

# 1 Targeted V1 comodulation supports task-adaptive sensory decisions

2 Caroline Haimerl<sup>5</sup> \*, Douglas A. Ruff<sup>2</sup>, Marlene R. Cohen<sup>2</sup>, Cristina Savin<sup>1,3</sup> †, and Eero P.  
3 Simoncelli<sup>1,3,4</sup> †

4 <sup>1</sup> Center for Neural Science, New York University, New York, NY 10003, USA

5 <sup>2</sup> Department of Neurobiology, University of Chicago, Chicago, IL 60637, United States

6 <sup>3</sup> Center for Data Science, New York University, New York, NY 10011, USA

7 <sup>4</sup> Flatiron Institute, Simons Foundation, New York, NY 10010, USA

8 <sup>5</sup> Champalimaud Centre for the Unknown, Lisboa, Portugal

9 \* Corresponding author

10 † Co-senior authors

## 11 Abstract

12 Sensory-guided behavior requires reliable encoding of stimulus information in neural populations, and task-specific  
13 readout through a selective combination of these responses. The former has been the topic of extensive study,  
14 but the means by which the brain achieves the latter remain poorly understood, especially when adapting to  
15 changing task demands. Here we introduce a novel theory for adaptive sensory processing based on functionally-  
16 targeted stochastic modulation. We find that responses of neurons in area V1 of monkeys performing a visual  
17 orientation discrimination task exhibit low-dimensional, rapidly fluctuating gain modulation, which is stronger  
18 in neurons that are most informative for the current behavioral task. We propose that this modulation serves  
19 as a label that supports adaptive downstream readout. Our theoretical and experimental results show that V1  
20 modulation can be used to decode from neural activity after only a small number of training trials, consistent  
21 with observed behavior. In a hierarchical visual neural network model, the modulator-induced labels are learned  
22 quickly and accompany task information across several processing stages to finally guide decisions. Consistent  
23 with this model, we find that the V1 modulatory signal is also present in the activity of simultaneously recorded  
24 MT units, and that its label of task information is preserved. Our findings provide evidence for a novel mechanism  
25 for task-adaptive information routing in the brain, through targeted co-modulation.

## 26 1 Introduction

27 Humans and animals are able to flexibly adapt their behavior according to ever-changing sensory input and goals.  
28 In the brain, sensory information is transformed through hierarchical stages of computation, building increasingly  
29 complex feature maps. Yet, decisions can rely on local stimulus attribute, which requires not just preserving this  
30 information throughout the processing hierarchy, but also selectively choosing which aspects of the representation to  
31 read out [1]. Consider a decision about local visual orientation. This information is explicitly represented in primary  
32 visual cortex (V1), where neurons respond selectively to specific orientations at specific locations in the visual field.  
33 However, decisions are not made in V1 – visual orientation signals undergo a sequence of transformations, presumably  
34 mixing with task irrelevant information (other features of the stimulus or information from other spatial locations)  
35 before reaching decision areas. How do areas downstream of V1 access task-relevant sensory information to flexibly  
36 guide behavior?

37 The problem of flexible sensory decision making has been studied from different perspectives. First, within the  
38 traditional “ideal observer” framework, statistically optimal decoders can be constructed from a complete description  
39 of response properties the encoding population, as they pertain to the task. These provide performance upper bounds  
40 for behavior [2–8], but fail to explain how a downstream circuit –with limited knowledge of each upstream neuron’s  
41 stimulus-response and noise properties– can construct such a readout [9]. Second, attentional boosts in the activity  
42 of the relevant neurons are believed to highlight task-informative sensory information for downstream processing  
43 [10–12]. However, this early-stage encoding selection may be insufficient to ensure the preferential transmission of  
44 task-specific information across a complex processing hierarchy [13]. Some have argued that the behavioral benefits  
45 of attention are largely due to effective contextual readouts [14], which may explain instances where behavioral-level  
46 benefits due to attention can be experimentally dissociated from increases in firing rates [15]. Finally, recurrent  
47 dynamics in prefrontal cortex can support context-dependent selection and integration of visual stimuli [16]. This  
48 has been demonstrated for cued switching between anatomically segregated stimulus features (such as color and  
49 motion), but it is not clear how this mechanism could generalize to the task of making decisions based on different  
50 local orientations and in the absence of an explicit cue. We also don’t know how the brain could learn the dynamics  
51 required for such late selection, from limited task experience.

52 Here we introduce a novel theory in which a stochastic modulatory signal induces shared variability in neural  
 53 responses to serve as a label of task relevance. We examine its implications in the context of a change detection  
 54 experiment in non-human primates [17; 18], with blocked task switches. We find that V1 neural responses exhibit  
 55 fluctuations that can be captured with a shared modulator that preferentially targets task-informative neurons. This  
 56 task-dependent covariability acts as a functional label that can be used to guide decoding, and can be learned within  
 57 a handful of trials, facilitating fast readout from the population. By studying stochastic modulation in an artificial  
 58 neural network model of the visual hierarchy, we find that task information can be read out using the modulator  
 59 label after additional stages of processing and with minimal amounts of task-specific feedback. As predicted by the  
 60 model, the V1 modulatory signal is also present in MT units, most strongly in task-informative MT units. These  
 61 results support the hypothesis that the task-specific labeling propagates through the visual hierarchy in parallel to  
 62 stimulus information, facilitating downstream decisions and actions.

## 63 2 Results

64 Monkeys were trained to detect a small change in orientation of a Gaussian-windowed drifting sine grating (Fig. 1A),  
 65 and spiking responses of neurons in their primary visual cortex (V1) and middle temporal area MT were recorded  
 66 simultaneously (Fig. 1B). Two to three gratings were present simultaneously, at high or low contrast levels, and  
 67 spontaneously changed their orientation. The animals were rewarded only for responding to changes of one of  
 68 these, with the others acting as distractors. The location of the relevant stimulus was fixed within each block of  
 69 trials, switching randomly between blocks throughout an experimental session. The two possible orientations of  
 70 the stimulus also switched between blocks. Monkeys were able to quickly adjust to these switches [18], reaching  
 71 asymptotic performance levels after a handful of trials (Fig. 1C-D). We aim to explain how the brain achieves this  
 72 impressive combination of accuracy and flexibility.

### 73 2.1 Encoding of local visual orientation in a V1 population

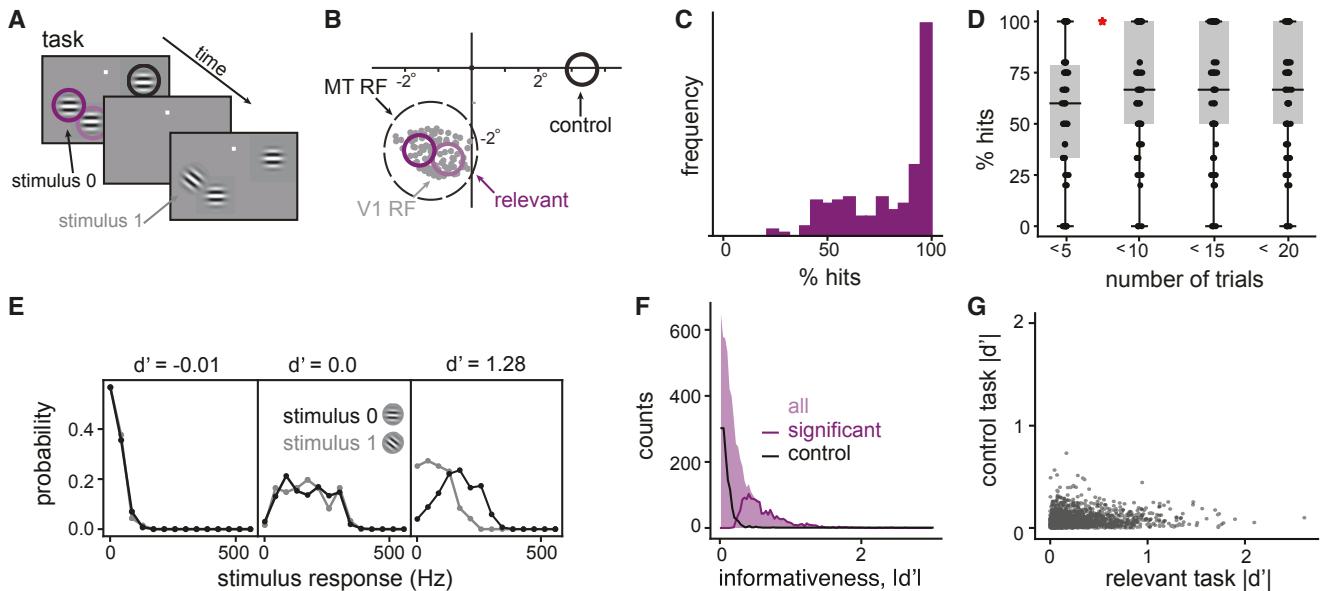


Figure 1: An orientation discrimination task with distractors. A) In each block of trials, 2-3 drifting gratings flash on and off on screen and can change their orientation. One stimulus is selected as relevant, and the monkey must report changes in its orientation with a saccadic eye movement. B) The recorded population of V1 neurons has receptive field centers (gray) within the receptive field of a simultaneously recorded MT unit. Two of the three stimuli locations are within the MT unit's receptive field (“relevant” - purple) and one is in the opposite hemifield (“control” - black). C) Distribution of behavioral performance across blocks, quantified by the % hits. D) Behavioral performance as a function of time within a block, binned using 5 consecutive trials; the boxes mark 25 and 75% quantiles, points indicate different blocks and the red star indicates a significant difference in means (relative two-sided t-test,  $p = 0.015$ ). E) The distribution of firing rates over all stimulus presentations, to each of the two task stimuli for three example neurons with different  $d'$  values. F)  $|d'|$  distribution, over all blocks of relevant tasks and all V1 neurons (shade). Lines indicate sub-distributions of neurons with significant informativeness (purple), and neurons in the control task (black). G) Relationship between the informativeness values in relevant and control tasks. A and B adapted from [18].

74 Neurons in V1 respond selectively to the local orientation of visual stimuli, and the selectivities of the full population span all orientations and visual field locations, in a topographical organization on the cortical sheet. In  
75 the experiment, individual grating stimuli are roughly matched to V1 receptive field (RF) sizes at the eccentricity  
76 at which recordings are performed, and orientation changes are relatively small (10-45°, see [18]), which restricts  
77 relevant stimulus information to a small subset of V1 neurons whose responses change with the stimulus orientation.  
78 As nearly all visual information passes through V1 [19], the behavior of the monkey must rely on the responses of  
79 this subset (the same throughout a block), while ignoring the background chatter of activity from the remainder of  
80 the population. Moreover, since downstream decision-making areas do not have access to V1 responses directly, the  
81 relevant information must be traced as it progresses through various stages of visual processing.

82 Two of the three stimulus locations were chosen so as to overlap the RFs of the recorded V1 population (Fig. 1A).  
83 When one of those locations is task-relevant, we expect a subset of the recorded neurons to provide information for  
84 the animal's decision ("relevant tasks"). In contrast, the neurons should be uninformative when the third stimulus  
85 location is task-relevant as it lies in the opposite hemisphere ("control task"; Fig. 1A). We quantified the task-  
86 informativeness of each V1 unit as the absolute difference in mean responses for the two orientations relative to  
87 response standard deviation ( $|d'|$ ). Figure 1E shows the relationship between informativeness and responsiveness for  
88 three representative examples. First, a large number of units are weakly responsive to both stimulus orientations  
89 (for instance, because their RFs did not overlap the stimulus location or because their preferred orientation was too  
90 different from the relevant stimuli) and consequently cannot be informative about stimulus identity (Fig. 1E, left).  
91 Second, some units respond strongly but similarly to both stimuli (Fig. 1E, middle), showing that responsiveness is  
92 necessary but not sufficient for task-relevance. Third, some units respond strongly to only one of the two stimuli  
93 and hence have high informativeness (Fig. 1E, right). Overall, for each relevant task block, a modest proportion of  
94 the recorded V1 units are significantly informative (monkey 1: 25.8%, monkey 2: 18.4%; non-parametric test, see  
95 Methods), whereas only 2.4% and 6% of units are significantly informative in the control task (Fig. 1F). Neurons that  
96 are most informative in either of the relevant tasks have low  $|d'|$  in the control task, reflecting their task-specificity  
97 (Fig. 1G). Across the two relevant tasks, unit informativeness is similar (61% of significant neurons are informative  
98 in both relevant tasks) because of the close proximity of the two relevant stimulus locations.

99 Within each task block, a different subset of V1 neurons carries task-relevant information. In order to make accurate  
100 decisions, a downstream circuit has to read out selectively from those, ignoring the rest. Moreover, the determination  
101 of this relevant subpopulation happens quickly: the monkey's performance reaches asymptotic levels roughly 5 trials  
102 after each task change (Fig. 1D). How can this flexible routing of information be achieved? Since basic response  
103 statistics such as mean or variance do not differ much between informative and uninformative neurons (Suppl. S1),  
104 they cannot guide this selection. Instead, we propose that task-specific structure in the joint statistics of neuronal  
105 responses [17; 20; 21] are key to understanding flexible readout.

## 107 2.2 A targeted shared stochastic modulator in V1

108 Neural responses fluctuate from trial to trial. Some of this variability is neuron-specific, but some is correlated  
109 across neurons, driven by circuit dynamics [22-24]. To determine the structure of co-variability, we fitted a modu-  
110 lated stimulus response model ("modulated-SR model") to the recorded population of V1 neurons in each block, using  
111 a Poisson latent dynamical system (PLDS, see Methods and [25]), which jointly estimates the stimulus drive to each  
112 unit and the shared, within-trial variability across the population (Fig. 2A, B). The stimulus response component  
113 ("SR model") accounts for stimulus-induced transients across multiple time bins of 50ms, with time-specific param-  
114 eters for each contrast condition (see Methods for details) and independent Poisson noise. The shared, within-trial  
115 variability is assumed to arise from a low dimensional dynamic stochastic signal, which multiplicatively modulates  
116 the stimulus responses of all units, with neuron-specific modulator coupling strengths. This statistical framework  
117 allows us to probe the existence, dimensionality, and structure of shared modulation in each block, in a way that  
118 simpler dimensionality reduction methods cannot achieve (Suppl. S2).

119 We found that 91% of blocks are better fit by the modulated-SR model than by the SR model alone (Fig. 2C).  
120 Moreover, varying the dimensionality of the modulator reveals that 72% of blocks are best described by a one-  
121 dimensional modulator (Fig. 2D; see Methods). For consistency, we restricted subsequent analyses to these blocks.  
122 The extracted modulator is unrelated to contrast variations in the stimulus (Suppl. S3) and fluctuates within and  
123 across trials at a fairly rapid timescale (Fig. 2B), with no evidence of oscillatory structure. The average estimated  
124 time scale of the fluctuations is 75ms (Fig. 2E) – faster than the average trial duration (3s) as well as the individual  
125 stimulus duration (200ms), and approaching the time resolution of spike count binning (50ms). This fast time  
126 scale, together with the unimodal marginal statistics of the estimated modulator (Suppl. S4), differentiate it from  
127 previously reported on-off dynamics [26].

128 The improvement in fit quality obtained by including the modulator varies across units (Fig. 2C), but is most  
129 prominent in task-informative neurons (Fig. 2F), suggesting that they may be more strongly modulated. A non-  
130 parametric comparison revealed that task-informative neurons have larger coupling weights (i.e. stronger modulation)

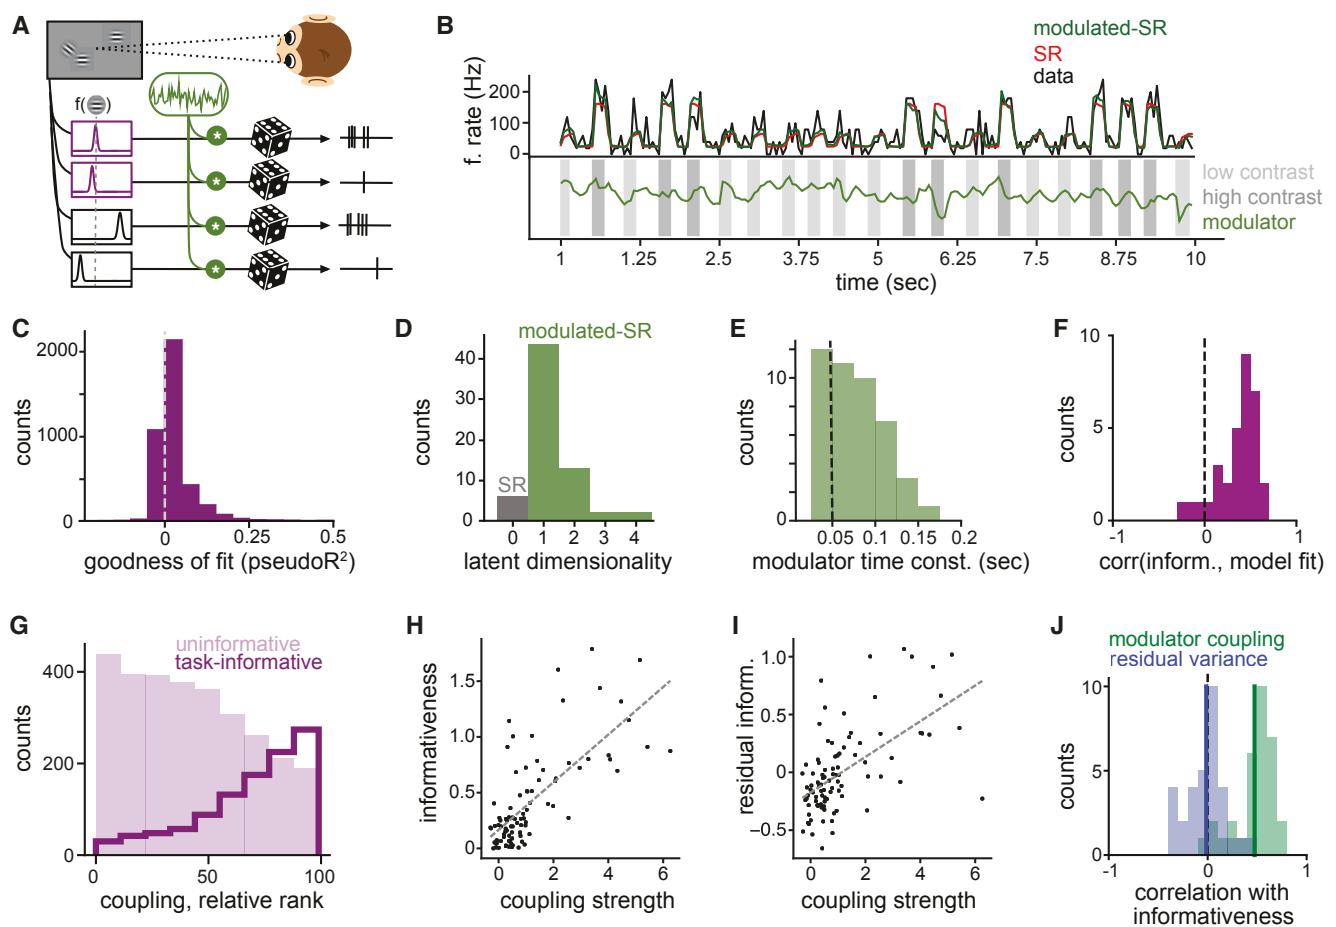


Figure 2: Estimating V1 modulator. A) An illustration of the modulated stimulus response model: Each neuron's tuning function is modulated by a time-varying shared source of multiplicative noise (green), with spiking modeled by a Poisson process. B) An example unit's activity over concatenated test trials of a block and the corresponding prediction of the SR model and the modulated-SR model. Bottom row shows the estimated trajectory of the modulator. C) The distribution of pseudo- $R^2$  values over all neurons in blocks that were best fitted by a 1-dimensional modulated-SR model. D) Summary of the dimensionality of best fitted models across relevant tasks. E) The distribution of estimated time constants over all blocks that were best fitted by a 1-dimensional modulated-SR model. F) Distribution of the correlations between the individual unit's model fit (pseudo- $R^2$ ) and their informativeness. (78% of blocks have significant positive correlations between informativeness and model fit, Spearman  $r$ ,  $p < 0.05$ ) G) Relative population rank of modulator coupling strength (within each block) for significantly informative (dark purple line) and uninformative (light purple shading) neurons. H) Informativeness vs. coupling strength for an example block. I) Residual informativeness (unexplained by linear effects of mean firing) vs. coupling strength in same example as H. J) Distribution of correlation coefficients obtained by partial correlation analysis across blocks (green, 84% of blocks significant Spearman  $r$ ) and a similarly obtained distribution that uses the modulated-SR model residual response variance as a proxy for neuron individual variance (blue).

than uninformative neurons (Fig. 2G). Although informativeness is correlated with the mean firing rate of a unit (Suppl. S5), a partial correlation analysis confirmed that firing rate differences cannot explain the inferred modulation targeting, as firing-rate-corrected informativeness and modulator couplings are significantly correlated in 84% of blocks (Spearman  $r$ ,  $\alpha = 0.05$ ; Fig. 2H-J). The increased variability in the task-relevant neurons (Suppl. S1) is primarily due to the modulation; residual variability unexplained by the modulated-SR model is generally not correlated with informativeness (Spearman  $r$  with  $\alpha = 0.05$ ; Fig. 2J); only 9% of blocks have significantly positive correlations between residual variability and informativeness (19% significantly negative). While most of this residual variability is neuron-specific, we also find weak, structured correlations in pairs of units which suggest additional sources of shared noise not captured by the model (Suppl. Fig. S2).

The modulator coupling is dissociable from traditional attentional effects on mean firing rate (Suppl. S7), which have been suggested to improve encoding precision of particular attended stimuli [27], and it cannot be explained by neural adaptation, as the degree of adaptation was uncorrelated with the quality of the fit of the modulated-SR model (Suppl. S8). Finally, the modulator structure cannot be explained by the fact that the response measurements are in the form of multi-unit spike counts (Suppl. S9). Overall, our analysis reveals that V1 responses are modulated by

145 a common fluctuating signal, and that the strength of this modulation in each unit reflects its task-informativeness.  
 146 From an encoding perspective, this seems counter-intuitive (Suppl. S10). Why would the brain inject noise specifically  
 147 in the few neurons that matter most?

### 148 2.3 Targeted modulation can facilitate decoding

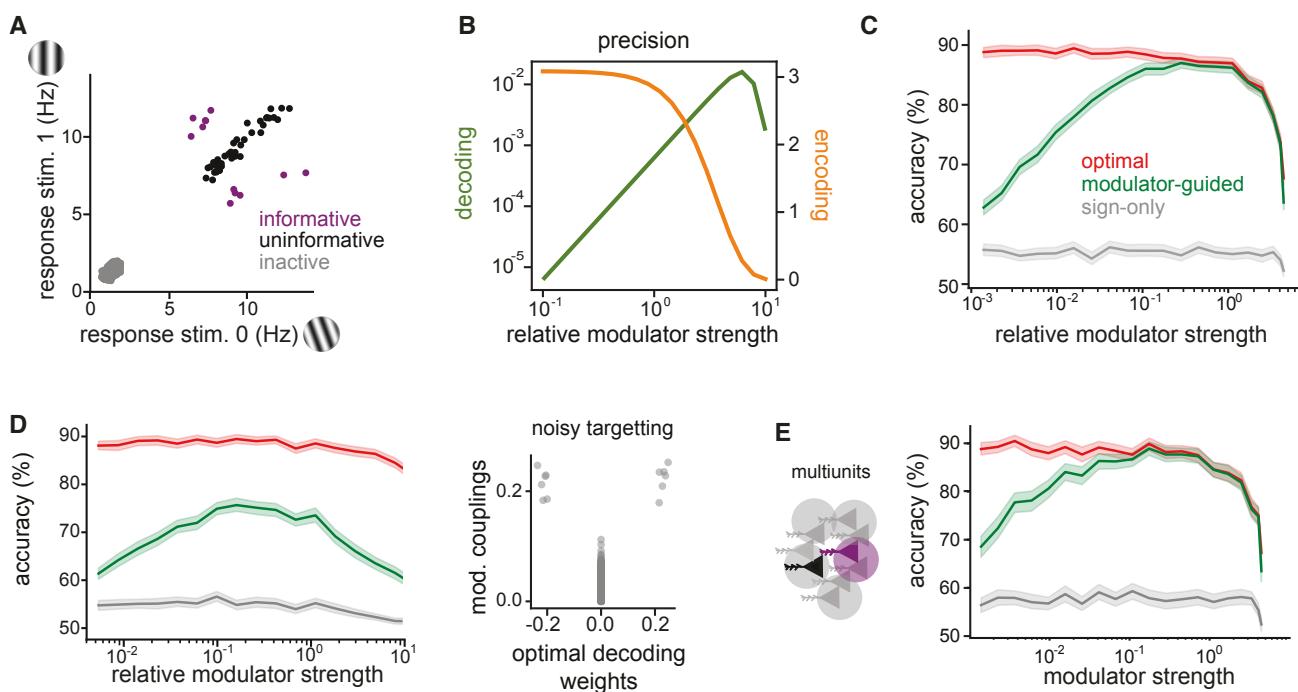


Figure 3: Theory of modulator guided decoding. A) The average response of neurons of the three subpopulations to two task stimuli. There are 12 informative, 38 uninformative and 4950 inactive neurons. B) Effects of increasing modulator strength on encoding and decoding, respectively, for modulator coupling weights equal to informativeness. Encoding is measured by the SNR, while decoding precision is quantified as the variance of the decoding weights of the modulator-guided decoder. C) Performance of three different decoders in simulations of a discrimination task with 1000 model V1 neurons, 50 informative, with increasing relative modulator strength (mean and 95% confidence interval). D) Same comparison as in C, but with the modulator coupling weights corrupted by Gaussian noise, as shown in the right panel. E) Decoder performance comparison for simulated multiunits, obtained by summing the activity of random pairs of neurons.

149 The modulator fluctuates rapidly, allowing any task information it provides to be accessed quickly, potentially on the  
 150 time scale of single trials. We hypothesize that the modulation serves to “label” the responses of the task-relevant  
 151 V1 subpopulation, so that downstream circuits can easily identify and use these signals.

152 To analyze the decoding process, we simulated an encoding model that captures the essence of the response properties  
 153 observed in the V1 data. For this, we use a variant of the modulated-SR model with static stimulus-dependent firing  
 154 rates, and one shared, temporally-independent stochastic modulator  $m_t$  (see Methods, and [9]):

$$k_{n,t}(s) \sim \text{Poisson}(\lambda_n(s) \exp(c_n m_t)), \quad (1)$$

155 where  $k_{n,t}(s)$  is the spike count of neuron  $n$  at time  $t$  in response to stimulus  $s$ ; the modulator  $m_t$  is drawn  
 156 independently from a Gaussian distribution with zero mean, and influences neuron  $n$  with coupling weight  $c_n$ ,  
 157 which is set to be proportional to the neuron’s task-informativeness. Finally, since the degree of modulation affects  
 158 not only variability but also mean responses, we explicitly correct for the mean increase to isolate the effects of  
 159 modulator-induced co-variability (see Methods).

160 Given this encoding model and a binary discrimination task,  $s \in \{0, 1\}$ , the ideal observer’s optimal decoder compares  
 161 a weighted sum of the neural responses with a modulator-specific decision threshold,  $c(m_t)$  (see Methods):

$$\sum_n a_n^{(\text{opt})} k_{n,t}(s) > q(m_t), \quad (2)$$

162 where  $a_n^{(\text{opt})} = \log(\lambda_n(1)) - \log(\lambda_n(0))$  denotes the optimal decoding weights. These are independent of the modulator

163 and equivalent to those derived from an independent Poisson model. The decoding weights are non-zero only for  
164 the small subpopulation of informative neurons (Fig. 3A, purple), with their signs indicating preference between the  
165 two stimulus alternatives. Zero weights eliminate active but uninformative (Fig. 3A, black) or inactive (Fig. 3A,  
166 grey) neurons.

167 The optimal decoder provides an upper bound on decoding performance given the encoding model, and motivates  
168 the use of a linear-threshold functional form for the readout. However, it uses weights that rely on full knowledge of  
169 each neuron's mean responses to the stimuli of the current task. The challenge for a downstream circuit is to find a  
170 way to approximate these weights, when provided only with incoming spikes, the task feedback, and potentially the  
171 modulator, but without explicit knowledge of the stimulus encoding model. How can the brain achieve this? The  
172 conventional means of learning decoding weights is regression. Although this is feasible for a small set of mostly  
173 informative neurons, the number of training examples needed for accurate weight estimation grows significantly with  
174 population size [28; 29]. So the behavioral flexibility exhibited by the monkeys precludes such a solution. Instead,  
175 we seek a heuristic that can be estimated quickly.

176 Consider first a decoder motivated by early work on neural binary discrimination [30]. The idea is to split all  
177 neurons into two sub-populations (“preferred” and “anti-preferred”) and then compare their average responses. This  
178 solution only assigns decoding signs ( $a_n^{\text{SO}} \in \{-1, 1\}$ ), which indicate relative stimulus preference, but ignores the  
179 relative importance of different neurons (there are no zero weights); we refer to this approach as the *sign-only* (SO)  
180 decoder. It can be learned quickly (Suppl. S10), but its performance falls as the fraction of informative neurons  
181 decreases (Suppl. S10): Since all neurons must be included, the noise from the uninformative neurons corrupts the  
182 decision signal. For realistically small fractions of informative neurons [2; 27], the SO decoder cannot match monkey  
183 performance (Suppl. S10).

184 To improve performance, the readout needs to consider the relative importance of individual neurons. A decoder can  
185 achieve this by estimating the amplitude of individual decoding weights. Since the relative strength of modulation  
186 of each neuron reflects the relative informativeness (by design  $c_n \propto |d'|$ ), we can define a *modulator-guided* (MG)  
187 decoder that sets its decoding weight amplitudes from temporal correlations of the modulator with each neuron's  
188 activity, which provide a simple estimate for  $c_n$ :

$$|a_n^{(\text{MG})}| \propto \frac{1}{T} \sum_t m_t k_{n,t}(s). \quad (3)$$

189 The MG decoder does not rely on knowledge of the response properties of the encoding population, but it assumes  
190 access to the modulator (e.g., it is a broadcast signal). This has important implications for learning the decoder;  
191 the MG weight estimates converge rapidly, on the time scale of the modulator fluctuations which are much faster  
192 than a trial (Sec. 2.2). Once the informative neurons have been identified, their decoding sign is determined based  
193 on explicit trial feedback, which only requires a handful of trials for small populations (Suppl. S10). For simplicity,  
194 the amplitude and sign were estimated separately here. Nonetheless, the two can also be learned jointly using a  
195 form of local online learning based on eligibility traces [31; 32] (Suppl. S11).

196 We compared the performance of different decoders in a binary discrimination task, based on simulated responses of  
197 a large population of V1 neurons with a small fraction of informative neurons (5%, Fig. 3A; see also Suppl. S10D for  
198 variations in percentage of informative neurons). The statistically optimal decoder corresponds to the ideal observer's  
199 solution, and thus provides an upper bound on achievable performance; the SO decoder provides a lower bound.  
200 The optimal decoder's accuracy deteriorates as the modulator increases in amplitude, corrupting the encoded signal  
201 (Fig. 3C). This reinforces the point that, unlike other forms of noise correlations [33; 29], the targeted, multiplicative  
202 noise is strictly detrimental for encoding (Suppl. S10). While the performance of the MG decoder is limited by this  
203 corruption as well, it also benefits from a stronger label in the informative neurons (Fig. 3B). Its performance follows  
204 an inverted U-shape as a function of modulation amplitude, reflecting the trade-off between these two opposing effects  
205 (Fig. 3C). MG decoding performance is maximized at an intermediate modulation amplitude, where it attains an  
206 accuracy close to that of the ideal observer, a result that is robust to variations in population size (Suppl. S10).

207 In practice, the performance of the MG decoder could depend on how strongly correlated the modulator couplings,  
208  $c_n$ , are with task-informativeness. To test the robustness of the MG decoder, we weakened the correlation between  
209 the modulator couplings,  $c_n$ , and task-informativeness by adding noise to  $c_n$ . We found that although performance  
210 decreases overall, the nonmonotonic dependence of the MG decoder performance on modulator strength is preserved  
211 (Fig. 3D). Given that our measurements mostly include multiunits, we also tested their impact on decoding and found  
212 that the results are qualitatively robust to such measurement noise (Fig. 3E). Interestingly, the optimal modulation  
213 amplitude generally shifts towards the range estimated from the data, suggesting that physiologically, the degree of  
214 modulation may be well-matched to the precision of the modulator targeting.

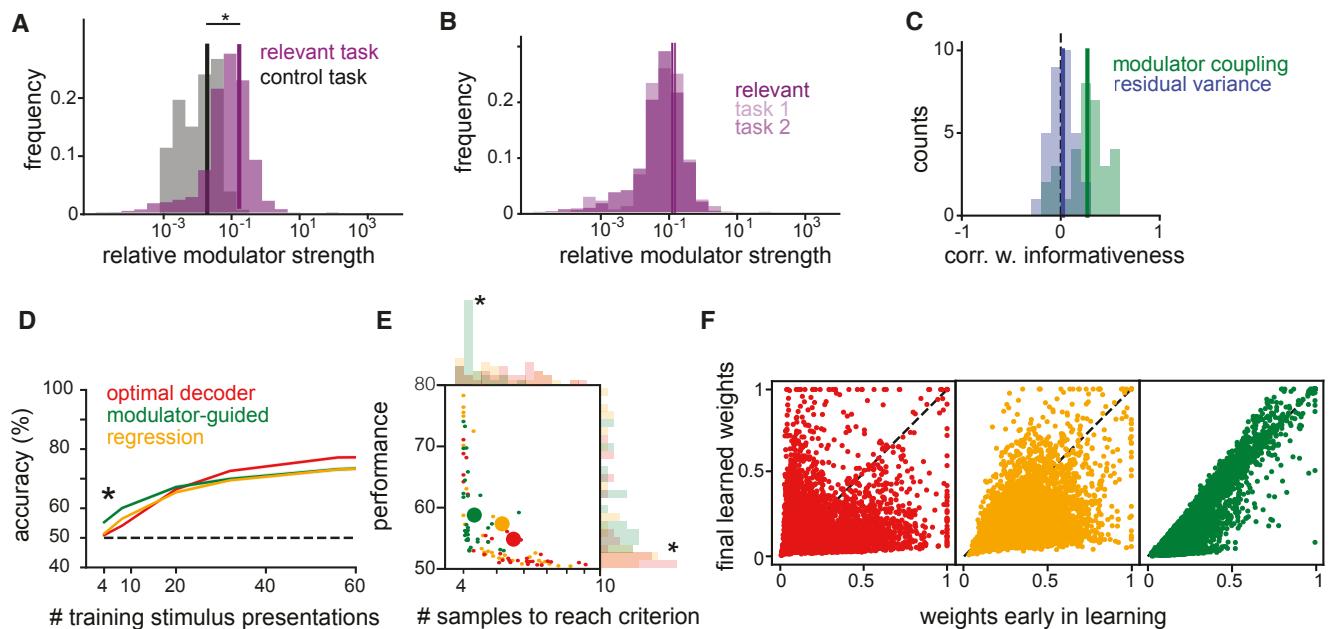


Figure 4: V1 modulator is task-specific and facilitates decoding. A) The distribution of relative modulator strengths across all relevant task blocks (purple) and all control task blocks (black). The star indicates significant difference between the two distributions (U-test,  $p < 0.001$ ). B) Same as A, but comparing across the two relevant tasks ( $p = 0.45$ ). C) The distribution of correlation coefficients between modulator coupling (green) or residual response variance (blue) and the residual behavioral relevance of a unit's activity (correlation with behavior), obtained by regressing out informativeness and mean firing rate. D) Decoding from the recorded V1 population: performance of different decoders or logistic regression for an example block population with increasing number of training samples (mean  $\pm$  SEM); star indicates significant differences between the optimal and the MG decoder. E) Performance with minimal training against minimal number of training samples (stimulus presentations) needed to reach above chance (50%) performance, for each block; stars indicate significant differences between the optimal and the MG decoder. F) Decoding weights estimated with maximum training (90% of all stimulus presentations) versus with minimal training (1%) for various decoders; same colors as D,E.

## 215 2.4 V1 modulator is task specific and facilitates decoding

216 In our experimental context, the theory predicts that the co-variability of neural responses should change based on  
 217 whether they are task-informative. Given that the recorded V1 population is informative in the relevant tasks but  
 218 not the control task (Fig. 1G), we expect differences in overall modulator strength across tasks and in individual  
 219 modulation strengths across neurons. Indeed, the overall strength of the estimated modulation significantly decreases  
 220 in the control task, both in absolute terms and relative to stimulus induced variations (Fig. 4A and Suppl. S12).  
 221 In comparison, the two relevant task conditions have indistinguishable statistics of overall modulation strength  
 222 (Fig. 4B). Our theory explains this difference as a change in labeling, from the recorded subpopulation that is  
 223 informative for the relevant tasks, to a different (unrecorded) subpopulation that is informative in the control task.

224 The comparison between the two relevant tasks is limited by the proximity of the two relevant stimulus locations, as  
 225 only few units are exclusively informative in one task (see Sec. 2.1). However, despite the reduced sample size, we  
 226 find a significant correlation between the difference in informativeness in the two relevant tasks and the difference  
 227 in coupling (Spearman correlation,  $r = 0.16$  with  $p < 0.05$ ), showing that units that are more informative in one of  
 228 the two tend to also have higher coupling in that task.

229 In our framework, decoding weights are approximated by estimating coupling strengths, and thus neurons with large  
 230 coupling (and thus strongly modulated) should have a stronger influence on behavior. Despite V1's early position  
 231 in the visual processing stream, we find this to be true in our data; 91% of blocks show significant correlations  
 232 (Spearman  $r$ ,  $\alpha = 0.05$ ) between modulator coupling and a unit's correlation with the monkey's behavior computed  
 233 as a  $d'$  of neural responses, with categories defined by the animal's choices rather than stimulus identity (see Methods).  
 234 Potential confounds in this analysis are not only overall firing rates, but also the informativeness of a unit, as the  
 235 most informative neurons would be expected to have a stronger influence on behavior [34; 35]. Nonetheless, even  
 236 after controlling for these confounds, it remains the case that units that are more modulated are the ones that are also  
 237 more predictive of behavior (Fig. 4C). This relationship is not present for the residual response variance (Fig. 4C).  
 238 Furthermore, we do not find a relationship with behavioral correlation in other shared noise sources (Suppl. S13),  
 239 which suggests that the shared modulator-induced fluctuations are particularly relevant for downstream processing.

240 The most direct prediction of the theory is the ability of the MG decoder to set appropriate decoding weights for the  
241 recorded V1 responses, and to do so rapidly, with limited data. To test these predictions, we decoded the stimulus  
242 identity from V1 responses using our heuristic MG decoder and compared its performance with that of the ideal  
243 observer for the estimated (modulated-SR) encoding model. When all available data is used for estimation, the MG  
244 decoder performance is close to that of the optimal decoder ( $\sim 80\%$  correct, which suggests that the strength and  
245 targeting precision of the estimated modulator is sufficient to guide decoding).

246 The optimal decoder provides an upper bound on decodability assuming perfect knowledge of the V1 response  
247 properties, but it can still perform poorly when the model is estimated from limited data; in fact, its performance is  
248 at chance in the low-data regime (Fig. 4D). Similarly, learning decoding weights directly through logistic regression  
249 requires many training trials before performing above chance (Fig. 4D). In contrast, the modulator-guided (MG)  
250 decoder finds informative units after only a few training examples, as it estimates the modulator coupling on the  
251 time scale of the modulator itself instead of that of trials. It outperforms the learned optimal decoder and logistic  
252 regression in the small training sample regime (comparing MG against either learned optimal or regression-based  
253 decoder significant; t-test,  $p < 0.0001$ , see Fig. 4D). We quantify this effect across all data and find that the MG  
254 decoder reaches above-chance performance significantly faster than the learned optimal decoder (t-test,  $p < 0.0001$ ,  
255 Fig. 4E) and that the performance attained with minimal training is significantly higher relative to that of the  
256 learned optimal decoder (t-test,  $p = 0.01$ ). The MG decoder also reaches above-chance performance significantly  
257 faster than a regression-based decoder (t-test  $p < 0.001$ ) and learned optimal and regression-based decoder do not  
258 differ significantly (t-test,  $p > 0.05$  for minimal training and performance). Our theory predicts that the advantage  
259 of the MG decoder lies in its ability to accurately estimate the decoding weights quickly. Indeed, we find a strong  
260 correlation between the MG decoding weights obtained with minimal training and those estimated from all available  
261 data, but this relationship does not hold for the learned optimal decoding weights or the regression weights (Fig. 4F).

262 Although significant, the difference in the number of trials required for above-chance performance may seem small.  
263 Nonetheless, it is likely that the benefits of modulation are substantially underestimated due to two experimental  
264 limitations. First, the recorded subpopulation is biased towards informative neurons since the stimuli are placed  
265 so as to drive these neurons. The animal must decode the information present in the entire V1 population, with a  
266 much lower percentage of informative neurons. Under such conditions, finding the few informative neurons from task  
267 feedback becomes even harder (Suppl. S10), and the benefits of modulation stronger. Second, the modulator may  
268 vary on a time scale faster than the stimulus-presentations of the experiment and model, which would allow an even  
269 faster estimation of the decoding weights (Eq. 3 could also be applied to single spikes). Finally, we found additional  
270 sources of co-variability not considered in the theory (measured as residual pairwise correlations, see Suppl. S6)  
271 which are consistent with previously documented effects of the task condition noise correlations [18]. These do not  
272 seem to interfere with the ability of the targeted modulator to facilitate decoding, suggesting that the theory is  
273 robust to deviations from the exact model assumptions. Overall, the benefits of the MG decoder for the V1 data  
274 provide strong support for the hypothesis that the brain could use task specific modulation to enable flexible task  
275 switching.

## 276 2.5 Learning modulator targeting in a hierarchical circuit

277 Visual information processing is hierarchical, and task-relevant information needs to propagate through several  
278 stages before reaching decision-making areas. Moreover, since receptive field sizes increase across stages of processing  
279 [36], localized task-specific information will diffuse in subsequent visual layers, making the task of identifying the  
280 subpopulation of relevant readout neurons even harder. Thus, the decoding problem identified in V1 persists, and  
281 likely worsens, in downstream areas. As a separate issue, while thus far we have assumed the correct modulator  
282 targeting to be already present in the circuit, the right degree of modulation for each neuron in a task needs to also  
283 be learned from experience. Can the modulator-guided readout still facilitate flexible and accurate task performance  
284 under these conditions?

285 To answer this question, we use an artificial neural network to model the visual processing hierarchy with a stochastic  
286 modulator and learned targeting. The first layer of the network consists of a V1-like encoding population with local-  
287 ized oriented filters, whose responses are then propagated through two processing layers of neurons with increasing  
288 RF size, and finally read out by a decision stage (Fig. 5A; details in Methods). To reflect previous experience, connec-  
289 tions between stages are pre-trained (via backpropagation), to solve a general image classification task (identifying  
290 handwritten digits [37] randomly positioned in different locations; Fig. 5B-C), in the absence of the modulator. As  
291 a result of this optimization, the model is capable of discriminating complex visual features.

292 Analogous to the V1 experiment, we use stochastic modulation to fine-tune this network to the task of discriminating  
293 the orientation of local gratings (Fig. 5D-E). After adjusting the decision circuit to the new data (see Methods for  
294 details), the network needs to perform a binary discrimination task involving two orientations at a fixed location  
295 (Fig. 5E). As in the actual experiment, distractors are placed at other locations in the image, something which  
296 the network has not encountered during the previous episodes of learning. We introduce shared, stochastic gain

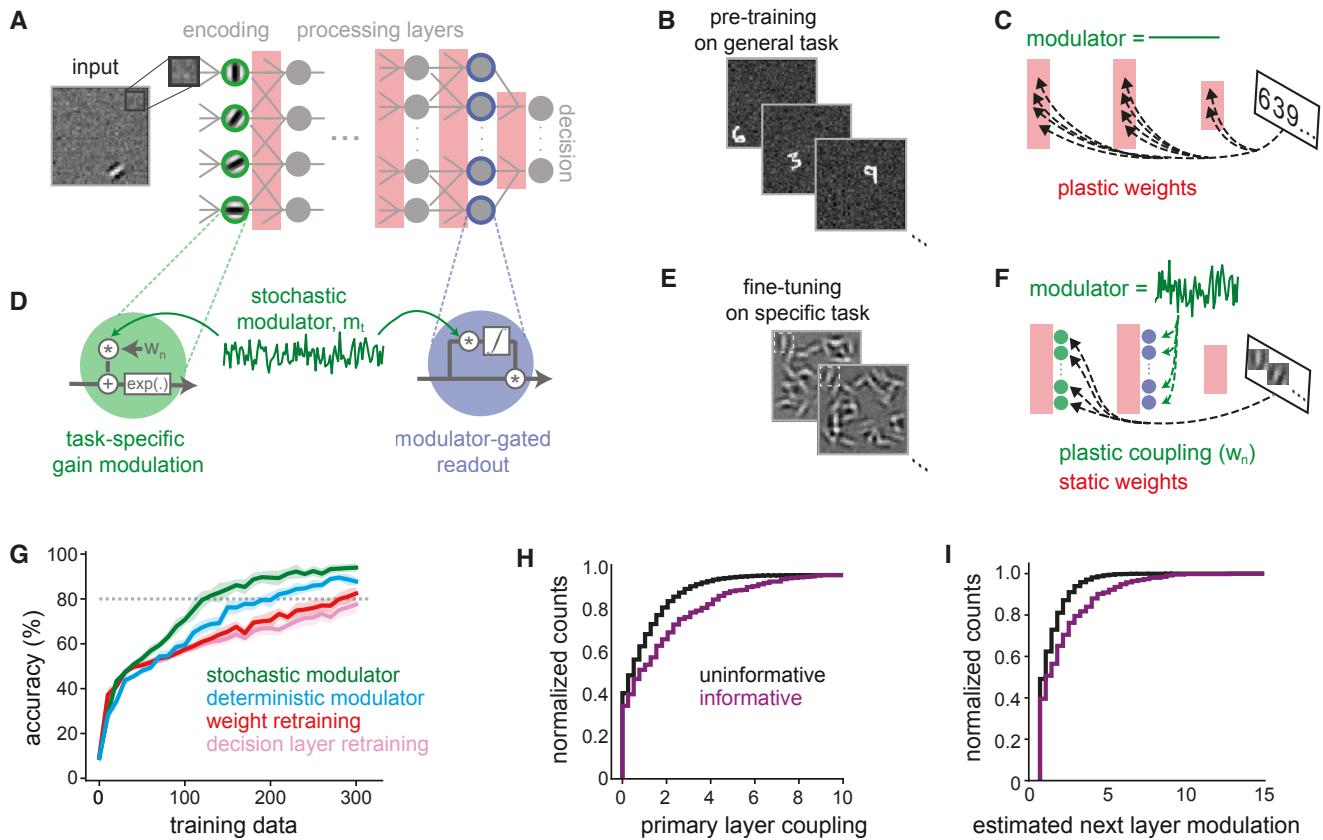


Figure 5: Learned stochastic modulation in a hierarchical network. A) Network with a primary encoding layer consisting of 2560 neurons with fixed Gabor filters with varying orientation and RF location, two locally connected processing layers and a fully connected decision layer. B) Pre-training on a spatially invariant version of the classic MNIST classification. C) Pre-training involves optimizing weights of the processing and the decision layer. D) A stochastic modulator  $m_t$  varies the gain of each primary layer neuron according to their coupling terms  $c_n$  (green). The fluctuations introduced at the primary layer guide the gain term of the input to the decision layer through a modulator-gated learning rule (blue). E) Task training involves binary discrimination of grating orientation at a particular location in the presence of distractors. F) Task-training involves learning the coupling terms  $c_n$  via task feedback and adjusting the modulator-gated readout accordingly. G) Performance of different decoding strategies, as a function of the amount of data used for task training. Grey dotted line indicates criterion performance. H) Distribution of task-optimized modulator coupling for most informative neurons (5% highest  $|d'|$  values) vs. all other neurons at the primary encoding layer. I) Estimated neuron-specific modulation at the first processing layer for most informative neurons vs. the rest.

297 modulation with neuron-specific coupling parameters in the primary encoding layer of the network (with the same  
 298 functional form as the original encoding model in Eq. 1, but without the Poisson noise; see Methods for details). This  
 299 injected variability accompanies the stimulus information across the processing layers. The responses of neurons in  
 300 the last layer are combined with gain terms  $g_n$ , which tune the readout of the decision circuit to the specific task  
 301 (Fig. 5D). As for the MG decoder in Eq. 3, these gains are adaptively computed using the correlations between  
 302 the individual neural responses and the modulator, which is again assumed to be available at the decision stage.  
 303 We optimize the modulator coupling strengths to maximize behavioral performance on the task, using explicit trial  
 304 feedback (via backpropagation). The general rationale is that if task-informative neurons can be modulator-labeled  
 305 in the V1 stage, then this labeling will be inherited downstream by exactly those neurons that receive their signal.  
 306 Thus their co-variability can guide decoding at the decision layer.

307 We assess the efficiency of the modulator-based solution by comparing it to two alternative models, both of which  
 308 adapt based on experience within the task, but which differ in their parameter complexity. At one extreme, we  
 309 consider a system that relearns the connection strengths between all layers de novo (“retraining”). This approach  
 310 corresponds conceptually to the regression model in Fig. 4. At the other extreme, we consider a fixed network  
 311 that only relearns the final readout weights (“readout only”). Retraining all network weights requires many training  
 312 examples to reach good performance (defined as > 80% accuracy; Fig. 5G), likely due to the high dimensionality  
 313 of the parameter space. Retraining only the decision layer results in poor performance, because the presence of  
 314 distractors renders the pre-trained representation insufficient for effective category discrimination. Compared to  
 315 alternative models, fine-tuning the network via the modulator substantially reduces the amount of task-training

316 required to reach criterion performance (Fig. 5G).

317 The improvement in performance of the modulator solution over regression-based relearning corresponds qualitatively  
 318 to what we found when decoding from the data in Fig. 4D). Nonetheless, one important distinction between this  
 319 hierarchical model and the previous MG decoder is that the modulator affects both the mean and the variance of  
 320 the V1-like encoding layer (see Methods). To disambiguate the effects of modulation on neural variability vs. mean  
 321 responses, we introduce a third model, which is parametrized and trained in the same way, but deterministically  
 322 boosts the gain of initial stage neurons [13], in the absence of stochastic modulation. We find that targeting of  
 323 deterministic gain modulation can be learned faster than retraining all the connections, but it does not reach the  
 324 same performance as the stochastic modulator given limited training. This suggests that the separation of stimulus  
 325 information and task relevance into the mean and variance of neural activity, respectively, further enhances the  
 326 identifiability of the stimuli at the decision stage.

327 When investigating the properties of the learned solution, we find that the learned couplings are highest for task-  
 328 informative neurons (5% highest  $|d'|$ ) in the primary encoding layer (Fig. 5H), as in the data (2F-J). Although the  
 329 modulator only affects the responses of these neurons directly, we find that informative neurons in the downstream  
 330 processing layer are also preferentially correlated with the modulator (Fig. 5I). This suggests that task relevance  
 331 propagates along the hierarchy in parallel to the stimulus information.

## 332 2.6 Modulator label is preserved in MT activity

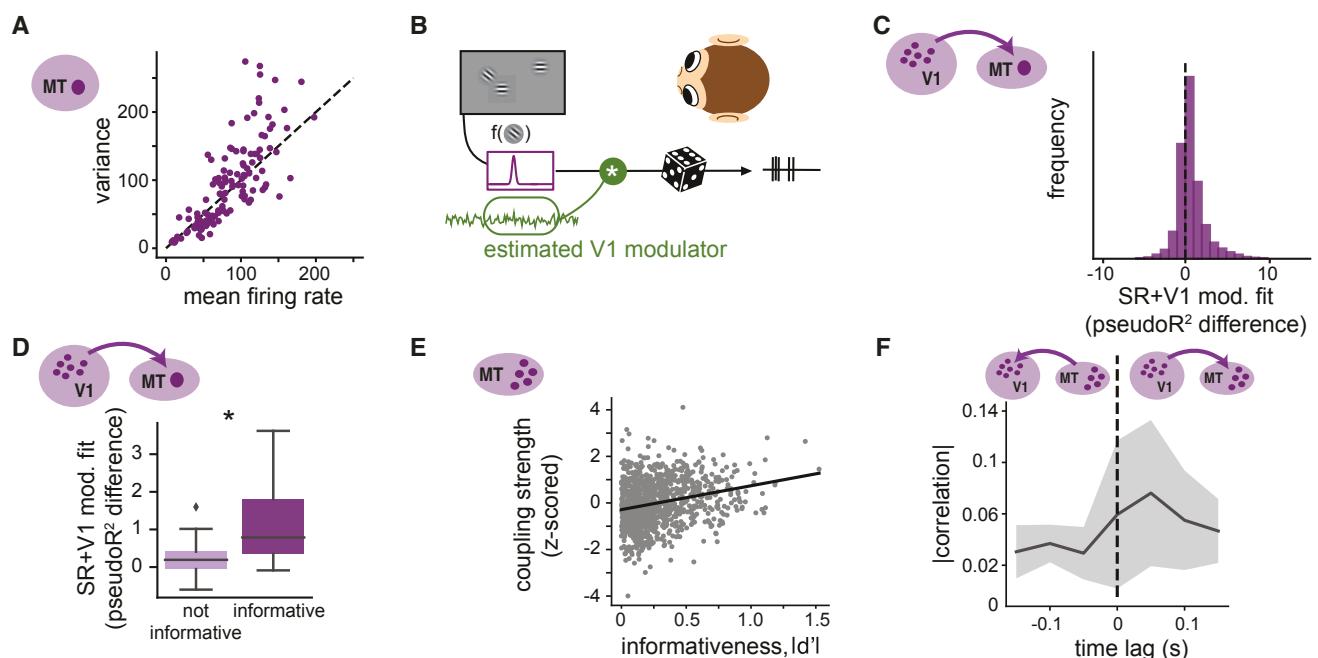


Figure 6: Effects of V1 modulator on simultaneously recorded MT units. A) Stimulus response variance as a function of mean firing for all MT units, and stimulus presentations. B) Schematic of the model; the spiking of each MT unit is specified by a tuning function potentially multiplicatively gated by the modulator estimated from V1 activity, with Poisson noise. C) Distribution of model fit ( $\text{pseudo-}R^2$ ) values obtained by comparing the log-likelihood of the SR+V1 modulation model against the SR model. D) Improvement in fit quality for the SR+V1 modulation model, grouping MT units into those with high informativeness values (50% with highest  $|d'|$ ) and those uninformative. Boxplot shows median and interquartile range. Black star indicates significant difference (t-test,  $p = 0.01$ ). E) A modulator is extracted from a population of MT cells. Shown are modulator couplings over informativeness in MT units over all 43 blocks. F) Correlations of the extracted V1 and MT modulators with positive (V1 before MT) and negative (MT before V1) time lag in seconds.

333 The model predicts that task-specific modulation introduced in V1 should label task-informative neurons in down-  
 334 stream areas. We look for signatures of such labeling in simultaneously recorded MT activity. MT neurons are  
 335 known to receive direct input from V1 [38] and selectively combine these afferents to construct their receptive field  
 336 properties, such as motion selectivity [39; 36]. Their receptive fields are larger and more complex, responding to  
 337 localized gratings with different combinations of position, speed and orientation [39; 40]. Given anatomical con-  
 338 siderations, we expect correlated activity in V1 to drive MT to some extent. What is specific to our theory is the  
 339 prediction that the degree of inherited modulation should reflect the task informativeness of individual MT units.

We find that responses of MT units that cover the two relevant stimulus locations (Fig. 1A) vary in their task-informativeness (Suppl. S15) and show different degrees of supra-Poisson variability (Fig. 6A), suggesting different levels of modulation [41]. The two measures are correlated across the MT units, with informative units having higher Fano factors (correlation coefficient of 0.48,  $p < 0.008$ ). To test whether the excess variability arises due to V1 modulation, we compared two models of MT activity. The first is based on the visual stimuli alone (“SR”); it resembles the V1 SR model, but includes stimulus drift direction (consistent with previous literature [39], drift direction did not have predictive power for the V1 units, see also [18], but has a strong effect on MT activity). The second model additionally conditions on the modulator estimated from the simultaneously recorded V1 units (“SR+V1 modulation”; Fig. 6B). The SR model provided a good fit for all MT units (Suppl. Fig. S6A), which is expected given that experimental stimuli were optimized to drive MT units. The inclusion of the V1-estimated modulator improved the fit for 73% of the MT units (measured as difference in pseudo- $R^2$ , see Methods; Fig. 6C). This effect is preferentially observed in task relevant units, which show a significantly larger model fit improvement relative to the uninformative units (t-test,  $p = 0.01$  Fig. 6D). Interestingly, this relationship was present only if the estimated V1 modulator showed significant targeting structure (significant Spearman correlations between coupling and informativeness); the few outlier blocks without structured V1 targeting could not explain MT variance (Suppl. S15).

The fact that both V1 and MT units are co-modulated as a function of their task informativeness is consistent with our theory, but does not exclude alternative patterns of information flow, such as top-down influences of MT on V1, or independent modulation of both areas from an external signal. To more directly address the nature of the modulation in MT we take advantage of a smaller set of MT population recordings (partly published in [18]). Despite the technical differences in recording procedure, this data recapitulates the same overall statistics, with 60% of the MT units having a significant part of their variability explained by the V1-estimated modulator. When independently extracting a modulator from the joint MT population responses (“SR+MT modulation”), we find that this population model better explains individual unit responses than the SR model (in 72 out of 73 blocks, Suppl. S16). The extracted modulator has consistent statistics across stimulus contrast variations (in 72% of blocks; Suppl. S16) and has similar time constants as those separately extracted in V1 (mean 61ms, s.d. 20ms). Lastly, there is a significantly positive correlation between modulator coupling and informativeness across blocks (Pearson  $r = 0.24$ ,  $p < 0.0001$ , Fig. 6E), suggesting that the same kind organization seen in V1 is qualitatively replicated in MT responses. Are these properties inherited from V1? We find that the cross-correlogram of the V1 and MT-extracted modulators is maximal at a time lag that is consistent with feedforward propagation from V1 to MT (Fig. 6F), although additional data and finer temporal precision will be required to more definitively understand this relationship. Altogether, our analysis of MT responses supports the idea that the modulation of task-relevant neurons in V1 is passed on to task-informative neurons in MT, allowing the propagation of labeling information towards decision areas.

### 3 Discussion

Humans and animals are impressive in their ability to respond rapidly and precisely to a variety of sensory stimuli, but the neural mechanisms supporting this flexibility remain poorly understood. We have presented a novel theory for flexible information readout, which uses a modulatory signal to induce shared response fluctuations in task-relevant cells, accompanying the task-relevant information as it propagates through subsequent stages of neural processing, and guiding decisions. We uncovered evidence for this labeling scheme in neural recordings from primate areas V1 and MT, obtained while the animals switch between local orientation discrimination tasks at different spatial locations. In particular, targeted modulation in V1 is sufficient to decode stimulus identity from neural responses after observing only a few trials. We also found evidence for the propagation of this modulator to informative neurons in downstream area MT.

The computational challenges faced by downstream circuits involved in decoding have been explored in seminal work by Shadlen and colleagues [30], who enumerated three potential factors that could reduce an animal’s behavioral performance compared to predictions of an optimal decoder (the “ideal observer”) operating on a hypothetical population of independent neurons: “*suboptimally stimulated neurons*” (in which the decoder includes irrelevant neurons in computing its decision), “*correlated noise*” (which worsens performance since it cannot be averaged out by the decoder), and “*pooling noise*” (additional noise in downstream circuits, whose contribution appears to be small [42]). The first factor has likely been underestimated in experimental data, since the recorded neurons are typically not representative of the full population due to experimental biases. As such, our conclusions regarding the benefits of targeted modulation for downstream readout are likely understated. The second factor, correlated noise, can either facilitate or impede stimulus encoding [23]. In particular, differential correlations, such as those reported in mouse V1 [43], are information-limiting. They restrict the encoding benefits that would otherwise arise from increasing population size [33] (but might support coding robustness [44]). Irrespective of correlation structure, identifying appropriate decoding weights using regression requires many trials [45], so flexible decoding remains a

397 problem. In contrast, although our modulator-induced correlations are also information limiting, their robustness  
398 to averaging enables the propagation of task relevance labels. Furthermore, their rapid time scale allows for the  
399 fast estimation of task-specific readouts. Finally, the changes introduced via the modulation are task-specific and  
400 ephemeral, allowing the circuit to instantly disengage from the task and revert to its original state, by reducing the  
401 strength of the modulator.

402 Top-down attention can facilitate sensory encoding, and has been shown to selectively affect neural responses,  
403 including increases in mean response [10–12], decreases in response variability [46], and decreases in noise correlations  
404 [27; 46; 47], all of which increase the signal-to-noise ratio (SNR) of the local sensory representation. These  
405 benefits for encoding are distinct from the modulatory effects we have explored here. They operate on the time  
406 scale of task conditions (minutes) or stimulus presentations (seconds), whereas the modulator that we estimate here  
407 fluctuates on a time scale of tens of milliseconds or faster. In addition, while attentional gain boosts are tuning-  
408 specific [47–49], we do not find evidence that they are specific to task-informative units (Suppl. S7). Moreover, the  
409 estimated modulator coupling is unrelated to the strength of attentional changes of the mean, suggesting that it  
410 may arise from separate mechanisms. This is consistent with effects of superior colliculus (SC) inactivation [15], and  
411 results documenting a similar dissociation between increases in mean and improvements in behavior over learning  
412 in V4 [50]. In the context of our theory, we hypothesize that SC inactivation may selectively disrupt the strength or  
413 targeting of modulation, affecting the propagation of task-relevant information to decision areas, a prediction that  
414 can be tested experimentally.

415 Our modulator is distinct from slow multiplicative, low-dimensional noise reported in other contexts [51; 52], which  
416 may serve other functional roles such as encoding uncertainty in visual areas. It is also distinct from gain changes  
417 due to fluctuations in attention which operate on the time scale of seconds [53]. Such signals are too slow to serve as a  
418 labeling mechanism of the type proposed here. Choice-related feedback signals have also been shown to modulate  
419 neural activity on a trial-by-trial basis, but they also occur on a slower time scale of several hundreds of milliseconds  
420 or seconds [54; 21]. The modulatory process of our theory does not replace, but coexists with these additional forms  
421 of gain modulation.

422 Shared oscillatory structure induces low-dimensional covariability and has been proposed as a mechanism for binding  
423 information across neurons [55]. The “communication through coherence” (CTC) theory [56; 57] formalizes this idea  
424 in an encoding-decoding framework, in which a top-down oscillatory modulator projects to both encoding neurons  
425 with the same feature selectivity, and to the decoding network that needs to read them out. The modulators we’ve  
426 extracted from our population recordings fail to show significant periodic structure. Beyond this, the CTC theory  
427 differs from our own in two important ways. First, oscillations target feature-selective rather than task-informative  
428 neurons [56]. These could be the same for a detection task, but differ for a discrimination such as that used in our  
429 experiment. Second, the CTC decoder uses a fixed (as opposed to a modulator-dependent) threshold, which we’ve  
430 shown to be suboptimal. Overall, the CTC framework describes a fixed labeling strategy based on tuning properties,  
431 while our theory proposes modulatory labeling adapted to task structure.

432 Some tasks, such as the context-dependent sensory evidence integration experiments by Mante and Sussillo [16], can  
433 achieve flexibility through the reorganization of late decision stages. We believe these mechanisms cannot explain  
434 flexibility in a low-level sensory discrimination task as presented here. First, numerical experiments using our  
435 hierarchical model demonstrate that it is particularly hard to achieve good performance in our task when adapting  
436 the readout alone. In addition, the recurrent dynamics supporting task switching are trained through extensive  
437 optimization [16] and although several proposals exist for the biological implementations of such learning [58], all  
438 require vast amounts of task experience. A final distinction is that our approach does not rely on an explicit  
439 context cue: the task relevance of sensory features is communicated solely through task feedback. Overall, multiple  
440 mechanisms for task-specific readout are likely to coexist in the brain and be engaged in a context dependent manner.

441 Our theory is agnostic to the source of the modulator and the circuit mechanisms underlying its task-specific  
442 targeting, but some previous studies provide potential clues. Changes in noise correlations across tasks could arise  
443 through either local circuit dynamics [24] or top-down mechanisms [21; 59], and later propagate to downstream  
444 regions. Given the sparsity of top-down connections relative to the full population size (at least, in V1), the  
445 reorganization of noise correlations likely needs to involve local recurrent dynamics, potentially taking advantage  
446 of its topographic organization. If this kind of spatially localized modulation was indeed an organizing principle of  
447 neural activity, it would predict that flexible decoding is most effective for tasks relying on sensory features that  
448 are localized in some brain area. Consistent with this idea, Nienborg and Cumming found that V1 neurons’ choice  
449 probability was significantly larger for orientation discrimination than for disparity discrimination, suggesting that  
450 V1 shows decision-related activity only if the task features are localized in the columnar organization [34]. Moreover,  
451 in a task involving higher order features, Koren et al. found neural variability was high in V4, but not V1 suggesting  
452 that the modulator could target later stages of processing depending on the task [60]. Regarding the physiological  
453 origins of our modulator, one potential source for low dimensional broadcast signals could be thalamic nuclei that  
454 integrate sensory and top-down information [61; 62]. Alternatively, it may be possible to eliminate the need for a  
455 copy of the modulator at the readout stage, by estimating the signal directly from the observable correlations in

456 population activity.

457 The lack of a biologically plausible theory of neural decoding is a fundamental shortcoming in our current under-  
 458 standing of neural computation. Resolving the puzzle of how sensory information is routed through brain regions  
 459 and extracted to perform specific tasks is critical for the study of sensory and cognitive dysfunction, including clinical  
 460 applications such as brain-computer interfaces (BCI) [63]. Moreover, flexible task-dependent information routing  
 461 poses a fundamental obstacle for the development of adaptive artificial intelligence systems. Our work provides a  
 462 novel framework for solving this problem, supported by both physiological data and computational theory.

## 463 4 Methods

### 464 Theoretical framework for decoding from a neural population

We simulated a binary discrimination task analogous to the experiment, which requires discriminating stimuli  $s = 0$  from  $s = 1$  on the basis of the activity of a population of  $N$  neurons. Neural responses are modeled as Poisson draws with a stimulus-dependent firing rate, which is itself modulated by a time-varying noisy signal,  $m_t$ , shared across neurons:

$$465 \quad k_{n,t}(s, m_t) \sim \text{Poisson}(\lambda_n(s) \exp(c_n m_t)), \quad (4)$$

466 where  $\lambda_n(s)$  is the stimulus response function of the neuron, and  $t$  indexes time within a stimulus presentation. The  
 467 modulator  $m_t$  is 1-dimensional i.i.d. Gaussian noise with zero mean and variance  $\sigma_m^2$ ; the nonlinearity  $\exp(\cdot)$  ensures  
 468 that the final firing rate is positive. The degree of modulation is neuron specific, parametrized by modulation  
 469 weights  $c_n$ , which we take to be proportional to the  $n$ -th neuron's ability to discriminate the two stimuli,  $c =$   
 470  $|\log(\lambda_n(1)) - \log(\lambda_n(0))|$ . We normalize responses by the expected increase in mean rate due to the modulator,  
 471  $\exp\left(\frac{\sigma_m^2 c^2}{2}\right)$  to compensate for systematic differences in mean firing rate due to modulation. The relative modulator  
 472 strength is defined as the ratio between modulator-induced and stimulus-induced variance.

473 Given this modulated Poisson encoding model, an ideal observer decides the stimulus based on the sign of the log  
 474 odds ratio, which reduces to comparing a weighted linear combination of the observed neural spike counts against a  
 modulator-dependent time-varying threshold (see also [9]):

$$475 \quad \sum_n a_n^{(\text{opt})} k_{n,t} > q^{(\text{opt})}(m_t), \quad (5)$$

476 with weights

$$477 \quad a_n^{(\text{opt})} = \log(\lambda_n(1)) - \log(\lambda_n(0)), \quad (6)$$

478 and time-varying threshold

$$479 \quad c^{(\text{opt})}(m_t) = - \sum_n \exp(m_t c_n) [\lambda_n(1) - \lambda_n(0)]. \quad (7)$$

The modulator-guided heuristic decoder assumes access to the modulator  $m_t$  and the neural responses  $k_{n,t}$ , and  
 learns approximate decoding weights based on co-fluctuations of the two within a trial:

$$480 \quad \left| a_n^{(\text{MG})} \right| = \frac{1}{T} \sum_t m_t k_{n,t}. \quad (8)$$

481 The sign of the decoding weight is separately estimated by comparing responses to the two stimuli (trial feedback;  
 482 see also [9] and Suppl. S10).

The sign-only decoder subtracts the summed responses of two subpopulations (i.e., a linear decoder with weights  
 $\pm 1$ ):

$$483 \quad a_n^{(\text{SO})} = \text{sign}(\lambda_n(1) - \lambda_n(0)). \quad (9)$$

### 484 Hierarchical information propagation with learned stochastic modulation

We use a 4 layer artificial neural network that maps an image stimulus with 3136 pixels into categories, corresponding  
 to 10 digits or different orientations. The first encoding layer includes neurons with fixed Gabor receptive fields.

The modulator affects encoding neurons through coupling terms  $c_n$ , which modulate the neuron's responses:

$$485 \quad h_{n,t}^{(0)} = \exp\left(\mathbf{w}_n^{(0)} \mathbf{s} + m_t c_n\right), \quad (10)$$

where  $h_{n,t}^{(0)}$  is the activity of neuron  $n$  in the encoding layer,  $\mathbf{w}_n^{(0)}$  are the weights from the input to this neuron. Neurons in the top layer include a multiplicative gain  $g_n \geq 0$ :

$$h_{n,t}^{(2)} = g_n \text{ReLU} \left( \mathbf{w}_n^{(2)} \mathbf{h}_t^{(1)} + b_n^{(2)} \right), \quad (11)$$

where  $b_n^{(2)}$  is a neuron-specific bias, optimized together with the weights  $\mathbf{w}_n^{(2)}$  during pre-training. The gain  $g_n$  is learned using the MG correlation rule:

$$g_n = \frac{1}{T} \sum_t m_t h_{n,t}^{(2)}(\mathbf{s}), \quad (12)$$

where  $h_{n,t}^{(2)}(\mathbf{s})$  denotes the activity at time  $t$  of neuron  $n$  in the last processing layer, in response to stimulus  $\mathbf{s}$ .

There are three stages of learning. 1) Pre-training optimizes all network weights to natural image statistics using a digit classification task (locally placed MNIST digits [37] with image presentation and pixel specific i.i.d. additive Gaussian noise), while  $m_t = 0$  and  $g_n = 1$ . 2) Learn an orientation discrimination readout from the neural responses of the fixed pretrained network (10 categories), when the input consists of single, local oriented gratings at various positions (14x14 positions). 3) Optimize the modulator targeting for an orientation discrimination task at one fixed task location, in the presence of distractors. The task involves binary discrimination of two oriented gratings with distractor gratings at other locations. At the fast time scale  $t$ , the modulator varies with 100 time points per stimulus presentation, i.i.d.  $m_t \sim \mathcal{N}(0, 0.1)$ , which drives gain changes in the last layer (Eq. 12). At the slow scale (stimulus presentations)  $m = 1$  and the coupling strengths  $c_n$  are optimized by backpropagation.

We compare the performance of our model (“stochastic modulator”, 2560 parameters for backpropagation, 7840 parameters including MG gain adjustment) to three controls: 1) full retraining of all connections (“weight retraining”, 256690 parameters), 2) retraining the decision layer weights (“decision layer retraining”, 78410 parameters), 3) all network weights are fixed, but the modulator is active  $m_t = 1$ , but constant, and the modulator coupling  $c_n$  are optimized for the task (“deterministic modulator”, 2560 parameters). In the first two approaches  $m_t = 0$  and  $g_n = 1$ .

## Population recordings in V1 and single units from MT

In experiments by Ruff and Cohen [18], two adult male rhesus monkeys performed a motion direction change detection task on one out of 2 – 3 oriented drifting gratings at high or low contrast on a screen. The task-relevant grating is indicated by a few instructional stimulus presentations, selected randomly for each block within the session (3 – 6 blocks per session). Most recording sessions analyzed use a 10 by 10 microelectrode array (Blackrock Microsystems) in area V1 and a recording chamber with access to area MT, allowing simultaneous recordings in the two areas (multiunit activity, details in [18]).

Two stimuli were positioned to drive the MT unit similarly and one stimulus was positioned outside of the MT RF. Within a block, changes in one out of the three stimuli had to be reported. In each trial, gratings flash on (200ms) and off (200-400ms) at the same orientation (repeated, stimulus 0) until a change occurs at an unknown time (target, stimulus 1). Stimuli vary in both contrast and orientation, randomly interleaved. We analyzed 67 blocks of 20 recording sessions across two monkeys where the task-relevant stimulus was positioned in the RF of the population (relevant tasks) and 20 blocks of 20 sessions where the stimulus outside of the RF was task-relevant (control task). Control and relevant task blocks were interleaved within a session. Neural populations may overlap across sessions.

We analyze 21 – 109 trials per block, where the monkey either detected the target (hit) or failed to detect it (miss). We discard trials where the monkey did not finish the task in a hit or miss and trials where one of the distractors changed orientation. This yields an average of 54 trials per block, each with several stimulus repeats and completed by a target presentation ( $s = 1$ , orientation-change). We only include blocks with a minimum of 20 valid trials (77 out of 90 blocks), as numerical simulations suggest 20 trials to be the minimum necessary to estimate informativeness reliably. Varying this criterion does not qualitatively change the results. The first stimulus in a trial was always removed to eliminate adaptation effects [27]. We only include units whose response to either one of the stimuli (presented individually) was at least 10% larger than baseline, to avoid inclusion of noise channels. On average 88 units ( $\sim 90\%$ ) in a block showed stimulus modulation for one of the two stimuli placed within the MT RF (min 52, max 95). We further exclude units with a Fano factor  $> 5$  standard deviations above the population average as this suggested especially many/diverse neurons in the unit and firing rates  $< 1\text{Hz}$  (in total 0 – 3 units were excluded per block).

## MT population recordings

An additional set of sessions (14 sessions with a total of 73 task blocks) in the same task had either exclusively MT recordings (24 channel probes) or simultaneous V1 and MT recordings, with stimuli placed to optimally drive the

525 MT units (but not necessarily V1).

## 526 Informativeness of a unit

527 The informativeness of a unit is quantified by  $d' = \left| \frac{\mu_0 - \mu_1}{\sqrt{0.5(\sigma_0^2 + \sigma_1^2)}} \right|$  where  $\mu_0$  and  $\sigma_0^2$ ,  $\mu_1$  and  $\sigma_1^2$  are the means  
 528 and variances of a unit's responses to the task-relevant stimulus 0 and stimulus 1, respectively. We compute  
 529 informativeness across all stimulus presentations in behaviorally correct trials of the same block. Significance is  
 530 assessed w.r.t. a null-distribution of  $d'$  values, constructed by comparing mean and variance of random subsets of  
 531 stimulus 0 responses ( $p=0.01$ ).

## 532 SR model

Stimulus effects are modeled with Linear-Nonlinear Poisson (LNP), taking into account effects of repeated stimulus presentations of stimulus 0, time varying in 50ms time bins and the effects of contrast (V1) or contrast+direction (MT). Orientation is not one of the stimulus dimensions as it does not change during the repeated stimulus presentation. Responses to target stimulus 1 are used only to compute informativeness and for decoding. Stimuli are parametrized by a one-hot encoding vector with 4 time-windows during 200ms stimulus presentation; this yields 8 stimulus dimensions for the contrast-specific V1 model, with one additional dimension indicating the stimulus drift direction in MT. We add one after-stimulus dimension to capture potential delayed effects of the stimulus presentation, and an offset for base firing:

$$\mathbf{k}_{n,t} \sim \text{Poisson}(\exp(\boldsymbol{\beta}_n \mathbf{s}_t)) \quad (13)$$

with spike counts measurements  $\mathbf{k}_n$ . Parameters  $\boldsymbol{\beta}_n$  are obtained by maximizing the log-likelihood of the data, separately for each block:

$$L(\boldsymbol{\beta}_n) = \sum_t -(\boldsymbol{\beta}_n \mathbf{s}_t)^T \mathbf{k}_{n,t} + \exp(\mathbf{1}^T \boldsymbol{\beta}_n \mathbf{s}_t) + \alpha \boldsymbol{\beta}_n^T \boldsymbol{\beta}_n. \quad (14)$$

533 The extended MT SR model includes the (normalized) V1 modulator as an additional predictive variable.

## 534 Modulated SR model

535 We use the framework of Poisson Linear Dynamical Systems (PLDS, [25; 20]), to model the temporal dependencies  
 536 within a trial while treating different trials as independent. The modulator terms of the PLDS are shared across  
 537 the population and influence each unit's activity through a linear mapping function  $\mathbf{C}$  (equivalent in meaning to the  
 538 coupling  $\mathbf{c}$  in the theory). This joint model has the form:

$$\mathbf{k}_t \sim \text{Poisson}(\exp(\mathbf{C}\mathbf{m}_t + \mathbf{B}\mathbf{s}_t)) \quad (15)$$

$$\begin{aligned} \mathbf{m}_{t+1} &= \mathbf{A}\mathbf{m}_t + \epsilon \\ \epsilon &\sim \mathcal{N}(0, \mathbf{Q}) \\ \mathbf{m}_0 &\sim \mathcal{N}(0, \mathbf{Q}_0) \end{aligned} \quad (16)$$

539 where the modulator  $\mathbf{m}_t$  at time  $t$  (within a trial across both stimulus presentation and inter-stimulus windows), is  
 540  $D$ -dimensional and the mapping  $\mathbf{C}$  is  $N \times D$ , with latent dimensionality  $D \ll N$ . Parameter  $\mathbf{A}$  implicitly defines the  
 541 modulator's time constant ( $\tau = -\frac{1}{\log(A)}$ , for 1d latents), while  $\mathbf{Q}, \mathbf{Q}_0$  define the noise covariance of the modulator.  
 542 The full model is fitted to data using the EM algorithm with a Laplace approximation for the E step (see [25]);  
 543 latent dimensionality is determined by model comparison ( $D = 0 - 4$ ).

## 544 Models validation and comparison

545 All models are 10-fold cross-validated, with model quality evaluated by 1) log-likelihood of test data (or the corre-  
 546 sponding leave one neuron out predictions from [64] for the PLDS, averaging over latent posterior uncertainty by  
 547 sampling), 2) variance explained by the model and 3) the pseudo- $R^2$  [65] which gives “the fraction of the maximum  
 548 potential log-likelihood gain (relative to the null model) achieved by the tested model”  $\frac{\log L(\hat{y}) - \log L(\bar{y})}{\log L(y) - \log L(\bar{y})}$ , where  $\hat{y}$  is  
 549 the estimation of the hypothesized model and  $\bar{y}$  is the null model. The null of the SR model had no stimulus-related  
 550 dimensions with average firing as the only explanatory variable. The SR model served as null for the PLDS.

551 For a fraction of the population the SR model ( 30% of neurons) does not improve prediction over a constant rate  
 552 model, suggesting that those neurons are not modulated by the stimulus. As expected, informative neurons show

553 significant improvements in fit quality from the SR model relative to the null (only 5% of informative neurons do  
554 not show improvements).

## 555 Modulator targeting

556 For Fig. 2G we computed the rank of each modulator coupling in its own block-specific population and compare  
557 the distribution of significantly informative to uninformative units. In Fig. 2H-J we used partial correlations to  
558 test for a relationship between unit's modulator coupling and task-informativeness in each block not explained by  
559 differences in overall firing rate. Specifically, we report the Spearman correlation between residual informativeness  
560 (after linearly regressing firing rate) and modulator coupling.

## 561 Modulator strength

562 When assessing the overall modulation strength, both the mapping  $\mathbf{C}$  and the modulator variance need to be  
563 considered jointly (as scaling up the mapping and decreasing the variance leaves results unchanged). We quantify  
564 the overall modulator strength as the variance of the modulator multiplied by the coupling norm  $\sqrt{\sum_n C_n^2}$ . The  
565 relative modulation strength is obtained by comparing to the stimulus drive, given by  $\sum_{n,i} \text{Var}(s_i B_{n,i})$  for each  
566 neuron  $n$ , where  $i$  indicates the stimulus dimension.

## 567 Linking behavioral choice to neural activity

568 We compute the difference in target-response between trials with correct target detection and those where the monkey  
569 missed the target, normalized by their variance  $|\frac{\mu_1 - \mu_2}{\sqrt{0.5 * (\sigma_1^2 + \sigma_2^2)}}|$  where  $\mu_{1,2}$  and  $\sigma_{1,2}^2$  are the means and variances of  
570 activity corresponding to the two choices, respectively. This provides an estimate of how involved a unit was in  
571 the choice of the animal. To asses the relationship with modulator strength we use a partial correlation with two  
572 covariates, firing rate and informativeness (by multivariate linear regression).

## 573 Decoding

574 We train each decoder on data that includes a balanced number of stimulus 0/1 presentations at high and low  
575 contrast, varying the size of the training set from the minimum 4 (one for each stimulus-contrast pair) to all  
576 available data. Decoder performance is tested on held out data. The optimal decoder uses maximum likelihood  
577 estimates (as in theory, with a 200ms decoding window), but based on estimated instead of ground truth parameters.  
578 It uses a constant threshold which is optimized on the training data. This is known to be suboptimal (Eq. 2), but is  
579 more robust to the noise in the data and therefore performs better in the limited data regime. The modulator-guided  
580 (MG) decoder estimates readout weights by taking the inner product between the unit's activity and the modulator  
581 values (Eq. 8, using 50ms bins), with signs determined from trial-level feedback, and a constant threshold.

## 582 References

583 [1] Kang, I. & Maunsell, J. H. The correlation of neuronal signals with behavior at different levels of visual cortex  
584 and their relative reliability for behavioral decisions. *Journal of Neuroscience* **40**, 3751–3767 (2020).

585 [2] Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A. A relationship between  
586 behavoiral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience* **13**, 87–100 (1996).

587 [3] Geisler, W. S. & Albrecht, D. G. Visual cortex neurons in monkeys and cats: Detection, discrimination, and  
588 identification. *Visual Neuroscience* **14**, 897–919 (1997).

589 [4] Dayan, P. & Abbott, L. F. *Theoretical neuroscience* (MIT Press, Cambridge, MA, 2005). 0-262-04199-5.

590 [5] Ma, W. J., Beck, J. M., Latham, P. E. & Pouget, A. Bayesian inference with probabilistic population codes.  
591 *Nature Neuroscience* **9**, 1432–8 (2006).

592 [6] Jazayeri, M. & Movshon, J. A. A new perceptual illusion reveals mechanisms of sensory decoding. *Nature* **446**,  
593 912–915 (2007).

594 [7] Graf, A. B. A., Kohn, A., Jazayeri, M. & Movshon, J. A. Decoding the activity of neuronal populations in  
595 macaque primary visual cortex. *Nature Neurosci.* **14**, 239–245 (2011).

596 [8] Berens, P. *et al.* A fast and simple population code for orientation in primate V1. *Journal of Neuroscience* **32**,  
597 10618–10626 (2012).

598 [9] Haimerl, C., Savin, C. & Simoncelli, E. Flexible information routing in neural populations through stochastic  
599 comodulation. *Advances in Neural Information Processing Systems* **32**, 14402–14411 (2019).

600 [10] Moran, J. & Robert, D. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**,  
601 782–784 (1985).

602 [11] McAdams, C. J. & Maunsell, J. H. R. Effects of attention on the reliability of individual neurons in monkey  
603 visual cortex proportionally and does not improve the selectivity of single neurons, as measured by the width  
604 of their tuning curve. *Neuron* **23**, 765–773 (1999).

605 [12] Treue, S. & Maunsell, J. H. Attentional modulation of visual motion processing in cortical areas MT and MST.  
606 *Nature* **382**, 539–541 (1996).

607 [13] Lindsay, G. W. & Miller, K. D. How biological attention mechanisms improve task performance in a large-scale  
608 visual system model. *eLife* 1–29 (2018).

609 [14] Pestilli, F., Carrasco, M., Heeger, D. J. & Gardner, J. L. Attentional enhancement via selection and pooling of  
610 early sensory responses in human visual cortex. *Neuron* **72**, 832–846 (2011).

611 [15] Zénon, A. & Krauzlis, R. J. Attention deficits without cortical neuronal deficits. *Nature* **489**, 434–437 (2012).

612 [16] Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent  
613 dynamics in prefrontal cortex. *Nature* **503** (2013). URL <https://www.nature.com/articles/nature12742.pdf>.

614 [17] Ruff, D. A. & Cohen, M. R. Stimulus dependence of correlated variability across cortical areas. *Journal of  
615 Neuroscience* **36**, 7546–7556 (2016).

616 [18] Ruff, D. A. & Cohen, M. R. Attention increases spike count correlations between visual cortical areas. *The  
617 Journal of neuroscience : the official journal of the Society for Neuroscience* **36**, 7523–34 (2016).

618 [19] Felleman, D. J. & Van Essen, D. C. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral  
619 Cortex* **1**, 1–47 (1991).

620 [20] Rabinowitz, N. C., Goris, R. L., Cohen, M. R. & Simoncelli, E. P. Attention stabilizes the shared gain of V4  
621 populations. *eLife* 1–24 (2015).

622 [21] Bondy, A. G., Haefner, R. M. & Cumming, B. G. Feedback determines the structure of correlated variability  
623 in primary visual cortex. *Nature Neuroscience* **21**, 598–606 (2018).

624 [22] Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nature  
625 Reviews Neuroscience* **7**, 358–366 (2006). URL <http://www.nature.com/articles/nrn1888>.

626 [23] Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nature neuroscience* **14**, 811–819  
627 (2011). URL <http://dx.doi.org/10.1038/nn.2842>.

629 [24] Huang, C. *et al.* Circuit models of low-dimensional shared variability in cortical networks. *Neuron* **101**, 1–12  
630 (2019).

631 [25] Macke, J. H., Buesing, L. & Sahani, M. Estimating State and Parameters in State Space Models of Spike  
632 Trains. In *Advanced State Space Methods for Neural and Clinical Data*, 137–159 (Cambridge University Press,  
633 2015).

634 [26] Engel, T. A. *et al.* Selective modulation of cortical state during spatial attention. *Science* **354** (2016).

635 [27] Cohen, M. R. & Maunsell, J. H. Attention improves performance primarily by reducing interneuronal corre-  
636 lations. *Nature Neuroscience* **12**, 1594 (2009).

637 [28] Hair, J. F., Black, W. C., Babin, B. J. & Anderson, R. E. *Multivariate Data Analysis* (Pearson Education  
638 Limited, Essex, 2014), 7th edn.

639 [29] Kanitscheider, I., Coen-Cagli, R. & Pouget, A. Origin of information-limiting noise correlations. *Proceedings of  
640 the National Academy of Sciences of the United States of America* **112**, E6973–82 (2015).

641 [30] Shadlen, M. N., Britten, K. H., Newsome, W. T. & Movshon, J. A. A computational analysis of the relationship  
642 between neuronal and behavioral responses to visual motion. *Journal of Neuroscience* **76**, 1486–1510 (1996).

643 [31] Izhikevich, E. M. Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral  
644 Cortex* **17**, 2443–2452 (2007).

645 [32] Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D. & Brea, J. Eligibility traces and plasticity on behavioral  
646 time scales: experimental support of NeoHebbian Three-Factor Learning Rules. *Frontiers in Neural Circuits*  
647 **12**, 1–16 (2018). 1801.05219.

648 [33] Moreno-Bote, R. *et al.* Information-limiting correlations. *Nature Neuroscience* **17**, 1410–1417 (2014). arXiv:  
649 1411.3159v1.

650 [34] Nienborg, H. & Cumming, B. G. Decision-related activity in sensory neurons may depend on the columnar  
651 architecture of cerebral cortex. *Journal of Neuroscience* **34**, 3579–3585 (2014).

652 [35] Haefner, R. M., Gerwinn, S., Macke, J. H. & Bethge, M. Inferring decoding strategies from choice probabilities  
653 in the presence of correlated variability. *Nature neuroscience* **16**, 235–242 (2013).

654 [36] Born, R. T. & Bradley, D. C. Structure and function of visual area MT. *Annu Rev Neurosci.* 157–89 (2005).

655 [37] LeCun, Y. & Cortes, C. MNIST handwritten digit database. <http://yann. lecun. com/exdb/mnist/> (2010).  
656 URL <http://yann. lecun. com/exdb/mnist/>.

657 [38] Maunsell, J. H. R. & Van Essen, D. C. The connections of the middle temporal visual area (MT) and their  
658 relationship to a cortical hierarchy in the macaque monkey. *Journal of Neuroscience* **3**, 2563–2586 (1983). URL  
659 <http://www. jneurosci. org/content/jneuro/3/12/2563. full. pdf>.

660 [39] Movshon, J. A., Adelson, E. H., Gizzi, M. S. & Newsome, W. T. The analysis of moving visual patterns. In  
661 Chagas, C., Gattass, R. & Gross, C. (eds.) *Experimental Brain Research Supplementum II: Pattern Recognition  
662 Mechanisms*, 117–151 (Springer-Verlag, New York, 1986).

663 [40] Simoncelli, E. P. & Heeger, D. J. A model of neuronal responses in visual area MT. *Vision Research* **38**,  
664 743–761 (1998).

665 [41] Goris, R. L., Movshon, J. A. & Simoncelli, E. P. Partitioning neuronal variability. *Nature Neuroscience* **17**,  
666 858–865 (2014).

667 [42] Osborne, L. C., Lisberger, S. G. & Bialek, W. A sensory source for motor variation. *Nature* **7057** (2005).

668 [43] Rumyantsev, O. I. *et al.* Fundamental bounds on the fidelity of sensory cortical coding. *Nature* **580**, 100–105  
669 (2020). URL <http://dx. doi. org/10. 1038/s41586-020-2130-2>.

670 [44] Pitkow, X., Liu, S., Angelaki, D. E., DeAngelis, G. C. & Pouget, A. How Can Single Sensory Neurons  
671 Predict Behavior? *Neuron* **87**, 411–423 (2015). URL <http://linkinghub. elsevier. com/retrieve/pii/S0896627315005966>. 15334406.

672 [45] Kanitscheider, I., Coen-Cagli, R., Kohn, A. & Pouget, A. Measuring Fisher Information Accurately in Correlated  
673 Neural Populations. *PLoS Computational Biology* **11**, 1–27 (2015).

674 [46] Mitchell, J. F., Sundberg, K. A. & Reynolds, J. H. Spatial attention decorrelates intrinsic activity fluctuations  
675 in macaque area V4. *Neuron* **63**, 879–888 (2009).

676

677 [47] Ruff, D. A. & Cohen, M. R. Attention can either increase or decrease spike count correlations in visual cortex.  
678 *Nature neuroscience* **17**, 1591–7 (2014).

679 [48] Maunsell, J. H. & Cook, E. P. The role of attention in visual processing. *Philos Trans R Soc Lond B Biol Sci.*  
680 **357**, 1063–72 (2002).

681 [49] Treue, S. & Martínez Trujillo, J. C. Feature-based attention influences motion processing gain in macaque  
682 visual cortex. *Nature* **399**, 575–579 (1999).

683 [50] Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J. & Cohen, M. R. Learning and attention reveal a general  
684 relationship between neuronal variability and perception. *Science* **465**, 1–28 (2017).

685 [51] Hénaff, O. J., Boundy-Singer, Z. M., Meding, K., Ziembka, C. M. & Goris, R. L. Representation of visual  
686 uncertainty through neural gain variability. *Nature Communications* **11** (2020).

687 [52] Festa, D., Aschner, A., Davila, A., Kohn, A. & Coen-Cagli, R. Neuronal variability reflects probabilistic  
688 inference tuned to natural image statistics. *bioRxiv* (2020).

689 [53] Denfield, G. H., Ecker, A. S., Shinn, T. J., Bethge, M. & Tolias, A. S. Attentional fluctuations induce shared  
690 variability in macaque primary visual cortex. *Nature Communications* **9** (2018).

691 [54] Engel, T. A., Chaisangmongkon, W., Freedman, D. J. & Wang, X.-J. Choice-correlated activity fluctuations  
692 underlie learning of neuronal category representation. *Nature communications* **6**, 6454 (2015). URL <http://www.nature.com/doifinder/10.1038/ncomms7454>.

693 [55] Singer, W. Neuronal synchrony: A versatile code review for the definition of relations? *Neuron* **24**, 49–65  
694 (1999).

695 [56] Akam, T. E. & Kullmann, D. M. Efficient "communication through coherence" requires oscillations structured  
696 to minimize interference between signals. *PLoS Computational Biology* **8** (2012).

697 [57] Akam, T. & Kullmann, D. M. Oscillatory multiplexing of population codes for selective communication in the  
698 mammalian brain. *Nature Reviews Neuroscience* **15**, 111–122 (2014). 1309.2848v1.

699 [58] Marschall, O., Cho, K. & Savin, C. A unified framework of online learning algorithms for training recurrent  
700 neural networks. *The Journal of Machine Learning Research* **21**, 5320–5353 (2020).

701 [59] Haefner, R. M., Berkes, P. & Fiser, J. Perceptual Decision-Making as Probabilistic Inference by Neural Sampling.  
702 *Neuron* **90**, 649–660 (2016).

703 [60] Koren, V., Andrei, A. R., Hu, M., Dragoi, V. & Obermayer, K. Pairwise Synchrony and Correlations Depend  
704 on the Structure of the Population Code in Visual Cortex. *Cell Reports* **33**, 108367 (2020). URL <https://doi.org/10.1016/j.celrep.2020.108367>.

705 [61] Sampathkumar, V., Miller-Hansen, A., Sherman, S. M. & Kasthuri, N. Integration of signals from different  
706 cortical areas in higher order thalamic neurons. *Proceedings of the National Academy of Sciences* **118** (2021).

707 [62] Purushothaman, G., Marion, R., Li, K. & Casagrande, V. A. Gating and control of primary visual cortex by  
708 pulvinar. *Nature neuroscience* **15**, 905–912 (2012).

709 [63] Andersen, R. A., Musallam, S. & Pesaran, B. Selecting the signals for a brain-machine interface. *Current*  
710 *opinion in neurobiology* **14**, 720–726 (2004).

711 [64] Yu, B. M. *et al.* Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Popu-  
712 lation Activity. *Journal of Neurophysiology* **102**, 614–635 (2009).

713 [65] Benjamin, A. S. *et al.* Modern machine learning outperforms GLMs at predicting spikes. *bioRxiv* (2017).