



Cholinergic shaping of neural correlations

Victor Minces^{a,1}, Lucas Pinto^{b,2}, Yang Dan^b, and Andrea A. Chiba^{a,c}

^aDepartment of Cognitive Science, University of California, San Diego, La Jolla, CA 92093; ^bDivision of Neurobiology, Department of Molecular and Cell Biology, Helen Wills Neuroscience Institute, Howard Hughes Medical Institute, University of California, Berkeley, CA 94720-3370; and ^cProgram in Neurosciences, University of California, San Diego, La Jolla, CA 92093

Edited by Thomas D. Albright, The Salk Institute for Biological Studies, La Jolla, CA, and approved April 21, 2017 (received for review January 10, 2017)

A primary function of the brain is to form representations of the sensory world. Its capacity to do so depends on the relationship between signal correlations, associated with neuronal receptive fields, and noise correlations, associated with neuronal response variability. It was recently shown that the behavioral relevance of sensory stimuli can modify the relationship between signal and noise correlations, presumably increasing the encoding capacity of the brain. In this work, we use data from the visual cortex of the awake mouse watching naturalistic stimuli and show that a similar modification is observed under heightened cholinergic modulation. Increasing cholinergic levels in the cortex through optogenetic stimulation of basal forebrain cholinergic neurons decreases the dependency that is commonly observed between signal and noise correlations. Simulations of correlated neural networks with realistic firing statistics indicate that this change in the correlation structure increases the encoding capacity of the network.

acetylcholine | neural coding | neural correlations | neuromodulation | sensory processing

Cortical network responses are shaped by neuromodulators to efficiently code the sensory world. For example, numerous studies have shown that cortical levels of the neurotransmitter acetylcholine impact behavioral and neural responses to sensory stimuli, improving discrimination performance (1–3) and generally increasing the amplitude of the neural responses (3–10). Although the response amplitude can contribute to the amount of information encoded by individual neurons, the encoding capacity of whole networks can be profoundly shaped by the neural correlations (11).

Activity dependencies across neuronal pairs can be analyzed in terms of the correlation of their total activity (total correlations), in terms of the similarity of their receptive fields (signal correlations), or in terms of the similarity in the neurons' trial-to-trial variability (noise correlations) (12, 13). Whereas previous work analyzed the influence of acetylcholine on total correlations (4) or noise correlations (14), it is not clear how these variables alone might influence neural coding (15). Theoretical research indicates that encoding capacity depends on the details of the correlation structure, analyzed in terms of the relationship between signal and noise correlations (12, 15–18). Consistent with the latter observation, two studies have shown that attention (19) and learning (17) alter the relationship between signal and noise correlations in a manner that is thought to increase encoding capacity. Despite the established importance of the relationship between signal and noise correlations for neural coding, previous work has not analyzed how this relationship is affected by cholinergic modulation.

For the purpose of evaluating the effect of increases in cortical acetylcholine on the correlation structure of the visual cortical network, our current study uses optogenetic stimulation of cortically projecting cholinergic neurons located in the basal forebrain. Specifically, we analyze the effect of cholinergic stimulation on the amplitudes of both signal and noise and on the relationship between signal and noise correlations. We find that increasing cholinergic input to the cortex decreases the slope between signal and noise correlations, in a manner consistent with changes observed following behavioral manipulations (17, 19). To understand the impact of this change in the correlation structure on the capacity

of the network to encode information, we use simulations of correlated neural networks with Poisson statistics. We find that the decrease in the correlations' slope increases the encoding capacity of the network.

Results

In this paper, we present a unique analysis of data from a previously reported experiment (3). The experiment measured the activity of mouse visual cortex neurons to repeated presentations of naturalistic movie sequences with or without concomitant optogenetic stimulation of basal forebrain cholinergic neurons (Fig. 1A). The goal of our work is to analyze the specific effect of cholinergic activity on both signal and noise correlations, as well as on the relationship between the two. Here, the neural signal is defined as the average number of spikes in response to a movie segment, and the neural noise is defined as the residual around the signal (12). This broad definition of signal implies that we are focusing on the analysis of the encoding of whole naturalistic images rather than certain individual features of the image.

We estimated the neural signal corresponding to a given neuron as the number of spikes in a time bin, averaged over repeated presentations of a movie (Fig. 1B; see Methods for precise mathematical definitions). We estimated the noise, in each presentation, as the residual activity around the signal (12, 16). Given that the estimation of the signal and the noise are dependent on the duration of an individual time bin, all of the data reported here are calculated for several bin sizes, corresponding to divisors of the total movie length. Once we estimated the signal and the noise, we calculated the signal-to-noise ratios, the signal correlations, and the noise correlations (Methods). We calculated these values from the activity of 113 visual cortex neurons, and 793 neural pairs, from nine animals (Methods) and compared the values across two conditions (with or without optogenetic cholinergic stimulation).

Significance

The capacity of a network of neurons to represent the sensory world depends not only on the way individual neurons respond to sensory stimuli but also on the similarity of activity across neurons comprising the network. This similarity can be quantified as the correlation of activity across neuronal pairs, or neural correlations. This work presents a unique finding, demonstrating that a neuromodulator, in this case acetylcholine, can produce specific changes in the correlation structure of the cortical network that ultimately increase the encoding capacity of the network.

Author contributions: V.M., L.P., Y.D., and A.A.C. designed research; V.M., L.P., Y.D., and A.A.C. performed research; V.M., L.P., Y.D., and A.A.C. analyzed data; and V.M. and A.A.C. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. Email: vminces@mail.ucsd.edu.

²Present address: Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08544.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1621493114/-DCSupplemental.

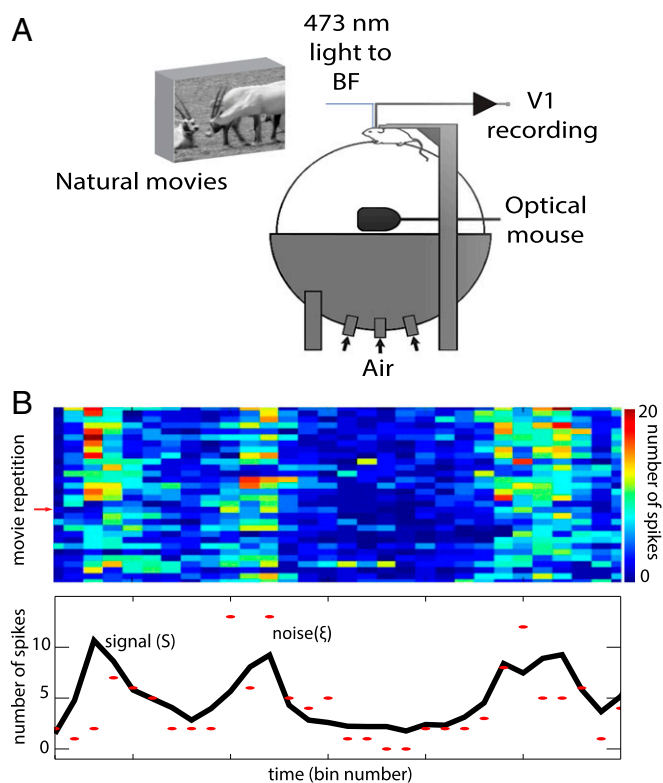


Fig. 1. Experimental setup, signal and noise estimation. (A) Experimental setup, adapted from ref. 3. (B) We presented 30 repetitions of a movie sequence and calculated the spike count in consecutive time bins (Top). The signal (black line, Bottom) is estimated as the average number of spikes over all repetitions. The red markers represent the spike counts corresponding to a given repetition (marked with a red arrow on the Top). The noise is estimated as the difference between the signal and the spike count. From the estimations of the signal and the noise, we calculate the signal and noise correlations of neuronal pairs (see Methods for a precise mathematical definition of signal, noise, amplitudes, and correlations).

We define the signal-to-noise ratio as $SN = \sigma_S^2 / \sigma_\xi^2$, where σ_S^2 is the variance of the signal, or signal amplitude, and σ_ξ^2 is the variance of the noise, or noise amplitude. This definition of signal to noise can be understood as the ratio of the dynamic range of the neuronal responses over the amplitude of the noise. It is directly linked to the information capacity of a single neuron (20) and has been proposed as an appropriate measure to quantify neural coding (21). We found that optogenetic stimulation of cholinergic neurons increased the amplitude of the neuronal signal (sign test, $P < 0.001$ across bin sizes; Table S1 and Fig. 2) but left the amplitude of the noise intact (sign test, $P > 0.1$ across bin sizes; Table S1 and Fig. 2), thereby increasing the signal-to-noise ratio ($P < 0.005$ across all bin sizes; Table S1). This specific effect of acetylcholine on the amplitude of the signal, but not the noise, has not been previously reported. The observed increase in signal-to-noise ratio indicates that acetylcholine augments the encoding capacity of individual neurons. The encoding capacity of the network, however, also depends on the correlation structure of the network.

Despite their putative importance in shaping the encoding capacity of a network, signal correlations are frequently overlooked, and this study analyzes their modulation by acetylcholine. We calculated the signal correlation between each pair of neurons as the Pearson correlation between the pair's signals across all time bins (12, 13, 16). We did not find an effect of cholinergic stimulation on signal correlations (sign test; except for one bin size, all P values are above 0.1; Table S2 and Fig. 3). We did, however, observe a decrease in the magnitude of noise correlations (sign test, $P < 0.001$ across all bin sizes; Table S2

and Fig. 3), although the effect is very small compared with a previous report (14) (Discussion). The noise correlation between a pair of neurons was defined as the Pearson correlation between the pair's noise across all time bins and stimulus (movie) repetitions. (See Methods for details.)

Another important factor determining the encoding capacity of the network is the relationship between signal and noise correlations. There have been several reports indicating that signal and noise correlations are related (13, 22) such that neuronal pairs with similar receptive fields (high signal correlations) tend to have larger common variability (high noise correlations). Although this phenomenon is well documented, only recently it was discovered that the tight association between signal and noise correlations decreases under conditions of learning and attention (17, 19). The association between signal and noise correlations was quantified (18) by the correlations' slope, defined as the slope in the signal correlations vs. noise correlations graph (Fig. 4). This quantification is biologically relevant given that existing theoretical work suggested (but see below) that a decrease in the correlations' slope is associated with greater encoding capacity by the neuronal pool (18). We, therefore, examined the effects of cholinergic modulation on the correlations' slope. Our analysis revealed that acetylcholine decreases the correlations' slope ($P < 0.001$ across all bin sizes; Table S3 and Fig. 4), thereby suggesting an increase in the encoding capacity of the cortical network.

Although previous work suggests that a decrease in the correlation slope is beneficial for network encoding (18), such work uses simulations in which all neurons have the same signal-to-noise ratio and similar (homogeneous) receptive fields; both are assumptions that can limit the validity of the results (23). To better understand the role of neural correlations in network

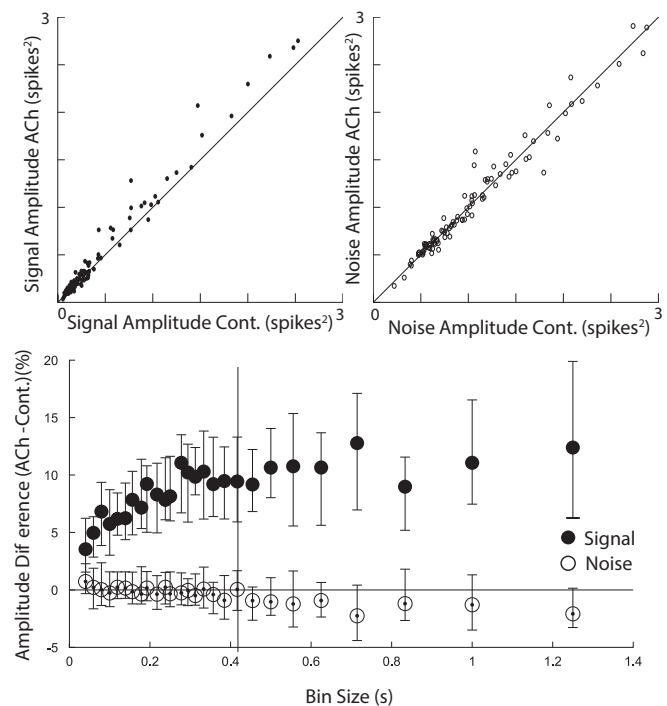


Fig. 2. Acetylcholine increases the signal amplitude while leaving the noise amplitude intact. (Top) Signal and noise amplitudes of the response of neurons in the visual cortex to naturalistic movie sequences, under cholinergic stimulation (ACh) or control (Cont.) conditions. Example taken from a representative bin size corresponding to the vertical bar shown at the Bottom. (Bottom) Percent difference of signal and noise amplitudes in the ACh vs. control conditions for various bin sizes, depicted is the median. Cholinergic activity significantly increases signal amplitudes (sign test, $P < 0.001$).




Fig. 3. Acetylcholine moderately decreases the noise correlations. (Top) Signal and noise correlations for the cholinergic stimulation (ACh) or control (Cont.) conditions (example for a representative bin size, as indicated by the vertical bar shown at the Bottom). (Bottom) Difference in signal and noise correlations in the ACh vs. Cont. conditions for all bin sizes. Plotted are the median differences. ACh significantly reduces noise correlations (sign test, $P < 0.001$), although the magnitude of the reduction is small.

encoding, we developed a technique that allows quantification of the network's encoding efficiency in terms of arbitrarily defined signal and noise covariance matrices, and that is applied to networks in which the neurons have Poisson statistics (*SI Methods*). Because of limitations in the software and running time of the computationally heavy simulation, we limit the size of the simulated network to 25 neurons. The essence of the technique is to generate a set of correlated signals and to use those signals to generate realizations of neural activity with correlated-Poisson (24) statistics. Furthermore, we used Bayesian methods to estimate the probability of a stimulus given the neural response and quantifying the accuracy of this "ideal estimator." This technique can be applied to investigate encoding efficiency for any possible covariance matrices. In this work, we use it to evaluate the effect of the correlations' slope on encoding efficiency. To do this, we generate a signal covariance matrix, and then we generate a noise covariance matrix such that the correlations' slope is high (*SI Methods*), and subsequently shuffle (to various degrees) the neurons' identities in the noise correlations matrix. This manipulation has the advantage of changing the correlations' slope while keeping the internal structure of the signal and the noise correlations intact. The accuracy of the estimator is quantified as the mean equivocation (*SI Methods*) over a set of realizations; a similar result is obtained when we use the success rate of a maximum likelihood predictor. For all simulation parameter sets chosen, we robustly observe that, as the correlations' slope decreases, the encoding efficiency increases ($P < 0.01$). Fig. 5 shows simulation results for five randomly chosen sets of parameters.

Studies considering neuronal models with Gaussian noise indicate that the effect of a given correlation structure in neural encoding might diminish as the size of the network increases (23, 25, 26). Therefore, we decided to explore how the effect of the correlations' slope depends on the size of the network. Repeating our analysis in networks of different sizes indicated that, up to our size limit, the effect of the correlations' slope increases

as the size of the network increases ($P < 0.05$ for each set of parameters; *SI Methods* and Fig. S1).

In summary, during the animals' passive observance of visual stimuli, optogenetic stimulation of cholinergic basal forebrain neurons modulated the cortical network by (i) increasing the signal-to-noise ratio of individual neurons, (ii) leaving the signal correlations intact, (iii) decreasing noise correlations, and (iv) reducing the dependence between signal and noise correlations (correlations' slope) in a manner that increases encoding efficiency. Thus, our analyses demonstrated that acetylcholine can modulate the statistical dependencies of cortical neurons to increase the encoding capacity of the network.

Discussion

By using optogenetic techniques to activate only basal forebrain cholinergic neurons while also recording from visual cortex neurons, our work allows a fine-grained investigation of the role of cortical acetylcholine on sensory representation. A large body of work addressing the effect of acetylcholine on neural responses was done in anesthetized animals (4, 7, 14, 27), where baseline cholinergic activity is low (28), making it difficult to interpret how the results generalize to more ecologically valid situations. Previous studies also used global stimulation of the basal forebrain (29–33), a technique that activates the entire population of basal forebrain cell types including cortically projecting cholinergic neurons in addition to GABAergic neurons projecting to cortex (for review, see ref. 34). Some work used iontophoretic application of cholinergic agonists (27), but it is not clear to what degree the amount of agonist liberated emulates a relevant physiological condition.

Our results support previous studies indicating that cholinergic activity produces an increase in the signal-to-noise ratio of individual cortical neurons (35–37). However, the definition of signal to noise used in such studies has been questioned for its lack of relevance as a descriptor of the encoding efficiency of the

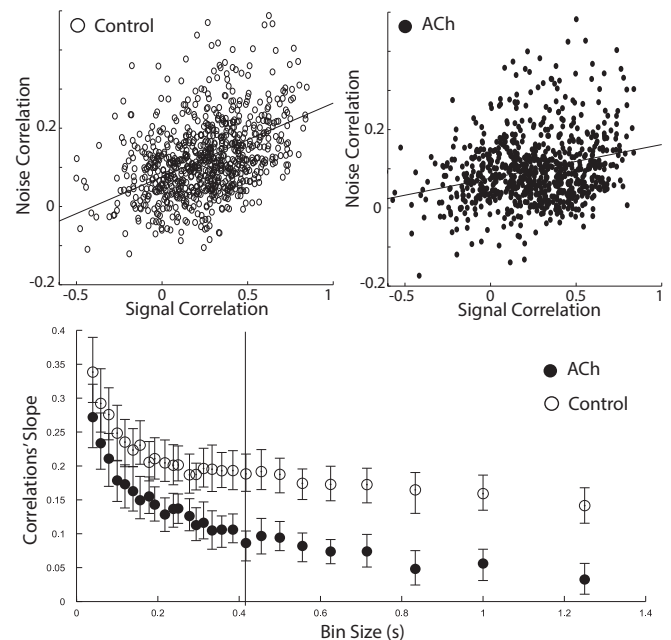


Fig. 4. Acetylcholine decreases the correlations' slope. We measured the slope between signal correlations and noise correlations. Recent work suggests that this variable is fundamental in determining the encoding efficiency of a network. (Top) Relationship between signal and noise correlations in the cholinergic stimulation (ACh) and control conditions (example for a representative bin size, as indicated by the vertical bar shown at the Bottom). (Bottom) The correlations' slope is depicted across bin sizes; acetylcholine decreases the correlations' slope ($P < 0.005$).

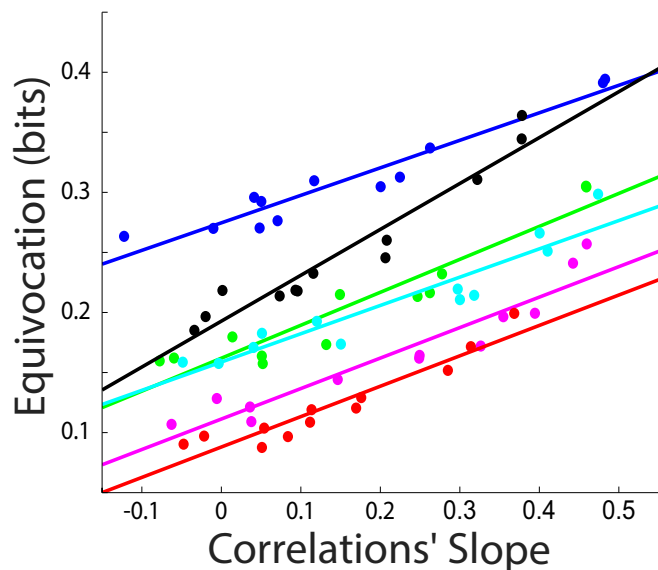


Fig. 5. Encoding efficiency decreases as the correlations' slope increases. We developed a method to study encoding efficiency as a function of the signal and noise covariance matrices in populations of correlated Poisson neurons. The figure shows the average equivocation, measured in bits, of an ideal estimator; a larger equivocation can be associated with a smaller encoding efficiency. Each color corresponds to a choice of simulation parameters; lines represent linear regressions. Across all sets of parameters studied, we observe that, as the correlations' slope increases, the encoding efficiency decreases ($P < 0.01$ for each set of parameters).

neurons (21). In this work, we use a definition of signal to noise that more correctly reflects the capacity of single neurons to encode information (20, 21).

The capacity of a network to encode information depends on the neural correlations; however, the influence of acetylcholine on the correlation among neurons has not been thoroughly explored. Activation of cortical acetylcholine decreased the correlation in the total activity of neuronal pairs (4), which depends not only on the signal and noise correlations but also on the signal-to-noise ratios (38). Acetylcholine also decreased the noise correlations (14); our results are consistent with this decrease, although the effect we observe is very small. This difference might be related to the fact that the referenced work (14) used a different species and brain region, and it was carried out in anesthetized animals. The latter difference is particularly important because in anesthetized preparations cholinergic levels are low and correlations are high (39, 40). Decreases in noise correlations have also been shown to be induced by increased attention (41, 42). However, the effect of these reductions on neural encoding is not entirely clear (15).

Signal correlations are presumably an important factor in shaping the encoding capacity of the network. However, there is a paucity of data regarding their modulation by acetylcholine. Signal correlations, however, can be indirectly related to the relationship between the neurons' receptive fields (13). If all neurons had sharper fields, the overlap in their activity would be smaller and the magnitude of the signal correlations would also be smaller. Thus, the lack of effect on signal correlations that we observe is consistent with the varied effect of acetylcholine on receptive fields found in previous work (35).

Seminal work regarding the influence of noise correlations on neural encoding indicates that these correlations are particularly harmful when they parallel the similarity between the receptive fields of corresponding neurons (15, 25). This is in accordance with work by Gu et al. (18) suggesting that a high correlations' slope is associated with poor encoding efficiency. More recent studies indicate that noise correlations are limiting for neural

coding when they are proportional to the product of the derivatives of the tuning curves (43), known as differential correlations. This measure has been associated with the signal correlation (44), as it reflects the local similarity between tuning curves.

The work cited above focuses on studying how the network is encoding a continuous parameter (such as an orientation angle). It is not clear, however, how that framework can be applied to more naturalistic situations in which the network is encoding a large combination of nominal and parametric variables. When studying how the network represents a parameter, investigators typically measure representation similarity across neurons in terms of tuning-curve similarity. A more naturalistic approach calls for the use of signal correlations, which quantify the similarity—across neurons—of the signals' probability distribution. The idea that encoding efficiency depends on the relationship between signal and noise correlations has existed for a long time (16, 45, 46). Work expressing that view analyzes discrimination between two stimuli using a single pair of neurons, without consideration of how their conclusions can be extrapolated to systems with many neurons and many stimuli.

Here, we present a method to analyze encoding efficiency in terms of full signal and noise covariance matrices, thereby providing a more holistic and accurate framework to understand neural coding. In this work, we do not assume that the noise is normally distributed (15, 23, 25, 47); rather, we study a model in which neurons respond to the stimuli in a discrete manner (48) with a Poisson probability distribution (38). This distinction is of importance because theoretical studies show that taking into account the discrete nature of neural activity influences the way the brain encodes stimuli (38, 48). When applying this more realistic method to investigate the role of the correlations' slope in neural encoding, we find that, as previous literature suggests (18), the encoding efficiency increases as the correlations' slope decreases.

A limitation of studies considering correlated discrete neural statistics is that the size of the networks investigated is small (38, 48); in this case, we use networks of 25 neurons. The observation that the magnitude of the effect we observe increases with the group size, suggests that the effect is also present in larger networks. We note, however, that models using neurons with Gaussian noise suggest that results pertaining to small networks, as the ones treated in this and previous work, may not generalize to very large networks (15, 23, 25, 26).

Our selection of a Poisson model is motivated by its widespread use and its consistency with the discrete nature of neural activity. Recent work (49) indicates that a more nuanced model can better capture some characteristics of neural activity. Such a model can be characterized as instantaneously Poisson, in the sense that it is Poisson at any given moment, with the neurons' rate being modulated by a gain factor that changes slowly over time and is shared across large swaths of the cortex. Thus, our Poisson model is not incompatible with the latter. The results presented here are valid as long as we consider that, at any given moment, our ideal decoder has information about the gain factor (26).

Despite the literature's suggestion that a tight association between signal and noise correlations is particularly bad for encoding, few studies have examined how this association is influenced by behavioral contingencies. One study of avian auditory cortex (17) indicates that learning and behavioral relevance of the auditory stimuli alter the relationship between signal and noise correlations, inverting the sign of the correlations' slope and thus making encoding more efficient. A more recent investigation (19) does not study signal correlations per se, but uses an alternative measure of receptive field similarity, showing that attention influences neural firing in such a way that the noise correlation of neuronal pairs with similar receptive fields decreases, whereas the noise correlation corresponding to neuronal pairs with dissimilar receptive fields increases. Here, we report that an analogous effect is elicited by enhanced cholinergic activation, suggesting acetylcholine as a potential candidate for shaping the cortical neural dynamics associated with attention and learning.

Methods

The neural data presented in this paper have been published elsewhere (3). However, all of the analyses and figures presented herein are our results generated by a unique quantitative approach.

Animals and Surgery. All procedures were approved by the Animal Care and Use Committee at the University of California, Berkeley. Experiments were performed on both male and female adult mice (2–6 mo old) weighing 20–45 g. The animals were housed on a 12/12-h light/dark cycle in cages of up to five animals before the implants, and individually after the implants. Data came from a total of nine *ChAT-ChR2(H134R)-EYFP* mice (50) (line 6; The Jackson Laboratory; *B6.Cg-Tg(Chat-COP4*H134RIEYFP)6Gfng/J*; stock number 014546).

Details about the surgical procedures and recordings can be found elsewhere (3). Briefly, we implanted a stainless-steel head plate under isoflurane anesthesia (5% induction and 1.5% maintenance). Moreover, a reference epidural screw was implanted above the left frontal cortex, and a half-drilled craniotomy was made to mark the location of the monocular region of the right V1, which was then sealed with a silicone elastomer (Kwik-Cast; World Precision Instruments). The animals were also implanted with a cannula to optogenetically target the basal forebrain (coordinates from bregma: anteroposterior, –0.5 mm; mediolateral, 1.8 mm; dorsoventral, 3.8 mm). Temperature was kept at 37 °C throughout the procedure using a heating pad, and the mice received two doses of buprenorphine (0.05 mg/kg, one before surgery and the other 6–8 h later) and supplementary analgesia with meloxicam (5 mg/kg) as necessary. They were allowed to recover for at least a week before recordings.

At the start of each recording session, the mouse was placed on a spherical treadmill under light isoflurane anesthesia. We performed a craniotomy ~300 μm in diameter over V1, preserving the dura. A laminar silicon probe was then inserted (~800 μm long, with up to 32 sites spaced by 50 μm; NeuroNexus Technologies; models polytrode 1B, 1C, or poly2). An optic fiber was then inserted through the implanted cannula to target the basal forebrain. The mouse was removed from anesthesia and allowed to recover for at least 45 min before recording.

Optogenetic Stimulation. Laser light was delivered to the basal forebrain via an optic fiber 200 μm in diameter (Thorlabs) inserted through and protruding 0.5 mm beyond the implanted cannula. We used a 473-nm DPSS laser (Crystalaser or Shanghai Laser and Optics Century Company) at a power of 1–3 mW at the fiber tip, delivered in 5-s square pulses. The laser was controlled by TTL pulses generated by the amplifier (Tucker-Davis Technologies).

Visual Stimulation. Visual stimuli were generated with a GeForce 7300 Graphics card (NVIDIA) in a PC running custom-written software and presented on a gamma-corrected 7" LCD monitor (Xenarc Technologies; maximal luminance, 250 cd/m²) with a refresh rate of 75 Hz. The monitor was placed 10 cm away from the left eye. All stimuli were presented in a 50° × 50° region centered at the average receptive field location of all simultaneously recorded units. The natural movie stimuli consisted of three 5-s clips selected from the van Hateren natural movie database (51). Each image was repeated for three frames, resulting in an effective frame rate of 25 Hz. Each trial started with 1 s of gray screen, followed by 1 s of the first movie frame, 5 s of movie, and 1 s of the last frame. Each movie was repeated 30 times in three blocks. Cholinergic stimulation and control conditions were interleaved.

Electrophysiology, Spike Sorting, and Neuron Selection. Spiking activity was recorded using a 32-channel TDT RZ5 (Tucker-Davis Technologies). Signals were filtered at 0.6–6 kHz and stored as raw voltage traces at 25 kHz. Spikes were detected off-line with custom-written software. We grouped nearby channels of the silicon probe into groups of three or four and performed semiautomatic spike sorting using Klusters (52). Spike clusters were considered single units if their autocorrelograms had a 2-ms refractory period and their cross-correlograms with other clusters did not have sharp peaks within ±2 ms

of zero lag. We excluded cells with average firing rates <0.5 Hz. We included all neurons in the analysis, whether or not they had a clear response to the stimuli. We took this approach considering that a neuron can contribute to population encoding even if it does not have an individual response to the stimuli. The results are equivalent, when we consider only neurons that respond to the stimuli, as assessed by an ANOVA with time bin as independent variable, neural activity as dependent variable, and inclusion criterion $P < 0.05$.

Calculation of Signal-to-Noise Ratios, Signal Correlations, and Noise Correlations. The three movie sequences were concatenated, thus generating a single 15-s sequence that is repeated 30 times. The number of spikes was counted on consecutive nonoverlapping time bins.

We designated $r_{i,t}^j$ as the number of spikes that neuron i fires on the j th repetition of the movie for bin centered at time t . We estimated the signal corresponding to neuron i and time t as the average of the spike counts taken over the 30 repetitions: $S_t^i = \langle r_{i,t}^j \rangle_j$. We then estimated the noise, for each neuron, time bin, and repetition, as the actual spike count minus the signal: $\xi_{i,t}^j = r_{i,t}^j - S_t^i$. We estimated the signal covariance between two neurons as the unbiased covariance of the signals calculated over all time bins $\sigma_{S_i S_j}^2 = (1/(T-1)) \sum_{t=1}^T (S_t^i - \langle S_t^i \rangle) (S_t^j - \langle S_t^j \rangle)$ and the noise covariance as the unbiased covariance over all time bins and repetitions $\sigma_{\xi_i \xi_j}^2 = (1/(JT-1)) \sum_{t=1}^T \sum_{j=1}^J \xi_{i,t}^j \xi_{j,t}^i$ (12). From the covariance matrices, we calculated the signal and noise amplitudes as the respective variances ($\sigma_S^2 = \sigma_{S_i S_i}^2$ and $\sigma_{\xi}^2 = \sigma_{\xi_i \xi_i}^2$); and the correlations as the covariance normalized by the product of the SDs.

To calculate the correlations' slope, we use total minimum squares, which minimizes the distance of the regression slopes to the data points and, therefore, makes no distinction between independent and dependent variables. The results are similar when using an ANOVA treating signal correlations as independent variables and noise correlations as dependent variables, as done in refs. 17 and 18.

Cholinergic Stimulation vs. Control Conditions. To calculate the effect of cholinergic stimulation on the amplitude of the signals, we calculated their difference in the cholinergic vs. control conditions. As the distribution of differences (across all neurons) is not normal (Lilliefors test, $P < 0.01$), we report the medians and use the sign test to assess for statistical differences. The distribution of the differences of noise amplitudes, the differences in signal correlations, and the differences in noise correlations is also not normal (Lilliefors test, $P < 0.01$). Therefore, again, we report the medians and use the sign test to assess for statistical differences. To calculate error bars, we use bootstrap methods. We randomly draw neurons with replacement and extract the variable of interest, repeating this process 10,000 times. The error bars in the figures represent the 95% confidence interval of the bootstrap distribution. The estimation errors in Tables S1–S3 represent the SD of the bootstrap distributions.

To calculate statistical differences across conditions in the correlations' slope, we use a bootstrap method. We draw neuron pairs with replacement and calculate the total least-squares regression, repeating this procedure 10,000 times. To calculate the P value, we perform an unpaired t test of the resulting distributions. The error bars shown in the figures represent the 95% confidence interval of the bootstrap distribution.

Data and Code. Data and Matlab code for the simulations can be accessed at <https://github.com/victorminces/neural-coding>.

ACKNOWLEDGMENTS. We thank Anita Disney, Andrew Alexander, and Laleh Quinn for suggestions and help with the manuscript. This work was supported by NIH Grant R01 EY018861, National Science Foundation (NSF) Grant 22250400-42533 (to Y.D.), Ruth L. Kirschstein National Research Service Award F31NS084696 from the National Institute of Neurological Disorders and Stroke (to L.P.), NSF Science of Learning Center Grant SMA 1041755 to the Temporal Dynamics of Learning Center (to V.M. and A.A.C.), and National Institute of Mental Health R01 MH110514 01 (to A.A.C.).

- Baxter MG, Chiba AA (1999) Cognitive functions of the basal forebrain. *Curr Opin Neurobiol* 9:178–183.
- Mincses VH, Alexander AS, Datlow M, Alfonso SI, Chiba AA (2013) The role of visual cortex acetylcholine in learning to discriminate temporally modulated visual stimuli. *Front Behav Neurosci* 7:16.
- Pinto L, et al. (2013) Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat Neurosci* 16:1857–1863.
- Goard M, Dan Y (2009) Basal forebrain activation enhances cortical coding of natural scenes. *Nat Neurosci* 12:1444–1449.
- Roberts MJ, et al. (2005) Acetylcholine dynamically controls spatial integration in marmoset primary visual cortex. *J Neurophysiol* 93:2062–2072.

- Castro-Alamancos MA (2004) Dynamics of sensory thalamocortical synaptic networks during information processing states. *Prog Neurobiol* 74:213–247.
- Disney AA, Aoki C, Hawken MJ (2007) Gain modulation by nicotine in macaque v1. *Neuron* 56:701–713.
- Kawai H, Lazar R, Metherate R (2007) Nicotinic control of axon excitability regulates thalamocortical transmission. *Nat Neurosci* 10:1168–1175.
- Metherate R, Cox CL, Ashe JH (1992) Cellular bases of neocortical activation: Modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. *J Neurosci* 12:4701–4711.
- Herrero JL, et al. (2008) Acetylcholine contributes through muscarinic receptors to attentional modulation in V1. *Nature* 454:1110–1114.

11. Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143.
12. Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13:2758–2771.
13. Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819.
14. Thiele A, Herrero JL, Distler C, Hoffmann K-P (2012) Contribution of cholinergic and GABAergic mechanisms to direction tuning, discriminability, response reliability, and neuronal rate correlations in macaque middle temporal area. *J Neurosci* 32:16602–16615.
15. Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101.
16. Averbach BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
17. Jeanne JM, Sharpee TO, Gentner TQ (2013) Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron* 78:352–363.
18. Gu Y, et al. (2011) Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron* 71:750–761.
19. Ruff DA, Cohen MR (2014) Attention can either increase or decrease spike count correlations in visual cortex. *Nat Neurosci* 17:1591–1597.
20. Cover TM, Thomas JA (1991) *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing (Wiley, New York).
21. Disney AA, Schultz SR (2004) Hallucinations and acetylcholine: Signal or noise? *Behav Brain Sci* 27:190–191.
22. Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
23. Ecker AS, Berens P, Tolias AS, Bethge M (2011) The effect of noise correlations in populations of diversely tuned neurons. *J Neurosci* 31:14272–14283.
24. Macke JH, Berens P, Ecker AS, Tolias AS, Bethge M (2009) Generating spike trains with specified correlation coefficients. *Neural Comput* 21:397–423.
25. Sompolinsky H, Yoon H, Kang K, Shamir M (2001) Population coding in neuronal systems with correlated noise. *Phys Rev E Stat Nonlin Soft Matter Phys* 64:051904.
26. Kanitscheider I, Coen-Cagli R, Pouget A (2015) Origin of information-limiting noise correlations. *Proc Natl Acad Sci USA* 112:E6973–E6982.
27. Tremblay N, Warren RA, Dykes RW (1990) Electrophysiological studies of acetylcholine and the role of the basal forebrain in the somatosensory cortex of the cat. II. Cortical neurons excited by somatic stimuli. *J Neurophysiol* 64:1212–1222.
28. Dringenberg HC, Vanderwolf CH (1998) Involvement of direct and indirect pathways in electrocorticographic activation. *Neurosci Biobehav Rev* 22:243–257.
29. Metherate R, Weinberger NM (1990) Cholinergic modulation of responses to single tones produces tone-specific receptive field alterations in cat auditory cortex. *Synapse* 6:133–145.
30. Edeline JM, Hars B, Maho C, Hennevin E (1994) Transient and prolonged facilitation of tone-evoked responses induced by basal forebrain stimulations in the rat auditory cortex. *Exp Brain Res* 97:373–386.
31. Metherate R, Ashe JH (1991) Basal forebrain stimulation modifies auditory cortex responsiveness by an action at muscarinic receptors. *Brain Res* 559:163–167.
32. Golmayo L, Nuñez A, Zaborszky L (2003) Electrophysiological evidence for the existence of a posterior cortical-prefrontal-basal forebrain circuitry in modulating sensory responses in visual and somatosensory rat cortical areas. *Neuroscience* 119:597–609.
33. Berg RW, Friedman B, Schroeder LF, Kleinfield D (2005) Activation of nucleus basalis facilitates cortical control of a brain stem motor program. *J Neurophysiol* 94:699–711.
34. Chiba AA, Quinn L (2007) The basal forebrain and memory. *Memory Systems, Vol 3, Learning and Memory: A Comprehensive Reference*, ed Eichenbaum H (Elsevier, Oxford), pp 281–302.
35. Zinke W, et al. (2006) Cholinergic modulation of response properties and orientation tuning of neurons in primary visual cortex of anaesthetized Marmoset monkeys. *Eur J Neurosci* 24:314–328.
36. Sato H, Hata Y, Masui H, Tsumoto T (1987) A functional role of cholinergic innervation to neurons in the cat visual cortex. *J Neurophysiol* 58:765–780.
37. Sillito AM, Kemp JA (1983) Cholinergic modulation of the functional organization of the cat visual cortex. *Brain Res* 289:143–155.
38. Pitkow X, Meister M (2012) Decorrelation and efficient coding by retinal ganglion cells. *Nat Neurosci* 15:628–635.
39. Dringenberg HC, Vanderwolf CH (1997) Neocortical activation: Modulation by multiple pathways acting on central cholinergic and serotonergic systems. *Exp Brain Res* 116:160–174.
40. Ecker AS, et al. (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82:235–248.
41. Cohen MR, Maunsell JHR (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600.
42. Mitchell JF, Sundberg KA, Reynolds JH (2009) Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* 63:879–888.
43. Moreno-Bote R, et al. (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417.
44. Hu Y, Zylberberg J, Shea-Brown E (2014) The sign rule and beyond: Boundary effects, flexibility, and noise correlations in neural population codes. *PLoS Comput Biol* 10:e1003469.
45. Johnson KO (1980) Sensory discrimination: Decision process. *J Neurophysiol* 43:1771–1792.
46. Oram MW, Földiák P, Perrett DI, Sengpiel F (1998) The “ideal homunculus”: Decoding neural population signals. *Trends Neurosci* 21:259–265.
47. Moreno-bote R, et al. (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417.
48. Tkacik G, Prentice JS, Balasubramanian V, Schneidman E (2010) Optimal population coding by noisy spiking neurons. *Proc Natl Acad Sci USA* 107:14419–14424.
49. Goris RLT, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17:858–865.
50. Zhao S, et al. (2011) Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. *Nat Methods* 8:745–752.
51. van Hateren JH, van der Schaaf A (1998) Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc Biol Sci* 265:359–366.
52. Hazan L, Zugaro M, Buzsáki G (2006) Klusters, NeuroScope, NDManager: A free software suite for neurophysiological data processing and visualization. *J Neurosci Methods* 155:207–216.