

Optogenetic Activation of an Inhibitory Network Enhances Feedforward Functional Connectivity in Auditory Cortex

Liberty S. Hamilton,¹ Jascha Sohl-Dickstein,² Alexander G. Huth,¹ Vanessa M. Carels,¹ Karl Deisseroth,³ and Shaowen Bao^{1,*}

¹Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA 94720, USA

²Redwood Institute for Theoretical Neuroscience, University of California, Berkeley, Berkeley, CA 94720, USA

³Departments of Bioengineering and Psychiatry, Howard Hughes Medical Institute, Stanford University, Stanford, CA 94305, USA

http://dx.doi.org/10.1016/j.neuron.2013.08.017

SUMMARY

The mammalian neocortex is a highly interconnected network of different types of neurons organized into both layers and columns. Overlaid on this structural organization is a pattern of functional connectivity that can be rapidly and flexibly altered during behavior. Parvalbumin-positive (PV+) inhibitory neurons, which are implicated in cortical oscillations and can change neuronal selectivity, may play a pivotal role in these dynamic changes. We found that optogenetic activation of PV+ neurons in the auditory cortex enhanced feedforward functional connectivity in the putative thalamorecipient circuit and in cortical columnar circuits. In contrast, stimulation of PV+ neurons induced no change in connectivity between sites in the same layers. The activity of PV+ neurons may thus serve as a gating mechanism to enhance feedforward, but not lateral or feedback, information flow in cortical circuits. Functionally, it may preferentially enhance the contribution of bottom-up sensory inputs to perception.

INTRODUCTION

Neurons communicate with each other in dynamically modulated circuits. Functional connectivity, a measure of interactions between neurons in these circuits, can change gradually during learning (McIntosh and Gonzalez-Lima, 1998) and formation of long-term memories, or it can change rapidly, depending on behavioral context and cognitive demands. While the mechanisms underlying long-term network plasticity have been extensively documented, those underlying rapid modulation of functional connectivity remain largely unknown. At the network level, functional connectivity is affected by up-down and oscillatory states of the neural network (Gray et al., 1989). Cortical inhibition plays a key role in this process (Cardin et al., 2009; Sohal et al., 2009; Womelsdorf et al., 2007). Parvalbumin-positive (PV+) interneurons, which make up more than half of the inhibitory neurons in the cortex (Celio, 1986), are particularly important as they provide strong feedforward and feedback inhibition that can synchronize the cortical network (Cardin et al., 2009; Fuchs et al., 2007; Isaacson and Scanziani, 2011; Sohal et al., 2009). Their precise influence on cortical networks during sensory processing, however, remains unclear. In particular, it is unknown how PV+ neurons may differentially modulate responses in different layers of the neocortex and how the anatomical organization of the cortex affects the flow of information through these networks.

Histological studies have shown that the cortex consists of defined layers with vertical projections between those layers (Lee and Winer, 2008; Linden and Schreiner, 2003; Winer and Lee, 2007). Functional connectivity within cortical networks has traditionally been investigated by measuring the cross-correlation between the spike trains of pairs of neurons (Douglas et al., 1989; Douglas and Martin, 1991). Still, little is known about functional connectivity under sensory stimulation or about the role of inhibition in the cortical network. We combine multiple computational approaches with optogenetic activation of PV+ neurons to determine how inhibitory activity modulates network connectivity within and across layers and columns of the cortex.

RESULTS AND DISCUSSION

We targeted expression of the light-sensitive channel channelrhodopsin-2 (ChR2) to PV+ neurons in the mouse auditory cortex (Figure 1A), using a Cre-dependent adeno-associated virus (Sohal et al., 2009). One month posttransfection, we recorded neural responses with a 4 \times 4 polytrode in putative L2/3 through L4 of the primary auditory cortex (Figure 1B) while playing pure tones to the contralateral ear and stimulating PV+ cells with blue light (Figure 1C).

Using Ising Models to Recover Functional Connectivity in Cortical Circuits

Functional connectivity between the recorded sites was quantified using Ising models, which have previously been used to model neural interactions in many different systems (Ganmor et al., 2011a, 2011b; Köster et al., 2012; Marre et al., 2009; Ohiorhenuan et al., 2010; Roudi et al., 2009a; Schaub and Schultz, 2012; Schneidman et al., 2006; Shlens et al., 2006, 2009; Tang et al., 2008). The Ising model describes the coupling (a measure

^{*}Correspondence: sbao@berkeley.edu



Figure 1. Viral Expression, Recording Set-Up, and Responses to Pure Tone and Optogenetic Stimulation

(A) We injected PV-Cre mice with 1 μ l of a Creinducible adeno-associated virus in the right auditory cortex that resulted in transfection of the light-sensitive ion channel ChR2 in PV+ cells. Histology confirmed the colocalization of ChR2eYFP (left) to PV+ cells in the auditory cortex (center, immunostained with Alexa 594, and right, [see Merge]). Approximately 58% of PV+ cells were transfected with ChR2 (white arrows indicate examples of colocalization). Scale bars, 50 μ m. (B) Schematic depicting recording set-up. A 4 x 4 silicon polytrode was lowered orthogonally to the

cortex so that the deepest sites were located at a depth of ${\sim}500~\mu m$. A 200- μm -diameter optic fiber coupled to a 473 nm blue laser was positioned parallel to the polytrode, 1–2 mm above the cortex to provide optogenetic stimulation during 50% of the trials.

(C) Light and sound stimulus conditions for example trials and corresponding spike raster plot. Input to the Ising model was a binary matrix

including the light condition at each time point (blue bars represent the time during which the light was on and PV+ cells were being stimulated), the frequency of the pure tone stimulus that was presented at each time (represented by pink bars), and the spike data for each channel, binned at 5 ms. Sound and light conditions were randomly interleaved for each 1 s trial.

of functional connectivity) between pairs of recording sites and between recording sites and external stimuli based on observed population firing patterns and corresponding stimuli (Figures 1B and 1C). Because all pairwise interactions are fitted simultaneously, Ising models are less prone to false-positive interactions that are inherent to traditional correlation analysis (Schneidman et al., 2006). For example, in a fully connected Ising model (see Experimental Procedures), the strongest coupling to sounds occurred in rows 3 and 4 (Figure 2A), corresponding to the thalamorecipient layers. By contrast, traditional correlation analysis indicated strong connectivity between sounds and sites in all rows (Figure 2B). This false-positive connectivity between sounds and activity in rows 1 and 2 is due to the absence of site-to-site interactions in the correlation analysis. In a reduced Ising model where recording sites were coupled to sound but not to each other, which we call the independent neurons model, positive couplings between neural activity and the sound stimulus were also present in all recorded layers and did not differ across depth (Figure 2C; p = 0.55, Kruskal-Wallis analysis of variance [ANOVA]). Furthermore, pairwise correlations were more tightly correlated with couplings in the independent neurons model than in the fully connected model (Figure 2D; correlations resampled 100,000 times, with the difference in correlations significant at bootstrapped p < 1 \times 10⁻⁴). The fully connected model showed significantly higher log-likelihood on held-out data than the independent model (Figure 2E; p = 0.013, Wilcoxon signed-rank test), suggesting a significant contribution of site-tosite interactions to neuronal activity.

The Ising model can discover spatial structure within the network despite no prior knowledge of spatial locations of the polytrode recording sites. In the fully connected Ising model, coupling was stronger in the vertical and horizontal directions than in the diagonal directions (Figure 2F), presumably due to neuronal projections within cortical columns and layers. In addition, coupling decreased more rapidly with vertical than with horizontal distance—sites up to 375 μ m apart horizontally were still more strongly coupled than sites 300 μ m away vertically (p = 4.3 × 10⁻⁶; Wilcoxon rank sum test). Such connectivity structure was much less prominent in the pairwise correlations (Figure 2G; ratio of column and layer to diagonal couplings = 1.26 ± 0.03 for correlations, 2.16 ± 0.20 for couplings; p = 0.001, Wilcoxon rank sum test). Thus, although the model is blind to the relative locations of the recording sites, the fully connected Ising model recovered known layer and column circuitry (Linden and Schreiner, 2003; Mountcastle, 1957).

Optogenetic Activation of PV+ Neurons Enhances Functional Connectivity

Using the fully connected Ising model, we analyzed how optogenetic activation of PV+ neurons influences functional connectivity in laminar, columnar, and thalamic input circuits of the primary auditory cortex. In keeping with PV+ neurons providing inhibitory input to connected pyramidal cells, we saw an overall reduction of the Ising model bias term in "light-on" trials, reflecting reduced firing rates in all rows (Figure 3A; Bonferroni-corrected p = 0.003, p = 0.0002, $p = 8.4 \times 10^{-6}$, and $p = 8.7 \times 10^{-5}$ for rows 1, 2, 3, and 4, respectively, Wilcoxon signed-rank tests). Furthermore, we found that stimulating PV+ neurons led to increases in vertical connectivity between sites within the same vertical column (Figure 3B; Bonferroni-corrected p = 0.01 and $p = 1 \times 10^{-4}$ for coupling between sites within the same column, two and three rows away, respectively, Wilcoxon signed-rank tests) but did not change horizontal connectivity within layers (p > 0.05 for all comparisons, Wilcoxon signed-rank tests). Coupling between neural activity and sounds increased for sites in rows 3 and 4 during stimulation of PV+ neurons (Figures 3C



Figure 2. The Ising Model Recovers Canonical Cortical Structure Not Observed with Traditional Correlation Analysis

(A) Mean sound-to-site coupling derived from one set of polytrode recording data using the fully connected model. When connections between sites are taken into consideration, the strongest couplings with sounds are seen in rows 3 and 4, corresponding to the putative thalamic input layers.
(B) Mean sound-to-site correlation from the same data set as in (A). Neural activity on all rows of the polytrode was positively correlated with sound presentation.

(C) Mean sound-to-site coupling derived from the same data set as in (A) using the independent neurons model. When we assume the sites are not connected to each other, the coupling between neural activity and sound does not change as a function of depth (electrode row).

(D) Pairwise correlations as a function of Ising model couplings for the fully connected model of connectivity for all subjects and locations (left, Spearman ρ = 0.61). Pairwise correlations as a function of Ising model couplings for the independent neurons model of connectivity for all subjects and locations (right, Spearman ρ = 0.97). When connections between neurons are considered, couplings are more distinct from correlations.

(E) Ising model performance for the fully connected model and the independent neurons model. The fully connected model showed the highest log-likelihood on held-out data. Data are presented as mean \pm SEM. **p < 0.01.

(F) Coupling between pairs of recording sites in the fully connected model as a function of horizontal and vertical distance collapsed across all locations. Laminar and columnar structure emerges from the model in the form of strong positive couplings for horizontal and vertical distances = 0, compared to couplings between sites situated diagonally (horizontal and vertical distance \neq 0), which are weaker.

(G) Mean correlation between sites as a function of horizontal and vertical distances collapsed across all locations. The correlation falls off strongly with vertical distance but remains more uniform across horizontal distance.

and 3D; Bonferroni-corrected p = 0.0003 and p = 8×10^{-13} for the third and fourth rows, respectively, Wilcoxon signed-rank tests). These sites were likely located in the thalamorecipient input layers (layer 4 and deep layer 3). The increase in soundto-site coupling in putative thalamorecipient layers was not an artifact of the response window selection (Figure S1 available online). Our findings indicate that activation of PV+ neurons results in enhanced functional connectivity, specifically in thalamocortical input and cortical columnar circuits.

Stimulation of PV+ Neurons Enhances Functional Connectivity in the Feedforward Direction

While the Ising model uncovers altered functional connectivity with inhibitory neuron stimulation, it is agnostic to the direction in which these changes occur. For example, the increased coupling within cortical columns during activation of PV+ neurons could be in the feedforward, feedback, or both directions. To address this issue, we used vector autoregression (VAR) to derive a linear model that described how activity in one site was modulated by spikes in other sites as a function of time delay (Figure 4A; see Experimental Procedures for details). Unlike the Ising model, which describes dynamics within a fixed time bin, this model considers how inputs from different rows at different times affect the neural responses in a given time. Prediction of one site's activity using the population activity was significantly better during the "light-on" than during the "light-off" epochs (Figure 4B; Wilcoxon signed-rank test, $p = 4.0 \times 10^{-10}$).

We then examined the contribution of each site to predicting the activity of another site (i.e., the weight function in the linear



Figure 3. Optogenetic Activation of PV+ Neurons Decreases Average Firing Rate while Increasing Functional Connectivity within Columns and from Sounds to Sites Located in Putative Thalamorecipient Layers

(A) Bias term for the Ising model for each site in different rows on the polytrode, representing a proxy for the intrinsic firing rate of each site. Light stimulation significantly reduced the bias term in all layers, indicating an overall reduction of firing rate with stimulation of PV+ neurons.

(B) Site-to-site couplings as a function of the distance between the sites. Couplings are plotted between sites one, two, and three nodes away within a column or within a layer (see node diagrams below the x axis). Light stimulation significantly increased couplings within a column at distances two and three nodes away and did not change couplings within a layer.

(C) Sound-to-site coupling averaged across sites in each row. During optogenetic stimulation, coupling between neural activity and sound stimulation increased in the putative thalamorecipient layers (rows 3 and 4).

(D) Examples of sound-to-site coupling in sites recorded from three different animals, plotted as a function of tone frequency and electrode row. In all cases, there was an increase in coupling to sounds in the putative thalamic input layers (rows 3 and 4) during light stimulation of PV+ neurons. Some locations showed decreases in coupling to sounds in the superficial layers during light stimulation, but this effect was not consistent across the population; see (C).

*p < 0.05, **p < 0.01, and ***p < 0.001, Bonferroni corrected for multiple comparisons. Error bars and black marker in (A) through (C) depict mean ± SEM. See also Figure S1.

model as a measure of functional connectivity; Figures 4C and 4D). In general, neural activity was more strongly modulated by activity of sites in the same cortical layers rather than in different layers. However, these weights were not significantly altered by activation of PV+ neurons (Figure 4D, diagonal subplots). By contrast, PV+ neuron activation significantly increased the weights for row 4 sites in predicting the activity of more superficial sites within a time window between 6 and 12 ms (Figure 4D, far right subplots). There was also a small trend (not significant) of increased excitatory drive from row 3 to row 4, consistent

with the primary input layer to auditory cortex arising in deep layer 3 and propagating information to layer 4 (Smith and Populin, 2001). Furthermore, inhibitory influences from superficial row 1 on activity in row 3 were lessened with PV+ neuron stimulation (Figure 4D, first column, third row subplot), suggesting that the normal feedback inhibition from superficial layers is altered when PV+ neurons inhibit those cells. The double dissociation between the stronger baseline intralayer influences and the light-activated increase for cross-layer influences supports our findings from the Ising model analysis that the activation of



Figure 4. PV+ Neuron Activation Increases Functional Connectivity in the Feedforward Direction

(A) Schematic of the VAR analysis. To model the interaction between the activity of sites at different depths, we fit a VAR model that predicts the activity for each site on the polytrode based on the activity of all other contact sites. In this schematic, neural activity in site 1 is predicted from activity in sites 2, 3, and 4. We fit a linear model that describes how spikes in each site modulate spiking in the predicted site at different time delays. For example, spikes in the teal trace often precede spikes in the black trace, so positive weights and apply an output nonlinearity to improve response predictions.

(B) Prediction performance of VAR models is significantly higher when modeling spike trains during "light-on" epochs compared to "light-off" epochs ($p < 4.0 \times 10^{-10}$, Wilcoxon signed-rank test). Prediction performance was assessed as the correlation between the response predicted by the VAR model and the actual response from held-out data. We used resampling to determine which models performed significantly better than chance, indicated in black. Gray points show VAR models that were excluded from further analysis. The red dashed line is the unity line y = x.

(C) Representative VAR model weights from one subject for sites on one shank of the polytrode. Each plot shows how the activity in sites from rows 1–4 modulates the activity of sites 1, 5, 9, and 13, respectively (on the leftmost shank, see Figure 1B). Excitatory drive is shown in red; inhibitory drive is shown in blue. Dashed lines separate channels in different rows on the polytrode.

(D) Average VAR model weights (β_i from Equation 4) collapsed across all subjects and across

contact sites in the same row show that light stimulation of PV+ cells results in stronger and longer excitatory drive from row 4 to rows 1 and 2. VAR model weights were significantly higher for "light-on" epochs (Bonferroni-corrected p value indicated by color on horizontal line, Wilcoxon signed-rank test) when predicting rows 1–2 from activity in row 4. Row 2 also showed significantly more suppression from the putative thalamorecipient layer before the window of strong excitation starting at \sim 12 ms. When predicting row 3 activity from activity in superficial row 1, we found that stimulation of PV+ neurons reduced the inhibitory influence of row 1 on row 3, as evidenced by regression weights closer to zero. The strong increase in drive from row 4 to more superficial layers and the small regression weights in the feedback direction shows that stimulation of PV+ neurons increases drive in the feedforward direction. Data are presented as mean regression weight across all sites and subjects ± SEM. See also Figure S2.

PV+ neurons specifically increases intracolumn functional connectivity. The increased contribution of activity in row 4 to firing in superficial rows during light stimulation further suggests that the enhanced functional connectivity is in the feedforward direction.

Qualitatively similar results were also observed when fitting the data in a generalized linear model (GLM) with an exponential nonlinearity (see Pillow et al., 2008), although predictive performance of the GLM as assessed by correlation was worse than the VAR model for both "light-off" and "light-on" conditions (p < 0.0001, Wilcoxon signed-rank tests; see Supplemental Experimental Procedures; Figure S2). These results, together with the increased sound-to-site coupling in the feedforward thalamocortical circuit, suggest that activation of auditory cortical PV+ neurons may facilitate bottom-up information flow in the feedforward direction.

Activation of PV+ Neurons Increases Detection Signal-to-Noise Ratio in Single Recording Sites

Previous studies have shown that optogenetic activation of PV+ neurons enhances stimulus feature selectivity and increases the signal-to-noise ratio (SNR) in cortical neurons (Atallah et al., 2012; Lee et al., 2012; Sohal et al., 2009; Wilson et al., 2012). In our study, light activation of PV+ neurons induced strong suppression of spontaneous firing and weak reduction of tone-evoked responses (mean percent suppression \pm SEM = 31.77% \pm 0.03% for spontaneous activity and 18.57% \pm 0.03% for evoked activity; see Figures 5A and 5B for examples of peristimulus time histograms and receptive fields). This led to an increase in the detection SNR (mean detection SNR \pm SEM = 6.13 \pm 0.73 for "light-on" versus 3.17 \pm 0.21 for "light-off" trials, p = 0.005 Wilcoxon signed-rank test, Figure 5C). In addition, PV+ neuron stimulation significantly narrowed





Figure 5. PV+ Neuron Activation Strongly Suppresses Spontaneous Activity, Weakly Suppresses Sound-Evoked Activity, and Improves Detection SNR of Responses

(A) Example peristimulus time histogram of multiunit activity. Trials with no light stimulation (left) showed strong responses to pure tones played at 0.5 s (red arrow). When blue light was used to stimulate PV+ neurons between 0.25 and 0.75 s (right), responses to sound were still observed, but the overall evoked response was reduced. Spontaneous activity was reduced throughout the duration of the light stimulus.

(B) Example receptive fields during "light-off" (left) and "light-on" (right) trials separately for the same site in (A). Each pixel in the plot represents the average number of spikes evoked by a stimulus at a particular frequency and intensity level. The reduction in spontaneous rate is evident in the decreased background spikes outside of the V-shaped receptive field.

(C) Detection SNR (measured as a [number evoked – number of spontaneous spikes]/[number of spontaneous spikes]) during "light-off" (gray) and "light-on" (blue) trials. Optogenetic activation of PV+ neurons results in a significant increase in the SNR compared to the no-stimulation condition. Data are presented as mean SNR \pm SEM.

(D) Receptive field bandwidths 20 dB above threshold (left) and receptive field thresholds (right) for receptive fields calculated from "light-off" (no stimulation of PV+ neurons) and light-on (stimulation of PV+ neurons) trials separately. Lines are colored according to whether a reduction (blue), an increase (red), or no change (black) in bandwidth was observed. We observed a significant reduction in receptive field bandwidths during PV+ neuron stimulation (p < 0.001, Wilcoxon signed-rank test) with no change in threshold. This indicates that stimulating PV+ neurons increases stimulus selectivity by narrowing the range of stimuli to which a site responds.

*p < 0.05, **p < 0.01, and ***p < 0.001, Bonferroni-corrected Wilcoxon signed-rank tests. See also Figure S3.

receptive field bandwidths (p < 0.001, Wilcoxon signed-rank test) without changing response thresholds at the characteristic frequency (p = 0.79, Wilcoxon signed-rank test, Figure 5D). In sham-injected control mice not expressing ChR2, light stimulation did not cause any significant change in response properties (Figure S3).

Reductions in Spontaneous Activity Alone Do Not Account for Functional Connectivity Changes

To test the possibility that reduced spontaneous activity and increased detection SNR (Figures 5A-5C) caused the observed increases in site-to-site coupling (Figure 3B), we randomly removed 20%-80% of spikes recorded in "light-off" trials to mimic the effects of stimulation of PV+ neurons with light and reconducted the Ising model analysis (see Experimental Procedures). The mean site-to-site coupling strength was not increased by the random reduction of spontaneous and evoked spikes (Figure 6A) but rather was reduced in sites one node away within the same column (p < 0.001 for all comparisons, Bonferroni-corrected Wilcoxon signed-rank tests). No changes to between-site coupling two and three sites away within the column were seen (Bonferroni-corrected p > 0.05, Wilcoxon signed-rank tests), even with reductions in activity that were far larger than the suppression caused by PV+ neuron stimulation (~32% suppression on average). There was also no change in sound-to-site coupling with these manipulations (Figure 6B). Finally, to determine if the altered site-to-site coupling strength was due to changes in evoked activity, we removed sound-evoked spikes and reconducted the analysis with only the (unaltered) spontaneous activity. The coupling strength was still higher during activation of the PV+ neurons (Figure 6C; Bonferroni-corrected p = 0.002 and p = 0.0002 for sites 2 and 3 away within a column, respectively, Wilcoxon signed-rank tests). These results indicate that the enhanced coupling by PV+ neuron activation was not due to the increased detection SNR or reduced baseline activity. Rather, it reflects the state of the circuit connectivity and is independent of sensory stimulation and responses.

Possible Mechanisms for the Enhancement of Functional Connectivity by PV+ Neurons

In this study, we quantified functional connectivity in the auditory cortex with coupling from the Ising model and the weight function from vector autoregression. Both measures elucidate how the activity of a neuron or the presentation of a sound stimulus drives the firing of a target neuron. The specific mechanisms underlying the modulation of functional connectivity by PV+ neurons were not investigated in the present study but could involve the modulation of synaptic connections and changes in global network states. For example, synaptic efficacy can be rapidly altered by the prior synaptic activity (Zucker and Regehr, 2002), which is likely influenced by the activity of PV+ neurons. Alternatively, by synchronizing network activity (Cardin et al., 2009; Sohal et al., 2009), PV+ neurons could set target neurons in a more excitable state when the projection neuron fires, thus enhancing their functional connectivity. The effects on column rather than layer connections may be related to anisotropic projection patterns of PV+ neurons (Packer and Yuste, 2011), whereby PV+ neurons preferentially inhibit pyramidal neurons



Figure 6. Effect of Manipulation of Spontaneous and Evoked Activity on Ising Model Couplings

(A) Random reduction of spikes in all layers during "light-off" trials does not increase site-to-site couplings in the Ising model as did PV+ neuron stimulation. Instead, significant reductions in site-to-site couplings are seen in both column and layer couplings one node away for several levels of firing rate reduction (20%, 50%, or 80% reduction). located in the same vertical columns over distances 200 μm and greater.

Potential Limitations

While both the Ising model and the VAR models allow us to analyze the relative changes to within- versus between-layer connectivity with PV+ neuron stimulation, some caution should be taken when interpreting these functional connections in terms of synaptic interactions. With extracellular recordings, it is not possible to reconstruct the synaptic connections between recorded (or stimulated) neurons. Coupling between neurons should be considered as a functional description rather than an anatomical one. For example, researchers have found that coupling weights in the Ising model do not necessarily correspond to synaptic connections in the network (Roudi et al., 2009b). The strength of the Ising model lies in its ability to distinguish direct from indirect interactions; for example, in finding direct stimulus input to rows 3 and 4, representing the thalamorecipient layer. However, the symmetric nature of Ising model couplings means that directed interactions, such as combined excitatory/inhibitory influences (cell A excites cell B, but B inhibits A), cannot be uncovered. The VAR model addresses some of these caveats, since it can quantify directional interactions between recording sites and describe how neuronal firing is affected in different time periods. Our model shows that strong feedforward drive is enhanced by stimulation of PV+ neurons. whereas feedback from superficial to putative thalamic input layers is not affected.

Implications for Sensory Perception

Specific enhancement of feedforward connectivity has important implications in processing sensory information. Sensory perception is derived from both bottom-up sensory inputs and top-down stimulus expectations (Kording and Wolpert, 2004; Stocker and Simoncelli, 2006). Previous theoretical work indicates that sensory cortical neurons could integrate multiple sources of information by linear summation of population responses activated by each source. To achieve optimal integration under this scheme, however, the weight placed on each information source must be dynamically adjusted according to the quality of the information and task demands (Ma et al., 2006). It is interesting that the firing rate of fast-spiking neurons-likely, PV+ neurons (Kawaguchi and Kubota, 1998; Toledo-Rodriguez et al., 2004)-appears to increase with demand of attention to external stimuli (Chen et al., 2008; Mitchell et al., 2007). Our results show that activation of PV+ neurons preferentially emphasizes bottom-up sensory information by increasing feedforward connectivity without changing the weight on top-down information presumably supplied through lateral or feedback connections. Thus, PV+ neurons may play an important role in optimal integration of sensory information

⁽B) Random reduction of spikes in all layers does not alter sound-to-site couplings.

⁽C) Removing sound-evoked periods from the spike trains does not change the effect of stimulation of PV+ neurons on site-to-site couplings, highlighting the importance of network connectivity over stimulus input. Data are presented as mean \pm SEM.

with top-down expectations in sensory perception. These results could inform future work on mechanisms of sensory pathologies in patients with autism and schizophrenia, both of which are associated with PV+ neuron dysfunction (Gandal et al., 2012; Gonzalez-Burgos and Lewis, 2012).

EXPERIMENTAL PROCEDURES

Subjects

The University of California, Berkeley (UC Berkeley), Animal Care and Use Committee approved all procedures. Subjects included 11 adult PV-Cre mice (strain B6;129P2-Pvalbtm1(cre)Arbr/J; Jackson Laboratory), aged approximately postnatal day 100 (~P100) at the time of recordings. Eight mice received an injection in the right auditory cortex at ${\sim}P60$ with 1 ${\mu}l$ of a Cre-dependent adeno-associated viral vector carrying a double-floxed inverted copy of the light-sensitive cation channel channelrhodopsin-2 [pAAV-EF1a-DIO-hChR2(H134R)-EYFP-WPRE-pA; 8 \times 10¹² viral particles per milliliter, University of North Carolina Vector Core] using a glass micropipette (Drummond Wiretrol, 10 µl) attached to a Quintessential Stereotaxic Injector (Stoelting) and procedures described elsewhere (Cardin et al., 2010). To control for the effect of light stimulation or heating of the cortex in general, we injected three mice with saline using the same protocol. A small burr hole (0.7 mm in diameter) was made over the right auditory cortex (1.75 mm rostral to lambda on the temporal ridge; Franklin and Paxinos, 2008), and virus (or saline) was delivered through a small durotomy. Each injection was performed in two stages, with 0.5 μ l of virus injected at a depth of 500 μ m and the remaining 0.5 μl injected at 250 $\mu m,$ at a rate of 0.1 $\mu l/s.$ Recordings were obtained after an infection period of approximately 1 month to ensure adequate expression of ChR2 throughout the auditory cortex, which was confirmed using green fluorescent protein (GFP) fluorescence goggles (BLS Ltd.).

Immunohistochemistry

Mice were perfused transcardially with cold 0.01 M PBS (pH = 7.4) followed by 4% paraformaldehyde (PFA) in 0.01 M PBS. Brains were postfixed in 4% PFA for 12 hr and then cryoprotected in 30% sucrose for 18 hr. Free-floating sections (50 μ m) were cut using a cryostat (Leica CM3050). Every other section was incubated with blocking solution (10% normal goat serum in 0.01 M PBS with 0.1% Triton X-100) for 45 min at room temperature (20°C) and then incubated in primary antibody (PV 25 rabbit anti-parvalbumin, Swant, 1:4,000 dilution) at 4°C overnight. The next day, slices were incubated in secondary antibody (Alexa 594 goat anti-rabbit immunoglobulin G, Invitrogen, 1:200 dilution) for 1.5 hr at room temperature. Sections were mounted on gelatin-subbed glass slides with Fluoromount-G (Southern Biotech) and coverslipped. Adjacent sections not stained for PV were washed in 0.01 M PBS, mounted on slides and left to dry for 48 hr. They were then Nissl stained with 0.5% cresyl (w/v) for identification of individual cortical layers.

Quantification of Virus Expression

We scored the spread of the virus by hand by analyzing each 50 μm coronal section for the presence of enhanced yellow fluorescent protein (eYFP) fluorescence using a Zeiss LSM 780 34-channel AxioExaminer fixed stage upright confocal microscope (UC Berkeley Molecular Imaging Center). Colocalization of ChR2-eYFP to PV+ cells was analyzed by acquiring confocal images and identifying cells from each fluorescence channel by hand using ImageJ's cell counter plug-in.

Electrophysiological Recording and Stimuli

The right auditory cortex was mapped for each mouse under anesthesia using a cocktail of ketamine (100 mg/kg) and xylazine (10 mg/kg) and procedures described previously (Han et al., 2007). Following deflection of the temporal muscle, exposure of the auditory cortex, and removal of the dura mater, we performed a coarse mapping with tungsten electrodes (FHC) to determine the location of primary auditory cortex based on rostrocaudal tonotopy and short spike latencies. We then recorded extracellular multiunit neural activity in putative layer 2/3 through layer 4 of the right primary auditory cortex (Franklin and Paxinos, 2008; Oviedo et al., 2010) using a 4 × 4 silicone polytrode

(NeuroNexus A4×4-3mm-100-125-177). Fourteen of 16 channels showed normal impedance measurements and were included in the analysis. A total of 350 multiunit sites (294 from ChR2-transfected animals, 56 from saline-injected controls) were used in our analyses. The polytrode was oriented parallel to the tonotopic axis and lowered orthogonally to the cortex so that the deepest contact sites were at a depth of approximately 500 μ m from the pial surface (Figure 1A). The extracellular signal was obtained using a TDT amplifier connected to TDT RX5 hardware (Tucker Davis Technologies), using a sampling rate of 25 kHz. Spike times were calculated by thresholding the extracellular signal at 1.5 times the SD of the SNR and bandpass filtering between 300 and 3,000 Hz, and they were logged using custom software running on a Windows XP computer. Sound stimuli were presented to the contralateral ear through an electrostatic cannulated speaker (EC1, Tucker Davis Technologies) controlled by TDT RX6 hardware and calibrated to ensure less than 3% spectral distortion and a flat output (<3 dB deviation) from 4 to 75 kHz (Brüel and Kjær microphone, preamplifier, and conditioning amplifier, with SigCal32 software). Sound stimuli were pure tones generated using MATLAB (25 ms length with 5 ms squared-cosine ramp, sampling rate, 156.25 kHz) played from 4 kHz to 75 kHz in 0.2 octave steps, for a total of 22 frequencies. Sounds were presented at six different loudness levels (20-70 dB SPL, 10 dB spacing) in a pseudorandom order with a 1 Hz repetition rate, and each frequency-intensity pair was repeated three times. For the 50 dB level, stimuli were presented an additional 12 times to obtain higher resolution data at this intermediate level. For each 1 s trial, a tone pip would play at 500 ms into the trial. For half of the trials, we stimulated ChR2-transfected PV+ neurons using a 500 ms pulse of 473 nm blue laser light (Shanghai Laser and Optics Century Co., model BL473T3) coupled to a 200 µm optic fiber (ThorLabs, BFL37-200) beginning at 250 ms into the trial and controlled by a transistor-transistor logic (TTL) pulse delivered by the RX5 hardware. This stimulation protocol results in the continuous spiking of the PV+ neurons throughout the duration of the light pulse (Zhao et al., 2011). The laser output was calibrated using a power meter (ThorLabs, PM100D with sensor S120C and neutral density filter NE03A-A) to deliver light at an intensity of 1.2 mW, or \sim 40 mW/mm². This light intensity was chosen as the minimal light level that induced silencing of cortical activity throughout the light stimulation period. Photoelectric light artifacts (sharp transients locked to the onset of the light stimulus) were removed by excluding time points immediately surrounding the light onset (Cardin et al., 2010). Classical receptive fields were calculated for "light-on" and "light-off" trials separately by counting the number of spikes elicited by each frequency-intensity pair in a window defined by the peak of the poststimulus time histogram. Receptive field thresholds were defined as the minimum sound intensity required to evoke a response (the intensity at the tip of the V-shaped receptive field). The receptive field bandwidths were calculated as the width of the frequency response in octaves 20 dB above the intensity threshold. Detection SNR was defined as (number of evoked spikes - number of spontaneous spikes)/(number of spontaneous spikes) for "light-on" and "light-off" epochs separately.

Ising Model Fitting

Binary matrices of the sound stimulus condition and spiking data for "light-off" and "light-on" trials were separately used as input to the model. Spiking activity from each of the contact sites on the polytrode was binned into 5 ms time epochs so that, for each bin, if a spike occurred during those 5 ms, the value was set to 1, and if no spikes occurred, the value was set to 0. The bin duration of 5 ms was chosen according to the average cross-correlation between all pairs of recording sites, which showed that the mean cross-correlation was approximately three times the SD of the baseline correlation at 5 ms. The sound stimulus matrix consisted of values for 22 frequency bins for each of t time points, with values of 1 when a stimulus at a given frequency was present and 0 when the stimulus was absent. Since cortical responses to sound occur with a delay relative to stimulus onset, we set the sound matrix to 1 for a window starting at 15 ms after the onset of the sound and ending at 50 ms after sound onset, corresponding to when the cross-correlation between the sound stimulus onset and neural spiking responses reached approximately three times the SD of the baseline cross-correlation. Fitting separate sound-to-site couplings for each time delay relative to the stimulus onset (from 0 to 100 ms after sound onset, see Supplemental Experimental Procedures) did not change our result (Figure S1). Only responses to the three highest decibel levels were used in the model (50, 60, and 70 dB). For each polytrode site, trials were randomized, the full data matrix was split into ten equal chunks, and each model was estimated by holding out one of the data chunks, training on the remaining 90% of the data, and repeating this process ten times for each possible training and validation set. This method, called 10-fold cross-validation (Kohavi, 1995), was used to ensure an accurate estimate of the log-likelihood that is more robust to noise in the data.

The stimulus-conditioned Ising model is defined as follows:

$$\rho(\mathbf{x}|\mathbf{s};\mathbf{J},\mathbf{W}) = \frac{1}{Z(\mathbf{s},\mathbf{J},\mathbf{W})} \exp(\mathbf{x}^{\mathsf{T}}\mathbf{J}\mathbf{x} + \mathbf{x}^{\mathsf{T}}\mathbf{W}\mathbf{s}), \quad \text{(Equation 1)}$$

where $\mathbf{x} \in \{0,1\}^N$ is the binary spike pattern for each time point, N is the number of recording sites, $\mathbf{J} \in \mathbb{R}^{N \times N}$ is the site-to-site coupling matrix, $\mathbf{W} \! \in \! \mathbb{R}^{N \times M}$ is the sound-to-site coupling matrix, M is the number of stimulus dimensions (in this case, sound frequencies presented), and $\mathbf{s} \in \{0, 1\}^{M}$ is the stimulus input vector. A positive coupling value $J_{ij} > 0$ indicates that sites x_i and x_j tend to be active simultaneously, while a coupling $J_{ij} < 0$ indicates that when a spike occurs at site x_i , x_j will be more likely to remain silent and vice versa. Similarly, a positive sound-to-site coupling value W_{ii} indicates that spiking in x_i tends to increase during presentation of stimulus s_j , while a negative sound-to-site coupling value of W_{ii} indicates that spiking in x_i is suppressed during presentation of s_i. Both site-to-site and sound-to-site couplings are unitless (much like linear regression weights, for example), with the magnitude of coupling indicating the strength of the relationship between their firing patterns. We were interested in the effect of the light stimulus condition on coupling, so separate coupling matrices ${\bf J}$ and ${\bf W}$ were trained for the "light-off" and "light-on" trials.

To estimate the couplings, we used minimum probability flow learning (MPF) (Schaub and Schultz, 2012; Sohl-Dickstein et al., 2011a, 2011b) to minimize an L1 regularized version of the MPF objective function,

$$\begin{split} \mathcal{K}(\mathbf{J}, \ \mathbf{W}) = & \frac{1}{T} \sum_{\{\mathbf{x}, \mathbf{s}\}} \sum_{\mathbf{x}' \in \mathcal{N}(\mathbf{x})} \exp\left(\frac{1}{2} [E(\mathbf{x}|\mathbf{s}; \ \mathbf{J}, \ \mathbf{W}) - E(\mathbf{x}'|\mathbf{s}; \ \mathbf{J}, \ \mathbf{W})]\right) \\ & + \lambda(\|\mathbf{W}\|_1 + \|\mathbf{J}\|_1) \end{split}$$

(Equation 2)

where the sum over {**x**, **s**} indicates a sum over all training observations, the neighborhood $\mathcal{N}(\mathbf{x})$ includes all states which differ from **x** by a single bitflip, and the single state in which all bits are flipped, $E(\mathbf{x}|\mathbf{s}; \mathbf{J}, \mathbf{W}) = -\mathbf{x}^T \mathbf{J} \mathbf{x} - \mathbf{x}^T \mathbf{W} \mathbf{s}$ is the energy function of the Ising model, λ is the regularization strength, and T indicates the total number of training samples (in 5-ms binned time points).

The L1 regularization term $\lambda(\|\mathbf{J}\|_1 + \|\mathbf{W}\|_1)$ was included to prevent overfitting to training data. Lambda (λ) was chosen by cross-validation from ten values logarithmically spaced between 10^{-7} and 10^{-2} . Cross-validation was performed by holding out 20% of the training data, training the model using the remaining 80%, repeating this five times, and choosing the λ with the best average log-likelihood across all light conditions and all sites. The choice of λ had little effect on the log-likelihoods of the model fit for "light-off" trials, but there was improvement for the "light-on" models at intermediate λ values. Thus, we chose to use the same value of λ regardless of light condition. Lambda (λ) was set to 5.9 × 10⁻⁵.

Following selection of the regularization parameter, we fit the model using all of the training data, and the model log-likelihood, conditioned on the stimulus, was tested on the held out validation set. This was repeated ten times for different validation sets, using the same regularization parameter. Coupling matrices shown in the figures are taken from the cross-validation iteration with the highest conditional likelihood on the validation set. We evaluated model likelihoods on held-out data,

$$\log L = \frac{1}{T} \sum_{\{\mathbf{x}, \mathbf{s}\}} \log p(\mathbf{x} | \mathbf{s}; \mathbf{J}, \mathbf{W})$$
 (Equation 3)

The normalization constant **Z**(**s**, **J**, **W**) required in the calculation of p(**x** | **s**; **J**, **W**) (Equation 1) was computed by exhaustive summation over all 2^{14} possible spiking states.

To test the effect of lowered baseline activity on Ising model couplings, we removed 20%, 50%, and 80% of spikes in all rows. Spikes were removed at random for each channel separately and included both spontaneous and evoked data. We then reran the Ising model for the new manipulated spike data using cross-validation as before and tested performance on a held-out set that had been manipulated similarly (20%–80% spikes removed).

To test the effect of evoked activity, we removed all time points between 15 and 50 ms after sound stimulus onset for each trial and fixed sound couplings to zero while training the model.

VAR Modeling

To investigate the directionality of functional connections in the auditory cortex, we used VAR models (Lütkepohl, 2005), which are linear models that relate multiple, dependent time series, such as spike trains, by a sum of linear weights. Such models allow us to predict the spiking activity of each site in the polytrode as a function of the previous spiking activity at all other sites. We fit models of the form:

$$\widehat{x}_{i}(t) = \overline{x}_{i} + \sum_{j=1}^{N} \sum_{\tau=1}^{T} \beta_{i}(\tau, j) x_{j}(t-\tau)$$
 (Equation 4)

where $\hat{x}_i(t)$ is the estimated response at recording site *i* at time *t*, \bar{x}_i is the baseline firing rate of that site, β_i is a matrix of linear weights for the *N* simultaneously recorded sites over each of *T* time delays, and $x_j(t - \tau)$ is the response at recording site *j* at a given time in the past, $(t - \tau)$. *T* is the total number of time delays included in the analysis, and *N* is the total number of simultaneously corded sites. We used delays up to 40 ms for each set of 14 simultaneously recorded sites. Those familiar with spectrotemporal receptive field (STRF) estimation will recognize this model as being essentially identical to a STRF (Aertsen and Johannesma, 1981; Theunissen et al., 2001; Wu et al., 2006), with the difference being that neural activity is predicted from other activity in the network rather than by a parameterization of the external stimulus.

To solve for the VAR weights, we used ridge regression, which is less prone to overfitting than ordinary least-squares. Ridge regression, also known as L2-penalized or Tikhonov regularization, minimizes the mean squared error between the actual and estimated response while constraining the L2 norm of the regression weights. The strength of the L2 penalty is determined by the ridge parameter, $\lambda \ge 0$, where larger values of λ result in greater shrinkage of the weights (Asari and Zador, 2009; Machens et al., 2002; Wu et al., 2006). In ridge regression, we minimize the following error function:

$$E(\beta_{i}) = \|x_{i}(t) - \hat{x}_{i}(t)\|_{2}^{2} + \lambda \|\beta_{i}\|_{2}^{2}$$
 (Equation 5)

where $x_i(t)$ is the true response of site *i* at time *t* and the estimated response $\hat{x}(t)$ is given by Equation 4. We estimated VAR weights using 80% of the data as a training set. Of the remaining 20% of the data, half was used for fitting the ridge parameter (10%) and half was used as a validation set to assess model performance (10%). The same recordings used in the Ising model were used in these analyses. Input to the model consisted of the binary spike trains binned at 2 ms for each of the channels on the polytrode. Separate models were fit for "light-on" and "light-off" trials. To find the optimal ridge parameter, we tested ten logarithmically spaced ridge parameters between 10⁻² and 10⁵ and then selected the value that resulted in the highest average correlation on the (ridge) test set across all sites on the polytrode and both "light-on" and "light-off" models.

Since a neuron's output is strictly positive and may scale nonlinearly with the input stimulus, we added an output nonlinearity to the model fitting (Wu et al., 2006). This nonlinearity does not change the weights of the model but rather rescales the response predicted by the linear model to more accurately match the true response. We fit the nonlinearity as a univariate cubic spline that minimized the mean squared error between the actual and predicted responses on the training data. For both "light-on" and "light-off" models, adding the output nonlinearity significantly increased the predictive performance of the model ($p = 4.6 \times 10^{-10}$ and $p = 4.4 \times 10^{-16}$ for "light-off" and "light-on," respectively, Wilcoxon signed-rank test), though these increases were quite small (0.6% ± 0.1% increase for "light off," and 1.5% ± 0.1% increase for "light on"). The

increase in correlation was significantly higher for "light on" over "light off" ($p = 6.4 \times 10^{-13}$, Wilcoxon rank sum test), which is likely due to the overall lower firing rate during "light-on" trials.

VAR model validation was performed by calculating the correlation coefficient between the response predicted by the model and the actual response on the held-out validation set. Significance of the correlation between predicted and actual responses was determined using resampling. The predicted response was randomly reshuffled 100,000 times, and the correlation between the shuffled prediction and actual response was computed. Reshuffling was done using 526 ms (263 time bin) segments to preserve local temporal statistics (this length was chosen to limit accidental alignment of the 1,000 ms stimulation protocol across shuffled samples). The p value of the model prediction was then computed as the fraction of the 100,000 shuffled correlations that were higher than the actual correlation.

Statistical Tests

To test differences in coupling, we used Wilcoxon rank sum tests (for comparing independent groups) or Wilcoxon signed-rank tests (for comparing paired groups) and corrected for multiple comparisons using Bonferroni correction. Parametric tests were not used because it was determined that the data being compared were not Gaussian distributed (Lilliefors test). Resampling techniques were used to obtain confidence intervals on correlation coefficients. Spearman rank correlations were used to test relationships between monotonically but not linearly related data, such as correlations and couplings in Figure 2D. Values are reported as mean \pm SEM unless otherwise stated.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and three figures and can be found with this article online at http://dx.doi. org/10.1016/j.neuron.2013.08.017.

AUTHOR CONTRIBUTIONS

L.S.H. and S.B. contributed to the study design. L.S.H. collected the data and performed the electrophysiological experiments. V.M.C. and L.S.H. performed the immunohistochemistry and histology. L.S.H., J.S.D., and A.G.H. wrote code to fit the models and analyzed the data. K.D. provided the original ChR2 construct. L.S.H. and S.B. wrote the manuscript. All authors discussed and commented on the manuscript.

ACKNOWLEDGMENTS

The authors thank Fritz Sommer, Michael DeWeese, Daniel Feldman, David Larue, Asako Miyakawa, Robert Gibboni, Aaron Koralek, and Colleen Kirkhart for helpful comments and advice. The authors also thank Seung-Hee Lee for advice on virus injection procedures and histology and Trevor Flynn for assistance in cell counting analyses. This work was supported by the National Institute on Deafness and Other Communication Disorders (DC009259), a William Orr Dingwall Neurolinguistics Dissertation Fellowship (to A.G.H), and a National Science Foundation Graduate Research Fellowship (to L.S.H.).

Accepted: August 19, 2013 Published November 20, 2013

REFERENCES

Aertsen, A.M., and Johannesma, P.I. (1981). The spectro-temporal receptive field. A functional characteristic of auditory neurons. Biol. Cybern. *42*, 133–143. Asari, H., and Zador, A.M. (2009). Long-lasting context dependence constrains neural encoding models in rodent auditory cortex. J. Neurophysiol. *102*, 2638–2656.

Atallah, B.V., Bruns, W., Carandini, M., and Scanziani, M. (2012). Parvalbuminexpressing interneurons linearly transform cortical responses to visual stimuli. Neuron 73, 159–170. Cardin, J.A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., and Moore, C.I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. Nature *459*, 663–667.

Cardin, J.A., Carlen, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.H., and Moore, C.I. (2010). Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2. Nat. Protoc. *5*, 247–254.

Celio, M.R. (1986). Parvalbumin in most gamma-aminobutyric acid-containing neurons of the rat cerebral cortex. Science 231, 995–997.

Chen, Y., Martinez-Conde, S., Macknik, S.L., Bereshpolova, Y., Swadlow, H.A., and Alonso, J.M. (2008). Task difficulty modulates the activity of specific neuronal populations in primary visual cortex. Nat. Neurosci. *11*, 974–982.

Douglas, R.J., and Martin, K.A. (1991). A functional microcircuit for cat visual cortex. J. Physiol. *440*, 735–769.

Douglas, R.J., Martin, K.A.C., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. Neural Comput. *488*, 480–488.

Franklin, K.B.J., and Paxinos, G. (2008). The Mouse Brain in Stereotaxic Coordinates, Third Edition. (Boston: Elsevier/Academic Press).

Fuchs, E.C., Zivkovic, A.R., Cunningham, M.O., Middleton, S., Lebeau, F.E., Bannerman, D.M., Rozov, A., Whittington, M.A., Traub, R.D., Rawlins, J.N., et al. (2007). Recruitment of parvalbumin-positive interneurons determines hippocampal function and associated behavior. Neuron *53*, 591–604.

Gandal, M.J., Nesbitt, A.M., McCurdy, R.M., and Alter, M.D. (2012). Measuring the maturity of the fast-spiking interneuron transcriptional program in autism, schizophrenia, and bipolar disorder. PLoS ONE 7, e41215.

Ganmor, E., Segev, R., and Schneidman, E. (2011a). The architecture of functional interaction networks in the retina. J. Neurosci. *31*, 3044–3054.

Ganmor, E., Segev, R., and Schneidman, E. (2011b). Sparse low-order interaction network underlies a highly correlated and learnable neural population code. Proc. Natl. Acad. Sci. USA *108*, 9679–9684.

Gonzalez-Burgos, G., and Lewis, D.A. (2012). NMDA receptor hypofunction, parvalbumin-positive neurons, and cortical gamma oscillations in schizo-phrenia. Schizophr. Bull. *38*, 950–957.

Gray, C.M., Konig, P., Engel, A.K., and Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. Nature *338*, 334–337.

Han, Y.K., Köver, H., Insanally, M.N., Semerdjian, J.H., and Bao, S. (2007). Early experience impairs perceptual discrimination. Nat. Neurosci. *10*, 1191– 1197.

Isaacson, J.S., and Scanziani, M. (2011). How inhibition shapes cortical activity. Neuron 72, 231–243.

Kawaguchi, Y., and Kubota, Y. (1998). Neurochemical features and synaptic connections of large physiologically-identified GABAergic cells in the rat frontal cortex. Neuroscience *85*, 677–701.

Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. Proceedings of the 14th International Joint Conference on Artificial Intelligence *2*, 1137–1143.

Kording, K.P., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. Nature 427, 244–247.

Köster, U., Sohl-Dickstein, J., and Olshausen, B. (2012). Higher order correlations within cortical layers dominate functional connectivity in microcolumns. arXiv, arXiv:1301.0050, http://arxiv.org/abs/1301.0050.

Lee, C.C., and Winer, J.A. (2008). Connections of cat auditory cortex: III. Corticocortical system. J. Comp. Neurol. 507, 1920–1943.

Lee, S.H., Kwan, A.C., Zhang, S., Phoumthipphavong, V., Flannery, J.G., Masmanidis, S.C., Taniguchi, H., Huang, Z.J., Zhang, F., Boyden, E.S., et al. (2012). Activation of specific interneurons improves V1 feature selectivity and visual perception. Nature 488, 379–383.

Linden, J.F., and Schreiner, C.E. (2003). Columnar transformations in auditory cortex? A comparison to visual and somatosensory cortices. Cereb. Cortex 13, 83–89.

Lütkepohl, H. (2005). New Introduction to Multiple Time Series Analysis. (Berlin: Springer).

Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. Nat. Neurosci. 9, 1432–1438.

Machens, C.K., Wehr, M., and Zador, A.M. (2002). Spectro-temporal receptive fields of subthreshold responses in auditory cortex. In Advances in Neural Information Processing Systems, *Volume 15*, (Cambridge, MA: The MIT Press).

Marre, O., El Boustani, S., Fregnac, Y., and Destexhe, A. (2009). Prediction of spatiotemporal patterns of neural activity from pairwise correlations. Phys. Rev. Lett. *102*, 138101.

McIntosh, A.R., and Gonzalez-Lima, F. (1998). Large-scale functional connectivity in associative learning: interrelations of the rat auditory, visual, and limbic systems. J. Neurophysiol. *80*, 3148–3162.

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

Mountcastle, V.B. (1957). Modality and topographic properties of single neurons of cat's somatic sensory cortex. J. Neurophysiol. *20*, 408–434.

Ohiorhenuan, I.E., Mechler, F., Purpura, K.P., Schmid, A.M., Hu, Q., and Victor, J.D. (2010). Sparse coding and high-order correlations in fine-scale cortical networks. Nature *466*, 617–621.

Oviedo, H.V., Bureau, I., Svoboda, K., and Zador, A.M. (2010). The functional asymmetry of auditory cortex is reflected in the organization of local cortical circuits. Nat. Neurosci. *13*, 1413–1420.

Packer, A.M., and Yuste, R. (2011). Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? J. Neurosci. *31*, 13260–13271.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999.

Roudi, Y., Tyrcha, J., and Hertz, J. (2009a). Ising model for neural data: model quality and approximate methods for extracting functional connectivity. Phys. Rev. E Stat. Nonlin. Soft Matter Phys. 79, 51915.

Roudi, Y., Nirenberg, S., and Latham, P.E. (2009b). Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. PLoS Comput. Biol. *5*, e1000380.

Schaub, M.T., and Schultz, S.R. (2012). The Ising decoder: reading out the activity of large neural ensembles. J. Comput. Neurosci. 32, 101–118.

Schneidman, E., Berry, M.J., 2nd, Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. Nature *440*, 1007–1012.

Shlens, J., Field, G.D., Gauthier, J.L., Grivich, M.I., Petrusca, D., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2006). The structure of multi-neuron firing patterns in primate retina. J. Neurosci. *26*, 8254–8266.

Shlens, J., Field, G.D., Gauthier, J.L., Greschner, M., Sher, A., Litke, A.M., and Chichilnisky, E.J. (2009). The structure of large-scale synchronized firing in primate retina. J. Neurosci. 29, 5022–5031.

Smith, P.H., and Populin, L.C. (2001). Fundamental differences between the thalamocortical recipient layers of the cat auditory and visual cortices. J. Comp. Neurol. *436*, 508–519.

Sohal, V.S., Zhang, F., Yizhar, O., and Deisseroth, K. (2009). Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. Nature *459*, 698–702.

Sohl-Dickstein, J., Battaglino, P., and DeWeese, M.R. (2011a). Minimum probability flow learning. Proceedings of the 28th International Conference on Machine Learning, Bellevue, WA, 905–912.

Sohl-Dickstein, J., Battaglino, P.B., and DeWeese, M.R. (2011b). New method for parameter estimation in probabilistic models: minimum probability flow. Phys. Rev. Lett. *107*, 11–14.

Stocker, A.A., and Simoncelli, E.P. (2006). Noise characteristics and prior expectations in human visual speed perception. Nat. Neurosci. 9, 578–585.

Tang, A., Jackson, D., Hobbs, J., Chen, W., Smith, J.L., Patel, H., Prieto, A., Petrusca, D., Grivich, M.I., Sher, A., et al. (2008). A maximum entropy model applied to spatial and temporal correlations from cortical networks in vitro. J. Neurosci. *28*, 505–518.

Theunissen, F.E., David, S.V., Singh, N.C., Hsu, A., Vinje, W.E., and Gallant, J.L. (2001). Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. Network *12*, 289–316.

Toledo-Rodriguez, M., Blumenfeld, B., Wu, C., Luo, J., Attali, B., Goodman, P., and Markram, H. (2004). Correlation maps allow neuronal electrical properties to be predicted from single-cell gene expression profiles in rat neocortex. Cereb. Cortex *14*, 1310–1327.

Wilson, N.R., Runyan, C.A., Wang, F.L., and Sur, M. (2012). Division and subtraction by distinct cortical inhibitory networks in vivo. Nature 488, 343–348.

Winer, J.A., and Lee, C.C. (2007). The distributed auditory cortex. Hear. Res. 229, 3–13.

Womelsdorf, T., Schoffelen, J.M., Oostenveld, R., Singer, W., Desimone, R., Engel, A.K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. Science *316*, 1609–1612.

Wu, M.C., David, S.V., and Gallant, J.L. (2006). Complete functional characterization of sensory neurons by system identification. Annu. Rev. Neurosci. 29, 477–505.

Zhao, S., Ting, J.T., Atallah, H.E., Qiu, L., Tan, J., Gloss, B., Augustine, G.J., Deisseroth, K., Luo, M., Graybiel, A.M., et al. (2011). Cell type-specific channelrhodopsin-2 transgenic mice for optogenetic dissection of neural circuitry function. Nat. Methods *8*, 745–752.

Zucker, R.S., and Regehr, W.G. (2002). Short-term synaptic plasticity. Annu. Rev. Physiol. *64*, 355–405.