Neuron
Article

Associative Learning Enhances Population Coding by Inverting Interneuronal Correlation Patterns

James M. Jeanne,[1,2,5,6] Tatyana O. Sharpee,[1,3,5] and Timothy Q. Gentner[1,2,4,*]
[1]Neurosciences Graduate Program
[2]Department of Psychology
[3]Center for Theoretical Biological Physics
[4]Kavli Institute for Brain and Mind
University of California San Diego, La Jolla, CA 92093, USA
[5]Computational Neurobiology Laboratory, Salk Institute for Biological Studies, La Jolla, CA 92093, USA
[6]Present address: Department of Neurobiology, Harvard Medical School, Boston, MA 02115, USA
*Correspondence: tgentner@ucsd.edu
http://dx.doi.org/10.1016/j.neuron.2013.02.023

SUMMARY

Learning-dependent cortical encoding has been well described in single neurons. But behaviorally relevant sensory signals drive the coordinated activity of millions of cortical neurons; whether learning produces stimulus-specific changes in population codes is unknown. Because the pattern of firing rate correlations between neurons—an emergent property of neural populations—can significantly impact encoding fidelity, we hypothesize that it is a target for learning. Using an associative learning procedure, we manipulated the behavioral relevance of natural acoustic signals and examined the evoked spiking activity in auditory cortical neurons in songbirds. We show that learning produces stimulus-specific changes in the pattern of interneuronal correlations that enhance the ability of neural populations to recognize signals relevant for behavior. This learning-dependent enhancement increases with population size. The results identify the pattern of interneuronal correlation in neural populations as a target of learning that can selectively enhance the representations of specific sensory signals.

INTRODUCTION

Vertebrate behaviors, from perception to action, are mediated by large ensembles of neurons (Averbeck et al., 2006). Learning, in turn, enables long-term changes in behavior by altering associations between specific sensory stimuli, actions, and the outcomes of those actions. Flexible neural representations in higher-order sensory cortical areas are believed to underlie these learned associations (Reed et al., 2011). Consistent with this, changes in single-neuron representations for behaviorally relevant stimuli are well documented (Blake et al., 2002, 2006; Gentner and Margoliash, 2003; Jeanne et al., 2011; Meliza and Margoliash, 2012; Thompson and Gentner, 2010; Thompson et al., 2013). In contrast, little is known about how, or even if, learning might act on the neural ensemble representations by changing emergent properties not observable in single neurons. One such property that learning could target is the correlation between neural firing rates, which can dramatically influence the fidelity of a population code (Cohen and Maunsell, 2009; Cohen and Kohn, 2011). Here, we examine how associative learning influences the stimulus-specific pattern of interneuronal correlations and encoding among neural ensembles in a high-level auditory region in the songbird brain.

Neurons are inherently noisy: multiple presentations of an identical sensory stimulus do not produce identical responses (Huber et al., 2008). Pooling responses across distributed populations of similarly tuned neurons can enhance encoding fidelity by averaging out this response variability (known as "noise correlation"), but only the component of this noise that is independent between neurons (Zohary et al., 1994). Neural variability, however, is rarely independent between neurons. Throughout the cortex, values of noise correlation tend to be broadly distributed, being small but positive on average (Cohen and Kohn, 2011). Consequently, noise correlations are traditionally thought to limit the value of population response pooling.

The effects of noise correlations, however, can be diverse. Most cortical circuits contain neurons with heterogeneous tuning functions. In such circuits, noise correlations can either enhance or impair coding fidelity, depending on how the noise correlation relates to tuning similarity (known as "signal correlation") for each pair of neurons (Abbott and Dayan, 1999; Averbeck et al., 2006; Cafaro and Rieke, 2010; Gu et al., 2011; Wilke and Eurich, 2002). Compared to independent noise, positively correlated noise between two similarly tuned neurons impairs encoding because no form of response pooling can attenuate the shared noise without simultaneously attenuating the signal (Bair et al., 2001; Shadlen et al., 1996; Shadlen and Newsome, 1998; Zohary et al., 1994). In contrast, positively correlated noise between two oppositely tuned neurons can improve encoding because subtracting one response from the other can both attenuate the shared noise and strengthen the signal (Romo et al., 2003).

In the constituent pairs of large neural populations in the cortex, noise correlations tend to positively covary with signal correlations (Bair et al., 2001; Cohen and Maunsell, 2009; Gu et al., 2011; Kohn and Smith, 2005). Such a correlation structure reduces population coding fidelity relative to independent noise

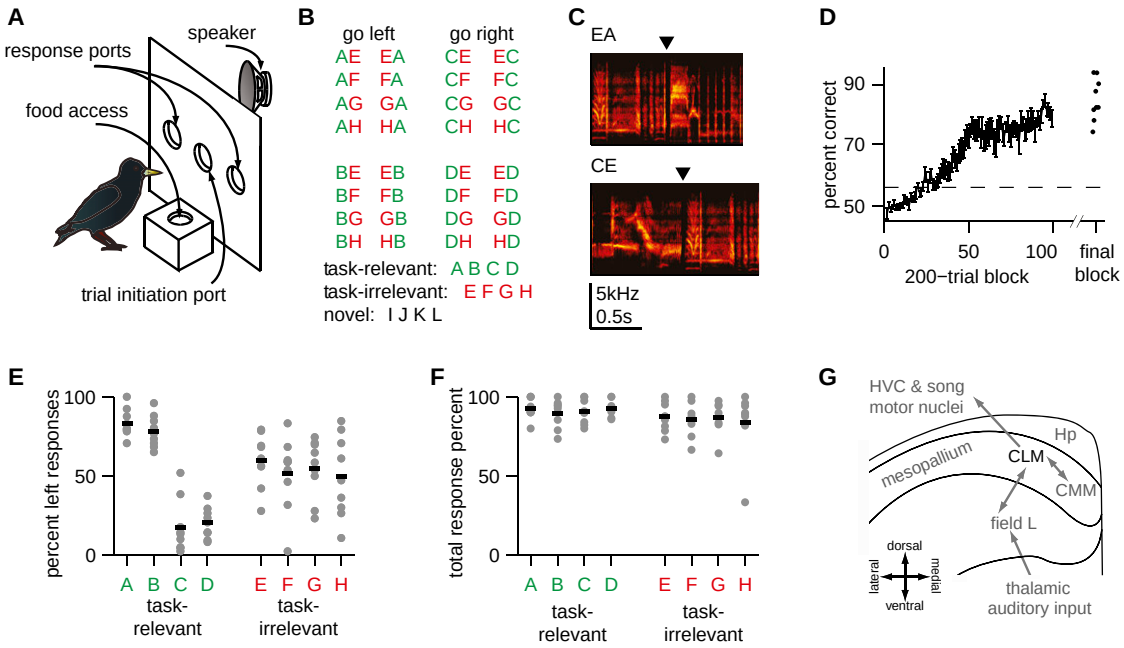

Figure 1. Experimental Design

(A) Schematic of behavioral training apparatus.

(B) Stimulus design. Each letter denotes a single motif. Task-relevant motifs (green) indicated whether to respond left or right. Task-irrelevant motifs (red) were paired in sequence with task-relevant motifs and occurred with equal probability in left stimuli and right stimuli. Novel motifs (black) were never presented during behavioral training but were presented during neural recording.

(C) Example "go left" stimulus (top) and "go right" stimulus (bottom). Arrowheads denote 20 ms silent periods between motifs.

(D) Mean (±SEM) acquisition curve showing increase in performance with training. Blocks are nonoverlapping. Dots at right denote behavioral performance for each bird during the 200 trials immediately prior to neural recording.

(E) Summary of responses to probe trials. Each probe trial was a single motif presented in isolation (i.e., not paired) and not reinforced. Data shown here are the percentages of all trials that elicited left responses. Each point represents a single bird. Task-relevant motifs in isolation show a strong learned association with left and right responses. Task-irrelevant motifs in isolation show no learned association, with each motif eliciting similar responses, on average. Although motifs were counterbalanced between birds, the labels for task-relevant and task-irrelevant motifs are given as A–D and E–H, respectively, for ease of display and interpretation. The mean task-irrelevant responses were intermediate to the go left and go right responses for all subjects.

(F) As in (E) but for the total number of responses (both left and right). Task-irrelevant motifs elicited response rates nearly as high as task-relevant motifs. Birds thus associate each task-relevant motif with a pecking in a specific port but associate each task-irrelevant motif only with the general response of pecking.

(G) Schematic of CLM within the avian auditory forebrain circuitry. Arrows denote major projection patterns. Hp, hyperpallium; CMM, caudomedial mesopallium. See also Figure S2.

because the similarly tuned pairs tend to have high noise correlation and dissimilarly tuned pairs tend to have low noise correlation (Gu et al., 2011). Conversely, an inverted correlation structure in which noise correlations negatively covary with signal correlations can yield higher-fidelity population representations relative to independent noise (see Figure S1 available online) (Averbeck et al., 2006; Gu et al., 2011). Thus, the relationship between signal and noise correlations has considerable impact on the fidelity of neural representations that cannot be predicted from the average responses of individual neurons. We hypothesized that learning-driven plasticity in the population correlation structure could provide a mechanism for selectively strengthening neural representations of important sensory signals.

To test this hypothesis, we investigated the effect of associative learning of natural song components ("motifs") on the relationship between signal and noise correlations in a higher-order auditory cortical area of the songbird brain. We found that learning inverted the relationship between signal and noise correlations in auditory cortex. Remarkably, this effect was restricted to the subset of motifs that explicitly guided the subjects' learned behaviors ("task-relevant" motifs). Equally familiar motifs that did not guide behavior ("task-irrelevant" motifs) and novel motifs elicited the canonical positive relationship between signal and noise correlations. This plasticity in the correlation structure yielded a modest, but significant, enhancement to the encoding fidelity of task-relevant motifs by pairs of neurons. The magnitude of this enhancement, however, grew larger for larger populations. These results reveal the interneuronal correlation structure as a target for learning-dependent enhancement of sensory encoding.

RESULTS

Associative Learning of Motifs

To understand how learning influences interneuronal correlations and sensory encoding by neural populations, we first trained European starlings (*Sturnus vulgaris*) to associate specific motifs with behaviors that led to reward (Figures 1A–1D;

see Experimental Procedures). In the wild, recognition of learned motifs underlies behaviors such as mate attraction and resource defense (Eens, 1997; Gentner and Hulse, 2000). In the laboratory, we controlled motif recognition with a two-alternative choice operant task. On each trial during training, birds heard a pair of sequentially ordered motifs (e.g., Figure 1C). One motif in the pair (referred to as "task relevant") always signaled the correct behavioral response for the trial (i.e., whether to peck at the left or right port to receive food) and the other motif (referred to as "task irrelevant") never signaled the correct response (Figure 1B). The task-relevant motif could be presented as either the first or the second motif in the pair. The task relevance of any given motif was held constant within a bird and counterbalanced across birds. All the training motifs were equally associated with food reward. All birds (n = 9) learned to perform this task accurately (Figure 1D).

To verify that learned behavior depended on the task relevance of the motifs rather than the association with reward, we tested the birds' behavioral responses to each motif in isolation (i.e., not paired; Experimental Procedures). As expected, each of the single task-relevant motifs evoked responses primarily to a single port, following the learned responses from training (Figure 1E). In contrast, responses to the task-irrelevant motifs were evenly distributed between both response ports (Figure 1E). Importantly, overall response rates for all the motifs were similarly high (Figure 1F). Thus, all of the motifs made familiar during operant training are associated with the general behavior of pecking (or perhaps the common outcome of that behavior, namely food), but only the task-relevant motifs are associated with the specific choice of pecking either left or right. The primary difference between the task-relevant and task-irrelevant motifs was thus the learned association between motifs and explicit behavioral choices.

Single Neuron Response Properties in CLM

After training, we recorded the simultaneous activity of multiple well-isolated single neurons in the caudolateral mesopallium (CLM) in response to task-relevant and task-irrelevant motifs and a third set of novel motifs under urethane anesthesia (Figures S2A–S2P; Experimental Procedures). CLM is a higher-order auditory region in the songbird cortex that is specialized for processing learned songs (Jeanne et al., 2011) and projects auditory information into the vocal premotor region HVC (Bauer et al., 2008) (Figure 1G).

Because connectivity and response properties within neural populations depend on cell type (Constantinidis and Goldman-Rakic, 2002; Hofer et al., 2011; Lee et al., 1998), we divided our data set into wide spiking (WS) and narrow spiking (NS) neurons on the basis of action potential width (Barthó et al., 2004; Mitchell et al., 2007; Experimental Procedures; Figures S2Q–S2S). We focus on WS neurons (n = 176 pairs from 98 single neurons) because our sample of NS neurons was not sufficient (n = 17 pairs from 36 single neurons) to perform reliable analysis.

Presentation of motifs evoked complex responses from individual neurons in CLM. Figure 2 shows the responses of two (simultaneously recorded) neurons to the presentation of task-relevant motifs (Figure 2A) and task-irrelevant motifs (Figure 2B). As was typical across our data set, these example neurons responded with different mean firing rates to different motifs and had considerable trial-to-trial variability. On average, firing rates were modestly higher for task-relevant motifs (3.03 ± 0.38 Hz) than for task-irrelevant motifs (2.74 ± 0.33 Hz, Wilcoxon signed-rank test, p = 0.0080) and were similar between task-irrelevant motifs and novel motifs (2.80 ± 0.34 Hz). This finding is consistent with previous reports that song recognition learning alters encoding by single neurons in CLM (Jeanne et al., 2011) and neighboring regions (Gentner and Margoliash, 2003; Thompson and Gentner, 2010). The modulation of single-neuron firing rates is subtle, however, especially in light of the animals' robust changes in behavioral performance over training (Figure 1D) and differential responding to relevant and irrelevant motifs after training (Figures 1E and 1F). Sensory representations may also be distributed across multiple neurons, however, and features that cannot be observed in single neurons in isolation can have profound effects on sensory encoding. We thus hypothesized that learned task relevance influenced interneuronal correlations, a distributed neural feature.

Noise and Signal Correlations in CLM

Learning is known to alter noise correlations in cortical brain regions (Gu et al., 2011). We thus asked whether noise correlations between pairs of CLM neurons during stimulation with motifs depended on the task relevance of the motif. Figures 2E and 2F show the individual trial spike counts (normalized by the Z score to measure noise correlations independently from signal correlations) of the same two neurons from Figures 2A and 2B (Experimental Procedures). The task-relevant motifs elicited nearly uncorrelated responses from this pair (Pearson correlation coefficient, r = 0.01), while the task-irrelevant motifs elicited responses between the pair that were positively correlated (r = 0.20), meaning that when one neuron fired more spikes than average, the other neuron was likely to do so as well. This effect, however, was not observed in all neuron pairs. Figures 2I and 2J show a second example pair in which noise correlations were very similar between task-relevant and task-irrelevant motifs.

To investigate potential differences in the population, we compared noise correlations between all three classes of motif (task-relevant, task-irrelevant, and novel) for all pairs of simultaneously recorded neurons. Consistent with previous reports (Cohen and Kohn, 2011; Gu et al., 2011; Kohn and Smith, 2005; Zohary et al., 1994), we observed broad distributions of noise correlations that had small, but positive, mean values (task relevant: 0.082 ± 0.012; task irrelevant: 0.100 ± 0.012; novel: 0.087 ± 0.012; Figure 3A). Surprisingly, there were no differences in the mean noise correlation between motif classes (repeated-measures ANOVA, p = 0.21; Figures 3A and 3C). A difference in mean noise correlation by itself is thus unlikely to contribute to learning-dependent differences in population coding of motifs in CLM.

Because learning can alter the receptive fields of cortical sensory neurons, we asked whether signal correlation between pairs of CLM neurons depends on task relevance of motifs. As with noise correlations, the effects of task relevance on signal correlations were variable. While the first example pair does

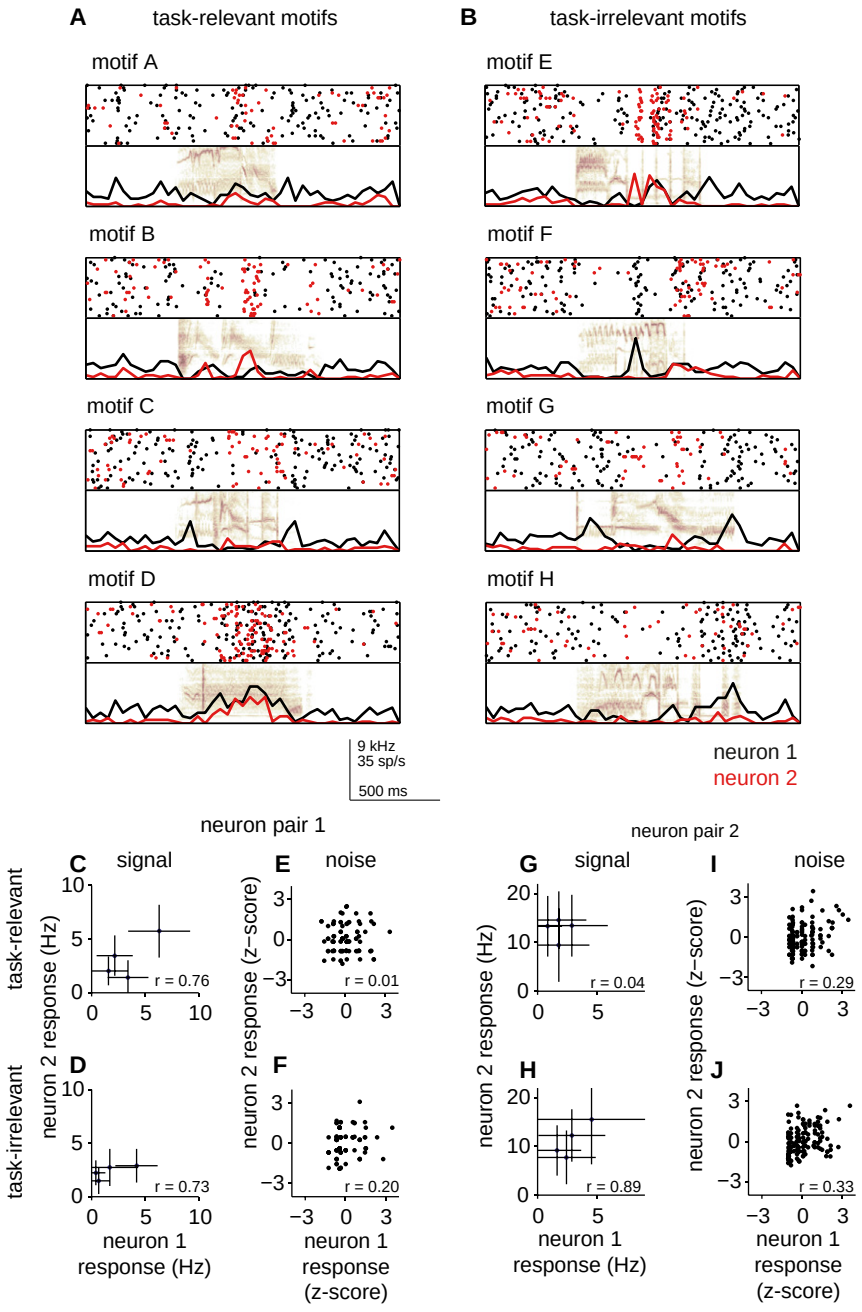

Figure 2. Example Responses from Pairs of Simultaneously Recorded CLM Neurons

(A and B) Spiking activity in response to task-relevant motifs (A) and task-irrelevant motifs (B). The top half of each panel shows spike rasters for repeated trials. Black dots denote spikes from neuron 1 and red dots denote spikes from neuron 2. The bottom half of each panel shows PSTHs for each neuron superimposed over the spectrogram of the motif stimulus.

(C and D) Mean ± SD (averaged over the repeated trials) responses of the same neuron pair from (A) and (B) to the four task-relevant motifs (C) and to the four task-irrelevant motifs (D).

(E and F) Normalized responses to each motif showing no noise correlations for task-relevant motifs (E) and positive noise correlations for task-irrelevant motifs (F).

(G–J) Same as (C)–(F) but for a second example pair of CLM neurons.

not show a considerable difference in signal correlations between task-relevant and task-irrelevant motifs (Figures 2C and 2D), the second example pair shows a large difference (Figures 2G and 2H).

We investigated whether signal correlations exhibited a systematic relationship with task relevance. We observed a broad distribution of signal correlation values for all three motif classes, indicative of the large range of tuning within CLM (Figure 3B). However, we found no evidence that task relevance influenced the magnitude of signal correlations (Figure 3B; Friedman test, $p = 0.18$). Thus, as with noise correlation, a difference in the mean signal correlation by itself is unlikely to contribute to any learning-dependent differences in population coding.

The average magnitude of the noise (or signal) correlation is less critical to encoding, however, than the relationship between the noise and signal correlation (Averbeck et al., 2006; Gu et al., 2011; Wilke and Eurich, 2002). Although no form of response pooling can dissipate positive noise correlations between similarly tuned neurons (positive signal correlation), subtractive pooling can dissipate positive noise correlations between dissimilarly tuned neurons (negative signal correlation). Thus, learning could improve population coding by altering the relationship between the signal correlation and noise correlation.

Learning-Dependent Relationship between Signal and Noise Correlations

To test whether the relationship between signal and noise correlations depends on task relevance, we directly compared these two measures for each pair of neurons in our data set. The example neurons depicted in Figures 2C–2J suggest that although task relevance can influence both signal and noise correlations, it does so following a specific relationship. We thus asked whether noise correlations systematically covary with signal correlations, and whether this depends on task relevance.

We found that each class of motifs exhibited a correlation between signal and noise correlations, but the sign of this relationship depended on task relevance. For task-relevant motifs, this relationship was negative (Spearman correlation coefficient: $r = -0.15$, $p = 0.051$, Figure 4A): larger signal

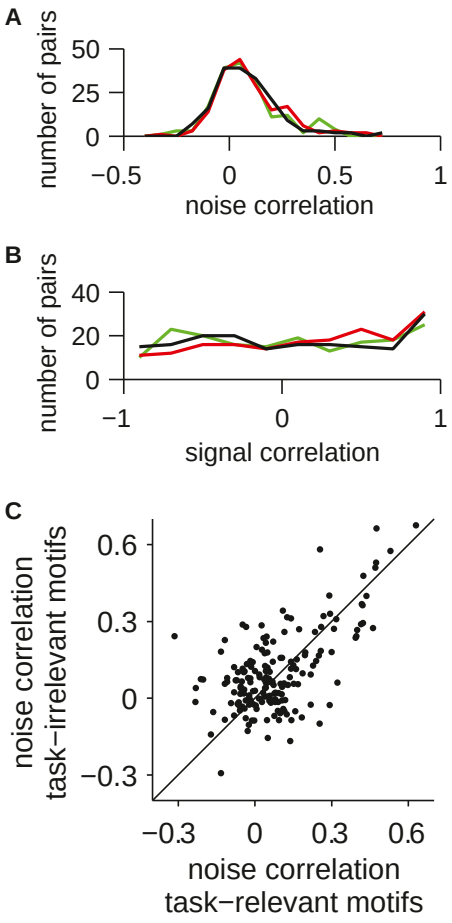

Figure 3. Task Relevance Does Not Influence Average Noise or Signal Correlations

(A) Distributions of noise correlations for task-relevant motifs (green), task-irrelevant motifs (red), and novel motifs (black).

(B) Same as (A) but for signal correlations.

(C) Scatterplot comparing noise correlations for task-relevant motifs with noise correlations for task-irrelevant motifs. Each point denotes one pair. See also Figure S3.

correlations were accompanied by smaller noise correlations. For task-irrelevant and novel motifs, in contrast, the relationship was positive (task irrelevant: $r = 0.19$, $p = 0.012$; novel: $r = 0.23$, $p = 0.0022$; Figures 4B and 4C): larger signal correlations were accompanied by larger noise correlations. The difference between these relationships was highly significant (ANCOVA motif class × regression slope interaction, $p = 7.9 \times 10^{-5}$). In contrast, we found no effects of learning on the relationship between mean firing rate and noise correlation and the relationship between distance between neurons and noise correlation (Figures S3A and S3B).

The relationship between signal correlation and noise correlation thus depends strongly on the learned task relevance of the motif. This dependence is particularly apparent in neuron pairs that have strong (either positive or negative) signal correlations (Figures 4D and 4E). Among neuron pairs with strong positive signal correlations (>0.4), the task-irrelevant and novel motifs evoked significantly larger noise correlations than the task-relevant motifs (Kruskal-Wallis test, $p = 0.0038$; Figure 4D). In contrast, among neuron pairs that had large negative signal correlations (<−0.4), the task-irrelevant and novel motifs evoked significantly weaker noise correlations than the task-relevant motifs (Kruskal-Wallis test, $p = 0.032$; Figure 4E). These observations suggest that learning selectively alters the relationship between the signal and noise correlations for signals that are behaviorally relevant.

Relationship between Correlation Structure and Motif Coding in Neuron Pairs

Previous work has demonstrated that the relationship between signal and noise correlations can dramatically affect population coding (Gu et al., 2011). Thus, we asked whether the observed differences in correlation structure influence the coding of motifs by CLM neurons. Because motif identity is best regarded as a nominal, or categorical, variable (i.e., motifs cannot be easily described with a small number of parameters), we use multinomial logistic regression (MNLR) to find the optimal classifier that maximizes the predictability of motif identity from the firing rates of multiple neurons (Long, 1997) (Experimental Procedures). This technique is particularly well suited to our data because it can be applied to any number of neurons and any number of nominal categories (i.e., motifs). Figure 5 depicts the optimal classifier for a single pair of neurons responding to task-relevant motifs and shows that the classification boundaries follow many of the firing rate patterns that distinguish the motifs. In the example case (Figure 5), the classifier correctly predicted motif identity with 51% accuracy (far better than chance performance of 25%). The MNLR model provides a rigorous quantification of the ability of CLM neurons to discriminate between different motifs.

Using the MNLR model, we first asked whether the relationship between signal and noise correlations benefitted motif discrimination performance. We expressed this potential benefit as the "classification ratio," which is simply the ratio of the MNLR classification accuracy with correlations intact to the classification accuracy without correlations (i.e., with trials shuffled, Experimental Procedures). Classification ratios greater than one indicate that correlations improve encoding while classification ratios less than one indicate that correlations impair encoding. Consistent with theoretical predictions (Averbeck et al., 2006; Gu et al., 2011), we find that the effect of correlation on encoding depends strongly on the relationship between signal and noise correlations (Figure 6A). For neuron pairs with positive noise correlations, the classification ratio is larger for pairs with negative signal correlations than for pairs with positive signal correlations (t test, $p = 2.2 \times 10^{-10}$; Figure 6A). Conversely, for neuron pairs with negative noise correlations, the classification ratio is larger for pairs with positive signal correlations than for pairs with negative signal correlations (t test, $p = 0.044$; Figure 6A). Thus, our data demonstrate that the observed correlations improve encoding when signal and noise correlations are of opposite sign and impair encoding when signal and noise correlations are of the same sign (two-way ANOVA interaction term, $p = 1.8 \times 10^{-8}$).

This observation, combined with our finding that task-relevance alters the relationship between signal and noise

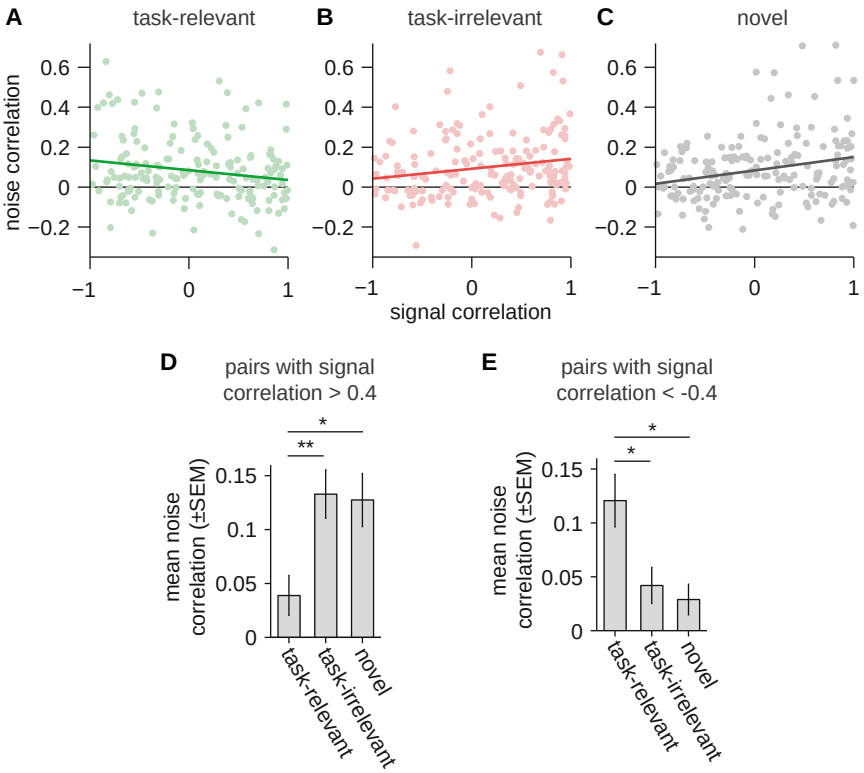

Figure 4. Task Relevance Alters the Relationship between Signal and Noise Correlations

(A–C) Scatter plots of signal and noise correlations for all pairs for responses to each class of motifs. Each point denotes one neuron pair. Black line depicts zero noise correlation; heavy colored lines are linear regression lines.

(D) Mean (±SEM) noise correlations for all pairs with signal correlation greater than 0.4 for relevant (n = 60 pairs), irrelevant (n = 72 pairs), and novel (n = 59 pairs) motifs.

(E) Mean (±SEM) noise correlations for all pairs with signal correlation less than −0.4 for relevant (n = 53 pairs), irrelevant (n = 39 pairs), and novel (n = 51 pairs) motifs. Wilcoxon rank-sum test, *p < 0.05, **p < 0.002.

correlations (Figure 4), suggests that task relevance may improve motif encoding among neuron pairs in CLM. To test this idea, we computed the mean classification ratio of all pairs for task-relevant, task-irrelevant, and novel motifs. We indeed found that task relevant motifs exhibit a higher classification ratio than the task-irrelevant or novel motifs (Friedman test, p = 0.018; Figure 6B), consistent with our observations of the correlation structure. Even in pairs of neurons, therefore, we find that the learning-dependent change in the correlation structure directly yields improved sensory coding of motifs.

Role of Correlation Structure on Encoding by Larger Populations

How does the correlation-dependent encoding in pairs of neurons translate into encoding by larger populations? Prior theoretical (Gu et al., 2011; Zohary et al., 1994) and experimental (Cohen and Maunsell, 2009) studies have demonstrated that even small changes in average noise correlations can have very large effects on neural encoding in populations as small as only 10 or 20 neurons. Furthermore, in larger populations, noise correlations can have an impact on encoding that is substantially greater than that from mean firing rates (Cohen and Maunsell, 2009; Mitchell et al., 2009). We thus asked whether the changes in correlations that we see in pairs of neurons yield larger effects in larger populations of neurons. Our data set makes it possible to test this explicitly because many of the pairs in our data set were actually recorded as sets of up to eight neurons.

Consistent with the idea that larger population sizes allow improved coding from a higher dimensionality of response space, we found that classification performance increased with population size for all classes of motifs (Figure 7A). Importantly, classification performance increased at a faster rate for task-relevant motifs than for either task-irrelevant or novel motifs (solid lines in Figure 7A). This observation could result either from learning-dependent changes to the underlying single-neuron response properties or from the changes to the correlation structure described above. To distinguish these two sources of increased performance, we compared the classification performance without correlations (i.e., with trials shuffled, which does not alter individual neuron responses) to that with correlations intact. Shuffling trials considerably reduces classification performance for task-relevant motifs, but not to the level of task-irrelevant or novel motifs (dashed lines in Figure 7A). This suggests that the enhanced coding fidelity for task-relevant motifs results both from single-neuron response properties and from correlations between neurons.

To isolate the effects of correlations on coding, we computed the classification ratio for each class of motif and for each population size (Experimental Procedures). We find that the effect of correlations on coding of task-irrelevant and novel motifs is small and does not depend on population size; in contrast, for task-relevant motifs, a modest effect in pairs of neurons grows considerably with population size (Figure 7B). Associative learning, therefore, can alter neural correlations in a way that dramatically improves sensory encoding in large neural populations but only for signals that are behaviorally relevant.

DISCUSSION

Associative learning inverts the relationship between signal correlation and noise correlation in pairs of CLM neurons. This inversion enhances population encoding of motifs associated with learned behavioral goals. Rather than affecting the overall magnitude of noise correlations, associative learning changes how noise correlations depend on signal correlations.

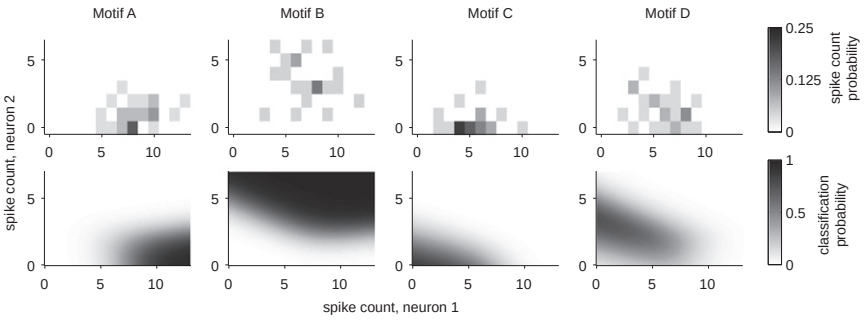

Figure 5. Motif Classification Using Multinomial Logistic Regression

Top row shows the two-dimensional response distributions of a pair of neurons for each of the four task-relevant motifs. Probability of observing each pair of responses is represented in grayscale. Bottom row depicts the optimal decoder model fit to the data. The probability of classifying a pair of responses from these two neurons as each of the four motifs is represented in grayscale. For this pair of neurons, the average probability of correct motif classification was 0.507.

Noise correlations are widely reported to covary with signal correlations (Cohen and Maunsell, 2009; Cohen and Newsome, 2008; Gu et al., 2011; Gutnisky and Dragoi, 2008; Hofer et al., 2011; Kohn and Smith, 2005). Although this relationship depends on cell type (Constantinidis and Goldman-Rakic, 2002; Hofer et al., 2011; Lee et al., 1998) and on behavioral context (Cohen and Newsome, 2008; Lee et al., 1998), it is generally positive. Positive relationships impair population encoding because common noise among similarly tuned neurons cannot be removed by pooling (Averbeck et al., 2006). In contrast, negative relationships can improve population coding because common noise among dissimilarly tuned neurons can be subtracted, which strengthens the signal while dissipating the noise.

To our knowledge, a negative relationship between signal and noise correlations has not previously been demonstrated. Theoretical studies, however, have predicted that changes to the sign of this relationship might underlie cognitive functions such as attention or learning (Oram et al., 1998). We provide experimental evidence to support this prediction: associative learning inverts this relationship, substantially enhancing population encoding of learned motifs. Importantly, our results show that learning enhances the population code in two ways: by changing single-neuron responses and by changing interneuronal correlations. Even with shuffled trials, we find that neural populations better distinguish between task-relevant motifs than between task-irrelevant or novel motifs (Figure 7A), demonstrating the plasticity of response properties of individual neurons. However, with correlations taken into account, the same neural populations discriminate between task-relevant motifs even better, without affecting discrimination of task-irrelevant or novel motifs (Figure 7). Thus, the relationship between the signal and the noise correlations acts in a stimulus-specific way to enhance the representation of only those signals made relevant by prior learning.

Neural Plasticity and Associative Learning

Psychologists have long recognized the wide range of associative relationships that can change as a result of learning—associations between different stimuli, between stimuli and responses and/or reward, and combinations of all these. Neuroscientists, for their part, have been relatively slow to explore these varied relationships. Prior studies of sensory plasticity in single auditory cortical neurons have pointed to the importance of actively engaging stimuli in the context of rewarded behaviors (Blake et al., 2006; Thompson and Gentner, 2010). This, in turn, supports the inference that associations between stimuli and reward (facilitated perhaps by attention or other cognitive processes) drive the observed sensory plasticity (Blake et al., 2006).

The design of the present training allows us to test this inference directly, because the role of reinforcement can be dissociated from the behavioral responses that lead to reinforcement. All the motifs used during training were heard equally often in the context of the task and paired equally with reinforcement, but only the task-relevant motifs signaled the correct behavioral response on each trial. Thus, any effect of learning mediated directly by reinforcement should apply to all of the training motifs. What we observe, however, is very different. Only the task-relevant motifs—those that birds learned to associate with a particular pecking location—elicited neural population activity with a negative relationship between signal and noise correlations. In contrast, task-irrelevant motifs—those that birds never learned to associate with a particular pecking location—elicited neural population activity with a positive correlation relationship indistinguishable from that elicited by novel motifs that birds never heard while awake. Thus, learning-dependent changes in the interneuronal correlation patterns depend on associations formed between stimuli and behavior, rather than experience, familiarity, or reward contingency. Reward is crucial, of course, in controlling responses (Herrnstein, 1961), but the role of the stimulus is to signal the appropriate action required to obtain that reward. In this context, which psychologists refer to as occasion setting (Schmajuk and Holland, 1998), we suggest that the neural representation of motifs in CLM is less a sensory trace than a predictive mapping of the learned behavioral response.

Comparison with Previous Studies

Understanding the CLM population representation as a product of sensory-motor learning may help to interpret our results in the context of other work involving different forms of learning. Indeed, a recent study found that perceptual learning did not alter the slope of the relationship between signal and noise correlations for neurons in the primate medial superior temporal area (Gu et al., 2011). This study differed from ours in multiple ways (e.g., species, brain region, and sensory modality) that make direct comparisons difficult, but one important difference is in the type of learning. Perceptual learning targets sensory acuity, forcing animals to resolve fine differences between previously indistinguishable low-dimensional stimuli. In contrast, the complex stimuli used in our study are likely to be easily

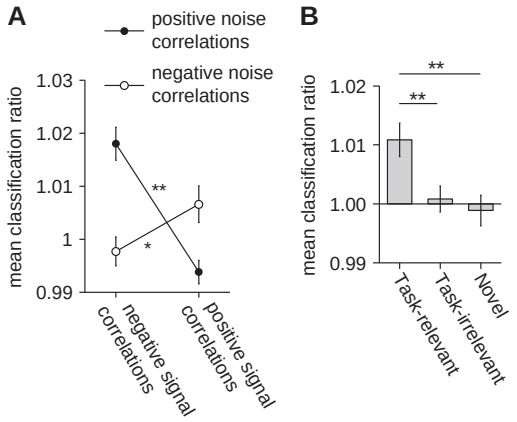

Figure 6. Relationship between Motif Decoding and Interneuronal Correlations

(A) The relationship between signal and noise correlations predicts the effect that correlations have on motif decoding. Neuron pairs are binned according to their signal and noise correlations; the relative performance of the decoder (ratio of correct classifications to correct classifications with trials shuffled) is plotted on the ordinate. The mean classification ratio is greater than 1 for pairs with negative signal and positive noise correlations and for pairs with positive signal and negative noise correlations. In contrast, the ratio is less than 1 for pairs with negative signal and negative noise correlations and for pairs with positive signal and positive noise correlations. t tests: $^{*}p < 0.05$; $^{**}p < 1 \times 10^{-9}$. (B) On average, the classification ratio for task-relevant motifs is greater than task-irrelevant motifs and novel motifs. t tests: $^{**}p < 0.01$. Data are reported as mean ± SEM. See also Figure S1.

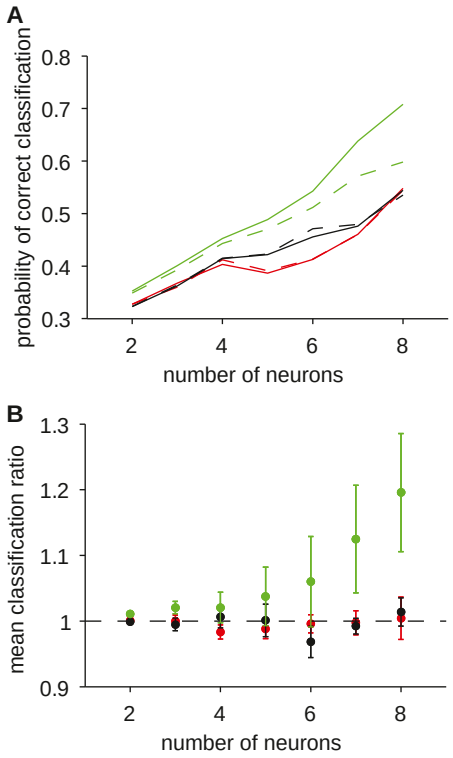

Figure 7. Effects of Correlations on Decoding of Task-Relevant Motifs Increase with Population Size

(A) Mean classification performance for populations ranging in size from two to eight neurons. Task-relevant, task-irrelevant, and novel motifs are denoted by green, red, and black, respectively. Solid lines show performance with correlations intact; dashed lines show performance without correlations (trials shuffled). Although increasing the population size generally increases classification performance even without correlations (dashed lines), the correlations have increasingly larger influences on the task-relevant motifs.

(B) Mean (±SEM) classification ratio for each class of motifs as a function of population size. Only the task-relevant motifs exhibit a positive relationship between population size and classification ratio. See also Figure S4.

discriminated from one another without intense training, and the predominant learning occurs in the strength of the association between stimuli and response (Figure 1E). Although some of the mechanisms of these two forms of learning may overlap (e.g., Law and Gold, 2008), the observed differences suggest that other mechanisms may be unique due to differing functional requirements.

A methodological difference between our study and that of Gu et al. (2011) is that we used anesthetized animals while Gu and colleagues used awake animals. We think it is highly unlikely that anesthesia could account for the differences between our results for two reasons. First, while noise correlations can, in principle, be influenced by fluctuations in the depth of anesthesia, they can also be influenced by internal factors in awake animals, such as fluctuations in alertness, attention, or motivation. Consistently, differences in noise correlation measurements between studies may be more likely to result from factors such as differences in the mean firing rate or the size of the temporal analysis window, than by differences in anesthetic (Cohen and Kohn, 2011), although more data are necessary. Second, and most important, even if anesthesia did influence the correlations we measured, this influence would apply equally to all three motif classes because our presentation of motifs during electrophysiology was fully randomized and all of our comparisons are within pairs (or populations) of neurons. In addition, we note that song-evoked responses in the starling forebrain are qualitatively quite similar between anesthetized and unanesthetized states, although some quantitative differences exist (Knudsen and Gentner, 2013; Meliza et al., 2010).

The most parsimonious explanation for our results, however, is that learning induces long-lasting changes to the neural circuitry that remain after training has concluded, even under anesthesia.

Possible Biophysical Mechanisms and Circuit Functions

The commonly observed positive relationship between signal and noise correlations is often accounted for by shared inputs that provide both signal and noise. In primary visual cortex, neurons that share receptive field properties are more likely to share thalamocortical afferent inputs (Alonso et al., 2001; Michalski et al., 1983). But a negative relationship would require a decorrelation of similarly tuned neurons and an increase in the correlation of dissimilarly tuned neurons. Simple feedforward inhibition circuits could, in theory, support both requirements. Recent modeling work has demonstrated that correlated noise in excitatory and inhibitory input can cancel each other, leading to decorrelated network states (Middleton et al., 2012; Renart et al., 2010). Complementary circuitry in which only excitatory inputs are correlated could preserve correlated noise in

dissimilarly tuned neurons (Figure S4). The plausibility of such circuitry is supported by a recent network model showing that learning-driven modulations to feedforward synaptic weights yields decreased noise correlations for similarly tuned cortical neurons and increased noise correlations for dissimilarly tuned neurons (Bejjanki et al., 2011). In addition, experimental work linking intracortical synaptic connectivity to noise correlations (Ko et al., 2011) suggests that local circuit mechanisms may also contribute to the relationship between signal and noise correlation. In our case, because the same population of CLM neurons can represent different stimuli using qualitatively different correlation structures, the circuitry (local or extrinsic) that controls the correlation structure must be flexible on a short timescale. Further experiments will be necessary to elucidate the circuitry that yields this stimulus-specific flexibility.

Our results also provide initial evidence that flexibility in the relationship between signal and noise correlations is cell type specific. For example, the correlations in the pooled population of NS-NS and NS-WS pairs did not exhibit the same effects as the WS-WS pairs (Figure S3C). This suggests that the plasticity of the correlation structure primarily exists within WS (putative excitatory) neurons, although more data are necessary. One possible explanation of this is that WS-WS pairs receive less common input than NS-NS pairs or NS-WS pairs, and thus their interneuronal correlations are most susceptible to modulation by local circuitry. Such an idea is supported by findings that noise correlations are higher among inhibitory interneurons than excitatory neurons (Constantinidis and Goldman-Rakic, 2002) and that the slope of the relationship between signal and noise correlations is much shallower for pairs of excitatory pyramidal neurons than for pairs of inhibitory parvalbumin-expressing neurons in primary visual cortex (Hofer et al., 2011).

Our results suggest that large neural populations in CLM better discriminate differences between task-relevant motifs than between task-irrelevant or novel motifs. CLM provides auditory information directly to HVC (Bauer et al., 2008), a region known to control song production (Long and Fee, 2008; Nottebohm et al., 1976). The enhanced population coding in CLM may influence the flow of auditory feedback into HVC during juvenile song learning and for adult song maintenance, two behaviors critical for survival, by selectively emphasizing the most important motifs. This possibility could be explored by chronically recording from CLM populations during these behaviors.

Conclusion

We demonstrate that the relationship between signal and noise correlations is a target of learning-dependent plasticity that can substantially enhance the representation of specific stimuli. Moreover, the effects of this plasticity on neural coding increase substantially with population size, becoming quite considerable once the population reaches five to six neurons. Our results support the longstanding hypothesis that these activity patterns underlie behaviorally relevant discrimination of sensory signals (Oram et al., 1998). The population correlation structure carries biologically significant information, independent of the activity in single neurons, crucial to the transformation of purely sensory codes into neural signals that ultimately drive learned behavior.

EXPERIMENTAL PROCEDURES

All procedures were carried out in accordance with the guidelines of the Institutional Animal Care and Use Committee at the University of California, San Diego.

Stimuli

We constructed all stimuli from 12 motifs (stereotyped multinote elements of natural starling song) recorded from the song repertoires of three adult European starlings. Motifs (565–957 ms long) were grouped into three sets: four motifs (A, B, C, and D) were labeled "task relevant," four motifs (E, F, G, and H) were labeled "task irrelevant," and four motifs (I, J, K, and L) were labeled "novel" (Figures S2A–S2L). For behavioral training, we presented a sequential pair of motifs for each trial (Figure 1B). Each pair contained exactly one relevant and one irrelevant motif, in either order, separated by a 20 ms silent gap (e.g., Figure 1C). This yielded 32 stimuli; the 16 containing motifs A or B were used as "left" stimuli and the 16 containing motifs C or D were used as "right" stimuli. All irrelevant motifs occurred with equal probability in both left and right stimuli. Novel motifs were never presented during training. To ensure that learning effects were not due to intrinsic acoustic differences between motifs, we counterbalanced motif assignment to task-relevant, task-irrelevant, and novel categories across birds. During neural recording sessions, we presented each of the 12 motifs in isolation (i.e., not paired).

Behavioral Training

Nine wild-caught adult European starlings (*Sturnus vulgaris*) were trained using a two-alternative choice operant conditioning paradigm (Gentner and Margoliash, 2003) to recognize the left and right stimuli described above and to associate them with food reward. Prior to training, none of the subjects had any exposure to these stimuli. All training took place inside a sound attenuation chamber with an operant response panel (Figure 1A). Birds initiated trials by inserting their beak into the center port of the response panel to start playback of one of the 32 stimuli from the speaker inside the chamber. After playback, birds had 2 s to respond by pecking in either the left or the right port. Incorrect responses (pecking the left port after a right stimulus or the right port after a left stimulus) were punished by extinguishing the lights and prohibiting trial initiation for 10–90 s. Correct responses were rewarded by a 2 s access to food on a fixed ratio reinforcement schedule. The number of consecutive correct trials required for reward was gradually increased over time from 1 to 5. A secondary reinforcer (flashing of LEDs on the response panel) was used on correct trials when the food reward was not delivered. Incorrect responses reset the fixed ratio counter. The fixed-ratio reinforcement ensured that subjects did not systematically ignore any stimuli and allowed for all stimuli to be presented an equal number of times. At the end of training, starlings were presented with random nondifferentially reinforced (with secondary reinforcer only) probe stimuli consisting of each of the eight training motifs in isolation (i.e., not paired) to obtain behavioral confirmation that all four task-relevant motifs were recognized (Figures 1E and 1F). Probe stimuli were randomly interleaved on 8%–20% of all trials during these probe sessions.

Electrophysiology

Starlings were anesthetized with urethane (20% by volume, 7–8 ml/kg) and head fixed to a stereotactic apparatus inside a sound-attenuating chamber. A small craniotomy was made dorsal to CLM, and multichannel silicon electrode arrays (177 μm^2 electrode surface area, 50 μm spacing, 1 × 16 and 1 × 32 electrode layout; NeuroNexus technologies) were inserted into CLM. For some subjects, only the 1 × 32 array was used (Figure S2M). Motif stimuli were presented free field from a speaker 30 cm from the bird at sound pressure levels matched to those during behavioral training (mean, 65 dB; peak, 96 dB). Electrode arrays were advanced while presenting the 12 motif stimuli until two or more auditory single units were isolated. Once single units were isolated, all 12 single motifs and the set of training motif pairs were presented pseudorandomly in blocks while the extracellular electrical activity was amplified (5,000 × gain; AM Systems), filtered (high pass, 300 Hz; low pass, 3–5 kHz), sampled (20 kHz), and saved digitally for offline analysis (Spike2; Cambridge Electronic Design).

Data Analysis

Putative action potentials in the recorded voltage traces were identified by amplitude and sorted into single units with principal components analysis on waveform shape using Spike2 software (Cambridge Electronic Design). Only large amplitude spike waveforms that formed a clear cluster in principal component space and that had very few refractory period violations were considered to be single units. In our sample, 99.3% (133/134) of all (Wide Spiking+Narrow Spiking) neurons had no refractory violations (interspike intervals of less than 1 ms) and one neuron had a single violation, which accounted for less than 0.005% of all measured ISIs for that neuron. Since presentation of task-relevant, task-irrelevant, and novel motifs was temporally interleaved, none of the effects reported here can be due to changes in neuron isolation or changes in anesthetic state. Because the recording sites on each multichannel array were only 50 μm apart, stereotrode sorts were used to further improve spike-sorting quality. All but one of the WS neuron pairs analyzed here were recorded from different electrode channels on the multichannel arrays. Omitting the one pair recorded from the same channel does not alter the main results. Only neurons that were driven by at least one motif were used in subsequent analyses.

All further analysis was performed using custom-written MATLAB (MathWorks) software. Spike shape classification was performed using spike width, the time from the initial trough to the subsequent peak. During recording, data from some birds were low-pass filtered at 3 kHz and others at 5 kHz. Because differences in this cutoff frequency can alter the spike shape (Vigneswaran et al., 2011), we applied a first-order low-pass Butterworth filter with cutoff frequency at 3 kHz to all spike shapes to equalize these differences. All mean spike waveforms were cubic spline interpolated to a 2.5 μs sampling interval. The filtering slightly increased the spike widths of all neurons. Thus, our threshold of 425 μs between WS and NS neurons is toward the upper end of the distribution of thresholds used in previous reports (Mitchell et al., 2007; Vigneswaran et al., 2011) but is conservative.

Because connectivity and correlation within neural populations depends on cell type (Constantinidis and Goldman-Rakic, 2002; Hofer et al., 2011; Lee et al., 1998), we divided our data set into wide spiking (WS) and narrow spiking (NS) neurons on the basis of action potential width (trough-to-peak duration; Figures S2Q–S2S) (Barthó et al., 2004; Mitchell et al., 2007). The distribution of action potential widths is bimodal (Hartigan's dip test, $p = 0.041$; Figure S2S) (Hartigan and Hartigan, 1985; Mitchell et al., 2007). Based on network interactions and correlations with intracellular properties, previous studies have established that WS and NS neurons correspond to excitatory principal neurons and inhibitory interneurons, respectively (Barthó et al., 2004; Harris et al., 2000; Tamura et al., 2004). Consistent with these classifications, our sample of NS neurons ($n = 36$) elicited significantly higher spontaneous firing rates (4.61 ± 0.76 Hz) than our sample of WS neurons ($n = 98$; 1.80 ± 0.21 Hz; Wilcoxon rank-sum test, $p = 1.03 \times 10^{-4}$). Because our sample of simultaneously recorded pairs of NS neurons was relatively small ($n = 17$ pairs), we focus our population analysis on pairs of WS neurons ($n = 176$ pairs from 6 birds).

Signal correlations were computed for each pair of neurons as the Pearson product-moment correlation coefficient between the mean (averaged over trials) firing rates to the four motifs within the task-relevant, task-irrelevant, and novel classes. Noise correlations were computed for each individual motif for each pair across trials then averaged for all motifs within a class. Because motifs were variable in duration (range: 565–957 ms, mean: 756 ms) and the size of the analysis window can affect measured correlation values (Cohen and Kohn, 2011), we use only the first 565 ms (the minimum motif duration) of each response in the analyses reported here. We note, however, that reanalyzing our data using the full duration of each motif yields similar patterns of correlations.

Motif discrimination ability was assessed using a predictive multinomial logistic regression model (Long, 1997). Logistic regression models the relationship between a set of continuous predictor variables (here the firing rates of multiple simultaneously recorded neurons) and a categorical output variable (here the motif identity), by representing the categorical variable as a probability. The multinomial logistic regression model is a generalization of logistic regression to more than two categorical variables. This is a minimal model consistent with the chosen set of observations (in this case the firing rate of neurons) that does not make any additional assumptions, and in particular does not assume that variations in the neural responses follow a Gaussian distribution (Graf et al., 2011). A similar approach is also used for conditional random fields in machine learning research (Lafferty et al., 2001) and for maximum noise entropy models in neuroscience (Fitzgerald et al., 2011). Given a set of neural responses X_i, the classifier produces the probability that this was caused by motif j as:

$$\Pr(Motif = j) = \frac{\exp\left(\sum_i \beta_{ji} X_i\right)}{\sum_{k=1}^{K} \exp\left(\sum_i \beta_{ki} X_i\right)}$$

where each β_{ji} is a set of coefficients fitted to the model by maximum likelihood estimation, with the index j (or k in the above equation) describing one of the possible K classification outputs and index i enumerating the neural responses among the n neurons in the population. This technique provides a convenient and mathematically optimal way to quantify how well a set of neurons can discriminate between multiple motifs. To find the coefficients, we used the MATLAB function *mnrfit*, and to find the probabilities of each motif from the model, we used the MATLAB function *mnrval* (Statistics toolbox, version 7.3, release 2010a). To avoid overfitting, we fit the model to 75% of the trials for each population and predicted motif identity for the remaining 25% of trials. This procedure was then repeated four times, to ensure that all trials in each population received a prediction. For each trial, the model predicted the probability that the set of firing rates resulted from each of the four motifs. To compute the probability of correct classification, the probability of predicting the correct motif was averaged over all trials and all motifs for each population.

Because we were interested in the net effect of correlations on motif discrimination, we needed an estimate of discrimination performance in the absence of correlations. To do this, we shuffled the trial ordering of each neuron in each data set, refit the model, and recomputed the probability of correct classification. This destroys trial-by-trial correlations (i.e., noise correlations), while leaving mean firing rates and signal correlations completely unaltered. To ensure that random correlations introduced by this process did not affect our analysis, we repeated the shuffling process 50 times and used the average probability of correct classification from these shuffles. We then computed the classification ratio as the probability of correct classification divided by the shuffled probability of correct classification.

To assess the relationship between the number of neurons and the probability of correct classification, we considered each simultaneously recorded population of size greater than or equal to 3 (range: 3 to 8 neurons; $n = 12$ populations). For a population of size k, we considered all possible subsets of the population of size 2 through $k - 1$. To avoid oversampling of the larger populations, we averaged the classification values for all subsets of a given size to a single data point. Thus for each population of size k, we had a single value for the probability of correct classification for subpopulations ranging from 2 to k. We then averaged the values for each subpopulation size together to generate the values in Figure 7.

Statistical Analysis

All data were tested for normality using the Lilliefors test evaluated at $p < 0.05$. When available, nonparametric tests were used when data were not normal. Central tendencies are reported as means ± SEM, except where noted.

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures and Supplemental Experimental Procedures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2013.02.023.

ACKNOWLEDGMENTS

We thank W. Kristan and D. Margoliash for comments on an earlier version of this manuscript and the members of the Gentner and Sharpee laboratories for conversations. This work was supported by a grant from the NIH (R01DC008358) to T.Q.G., grants from the NIH (R01EY019493), the Alfred P.

Sloan Foundation, the Searle Scholars Program, the Center for Theoretical Biological Physics (NSF), the W.M. Keck Foundation, the Ray Thomas Edwards Career award in Biomedical Sciences, and the McKnight Scholar Award to T.O.S., and by an NSF Graduate Research Fellowship and an Institute for Neural Computation (UCSD) Fellowship to J.M.J. J.M.J., T.O.S., and T.Q.G. designed research. J.M.J. performed research. J.M.J., T.O.S., and T.Q.G. analyzed data and wrote the paper.

Accepted: February 25, 2013
Published: April 24, 2013

REFERENCES

Abbott, L.F., and Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. Neural Comput. *11*, 91–101.

Alonso, J.M., Usrey, W.M., and Reid, R.C. (2001). Rules of connectivity between geniculate cells and simple cells in cat primary visual cortex. J. Neurosci. *21*, 4002–4015.

Averbeck, B.B., Latham, P.E., and Pouget, A. (2006). Neural correlations, population coding and computation. Nat. Rev. Neurosci. *7*, 358–366.

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

Barthó, P., Hirase, H., Monconduit, L., Zugaro, M., Harris, K.D., and Buzsáki, G. (2004). Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. J. Neurophysiol. *92*, 600–608.

Bauer, E.E., Coleman, M.J., Roberts, T.F., Roy, A., Prather, J.F., and Mooney, R. (2008). A synaptic basis for auditory-vocal integration in the songbird. J. Neurosci. *28*, 1509–1522.

Bejjanki, V.R., Beck, J.M., Lu, Z.L., and Pouget, A. (2011). Perceptual learning as improved probabilistic inference in early sensory areas. Nat. Neurosci. *14*, 642–648.

Blake, D.T., Strata, F., Churchland, A.K., and Merzenich, M.M. (2002). Neural correlates of instrumental learning in primary auditory cortex. Proc. Natl. Acad. Sci. USA *99*, 10114–10119.

Blake, D.T., Heiser, M.A., Caywood, M., and Merzenich, M.M. (2006). Experience-dependent adult cortical plasticity requires cognitive association between sensation and reward. Neuron *52*, 371–381.

Cafaro, J., and Rieke, F. (2010). Noise correlations improve response fidelity and stimulus encoding. Nature *468*, 964–967.

Cohen, M.R., and Kohn, A. (2011). Measuring and interpreting neuronal correlations. Nat. Neurosci. *14*, 811–819.

Cohen, M.R., and Maunsell, J.H. (2009). Attention improves performance primarily by reducing interneuronal correlations. Nat. Neurosci. *12*, 1594–1600.

Cohen, M.R., and Newsome, W.T. (2008). Context-dependent changes in functional circuitry in visual area MT. Neuron *60*, 162–173.

Constantinidis, C., and Goldman-Rakic, P.S. (2002). Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. J. Neurophysiol. *88*, 3487–3497.

Eens, M. (1997). Understanding the complex song of the European starling: an integrated ethological approach. Adv. Stud. Behav. *26*, 355–434.

Fitzgerald, J.D., Sincich, L.C., and Sharpee, T.O. (2011). Minimal models of multidimensional computations. PLoS Comput. Biol. *7*, e1001111.

Gentner, T.Q., and Hulse, S.H. (2000). Perceptual classification based on the component structure of song in European starlings. J. Acoust. Soc. Am. *107*, 3369–3381.

Gentner, T.Q., and Margoliash, D. (2003). Neuronal populations and single cells representing learned auditory objects. Nature *424*, 669–674.

Graf, A.B., Kohn, A., Jazayeri, M., and Movshon, J.A. (2011). Decoding the activity of neuronal populations in macaque primary visual cortex. Nat. Neurosci. *14*, 239–245.

Gu, Y., Liu, S., Fetsch, C.R., Yang, Y., Fok, S., Sunkara, A., DeAngelis, G.C., and Angelaki, D.E. (2011). Perceptual learning reduces interneuronal correlations in macaque visual cortex. Neuron *71*, 750–761.

Gutnisky, D.A., and Dragoi, V. (2008). Adaptive coding of visual information in neural populations. Nature *452*, 220–224.

Harris, K.D., Henze, D.A., Csicsvari, J., Hirase, H., and Buzsáki, G. (2000). Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J. Neurophysiol. *84*, 401–414.

Hartigan, J.A., and Hartigan, P.M. (1985). The dip test of unimodality. Ann. Stat. *13*, 70–84.

Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. J. Exp. Anal. Behav. *4*, 267–272.

Hofer, S.B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N.A., and Mrsic-Flogel, T.D. (2011). Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. Nat. Neurosci. *14*, 1045–1052.

Huber, D., Petreanu, L., Ghitani, N., Ranade, S., Hromádka, T., Mainen, Z., and Svoboda, K. (2008). Sparse optical microstimulation in barrel cortex drives learned behaviour in freely moving mice. Nature *451*, 61–64.

Jeanne, J.M., Thompson, J.V., Sharpee, T.O., and Gentner, T.Q. (2011). Emergence of learned categorical representations within an auditory forebrain circuit. J. Neurosci. *31*, 2595–2606.

Knudsen, D.P., and Gentner, T.Q. (2013). Active recognition enhances the representation of behaviorally relevant information in single auditory forebrain neurons. J. Neurophysiol. Published online January 9, 2013. http://dx.doi.org/10.1152/jn.00461.2012.

Ko, H., Hofer, S.B., Pichler, B., Buchanan, K.A., Sjöström, P.J., and Mrsic-Flogel, T.D. (2011). Functional specificity of local synaptic connections in neocortical networks. Nature *473*, 87–91.

Kohn, A., and Smith, M.A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. J. Neurosci. *25*, 3661–3673.

Lafferty, J., McCallum, A., and Pereira, F. (2001). Conditional random fields: probabilistic models for segmenting and labeling sequence data. C.E. Brodley and A.P. Danyluk, eds. Proceedings of the Eighteenth International Conference on Machine Learning (ICML '01), 282–289.

Law, C.T., and Gold, J.I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. Nat. Neurosci. *11*, 505–513.

Lee, D., Port, N.L., Kruse, W., and Georgopoulos, A.P. (1998). Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. J. Neurosci. *18*, 1161–1170.

Long, J.S. (1997). Regression Models for Categorical and Limited Dependent Variables (Thousand Oaks, CA: SAGE Publications).

Long, M.A., and Fee, M.S. (2008). Using temperature to analyse temporal dynamics in the songbird motor pathway. Nature *456*, 189–194.

Meliza, C.D., and Margoliash, D. (2012). Emergence of selectivity and tolerance in the avian auditory cortex. J. Neurosci. *32*, 15158–15168.

Meliza, C.D., Chi, Z., and Margoliash, D. (2010). Representations of conspecific song by starling secondary forebrain auditory neurons: toward a hierarchical framework. J. Neurophysiol. *103*, 1195–1208.

Michalski, A., Gerstein, G.L., Czarkowska, J., and Tarnecki, R. (1983). Interactions between cat striate cortex neurons. Exp. Brain Res. *51*, 97–107.

Middleton, J.W., Omar, C., Doiron, B., and Simons, D.J. (2012). Neural correlation is stimulus modulated by feedforward inhibitory circuitry. J. Neurosci. *32*, 506–518.

Mitchell, J.F., Sundberg, K.A., and Reynolds, J.H. (2007). Differential attention-dependent response modulation across cell classes in macaque visual area V4. Neuron *55*, 131–141.

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

Nottebohm, F., Stokes, T.M., and Leonard, C.M. (1976). Central control of song in the canary, Serinus canarius. J. Comp. Neurol. *165*, 457–486.

Oram, M.W., Földiák, P., Perrett, D.I., and Sengpiel, F. (1998). The 'Ideal Homunculus': decoding neural population signals. Trends Neurosci. *21*, 259–265.

Reed, A., Riley, J., Carraway, R., Carrasco, A., Perez, C., Jakkamsetti, V., and Kilgard, M.P. (2011). Cortical map plasticity improves learning but is not necessary for improved performance. Neuron *70*, 121–131.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K.D. (2010). The asynchronous state in cortical circuits. Science *327*, 587–590.

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

Schmajuk, N.A., and Holland, P.C. (1998). Occasion Setting: Associative Learning and Cognition in Animals (Washington, DC: American Psychological Association).

Shadlen, M.N., and Newsome, W.T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. J. Neurosci. *18*, 3870–3896.

Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J. Neurosci. *16*, 1486–1510.

Tamura, H., Kaneko, H., Kawasaki, K., and Fujita, I. (2004). Presumed inhibitory neurons in the macaque inferior temporal cortex: visual response properties and functional interactions with adjacent neurons. J. Neurophysiol. *91*, 2782–2796.

Thompson, J.V., and Gentner, T.Q. (2010). Song recognition learning and stimulus-specific weakening of neural responses in the avian auditory forebrain. J. Neurophysiol. *103*, 1785–1797.

Thompson, J.V., Jeanne, J.M., and Gentner, T.Q. (2013). Local inhibition modulates learning-dependent song encoding in the songbird auditory cortex. J. Neurophysiol. *109*, 721–733.

Vigneswaran, G., Kraskov, A., and Lemon, R.N. (2011). Large identified pyramidal cells in macaque motor and premotor cortex exhibit "thin spikes": implications for cell type classification. J. Neurosci. *31*, 14235–14242.

Wilke, S.D., and Eurich, C.W. (2002). Representational accuracy of stochastic neural populations. Neural Comput. *14*, 155–189.

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