with the same stimulus category (Fig. 1g). Because neurons
115   exhibited activity selective to the stimulus category (more left or right cues) (Fig. 1h), we were able
116   to decode the stimulus category from the population activity vector (Fig. 1i). We disrupted across-
117   neuron correlations by dividing the neural population into two non-overlapping randomly-chosen
118   pools of neurons of equal size and shuffling the trial labels separately for each pool within trials of
119   the same stimulus category. This shuffle scrambled within-trial relationships between pools and thus
120   removed correlations between pools (Methods). Importantly, decoding accuracy was higher when
121   correlations were disrupted by shuffling than in the unshuffled data with correlations intact, showing
122   that across-neuron correlations were also information limiting (Fig. 1i).

123

124   Whether noise correlations limit the information encoded by a neural population depends on how
125   they relate to signal correlations (the correlations between trial-averaged responses to individual
126   stimuli). We parametrized the relationship between signal and noise correlations in our data by
127   computing the angle[24,25], which we termed γ, between the axis of largest variation of neural responses
128   at fixed stimulus (noise correlations axis) and the axis of largest stimulus-related variation (signal
129   correlations axis) (Methods and Fig. 1e-j, inset). This angle was variable, but small on average
130   (γ=0.111π ± 0.001π and 0.098π ± 0.001π for the sound localization and evidence accumulation

3

131   datasets respectively, Fig. 1e, j). The smaller γ, the more noise correlations made the discrimination
132   of different stimuli more difficult by increasing the overlap between the stimulus-specific response
133   distributions to different stimuli (Fig. 1e, j, and Extended Data Fig. 2a, b). High γ values,
134   corresponding to information-enhancing correlations reported in some previous studies[26,27], were
135   present also in our dataset, but were not predominant (Fig. 1e, j).
136
137   Often it has been assumed that perceptual discrimination performance increases proportionally with
138   the amount of sensory information encoded in neural activity[28]. Accordingly, the fact that correlations
139   limit sensory information has been interpreted to imply that correlations hinder the ability to
140   discriminate sensory stimuli[6,29]. Specifically, across-time correlations have been proposed to limit the
141   benefit of integrating noisy information over time as is commonly considered for a speed-accuracy
142   tradeoff[30,31]. Further, across-neuron correlations are thought to lessen the benefit of averaging noisy
143   information across populations of neurons[6,7,9]. If correlations are detrimental to perceptual behaviors,
144   one would intuitively expect noise correlations to be lower when animals make correct choices and
145   higher when animals make errors. Contrary to this expectation, both across-time and across-neuron
146   correlations were higher in correct trials than in error trials (Fig. 1b, g), despite limiting information
147   in population activity. This finding leads to the paradoxical suggestion that correlations limit
148   information encoded by a neural population but at the same time may be beneficial for making
149   accurate choices.
150
151   **A simple model of encoding and readout to test how correlations affect task performance**
152   To reconcile the observations of correlations being information limiting, but at the same time being
153   higher in correct trials, we developed a simple mathematical model that incorporated both the
154   encoding of stimulus information and the readout of this information to form a choice. We compared
155   two alternative views of how information in population activity may be used to perform a stimulus
156   discrimination task. In the traditional view, the accuracy of a choice is proportional to the amount of
157   information in a neural population, and thus information-limiting correlations constrain task
158   performance. Alternatively, we considered that a choice could depend on both stimulus information
159   and features of neural activity that emerge from correlations, in particular the consistency of
160   information across time and neurons in a population.
161
162   We simulated a perceptual discrimination task with two possible stimuli that had to be converted into
163   two possible corresponding choices ($c = 1$ for $s = 1$ and $c = -1$ for $s = -1$) (Fig. 2a). We simulated
164   neural activity generically for two features r_1 and r_2, which could represent neural activity at different
165   points in time (i.e. for across-time correlations) or activity of different neurons (i.e. for across-neuron
166   correlations) (Fig. 2c). For each feature alone, higher-than-average values were simulated to indicate
167   one sensory stimulus ($s = 1$), and lower-than-average values were simulated to indicate the opposite
168   sensory stimulus ($s = -1$), meaning that the two features showed positive signal correlations. In
169   addition, we simulated noise correlations between r_1 and r_2, such that the noise axis was closely
170   aligned to the signal axis ($\gamma = 0.08\pi$, close to the mean values of both experimental datasets). Because
171   of the close signal-noise alignment, noise correlations increased the overlap between the stimulus-
172   specific response distributions (cf. blue and green ellipses in Fig. 2c, left panel vs right panel) and
173   decreased the stimulus information encoded by the two features jointly (Fig. 2d).
174

175 We then considered the readout stage of our model. Commonly, the choice is expected to follow the
176 decoded stimulus. However, because of the apparent importance of correlations for accurate choices
177 in our experimental data, we hypothesized that the readout of stimulus information to inform choices
178 might utilize particular aspects of population activity imposed by correlations[1]. Intuitively,
179 correlations imply that there is greater consistency in the neural population representations (Fig. 2e).
180 In our model, we defined consistency as a single-trial measure of similarity between the stimuli that
181 are decoded from features r_1 and r_2 separately. In our simulations, a stimulus representation in a trial
182 was classified as consistent when features r_1 and r_2 both signaled the same stimulus (i.e. both features
183 higher than average and thus both signaling $s = 1$, top right quadrant of panels in Fig. 2c; or both
184 lower than average and both signaling $s = -1$, bottom left quadrant of panels in Fig. 2c).

186 We simulated binary task discrimination choices with two alternative readout models, formulated as
187 logistic models of the dependence of choice on several features of stimulus encoding. In the first
188 model, that we termed the consistency-independent readout, the simulated choice in each trial
189 depended only on the stimulus decoded from the two features jointly (Fig. 2f, right panel). This case
190 followed the traditional assumption[29] that the choice reflects the stimulus decoded from the full
191 population activity (Fig. 2f, readout map superimposed to left and middle panels). Note that, in our
192 experimental data, we did not observe a perfect correspondence between the stimulus decoded from
193 neural activity and the choice of the mouse (the fraction of times the mouse's choice matched the
194 decoded stimulus was $61.0\% \pm 0.2\%$ in the sound localization dataset and $91.1\% \pm 0.1\%$ in the
195 evidence accumulation dataset; Supplemental Information S7). Therefore, in the model, we set the
196 average probability that, in a given trial, the choice matched the decoded stimulus, which we termed
197 "readout efficacy", to a value smaller than 100%.

199 In the second readout model, which we termed the enhanced-by-consistency readout, the choice in
200 each trial depended not only on the stimulus decoded from both features jointly, but also on the
201 consistency of the stimulus decoded from the features separately (Methods) (Fig. 2i, right panel). If
202 r_1 and r_2 reported consistent information about the stimulus, this readout was more likely to use the
203 stimulus encoded in neural activity to inform the choice. This effect resulted from positive
204 coefficients assigned to the interaction terms between the decoded stimulus and consistency (see
205 Methods and Fig. 2i, right panel). In other words, the readout efficacy was higher when the two
206 features were consistent (Fig. 2i, readout map superimposed to left and middle panel). Importantly,
207 the average readout efficacy of this model was also smaller than 100%, in agreement with
208 experimental data, and was matched to the readout efficacy of the consistency-independent readout
209 model.

211 For the consistency-independent readout, correlated activity resulted in worse task performance
212 compared to activity in which correlations were absent (Fig. 2h). This was expected because with this
213 readout task performance is directly proportional to the level of stimulus information, with higher
214 information resulting in higher performance. Further, noise correlations were slightly higher on
215 simulated error trials than on correct trials (Fig. 2g), which was notably inconsistent with our
216 experimental observations (Fig.1b, g). When considering the enhanced-by-consistency readout, noise
217 correlations were higher in correct trials than error trials (Fig. 2j), matching our experimental
218 findings. This readout also resulted in higher task performance in the presence of noise correlations
219 than in the absence of correlations (Fig. 2k). Critically, this result shows that the enhanced-by-

220 consistency readout mechanism can reconcile our experimental observations by providing a scenario
221 in which correlations limit information, thus producing more decoding errors, while at the same time
222 providing a benefit to task performance.

223

224 The ability of the enhanced-by-consistency readout to generate a task performance that overturns
225 information encoding deficits was affected by how noise correlations both create consistency and
226 limit information, which are in turn determined by the strength and structure of correlations. We
227 explored the model parameters in depth (Extended Data Fig. 2, Supplementary Information) and
228 found a range of parameters in which the enhanced-by-consistency readout produced higher task
229 performance with correlations even when they limit information. This happened when correlations
230 increased the proportion of trials with correctly decoded and consistent information more than they
231 increased the proportion of trials with incorrectly decoded and consistent information, with respect
232 to uncorrelated data. This condition was met when the angle between signal and noise was small but
233 non-null, as in our experimental data (Extended Data Fig. 2b-e, Supplementary Information, Fig. 1e,
234 j).

235

236 **Correlations and consistency in population activity contribute to behavioral choices**
237 Together, our model results highlight the critical importance of estimating the effect of correlations
238 not only on stimulus encoding, but also on the readout of single-trial activity to inform a choice. We
239 therefore used our experimental measurements of neural activity to test for signatures of an enhanced-
240 by-consistency readout. A key prediction of this readout is that the mouse's performance should
241 depend not only on whether stimulus encoding is correct or incorrect but also on the consistency of
242 stimulus information across neurons and across time. In our experimental data, we defined
243 consistency as the single-trial similarity between the stimuli that are decoded from population activity
244 at different points in time (across-time consistency) or between the stimuli that are decoded from
245 separate neuronal pools in the same time window (across-neuron consistency). For example, across-
246 time consistency would be present in a trial if the population activity at time point 1 signaled the same
247 stimulus category (e.g. s = 1) as the population activity at time point 2. We calculated the mouse's
248 performance for four subclasses of trials, defined by correctness and consistency of the stimulus
249 encoded in neural activity in a given trial. In both experimental datasets, the mouse's task
250 performance was higher for trials with correct stimulus decoding than for incorrectly decoded trials,
251 suggesting that the stimulus information carried by PPC neurons was used to inform behavioral
252 choices (Fig. 3a, d, right panels). Also, the mouse's task performance was higher for trials with
253 consistent information across time or across neurons, supporting the idea that the consistency of
254 neural population information was critical for accurate choices (Fig. 3a, d, left panels). Critically, in
255 trials with correctly decoded stimulus information, the mouse's task performance was higher when
256 information was consistent than when it was inconsistent, both across neurons and across time.
257 Further, in trials with incorrectly decoded stimulus information, task performance was lower on
258 consistent trials than on inconsistent trials (Fig. 3a, d, right panels). These findings indicate that the
259 stimulus information encoded in PPC was read out in a manner that amplified the effect of the
260 encoded stimulus information on choice in consistent trials, both when the information was correct
261 or incorrect.

262

263 To rule out that differences in a mouse's task performance between consistent and inconsistent trials
264 might be due to higher stimulus information in consistent trials, we sorted trials according to the

265 stimulus information level and verified that performance in correctly (respectively incorrectly)
266 decoded trials was still higher (respectively lower) when information was consistent across neurons
267 or across time (Methods, Supplementary Information, and Extended Data Fig. 4c-e, g-i).
268
269 We further examined the possibility of an enhanced-by-consistency readout by developing an
270 analytical understanding of how choices were made by the mouse. We used logistic regression to
271 relate PPC activity to the mouse's choices. We expressed a mouse's choice on a given trial as a
272 function of features of the recorded neural activity: the stimulus decoded from the full PPC population
273 activity and its interaction with the across-time or across-neuron consistency (Methods). In addition,
274 in the choice regression, we included a predictor for the experimenter-defined stimulus presented to
275 the mouse and a bias term. These two terms accounted for the stimulus-related and stimulus-unrelated
276 information carried by sources other than the recorded neurons, such as non-recorded neurons. The
277 inclusion of these terms therefore allowed us to test how much the stimulus information in the
278 recorded neural activity explained the mouse's choice beyond what could be explained by other
279 sources.
280
281 The average regression coefficients for the stimulus decoded from neural activity and the consistency-
282 dependent interaction terms were positive (Fig. 3c, f, left), indicating that the stimulus information in
283 neural activity, including its consistency, impacted the mouse's choice. Positive coefficients for the
284 consistency-dependent terms indicates that the readout of PPC activity performed similarly to the
285 enhanced-by-consistency readout model from Figure 2: that is, the probability that the choice matched
286 the stimulus decoded from neural activity was higher in the presence of consistency (Fig. 3c, f, right).
287 We tested the specific contribution of the neural-based predictors in explaining the mouse's choices
288 by fitting the logistic regression after shuffling the values of these predictors across trials (Fig. 3b, e).
289 Shuffling all neural-based predictors ("No Neural" choice regression in Fig. 3b, e) made it harder to
290 predict a mouse's choice, again demonstrating that neural activity contributed to choices. Moreover,
291 shuffling only the neural consistency values ("No Cons" choice regression in Fig. 3b, e), while
292 leaving intact the stimulus decoded from neural activity and the experimenter-determined stimulus,
293 resulted in worse predictions of a mouse's choices. Neural consistency thus seemingly contributed to
294 the mouse's choice.
295
296 To rule out that the modulation of the readout by consistency might just reflect differences in overall
297 stimulus information levels between consistent and inconsistent trials, we verified that consistency
298 provided a similar contribution to predicting a mouse's choices when we used a more sophisticated
299 logistic regression in which we included the magnitude of the stimulus information, instead of only
300 the identity of the decoded stimulus (Supplementary Information and Extended Data Fig. 4b, f).
301 Further, to control for and discount potential contributions from running-related neural activity, we
302 verified that neural consistency also contributed to predicting choices when adding to the regression
303 the consistency of the mouse's running speed and direction (Supplementary Information and
304 Extended Data Fig. 5c-e).
305
306 **The enhanced-by-consistency readout of PPC activity can overturn the information-limiting**
307 **effect of correlations**
308 Thus far, our results show that across-time and across-neuron consistency in the experimental data
309 have an impact on a mouse's choices. We next examined the implications of this finding for mouse

7

310 task performance, in the presence or absence of experimentally-measured information-limiting
311 correlations. Because correlations cannot be removed with current experimental approaches, we
312 instead created a set of simulated choices using the experimentally-fit logistic choice regression from
313 Figure 3 (Methods). We used either trials with real neural activity or trials with neural activity
314 shuffled to disrupt across-time or across-neuron correlations as input to the experimentally-fit choice
315 regression. As a reminder, shuffled activity had higher stimulus information (Fig. 1d, i and Fig. 4b,
316 d, left) and lower across-time and across-neuron consistency (Fig. 4b, d, right). We used the simulated
317 choices to provide an estimate of how well the mouse would have performed on the task with and
318 without correlations present (Fig. 4a).
319
320 We focused on the contribution of the recorded neural population to task performance. This
321 contribution was isolated by computing task performance from choices simulated using all predictors
322 extracted from experimental data, and then subtracting the task performance computed from choices
323 simulated after shuffling neural predictors values across trials (Methods). In the sound localization
324 dataset, we estimated that the recorded ~50 neurons increased task performance by ~3.5% (Fig. 4c).
325 Strikingly, the amount the recorded neurons increased task performance was greater when
326 correlations were intact than when correlations were removed by shuffling (Fig. 4c). This is especially
327 of note because stimulus information in the recorded neurons was lower with correlations intact (Fig.
328 4b, left), indicating that the enhanced-by-consistency feature of the experimentally-fit readout was
329 able to overcome the information-limiting effect of across-time correlations and improve task
330 performance. In the evidence accumulation data, the ~350 recorded neurons contributed ~25% of task
331 performance (Fig. 4e). This estimate was similar when correlations were intact or disrupted (Fig. 4e),
332 even though the stimulus information was lower when correlations were intact (Fig. 4d, left). This
333 result suggests that across-neuron correlations did not hinder task performance despite having an
334 information-limiting effect. Together, our results strongly indicate that correlations had an advantage
335 in the readout of stimulus information that overcame or compensated for their information-limiting
336 effects.
337
338 These results incorporate the overall impact of correlations on task performance by combining the
339 effects of the encoding and readout stages. To analyze the specific contribution of the readout, we
340 again simulated choices from the experimentally-fit choice regression, except we equalized the
341 stimulus information in the correlated and shuffled populations by selecting subsets of trials having
342 the same fraction of correctly decoded stimuli. With such matching, the correlated and shuffled trials
343 differed only in their neural consistency (Fig. 4f, h). For both datasets, the estimated contribution to
344 task performance of the recorded neurons was higher when correlations were intact than when they
345 were disrupted (Fig. 4g, i). This result shows that the readout of PPC activity was more efficient in
346 extracting information from correlated data and can thus improve task performance.
347
348 It is interesting to note that our results here indicate that the readout of stimulus information from
349 PPC activity is suboptimal. From the ~50 neurons recorded in the sound localization task, we were
350 able to decode the stimulus at ~60% correct (Fig. 4b, left). This population therefore could have
351 increased task performance by ~10% above chance if this stimulus information were read out
352 optimally. However, this population only increased task performance by ~3.5% (Fig. 4c). Similarly,
353 for the ~350 neurons recorded during the evidence accumulation task, we were able to decode the
354 stimulus at ~80% correct (Fig. 4d, left), and thus these neurons could have increased task performance

355   by ~30% above chance if all their stimulus information was read out. However, we estimated that
356   these neurons only increased task performance by ~25% (Fig. 4e). Therefore, in both datasets, the
357   recorded neurons apparently increased task performance by a smaller amount than would have been
358   possible if all their stimulus information was converted into choice. This result indicates that stimulus
359   information in PPC is read out for behavior, but not optimally or entirely.

361   We further considered whether the experimentally-fit choice readout, which shows characteristics
362   compatible with our enhanced-by-consistency readout model, is particularly well suited for reading
363   out the information contained in the recorded PPC activity, or if an alternative consistency-
364   independent readout would yield more correct choices. Our comparison was similar in spirit to that
365   of Figure 2f-k, except using experimental data. We first generated a simulated set of choices by
366   feeding real neural activity to the experimentally-fit regression that incorporated across-time and
367   across-neuron consistency. We also generated a second set of simulated choices using an alternative
368   choice regression that included only the decoded stimulus predictor, regardless of its consistency
369   (Fig. 4j, k, m). For fairness of comparison, the coefficients for this second choice regression were
370   selected to yield the same readout efficacy as for the experimentally-fit regression (Methods). The
371   estimated contribution of the recorded neurons to task performance was higher with the
372   experimentally-fit choice regression that used consistency than with the consistency-independent
373   choice regression with matched readout efficacy (Fig. 4l, n). These experimental findings, in
374   agreement with model predictions (Extended Data Fig. 3 and Supplementary Information), suggest
375   that the enhanced-by-consistency readout is well suited for forming behavioral choices in the
376   presence of information-limiting noise correlations, such as those found in PPC.

378   We replicated all analyses of across-neuron correlations in the sound localization dataset, that had
379   smaller numbers of neurons per session than the evidence accumulation dataset. We found that
380   across-neuron correlations enhanced the readout of stimulus information and task performance, even
381   though these correlations limited information (Supplementary Information, and Extended Data Fig.
382   6). We could not test the effect of the across-time correlations in the evidence accumulation dataset
383   because of challenges in dealing with time-varying stimulus evidence in that task.

385   Finally, in the sound localization experiments, we also had experimental data from auditory cortex.
386   Interestingly, in auditory cortex, relative to PPC, we observed weaker noise correlations, a smaller
387   information-limiting effect of correlations, and a lower impact of consistency on the reading out of
388   information (Supplementary Information and Extended Data Fig. 6, 7). Therefore, sensory and
389   association areas may differ in their levels of correlations, as we reported previously[17], such that an
390   enhanced-by-consistency readout may be most beneficial for reading out PPC activity and less
391   beneficial for reading out auditory cortex activity.

Discussion
394   Our results show that noise correlations limit information at the encoding stage, but surprisingly they
395   also enhance consistency in neural codes, which improves the readout stage. The trade-off of these
396   two effects defines the overall impact of correlations on task performance. We found that, strikingly,
397   noise correlations can enhance task performance despite limiting the information capacity of a neural
398   population.

400 Much work has emphasized that the information limiting effect of correlations in sensory areas may
401 be a major bottleneck for behavioral performance[1,6,7], even if the information-limiting effects of
402 correlations that inevitably arise from common inputs may be partly reduced by anti-correlated
403 fluctuations of excitatory and inhibitory synaptic conductances[32,33]. A largely separate set of
404 theoretical and biophysical work has alternatively proposed that correlations improve the propagation
405 of neural activity by mechanisms at the single neuron and network levels[10–12]. However, empirical
406 support for a role of correlations in facilitating the readout of population information to aid behavior
407 has been limited thus far[34,35]. First, biophysical work on signal propagation has not been connected
408 to information coding, as it has not specified whether the transmitted activity is informative and has
409 not distinguished between propagation of correct or incorrect information[13], and it has been seldom
410 connected to behavior. Second, although some studies have suggested a beneficial role of correlations
411 by reporting higher correlation levels during correct behaviors[34,36], these effects have not been
412 reported when correlations limit information encoding. Further, in the analysis of our model we
413 occasionally found conditions in which information-limiting noise correlations were stronger in
414 correct trials, without them having a positive effect on behavioral performance, suggesting that
415 reports of higher correlations in correct trials is indicative that consistency affects readout, but it is
416 not sufficient to imply a behavioral performance advantage of correlations (see Extended Data Figure
417 2 and Supplemental Information S8). Thus, from previous work it may be challenging to conclude
418 whether the advantages of correlations for signal propagation can overturn their information-limiting
419 effect.
420
421 Here we developed a formalism that bridges between the two separate research streams of information
422 encoding and signal propagation, to address how the negative encoding effects and the positive
423 readout effects of correlations intersect. Remarkably, in our data the advantages of correlations for
424 signal propagation were large enough to compensate and even overcome the negative encoding
425 effects. We unraveled a readout computation that could not be inferred from previous biophysical
426 activity propagation studies: that when correlations are information limiting the interaction between
427 the readout and encoding of stimulus information is such that it amplifies correctly decoded sensory
428 information more than incorrectly decoded information.
429
430 Our approach provides a novel and generally applicable framework to parametrically dissect the
431 contribution of correlated neural activity to perceptual behaviors. We anticipate that this approach
432 can be applied to different tasks and brain areas. Our initial observations comparing primary sensory
433 and association cortices suggests that the best trade-off between the effect of correlations on encoding
434 and readout may vary depending on the area. In sensory cortices, low correlations and a readout
435 insensitive to consistency may be advantageous to encode rapidly changing and high dimensional
436 sensory features, in order to represent many features of the sensory environment regardless of if they
437 are used for the immediate behavior at hand. This view is compatible with reciprocal relationships
438 between noise correlation levels and behavioral performance in sensory cortices[1]. In contrast,
439 association cortices that are closer to behavioral output may only need to encode a moderate amount
440 of behaviorally relevant sensory information but this information should have a strong impact on
441 behavior. In these areas, the price to be paid for reducing encoded information is smaller and the price
442 to be paid for losing signal in the propagation to behavior is higher. Thus, in association areas, the
443 best tradeoff may be to have some redundancy in the neural representation coupled with a readout
444 mechanism that uses this redundancy to enhance signal propagation to inform choice, as we found

445 here. Further, we anticipate that our general formalism will allow us to design direct causal tests of
446 the actual readout used in the brain during perceptual discrimination tasks. We propose that our
447 formalism together with holographic perturbations[37] will allow us to test the behavioral role of
448 correlations and consistency across neurons and across time.

450 Noise correlations can reflect interactions between cells, shared covariations due to common inputs,
451 general fluctuations in behavioral state or network excitability, or even variations with stimuli within
452 the same category[2]. This shared variability, independently of its origin, acts as noise (that cannot be
453 reduced by integrating information over more cells or longer times) from the decoding point of view,
454 but helps signal propagation by generating more consistent neural representations. Thus, our
455 conclusions hold independently of the specific biophysical origins of the observed noise correlations.

457 Most studies of neural coding implicitly or explicitly assume that the readout of sensory information
458 is optimal and interpret neural codes with higher sensory information as being more relevant for
459 perception[6,28,29]. Part of the reason for this assumption is that the presence and shape of non-optimality
460 are unknown. If the readout is not optimal, then neural codes with higher information are not
461 necessarily the most relevant ones for perception. Our study suggests (see Supplemental Information)
462 that a sub-optimality in the readout of our PPC data exists, because stimulus information in neural
463 activity seems not to be used in its entirety to produce accurate behavioral choices, as also predicted
464 by other studies[38,39]. Our work provides a measure of both the nature of readout non-optimality and
465 its implication for the behavioral relevance of a neural code. Previous work on the optimality of
466 readouts has examined whether a decoder of correlated population activity could be trained sub-
467 optimally to decode separately single-cell data and then join together their evidence[40–42]. Several of
468 these studies have shown that even relatively simple decoders trained on single cells can decode
469 stimulus information from population activity, so that correlations among cells do not greatly
470 complicate the extraction of information. However, unlike ours, these previous studies did not
471 examine the effect of correlations on determining behavioral choices. Taken together, emerging
472 evidence suggests that correlations do not necessarily greatly complicate the decoding of sensory
473 information and may offer advantages for turning sensory information into appropriate behavioral
474 choices.

483 **Author Contributions**
484 SP and CDH conceived and supervised the study. CAR and ASM acquired the experimental data.
485 MV and GP performed the computational work. SP, CDH, MV and GP wrote the manuscript, with
486 feedback from all authors.

References

1. 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).

2. Kohn, A., Coen-Cagli, R., Kanitscheider, I. & Pouget, A. Correlations and Neuronal Population Information. Annu. Rev. Neurosci. 39, 237–256 (2016).

3. 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).

4. Gawne, T. J. & Richmond, B. J. How Independent Are the Messages Temporal Cortical Neurons? by Adjacent Inferior. The Journal of Neuroscience 73, (1993).

5. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. Nature Reviews Neuroscience 7, 358–366 (2006).

6. Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature 370, 140–143 (1994).

7. Moreno-Bote, R. et al. Information-limiting correlations. Nat. Neurosci. 17, 1410–1417 (2014).

8. Bartolo, R., Saunders, R. C., Mitz, A. R. & Averbeck, B. B. Information limiting correlations in large neural populations. J. Neurosci. 2072–19 (2020).

9. Rumyantsev, O. I. et al. Fundamental bounds on the fidelity of sensory cortical coding. Nature 1–6 (2020). doi:10.1038/s41586-020-2130-2

10. Diesmann, M., Gewaltig, M. O. & Aertsen, A. Stable propagation of synchronous spiking in cortical neural networks. Nature 402, 529–533 (1999).

11. Zandvakili, A. & Kohn, A. Coordinated Neuronal Activity Enhances Corticocortical Communication. Neuron 87, 827–839 (2015).

12. Salinas, E. & Sejnowski, T. J. Correlated neuronal activity and the flow of neural information. Nature Reviews Neuroscience 2, 539–550 (2001).

13. Zylberberg, J., Pouget, A., Latham, P. E. & Shea-Brown, E. Robust information propagation through noisy neural circuits. PLoS Computational Biology 13(4), e1005497 (2017).

14. Abeles, M. Corticonics : neural circuits of the cerebral cortex. (Cambridge University Press, 1991).

15. 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).

16. Minderer, M., Brown, K. D. & Harvey, C. D. The Spatial Structure of Neural Encoding in Mouse Posterior Cortex during Navigation. Neuron 102, 232-248.e11 (2019).

17. Runyan, C. A., Piasini, E., Panzeri, S. & Harvey, C. D. Distinct timescales of population coding across cortex. Nature 548, 92-96 (2017). doi:10.1038/nature23020

18. Hanks, T. D. et al. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature 520, 220–223 (2015).

19. Morcos, A. S. & Harvey, C. D. History-dependent variability in population dynamics during evidence accumulation in cortex. Nat. Neurosci. 19, 1672–1681 (2016).

20. 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).

21. 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).

22. Goard, M. J., Pho, G. N., Woodson, J. & Sur, M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife 5, (2016).

23. Akrami, A., Kopec, C. D., Diamond, M. E. & Brody, C. D. Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature 554, 368–372 (2018).

24. Averbeck, B. B. & Lee, D. Effects of noise correlations on information encoding and decoding. J. Neurophysiol. 95, 3633–3644 (2006).

25. 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).

26. Romo, R., Hernández, A., Zainos, A. & Salinas, E. Correlated neuronal discharges that increase coding efficiency during perceptual discrimination. Neuron 38, 649–657 (2003).

27. Reich, D. S., Mechler, F. & Victor, J. D. Independent and redundant information in nearby cortical neurons. Science 294, 2566–2568 (2001).

28. Gold, J. I. & Shadlen, M. N. Neural computations that underlie decisions about sensory stimuli. Trends Cogn. Sci. 5, 10–16 (2001).

29. Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. Proc. Natl. Acad. Sci. U. S. A. 112, E6973–E6982 (2015).

30. Zariwala, H. A., Kepecs, A., Uchida, N., Hirokawa, J. & Mainen, Z. F. The limits of deliberation in a perceptual decision task. Neuron 78, 339–351 (2013).

31. Mazurek, M. E. & Shadlen, M. N. Limits to the temporal fidelity of cortical spike rate signals. Nat. Neurosci. 5, 463–471 (2002).

32. Renart, A. et al. The asynchronous state in cortical circuits. Science 327, 587–590 (2010).

33. Ecker, A. S. et al. Decorrelated neuronal firing in cortical microcircuits. Science 327, 584–587 (2010).

34. 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).

35. Histed, M. H. & Maunsell, J. H. R. Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. Proc. Natl. Acad. Sci. U. S. A. 111, E178-E187 (2014).

36. Francis, N. A. et al. Small Networks Encode Decision-Making in Primary Auditory Cortex. Neuron 97, 885-897 (2018).

37. Emiliani, V., Cohen, A. E., Deisseroth, K. & Häusser, M. All-optical interrogation of neural circuits. Journal of Neuroscience 35, 13917–13926 (2015).

38. Stringer, C., Michaelos, M. & Pachitariu, M. High precision coding in visual cortex. bioRxiv 679324 (2019). doi:10.1101/679324

39. O'Connor, D. H. et al. Neural coding during active somatosensation revealed using illusory touch. Nat. Neurosci. 16, 958–965 (2013).

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).

577   42.   Karpas, E., Maoz, O., Kiani, R. & Schneidman, E. Strongly correlated spatiotemporal
578   encoding and simple decoding in the prefrontal cortex. bioRxiv 693192 (2019).
579   doi:10.1101/693192
580   43.   Berens, P. CircStat : A MATLAB Toolbox for Circular Statistics . J. Stat. Softw. 31, 1–21
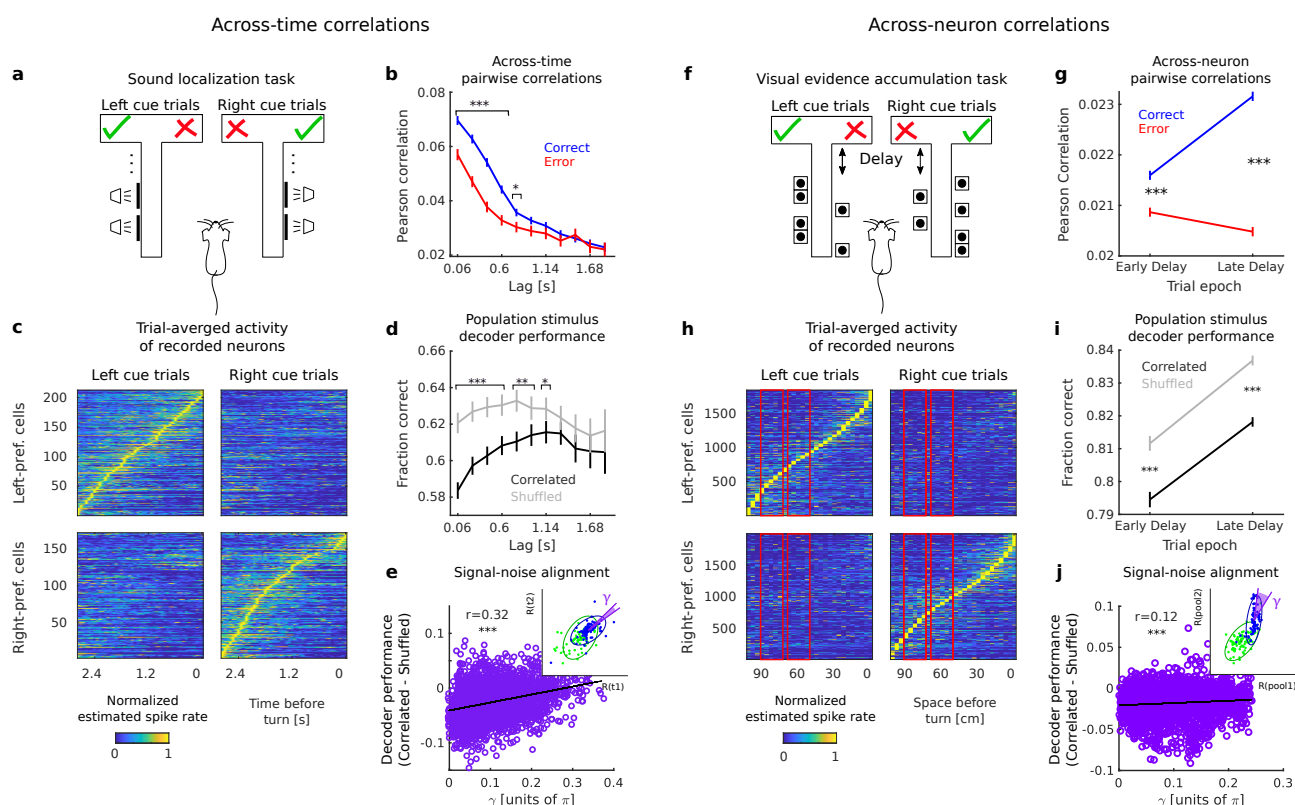581   (2009).
582

Figure 1: Across-time and across-neuron correlations in PPC responses limit sensory coding during perceptual discrimination tasks. a, Schematic of the sound localization task. Left/right sound category (speaker symbols) indicate the rewarded side of the maze (checkmark). **b**, Pearson's correlations in time-lagged activity of all simultaneously-recorded pairs of PPC cells at fixed stimulus, computed separately for correct and error trials. Results are averaged across all time point pairs separated by a given lag. Data are mean ± s.e.m. across all imaged cell pairs. ***P<0.001; *P<0.05, two-sided permutation test. **c**, Normalized mean estimated spike rate traces for all the PPC cells imaged during the sound localization task, divided by left-preferring (n=212) and right preferring (n=172) cells, aligned to the turn frame. Traces were normalized to the peak of each cell's mean estimated spike rate on preferred trials and sorted by peak time. **d**, Performance of a decoder of stimulus identity applied to joint population activity at two different time points, separately for real recorded (black) or trial-shuffled (gray) population vectors. Data are mean ± s.e.m. across sessions and time point pairs. ***P<0.001; **P<0.01; *P<0.05, two-sided permutation test. **e**, Correlation between difference in stimulus decoder performance from real and shuffled population vectors and misalignment angle γ between signal and noise. The inset shows an example of γ computation in the reduced 2D space of projected population activity vectors (see Methods). Ellipses indicate 90% confidence interval for best-fit bivariate Gaussian. ***P<0.001, circular-linear correlation[43]. **f**, Schematic of the evidence accumulation task. The rewarded side of the maze (checkmark) is the one identified by the most numerous visual cues (wall segments with black dots patterns). **g**, Pearson's correlations in time-averaged activity of all simultaneously-recorded pairs of PPC cells, computed separately for correct and error trials. Data are mean ± s.e.m. across all imaged cell pairs. ***P<0.001, two-sided permutation test. **h**, As in **b**, for all cells (left-preferring, n=1840; right-preferring: n=2000) imaged during the evidence accumulation task. Note that activity traces are averaged over spatial bins. Red rectangles indicate the trial epochs selected for analyses (Early delay, Late delay). **i**, Performance of a decoder of stimulus identity applied to joint population activity of two randomly-defined non-overlapping neuronal pools, separately for real recorded (black) or within-pool trial-shuffled (gray) population vectors. Data are mean ± s.e.m. across sessions and 100 pairs of neuronal pools. ***P<0.001, two-sided permutation test. **j**, As in **e**. For panels **d-e**, **i-j**, data are first averaged across trial splits.
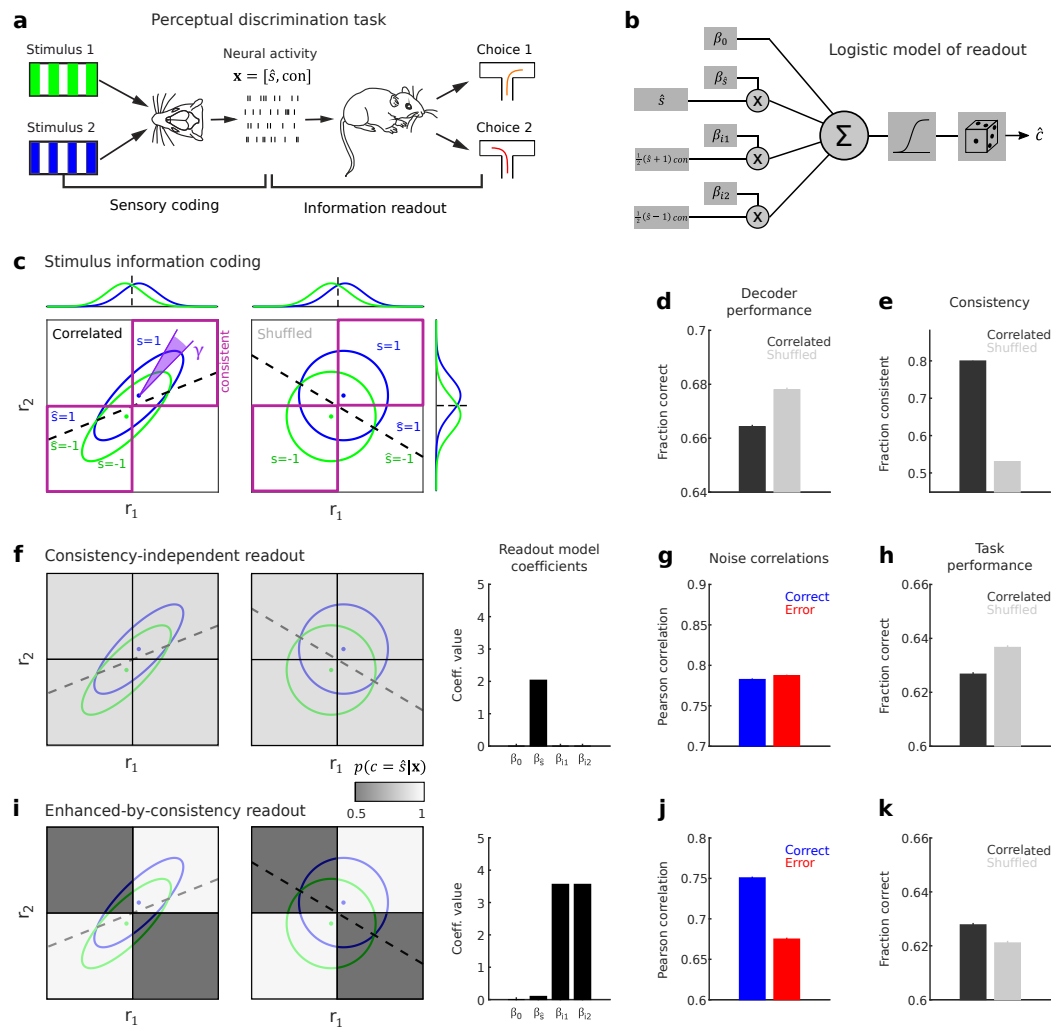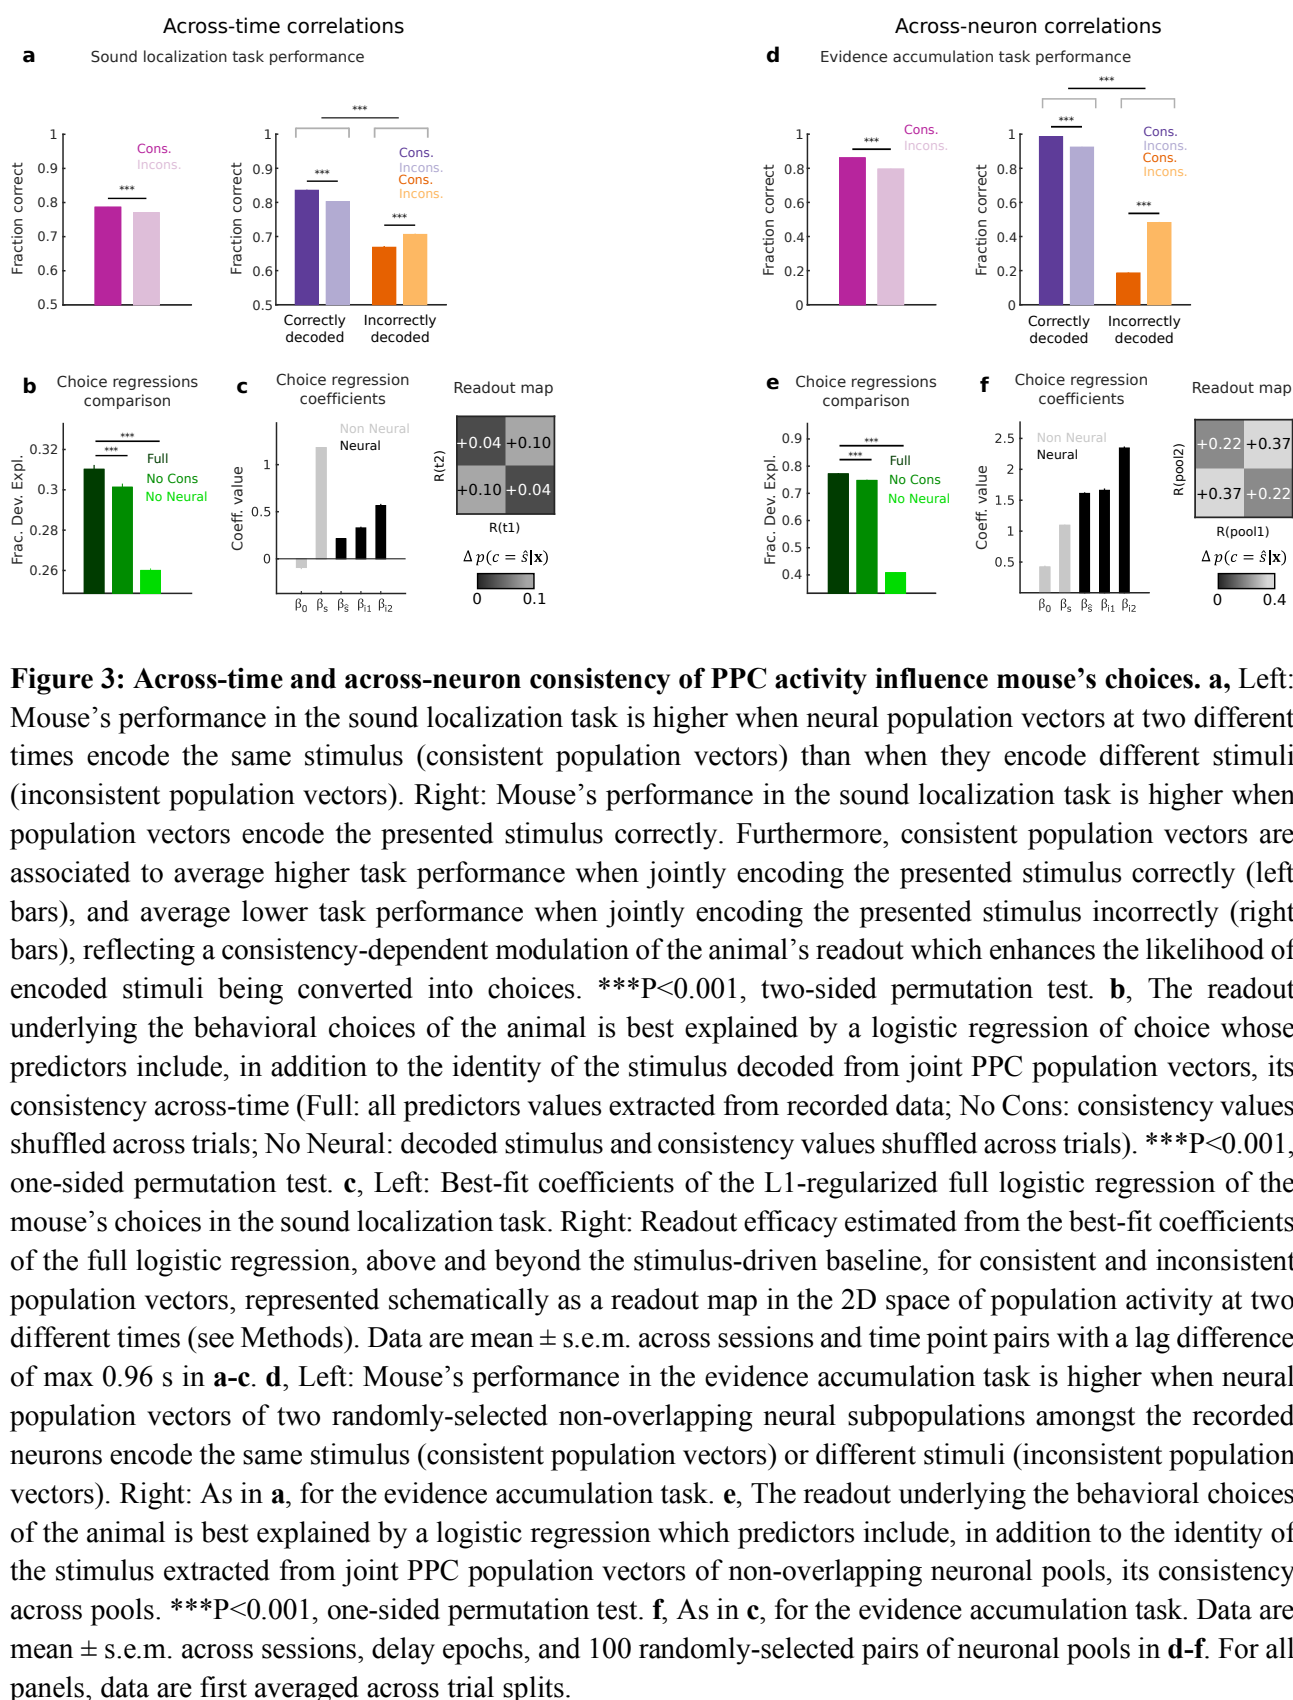
15

612
613
614 **Figure 2: A simple encoding-readout model of stimulus information shows that different readouts**
615 **determine different impacts of correlations on task performance. a,** Schematic showing the two
616 fundamental stages in the information processing chain for sensory perception, which are the constituting parts
617 of the simple model that we propose: sensory coding and information readout. Sensory coding refers to the
618 mapping from external stimuli (Stimulus 1, green; Stimulus 2, blue) to neural activity; information readout
619 refers to the mapping from neural activity to a behavioral choice (Choice 1, turn right; Choice 2, turn left).
620 Note that task-relevant neural activity is here recapitulated by two selected features, being the stimulus
621 predicted when observing a given population activity ($\hat{s}$) and its consistency (con) (see Methods). **b**, Schematic
622 of the readout model that was used to generate synthetic choices from the simulated neural activity patterns
623 (see Methods). Task predictors extracted from simulated neural activity are weighted by corresponding model
624 coefficients, summed and transformed non-linearly through a sigmoid function that outputs the probability of
625 making a given binary choice. The actual choice is generated by randomly drawing from a binomial
626 distribution with probability of single experiment success matching the output of the sigmoid function. **c**,
627 Example of simulated response distributions along two neural features (r_1, r_2) to two stimuli (s=1, s=-1),
628 modelled as bivariate Gaussians with different degree of correlation (correlated: ρ=.8, uncorrelated: ρ=0). For
629 the correlated example, noise correlations axis and signal correlations axis are closely but not perfectly aligned
630 (γ=0.08π). Ellipses indicate 95% confidence intervals. The dashed black line indicates the optimal stimulus
631 decoding boundary. Purple squares indicate regions of the 2D response space in which r_1 and r_2 encode
632 consistent stimulus information, i.e. they both code for the same stimulus (see Methods). Marginal response
633 distributions along r_1 and r_2 and the corresponding decoding boundaries are shown on the sides. **d**, Fraction
634 correct of a linear decoder of stimulus identity applied to simulated bivariate responses as in **c** (nr. trials =

635   1000, nr. simulations = 1000) is higher for uncorrelated responses (correlations limit the encoded stimulus
636   information). **e**, The fraction of trials in which r_1 and r_2 encode consistent stimulus information is higher for
637   correlated responses (correlations increase consistency). **f**, Consistency-independent readout represented as a
638   grayscale map in the r_1-r_2 response space is superimposed to the distributions in **a**. The shade of gray represents
639   the readout efficacy of transformation from encoded stimulus into choice (see Methods), which is the same for
640   consistent and inconsistent trials ($p(c = \hat{s}|\hat{s}, con = 1) = p(c = \hat{s}|\hat{s}, con = 0) = .89$). The corresponding
641   readout model coefficients are shown on the right. **g**, Pearson correlation coefficient of neural responses at
642   fixed stimulus is slightly higher in trials with incorrect choice than in correct trials, when readout is
643   consistency-independent and correlations are as in **c**. **h**, Task performance resulting from the consistency-
644   independent readout of simulated neural activity is higher for uncorrelated responses, thus reflecting the impact
645   of correlations on stimulus coding. **i**, Same as **f**, for enhanced-by-consistency readout (consistency modulation
646   index $\eta = 0.9$, see Methods). The readout efficacy differs between consistent and inconsistent trials ($p(c =$
647     $\hat{s}|\hat{s}, con = 1) = .97;\ p(c = \hat{s}|\hat{s}, con = 0) = .53$). **j**, Pearson correlation coefficient is higher in trials with
648   correct choice than in error trials, when readout is enhanced-by-consistency and correlations are as in **c**. **k**,
649   Task performance resulting from the enhanced-by-consistency readout of simulated neural activity is higher
650   for correlated than for uncorrelated responses, thus reversing the impact of correlations on stimulus coding.
651   Note that in the examples of this figure, model parameters did not match real data and were purely illustrative.
652
653

Figure 3: Across-time and across-neuron consistency of PPC activity influence mouse's choices. a, Left: Mouse's performance in the sound localization task is higher when neural population vectors at two different times encode the same stimulus (consistent population vectors) than when they encode different stimuli (inconsistent population vectors). Right: Mouse's performance in the sound localization task is higher when population vectors encode the presented stimulus correctly. Furthermore, consistent population vectors are associated to average higher task performance when jointly encoding the presented stimulus correctly (left bars), and average lower task performance when jointly encoding the presented stimulus incorrectly (right bars), reflecting a consistency-dependent modulation of the animal's readout which enhances the likelihood of encoded stimuli being converted into choices. ***P<0.001, two-sided permutation test. **b,** The readout underlying the behavioral choices of the animal is best explained by a logistic regression of choice whose predictors include, in addition to the identity of the stimulus decoded from joint PPC population vectors, its consistency across-time (Full: all predictors values extracted from recorded data; No Cons: consistency values shuffled across trials; No Neural: decoded stimulus and consistency values shuffled across trials). ***P<0.001, one-sided permutation test. **c,** Left: Best-fit coefficients of the L1-regularized full logistic regression of the mouse's choices in the sound localization task. Right: Readout efficacy estimated from the best-fit coefficients of the full logistic regression, above and beyond the stimulus-driven baseline, for consistent and inconsistent population vectors, represented schematically as a readout map in the 2D space of population activity at two different times (see Methods). Data are mean ± s.e.m. across sessions and time point pairs with a lag difference of max 0.96 s in **a-c**. **d,** Left: Mouse's performance in the evidence accumulation task is higher when neural population vectors of two randomly-selected non-overlapping neural subpopulations amongst the recorded neurons encode the same stimulus (consistent population vectors) or different stimuli (inconsistent population vectors). Right: As in **a**, for the evidence accumulation task. **e,** The readout underlying the behavioral choices of the animal is best explained by a logistic regression which predictors include, in addition to the identity of the stimulus extracted from joint PPC population vectors of non-overlapping neuronal pools, its consistency across pools. ***P<0.001, one-sided permutation test. **f,** As in **c**, for the evidence accumulation task. Data are mean ± s.e.m. across sessions, delay epochs, and 100 randomly-selected pairs of neuronal pools in **d-f**. For all panels, data are first averaged across trial splits.
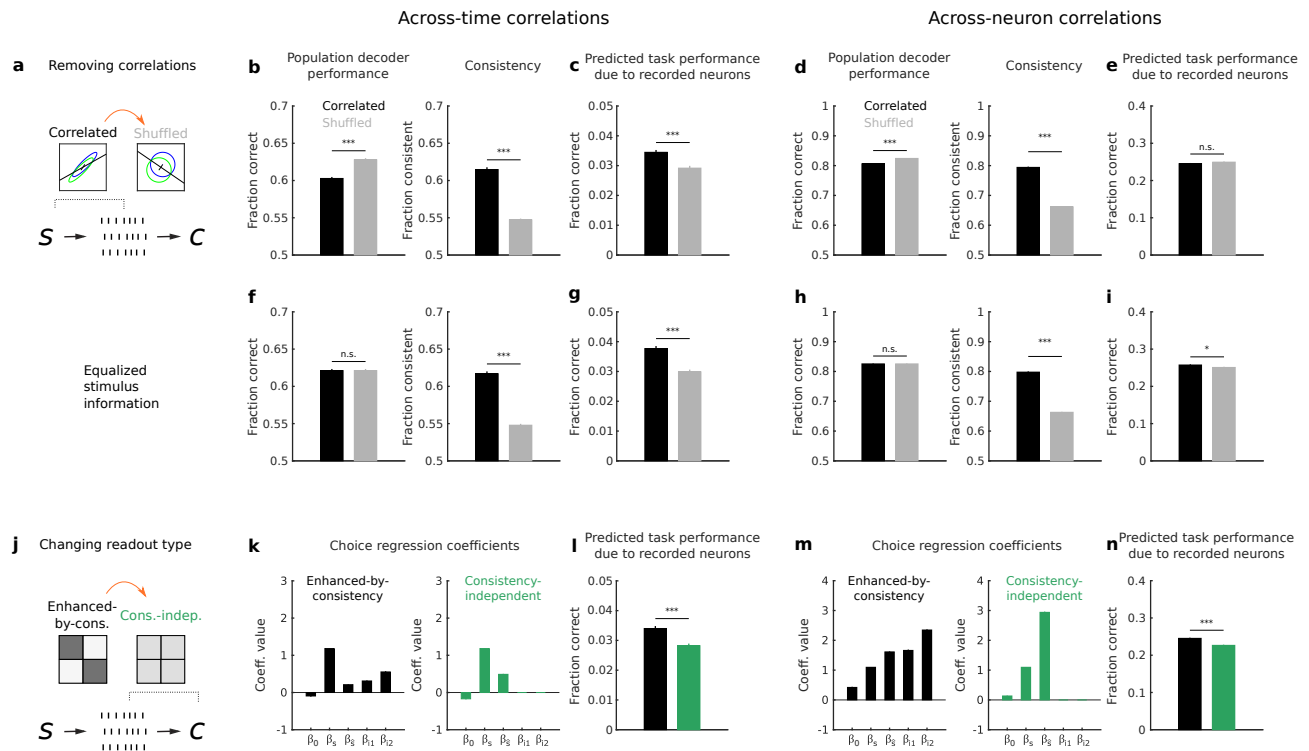
Figure 4: Simulated mouse's choices show that the experimentally-fit enhanced-by-consistency readout improves task performance in the presence of information-limiting correlations in PPC. a, To study the impact on task performance of across-time and across-neuron noise correlations, the best-fit logistic choice regression was fed with patterns of neural activity containing the original correlations, and with artificial neural activity patterns where correlations were removed by shuffling. **b**, Left: Performance of a decoder of stimulus identity applied to joint pairs of time-lagged real recorded population activity vectors (black) and artificial shuffled population activity vectors (grey). Right: Fraction of consistent stimuli decoded separately from real and artificial pairs of time-lagged population activity vectors. **c**, Task performance attributable to the recorded neural activity, computed from choices simulated through the best-fit logistic choice regression (see Methods). Task performance was computed in the presence (black) or absence (grey) of across-time correlations between time-lagged population activity vectors. **d-e**, As in **b-c**, for across-neuron correlations. **f-g**, As in **b-c**, after subsampling trials to equalize stimulus decoder performance between real and shuffled population activity vectors. **h-i**, As in **f-g**, for across-neuron correlations. **j**, To study the impact on task performance of the consistency-dependent predictors coefficients of the best-fit choice logistic regression, the real patterns of neural activity were fed to a consistency-independent logistic regression of choice which was matched for readout efficacy with the best-fit choice regression (see Methods). **k**, Coefficients of the best-fit enhanced-by-consistency choice regression (black) and the consistency-independent choice regression (green). The latter coefficients are obtained after appropriate manipulations of the best-fit coefficients (see Methods). **l**, Task performance attributable to the recorded neural activity, computed from choices simulated through the best-fit logistic choice regression (black) or through the consistency-independent choice regression (green), in the presence of across-time correlations. **m-n**, As in **k-l**, for across-neuron correlations. Data are mean ± s.e.m. across sessions and time point pairs with a lag difference of max 0.96 s in **b-c, f-g** and **k-l**. Data are mean ± s.e.m. across sessions, delay epochs, and 100 randomly-selected pairs of neuronal pools in **d-e, h-i** and **m-n**. ***P<0.001; *P<0.05; n.s. (no statistical significance), two-sided permutation test. For all panels, data are first averaged across trial splits.
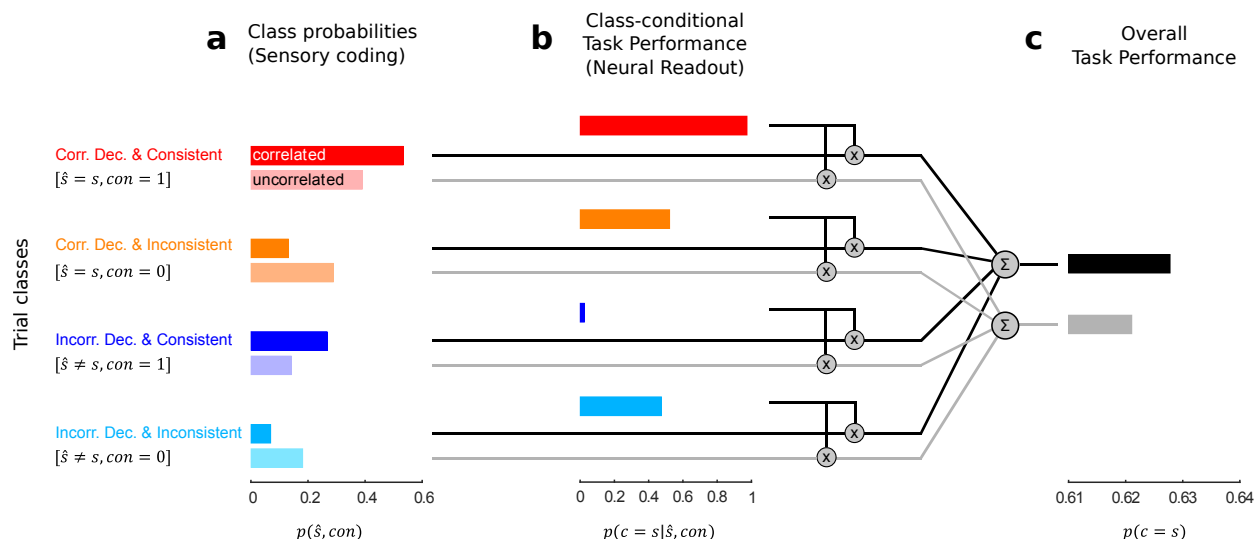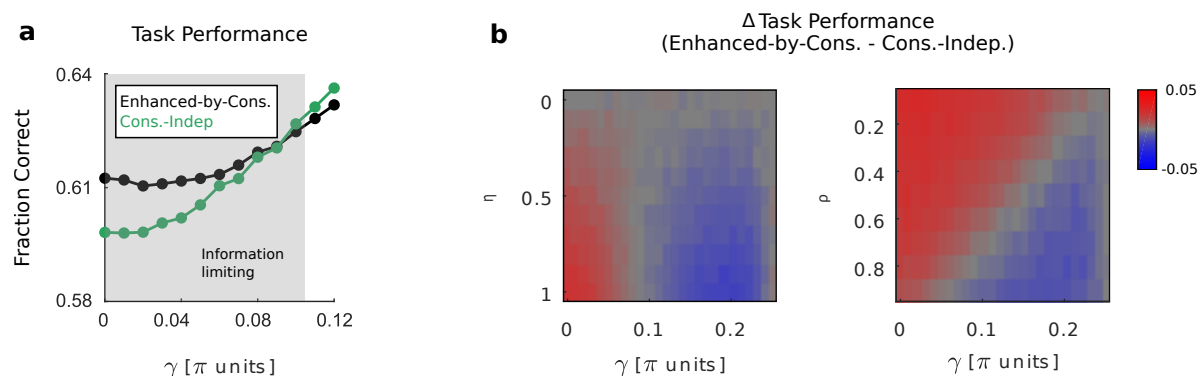
Extended Data Figure 1: Task performance as a sum of encoding-by-readout terms depending on correctness and consistency of encoded stimulus information. a. Probabilities for a set of n=200 trials drawn separately from the correlated (dark colors) and uncorrelated (light colors) distributions sketched in Fig. 2c to belong to one of the four trial classes defined by correctness of decoded stimulus and consistency. Correlations overall decrease the fraction of correctly decoded trials (red + orange bars). At the same time, correlations increase the fraction of consistent trials which encode the correct stimulus identity (red), but also the fraction of consistent trials which encode the incorrect stimulus identity (blue). **b**, Average task performance for each of the four classes of trials in **a**, determined by the enhanced-by-consistency readout model coefficients sketched in Fig. 2i. The enhanced-by-consistency readout is characterized by close-to-random task performance for inconsistent trials, close-to-totally-correct choices for correctly decoded consistent trials and close-to-totally-wrong choices for incorrectly decoded consistent trials. **c**, Overall task performance depends on both on the stimulus decoding correctness and consistency of trials characterized by a given correlation structure and on the class-conditional task performance imposed by the readout model. Note that under hypothesis of enhanced-by-consistency readout it is possible for correlations that decrease the overall fraction of correctly decoded trials to increase the overall fraction of correct choice trials (with respect to the uncorrelated equivalent).

20
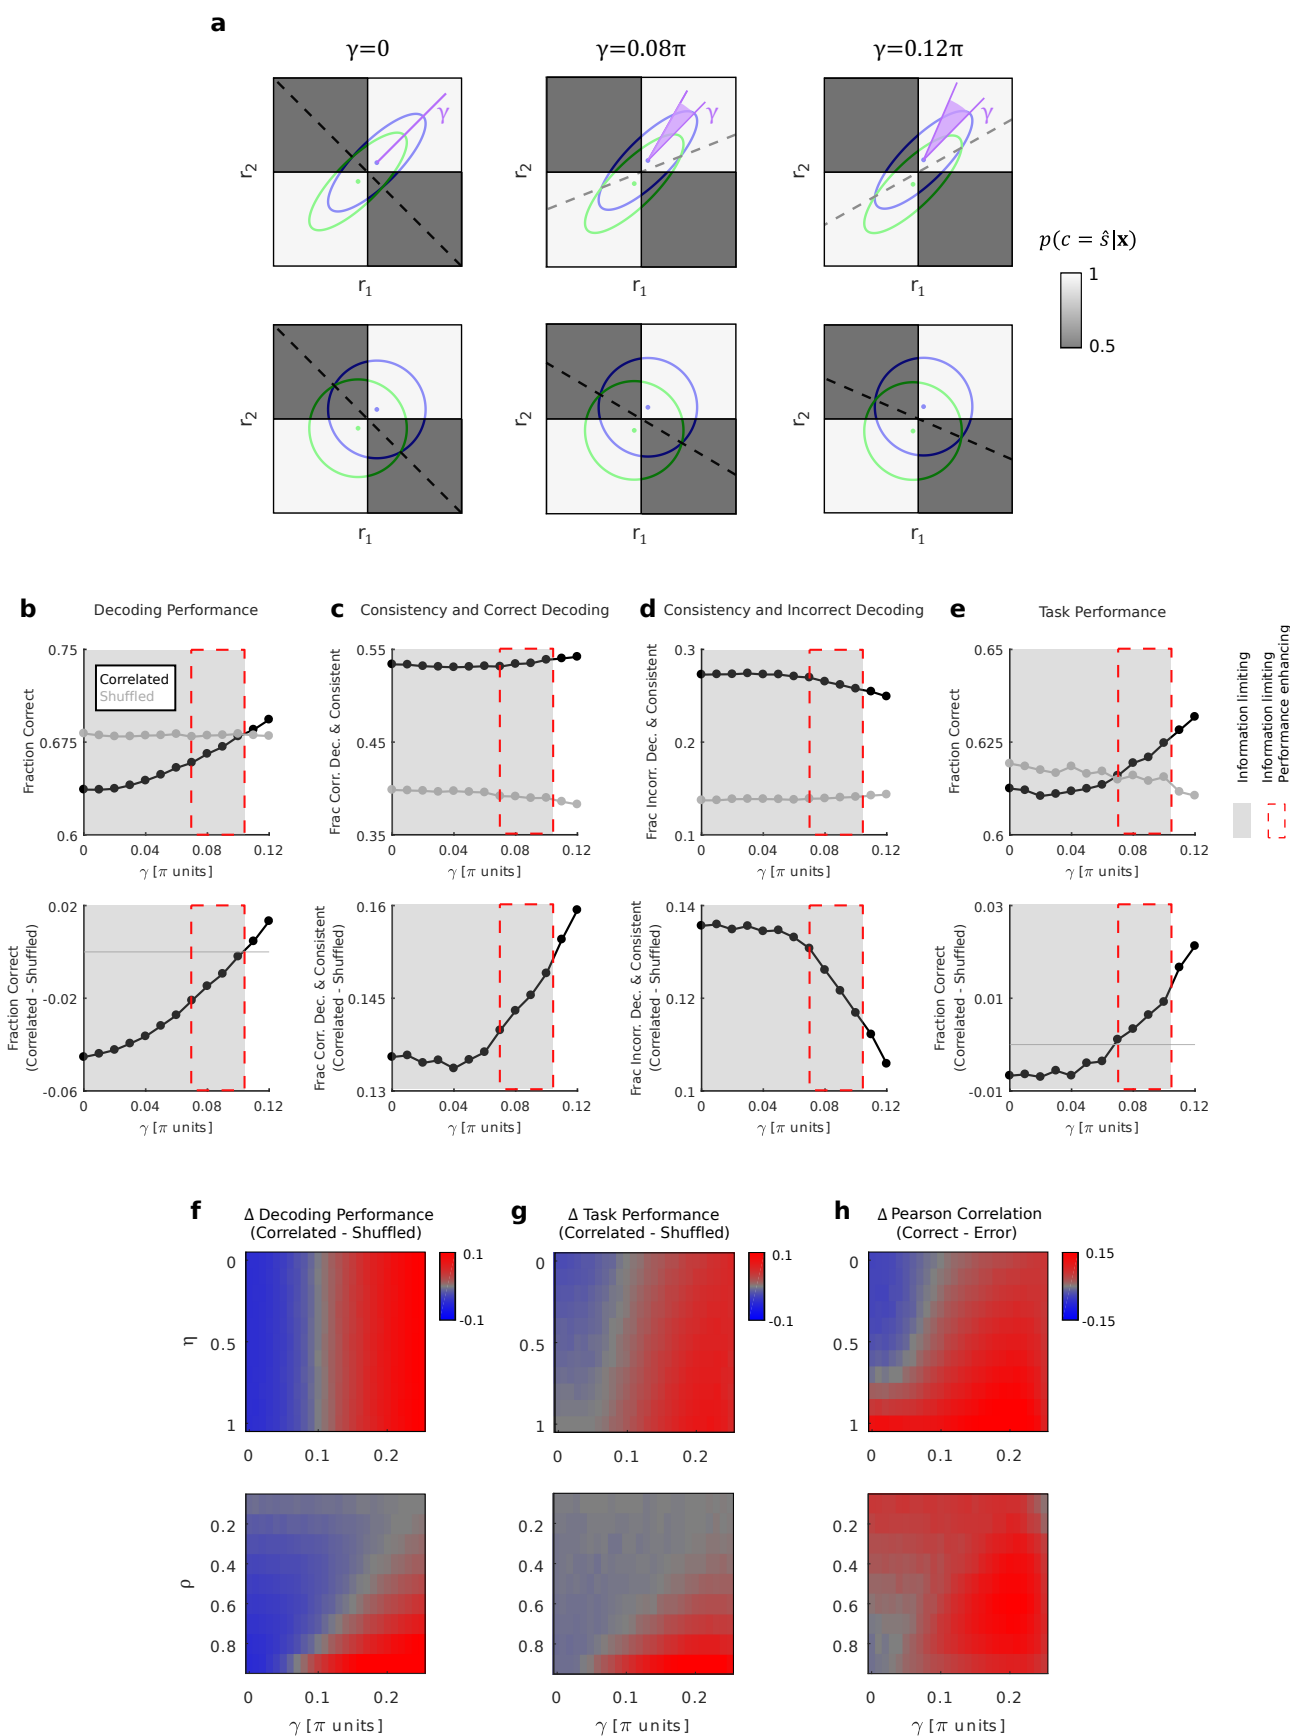
Extended Data Figure 2: Exploration of the parameter space of the two-feature encoding readout model, to investigate the conditions in which an enhanced-by-consistency readout produces an advantage for task performance for data characterized by information-limiting correlations. a, Example Gaussian

21

response distributions along two neural features (r_1,r_2) to two stimuli (s=1: blue, s=-1: green) at different degrees of alignment between signal and noise (γ) for correlated features (top, $\rho = 0.8$), and the corresponding shuffled distributions (bottom, $\rho = 0$). Ellipses indicate 95% confidence intervals. The dashed black line indicates the optimal decoding boundary. Note how, when increasing γ, the correlated response distributions to the two stimuli become less overlapped. **b**, Top: Performance of a linear decoder of stimulus identity applied to correlated (black) and uncorrelated (grey) responses for a range of γ values. The shaded grey region indicates values of γ for which correlations are information limiting. **c**, Top: The fraction of trials being correctly decoded and carrying consistent stimulus information increases at increasing γ for correlated but not for uncorrelated responses. **d**, Top: The fraction of trials being incorrectly decoded and carrying consistent stimulus information decreases at increasing γ for correlated but not for uncorrelated responses. **e**, Top: Task performance predicted by applying the enhanced-by-consistency readout model in Fig. 2i to correlated and uncorrelated responses for a range of γ values. The dashed red rectangle indicates values of γ for which information limiting correlations are advantageous for task performance. **b-e**, Bottom: Difference between correlated and uncorrelated responses, for the quantities shown in the top panels. **f**, Difference between stimulus decoding performance from correlated and uncorrelated responses for different combinations of model parameters values (γ: alignment between signal and noise. η: strength of consistency modulation in the readout. ρ: noise correlations strength). **g**, Difference between task performance predicted by applying the enhanced-by-consistency readout model in Fig. 2i to correlated and uncorrelated responses for different combinations of model parameters values. Note that there exist a range of model parameters for which correlations at the same time decrease decoding performance and improve task performance. **h**, Difference in Pearson correlation between correct choice and incorrect choice trials from correlated responses for different combinations of model parameters values. Where not otherwise specified, $\rho = 0.8$ and $\eta = 0.7$ in **b-h**.

Extended Data Figure 3: Exploration of the parameter space of the two-feature encoding readout model, to investigate the conditions in which an enhanced-by-consistency readout produces an advantage for task performance over a consistency-independent readout with matched readout efficacy. a, Task performance predicted by applying the enhanced-by-consistency readout in Fig. 2i (black) or a consistency-independent readout with matched readout efficacy (green, see Methods) to simulated correlated responses (as in Extended Data Fig. 2a, top), for different degrees of alignment between signal and noise (γ). The shaded grey region indicates values of γ for which correlations are information limiting. **b,** Difference between task performance predicted by applying the enhanced-by-consistency readout and the one predicted by applying the consistency-independent readout to correlated responses for different combinations of model parameters values (γ: alignment between signal and noise. η: strength of consistency modulation in the readout. ρ: noise correlations strength). Where not otherwise specified, $\rho = 0.8$ and $\eta = 0.7$ in **a-b**.

23
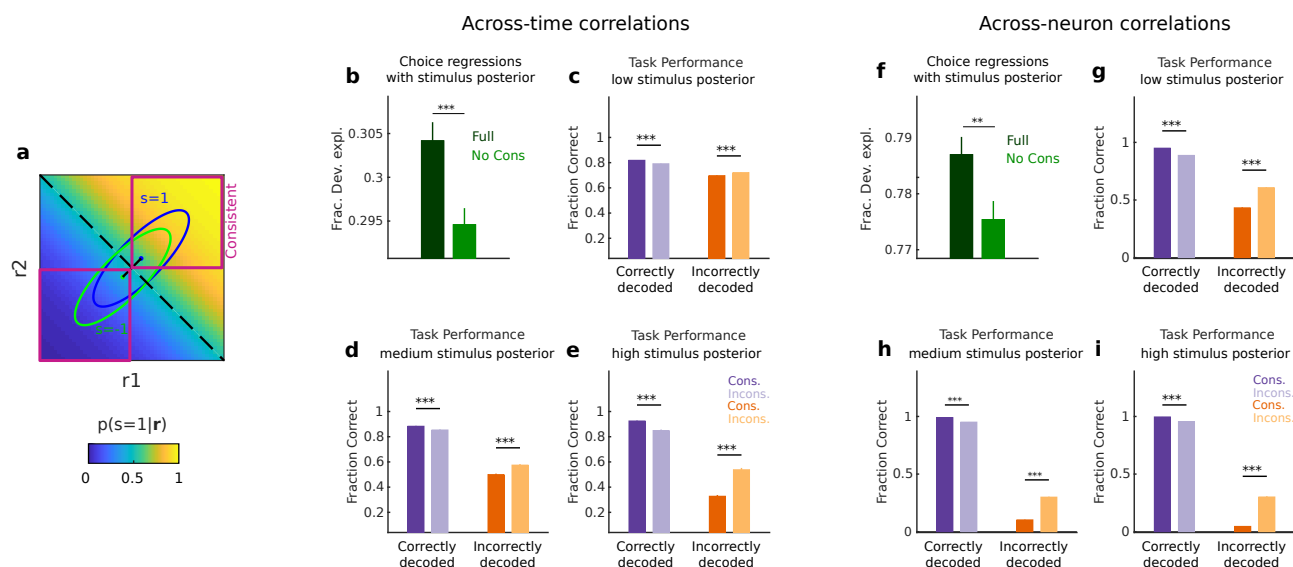
Extended Data Figure 4: The effect of neural consistency on the mouse's single trial choices cannot be explained by higher stimulus information associated to consistent neural representations. a, Schematic example showing response distributions along two neural features (r_1, r_2) to two stimuli ($s = -1$: green, $s = 1$: blue), which are best separated by the black dashed line, which represents the decoding boundary of a linear decoder trained on the observed responses. The background color encodes 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, the more "informative" $\mathbf{r}$ is about the stimulus. Note that, in the depicted example, 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. **b,** The readout underlying the choices of the animal is best explained by a choice regression that includes, in addition to the posterior probability of the stimulus extracted from joint PPC population vectors, its consistency across-time (Full: all predictors values, including stimulus posterior probability and posterior probability consistency, extracted from recorded data; No Cons: posterior probability consistency values shuffled across trials). ***P<0.001, one-sided permutation test. **c-e,** In order to control for possible confounding effects of the posterior probability in determining the behavioral relevance of across-time consistency, the computation in Fig 3b-c was repeated considering only subsets of trials having 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 ($|p(s = 1|\mathbf{r}) - 0.5| > 0.32$) posterior. Consistent population vectors result in average higher task performance when jointly encoding the presented stimulus correctly (left bars), and average lower task performance when jointly encoding the presented stimulus incorrectly (right bars), even when subsets of trials having approximately the same posterior are used. ***P<0.001, two-sided permutation test. **f-i,** as **b-e**, for across-neuron correlations. Data are plotted as mean ± s.e.m. across sessions and time point pairs with a lag difference of max 0.96 s in **b-e**. Data are plotted as mean ± s.e.m. across sessions, delay epochs, and 100 randomly-selected pairs of neuronal pools in **f-i**. For panels **b-i**, data are first averaged across trial splits.

24
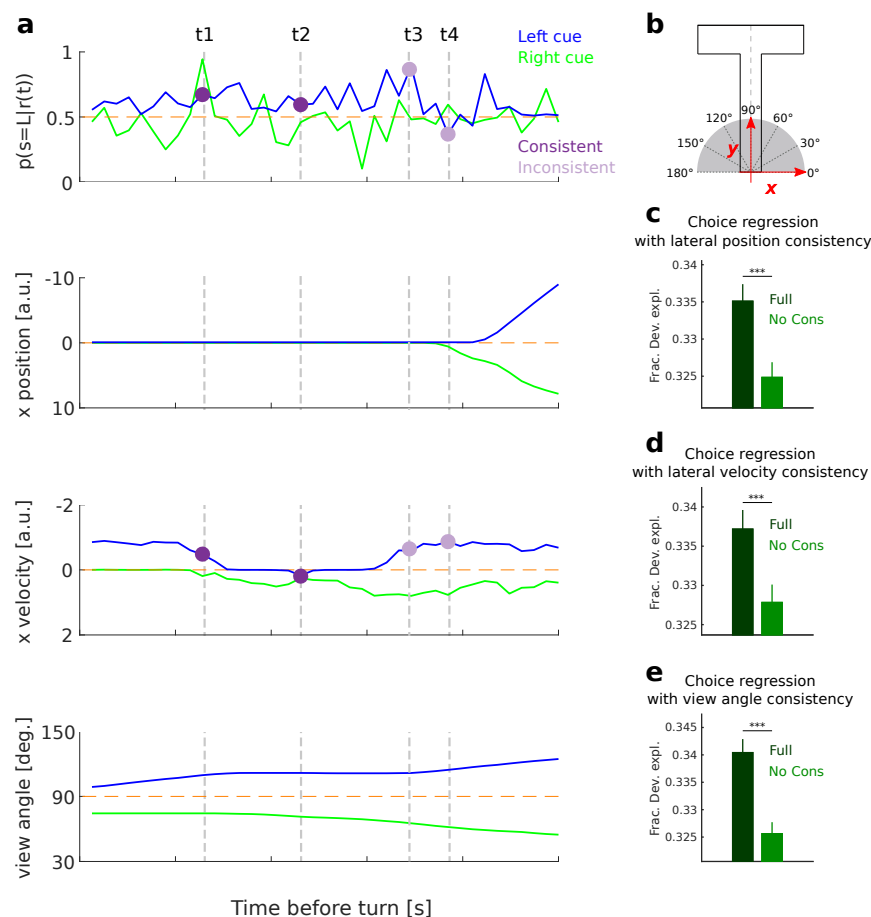
807
808
809 **Extended Data Figure 5: The role of neural consistency in the readout of PPC activity is not due to the**
810 **consistency of measured behavioral parameters. a,** The across-time evolution of the decoder posterior
811 probability of left stimulus presentation given the recorded PPC population activity (sound localization task
812 dataset) is shown along with the corresponding across-time evolution of a selection of three concurrently-
813 measured behavioral parameters (lateral position, lateral velocity, view angle), for an example left (blue) and
814 right (green) cue trial. Example time point pairs with consistent (t_1-t_2, dark purple) or inconsistent (t_3-t_4, light
815 purple) neural information are highlighted. Neural consistency is not necessarily associated to behavioral
816 consistency (e.g. when considering lateral running speed, t_1-t_2 are behaviorally inconsistent while t_3-t_4 are
817 behaviorally consistent). **b**, Schematic representation of the virtual T-maze with corresponding x-y coordinates
818 labelling and mouse's view angle (for a mouse oriented along the y axis). **c-e**, The readout underlying the
819 mouse's choices is best explained by a choice regression that includes, in addition to the consistency of
820 measured behavioral parameters (**c**: lateral position, **d**: lateral velocity, **e**: view angle), the consistency of the
821 stimulus identity extracted from neural activity (Full: all predictors values, including neural and behavioral
822 consistency, extracted from recorded data; No Cons: neural consistency values shuffled across trials). For
823 panels **c-e**: data are plotted as mean ± s.e.m. across sessions and time point pairs with a lag difference of max
824 0.96 s; data are first averaged across trial splits; ***P<0.001, one-sided permutation test.
825
826

25

Across-time correlations in PPC

Across-time correlations in AC
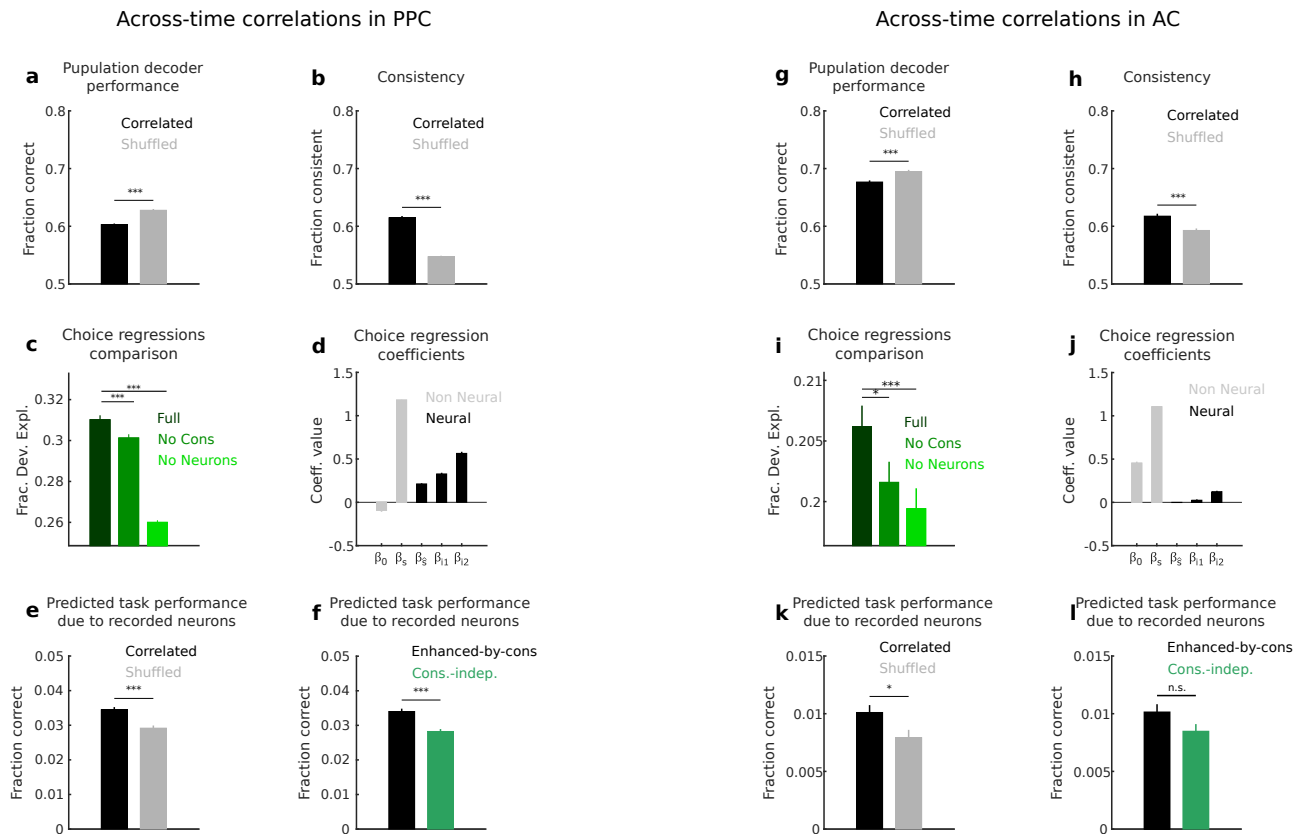
827
828
829 **Extended Data Figure 6: Across-time consistency of stimulus information in AC does not benefit task**
830 **performance as it does in PPC. a-f**, Summary of the main results of the analysis of across-time correlations
831 in PPC activity (sound localization task dataset), reported in Fig. 3, 4, useful for comparison with the
832 corresponding results of analysis of across-time correlations in AC activity (sound localization task dataset),
833 reported in **g-h. g**, Performance of a decoder of stimulus identity applied to joint population activity at two
834 different time points is higher in AC than in PPC. Across-time correlations limit the encoding of stimulus
835 information in AC just like in PPC. **h**, Across-time correlations in AC increase the consistency of encoded
836 stimulus information less than across-time correlations in PPC. ***$P<0.001$, two-sided permutation test. **i**,
837 Across-time consistency of stimulus information encoded in AC does not improve substantially the
838 performance of the logistic regression in predicting the animal's choices. ***$P<0.001$; *$P<0.05$, one-sided
839 permutation test. **j**, Best-fit beta coefficients of the full choice regression, showing that AC neural predictors
840 are assigned very low weights in the readout. **k**, Task performance that can be attributed to the recorded
841 neurons is much lower in AC than in PPC (~1% in AC, ~3.5% in PPC). Correlations in AC activity happen to
842 enhance task performance, but the effect is very small. **l**, Task performance which is attributable to the recorded
843 AC neural activity would not be substantially different if the behavioral readout was consistency-independent.
844 *$P<0.05$; n.s. (no statistical significance), two-sided permutation test. Data are mean ± s.e.m. across sessions
845 and time point pairs with a lag difference of max 0.96 s. Data are first averaged across trial splits.
846
847

26
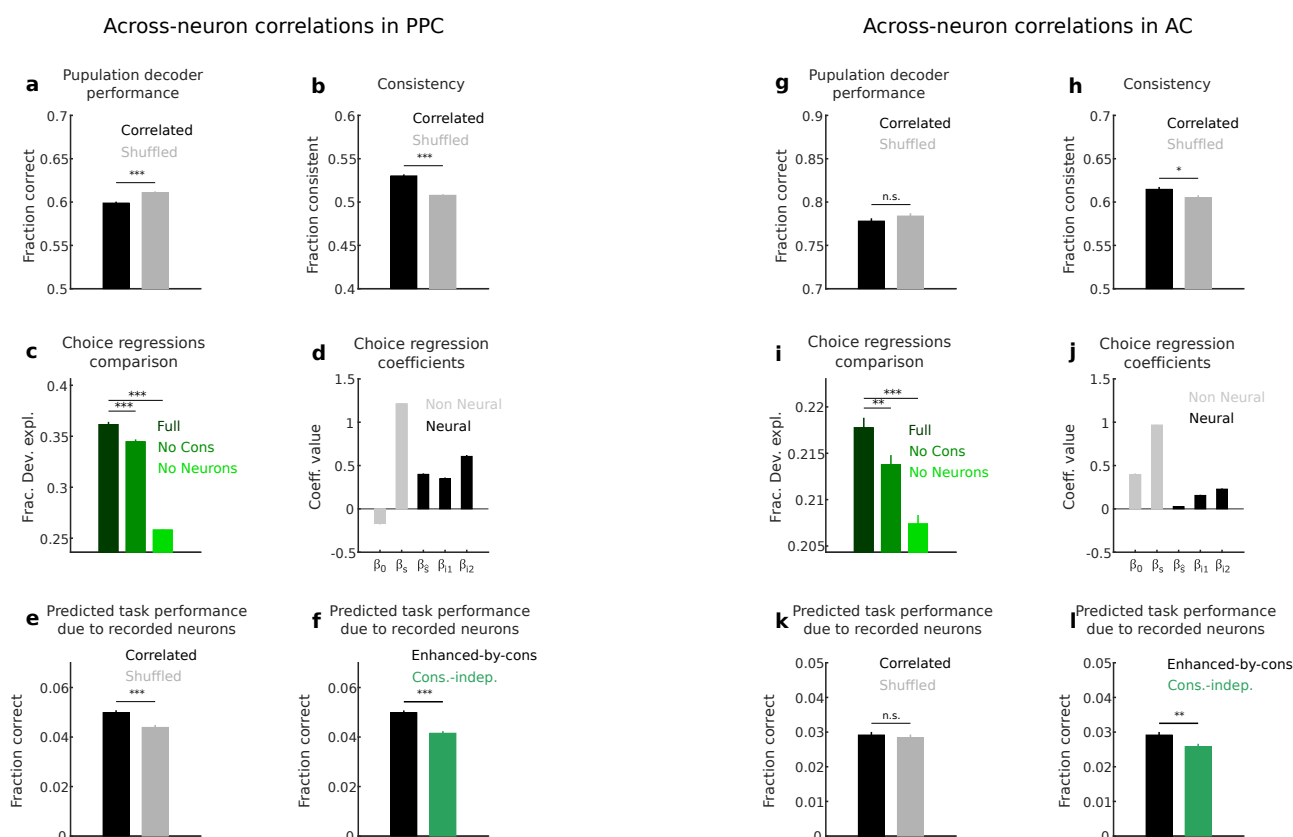
Extended Data Figure 7: In the sound localization dataset, across-neuron correlations in PPC are information limiting and provide an advantage for task performance. a, Across-neuron correlations in PPC limit the encoding of stimulus information during the sound localization task. **b**, Across-neuron correlations increase the consistency of encoded stimulus identity across neuronal subpopulations; however, the effect is less strong than that of across-time correlations (compare with Extended Data Fig. 6b). ***P<0.001, two-sided permutation test. **c**, The readout underlying the behavioral choices of the animal in the sound localization task is best explained by a choice regression that includes, in addition to the identity of the stimulus decoded from joint PPC population vectors, its consistency across neurons. ***P<0.001, one-sided permutation test. **d**, Best-fit beta coefficients of the full regression (which best explains the behavioral choices of the animal), showing that the across-neuron consistency plays a relevant role in the readout. **e**, Task performance in the sound localization task that can be attributed to the recorded neurons benefits from the presence of correlations across neurons in PPC. **f**, Task performance in the sound localization task that can be attributed to the recorded neurons would be lower if the readout of PPC activity was independent from consistency across-neurons. ***P<0.001, two-sided permutation test. **g-l,** Similarly to the case of across-time correlations (see Extended Data Fig. 6g-l), across-neuron correlations in AC are weak and so is their effect on stimulus information encoding (**g**, **h**) and predicted task performance (**k**). Across-neuron consistency seem to play a more important role in readout than across-time consistency in AC (**i**, **j**, **l**). Data are mean ± s.e.m. across sessions, time windows with a max duration of 0.96 s and 100 randomly-selected pairs of neuronal pools. Data are first averaged across trial splits.

27

870

Methods

872 No statistical methods were used to predetermine sample size. The experiments were not randomized.
873 The investigators were not blinded to allocation during experiments and outcome assessment.

874

Note on present study

876 This study represents an independent analysis of experiments described previously[17,19]. Although the
877 raw data are the same, the questions investigated and analytical tools are different, focusing here on
878 the analysis of the role of across-time and across-neuron correlations on information coding and
879 perceptual discrimination performance.

880

Subjects, behavioral task and two-photon imaging

882 A brief summary of the experimental procedures is provided here. A more detailed description can
883 be found in the original works[17,19]. All experimental procedures were approved by the Harvard
884 Medical School Institutional Animal Care and Use Committee.

885

886 Both experiments made use of a modified version of a visual virtual reality system that has been
887 described previously[44]. Head-restrained mice ran on a spherical treadmill, while images of a virtual
888 maze were projected on a half-cylindrical screen. Forward/backward translation in the maze was
889 controlled by treadmill changes in pitch, and rotation in the virtual environment was controlled by
890 roll of the treadmill. The virtual maze was constructed using the Virtual Reality Mouse Engine
891 (VirMEn) in MATLAB[45].

892

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

897

898 In the final version of the task that was used during imaging, mice ran down the stem of the virtual
899 T-maze, while sound stimuli were delivered from 8 possible locations (-90°, -60°, -30°, -15°, +15°,
900 +30°, +60°, +90°) using four electrostatic speakers positioned in a semicircular array, centered on the
901 mouse's head. The sound stimulus was activated when the mouse passed an invisible spatial threshold
902 at ~10 cm into the T-stem and originated from one of eight possible locations. The stimulus was
903 repeated after a 100 ms gap; repeats continued until the mouse reached the T-stem. Task difficulty
904 was modulated by the direction of the incoming stimulus, with the -90°/+90° trials being the easiest
905 ones and the -15°/+15° trials being the most difficult ones. In order to receive a reward (4 μl water),
906 mice had to judge the location of sound stimuli to be either on the left or right, and to report their
907 decisions by turning left or right at the T-intersection. A "reward tone" was played as the water reward
908 was delivered on correct trials (when the mouse had reached ~10 cm into the correct arm of the T-
909 maze), and a "no reward tone" was played when the mouse reached ~10 cm into the incorrect arm on
910 error trials. The inter-trial interval was 3 s on correct trials and 5 s on error trials. Mice performed
911 ~200 trials (range, 125–251) in a typical session over approximately 45–60 minutes.

912

913 Imaging was performed on alternating days from AC and PPC on the left hemisphere (PPC centered
914 at 2 mm posterior and 1.75 mm lateral to bregma; AC centered at 3.0 mm posterior and 4.3 mm lateral

28

915   to bregma). In a given session, ~50 neurons (range, 37-69) were simultaneously imaged using a two-

916   photon microscope (Sutter MOM) operating at 15.6 Hz frame rate and at 256 × 64 pixel resolution

917   (~ 250 μm × 100 μm). Imaging data were acquired at depths between 150 and 300 μm, corresponding

918   to layers 2/3. Seven AC fields of view and seven PPC fields of view from five mice were used for

919   analysis.

920

921   *Evidence accumulation task dataset*

922   Imaging data were acquired from five male C57BL/6J mice (The Jackson Laboratory), aged 8–10

923   weeks at the at the initiation of behavioral task training. Imaging began at least 4 weeks after viral

924   injection and was continued for up to 12 weeks.

925

926   In the final version of the task that was used during imaging, mice run down the stem of a virtual T-

927   maze with predominantly gray walls, encountering six visual cues (white wall segments with black

928   dots) at fixed locations. Each cue could appear on either the left or right wall, and only one cue was

929   visible at a time. To receive a reward (4 μl 10% sweetened condensed milk), mice had to determine

930   whether more cues were presented on the left or the right and, after a short stretch of maze without

931   additional cues (90 cm), turn at the T-intersection toward the direction that had more cues (left for 6-

932   0, 5-1, 4-2 trials; right for 2-4, 1-5, 0-6 trials). Task difficulty was modulated by varying the difference

933   between the number of left and right cues (net evidence). The sequence of cues was determined

934   randomly for each trial of a given net evidence. On trials having zero net evidence (3-3 trials), the

935   rewarded location was selected randomly. Following the completion of the trial, the screen changed

936   to black for the duration of the inter-trial interval (2 s for correct choice and 4 s for incorrect choice).

937   Mice performed ~300 trials (range, 231–414) in a typical session over approximately 45–60 minutes.

938

939   Imaging data were acquired from the left PPC (PPC centered at 2 mm posterior and 1.75 mm lateral

940   to bregma). In a given session, ~350 neurons (range, 188-648) were simultaneously imaged using a

941   custom-built two-photon microscope operating at ~30 Hz frame rate and at 512 × 512 pixel resolution

942   (~700 μm × 700 μm). Imaging data were acquired at depths between 100 and 200 μm below the dura.

943   Eleven fields of view from five mice were used for analysis.

944

945   **Imaging data processing**

946   Imaging frames were concatenated for motion correction, cell body identification, extraction of

947   fluorescence traces and estimation of relative spike rates. In brief, after motion correction[46],

948   correlations in fluorescence time series between pixels within ~60 μm were calculated. Fluorescence

949   sources (putative cells) were identified by applying a continuous-valued eigenvector-based

950   approximation of the normalized cuts objective[47] to the correlation matrix, followed by k-means

951   clustering segmentation, yielding binary masks for all identifiable fluorescence sources. To estimate

952   potential neuropil contamination, the cell body fluorescence signal was regressed against the signal

953   from surrounding pixels during the imaging frames when the cell of interest was not active, and then

954   neuropil contamination was removed during the ΔF/F calculation by subtracting a scaled version of

955   the neuropil signal from the cell body signal. All fluorescence traces were deconvolved to estimate

956   the relative spike rate in each imaging frame, which was expressed in arbitrary units given that the

957   estimate is relative to baseline activity without perfect knowledge of the fluorescence change

958   associated with a single spike[48].

959

Data inclusion and task epoch selection for encoding and readout analyses

Sound localization task dataset

Population activity data from the sound localization task dataset were used for the analysis of across-time correlations in PPC reported in the main text and for additional analyses of across-time correlations in AC and across-neuron correlations in PPC and AC reported in Supplementary Information.

For the analysis of across-time correlations in 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 work[17] 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 following analyses (see "Analysis of stimulus encoding and consistency").

For the analysis of across-time correlations in 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 post-stimulus 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 work[17] was excluded due to the large unbalance of left/right stimuli that were presented to the mouse across trials in that session.

For the analysis of across-neuron correlations in PPC, population activity data were temporally aligned to the imaging time frame of the turn and averaged over ten turn-aligned time windows of increasing length, comprising the {4:8: ... :36:39} frames preceding the turn imaging frame. The same six sessions used for the analysis of across-time correlations in PPC were used.

Evidence accumulation task dataset

Population activity data from the evidence accumulation task dataset were used for the analysis of across-neuron correlations in PPC reported in the main text. PPC population activity data were first grouped into spatial bins (3.75 cm/bin) covering the whole length of the T-maze (long and short arm) by averaging population activity in each bin (2 or 3 imaging frames per bin per trial). Population activity data were then further averaged over epochs of 4 spatial bins each (about 200 ms). We considered the same 10 epochs that were defined in the original work[19]. We plotted results corresponding to the 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 (more precisely, these epochs correspond to the four spatial bins beginning respectively 15 and 37.5 cm after the offset of the final cue). Therefore, it is reasonable to assume that these epochs are involved in the formation of the animal's decision. All 11 sessions in the original work[19] were used. We discarded 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. Stimulus category was defined as a binary variable (left/right). 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°, -15°; right: +15°, +30°, +60°, +90°). 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, 4-2 trials; right: 2-4, 1-5, 0-6 trials).

The selectivity index[15] (SI) was quantified as:

$$SI = \frac{mean\ \Delta F/F_{right\ trials} - mean\ \Delta F/F_{left\ trials}}{mean\ \Delta F/F_{right\ trials} + mean\ \Delta F/F_{left\ trials}} \ . \tag{1}$$

Cells that had a selectivity index greater than zero were considered right-preferring cells and those that had a selectivity index less than zero were considered left-preferring cells.

Pairwise noise correlations

We quantified across-time and across-neuron pairwise noise correlations separately for trials in which mice made correct or error choices. In order to discount possible differences in correlation related to trials numerosity, we subsampled trials to equalize the number of correct and error trials in each recorded session. The subsampling procedure was repeated 20 times, and results were averaged over the 20 subsamples.

Across-time correlations

We quantified across-time pairwise correlations by computing, for each pair of recorded neurons and for each pair of time points in the trial epoch selected for analyses (see "Data inclusion and task epoch selection for encoding and readout analyses"), 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 sharing the same stimulus category. Results were averaged first across all time point pairs sharing the same lag between each other, then across stimuli and finally across trials subsamples.

Across-neuron correlations

We quantified across-neuron pairwise correlations by computing the Pearson correlation for each pair of neurons recorded in a single session, across trials sharing the same stimulus category. Results were averaged first across stimuli and then across trials subsamples.

Analysis of stimulus encoding and consistency

For encoding and consistency analyses, we considered information about stimulus category (see "Selectivity of single cells to 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 classifier (C-SVC) with a linear basis function kernel[49,50]. The C-SVC was implemented with custom MATLAB software, by building on the libsvm library[51]. 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 10 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-SVC 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 3-fold 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-SVC[52].

Across-time consistency

For the across-time correlations analysis, we decoded the most likely stimulus from the population activity at each time point in the trial epoch selected for analyses (see "Data inclusion and task epoch selection for encoding and readout analyses") and from the concatenated population activity of every pair of time points. Any two time points were defined to be consistent if the stimuli decoded from the two time points individually coincided.

Across-neuron consistency

For the across-neuron correlations analysis, we first split the neuronal population recorded in each session into two randomly-selected, equally-sized, disjoint pools of neurons. We then decoded the most likely stimulus from the population activity of each pool individually and from the concatenated population activity of both pools (note that this corresponds to the activity of the whole recorded neuronal population). Two pools were defined to be consistent if the stimuli decoded from the two pools individually coincided. The random assignment of the neurons to the two pools was repeated 100 times.

Quantifying the angle between the signal and noise axes

We quantified, in the neural population response space, the angle γ between the direction of maximum stimulus variation (signal correlations axis) and the direction of maximum noise variation (noise correlations axis)[24,25]. The signal correlations axis was computed as the slope of the linear regression of the trial-averaged responses at fixed stimulus. The noise correlations axis was computed as the slope of the linear regression of the trial-by-trial neural responses at fixed stimulus, averaged across stimuli. γ takes values between 0 and $\pi/2$.

Across-time correlations

For across-time correlations, multidimensional population responses at a given time point were first reduced to scalars by computing the projection along the direction of maximum stimulus variation in the space spanned by the population responses at the selected time point. Then, for each pair of time points in the trial epoch selected for analyses (see "Data inclusion and task epoch selection for encoding and readout analyses"), γ was computed in the 2D space defined by the projections of the population responses at the two time points.

Across-neurons correlations

For across-neuron correlations, we considered the same 100 randomly-defined pairs of neuronal pools that were used for encoding and consistency analysis (see "Analysis of stimulus encoding and

32

1093   consistency"). Multidimensional population responses of a given neuronal pool were first reduced to
1094   scalars by computing the projection along the direction of maximum stimulus variation in the space
1095   spanned by the population responses of the selected neuronal pool. Then, for each randomly-defined
1096   pair of disjoint pools, γ was computed in the 2D space defined by the projections of the population
1097   responses of the two pools.

1098

1099   **Mathematical model of encoding and readout with two neural features**
1100   We developed a simple two-features mathematical model of encoding and readout in a simulated
1101   stimulus discrimination task, to explore how possibly distinct roles of correlations in information
1102   encoding and information readout interact and impact task performance.

1103

1104   *Neural encoding (stimulus-response) model*
1105   As far as the encoding model is concerned, we simulated the response of two generic features of
1106   neural activity (r_1, r_2) to two stimuli ($s = -1$, $s = 1$). Neural response distributions for a given stimulus
1107   are modelled as bivariate Gaussians $N(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_s)$, with mean and covariance matrix given by:

1108

$$\boldsymbol{\mu}_s = sign(s) * d * \begin{bmatrix} \cos{(\gamma + \frac{\pi}{4})} \\ \sin{(\gamma + \frac{\pi}{4})} \end{bmatrix}$$

$$(2)$$

$$\boldsymbol{\Sigma}_s = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix},$$

1109

1110   where σ_1^2 and σ_2^2 represent the variance of the neural responses along r_1 and r_2 respectively, while ρ
1111   determines the correlation between r_1 and r_2. For simplicity, in our simulations we chose $\sigma_1 = \sigma_2 =$
1112 σ and $\boldsymbol{\Sigma}_{s1} = \boldsymbol{\Sigma}_{s2}$. We selected the means of the two response distributions ($\boldsymbol{\mu}_{s1}$, $\boldsymbol{\mu}_{s2}$) to be
1113   symmetrically located around the origin in the 2D space defined by r_1 and r_2, at a distance d. Taken
1114   together, d and σ control the overlap between the two response distributions (in all our simulations
1115   we arbitrarily set $d = \sqrt{0.02}$ and $\sigma = 0.3$). The parameter γ controls the alignment between the signal
1116   correlations axis and the noise correlations axis, which are defined as the axis of maximum variability
1117   of trial-averaged responses at fixed stimulus and the axis of maximum trial-by-trial variability at fixed
1118   stimulus, averaged across stimuli, respectively. In our simulations, we varied γ by varying the signal
1119   correlations axis orientation (see Extended Data Fig. 2 and Supplementary Information): for $\gamma = 0$,
1120 $\boldsymbol{\mu}_{s1}$ and $\boldsymbol{\mu}_{s2}$ are located on the bisector of the first and third quadrants, resulting in perfect alignment
1121   between signal and noise; at increasing values of γ, the signal correlations axis direction is
1122   progressively rotated counterclockwise, resulting in more and more misalignment with the noise
1123   correlations axis; the maximum misalignment is reached when $\gamma = \frac{\pi}{4}$ and signal and noise directions
1124   are orthogonal. We quantified the amount of stimulus information carried by the simulated neural
1125   responses as the fraction correct predictions of a linear decoder of stimulus identity applied to the
1126   simulated (r_1, r_2) responses. We further applied the same linear decoder to the responses along r_1 and
1127 r_2 separately for the computation of consistency. For the simulations in Fig. 2, Extended Data Fig. 2
1128   and Extended Data Fig. 3 we performed, for a given choice of model parameters, 1000 simulations
1129   with 500 trials per stimulus.

33

130

131 *Behavioral decision model*

132 As far as the readout model is concerned, we simulated the process of generating a binary choice in

133 each trial from neural activity through a logistic regression model:

134

$$logit(p(c = 1|\,\mathbf{x})) = \beta_0 + \beta_{\hat{s}}\hat{s} + \frac{\beta_{i1}}{2}(\hat{s} + 1)con + \frac{\beta_{i2}}{2}(\hat{s} - 1)con \tag{3}$$

135

136 where:

137 • $\hat{s}$ is a binary variable ($\hat{s} = -1,\ \hat{s} = 1$) which encodes the stimulus decoded from the

138 concatenated activity of the two neural features;

139 • *con* is a binary variable that is 1 if the stimuli decoded individually from each neural feature

140 are the same, and 0 otherwise;

141 • $\mathbf{x}$ indicates the entire set of predictors ($\hat{s}, con$).

142

143 The model coefficients $\beta_0, \beta_{\hat{s}}, \beta_{i1}$ and β_{i2} control the relative impact of the different model predictors

144 on the simulated choice. The values for the model coefficients were set as follows. We first defined

145 a consistency modulation index η, ranging from 0 to 1, to control the relative strength of neural

146 consistency in the readout. We then derived the readout efficacy, which we defined as the probability

147 of conversion from $\hat{s}$ to c, for each of the four possible combinations of predictors values, from the

148 modulation index η and a reference readout efficacy $\alpha(\hat{s})$ as:

149

$$p(c = \hat{s}|[\hat{s}, con]) = \begin{cases} \alpha(\hat{s}) + \eta(1 - \alpha(\hat{s})), & con = 1 \\ \alpha(\hat{s}) - \eta(\alpha(\hat{s}) - 0.5), & con = 0 \end{cases} \tag{4}$$

150

151 where $\alpha(\hat{s})$ takes values between .5 and 1. For the simulations in Fig. 2, Extended Data Fig. 2 and

152 Extended Data Fig. 3, we arbitrarily set $\alpha(0) = \alpha(1) = 0.75$. Given the readout efficacy values from

153 equation (4), we used equation (3) to compute the model coefficients corresponding to the chosen

154 modulation index η.

155

156 **Logistic regression of the mouse's choice**

157 To study the relevance of selected features of recorded neural population activity on the mouse's

158 choices, we fit to the recorded neural activity a logistic regression of the choice made by the animal

159 in each trial (left/right turns). For each trial, we considered the choice c made by the mouse ($c =$

160 left/right), the presented stimulus s ($s =$ left/right), and the neural population activity for each pair of

161 time points (across-time correlations analysis) or for each pair of randomly-selected neuronal pools

162 (across-neuron correlations analysis) that were used for encoding and consistency analyses (see

163 "Analysis of stimulus encoding and consistency"). For each session, trial split, and pair of time points

164 or neuronal pools, we fitted the following logistic regression:

165

$$logit(p(c = left|\,\mathbf{x})) = \beta_0 + \beta_s s + \beta_{\hat{s}}\hat{s} + \frac{\beta_{i1}}{2}(\hat{s} + 1)con + \frac{\beta_{i2}}{2}(\hat{s} - 1)con \tag{5}$$

166

167 where

34

- $\hat{s}$ is a binary variable ($\hat{s}$ = left/right) that encodes the most likely stimulus decoded from the concatenated activity of two time points or neuronal pools;
- con is a binary variable that is 1 if the stimuli decoded individually from each time point or neuronal pool are the same, and 0 otherwise;
- $\mathbf{x}$ indicates the entire set of neural ($\hat{s}, con$) and non-neural (s) predictors.

Logistic regression fitting was implemented with custom Python software, by building on the *statsmodel* module[53]. Note that we used the convention (left = 1, right = -1) to map the categorical variables $s, \hat{s}, c$ to numerical variables suitable for logistic fitting. For each trial split defined for encoding and consistency analyses (see "Analysis of stimulus encoding and consistency"), the logistic regression was fit on trials belonging to the testing set, for which there exists a prediction of the most likely stimulus from population activity for each individual and pair of time points or neuronal pools. The logistic regression was fit using L1-regularized maximum likelihood, for a conservative selection of the predictors of the mouse's choice. The regularization hyperparameter (λ) was selected by maximizing 3-fold cross-validated fraction of deviance explained (FDE, see "Predictive performance and cross validation"). The few instances where the maximum-likelihood optimization did not converge numerically were not considered for further analyses.

For control analyses, we fitted mouse's choices with more complex choice regressions that included other predictors on top of those described in equation (5).

In order 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 which included additional behavioral consistency-dependent predictors:

$$logit\big(p(c = left|\,\mathbf{x})\big) \qquad\qquad (6)$$
$$= \beta_0 + \beta_s s + \beta_{\hat{s}}\hat{s} + \frac{\beta_{i1}}{2}(\hat{s}+1)con + \frac{\beta_{i2}}{2}(\hat{s}-1)con + \frac{\beta_{i3}}{2}(\hat{s}+1)con_b$$
$$+ \frac{\beta_{i4}}{2}(\hat{s}-1)con_b$$

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 (see Extended Data Fig. 5 and Supplementary Information). Two values of lateral running velocity or lateral position taken 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°.

In order to resolve the genuine role of across-time or across-neurons consistency in explaining the mouse's choices, above and beyond the overall level of stimulus information encoded in neural activity, we fitted mouse's choices with a logistic regression where the discrete binary variable $\hat{s}$ in equation (5) was replaced with the corresponding continuous value of the decoder stimulus posterior probability $p(s = left|\boldsymbol{r})$ (see "Analysis of stimulus encoding and consistency"). This choice

1206   regression, differently from the one of equation (5), also accounts for the confidence associated to a
1207   given stimulus decoded from the concatenated activity of two time points or two neuronal pools.
1208

1209   *Predictive performance and cross validation*
1210   To test for the significance of the contribution of a predictor (or set of predictors) in explaining the
1211   mouse's choices, we compared the predictive performance of the regression fitted (and tested) on
1212   predictors values extracted from the recorded data with the predictive performance of a regression
1213   fitted (and tested) on artificial data in which the values of the predictor (or set of predictors) of interest
1214   were shuffled across trials. Predictive performance was quantified as the cross-validated fraction of
1215   deviance explained (FDE). FDE was evaluated with 3-fold cross validation, thus addressing concerns
1216   of overfitting. For each fold, we computed the log-likelihood l of the test data given the values of the
1217 β coefficients of the regression fit on the training data. To calculate a reference null value for the log-
1218   likelihood, we computed the log-likelihood l_0 of the test data given the value of the coefficient β_0 of
1219   an intercept-only regression fit on the training data. The FDE was then defined as:
1220

$$FDE = 1 - l/l_0 \tag{7}$$

1221

1222   **Estimating the impact of across-time and across-neuron correlations on task performance**
1223   To estimate the impact of across-time and across-neuron correlations on mouse's task performance,
1224   we generated synthetic choices using the experimentally-fit choice regression of equation (5). As
1225   input to the regression, we provided either predictors extracted from the real recorded neural data,
1226   which included across-time and across-neuron correlations, or predictors extracted from hypothetical
1227   neural data where correlations of interest were artificially removed by shuffling (see "Removing
1228   across-time or across-neuron noise correlations by shuffling"). We used this analytic approach
1229   because current experimental methods cannot remove correlations online during task performance,
1230   and thus we had to estimate effects with post-hoc removal of correlations.
1231

1232   Technically, given the discrete nature of the predictors used in the choice regression of equation (5),
1233   we did not need to generate synthetic choices for the computation of the task performance $p(c = s)$,
1234   which is the probability that the choice c matches the presented stimulus s. Rather, we simply
1235   computed the probabilities $p(\mathbf{x})$ for all possible combination of predictors values X, we multiplied
1236   them with the corresponding readout probabilities $p(c = s|\mathbf{x})$ obtained from the logistic choice
1237   regression, and then summed over all possible combinations of predictors values X as follows (see
1238   Extended Data Fig. 1 for a schematic representation of this operation):
1239

$$p(c = s) = \sum_{\mathbf{x} \in X} p(c = s|\mathbf{x}) * p(\mathbf{x}). \tag{8}$$

1240

1241   Note that the same readout probabilities $p(c = s|\mathbf{x})$ were used for the computation of task
1242   performance from both real and shuffled neural data.
1243

1244   We further isolated the part of task performance that can be attributed to the (real or shuffled) neural
1245   activity, by subtracting from the total estimated task performance a baseline non-neural estimated
1246   task performance. The non-neural task performance was computed by generating choices (or applying

36

1247 equation (8)) after shuffling the values of neural predictors (i.e. predictors defined from features of
1248 neural activity) across trials, while keeping the relationship between non-neural predictors (i.e.
1249 predictors that did not depend on neural activity) and mouse's choices fixed. Note that the shuffling
1250 procedure allowed to discard uninteresting differences between task performance from real neural
1251 data and that from shuffled neural data that are due to intrinsic differences between the marginal
1252 distributions of the neural predictors. The estimated value of task performance due to the recorded
1253 neurons is by definition non-negative and it is identically zero when the coefficients associated to
1254 neural predictors are zero, as we intuitively require to a measure of the behavioral impact of the neural
1255 predictors.

1256

1257 **Computation of readout efficacy of the transformation from stimulus information to choice**
1258 We termed "readout efficacy" $p(c = \hat{s})$ the probability of conversion from stimulus encoded in
1259 neural activity to choice associated to a given readout. Given the discrete nature of the predictors used
1260 in the choice regression of equation (5), similarly to what we did for the estimation of task
1261 performance (see "Estimating the impact of across-time and across-neuron correlations on task
1262 performance"), we computed the efficacy of the enhanced-by-consistency readout as:

1263

$$p(c = \hat{s}) = \sum_{\mathbf{x} \in \mathbf{X}} p(c = \hat{s}|\mathbf{x}) * p(\mathbf{x}), \tag{9}$$

1264

1265 where X represents the set of all possible combinations of predictors values and $p(c = \hat{s}|\mathbf{x})$ are
1266 obtained from the logistic choice regression.

1267

1268 To generate the readout maps in Fig. 3c, f we computed, separately for consistent and inconsistent
1269 trials, readout efficacy as a deviation from a stimulus-driven baseline, which is the average probability
1270 of choice being left or right when the presented stimulus is left or right, as:

1271

$$\Delta\mathrm{p}(c = \hat{s}|con) = \sum_{\mathbf{x} \in X} \Delta p(c = \hat{s}|[\mathbf{x}, con]) * p(\mathbf{x}|con). \tag{10}$$

1272

1273 In equation (10), X represents the set of all possible combinations of $[s, \hat{s}]$ predictors values and
1274 $\Delta p(c = \hat{s}|[\mathbf{x}, con]) = p(c = \hat{s}|[s, \hat{s}, con]) - p(c = \hat{s}|s)$.

1275

1276 **Matching enhanced-by-consistency and consistency-independent readouts in terms of efficacy**
1277 To quantify the impact on task performance of the enhanced-by-consistency experimentally-
1278 measured readout, we compared the task performance predicted by the experimentally-fit choice
1279 regression to the one predicted by a consistency-independent choice regression of the form:

$$logit\big(p(c = left|\mathbf{x})\big) = \beta_0' + \beta_s's + \beta_{\hat{s}}'\hat{s}. \tag{11}$$

1280

1281 For a fair comparison, the values of the coefficients β_0', β_s' and $\beta_{\hat{s}}'$ were chosen so that the two
1282 readouts were matched in terms of readout efficacy (see "Computation of readout efficacy in the
1283 transformation from stimulus information to choice"). Note how the readout efficacy $p(c = \hat{s})$ in

37

1284 equation (9) depends on both the readout probabilities $p(c = \hat{s}|\mathbf{x})$ and the predictors probabilities
1285 $p(\mathbf{x})$. Thus, the condition of matched readout efficacy is specific for a given input to the regressions.
1286
1287 Concretely, in order to compute the values of the coefficient of the consistency-independent choice
1288 regression in (11), we imposed the following conditions:
1289

$$
\begin{cases}
p_{cons.-indep.}(c = 0, \hat{s} = 0) = p_{enhanced-by-cons.}(c = 0, \hat{s} = 0) & (12a) \\
p_{cons.-indep.}(c = 1, \hat{s} = 1) = p_{enhanced-by-cons.}(c = 1, \hat{s} = 1) & (12b) \\
\beta'_s = \beta_s & (12c)
\end{cases}
$$

1290
1291 then appropriately plugged equations (5) and (11) into (12a) and (12b), and solved for β'_0 and $\beta'_{\hat{s}}$.
1292

1293 **Removing across-time or across-neuron noise correlations by shuffling**
1294 To study the role of across-time and across-neuron correlations on stimulus encoding and task
1295 performance, we generated artificial neural population vectors in which we removed correlations by
1296 a shuffling procedure.
1297

1298 *Across-time correlations analysis*
1299 Across-time correlations between population vectors at different time points were removed by
1300 shuffling trial identities independently for each of the two population vectors within subsets of trials
1301 of the same stimulus category. This procedure ensured that "across-time signal correlations" were
1302 maintained, while "across-time noise correlations" were disrupted. Note that single-cell
1303 autocorrelations were also disrupted with this shuffling procedure.
1304

1305 *Across-neuron correlations analysis*
1306 Across-neuron correlations between two neuronal pools were removed by shuffling trial identities
1307 independently for each pool within subsets of trials of the same stimulus category. This procedure
1308 ensured that "across-neuron signal correlations" were maintained, while "across-neuron noise
1309 correlations" were disrupted. Note that classical signal correlations between any pairs of neurons
1310 from the two disjoint pools were thus maintained, while classical noise correlations between any pairs
1311 of neurons from the two disjoint pools were disrupted.
1312

1313 Shuffling was performed separately within trials of the training and the testing set.
1314

1315 **Data and code availability**
1316 The data and code that support the findings of this study are available from the corresponding authors
1317 upon reasonable request.
1318

1319 **Methods references**
1320

1321 44.   Harvey, C. D., Collman, F., Dombeck, D. A. & Tank, D. W. Intracellular dynamics of
1322   hippocampal place cells during virtual navigation. Nature 461, 941–946 (2009).
1323 45.   Aronov, D. & Tank, D. W. Engagement of Neural Circuits Underlying 2D Spatial Navigation
1324   in a Rodent Virtual Reality System. Neuron 84(2), 442-456 (2014).
1325   doi:10.1016/j.neuron.2014.08.042

46. Greenberg, D. S. & Kerr, J. N. D. Automated correction of fast motion artifacts for two-photon imaging of awake animals. J. Neurosci. Methods (2009). doi:10.1016/j.jneumeth.2008.08.020

47. Shi, J. & Malik, J. Normalized cuts and image segmentation. IEEE Trans. Pattern Anal. Mach. Intell. (2000). doi:10.1109/34.868688

48. Vogelstein, J. T. et al. Fast nonnegative deconvolution for spike train inference from population calcium imaging. J. Neurophysiol. 104, 3691–3704 (2010).

49. Boser, B. E., Guyon, I. M. & Vapnik, V. N. A training algorithm for optimal margin classifiers. in Proceedings of the fifth annual workshop on Computational learning theory - COLT '92 144–152 (ACM Press, 1992). doi:10.1145/130385.130401

50. Cortes, C. & Vapnik, V. Support-vector networks. Mach. Learn. 20, 273–297 (1995).

51. Chang, C.-C. & Lin, C.-J. LIBSVM: A Library for Support Vector Machines. (2001).

52. Lin, H.-T., Lin, C.-J. & Weng, R. C. A note on Platt's probabilistic outputs for support vector machines. Mach. Learn. 68, 267–276 (2007).

53. Seabold, S. & Perktold, J. Statsmodels: Econometric and Statistical Modeling with Python. PROC. OF THE 9th PYTHON IN SCIENCE CONF (2010).