Noise correlations for faster and more robust learning

Matthew R. Nassar[1,2], Apoorva Bhandari[1,3]

1. Robert J. & Nancy D. Carney Institute for Brain Science, Brown University, Providence RI 02912-1821, USA
2. Department of Neuroscience, Brown University, Providence RI 02912-1821, USA
3. Department of Cognitive, Linguistic, and Psychological Sciences, Providence RI, 02912-1821

Acknowledgements:

We would like to thank Michael Frank, Drew Linsley, Chris Moore and Jan Drugowitsch for helpful discussion. This work was funded by NIH grants F32MH102009 and K99AG054732 (MRN), NINDS R21NS108380 (AB). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Competing interests:

The authors have no financial or non-financial conflicts of interest related to this work.

Abstract:

Distributed population codes are ubiquitous in the brain and pose a challenge to downstream neurons that must learn an appropriate readout. Here we explore the possibility that this learning problem is simplified through inductive biases implemented through stimulus-independent noise correlations that constrain learning to task-relevant dimensions. We test this idea with a neural network model of a perceptual discrimination task in which the correlation among similarly tuned units can be manipulated independently of overall population signal-to-noise ratio. Higher noise correlations among similarly tuned units led to faster and more robust learning favoring homogenous weights assigned to neurons within a functionally similar pool and could emerge naturally with Hebbian learning. When multiple discriminations were learned simultaneously, noise correlations across relevant feature dimensions sped learning whereas those across irrelevant feature dimensions slowed it. These results suggest that noise correlations may serve constrain learning to appropriate dimensions to optimize readout learning.

Introduction:

The brain represents information using distributed population codes in which particular feature values are encoded by large numbers of neurons. One advantage of such codes is that a pooled readout across many neurons can effectively reduce the impact of stimulus-independent variability (noise) in the firing of individual neurons (Averbeck, Latham, & Pouget, 2006; Moreno-Bote et al., 2014; Pouget, Dayan, & Zemel, 2000). However, the extent to which this benefit can be employed in practice is constrained by noise correlations, or the degree to which stimulus-independent variability is shared across neurons in the population (Averbeck et al., 2006; Beck, Ma, Pitkow, Latham, & Pouget, 2012; Kanitscheider, Coen-Cagli, & Pouget, 2015). In particular, positive noise correlations between neurons that share the same stimulus tuning can reduce the amount of decodable information in the neural population (Averbeck et al, 2006; Moreno-Bote et al., 2014; Hu et al., 2014). Despite their detrimental effect on encoding, noise correlations of this type are reliably observed (Cohen & Kohn, 2011; Law & Gold, 2009). Indeed, noise correlations between neurons are dynamically enhanced under conditions where those neurons provide evidence the same response in a perceptual categorization task (Cohen & Newsome, 2008) raising the question of whether they serve a function.

At the same time, learning to effectively readout a distributed code also poses a significant challenge. Learning the appropriate weights for potentially tens of thousands of neurons in a low signal-to noise regime is a difficult, high-dimensional problem, requiring a very large number of learning trials and entailing considerable risk of “over fitting”. Nonetheless, people and animals can often rapidly learn to perform perceptual discrimination tasks, albeit with performance that does not approach theoretically achievable levels (Cohen & Newsome, 2008; Hawkey, Amitay, & Moore, 2004; Stringer, Michaelos, & Pachitariu, 2019). In comparison, deep neural networks capable of achieving human level performance typically require a far greater number of learning trials than would be required by humans and other animals (Cohen & Newsome, 2008; Tsividis, Pouncy, Xu, Tenenbaum, & Gershman, 2017). This raises the question of how brains implement inductive biases that enable efficient learning in high dimensional spaces.

One possible answer might be that noise correlations serve the purpose of reducing the effective dimensionality of learning problems. For example, perceptual stimuli often contain a large number of features that may be irrelevant to a given categorization. At the level of a neural population, individual neurons may differ in the degree to which they encode task irrelevant information, thus making the learning problem more difficult. In principle, noise correlations in the relevant dimension could reduce the effects of this variability on learned readout. Such an explanation would be consistent with computational analyses of Hebbian

learning rules (Adibi, McDonald, Clifford, & Arabzadeh, 2013; Averbeck & Lee, 2003; Bair, Zohary, & Newsome, 2001; Cohen & Maunsell, 2009; Ecker et al., 2010; Gu et al., 2011; Huang & Lisberger, 2009; Maynard et al., 1999; Oja, 1982; Zohary, Shadlen, & Newsome, 1994), which can both facilitate faster and more robust learning (Cohen & Newsome, 2008; Krotov & Hopfield, 2019), and in turn may induce noise correlations. We propose that faster learning of an approximate readout is made possible through low dimensional representations that share both signal and noise across a large neural population. In particular, we hypothesize that representations characterized by enhanced noise correlations among similarly tuned neurons can improve learning by focusing adjustments of the readout onto task relevant dimensions.

We explore this possibility using neural network models of a two alternative forced choice perceptual discrimination task in which the correlation among similarly tuned neurons can be manipulated independently of the overall population signal-to-noise ratio. Within this framework, noise correlations, which can be learned through Hebbian mechanisms, speed learning by forcing learned weights to be similar across pools of similarly tuned neurons, thereby ensuring learning occurs over the most task relevant dimension. We extend our framework to a cued multidimensional discrimination task and show that dynamic noise correlations similar to those observed in vivo, speed learning by constraining weight updates to the relevant feature space. These results demonstrate that when information is extrinsically limited, noise correlations can make learning faster and more robust by controlling the dimensions over which learning occurs.

Results:

We examine how noise correlations affect learning in a simplified neural network where the appropriate readout of hundreds of weakly tuned units is learned over time through reinforcement. In order to isolate the effects of noise correlations on learning, rather than their effects on other factors such as representational capacity, we consider population encoding schemes at the input layer that can be constrained to a fixed signal-to-noise ratio. This assumption differs from previous work on noise correlations where the *variance* of the neural population is assumed to be fixed and covariance is changed to produce noise correlations, thereby affecting the representational capacity of the population (Figure 1A; (Averbeck et al., 2006; Beck et al., 2012; Kanitscheider et al., 2015; Moreno-Bote et al., 2014)). Under our assumptions, a fixed signal-to-noise ratio can be achieved for any level of covariance by scaling the variance (Figure 1B; equations 1-3), or, alternately scaling the magnitude of the signal (not shown). While we do not discount the degree to which noise correlations affect the encoding potential of neural populations, we believe that in many cases the

relevant information is limited by extrinsic factors (eg. the stimulus itself, or upstream neural populations providing input (Averbeck et al., 2006; Beck et al., 2012; Kanitscheider et al., 2015; Moreno-Bote et al., 2014). In such conditions, it is not possible to maximize encoding potential as this would be tantamount to the population “creating new information”. Our framework can be thought of as testing how best to format limited available information in a neural population in order to ensure that an acceptable readout can be rapidly and robustly learned.

We propose that within this framework, noise correlations of the form that have previously been shown to limit encoding are beneficial because they constrain learning to occur over the most relevant dimensions. In general, a linear readout can be thought of as hyperplane serving as a classification boundary in an N dimensional space, where N reflects the number of neurons in a population. Learning in such a framework involves useful adjustments of the hyperplane in the dimension that best discriminates signal from noise (central arrows in figure 1C&D) but also in adjustments in dimensions orthogonal to the relevant one, leading to “twisting” of the hyperplane (curved arrows in figure 1C&D) that could potentially impair performance. Our motivating hypothesis is that by focusing population activity into the task relevant dimension, noise correlations can increase the increase the fraction of hyperplane adjustments that occur in the task relevant dimension (Fig 1D).

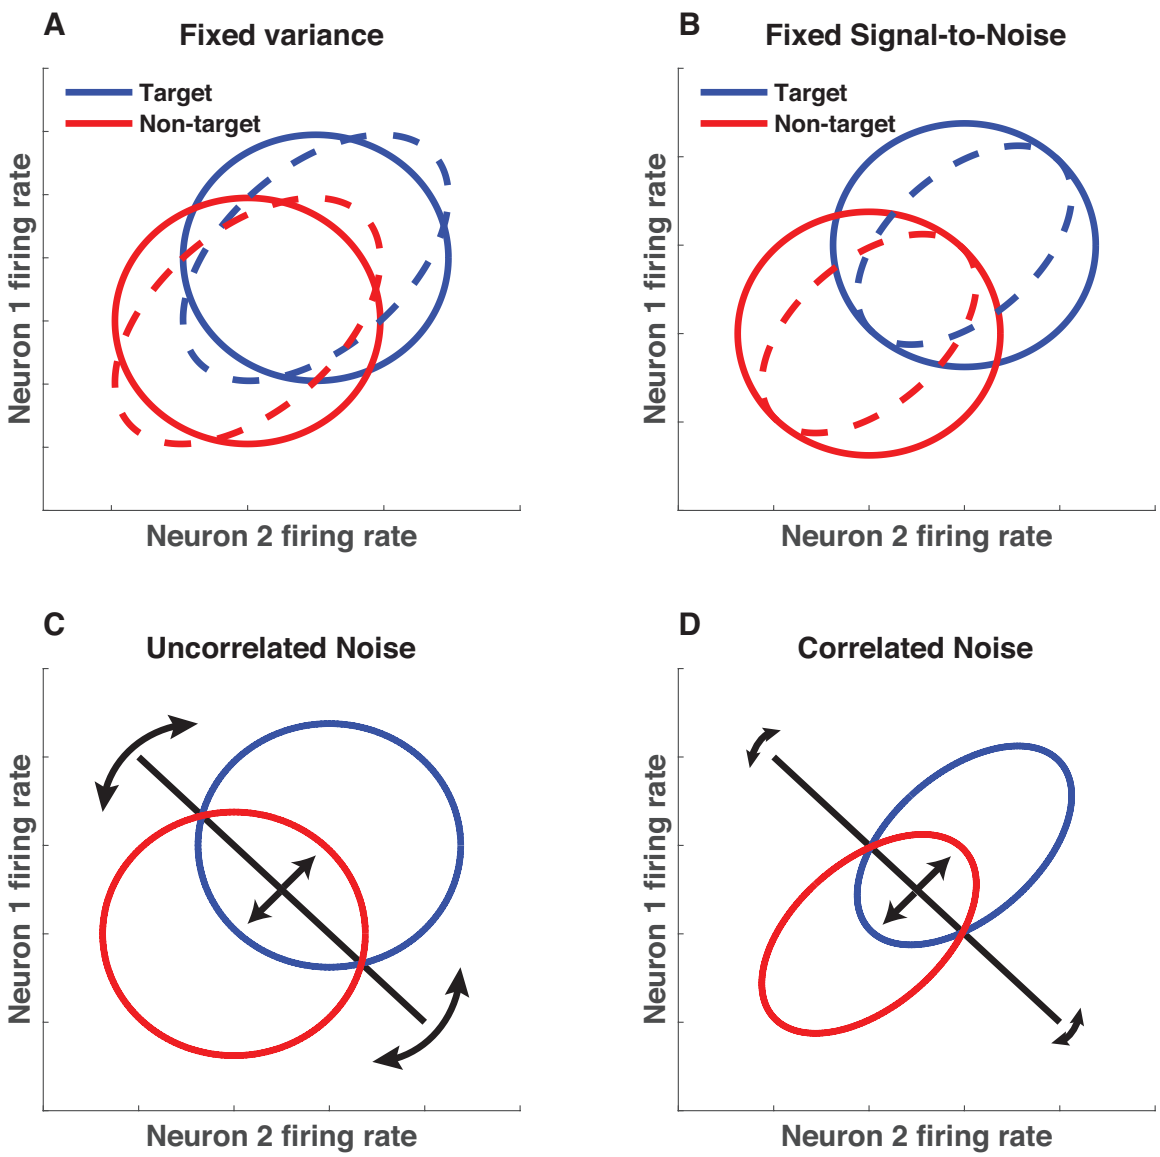

Figure 1: Modeling noise correlations in under extrinsic constraint on signal-to-noise ratio. A) Previous work has modeled noise correlations by assuming that population variance is fixed and that covariance is manipulated to produce noise correlations. Under such assumptions, the firing rate of two similarly tuned neurons is plotted in the absence (solid) or presence (dotted) of information-limiting noise correlations. **B)** Here we assume that the signal-to-noise ratio of the neural population is limited to a fixed value such that noise correlations between similarly tuned neurons do not affect theoretical performance. Thus, the percent overlap of blue (target) and red (non-target) activity profiles does not differ in the presence (dotted) or absence (solid) of noise correlations. **C&D)** Under this assumption, noise correlations among similarly tuned neurons could compress the population activity to a plane orthogonal to the optimal decision boundary, thereby minimizing boundary adjustments in irrelevant dimensions (**C**) and maximizing boundary adjustments on relevant ones (**D**).

In order to test this hypothesis, we constructed a fully connected two-layer feedforward neural network in which input layer units responded to one of two

stimulus categories (pool 1 & pool 2) and each output unit produced a response consistent with a category perception (orange/blue units in Fig 2A). On each trial, the network was presented with one stimulus at random, and input unit firing for each pool was drawn from a multivariate Gaussian in which we manipulated the covariance while constraining the population signal-to-noise ratio. Output units were activated according to a weighted average of inputs and a response was selected according to output unit activations. On each trial, weights to the selected action were adjusted according to a reinforcement learning rule that strengthened connections that facilitated a rewarded action and weakened connections that facilitated an unrewarded action (Cohen & Kohn, 2011; Law & Gold, 2009).

Noise correlations led to faster and more robust learning of the appropriate stimulus-response mapping. All neural networks learned to perform the requisite discrimination, but neural networks that employed correlations among similarly tuned neurons learned more rapidly (Fig 2B). After learning, networks that employed such noise correlations assigned more homogenous weights to input units of a given pool than did networks that lacked noise correlations (compare Fig 2C&D). This led to better trained-task performance (Fig 2E; Pearson correlation between noise correlations and test performance: $R = 0.29$, $p < 10e-50$) and greater robustness to adversarial noise profiles (Fig 2F; $R = 0.81$, $p < 10e-50$) in the networks that employed noise correlations. Critically, these learning advantages emerged despite the fact that optimal readout of all networks achieved similar levels of performance and robustness (Fig 2E&F, red).

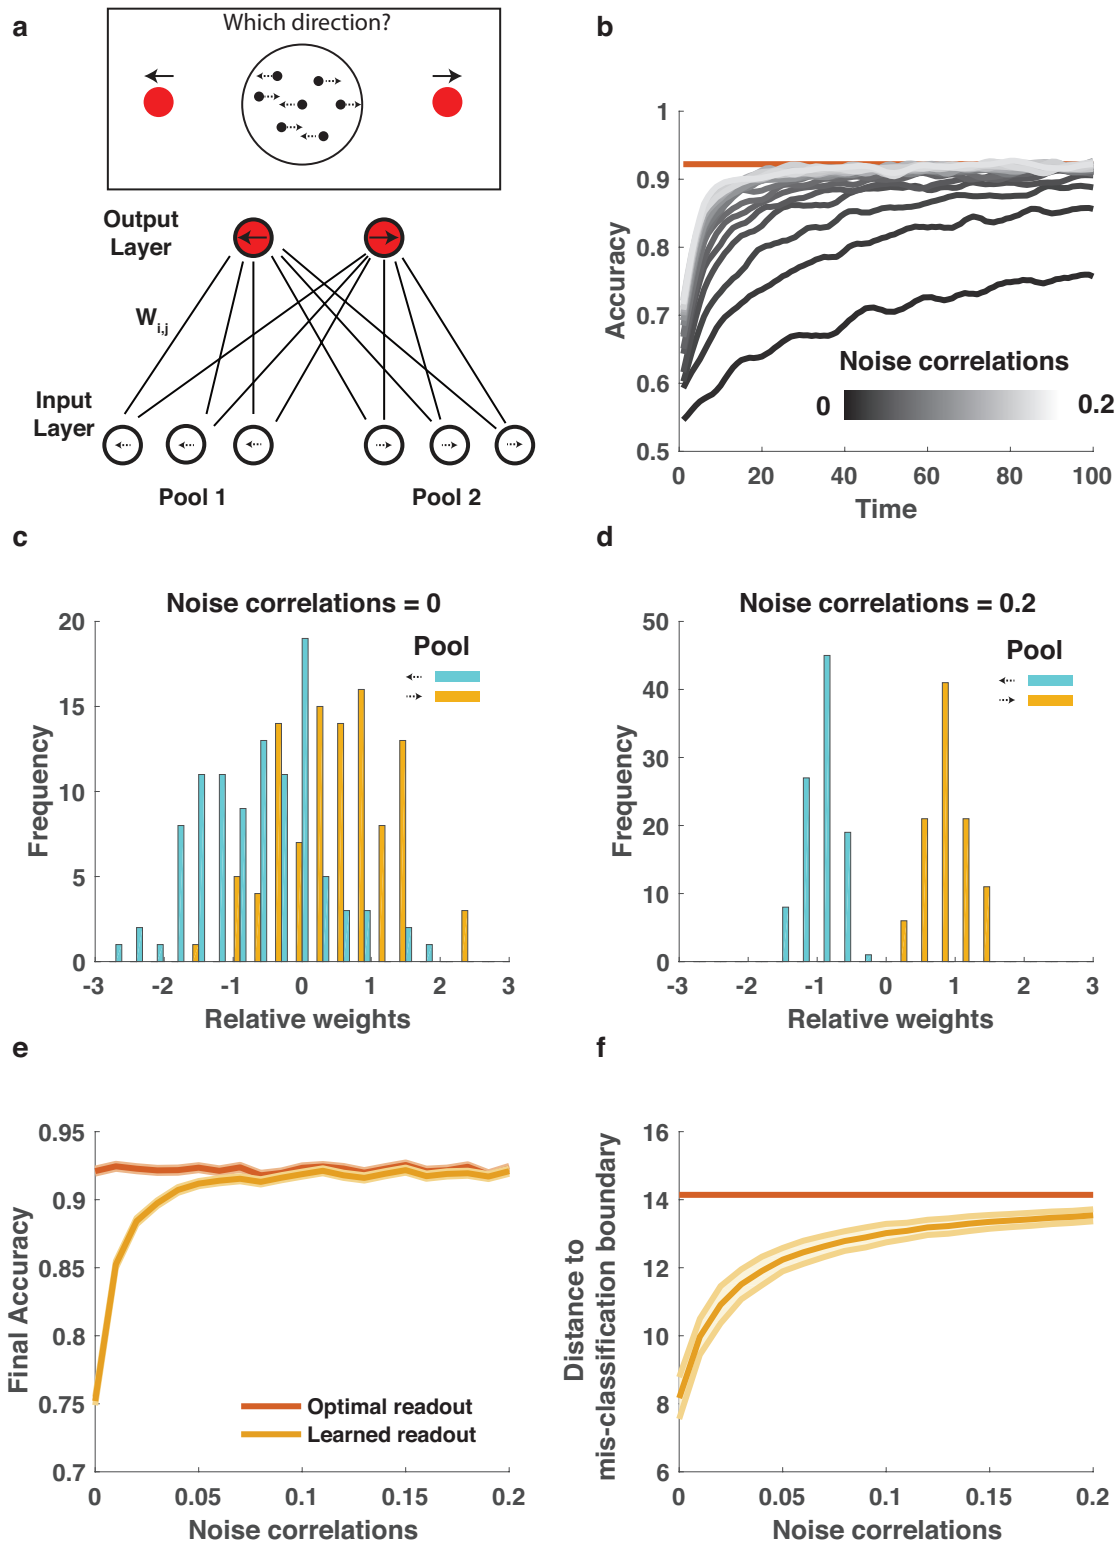

Figure 2: Correlated noise within similarly tuned populations leads to faster and more robust learning of a perceptual discrimination. A) A two-layer feed-forward neural network

was designed to solve a two alternative forced choice motion discrimination task at or near perceptual threshold. Input layer contains two pools of neurons that provide evidence for alternate percepts (eg. leftward motion versus rightward motion) and output neurons encode alternate courses of actions (eg. saccade left versus saccade right). Layers are fully connected with weights randomized to small values and adjusted after each trial according to rewards (see methods). **B)** Average learning curves for neural network models in which population signal-to-noise ratio in pools 1 and 2 were fixed, but noise correlations (grayscale) were allowed to vary from small (dark) to large (light) values. **C&D)** Weight differences (Orange output – Blue output) for each input unit (color coded according to pool) after 100 timesteps of learning for low (**C**) and high (**D**) noise correlations. **E**) Accuracy in the last 20 training trials is plotted as a function of noise correlations for learned readouts (orange) and optimal readout (red). Lines/shading reflect Mean/SEM. F) The shortest distance, in terms of neural activation, required to take the mean input for a given category (eg. blue or orange) to the boundary that would result in misclassification is plotted for the final learned (orange) and optimal (red) weights for each noise correlation condition (abscissa). Lines/shading reflect Mean/SEM.

Given that noise correlations implemented in our previous simulation, like those observed in the brain, depended on the tuning of individual units, we tested whether such noise correlations might be produced via Hebbian plasticity. Specifically, we considered an extension of our neural network in which an additional intermediate layer is included between input and output neurons (Fig 3a). Input units were again divided into two pools that differed in their encoding, but variability was uncorrelated across neurons within a given pool. Connections between the input layer and intermediate layer were initialized such that each input unit strongly activated one intermediate layer unit, and shaped over time using a Hebbian learning rule that strengthened connections between co-activated neuron pairs. Despite the lack of noise correlations in the input layer of this network (Fig 3b; mean[std] in pool residual correlation = 0.0015[0.10]), neurons in the intermediate layer developed tuning-specific noise correlations of the form that were beneficial for learning in the previous simulations (Fig 3c; mean[std] in pool residual correlation = 0.55[0.07]; *t*-test on difference from input layer correlations $t = 443$, $dof = 19800$, $p < 10e\text{-}50$).

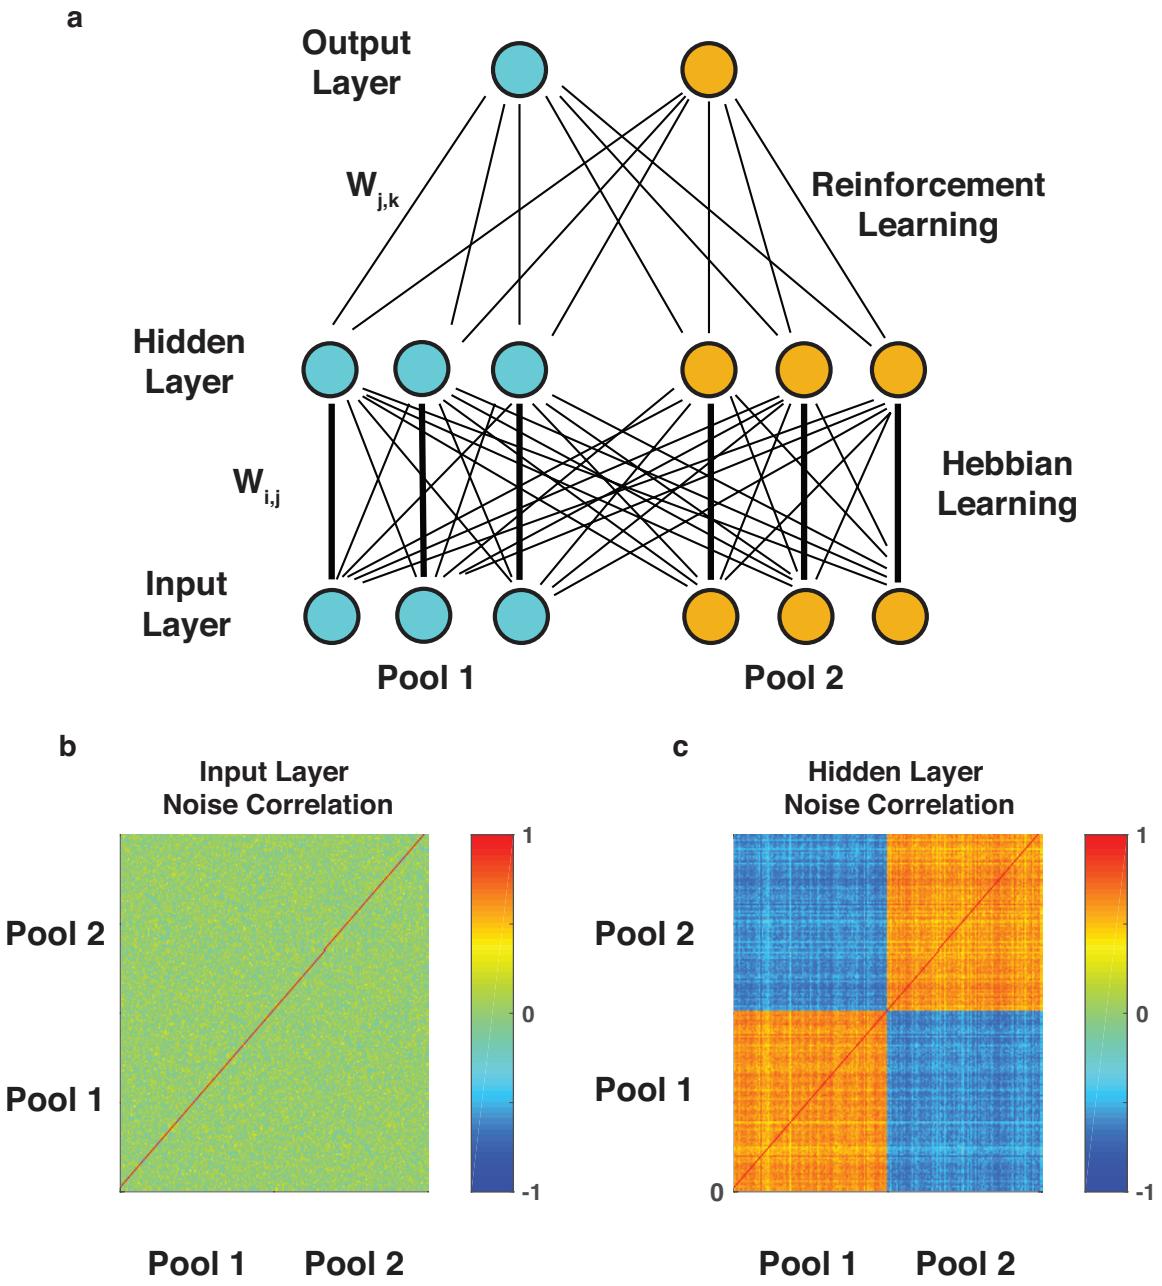

Figure 3: Hebbian learning produces correlations within similarly tuned populations in a perceptual discrimination task. A) Three-layer neural network architecture. Input layer feeds forward to hidden layer, which is fully connected to an output layer. Input layer provides uncorrelated inputs to hidden layer through projection weights that are adjusted according to a Hebbian learning rule. **B&C**) Noise correlations observed in hidden layer units at the beginning (**B**) and end (**C**) of training.

In order to understand how noise correlations might impact learning in mixed encoding populations, we extended our perceptual discrimination task to include two directions of motion discrimination (eg. up/down and left/right). On each trial, a cue indicated which of two possible motion discriminations should be performed (Fig 4a, left; (Cohen & Newsome, 2008; Kanitscheider et al., 2015; Kohn, Coen-Cagli, Kanitscheider, & Pouget, 2016)). We extended our neural network to include four populations of one hundred input units, each population encoding a conjunction of motion directions (up-right, up-left, down-right, down-left; Fig 4a; input layer). Two additional inputs provided a perfectly reliable "cue" regarding the relevant feature for the trial (Fig 4a; task units). Four output neurons encoded the four possible responses (up, left, down, right) and were fully connected to the input layer (Fig 4a; output layer). Task units were hard wired to eliminate irrelevant task responses, but weights of input units were learned over time as in our previous simulations.

Learning performance in the two-feature discrimination task depended not only on the level of noise correlations, but also on the type. As in the previous simulation, adding noise correlations to each individual population of identically tuned units led to faster learning of the appropriate readout (Figure 4B&C, compare blue and yellow; Figure 4D&E, vertical axis; mean[std] accuracy across training: 0.53[0.05] and 0.614[0.08] for minimum (0) and maximum (0.2) in pool correlations, t-test for difference in accuracy: $t = 95$, $dof = 19998$, $p < 10e\text{-}50$).

However, the more complex task design also allowed us to test whether dynamic trial-to-trial correlations might further facilitate learning. Specifically, correlations that increase shared variability among units that contribute evidence to the same response have been observed previously (Cohen & Newsome, 2008; Gu et al., 2011; Ni, Ruff, Alberts, Symmonds, & Cohen, 2018), and could in principle focus learning on relevant dimensions (Fig 1c&d) even when those dimensions change from trial to trial. Indeed, adding correlations among separate pools that share the same encoding of the relevant feature (eg. UP on a vertical trial) led to faster learning (Fig 4b; mean[std] training accuracy for model with relevant pool correlations: 0.64[0.09], *t*-test for difference from in pool correlation only model: $t = 22$, $dof = 19998$, $p < 10e\text{-}50$) and weights that more closely approached the optimal readout (Fig 4e, horizontal axis). In contrast, when positive noise correlations were introduced across separate encoding pools that shared the same tuning for the irrelevant dimension on each trial (eg. UP on a horizontal trial) learning was impaired dramatically (Fig 4c; mean[std] training accuracy for model with irrelevant pool correlations: 0.51[0.05], *t*-test for difference from in pool correlation only model: $t = -112$, $dof = 19998$, $p < 10e\text{-}50$) and learned weights diverged from the optimal readout (Fig 4f, horizontal axis). Model performance differences were completely attributable to learning the readout, as all models performed similarly when using the optimal readout (Fig4S1).

In order to test the idea that noise correlations might focus learning onto relevant dimensions, we extracted weight updates from each trial and projected these updates into a two-dimensional space where the first dimension captured the relative sensitivity to leftward versus rightward motion and the second dimension captured relative sensitivity to upward versus downward motion. In the model where input units were only correlated within their identically tuned pool, weight updates projected in all directions more or less uniformly (Fig 4g), and did not differ systematically across trial types (vertical versus horizontal). However, dynamic noise correlations that shared variability across the relevant dimension tended to push weight updates onto the dimension that was most relevant for a given trial (Fig 4f; *t*-test for difference in the magnitude of updating in up/down and left/right dimensions across conditions [up/down – left/right]: $t = 3.4$, $dof=98$, $p = 0.001$). In contrast, dynamic noise correlations that shared variability across the irrelevant dimension tended to push weight updates onto the wrong dimension (Fig 4h; t-test for difference in the magnitude of updating in up/down and left/right dimensions across conditions [up/down – left/right]: $t = -9.5$, $dof=98$, $p = 10e-14$). Both of these trends were consistent across simulations, providing an explanation for the performance improvements achieved by relevant noise correlations (projection of learning onto appropriate dimension) and performance impairments produced by irrelevant noise correlations (projection of learning onto inappropriate dimension).

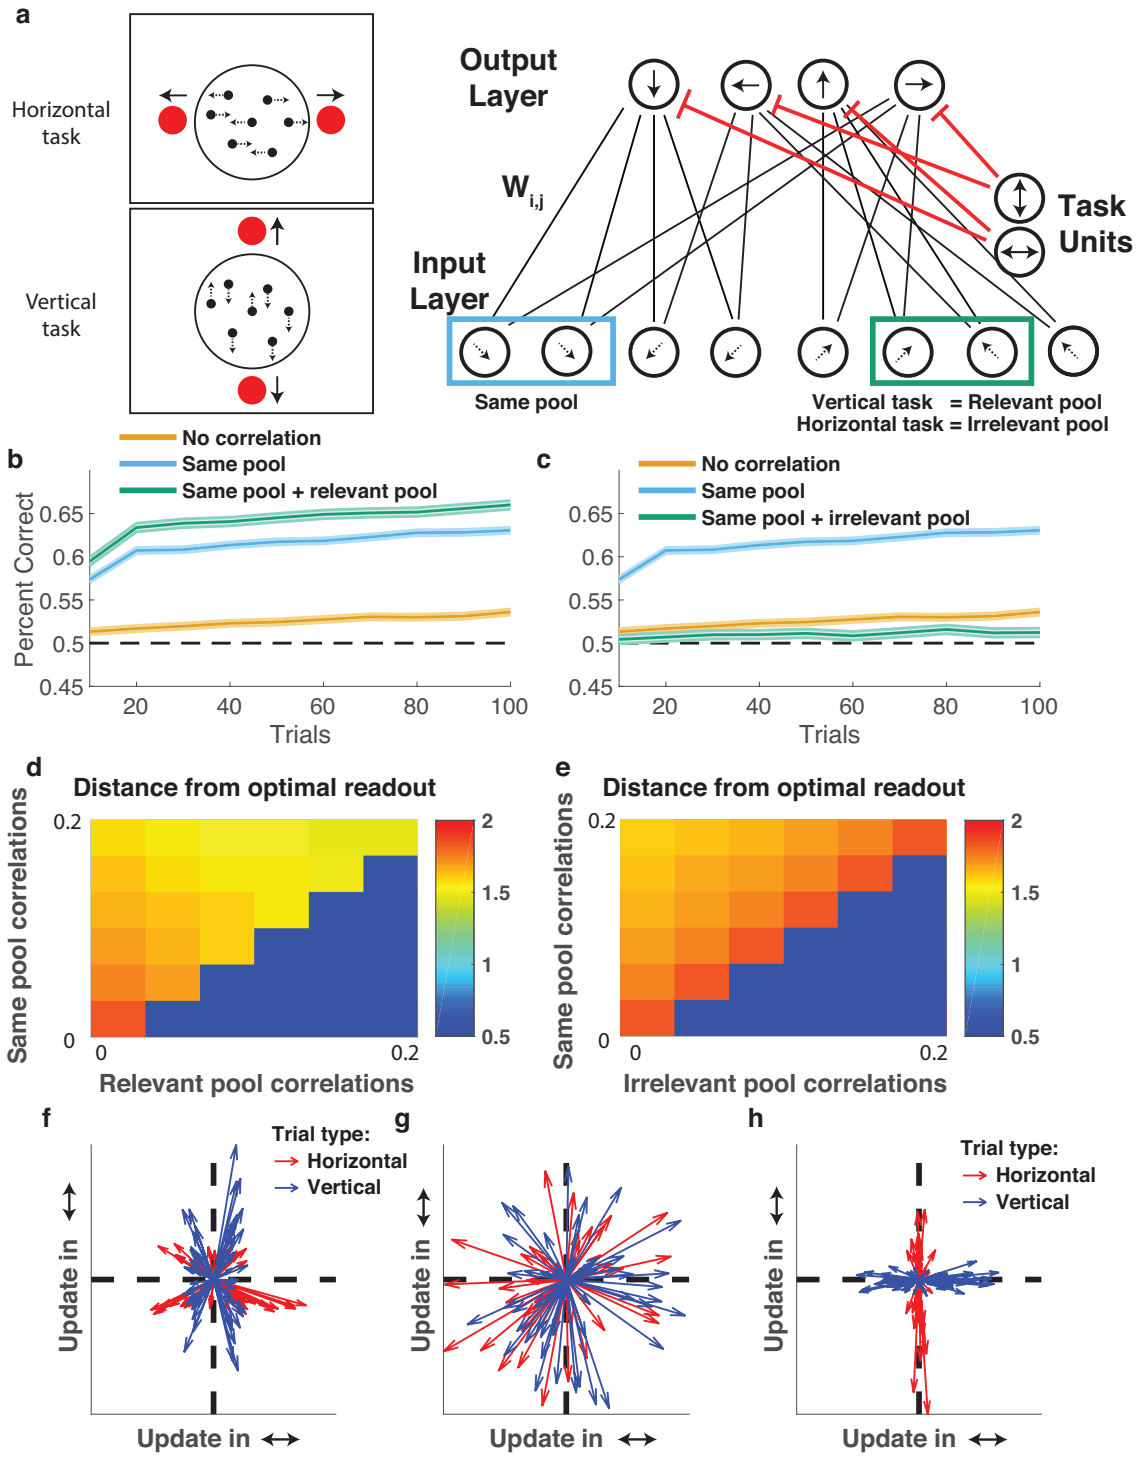

Figure 4: Task dependent noise correlations affect learning speed by projecting learning onto specific feature dimensions. A) A neural network was trained to perform two interleaved motion discrimination tasks (left; (Cohen & Newsome, 2008)). Network schematic (right) depicts two-layer feedforward network in which each population of input units represents two dimensions of motion (up versus down, and left versus right), and output units produce responses in favor of alternative actions (up, down, left, right). Two additional input units provide cue information that

biases output units to produce an output corresponding to the discrimination appropriate on this trial (eg. color or texture). Noise correlations were manipulated among 1) identically tuned neurons (blue rectangle; same pool), 2) neurons that have similar encoding of the task relevant feature (green rectangle pair in vertical trials; relevant pool), and 3) neurons that have similar encoding of the task irrelevant feature (green rectangle pair in horizontal trials; irrelevant pool). **B&C**) Learning curves showing accuracy (ordinate) over trials (abscissa) for models 1) lacking noise correlations (orange), 2) containing noise correlations that are limited to neurons that have same tuning for both features (same pool; blue), 3) containing same pool noise correlations along with correlations between neurons in different pools that have the same tuning for the task-relevant feature (in pool+rel pool; green in **B**), and 4) containing in-pool noise correlations along with correlations between neurons in different pools that have the same tuning for the task irrelevant feature (in pool+irrel pool; green in **C**). **D&E**) Distance between learned weights and the optimal readout (color) for models that differ in their level of "in pool" correlations (ordinate, both plots), "relevant pool" correlations (abscissa, **D**), and "irrelevant pool" correlations (abscissa, **E**). **F,G,H**) Weight updates for example learning sessions were projected into a two dimensional space in which net updates to the relative contribution of color units (eg. blue versus orange) is represented on the abscissa and updates to the relative contribution of texture unites (eg. striped versus unstriped) is represented on the ordinate. Arrows reflect single trial weight updates and are colored according to the trial type (red = color discrimination, blue = texture discrimination). Weight updates for a model with only "in pool" correlations look similar across trial types (**G**), but weight updates for a model with "relevant pool" correlations indicate more weight updating on the relevant feature (**F**), whereas the opposite was observed in the case of "irrelevant pool" correlations (**H**).

Discussion:

Taken together, our results suggest that in settings where the population signal-to-noise ratio is limited by external factors and relevant task representations are low dimensional, noise correlations can make learning faster and more robust by focusing learning on the most relevant dimensions. We demonstrate this basic principle in a simple perceptual learning task (Fig 2), where beneficial noise correlations between similarly tuned units could be produced through a simple Hebbian learning rule (Fig 3). We extended our framework to a contextual learning task to demonstrate that dynamic noise correlations that bind task relevant feature representations facilitate faster learning (Fig 4b&d) by pushing learning onto task-relevant dimensions (Fig 4f). Given the pervasiveness of noise correlations among similarly tuned sensory neurons (Adibi et al., 2013; Averbeck & Lee, 2003; Bair et al., 2001; Cohen & Maunsell, 2009; Ecker et al., 2010; Gu et al., 2011; Huang & Lisberger, 2009; Maynard et al., 1999; Zanto & Gazzaley, 2009; Zohary et al., 1994), and that the noise correlations dynamics beneficial for learning in our simulations are similar to those that have been observed *in vivo* (Akaishi, Kolling, Brown, & Rushworth, 2016; Cohen & Newsome, 2008; Leong, Radulescu, Daniel, DeWoskin, & Niv, 2017), we interpret our results as suggesting that "information limiting" noise correlations are a feature of neural coding architectures that ensures efficient readout learning, rather than a bug that limits encoding potential.

This interpretation rests on several assumptions in our model. Of particular importance is the assumption that signal-to-noise ratio of our populations is fixed, meaning that our manipulation of noise correlations can focus variance on specific dimensions without gaining or losing information. This assumption reflects conditions in which information is limited at the level of the inputs to the population, for instance due to noisy peripheral sensors (Beck et al., 2012; Kanitscheider et al., 2015; Mack, Preston, & Love, 2019). In such conditions, even with optimal encoding, population information saturates at an upper bound determined by the information available in the inputs to the population. Therefore, fixing the signal-to-noise ratio enabled us to examine the effect of noise correlations on downstream processes that learn to read-out the population code in the absence of any influence of noise correlations on the quantity of information contained within that population code.

Previous theoretical work exploring the role of noise correlations in encoding has typically assumed that single neurons have a fixed variance, such that tilting the covariance of neural populations towards or away from the dimension of signal encoding would have a large impact on the amount of information that can be encoded by a population (Fig 1a; (Averbeck et al., 2006; Bondy, Haefner, & Cumming, 2018; Haefner, Pietro Berkes, & Fiser, 2016; Lange, Chattoraj, Beck, Yates, & Haefner, 2018; Moreno-Bote et al., 2014)). Such assumptions lead to the idea that positive noise correlations among similarly tuned neurons limit encoding potential, raising the question of why they are so common in the brain (Cohen & Kohn, 2011; Cohen & Maunsell, 2009; 2011; Doiron, Litwin-Kumar, Rosenbaum, Ocker, & Josić, 2016; Herrero, Gieselmann, Sanayei, & Thiele, 2013; Mitchell, Sundberg, & Reynolds, 2009). In considering the implications of this framework, one important question is: if information encoded in the population can be increased by changing the correlation structure among neurons, where does this additional information come from? In some cases, the neural population in question may indeed receive sufficient task relevant information from upstream brain regions to reorganize its encoding in this way, but in other cases it is likely that information is limited by the inputs to a neural population (Downer, Niwa, & Sutter, 2015; Kanitscheider et al., 2015; Kohn et al., 2016; Ruff & Cohen, 2014). In cases where incoming information is limited, increasing representational capacity is not possible, and formatting information for efficient readout is essentially the best that the population code could do. Here we show that the noise correlations that have previously been described as “information limiting” are exactly the type of correlations that make learning the appropriate readout more efficient under such conditions.

Jointly considering these antagonistic perspectives on noise correlations provides a more nuanced view of how neural representations are likely optimized for learning. In order to optimize an objective function, a neural population can

reduce correlated noise in task relevant dimensions to increase its representational capacity up to a level constrained by its inputs. But once the population is fully representing all task relevant information that has been provided to it, it can additionally optimize representations by pushing as much variance onto task relevant dimensions as possible, thereby affording efficient learning in downstream neural populations. In short, optimization of a neural population code does not occur in a vacuum, and instead depends critically on both upstream (eg. input constraints) and downstream (eg. readout) neural populations. Through this view, if a neural population is *not* fully representing the decision relevant information made available to it, then learning could improve the efficiency of representations by reducing rate limiting noise correlations as has been observed in some paradigms (Gu et al., 2011; Ni et al., 2018). In contrast, once available information is fully represented, readout learning could be further optimized by reformatting population codes such that variability is shared across neurons with similar tuning for the relevant task feature, producing the sorts of dynamic noise correlations that have been observed in well trained animals (Cohen & Newsome, 2008).

In addition to key assumptions about an external limitation on signal-to-noise, our modeling included a number of simplifying assumptions that are unlikely to hold up in real neural populations. For example, we consider discrete pools of identically tuned neurons, rather than the heterogeneous populations observed in sensory cortical regions of the brain. A primary goal of our work was to identify the computational principles that control the speed at which readout can be learned, and our simplified populations are considerably more tractable and transparent than realistic neural populations. The principles that we identify here are certainly at play in real neural populations, albeit with implications that are far less transparent. We hope that our simplified results pave the way for future work to assess nuances that can emerge in mixed heterogeneous populations, or in more realistic architectures that go beyond the simple feed forward flow of information considered here.

Relation to attentional effects on noise correlations

In broad strokes, our finding that manipulation of noise correlations can focus variance on specific dimensions when controlling the degree to which task-related information is gained or lost is in line with specific models of attention. In particular, noise in task irrelevant dimensions might be considered in the same light that is often cast on suppression of task irrelevant dimensions by attentional mechanisms (Devauges & Sara, 1990; Lapiz & Morilak, 2006; Zanto & Gazzaley, 2009), in particular for purposes of accurate credit assignment (Akaishi et al., 2016; Joshi, Li, Kalwani, & Gold, 2016; Leong et al., 2017; Reimer et al., 2016). One possibility is that compressed low-dimensional task representations in higher-order decision regions (Joshi & Gold, n.d.; Mack et al., 2019; Vinck,

Batista-Brito, Knoblich, & Cardin, 2015) may pass accumulated decision related information back to sensory regions in order to approximate Bayesian inference (Bondy et al., 2018; Bouret & Sara, 2005; Haefner et al., 2016; Lange et al., 2018). As task relevant features are learned, such a process would promote noise correlations between neurons coding those relevant features. While noise correlations are measured as variability that cannot be accounted for according to specific task conditions, it is difficult to know whether this variability might relate to other internal or external factors that are not recorded during experimental procedures, and thus whether noise might emerge through a suppression of orthogonal task-unrelated representations.

One observation that seems at odds with this interpretation is that manipulations of attention that cue a particular location or feature tend to decrease noise correlations among neurons that encode that location or feature (Cohen & Maunsell, 2009; 2011; Doiron et al., 2016; Herrero et al., 2013; Kanitscheider et al., 2015; Kohn et al., 2016; Mitchell et al., 2009). The effects of attentional cuing on noise correlations are dynamic in that cues change from one trial to the next, and contextual, in that noise correlations are reduced most dramatically among neurons that contribute evidence toward the same response in a manner consistent with increasing the amount of task relevant information in the population code (Downer et al., 2015; Ruff & Cohen, 2014; Shadlen & Newsome, 1998). Our model does not account for these attentional effects, as we intentionally constrained the signal-to-noise ratio of our neural populations, thereby eliminating any potential changes in information encoding potential. However, we hope that our work motivates future studies to jointly consider the impacts of noise correlations on both learning and immediate performance in order to better understand the potentially competing imperatives that the brain faces in dynamically controlling the correlation structure of its own representations (see (Haimerl, Savin, & Simoncelli, 2019) for one attempt to do so).

Model predictions

Our work shows that noise correlations can focus the gradient of learning onto the most appropriate dimensions. Thus, our model predicts that the degree to which similarly tuned neurons are correlated during a perceptual discrimination should be positively related to performance improvements experienced on subsequent discriminations. In contrast, our model predicts that the degree of correlation between neurons that are similarly tuned to a task irrelevant feature should control the degree of learning on irrelevant dimensions, and thus negatively relate to performance improvements on subsequent discriminations. These predictions are strongest for the earliest stages of learning where weight adjustments are critical for subsequent performance, but they may also hold for later stages of learning, when correlations on irrelevant dimensions, including

independent noise channels, could potentially lead to systematic deviations from optimal readout (Fig 2f, 4d&e).

One interesting special case involves tasks where the relevant dimension changes in an unsignaled manner (Birrell & Brown, 2000; Haefner et al., 2016). In such tasks, noise correlations on the previously relevant dimension would, after such an “extradimensional shift”, force gradients into a task-irrelevant dimension and thus impair learning performance. Interestingly, learning after extra-dimensional shifts can be selectively improved by enhancing noradrenergic signaling (Devauges & Sara, 1990; Lapiz & Morilak, 2006), which leads to increased arousal (Joshi et al., 2016; Reimer et al., 2016) and decreased cortical pairwise noise correlations in sensory and higher order cortex (Joshi & Gold, n.d.; Vinck et al., 2015). While these observations have been made in different paradigms, our model suggests that the reduction of noise correlations resulting from increased sustained levels of norepinephrine after an extradimensional shift (Bouret & Sara, 2005) could mediate faster learning by expanding the dimensionality of the learning gradients (compare figure 4G to 4f) to consider features that have not been task-relevant in the past.

Origins of useful noise correlations

One important question stemming from our work is how noise correlations emerge in the brain. This question has been one of longstanding debate, largely because there are so many potential mechanisms through which correlations could emerge (Kanitscheider et al., 2015; Kohn et al., 2016). Noise correlations could emerge from convergent and divergent feed forward wiring (Shadlen & Newsome, 1998), local connectivity patterns within a neural population (Hansen, Chelaru, & Dragoi, 2012; Smith, Jia, Zandvakili, & Kohn, 2013), or top down inputs provided to separately to different neural populations (Haefner et al., 2016). Here we show that static noise correlations that are useful for perceptual learning emerge naturally from Hebbian learning in a feedforward network. While this certainly suggests that useful noise correlations could emerge through feed forward wiring, it is also possible to consider our Hebbian learning as occurring in a one-step recurrence of the input units, and thus the same data support the possibility of noise correlations through local recurrence. The context dependent noise correlations that speed learning (Fig 4), however, would not arise through simple Hebbian learning. However, such correlations could potentially be produced through selective top-down signals from the choice neurons, as has been previously proposed (Bondy et al., 2018; Haefner et al., 2016; Lange et al., 2018; Wimmer et al., 2015). Moreover, top-down input may selectively target neuronal ensembles produced through Hebbian learning (Collins & Frank, 2013). While previous work has suggested that such a mechanism could be adaptive for accumulating information over the course of a decision (Haefner et al., 2016), our work demonstrates that the same mechanism could effectively be used to tag

relevant neurons for weight updating between trials, making efficient use of top-down circuitry. Haimerl et al. recently made a similar point, showing that stochastic modulatory signals shared across task-informative neurons can serve to tag them for a decoder (Haimerl et al., 2019).

Noise correlations as inductive biases

Artificial intelligence has undergone a revolution over the past decade leading to human level performance in a wide range of tasks (Mnih et al., 2015). However, a major issue for modern artificial intelligence systems, which build heavily on neural network architectures, is that they require far more training examples than a biological system would (Hassabis, Kumaran, Summerfield, & Botvinick, 2017). This biological advantage occurs despite the fact that the total number of synapses in the human brain, which could be thought of as the free parameters in our learning architecture, is much greater than the number of weights in even the most parameter-heavy deep learning architectures. Our work provides some insight into why this occurs; correlated variability across neurons in the brain constrain learning to specific dimensions, thereby limiting the effective complexity of the learning problem (Fig 4f-g). We show that, for simple tasks, this can be achieved using Hebbian learning rules (Fig 3), but that contextual noise correlations, of the form that might be produced through top-down signals (Haefner et al., 2016), are critical for appropriately focusing learning in more complex circumstances. In principle, algorithms that effectively learn and implement noise correlations might reduce the amount of data needed to train AI systems by limiting degrees of freedom to those dimensions that are most relevant. Furthermore, our work suggests that large scale neural recordings in early stages of learning complex tasks might serve as indicators of the inductive biases that constrain learning in biological systems.

In summary, we show that under external constraints of task-relevant information, noise correlations that have previously been called “rate limiting” can serve an important role in constraining learning to task-relevant dimensions. In the context of previous theory focusing on representation, our work suggests that neural populations are subject to competing forces when optimizing covariance structures; on one hand reducing correlations between pairs of similarly tuned neurons can be helpful to fully represent available information, but increasing correlations among similarly tuned neurons can be helpful for assigning credit to task relevant features. We believe that this view of the learning process not only provides insight to understanding the role of noise correlations in the brain, but opens up the door to better understand the inductive biases that guide learning in biological systems.

Methods:

Learning readout in perceptual learning task

Simulations and analyses were performed using a simplified and statistically tractable two-layer neural network (Fig 2a). The input layer consisted of two pools of 100 units that were each "tuned" to one of two motion directions (left, right). On each trial normalized firing rates for the neural population were drawn from a multivariate normal distribution that was specified by a vector of stimulus-dependent mean firing rates (signal: +1 for preferred stimulus, -1 for non-preferred stimulus) and a covariance matrix. All elements of the covariance matrix corresponding to covariance between units that were "tuned" to different stimuli were set to zero. The key manipulation was to systematically vary the magnitude of diagonal covariance components (eg. noise in the firing of individual units) and the "same pool" covariance elements (eg. shared noise across identically tuned neurons) while maintaining a fixed level of variance in the summed population response for each pool:

$$\sigma_{pool}^2 = n\sigma_{unit}^2 + n(n-1)Cov(within\ pool) \quad Eq.1$$

Where σ_{pool}^2 is the variance on the sum of normalized firing rates from neurons within a given pool, n is the number of units in the pool and the within pool covariance ($Cov(within\ pool)$) specifies the covariance of pairs of units belonging to the same pool. The signal to noise ratio for the population response was fixed to one. Given this constraint, the fraction of the total population noise that was shared across neurons was manipulated as follows:

$$\sigma_{unit}^2 = \frac{\sigma_{pool}^2}{n + n(n-1)\phi} \quad Eq.2$$

$$Cov(within\ pool) = \phi\sigma_{unit}^2 \quad Eq.3$$

Where ϕ reflects the fraction of noise that is correlated across units, which we refer to in the text as noise correlations. Noise correlations were manipulated across values ranging from 0 to 0.2 for simulations.

The output layer contained one unit for each pool in the input layer, and was fully connected to the input units in a feedforward manner. Output units were activated on a given trial according to weighted function of their inputs:

$$F_j = \sum_i F_i w_{i,j} \quad Eq.4$$

Actions were selected as a softmax function of output firing rates:

$$p(A_j) = \frac{e^{\beta F_j}}{\sum_k e^{\beta F_k}} \quad Eq.5$$

where β is an inverse temperature, which was set to a relatively deterministic value (10000). Learning was implemented through reinforcement learning of weights to the selected output neuron:

$$\Delta w_{i,j} = \alpha \delta F_i \quad Eq.6$$

Where F_i is the normalized firing rate of the ith neuron, δ is the reward prediction error experienced on a given trial [+0.5 for correct trials and -0.5 for error trials], and α is a learning rate (set to 0.0001 for simulations in figure 2). The network was trained to correctly identify two stimuli (each of which was preferred by a single pool of input neurons) over 100 trials (the last 20 trials of which were considered testing). Simulations were repeated 1000 times for each level of ϕ and performance measures were averaged across all repetitions. Mean accuracy per trial across all simulations was convolved with a Gaussian kernel (standard deviation = 0.5 trials) for plotting in figure 2b. Mean accuracy across the final 20 trials was used as a measure of final accuracy (figure 2e). Statistics on model performance were computed as Pearson correlations between noise correlations ϕ and performance measures across all simulations and repetitions.

Hebbian learning of noise correlations in three layer network

We extended the two-layer feedforward architecture described above to include a third hidden layer in order to test whether Hebbian learning could facilitate production of noise correlations among similarly tuned neurons. The input layer was fully connected to the hidden layer, and each layer contained 200 neurons. In the input layer, neurons were tuned (100 leftward, 100 rightward) as described above, with ϕ set to zero (eg. no noise correlations). Weights to the hidden layer were initialized to favor one-to-one connections between input layer units and

hidden layer units by adding a small normal random weight perturbation (mean=0, standard deviation = 0.01) to an identity matrix. During learning, weights between the input and hidden layer were adjusted according to a normalized Hebbian learning rule:

$$\Delta W = \alpha_{hebb} F_1' F_2 \quad Eq.7$$

Where F_1' is a normalized vector of firing rates corresponding to the input layer and F_2 is a normalized vector of firing rates corresponding to the hidden layer units. The learning rate for Hebbian plasticity (α_{hebb}) was set to 0.00005 for simulations in figure 3. The model was "trained" over 100 trials in the same perceptual discrimination task described above and an additional 100 trials of the task were completed to measure emergent noise correlations in the hidden layer. Noise correlations were measured by regressing out variance attributable to the stimulus on each trial, and then computing the Pearson correlation of residual firing rate across each pair of neurons for the 100 testing trials (Figure 3b&c).

Learning readout in multiple discrimination task

In order to test the impact of contextual noise correlations on learning (Cohen & Newsome, 2008), the perceptual discrimination task was extended to include two dimensions and two interleaved trial types: one in which an up/down discrimination was performed (vertical), and one in which a right/left discrimination was performed (horizontal). Each trial contained motion on the vertical axis (up or down) and on the horizontal axis (left or right), but only one of these motion axes was relevant on each trial as indicated by a cue.

In order to model this task we extended our two-layer feed-forward network to include 4 populations of input units, 4 output units, and 2 task units. Each population of 100 input units encoded a conjunction of the movement directions (up-right, up-left, down-right, down-left). On each trial, the mean firing rate of each input unit population was determined according to their tuning preferences:

$$\mu = V + H \quad Eq.8$$

Where V was +1/-1 for trials with the preferred/anti-preferred vertical motion direction H was +1/-1 for trials with the preferred/anti-preferred horizontal motion direction. Firing rates for individual neurons were sampled from a multivariate Gaussian distribution with mean μ and a covariance matrix that depended on trial type (vertical versus horizontal) and the level of same pool, relevant pool, and irrelevant pool correlations.

In order to create a covariance matrix, we stipulated a desired level standard error of the mean summed population activity (SEM=20 for simulations in figure 4) and determined the summed population variance that would correspond to that value (σ^2_{pool}). We then determined the variance on individual neurons that would yield this population response under a given noise correlation profile as follows:

$$\sigma^2_{unit} = \frac{\sigma^2_{pool}}{n + n(n-1)\phi_{same} + n^2\phi_{relevant} - n^2\phi_{irrelevant}} \quad Eq.9$$

Where ϕ_{same} is the level of same pool correlations (range: 0-0.2 in our simulations), $\phi_{relevant}$ is the level of relevant pool correlations (range: 0-0.2 in our simulations), $\phi_{irrelevant}$ is the level of irrelevant pool correlations (range: 0-0.2 in our simulations. Note that increasing the same pool or in pool correlations reduces the overall variance in order to preserve the same level of variance on the task relevant dimension in the population response, but that increasing irrelevant pool correlations has the opposite effect. Covariance elements of the covariance matrix were determined as follows:

$$Cov(same\ pool) = \phi_{same}\sigma^2_{unit} \quad Eq.10$$

$$Cov(relevant\ pool) = \phi_{relevant}\sigma^2_{unit} \quad Eq.11$$

$$Cov(irrelevant\ pool) = \phi_{irrelevant}\sigma^2_{unit} \quad Eq.12$$

Variance and covariance values above were used to construct a covariance matrix for each trial type (vertical/horizontal) as depicted in figure 5.

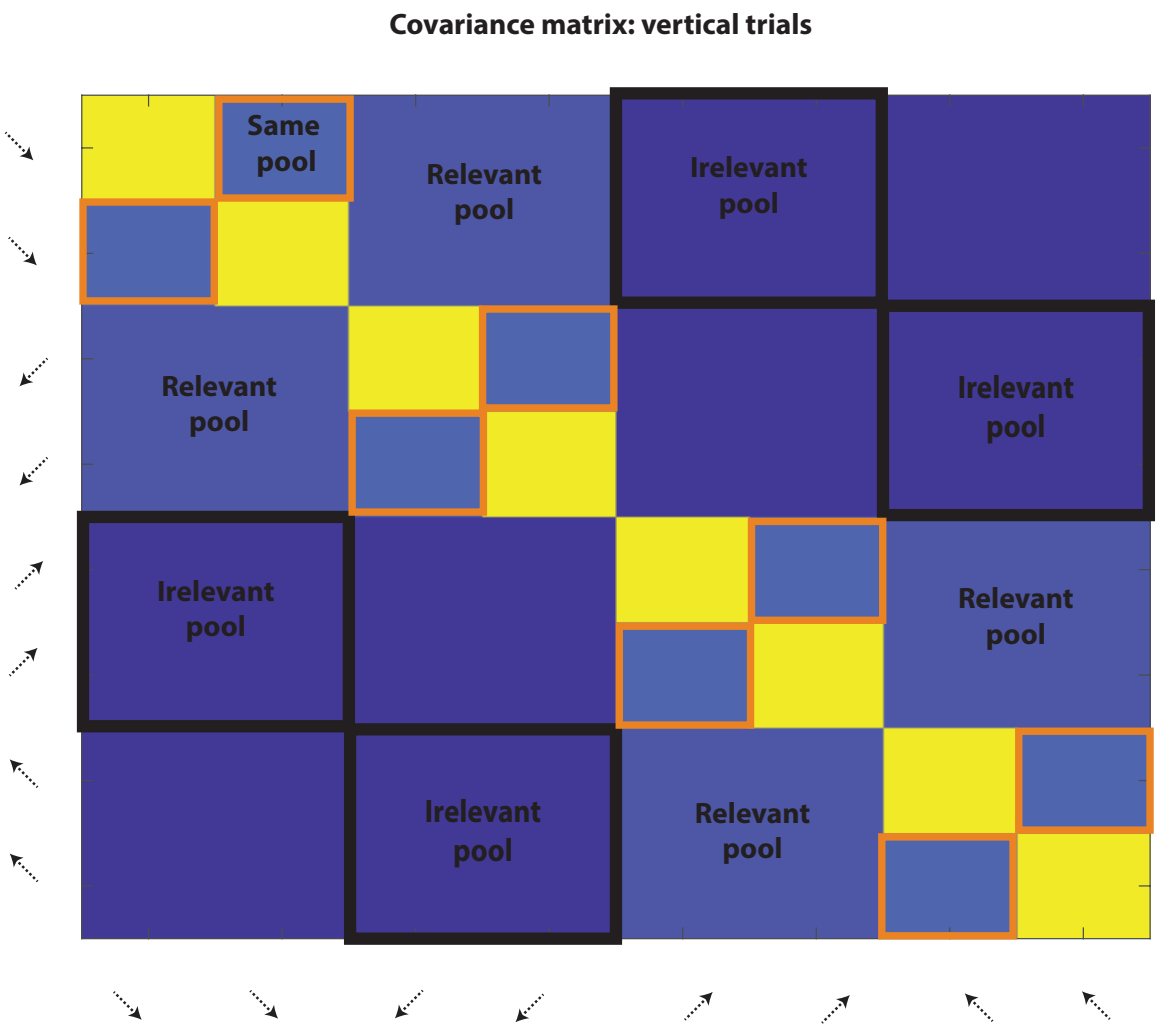

Figure 5: Schematic of covariance matrix for two-dimensional motion discrimination task. Same pool correlations are controlled by covariance elements between neurons with identical tuning (orange boxes). Relevant pool correlations are controlled by covariance elements between neurons that are similarly tuned to the task-relevant feature. Task irrelevant correlations are controlled by covariance elements between neurons that are similarly tuned to the task-irrelevant feature. The covariance matrix shown here is for a vertical trial – on a horizontal trial the irrelevant pool and relevant pool locations would be reversed. Covariance elements for pairs of neurons that differed in tuning on both dimensions were set to zero. Each input population has been depicted as two units here for presentation purposes. Background color reflects the case where same pool correlations = 0.2 and relevant pool correlations = 0.1.

Output units corresponded to the four possible task responses (up, down, left, right) and were activated according to a weighted sum of their inputs as described previously. Task units were modeled as containing perfect information about the task cue (vertical versus horizontal) and were modeled to completely inhibit the responses of the irrelevant output units. Decisions were made on each trial by selecting the output unit with the highest activity level. Weights to chosen output unit were updated using the same reinforcement learning procedure that was used in the two alternative perceptual learning task.

Supplementary figures:

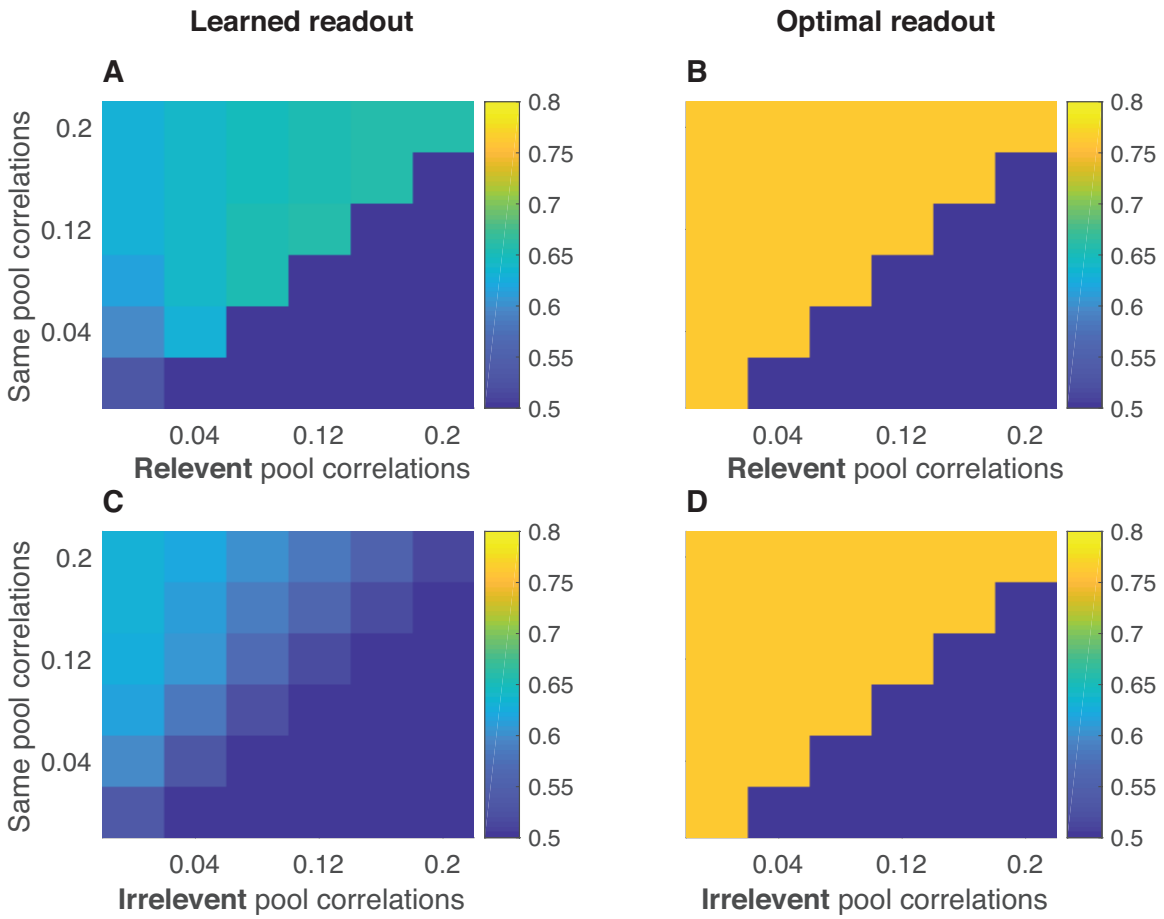

Figure 4S1: **Noise correlations affect speed of learning, but not performance using optimal readout**. **A)** Mean test accuracy (color) of all models spanning the range of in pool correlations (abscissa) and relevant pool correlations (ordinate). **B)** Mean accuracy of same models using optimal readout, rather than the learned readout. **C)** Mean test accuracy (color) of all models spanning the range of in pool correlations (abscissa) and irrelevant pool correlations (ordinate). **D)** Mean accuracy of same models using optimal readout, rather than the learned readout.

Adibi, M., McDonald, J. S., Clifford, C. W. G., & Arabzadeh, E. (2013). Adaptation improves neural coding efficiency despite increasing correlations in variability. *Journal of Neuroscience*, *33*(5), 2108–2120. http://doi.org/10.1523/JNEUROSCI.3449-12.2013

Akaishi, R., Kolling, N., Brown, J. W., & Rushworth, M. (2016). Neural Mechanisms of Credit Assignment in a Multicue Environment. *Journal of Neuroscience*, *36*(4), 1096–1112. http://doi.org/10.1523/JNEUROSCI.3159-15.2016

Averbeck, B. B., & Lee, D. (2003). Neural noise and movement-related codes in

the macaque supplementary motor area. *Journal of Neuroscience*, *23*(20), 7630–7641.

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*(5), 358–366. http://doi.org/10.1007/978-3-662-02728-8

Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *Journal of Neuroscience*, *21*(5), 1676–1697.

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Perspective. *Neuron*, *74*(1), 30–39. http://doi.org/10.1016/j.neuron.2012.03.016

Birrell, J. M., & Brown, V. J. (2000). Medial frontal cortex mediates perceptual attentional set shifting in the rat. *Journal of Neuroscience*, *20*(11), 4320–4324.

Bondy, A. G., Haefner, R. M., & Cumming, B. G. (2018). Feedback determines the structure of correlated variability in primary visual cortex. *Nature Publishing Group*, 1–15. http://doi.org/10.1038/s41593-018-0089-1

Bouret, S., & Sara, S. J. (2005). Network reset: a simplified overarching theory of locus coeruleus noradrenaline function. *Trends in Neurosciences*, *28*(11), 574–582. http://doi.org/10.1016/j.tins.2005.09.002

Cohen, M. R., & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature Publishing Group*, *14*(7), 811–819. http://doi.org/10.1038/nn.2842

Cohen, M. R., & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Publishing Group*, *12*(12), 1594–1600. http://doi.org/10.1038/nn.2439

Cohen, M. R., & Maunsell, J. H. R. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*, *70*(6), 1192–1204. http://doi.org/10.1016/j.neuron.2011.04.029

Cohen, M. R., & Newsome, W. T. (2008). Context-Dependent Changes in Functional Circuitry in Visual Area MT. *Neuron*, *60*(1), 162–173. http://doi.org/10.1016/j.neuron.2008.08.007

Collins, A. G. E., & Frank, M. J. (2013). Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychological Review*, *120*(1), 190–229. http://doi.org/10.1037/a0030852

Devauges, V., & Sara, S. J. (1990). Activation of the noradrenergic system facilitates an attentional shift in the rat. *Behavioural Brain Research*, *39*(1), 19–28.

Doiron, B., Litwin-Kumar, A., Rosenbaum, R., Ocker, G. K., & Josić, K. (2016). The mechanics of state-dependent neural correlations. *Nature Publishing Group*, *19*(3), 383–393. http://doi.org/10.1038/nn.4242

Downer, J. D., Niwa, M., & Sutter, M. L. (2015). Task engagement selectively modulates neural correlations in primary auditory cortex. *Journal of Neuroscience*, *35*(19), 7565–7574. http://doi.org/10.1523/JNEUROSCI.4094-

14.2015
Ecker, A. S., Berens, P., Keliris, G. A., Bethge, M., Logothetis, N. K., & Tolias, A. S. (2010). Decorrelated neuronal firing in cortical microcircuits. *Science*, *327*(5965), 584–587. http://doi.org/10.1126/science.1179867
Gu, Y., Liu, S., Fetsch, C. R., Yang, Y., Fok, S., Sunkara, A., et al. (2011). Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron*, *71*(4), 750–761. http://doi.org/10.1016/j.neuron.2011.06.015
Haefner, R. M., Pietro Berkes, & Fiser, J. (2016). Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron*, *90*(3), 649–660. http://doi.org/10.1016/j.neuron.2016.03.020
Haimerl, C., Savin, C., & Simoncelli, E. P. (2019). Flexible and accurate decoding of neural populations through stochastic comodulation. *bioRxiv*, *21*(1), 598. http://doi.org/10.1016/S0896-6273(00)80821-1
Hansen, B. J., Chelaru, M. I., & Dragoi, V. (2012). Correlated variability in laminar cortical circuits. *Neuron*, *76*(3), 590–602. http://doi.org/10.1016/j.neuron.2012.08.029
Hassabis, D., Kumaran, D., Summerfield, C., & Botvinick, M. (2017). Neuroscience-Inspired Artificial Intelligence. *Neuron*, *95*(2), 245–258. http://doi.org/10.1016/j.neuron.2017.06.011
Hawkey, D. J. C., Amitay, S., & Moore, D. R. (2004). Early and rapid perceptual learning. *Nature Publishing Group*, *7*(10), 1055–1056. http://doi.org/10.1126/science.1589770
Herrero, J. L., Gieselmann, M. A., Sanayei, M., & Thiele, A. (2013). Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron*, *78*(4), 729–739. http://doi.org/10.1016/j.neuron.2013.03.029
Huang, X., & Lisberger, S. G. (2009). Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth-pursuit eye movements. *Journal of Neurophysiology*, *101*(6), 3012–3030. http://doi.org/10.1152/jn.00010.2009
Joshi, S., & Gold, J. I. (n.d.). Context-Dependent Relationships between Locus Coeruleus Firing Patterns and Coordinated Neural Activity in the Anterior Cingulate Cortex. *bioRxiv*. http://doi.org/10.1126/science.281.5379.985
Joshi, S., Li, Y., Kalwani, R. M., & Gold, J. I. (2016). Relationships between Pupil Diameter and Neuronal Activity in the Locus Coeruleus, Colliculi, and Cingulate Cortex. *Neuron*, *89*(1), 221–234. http://doi.org/10.1016/j.neuron.2015.11.028
Kanitscheider, I., Coen-Cagli, R., & Pouget, A. (2015). Origin of information-limiting noise correlations. *Proceedings of the National Academy of Sciences*, *112*(50), E6973–E6982. http://doi.org/10.1038/nn.3309
Kohn, A., Coen-Cagli, R., Kanitscheider, I., & Pouget, A. (2016). Correlations and Neuronal Population Information. *Annual Review of Neuroscience*, *39*, 237–256. http://doi.org/10.1146/annurev-neuro-070815-013851
Krotov, D., & Hopfield, J. J. (2019). Unsupervised learning by competing hidden

units. *Proceedings of the National Academy of Sciences*, *116*(16), 7723–7731. http://doi.org/10.3389/fncom.2015.00099

Lange, R. D., Chattoraj, A., Beck, J. M., Yates, J. L., & Haefner, R. M. (2018). A confirmation bias in perceptual decision-making due to hierarchical approximate inference. *bioRxiv*. http://doi.org/10.1016/j.neuron.2012.09.015

Lapiz, M. D. S., & Morilak, D. A. (2006). Noradrenergic modulation of cognitive function in rat medial prefrontal cortex as measured by attentional set shifting capability. *Neuroscience*, *137*(3), 1039–1049. http://doi.org/10.1016/j.neuroscience.2005.09.031

Law, C.-T., & Gold, J. I. (2009). Reinforcement learning can account for associative and perceptual learning on a visual-decision task. *Nature Neuroscience*, *12*(5), 655–663. http://doi.org/10.1038/nn.2304

Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic Interaction between Reinforcement Learning and Attention in Multidimensional Environments. *Neuron*, *93*(2), 451–463. http://doi.org/10.1016/j.neuron.2016.12.040

Mack, M. L., Preston, A. R., & Love, B. C. (2019). Ventromedial prefrontal cortex compression during concept learning. *Nature Communications*, 1–11. http://doi.org/10.1038/s41467-019-13930-8

Maynard, E. M., Hatsopoulos, N. G., Ojakangas, C. L., Acuna, B. D., Sanes, J. N., Normann, R. A., & Donoghue, J. P. (1999). Neuronal interactions improve cortical population coding of movement direction. *Journal of Neuroscience*, *19*(18), 8083–8093.

Mitchell, J. F., Sundberg, K. A., & Reynolds, J. H. (2009). Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron*, *63*(6), 879–888. http://doi.org/10.1016/j.neuron.2009.09.013

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. http://doi.org/10.1038/nature14236

Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014). Information-limiting correlations. *Nature Publishing Group*, *17*(10), 1410–1417. http://doi.org/10.1038/nn.3807

Ni, A. M., Ruff, D. A., Alberts, J. J., Symmonds, J., & Cohen, M. R. (2018). Learning and attention reveal a general relationship between population activity and behavior. *Science*, *359*(6374), 463–465. http://doi.org/10.1126/science.aao0284

Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 1–7.

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125–132. http://doi.org/10.1038/35039062

Reimer, J., McGinley, M. J., Liu, Y., Rodenkirch, C., Wang, Q., McCormick, D. A., & Tolias, A. S. (2016). Pupil fluctuations track rapid changes in adrenergic and cholinergic activity in cortex. *Nature Communications*, *7*, 13289.

http://doi.org/10.1038/ncomms13289

Ruff, D. A., & Cohen, M. R. (2014). Attention can either increase or decrease spike count correlations in visual cortex. *Nature Publishing Group*, *17*(11), 1591–1597. http://doi.org/10.1038/nn.3835

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of Neuroscience : the Official Journal of the Society for Neuroscience*, *18*(10), 3870–3896.

Smith, M. A., Jia, X., Zandvakili, A., & Kohn, A. (2013). Laminar dependence of neuronal correlations in visual cortex. *Journal of Neurophysiology*, *109*(4), 940–947. http://doi.org/10.1038/370140a0

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

Tsividis, P., Pouncy, T., Xu, J. L., Tenenbaum, J. B., & Gershman, S. J. (2017). Human Learning in Atari. *2017 AAAI Spring Symposium Series, Science of Intelligence: Computational Principles of Natural and Artificial Intelligence*, 1–4.

Vinck, M., Batista-Brito, R., Knoblich, U., & Cardin, J. A. (2015). Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding. *Neuron*, *86*(3), 740–754. http://doi.org/10.1016/j.neuron.2015.03.028

Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nakajima, M., Deisseroth, K., & Halassa, M. M. (2015). Thalamic control of sensory selection in divided attention. *Nature*, *526*(7575), 705–709. http://doi.org/10.1097/WNR.0b013e3282f4cb3a

Zanto, T. P., & Gazzaley, A. (2009). Neural Suppression of Irrelevant Information Underlies Optimal Working Memory Performance. *Journal of Neuroscience*, *29*(10), 3059–3066. http://doi.org/10.1523/JNEUROSCI.4621-08.2009

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, *370*(6485), 140–143. http://doi.org/10.1038/370140a0