



Cortical ensembles selective for context

Jordan P. Hamm^{a,1}, Yuriy Shymkiv^a, Shuting Han^a, Weijian Yang^{a,2}, and Rafael Yuste^a

^aDepartment of Biological Sciences, Columbia University, New York, NY 10027

Edited by Ranulfo Romo, National Autonomous University of Mexico, Mexico City, D.F., Mexico, and approved February 26, 2021 (received for review December 19, 2020)

Neural processing of sensory information is strongly influenced by context. For instance, cortical responses are reduced to predictable stimuli, while responses are increased to novel stimuli that deviate from contextual regularities. Such bidirectional modulation based on preceding sensory context is likely a critical component or manifestation of attention, learning, and behavior, yet how it arises in cortical circuits remains unclear. Using volumetric two-photon calcium imaging and local field potentials in primary visual cortex (V1) from awake mice presented with visual “oddball” paradigms, we identify both reductions and augmentations of stimulus-evoked responses depending, on whether the stimulus was redundant or deviant, respectively. Interestingly, deviance-augmented responses were limited to a specific subset of neurons mostly in supragranular layers. These deviance-detecting cells were spatially intermixed with other visually responsive neurons and were functionally correlated, forming a neuronal ensemble. Optogenetic suppression of prefrontal inputs to V1 reduced the contextual selectivity of deviance-detecting ensembles, demonstrating a causal role for top-down inputs. The presence of specialized context-selective ensembles in primary sensory cortex, modulated by higher cortical areas, provides a circuit substrate for the brain’s construction and selection of prediction errors, computations which are key for survival and deficient in many psychiatric disorders.

predictive coding | circuits | neocortex

In the mammalian brain, sensorineural processing is significantly influenced by context. For instance, sensory processing circuits tend to suppress processing of predictable or “redundant” stimuli (e.g., tree branches in a forest), while amplifying responses to contextually salient stimuli (like an airborne predator). This strategy conserves energy, while also allowing stimuli with potential behavioral or survival relevance to “stand out” and garner additional neuronal resources.

In a healthy neocortex, neural responses to incident sensory stimuli are modulated by past experience on short (0.01 to 10 s) as well as long (>10 s) time scales. For instance, in a classic sensory “oddball” paradigm, repetition of a given stimulus at a rate of 0.1 to 2 Hz results in a phenomenon termed “stimulus specific adaptation” (SSA), wherein response magnitudes decrease rapidly to the repeated or “adapted” stimulus (1–3). In contrast, when a stimulus deviates from the established contextual regularities (e.g., probability of occurrence), a phenomenon termed “deviance detection” (DD) is observed, wherein responses are augmented beyond the typical magnitude observed in a neutral context (1).

Theoretical work has sought to explain these phenomena in a “predictive coding” framework, wherein a generative model of the environment, embedded in an increasingly hierarchical cortical network, serves to suppress sensory cortical responses to predictable stimuli while allowing responses to contextually deviant stimuli (i.e., prediction errors) to propagate and update the model (4). Experimentally, these phenomena have been studied at either the single-neuron level in animals (1–3) or at the brain-wide level with electroencephalography (EEG) or magnetoencephalography (MEG) in humans (5, 6). Resulting work highlights a role of *N*-methyl-D-aspartate (NMDA) receptors (7) as well as gamma-Aminobutyric acid (GABA)-ergic interneurons for SSA (3) and DD (1).

Still, much remains unclear about how DD and SSA are computed within cortical circuits (8). For instance, are response decreases or increases (to redundancy and deviance, respectively) expressed evenly within responsive neurons in a local cortical region, mirroring the gross-level readout of EEG event-related potentials, or is there some division of labor within a cortical region, with subsets of neurons expressing DD or “prediction error” (9)? A recent study in mouse primary visual cortex (V1) investigating sensory–motor mismatch suggests that a subset of layer 2/3 neurons respond selectively when visual stimuli do not match locomotor predictions (10). Whether and how this finding applies to sensory–sensory mismatch, that is, when incident stimuli do not match purely sensory-based predictions (i.e., as studied with the classic “oddball” paradigm) is unknown. Furthermore, whether and how such “prediction error” cells are selectively influenced by top-down inputs (i.e., backward projections) remains unknown as well. Such questions, and prediction error in general, carry major clinical significance since cortical SSA and DD are characteristically reduced in individuals with psychotic disorders (6, 11), potentially due to diminished integration of top-down modulation (12) or destabilized local ensembles in V1 (13–15).

Utilizing a standard sensory “oddball” paradigm commonly used in clinical neuropsychiatry, we show with volumetric two-photon calcium imaging in awake mouse V1 that, among all V1 neurons responsive to a given visual stimulus, a subset (about a third) respond selectively when the stimulus is contextually deviant. This strong deviance preference remains stable across trials and subsequent experiments, while the rest of the visually driven neurons show largely absent contextual modulation. Such “deviance detectors” display high intragroup activity correlations even in the absence of direct visual stimulation, suggesting that they

Significance

In perceiving one’s environment, stimuli are processed not in isolation but in the context in which they appear. This contextual filter on perception, like noticing novel stimuli while ignoring repeated ones, is seemingly automatic and crucial for the survival of the organism. Further, novelty detection and contextual modulation are affected in diseases like schizophrenia, potentially undermining how affected individuals relate to the world. Here, we show that a specific subset of neurons in the mammalian cortex selectively detect when a stimulus deviates from predicted regularities. Our results identify a functional group of neurons in the brain which could be targeted for novel neurotherapies for diseases of perception and thought.

Author contributions: J.P.H. and R.Y. designed research; J.P.H., Y.S., S.H., and W.Y. performed research; J.P.H. and Y.S. analyzed data; and J.P.H. and R.Y. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Published under the PNAS license.

¹To whom correspondence may be addressed. Email: jhamm1@gsu.edu.

²Present address: Department of Electrical and Computer Engineering, University of California, Davis, CA 95616.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2026179118/-DCSupplemental>.

Published April 2, 2021.

form functional neuronal ensembles. “Deviance detectors” included mainly excitatory neurons and were most prevalent in superficial layers, consistent with theoretical models of predictive coding (9). Further, cortico-cortical inputs from prefrontal regions supported context processing in V1 by selectively modulating the activity of “deviance detector” ensembles bidirectionally, enhancing their responses to contextually deviant stimuli while reducing their responses to predictable stimuli, suggesting a mediating role for local V1 inhibitory and disinhibitory circuitry (1, 8). This is consistent with a theorized role of both prefrontal cortical (PFC) and sensory cortices in mismatch negativity and schizophrenia-related psychopathology (16, 17).

Results

Contextual Modulation of Neuronal Responses in V1. We first recorded the activity of populations of cortical neurons in awake mice (Fig. 1A; $n = 15$ experiments from 10 mice; 1,642 neurons) while they viewed a classic visual oddball paradigm (100% contrast, full-field square-wave gratings oriented 0 versus 90° or 45 versus 135°; “standard” or redundant stimuli 87.5% probability, deviant stimuli [opposite orientation] 12.5% probability; Fig. 1B). In what follows, we will refer to standard stimuli as redundant stimuli, in the sense that they contain no new information. Responses in the oddball paradigm were compared to a reversed run (e.g., deviant stimuli in first run become redundant stimuli in second run, 5 min later) and to a “many standards control” (1), wherein eight different orientations were presented at the same rate with 12.5% probability, such that stimuli were neither redundant nor contextually deviant (Fig. 1C). We first employed fast two-photon calcium imaging (30 Hz resonance scanning) with GCaMP6s expressed virally under the synapsin promoter (pan-neuronal) in layer 2/3 of V1, focusing on increases in fluorescence (Fig. 1D and E), to quantify neuronal responses in the same neurons to the same oriented stimulus under three separate contexts (redundant, deviant, and neutral or “equiprobable”). Analyses focused on robustly visually driven neurons showing Z-scored average responses (during the 0.5 s of stimulus presentation) which were greater than 1.67 SDs above the prestimulus baseline (one-tailed Z-test $P < 0.05$) to at least one of two orientation stimuli in at least one of three contexts (neutral, redundant, and deviant; $n = 841$ neurons, about 51% of total cells).

When averaged across all visually driven cells in layer 2/3 and all orientations, cortical responses were overall suppressed to redundant stimuli (SSA) and amplified to deviant stimuli (DD) relative to control, equiprobable stimulation (Fig. 1F–I). Combining across all neurons, these effects were significant [SSA; focusing on first redundant stimulus in the sequence; control (cntrl) > redundant (rdnt); $t^{\text{paired}}(839) = -2.57, P < 0.01$; DD; deviant (dev) > cntrl; $t^{\text{paired}}(839) = 3.25, P < 0.001$; Fig. 1J]. The local field potential (LFP), current source density (CSD) estimate, and time-frequency decomposition of the LFP confirmed the presence of SSA and DD (SI Appendix, Fig. S1). All together, this contextual modulation during the oddball paradigm in the aggregate responses replicates past work in visual and auditory cortices in mice (1, 18) and humans (11, 19).

A Subset of V1 Neurons Are Selective for Deviant Stimuli. While population averages showed both SSA and DD, it is possible that these modulations were expressed variably across the population. In particular, we wondered whether DD was present in all neurons responsive to a given stimulus, reflecting some global modulation of neuronal excitability (7), or whether a subset of neurons selectively signaled a deviant event in the train of stimuli, akin to a “prediction error” (4). For every neuronal response to a given stimulus orientation (e.g., 45°) we calculated two orthogonal variables: DD (i.e., responses to stimulus when contextually deviant minus responses to stimulus when it was equiprobable, i.e., control context) and SSA (i.e., responses to stimulus when it was [control

or deviant] minus responses when it was contextually [redundant]). A scatterplot of these neuronal responses (Fig. 2A) revealed clear clusters of cells, including a cluster of responses showing significant DD (i.e., upper half of plot). This cluster (boundaries confirmed with k -means, SI Appendix, Fig. S2A; henceforth termed “deviance detectors,” DDs) showed strong and consistent responses to the deviant stimulus (red; Fig. 2B) while neurons outside of this group (nondeviance detectors; nDDs) showed unchanged or smaller responses to the deviant stimulus relative to the control stimulus (black; Fig. 2B). This preference for deviant events could not be easily explained by a drift in the imaging field of view, as both DDs and nDDs showed clear activity during both runs (e.g., clear response to the onset of the visual liquid-crystal display (LCD), as in Fig. 2B, but not visual stimuli).

Interestingly, neurons identified as DDs were spatially intermixed among other neurons (Fig. 2C). In a given region of V1, about 23% of recorded neurons were visually responsive to a given stimulus orientation (i.e., greater than 1.67 SDs above prestimulus baseline), with about 8% of the total showing significant DD (response to deviants > -1.67 SDs above control; Fig. 2D). DDs to a given stimulus orientation (e.g., 90°) rarely showed significant activity or DD to the orthogonal orientation (e.g., 0°) in a given experiment, and this was partially true for nDDs as well (~6% and ~18%, respectively, showing activity greater than 1.67 to both orientations; SI Appendix, Fig. S2E and F). Thus, within a given population of V1 cells showing orientation selectivity, for every orientation there are subsets of both DDs and nDDs, suggesting that orientation selectivity does not predict or explain DD. Further, DD is expressed in an orientation selective manner, such that V1 cells exhibiting DD tend to signal deviance for only select orientations.

Given the documented variability and sparseness of V1 single neuron activity (20), even among ostensibly highly tuned cells (21, 22), it is possible that neurons spontaneously active during a minority of trials in this paradigm could masquerade as having a “context” preference in the grand average (e.g., neurons silent during most of the paradigm, except for during one “deviant” trial, could appear to be spurious “deviance detectors”). We sought to confirm context preferences in DDs with cross validation, by “training” the neuron groups (or identifying the cluster boundaries) with one set of data and statistically assessing their context preferences (e.g., DD and SSA) on the rest of the data. First, we identified neurons in the DD cluster after averaging activity across the even trials (i.e., second, fourth, sixth, etc. time the stimulus appeared during the experiment) and assessed their averages to the odd trials. Deviance detectors displayed strong preference for deviant stimuli during odd trials [Fig. 2E and F; $t^{\text{paired}}(371) = 7.27, P < 0.001$]; nDDs showed absent DD but significant SSA [Fig. 2G and H; $t^{\text{paired}}(488) = -2.42, P < 0.01$]. As a more rigorous step, we repeated experiments with the same stimuli, recording from the same groups of neurons, after a 30 min break. Focusing on neurons showing stable and significant responses, we found again that DDs retained their context preferences, showing significant DD [SI Appendix, Fig. S2B, G, and H; $t^{\text{paired}}(200) = 2.54, P < 0.01$] while nDDs did not show DD [SI Appendix, Fig. S2C, D, I, and J; $t^{\text{paired}}(307) = 0.9, P > 0.05$].

Although two separate nondeviance detector groups can be ostensibly identified in the omnibus scatterplot (Fig. 2A), these two neuronal groups did not retain distinctness or context selectivity after cross validation, even when the groups were determined with k -means analyses (carried out with two to six cluster solutions; i.e., $k = 2$ to 6; SI Appendix, Fig. S2C and D). Clusters exhibiting stable DD after cross validation, however, were present at all values of k (SI Appendix, Fig. S2D; dev > cntrl; $P < 0.01$). Thus, we focused primarily on deviance detectors in subsequent analyses.

DD Ensembles in V1. We sought to determine whether cells within the same clusters showed increased synchrony across imaging sessions, suggesting the presence of functional ensembles (23) or




Fig. 1. Visual cortex neurons are modulated by stimulus context in visual "oddball" paradigm. (A) Head-fixed awake mice viewed full-field visual grating stimuli in (B) a typical oddball run (stimuli are redundant [rdnt] or deviant [dev]) and a many-standards control run (stimuli are equally probable [cntrl]). (C) Responses to the same visual stimulus were compared across these three context conditions. (D) Two-photon calcium imaging quantified (E) individual neural activity in layer 2/3 of V1 cortex. (F) Averaged ongoing activity across all visually driven neurons indicate DD (deviant > control) and stimulus-specific adaptation (control > redundant) at the population level across trials (note, as stimulus order was variable across experiments, data streams have been rearranged for simplicity and comparison across cells). (G) Trial-average responses for individual neurons and (H) averaged across neurons to the same stimulus across stimulus contexts. (I) Trial-averaged responses across all responsive neurons to three stimulus contexts, (J) averaged again within the stimulation period. *** $P < 0.001$, ** $P < 0.01$, two-tailed t test; all error bars (H–J) reflect SEM.

local subcircuits (24) selective for stimulus context. Critical for the definition of ensembles is the presence of correlated firing, not only in response to a stimulus but also in ongoing, spontaneous activity. Correlation coefficients were calculated for all

cell pairs during the visual stimulus and during "rest" (i.e., 5 min interrun intervals). Coefficients (Pearson- r values) were compared to a time-shuffled surrogate dataset for each cell pair to convert them to a Z-score [Fig. 2I (15)]. An average value for all




Fig. 2. Neurons selective for deviant stimuli form correlated ensembles in V1. (A) Scatterplot of all neuronal responses, z scored in orthogonal contrast space: DD (deviant minus control) versus SSA ([deviant and control] minus redundant). A visually identifiable cluster of neurons showing strong DD is colored in red (i.e., “deviance detectors,” DDs), with others in black (i.e., “non-deviance detectors,” nDDs). (B) Single trial (ongoing) activity over the paradigm, averaged across DDs versus nDDs, z-scored $\Delta F/F$ (see note in Fig. 1). (C) Spatial distribution of DDs (red) among other neurons in an example experiment. (D) Relative proportion of DDs, nDDs, and nonresponsive neurons (mean \pm SD across mice). Averaged responses (E, F) DDs ($n = 373$ neurons from 10 mice) and (G, H) nDDs ($n = 490$ neurons from 10 mice; identified in even trials) to deviant, redundant, and control contexts (paired t tests; *** $P < 0.001$; * $P < 0.05$). (I) Z-normalized correlations among identified DD cell-cell pairs and among nDD cell-cell pairs, calculated over a subsequent experiment. (J, K) Averages of z-normalized correlation coefficients among DDs, nDDs, and between DDs and nDDs (two-sample t test; *** $P < 0.001$; dotted line indicates average correlations among all neurons recorded). (L) Example $\Delta F/F$ of three DDs and three nDDs from a single experiment.

neuron pairs across all conditions, regardless of stimulus or context preference, after this step was about $z = 3.02$, suggesting that some baseline correlation in activity was present among all V1 cells, potentially related to global fluctuations in arousal or locomotion (25, 26). However, DD neuron activity during stimulation and rest periods was significantly more correlated with other DD neurons than with nDDs [Fig. 2 J–L; stim period: $t(1,620) = 4.39, P < 0.001$; rest: $t(1,620) = 4.36, P < 0.001$] and more correlated than nDDs were correlated with other nDDs [stim period: $t(4,525) = 4.07, P < 0.001$; rest: $t(4,525) = 4.25, P < 0.001$]. We thus concluded that DD neurons are organized into coactive ensembles in mouse V1.

DD Present with Moving Stimuli and in Excitatory Neurons. Since our initial experiments used GCaMP6s, we wondered if the slower dynamics of this indicator could contaminate or distort our results. To examine this, as a follow-up study ($n = 3$ mice, 8 fields of view; layer 2/3) utilizing GCaMP6f (VGlut-promoter; excitatory neurons)

replicated SSA [$t^{\text{paired}}(419) = -2.22, P < 0.05$] and DD [dev > cont; $t^{\text{paired}}(419) = 3.56, P < 0.001$] at the population level (SI Appendix, Fig. S3 A–G) and again revealing a subset of neurons selective for contextual deviance (SI Appendix, Fig. S3F); DD was still strong and present in this subset (identified statistically, see below) after even/odd cross validation [$t^{\text{paired}}(141) = 8.43, P < 0.001$]. Interestingly, excitatory neuron DDs were clearly present in the cell-wise scatterplot (SI Appendix, Fig. S3F), and the percentage of these showing significant DD to a given stimulus was 6% (SI Appendix, Fig. S3H; i.e., response to deviant was 1.67 SDs greater than control, versus 8.3% when measuring all neurons), suggesting that a large portion of DDs are excitatory neurons. Also, in this follow-up experiment, mice were presented with moving gratings (two cycles per second, completing a whole cycle each trial; SI Appendix, Fig. S3B) instead of static gratings used in the main study, ruling out the possibility that the spatial frequency of our visual stimuli unevenly stimulated V1 neuronal receptive fields to generate spurious context preferences.

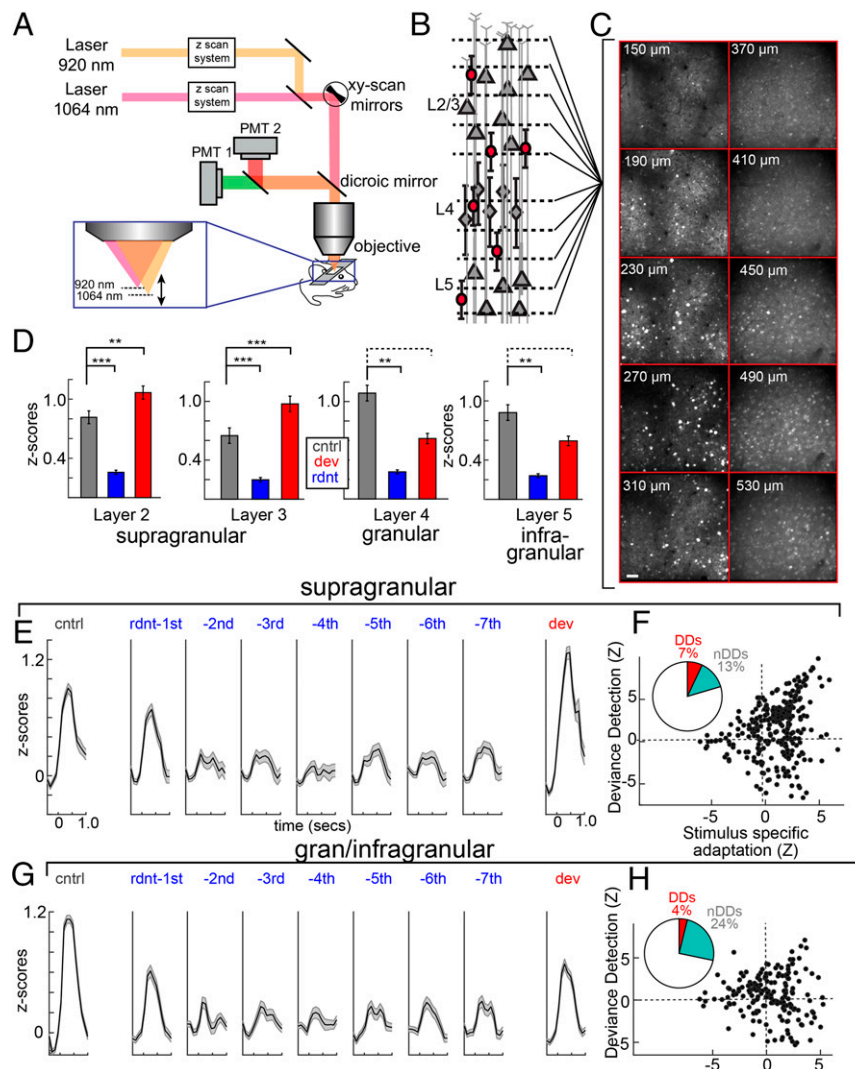


Fig. 3. DD is stronger in supragranular cortical layers. (A) Schematic of microscope setup involving two-excitation lasers with distinct z-sectioning/z-scanning strategies (SLM and ETL). (B) This enabled simultaneous imaging of neurons across 10 planes and ~400 μm of cortical depth. (C) Average images from one experiment at each depth. (Scale bar, 50 μm .) (D) Bar plots averaging across visually responsive neurons in each layer (right to left: $n = 610, n = 451, n = 529, n = 698$ neurons; when dev > cntrl $***P < 0.001, **P < 0.01$). (E) Averages of stimulus-evoked calcium transients across supragranular neurons (putative layer 2 to 3; 150 to 310 μm from surface), and (F) scatterplots showing averaged DD and SSA computations, as well as an overall pie chart showing the proportion of recorded neurons displaying no response to stimuli (white), DD (red), and responses but no DD (gray). G and H are the same as E and F but for infragranular layers (putative layers 4 to 5a; 370 to 530 μm).

DD Strongest in Supragranular Layers. To explore potential laminar difference in context modulation and DD, we also imaged deeper fields of view (GCaMP6f, moving gratings, experiments; 475 to 500 μm ; *SI Appendix, Fig. S3 H and I*; $n = 2$ mice). Infragranular neurons exhibited strong responses to stimuli (*SI Appendix, Fig. S3L*) as well as significant SSA [$t^{\text{paired}}(124) = -2.73$, $P < 0.01$] but, unlike layer 2/3 neurons, did not exhibit significant DD in the aggregate response [*SI Appendix, Fig. S3J*; SSA; DD; dev > cont; $t^{\text{paired}}(124) = 0.01$, $P = 0.99$]. Further, about 14% of neurons displayed significant visually driven activity, with only a subset of 3% of neurons displaying significant DD [*SI Appendix, Fig. S3K*; which survived even-odd cross validation, $t^{\text{paired}}(27) = 3.48$, $P < 0.01$].

This difference in DD across layers could result from laminar differences in connectivity (27) and agrees with models of prediction-error circuits in neocortical microcircuitry (9). One caveat is that our initial deep recordings were carried out after the supragranular recordings, bringing up the possibility of longer-term generalized adaptation effects previously demonstrated in this paradigm (1). We sought to systematically test the possibility of supragranular-biased DD by simultaneously recording neocortical populations across laminae within a column, utilizing two-color multiplane imaging [Fig. 3A (28, 29)] which enabled fast volumetric calcium imaging of large populations of neurons (400 to 2,000 cells) across 500 μm of depth (four mice, eight experiments) in V1. GCaMP6s was expressed in neurons in layers 1 through 3

(depth 150 to 350 μm), while jRGECO1b (a red calcium indicator) was expressed in neurons in layers 4 and 5 (depth 350 to 550 μm ; Fig. 3B and C). Imaging was carried out with two lasers (920 nm and 1,064 nm), which simultaneously scanned two different depths (putative layers 2 to 3, and 4 to 5, respectively) using a resonant scanner. The signal was separated by the emission wavelength of GCaMP6s and jRGECO1b while an electrically tunable lens and a spatial light modulator were implemented in the two beam paths to enable fast sequential scanning of different focal planes (30 to 40 μm plane separation across 150 to 530 μm depth range; Fig. 3A). A total of 10 imaging planes were imaged at a 10 Hz volume rate (Fig. 3B). Images were separated into layers (2 to 5) based on laminar boundaries and pre- and postimaging z-sectioning and later confirmed with histology (28).

All imaged layers showed robust visually driven responses. We first analyzed responses from layers 2 through 5: 150 to 230 μm (layer 2; 610 out of 2,462 cells responsive, 25%), 270 to 310 μm (layer 3; 451/1,669, 27%), 370 to 410 μm (layer 4; 529/1,580, 33%), and 450 to 530 μm (layer 5; 698/1,449, 48%). Confirming our initial experiments, aggregate responses of neurons averaged within superficial fields of view demonstrate significant SSA [Fig. 3D and E; layer 2: $t(609) = -8.30$, $P < 0.001$; layer 3: $t(450) = -5.64$, $P < 0.001$] and DD [layer 2: $t(609) = 2.60$, $P < 0.01$; layer 3: $t(450) = 2.76$, $P < 0.01$]. Neurons from deeper fields of view demonstrated SSA as well [Fig. 3D and G; layer 4: $t(528) = -9.88$, $P < 0.001$; layer 5: $t(697) = -7.89$, $P < 0.001$] but

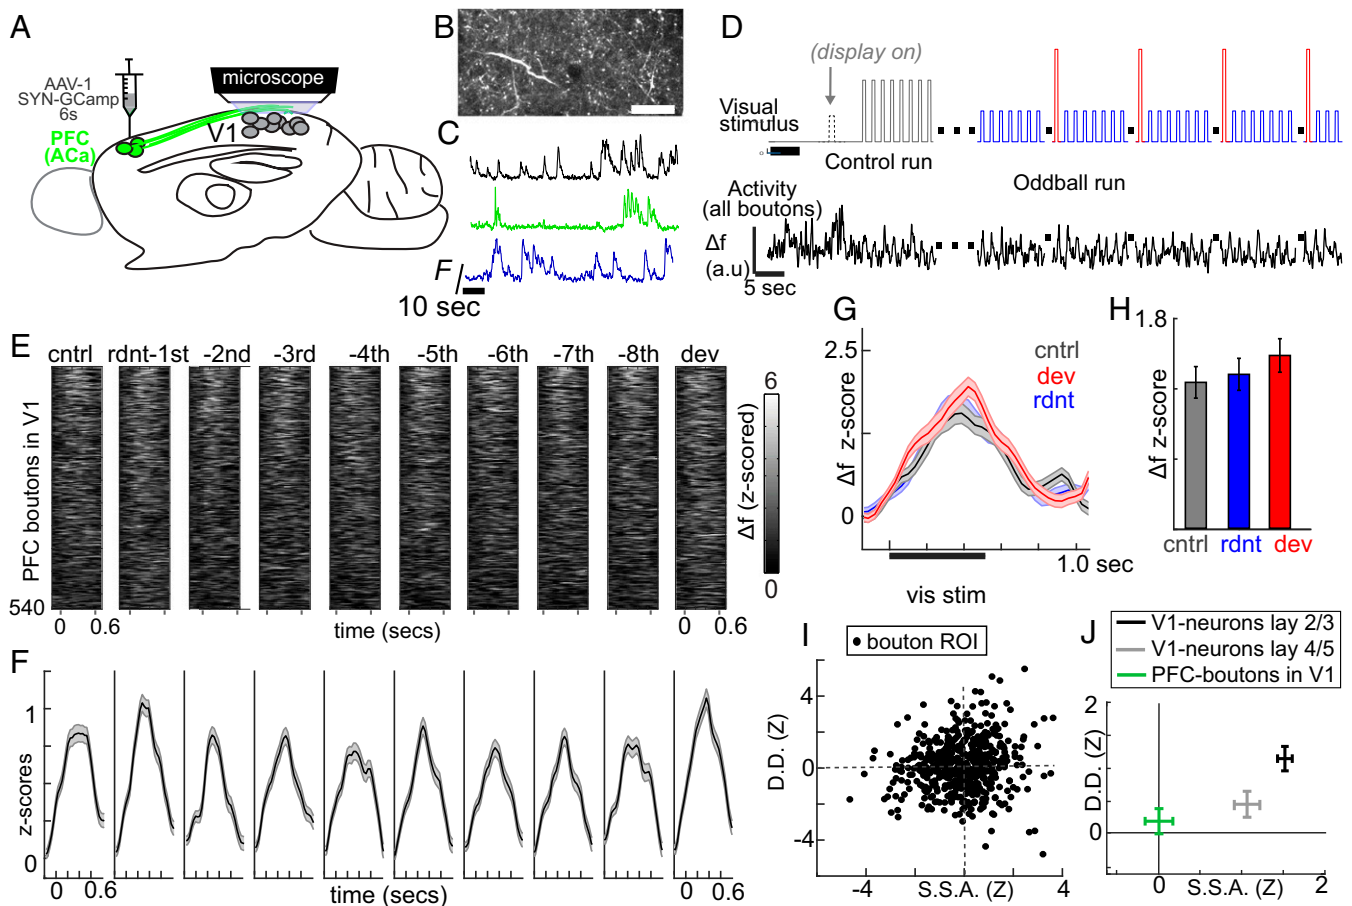


Fig. 4. Context responses of prefrontal inputs to V1. (A) Schematic of viral injection to mouse prefrontal region Aca, and (B) GCaMP6s expression in PFC axons in layer 1 of V1 (25 μm below dura) shown in an average projection (above) and (C) three example traces from boutons. (D) Average responses from responsive ROI combinations (putative-projection neurons). (E) Trial-average responses for individual neuronal responses ($n = 534$) to the same stimulus across context. (F–H) Averages of the same data from E. (I) Scatterplot of PFC neurons in SSA and DD space. (J) Comparisons of population averages of V1 supragranular (from Fig. 3F), V1 gran/infragranular (from Fig. 3H), and PFC ROIs (Fig. 4I) averaged within the stimulation period, also in SSA and DD space. ****** $P < 0.001$, ***** $P < 0.01$, two-tailed t test; all error bars (H–J) reflect SEM.

showed no significant DD. In fact, in infragranular layers, responses to deviant stimuli were significantly smaller than that during the control [layer 4: t (528) = -4.73 , $P < 0.001$; layer 5: t (697) = -3.02 , $P < 0.01$], suggesting that granular and infragranular layers express some non-stimulus-specific adaptation during the oddball paradigm, perhaps due to generalized increases in tonic inhibition affecting responses to all stimuli. Interestingly, this trend was also observed in the CSD profile from multielectrode recordings (SI Appendix, Fig. S1E).

Further, across experiments, supragranular layers displayed a significantly smaller proportion of overall cells responsive to the visual stimulus [Fig. 3 F and H; t (7) = -4.82 , $P < 0.01$], consistent with the known sparseness of this population (20) but larger proportion of cells qualifying as “deviance detectors” [Fig. 3 F and H; i.e., responses to deviant 1.67 SD larger than control; t (7) = 2.39 , $P < 0.05$]. In sum, these experiments demonstrate in the same mice, in the same cortical columns that DD neural ensembles are primarily present in supragranular layers.

Prefrontal Axons in Visual Cortex Respond to Visual Stimuli. The abundance of deviant detectors in supragranular layers could imply that they may be controlled or modulated by higher cortical areas, which preferentially target upper layers. In fact, mouse V1 receives direct top-down cortico-cortical inputs from executive brain regions, like prefrontal cortex area Cg1 [also known as anterior cingulate area, Aca (30)]. In a basic visual stimulation paradigm, these PFC-V1 inputs have a net inhibitory effect which aids in simple visual discrimination tasks (31), yet their role in modulating the activity of sensory areas in accord with context, particularly in the processing of deviant versus redundant events in the oddball paradigm, is unknown. This information could be critical for relating known deficits in context processing (e.g., attenuation of the mismatch negativity event-related potential) to PFC and cognitive dysfunction known to be present in individuals with schizophrenia, which has long been associated with PFC dysfunction (32).

We expressed GCaMP6s in mouse prefrontal cortex, targeted in stereotaxically identified Aca (Fig. 4A and SI Appendix, Fig. S4; $n = 7$ mice) and performed two-photon calcium imaging of the axonal boutons from PFC in layer 1 in the V1 [Fig. 4B and C (33, 34)]. We removed duplicate boutons with highly similar calcium traces within a single experiment in order to arrive at regions of interest (ROIs) putatively reflecting individual projection axons [see Materials and Methods (34)]. A total of 534 PFC ROIs (out of 808 recorded) showed consistent stimulus-driven activity across the control and oddball paradigms (Fig. 4D–G), but, unlike V1 neurons, PFC inputs to V1 did not show significant group-level SSA [t (533) = 0.86 , $P = 0.40$] or DD [t (533) = 0.97 , $P = 0.33$; Fig. 4H–J]. Further, a scatterplot of all ROIs suggests the absence of a clear subpopulation of deviance detectors within PFC projections (Fig. 4I).

Prefrontal Inputs Modulate Contextual Contrast in V1. Direct inputs to V1 from PFC-Aca showed visual stimulus-driven activity, which was consistent and unmodulated by stimulus context across the oddball, oddball-flip, and control paradigms, showing no clear population-wide DD or SSA (Fig. 4E–J). However, given the previously identified role of Aca in inhibiting V1 activity (31), we wondered whether PFC-Aca stimulus-driven inputs may nevertheless play a modulatory role affecting how V1 responds to stimulus context (i.e., DD or SSA), potentially by affecting some V1 subpopulations more than others (deviance detectors versus nDDs).

To test this, we expressed an optogenetic silencer, ArchT, in PFC neurons under the synapsin promoter (Fig. 5A). We illuminated the visual cortex with a cannula-directed 617 nm light-emitting diode (LED; ~ 1 mm radius, 4 mW/mm²) centered over a 2 mm craniotomy (Fig. 5A) and suppressed axons (Fig. 5B) from PFC. We confirmed this effect by tracking neural activity in PFC axons with GCaMP6s while suppressing them with coexpressed

ArchT and 2 s LED pulses while performing two-photon calcium imaging (Fig. 5C and D). Indeed, after photostimulation, PFC axons showed a clear decrease in calcium transients (averaged over cells and 20 trials, Fig. 5D; one mouse, 30 ROIs). Further, we recorded LFPs with 16-channel multielectrode arrays while suppressing axons from PFC locally with ArchT (SI Appendix, Fig. S5A; $n = 8$ mice) and compared this with recordings from mice expressing only GCaMP6s in PFC axons (SI Appendix, Fig. S5B; $n = 10$ mice; LED only controls). This technique showed that our illumination strategy also elicited a visually evoked potential present in the CSD plots even when ArchT was not expressed (SI Appendix, Fig. S5B), underscoring the necessity of an LED control condition. Nevertheless, the ArchT induced current 1) started immediately (< 6 ms) after the onset of the LED, while the visually induced response started 60 ms later; 2) had a distinctly different current source distribution (difference plot SI Appendix, Fig. S5C); and 3) was much stronger in superficial layers (SI Appendix, Fig. S5D), perhaps due to a combination of the fact that PFC axons terminate mostly in layer 1 and 6 (31) and the fact that the LED illumination was likely strongest in superficial tissue due to scattering. While interpreting this LFP/CSD distribution is complex, it is clear that the effect of ArchT stimulation in V1 was present, dramatic, and sufficient to suppress PFC axon activity.

We then imaged somas of V1 neurons expressing GCaMP6s in the oddball and control paradigms as described above (“baseline”). Then, we repeated this visual stimulation sequence in a subsequent run while illuminating V1 with 617 nm LED every odd trial in mice expressing ArchT in PFC neurons (“optostimulation” run; eight experiments in eight mice; Fig. 5E). We compared the baseline runs to subsequent optostimulation runs, focusing analyses on oddly numbered trials in each. First, pooling all responsive neurons with stable average responses to at least one stimulus in both runs (Fig. 5F, G, J, and K), a 2×2 repeated measures ANOVA indicated that PFC suppression had a significant effect on DD [control versus deviant; $F^{\text{optostim} \times \text{context}}$ (1,252) = 5.98 , $P < 0.05$]. At baseline, neurons displayed clear DD [t (252) = 3.11 , $P < 0.01$], yet, during PFC suppression, DD was apparently absent [t (252) = 0.32 , $P = 0.75$]. PFC suppression did not affect SSA [control versus redundant; $F^{\text{optostim} \times \text{context}}$ (1,252) = 1.14 , $P = 0.28$]. Electrophysiological recording of multiunit activity (MUA) via multielectrode arrays in the same paradigm confirmed this effect (SI Appendix, Fig. S6). The LED-only condition did not exhibit any effects of optostimulation on DD [$F^{\text{optostim} \times \text{context}}$ (1,254) = 0.104 , $P = 0.78$] or SSA [$F^{\text{optostim} \times \text{context}}$ (1,252) = 2.30 , $P = 0.13$], confirming that these effects of PFC suppression on V1 are not due to LED-induced artifact or inadvertent visual stimulation (Fig. 5H and I).

Thus, although the activity of PFC inputs to V1 was equal across contexts (Fig. 4), disrupting this top-down influence reduced DD in V1, suggesting that they provide modulatory input which enhances contrast (35), potentially across contexts. To understand how this may be carried out, we analyzed the effect of PFC suppression on individual contexts (control redundant, deviant) separately for deviance detectors and nDDs. Interestingly, the effects of PFC suppression were only significant on neurons labeled as deviance detectors (Fig. 5L; i.e., deviant minus control $> Z = 1.67$ in the baseline or the optostimulation run). PFC suppression increased responses of deviance detectors, a given stimulus during the control [t (122) = -2.17 , $P < 0.05$] and redundant contexts [t (122) = -2.07 , $P < 0.05$], but decreased responses to that stimulus when it was contextually deviant [t (122) = 2.08 , $P < 0.05$]. These effects were not present for nDDs [Fig. 5M; t (129) = -0.67 , $P = 0.50$, t (129) = -0.44 , $P = 0.65$, t (129) = -0.54 , $P = 0.58$] or during the LED-only control run [DDs: t (140) = 0.42 , $P = 0.67$, t (140) = -0.23 , $P = 0.81$, t (140) = 1.35 , $P = 0.18$; nDDs: t (113) = 1.13 , $P = 0.26$, t (113) = -1.34 , $P = 0.18$, t (113) = 0.51 , $P = 0.61$].

Interestingly, while it is apparent that PFC modulates V1 deviance-detector responses, these results also suggest that the contextual preferences of deviance-detector neurons may exist




Fig. 5. PFC enhances DD in V1. (A) Schematic of viral injections to ACA of PFC and V1, and (B) GCaMP6s expression in PFC axons in V1 (50 μm below dura). (C) Average responses from axon segments expressing ArchT during LED stimulation. (D) Axons from C averaged together. (E) Schematic of experimental paradigm. LED illuminated cortex every other trial during the oddball and control paradigms. (F) Average responses across V1 neurons ($n = 252$) before (Left; baseline) and during PFC axon suppression. (G) Activity from stimulus interval from F. H and I are the same as G but from neurons from control mice ($n = 254$; i.e., no ArchT). (J) Raster plot of individual neuron responses (from F) to the same stimulus orientation averaged across control, redundant, and deviant conditions at baseline and (K) during PFC disruption. (L) Average stimulus-evoked responses of deviance-detector neurons during LED-on versus LED-off trials compared to (M) nondeviance-detector neurons. (N) Scatterplot depicting average responses from DDs (red) and nDD neurons (black) for each mouse ($n = 8$). (O) Proportion of visually nonresponsive neurons, deviance detectors, and nDDs. Paired-sample t tests; $*P < 0.05$, $**P < 0.01$, $***P < 0.001$.

independently of top-down modulation. Though attenuated, basic DD was still present in these neurons during PFC suppression [Fig. 5 L and N; $t(122) = 5.43$, $P < 0.05$], and PFC did not affect the relative proportions of DDs when baseline and opto-stimulation runs were analyzed separately (Fig. 5O).

Discussion

Functional Ensembles for Prediction Error in Cortical Circuits. Here, we demonstrate the existence of functionally correlated groups of neurons in layer 2/3 of V1 which are selectively responsive to deviant stimuli, or contextually unpredictable stimuli, in the visual oddball paradigm. These “deviance detector” ensembles are spatially intermixed with nondeviance detector cells in V1, responding to the same identical stimulus yet differently, in different contexts.

The laminar distribution of these deviance detectors is consistent with theoretical descriptions of a predictive coding (4, 9), which suggests that “prediction error” is signaled in feed-forward projections from supragranular layers (i.e., V1 deviance detectors in our data) while contextual “predictions” are encoded in top-down back projections from infragranular neurons (i.e., PFC axons in our data). Our results support this framework by demonstrating that “prediction error” is expressed in only a subset of supragranular neurons and that top-down inputs modulate the contextual precision of these errors. Although the functional consequence of this parcellation of responses among layer 2/3 neurons is unclear, it is possible that these “deviance detectors” project to different brain areas, more diffusely, or more distally than nDDs in layer 2/3, passing along the prediction error signal to

key targets. These results extend a recent finding identifying layer 2/3 neurons selective for mismatches of locomotion and visual input (i.e., sensory–motor mismatch instead of sensory–sensory mismatch) (10) propose that DD neurons may exist as a general feature of V1, which responds to mismatches of both sensory and motor predictions. Additionally, our results demonstrate key properties of how these deviance detecting cells (putative “prediction error” neurons) relate to local V1 and top-down networks.

Context-Specific Ensembles Form Nonoverlapping Subnetworks. Interestingly, context modulation appears to be separate from stimulus feature selectivity, or orientation selectivity in this case. In other words, according to our results, each group of neurons which selectively respond to a given orientation would be subdivided into different ensembles, each responding differently to deviant stimuli. V1 neurons simultaneously exhibit selectivity along multiple physical attributes (36), and here, we show that contextual preferences represent another such dimension. On the other hand, while not all orientation selective neurons show DD, most DD neurons exhibit DD in an orientation-selective manner. While this suggests an intriguing hierarchical relationship between the two types of selectivity (context and orientation), more research is needed which thoroughly examines DD responses across a range of physical visual features.

Our finding of separate, spatially intermixed ensembles echoes previous demonstrations of nonoverlapping feedforward (layer 4 to 2/3) and interneuronal subnetworks of neurons in mouse neocortex (24, 37). Interestingly, while PFC disruption did modulate the activity of deviance detectors, the statistical preferences and approximate proportion (about 1/3) of these neurons did not change. As ACA is just one source of top-down input to V1, these results cannot rule out the possibility that DD neuron preferences are inherited in a fully top-down manner. On the other hand, the DD ensembles that we identify could be hardwired in local inhibitory and excitatory synaptic innervation patterns (8, 38) or could comprise a molecularly distinct cellular subtype (39). Future work could address this, targeting opto- or chemico-genetic tools to these subpopulations of cells, currently defined purely on their functional properties during this paradigm and/or cellular reconstructions or RNA sequencing from cells recorded *in vivo*.

Prefrontal Inputs Enhance Contextual “Contrast”. We also explored the areal contributions to DD in sensory cortex. ACA inputs to layer 1 of V1 showed strong visually driven activity but no contextual modulation in their stimulus-driven responses. That is, while the response levels of V1 layers 2 to 5 depended on the context in which a given visual stimulus was presented, PFC inputs mostly responded equally, on average, across contexts (Fig. 4 *I* and *J*). Interestingly, optogenetic suppression of PFC inputs to V1 showed that context processing in V1 is nevertheless modulated by PFC, as DD was strongly affected by PFC inhibition. Interestingly, this effect was selective to deviance detectors, suggesting either targeting specificity of PFC-V1 projections or, alternatively, since most top-down innervation tends to be modulatory rather than driving (35, 40), differential sensitivity to top-down modulation among cell groups in V1. The responsible circuit mechanisms for such an ensemble-specific bidirectional effect are likely numerous and will require systematic investigations aided by better methods for functional precision tagging of cell types (see above).

More generally, this role of PFC in bidirectionally modulating V1 activity in accordance with context is consistent with the notion that these backward projections send contextual predictions about the world (or the “causes” of sensory information) to V1 (4) and manifest as modulations of synaptic gain (41, 42). This prediction should 1) be more or less constant across trials once the contextual probabilities are clear (Fig. 4), 2) serve to suppress predicted responses (i.e., affecting reductions to control/redundant in Fig. 5 *L*) by targeting mainly local interneurons (31), and 3) consequently

enhance deviations from these predictions via a relaxation of disinhibitory or lateral inhibition among V1 neuron with separate stimulus preferences (i.e., affecting augmentation to deviant in Fig. 5 *L*). Thus, PFC may serve to enhance the “contextual contrast” or precision of the outputs of deviance detectors specifically, suppressing their activity during predictable stimulation and facilitating their activity to a nonpredicted stimulus (Fig. 5 *L* and *N*).

Understanding of such a role of PFC in contextual contrast modulation is particularly important given the substantial body of work focused on the role of prefrontal cortex in minimizing sensory distraction (43), especially during tasks requiring working memory (44). Indeed, hypofrontality seen in schizophrenia (45) or frontal lesions give rise to distractibility (46). The fact that PFC suppression in our experiments augmented the V1 outputs to nondeviant stimuli suggests a mechanistic explanation for these clinical phenomena. Given what has been observed in schizophrenia patients with EEG recordings (16, 47), a dysfunction of PFC inputs to V1, disorganization of local V1 ensembles (15), or altered inhibitory microcircuits (12) could each provide (nonmutually exclusive) explanations for mismatch deficits in the disorder.

More deeply exploring the role, if any, of non-V1 regions in giving rise to this DD ensemble remains a future goal. Further, as some proportion of PFC-V1 inputs terminate in layer 6 (31) and may send collaterals to other visual areas downstream from V1, we do not know the degree to which our manipulations affected these inputs (versus layer 1 inputs). It remains possible that this layer 6 portion of the circuitry plays a different, still uncharacterized role in V1 DD and ensemble membership. Also, it was recently reported that DD appears to increase in magnitude from sub- to primary- to frontal-cortical regions, particularly in the auditory domain (48). One intriguing possibility, which would conform well with predictive coding theories (9), is that the difference from our report lies in the type of neurons recorded; our study selectively recorded PFC neurons which send backward projections to V1 (which convey predictions, not prediction errors), while Casado-Román et al. recorded indiscriminately among PFC neurons (48).

Clinical Implications. These results highlight that deviance-selective, or “prediction error,” neurons are fundamental to visual cortical processing (24), carrying implications not just for sensory neuroscience but also for clinical neurophysiology. Indeed, the oddball stimulus presentation employed in this study is not only a simple and effective paradigm for understanding basic sensory processing, but it is also a well-established approach for studying context processing deficits in neuropsychiatric disorders such as schizophrenia and dementia (49–51). In particular, the DD measure is a known component of the classic “mismatch negativity” (5), which is not only one of the most reliable biomarkers of schizophrenia (52) but is also strongly predictive of generalized functioning in the disease (6).

In light of this, the results of the current paper provide an angle in understanding the basic neurobiology of this measure and its relationship with the pathophysiology of schizophrenia and other psychoses. Specifically, we show how context processing in oddball paradigms is not well captured by single-cell or gross-level recordings. Instead, contextual modulation in sensory cortex is manifest in the diverse activations of subnetworks of neuronal ensembles. Indeed, schizophrenia, while originating from a vast array of genetic and environmental causes, may be fundamentally a disorder of disorganized cortical ensembles (13, 15, 53).

Another key implication for psychiatry lies in our findings regarding the interaction of top-down inputs with superficial layers of sensory cortex. In a predictive coding framework, a distinction is generally made between descending or top-down predictions of content (which stimuli may occur) and context (the likelihood of a given stimulus/set of stimuli) (42). False perceptual inference in

psychotic disorders may result from a specific failure of contextual predictions. This is often described in terms of an imbalance between the precision of sensory data versus contextual predictions (from top-down sources), amplifying or attenuating the precision of sensory prediction errors (i.e., the mismatch between them) (54, 55). Psychologically speaking, this would represent a failure to attend away from—or ignore—certain stimuli or their attributes (41). Physiologically, the encoding of precision is thought to be mediated by synaptic gain control and balanced inhibitory networks (56) that implicate superficial pyramidal cells thought to encode prediction errors (e.g., deviance detectors).

Importantly, this formulation is consistent with the nonlinear pattern of effects seen after suppression of top-down inputs in our study; that is, PFC inputs to V1 deviance detectors appear to be modulatory in nature, corresponding to the precision and the gain control described above (42). Thus, the current results may provide direct empirical evidence—not only for the encoding of prediction errors in supragranular layers of visual cortex (9) but for the selective modulation of the gain afforded to V1 pyramidal cells by PFC inputs, which controls the precision of these prediction errors, altogether underwriting predictive coding accounts of psychosis, hallucinosis, and false perceptual inference (57).

In summary, our results reveal a functional parcellation of cortical circuits into specific ensembles of cortical neurons engaged in different contextual computations with the same inputs. These distinct context preferring ensembles, and DD cells in particular, provide a circuit mechanism underlying how the brain generates prediction errors and responds during the “oddball” paradigm, a paradigm of great importance for basic psychology and understanding the pathophysiology of major psychiatric disorders (5).

Materials and Methods

Animals, Surgery, and Training. All experimental procedures were approved by and carried out in accordance with Columbia University institutional animal care guidelines. Experiments were performed on adult C57BL/6 mice ($n = 44$; Jackson Laboratory, 22 to 32 g) at the age of postnatal day (P) P60 to P90. Virus injection, head plate fixation, and craniotomy were carried out in that order over the course of 4 wk. For virus injection, mice were anesthetized with isoflurane (initially 3% [partial pressure in air] and reduced to 1 to 2%). A small hole was drilled in the skull at (x, y) coordinates above left V1 (coordinates from lambda: $X = -2,500, Y = 200, Z = -250 \mu\text{m}; n = 21$) and left ACa (from bregma: $-300, 600, -900 \mu\text{m}; n = 19$) using a dental drill, taking care not to pierce the dura mater. A glass capillary pulled to a sharp micropipette (10 to 40 μm diameter) was advanced to the z coordinate (from pial surface) with the stereotaxic instrument, and 750 nl solution of 1:1 diluted AAV1-Syn-GCaMP6s (obtained from the University of Pennsylvania Vector Core; $n = 20$) or AAV1-CAG-ArchT-tdTomato (University of North Carolina [UNC] vector core; $n = 7$ mice) was injected over a 10 min period using an UltraMicroPump3 (UMP3) microsyringe pump (World Precision Instruments). Additionally, for the volumetric imaging experiments, in four mice, AAV1-Syn-jRGECO1b was injected at 500 μm deep in V1, together with AAV9-Syn-GCaMP6s at 250 μm . Immediately after injection, a titanium head plate was attached to the skull centered on V1 using dental cement. Mice were allowed to recover for at least 3 d in their home cage and given analgesics (5 mg/kg carprofen intraperitoneally; I.P.). Approximately 2 to 3 wk after virus injection, a glass cranial window was made in the skull over V1 or A1. Follow-up experiments (SI Appendix, Fig. S3) were done in TIGRE2.0 transgenic mice expressing GCaMP6f under Vglut1 promoter (Ai148D \times Slc17a7-IRES2-Cre) (58), yielding cortex-wide expression in excitatory cells. This transgenic approach generated homogeneous expression of GCaMP6f, which allowed for imaging larger populations of neurons in a given mouse, as well as deeper regions (e.g., layer 5). For the craniotomy, a circular opening was drilled ($3 \times 3 \text{ mm}$), the exposed brain washed with phosphate buffered saline (PBS) over a few minutes to limit bleeding from dura, submerged with 37 °C low-melting point agarose, covered with circular glass coverslip, and fixed with dental cement. Mice were accustomed to head fixation as previously described (1). During training sessions and prior to the first imaging session, mice viewed moving square-wave gratings for stimulus habituation.

Visual Stimulation. Visual stimuli were generated using the MATLAB (MathWorks) Psychophysics Toolbox and displayed on a liquid crystal display

monitor (19-inch diameter, 60 Hz refresh rate) positioned 15 cm from the right eye, roughly at 45° to the long axis of the animal (Fig. 1A). Stimuli were static or moving (two cycles per second) full-field square-wave gratings (100% contrast, 0.04 cycles per degree) oriented in one of two separate orientations for the oddball paradigm (45° and 135°; 0° and 90°; counter-balanced across mice) or in either orientation for the many-standards control (30, 45, 60, 90, 120, 135, 150, and 180°). Stimuli were presented for 500 ms followed by an interstimulus interval of 500 to 550 ms of black screen. In the oddball sessions, the “standard” stimulus was presented at a minimum of three sequential trials, followed by a linearly increasing probability of the “target” stimulus on each successive trial to yield an overall 12.5% probability of targets. These sessions lasted 10 min and were repeated with the “standard” and “target” stimuli reversed. In the many-standards sessions, stimuli of eight separate orientations each occurred at random with a 12.5% probability in a session of 10 min.

Two-Photon Calcium Imaging: V1 Single Plane ($n = 15$ Mice). The activity of cortical neurons was recorded by imaging fluorescence changes under a two-photon microscope (Bruker Ultima In Vivo; Billerica, MA) excited with a Ti:Sapphire laser (MaiTai DeepSee) tuned to 940 nm. The laser beam was intensity modulated with a Pockels cell (Conoptics 350-105, with 302 RM driver) and scanned with galvometers through a 20×0.95 numerical aperture (N.A.; Olympus) water immersion objective (Fig. 3A and B). To ensure stability of the imaging meniscus for long-duration imaging sessions, a small volume (approximately 1 mL) of Aquasonic ultrasound gel (Parker Laboratories Inc.) was centrifuged and dolloped onto the glass coverslip over the window. Scanning and image acquisition were controlled by Prairie View software (56.9 frames per second, $2 \times$ averaging, for 256×256 pixels, 268 ns dwell time per pixel, 170 to 250 μm beneath the pial surface). Mice and recordings were visually monitored by the experimenter to ensure they were awake during data collection. Locomotion was recorded with an infrared LED/photodarlington pair (Honeywell S&C HOA1877-003). When detected, frames or trials during locomotion periods were excluded along with the previous and subsequent 12 frames; this did not change the pattern of effects (15).

Two-Photon Calcium Imaging: Multiplane/Volumetric ($n = 4$ Mice). Two-photon volumetric-calcium-imaging experiments were carried out with a customized microscope as previously reported in detail (28). Two excitation lasers were used: Ti:Sapphire laser (Chameleon Ultra II, Coherent) at 920 nm, and an amplified fiber laser (Fianium) at 1,064 nm. An electrical tunable lens (EL-10-30-C-NIR-LD-MV, Optotune) was inserted in the 920 nm laser beam path and a spatial light modulator (SLM; HSP512, Meadowlark Optics) in the 1064 nm laser beam path, so that the focal plane of both lasers could be dynamically adjusted. The two laser beams were combined through a dichroic mirror, scanned by resonant galvanometric scanning mirrors at two different focal planes simultaneously, and rapidly shifted across five sequential depths to yield 10 volumes per second, each volume containing 10 planes with 256×256 pixels per plane. Emission fluorescence was collected through two separate photomultiplier tubes (PMTs), with a collection filter of $520 \pm 40 \text{ nm}$ for the green path and a $630 \pm 75 \text{ nm}$ for the red path.

Image Analysis. Imaging datasets were scored similarly to previous reports (1, 15, 21, 59). The raw images were processed to correct translational brain motion artifacts using an in-house plugin named “Moco” for ImageJ (60, 61). Then, cell ROIs were detected semiautomatically, with a principal components analysis (PCA) assisted selection of pixels for each cell, for each imaging session, and individually confirmed as described in our past work (1, 13, 15). Fluorescence of active cells was then calculated as the average across all pixels within this ROI minus the average of the pixels just outside the selected rectangle, termed the “halo,” which excluded pixels from nearby cell bodies. This subtraction removed background contamination from neuropil and nearby cells (59). For a subset of the data (the images acquired with the volumetric multiplane approach, $n = 4$ mice, eight experiments or the GCaMP6f experiments from SI Appendix, Fig. S3, $n = 3$ mice, 10 experiments), a slightly different approach was employed to identify the cell body ROI given the vast numbers of cells present. Cell centroids were manually initialized, then cell body masks, fluorescence “signals,” and background were automatically modeled using the constrained nonnegative matrix factorization algorithm (62), automated for MATLAB with scripts written in-house.

After these steps, each cell outline and trace was manually inspected, and cells with no apparent calcium transients were excluded from further analysis. For the axon imaging experiments, correlations of all remaining traces were computed, and ROI pairs with correlations above $r = 0.5$ were visually

inspected in order to rule out multiple ROIs being included in the analysis from the same cell. Traces of ROIs suspected to be from a previously scored ROI were removed. This led to the elimination of only about 5% of ROIs. The remaining traces were then filtered with a 1 s low-pass envelope (63). Finally, the discrete positive first derivative was scored as Δ -f (within-cell/single-cell comparisons). After averaging over trials, the 250 ms prestimulus average for all three stimulus types (control, redundant, and deviant) were combined, and the mean and SD was calculated on the Δ -f values in this 750 ms window for each cell. These values were used to compute a z-scored Δ -f for visualizing and combining activity across cells (1).

Multielectrode Recordings. Extracellular electrophysiological data are reported on 14 mice (8 female, 22 to 28g) undergoing the virus injection protocol described above. Recordings with 16-channel linear silicon probes (spaced at 50 μ m intervals; model a1x16-3mm50-177, Neuronexus Technologies) inserted perpendicular to the pial surface were carried out as previously described (1). Continuous data were acquired with a Plexon MiniDigi amplifier and software (Plexon Inc.). LFP signals were filtered from 0.1 to 300 Hz, sampled at 1 kHz, and analyzed as LFP and CSD (see below). MUA was sampled at 40 kHz, digitally filtered (300 to 5,000 Hz; bandpass least squares finite impulse response filter; FIR), rectified, and then downsampled to 1 kHz. The result was low-pass filtered at 100 Hz (least squares FIR) to estimate the local population spiking envelope.

LFP and CSD Processing and Analysis. LFP data were manually prescreened for excessive artifacts as previously described (1). For analysis of the LFP response, the channel of maximal negative deflection in the first 150 ms poststimulus (typically channel 12 to 16, or layer 4 to 5) was selected. Initial demonstration of SSA and DD were established via a “planned comparisons” approach with paired *t* tests (one-tailed) on mouse-wise averages. For SSA, we compared responses between redundant (using the fourth redundant in order to normalize for trial counts and for relative time during the run) and control stimuli in the early time range (40 to 110 ms) most likely to capture SSA (2). For deviance processing, we compared responses between deviant and control stimuli in the late time range (140 to 300 ms) (1). Trial numbers were equated between stimulus conditions. All SE bars in all figures reflect within-subjects/within-cells SE.

For descriptive purposes, average CSD was computed from the average LFP (SI Appendix, Fig. S1A) by taking the discrete second derivative across the electrode sites and interpolated to produce a smooth CSD map (1, 22). Putative laminar subregions (three channels) were defined based on CSD demarcations previously published and verified (1, 22).

Calcium Imaging Analysis: Single Cell Effects. Condition averages of normalized Δ -f values for redundant (first in sequence), deviant, and control stimuli were calculated separately for each stimulus type for each neuron. All analyses focused on the first eight trials to equate across conditions, cells, and mice with varying numbers of available trials (i.e., all mice and conditions had at least eight usable trials for each stimulus and condition after scoring). Initial analyses focused on neurons showing, during control or deviant conditions, an average poststimulus (0 to 0.5 s) response of 1.67 SD above prestimulus baseline (equating to a one-tailed *P* value of 0.05). Only responses to one stimulus orientation were considered for each cell (i.e., the orientation with greater magnitude). Initial demonstration of SSA and DD were established via a “planned comparisons” approach with paired *t* tests (two-tailed) on cell-wise averages of poststimulus activity (0 to 0.5 s). For statistical comparisons, we used paired-samples two-tailed *t* test on the same neurons across contexts unless otherwise specified. For SSA, we compared responses between redundant and control stimuli; for DD, we compared responses between deviant and control stimuli.

Calcium Imaging Analysis: Cluster Analyses. Average responses to control, redundant, and deviant stimuli averaged over trials and the stimulation

period (0 to 500 ms poststimulus). For creating scatterplots across conditions, we calculated two variables for each neuron and condition which were statistically orthogonal, differing slightly from the calculations described above: putative DD (response to deviant condition minus response to control condition) and SSA (response to [deviant plus control] conditions minus response to redundant condition). These variables represent orthogonal contrasts used only for plots and identifying subpopulations and were not used for statistical comparisons.

Plotting all neurons on these two variables revealed ostensibly three clusters of neuronal responses. Group demarcations were identified with *k*-means clustering analyses where *k* = 2 to 6 (SI Appendix, Fig. S2). After assigning each neuron to a given cluster, we computed their responses to stimulus contexts in a subsequent run (30 min later) as a cross validation step.

Optogenetic Suppression of PFC Axons. A 617 nm high-power LED coupled to the \varnothing 400 μ m Core Patch Cable (Thorlabs) was positioned around the cranial window as described previously (64). Light intensity at the tip was \sim 4 mW. The tip was coupled to the fiber optic cannula (Thorlabs) and positioned under a 20 \times water immersion objective (0.95 N.A.; Olympus), to indicate the center of the imaging window. The distance between the cranial window and the tip was adjusted so that the radius of the illumination field was \sim 1 mm. The imaging setup and the space between cranial window and the objective were completely enclosed with blackout fabric (Thorlabs). First, baseline runs (many-standards control, oddball, and oddball flipped) for each experiment were completed without LED illumination. Then, after 30 min, these runs were repeated with the same visual stimuli along with LED illumination occurring every other trial, starting 500 ms pre- to 500 ms postvisual stimulus onset.

Activity of visual cortical neurons was imaged and processed as described above. Only cells with stimulus driven responses greater than $z = 2.6$ ($P < 0.01$) to a single stimulus and context in both baseline and optostim were included. This is a higher cutoff than the other experiments in this study ($z > 1.67$; $P < 0.05$) and was chosen because PFC suppression led to more spontaneous activity, and thus, erroneously labeled “active” neurons. Importantly, this difference in cutoff essentially led to the same proportion of neurons being estimated as “responsive” as the other experiments (i.e., compare Fig. 5O with Fig. 2D).

Mice with PFC ArchT injections were compared to mice receiving only PFC GCaMP6s injections. Cell-level effects on SSA and DD were analyzed for PFC-suppression mice with a 2×2 repeated measures ANOVA on individual cells with CONTEXT (CONTROL; REDUNDANT [for SSA]) or (CONTROL; DEVIANT [for DD]) and OPTOSTIMULATION (without LED; with LED) as variable. Paired-samples *t* tests were used to describe interaction effects (i.e., the presence of SSA and DD before and during stimulation). For comparisons within neuron type (DD, nDD), paired-samples *t* tests were used to determine whether responses to specific stimuli (control, deviant, and redundant) were changed by the LED. The same statistical model was used for the control condition to rule out effects of LED illumination on its own.

Data Availability. All study data and code included in the paper are available in Figshare [https://figshare.com/articles/dataset/Figure_1-2_V1_neuron_calcium_data/14225735 (65) and https://figshare.com/articles/dataset/Figure_3_V1_neuron_calcium_data_multiplane/14225759 (66)].

ACKNOWLEDGMENTS. We thank members of the Yuste laboratory for comments. We thank the following for funding: National Institute of Mental Health Grants (F32-MH106265, R00MH115082, R01MH115900), National Eye Institute Grant (R01EY011787), National Institute of General Medical Sciences Grant (T32GM008798-17), Howard Hughes Medical Institute student research fellowship, Brains and Behavior Research Foundation Grant (19944), the Whitehall Foundation Grant (2019-05-44), and Burroughs Wellcome Fund Grant (CASI 1015761).

- J. P. Hamm, R. Yuste, Somatostatin interneurons control a key component of mismatch negativity in mouse visual cortex. *Cell Rep.* **16**, 597–604 (2016).
- I.-W. Chen, F. Helmchen, H. Lütke, Specific early and late oddball-evoked responses in excitatory and inhibitory neurons of mouse auditory cortex. *J. Neurosci.* **35**, 12560–12573 (2015).
- R. G. Natan *et al.*, Complementary control of sensory adaptation by two types of cortical interneurons. *eLife* **4**, e09868 (2015).
- K. Friston, A theory of cortical responses. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **360**, 815–836 (2005).
- R. Näätänen, The role of attention in auditory information processing as revealed by event-related potentials and other brain measures of cognitive function. *Behav. Brain Sci.* **13**, 201–233 (1990).
- G. A. Light, D. L. Braff, Mismatch negativity deficits are associated with poor functioning in schizophrenia patients. *Arch. Gen. Psychiatry* **62**, 127–136 (2005).
- D. C. Javitt, M. Steinschneider, C. E. Schroeder, J. C. Arezzo, Role of cortical N-methyl-D-aspartate receptors in auditory sensory memory and mismatch negativity generation: Implications for schizophrenia. *Proc. Natl. Acad. Sci. U.S.A.* **93**, 11962–11967 (1996).
- J. M. Ross, J. P. Hamm, Cortical microcircuit mechanisms of mismatch negativity and its underlying subcomponents. *Front. Neural Circuits* **14**, 13 (2020).
- A. M. Bastos *et al.*, Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
- R. Jordan, G. B. Keller, Opposing influence of top-down and bottom-up input on excitatory layer 2/3 neurons in mouse primary visual cortex. *Neuron* **108**, 1194–1206.e5 (2020).

11. J. Kremláček *et al.*, Visual mismatch negativity (vMMN): A review and meta-analysis of studies in psychiatric and neurological disorders. *Cortex* **80**, 76–112 (2016).
12. A. B. Van Derveer *et al.*, A role for somatostatin-positive interneurons in neuro-oscillatory and information processing deficits in schizophrenia. *Schizophr. Bull.*, sbaa184 (2020).
13. J. P. Hamm, Y. Shymkiv, J. Mukai, J. A. Gogos, R. Yuste, Aberrant cortical ensembles and schizophrenia-like sensory phenotypes in *Setd1a*^{fl/-} mice. *Biol. Psychiatry* **88**, 215–223 (2020).
14. J. D. Murray *et al.*, Linking microcircuit dysfunction to cognitive impairment: Effects of disinhibition associated with schizophrenia in a cortical working memory model. *Cereb. Cortex* **24**, 859–872 (2014).
15. J. P. Hamm, D. S. Peterka, J. A. Gogos, R. Yuste, Altered cortical ensembles in mouse models of schizophrenia. *Neuron* **94**, 153–167.e8 (2017).
16. S. Ranlund *et al.*, Impaired prefrontal synaptic gain in people with psychosis and their relatives during the mismatch negativity. *Hum. Brain Mapp.* **37**, 351–365 (2016).
17. A. J. Gaebler *et al.*, Auditory mismatch impairments are characterized by core neural dysfunctions in schizophrenia. *Brain* **138**, 1410–1423 (2015).
18. G. G. Parras *et al.*, Neurons along the auditory pathway exhibit a hierarchical organization of prediction error. *Nat. Commun.* **8**, 2148 (2017).
19. S. M. Haigh *et al.*, Reduced late mismatch negativity and auditory sustained potential to rule-based patterns in schizophrenia. *Eur. J. Neurosci.* **49**, 275–289 (2019).
20. C. C. H. Petersen, S. Crochet, Synaptic computation and sensory processing in neocortical layer 2/3. *Neuron* **78**, 28–48 (2013).
21. J. E. Miller, I. Ayzenshtat, L. Carrillo-Reid, R. Yuste, Visual stimuli recruit intrinsically generated cortical ensembles. *Proc. Natl. Acad. Sci. U.S.A.* **111**, E4053–E4061 (2014).
22. C. M. Niell, M. P. Stryker, Highly selective receptive fields in mouse visual cortex. *J. Neurosci.* **28**, 7520–7536 (2008).
23. R. Yuste, From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* **16**, 487–497 (2015).
24. Y. Yoshimura, J. L. M. Dantzker, E. M. Callaway, Excitatory cortical neurons form fine-scale functional networks. *Nature* **433**, 868–873 (2005).
25. C. M. Niell, M. P. Stryker, Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**, 472–479 (2010).
26. L. Pinto *et al.*, Fast modulation of visual perception by basal forebrain cholinergic neurons. *Nat. Neurosci.* **16**, 1857–1863 (2013).
27. R. J. Douglas, K. A. C. Martin, Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.* **27**, 419–451 (2004).
28. S. Han, W. Yang, R. Yuste, Two-color volumetric imaging of neuronal activity of cortical columns. *Cell Rep.* **27**, 2229–2240.e4 (2019).
29. W. Yang *et al.*, Simultaneous multi-plane imaging of neural circuits. *Neuron* **89**, 269–284 (2016).
30. S. Zhang *et al.*, Organization of long-range inputs and outputs of frontal cortex for top-down control. *Nat. Neurosci.* **19**, 1733–1742 (2016).
31. S. Zhang *et al.*, Selective attention. Long-range and local circuits for top-down modulation of visual cortex processing. *Science* **345**, 660–665 (2014).
32. L. D. Selemon, J. E. Kleinman, M. M. Herman, P. S. Goldman-Rakic, Smaller frontal gray matter volume in postmortem schizophrenic brains. *Am. J. Psychiatry* **159**, 1983–1991 (2002).
33. M. M. Roth *et al.*, Thalamic nuclei convey diverse contextual information to layer 1 of visual cortex. *Nat. Neurosci.* **19**, 299–307 (2016).
34. L. L. Glickfeld, M. L. Andermann, V. Bonin, R. C. Reid, Cortico-cortical projections in mouse visual cortex are functionally target specific. *Nat. Neurosci.* **16**, 219–226 (2013).
35. J. M. Hupé *et al.*, Cortical feedback improves discrimination between figure and background by V1, V2 and V3 neurons. *Nature* **394**, 784–787 (1998).
36. J. H. Marshel, M. E. Garrett, I. Nauhaus, E. M. Callaway, Functional specialization of seven mouse visual cortical areas. *Neuron* **72**, 1040–1054 (2011).
37. M. M. Karnani *et al.*, Cooperative subnetworks of molecularly similar interneurons in mouse neocortex. *Neuron* **90**, 86–100 (2016).
38. G. Silberberg, H. Markram, Disynaptic inhibition between neocortical pyramidal cells mediated by Martinotti cells. *Neuron* **53**, 735–746 (2007).
39. R. Yuste *et al.*, A community-based transcriptomics classification and nomenclature of neocortical cell types. *Nat. Neurosci.* **23**, 1456–1468 (2020).
40. S. M. Sherman, Thalamus plays a central role in ongoing cortical functioning. *Nat. Neurosci.* **19**, 533–541 (2016).
41. R. A. Adams, K. E. Stephan, H. R. Brown, C. D. Frith, K. J. Friston, The computational anatomy of psychosis. *Front. Psychiatry* **4**, 47 (2013).
42. K. Friston, Does predictive coding have a future? *Nat. Neurosci.* **21**, 1019–1021 (2018).
43. J. E. McDowell, K. A. Dyckman, B. P. Austin, B. A. Clementz, Neurophysiology and neuroanatomy of reflexive and volitional saccades: Evidence from studies of humans. *Brain Cogn.* **68**, 255–270 (2008).
44. Y. Ku, M. Bodner, Y.-D. Zhou, Prefrontal cortex and sensory cortices during working memory: Quantity and quality. *Neurosci. Bull.* **31**, 175–182 (2015).
45. D. R. Weinberger, M. S. Aloia, T. E. Goldberg, K. F. Berman, The frontal lobes and schizophrenia. [published erratum appears in *J. Neuropsychiatry Clin. Neurosci.* 1995 Winter;7(1):121]. *J. Neuropsychiatry Clin. Neurosci.* **6**, 419–427 (1994).
46. E. K. Miller, J. D. Cohen, An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* **24**, 167–202 (2001).
47. D. C. Javitt, Intracortical mechanisms of mismatch negativity dysfunction in schizophrenia. *Audiol. Neurotol.* **5**, 207–215 (2000).
48. L. Casado-Román, G. V. Carbajal, D. Pérez-González, M. S. Malmierca, Prediction error signaling explains neuronal mismatch responses in the medial prefrontal cortex. *PLoS Biol.* **18**, e3001019 (2020).
49. J. P. Hamm *et al.*, Family history of psychosis moderates early auditory cortical response abnormalities in non-psychotic bipolar disorder. *Bipolar Disord.* **15**, 774–786 (2013).
50. K. Farkas, G. Stefanics, C. Marosi, G. Csukly, Elementary sensory deficits in schizophrenia indexed by impaired visual mismatch negativity. *Schizophr. Res.* **166**, 164–170 (2015).
51. L. E. Ethridge *et al.*, Event-related potential and time-frequency endophenotypes for schizophrenia and psychotic bipolar disorder. *Biol. Psychiatry* **77**, 127–136 (2015).
52. M. A. Erickson, A. Ruffe, J. M. Gold, A meta-analysis of mismatch negativity in schizophrenia: From clinical risk to disease specificity and progression. *Biol. Psychiatry* **79**, 980–987 (2015).
53. E. T. Rolls, M. Loh, G. Deco, G. Winterer, Computational models of schizophrenia and dopamine modulation in the prefrontal cortex. *Nat. Rev. Neurosci.* **9**, 696–709 (2008).
54. E. Pellicano, D. Burr, When the world becomes ‘too real’: A Bayesian explanation of autistic perception. *Trends Cogn. Sci.* **16**, 504–510 (2012).
55. A. R. Powers, C. Mathys, P. R. Corlett, Pavlovian conditioning-induced hallucinations result from overweighting of perceptual priors. *Science* **357**, 596–600 (2017).
56. X.-J. Wang, J. H. Krystal, Computational psychiatry. *Neuron* **84**, 638–654 (2014).
57. P. Sterzer *et al.*, The predictive coding account of psychosis. *Biol. Psychiatry* **84**, 634–643 (2018).
58. T. L. Daigle *et al.*, A suite of transgenic driver and reporter mouse lines with enhanced brain-cell-type targeting and functionality. *Cell* **174**, 465–480.e22 (2018).
59. T.-W. Chen *et al.*, Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
60. C. A. Schneider, W. S. Rasband, K. W. Eliceiri, NIH image to ImageJ: 25 years of image analysis. *Nat. Methods* **9**, 671–675 (2012).
61. A. Dubbs, J. Guevara, R. Yuste, moco: Fast motion correction for calcium imaging. *Front. Neuroinform.* **10**, 6 (2016).
62. E. A. Pnevmatikakis *et al.*, Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron* **89**, 285–299 (2016).
63. Y. Masamizu *et al.*, Two distinct layer-specific dynamics of cortical ensembles during learning of a motor task. *Nat. Neurosci.* **17**, 987–994 (2014).
64. M. Agetsuma, J. P. Hamm, K. Tao, S. Fujisawa, R. Yuste, Parvalbumin-positive interneurons regulate neuronal ensembles in visual cortex. *Cereb. Cortex* **6**, 1–15 (2017).
65. J. Hamm, Figure 1-2, V1 neuron calcium data. *Figshare*. Dataset. <https://doi.org/10.6084/m9.figshare.14225735.v1>. Deposited 16 March 2021.
66. J. Hamm, Figure 3, V1 neuron calcium data multiplane. *Figshare*. Dataset. <https://doi.org/10.6084/m9.figshare.14225759.v1>. Deposited 16 March 2021.