Cracking the Function of Layers in the Sensory Cortex

Hillel Adesnik[1,2,*] and Alexander Naka[1,2]
[1]Department of Molecular and Cell Biology, University of California, Berkeley, Berkeley, CA, USA
[2]The Helen Wills Neuroscience Institute, University of California, Berkeley, Berkeley, CA, USA
*Correspondence: hadesnik@berkeley.edu
https://doi.org/10.1016/j.neuron.2018.10.032

Understanding how cortical activity generates sensory perceptions requires a detailed dissection of the function of cortical layers. Despite our relatively extensive knowledge of their anatomy and wiring, we have a limited grasp of what each layer contributes to cortical computation. We need to develop a theory of cortical function that is rooted solidly in each layer's component cell types and fine circuit architecture and produces predictions that can be validated by specific perturbations. Here we briefly review the progress toward such a theory and suggest an experimental road map toward this goal. We discuss new methods for the all-optical interrogation of cortical layers, for correlating *in vivo* function with precise identification of transcriptional cell type, and for mapping local and long-range activity *in vivo* with synaptic resolution. The new technologies that can crack the function of cortical layers are finally on the immediate horizon.

"At present we have no direct evidence on how the cortex transforms the incoming visual information. Ideally, one should determine the properties of a cortical cell, and then examine one by one the receptive fields of all the afferents projecting upon that cell."– Hubel and Wiesel, 1962, *Journal of Physiology*

A primary goal of cortical physiology is to explain how the cortex transforms incoming information to generate perceptions. More than half a century has passed since the above statement was made, but a detailed understanding of the mechanisms that mediate cortical transformations across the cortical layers remains remarkably incomplete. However, recent technological advances finally allow execution of the experiment that Hubel and Wiesel prescribed, as well as many other sophisticated assays that can overcome this conceptual challenge. First, we briefly review how existing data have motivated the available theories regarding the function of cortical layers, primarily with respect to sensory transformations. Next, we highlight the key data we lack that could confirm or invalidate these models or motivate new ones. Finally, we propose the new technologies and experiments that are needed to obtain the data that will allow us to arrive at a much more mechanistic, circuit-driven theory for the unique contributions of layer-specific circuits in sensory perception.

The cortical generation of sensory percepts can be thought of as a synthetic, hierarchical process or as one based largely in statistical inference. In the hierarchical model, neurons integrate their inputs to filter the sensory data and transform it into an output spike train that encodes features of the stimulus. A simple feedforward architecture composed of many neurons filtering their input in this manner should ultimately enable complex computations to mediate object identification and scene analysis (Hubel and Wiesel, 1962; Marr, 2010). The apparent feedforward architecture of the primate visual system might help explain why object recognition is fast (Thorpe et al., 1996). In the framework of statistical inference, cortical circuits encode a generative model of the sensory environment, and recurrent interactions between cortical processing stages compare the expectations of the internally generated model with incoming data from the sensory apparatus (Bastos et al., 2012).

Two of the most compelling examples of the synthetic process are the encoding of edge orientation in the primary visual cortex (Hubel and Wiesel, 1959, 1962) and that of object or face selectivity in the inferotemporal cortex (Bruce et al., 1981; Gross et al., 1972). The emergence of orientation tuning stands as one of the few concrete examples of a de novo transformation that occurs in a layer of the primary visual cortex (V1) and can be well explained by a simple feedforward model involving integration over a specific set of center-surround thalamic relay neurons (Hubel and Wiesel, 1962). Although the mechanistic details of an analogous feedforward circuit for the generation of face selectivity are lacking, one can conceptualize a similar process where neurons exhibiting increasingly sophisticated feature tuning are built by summing over neurons with more elementary filtering properties (e.g., edges to contours and contours to faces) (Chang and Tsao, 2017; Liu et al., 2016).

Based on this framework, one might expect that further de novo transformations would occur as sensory signals propagate through the layers of the cortex (e.g., from layer 4 to layer 2/3). However, although ample data collected across cortical areas are consistent with the synthetic model, remarkably few, if any, compelling examples of such transformations have been observed between cortical layers of a single sensory area such as V1. Orientation tuning, direction selectivity, and ocular dominance are all observable within layer 4 neurons (Hubel and Wiesel, 1962; Sun et al., 2016). In other cortical areas, such as the somatosensory cortex, we have arguably even less insight into the synthesis of new response properties (Brecht, 2017). Receptive fields in the rodent barrel cortex and in the cat visual cortex tend to grow or change shape across layers (Brecht et al., 2003; Martinez et al., 2005), and some

evidence supports the de novo generation of complex cells and contextual properties such as end-stopping between layer (L)4 and L2/3 in cats (Alonso and Martinez, 1998; Hubel and Wiesel, 1962; Martinez and Alonso, 2001). However, the striking lack of concrete examples akin to orientation tuning, at least outside of the monkey V1, implies that the laminar circuit in a single cortical area is not set up to generate new types of feature selectivity. Furthermore, this hierarchical framework fails to account for a wide range of context-dependent phenomena observed in cortical activity, nor does it provide a compelling explanation for the profuse amount of feedback connections from higher cortical areas to lower ones (Felleman and Van Essen, 1991).

In contrast, the alternative framework that sees sensory processing as probabilistic inference can explain these "top-down" phenomena. In this scheme, neurons in different cortical layers play unique roles in computing the conditional probabilities that a given pattern of afferent neural input represents a specific sensory stimulus (Bastos et al., 2012; Rao and Ballard, 1999). The core notion is that cortical neurons, moment to moment, compare afferent input from each preceding stage with an internal generative model of the sensory environment conveyed by top-down projections, a model based on both the recent past and accumulated experience. Predictions of this model are passed from higher to lower stages (both across layers in individual areas and between areas) through feedback connections, and neurons in earlier stages compare these predictions with errors indicated by deviation from the afferent, "bottom-up" sensory data. In one version of this theory, principal cells in superficial cortical layers of primary sensory areas encode prediction errors, whereas those in deeper layers encode conditional expectations from which predictions are made (Bastos et al., 2012). Inhibitory interneurons within each cortical layer might be critical for canceling errors (i.e., incompatible predictions) when the predictive model matches the sensory data. This conceptual framework is attractive because one can assign specific functions to different layers and cell types that should be experimentally testable. However, at present, data supporting the inferential model is limited (see, for example, Homann et al., 2017). An intriguing variant on this theme is a "body model" for the somatosensory cortex in which the goal of S1 is to generate mental simulations of planned body actions. This model also assigns specific functions to each layer: body simulation to L4, sensory memory storage to L2/3, motor memory storage in L5, and relay of the top-down drive from M1 through L6 and back to L4 (Brecht, 2017).

More recently, the canonical circuit has been conceptualized less in terms of layers and more in terms of cell types that occupy specific layers and are connected by cell type-specific pathways (Harris and Shepherd, 2015). Although layer and cell type are closely intertwined, this is an important distinction. A layer-centric view implicitly assumes that at least some basic cortical computations can be understood mechanistically by analyzing the activity of neurons in just one layer; the cell type-centered view assumes that we cannot achieve a satisfactory understanding of any computation without taking into account coordinated activity across multiple layers. Although we retain layer as an organizing concept for the purposes of this perspective, we note that the experimental approaches we outline below apply equally well to cracking the function of cortical sublaminae or cortical cell types. Nevertheless, we argue that the cortical literature regarding layers supports the notion that, in a specific set of contexts, ensembles of neurons located in just one layer are sufficient to mediate key cortical computations, such as fast sensorimotor transformations. However, under most conditions, such as those involved in generating conscious sensory percepts, the basic unit of cortical computation is a neuronal ensemble spread across multiple layers or spread across multiple layers and cortical areas.

Layer-Specific Features of "Canonical" Cortical Circuits
Our empirical knowledge of cortical layers from which we can build theories of their function comes from four types of exploration: anatomy of single cells and their projections, connectivity of pairs of cells according to their laminar location and cell type, physiological responses to sensory stimuli, and activation or suppression of neurons in discrete layers. Anatomy and connectivity are the most absolute in that they do not depend on brain state or type of sensory stimulation. They constrain the types of computations layers can perform and the dynamics they can exhibit but, on their own, provide limited insight into function. Conversely, physiological perturbations are much less absolute in that the resulting data will depend on the brain state and the context in which they were obtained. However, they should provide the most direct insight into the different functions of cortical layers.

The often-repeated (although just as often maligned) notion of the "canonical cortical microcircuit" is largely based on studies of anatomy and connectivity in rodents, primates, and cats (Gilbert, 1983). These data have converged on a core model where thalamic input drives activity in a feedforward and sequential fashion from L4 to L2/3 to L5 and out to other cortical and subcortical regions (Armstrong-James et al., 1992; Binzegger et al., 2004; Douglas and Martin, 1991; Lefort et al., 2009). Although numerous examples of alternate connections exist (e.g., thalamus to other layers, L5 to L2/3, and L4 to L5), the anatomy of these neurons (i.e., their intracortical axons and dendrites) matches well with paired electrophysiological recording and circuit mapping via optical approaches (see full citations below). Because these and other pathways in the cortex have been extensively reviewed elsewhere (Callaway, 1998; Douglas and Martin, 2004; Feldmeyer, 2012; Gilbert, 1983; Harris and Shepherd, 2015; Thomson and Lamy, 2007), we will focus on data revealing the physiology and functional effect of the principal excitatory neurons in different cortical layers (Figure 1). The long-range input/output logic of the canonical microcircuit is organized by layer. L4 neurons are thought to primarily target their local neighbors (Binzegger et al., 2004). The principal neurons of L2/3 are intratelencephalic (IT) cells, meaning that their long-range axons project only to targets within the telencephalon, such as other cortical areas and the striatum. L5, sometimes called the primary cortical output layer, harbors IT cells as well as pyramidal tract (PT) cells, which send widely divergent projections to subcortical areas. L6 contains corticothalamic (CT) cells, which provide a major feedback projection to the thalamus, as well as IT cells.

Figure 1. Architectural Principles of the Cortical Layers

(A) Diagram of the major excitatory cell types of the cortical layers, including L4 neurons that project locally, L2/3 and L5 IT (intratelencephalic) neurons that project intracortically, and L5 pyramidal tract (PT) and L6 corticothalamic (CT) neurons that project subcortically.

(B) Schematic of the "feedforward" architecture of the neocortex, emphasizing the increase in the specificity of translaminar targets across the layers. L6 circuitry has been excluded for clarity.

(C) Schematic of the change in population size of principal cells in each layer (left) and the corresponding sparsity of their population responses to sensory stimuli (right).

Corticocortical pathways (i.e, inter-areal) are often conceptualized as being either feedforward, lateral, or feedback pathways (Felleman and Van Essen, 1991; Gămănuţ et al., 2018). Layer plays a key organizing role in this inter-areal hierarchical scheme (D'Souza and Burkhalter, 2017). Interestingly, as one moves along the canonical pathway, translaminar connectivity becomes increasingly specific (Figure 1). The primary sensory thalamus, constituting the input stage of the hierarchy, provides highly divergent output impinging on cells in all cortical layers (Cruikshank et al., 2010; Petreanu et al., 2009). L4 neurons exhibit strong recurrent intralaminar connectivity (Binzegger et al., 2004) and broadcast their output to all other layers but do not provide feedback to the thalamus. Similarly, L2/3 has minimal feedback connectivity to L4 but connects densely to both types of pyramidal neurons in L5 (Adesnik and Scanziani, 2010; Lefort et al., 2009). In turn, L5 IT neurons exhibit dense, asymmetric connectivity onto PT neurons, and, finally, PT neurons appear to connect primarily only to other PT neurons, providing minimal feedback to any of the earlier layers (Yamawaki and Shepherd, 2015). An exception to this pattern of increasing selectivity is an ascending connection from L5 IT cells to L2/3 (Binzegger et al., 2004). As for L6, earlier models of the canonical microcircuit based primarily on data from monkeys and cats proposed that L6 receives major input from L5 and from superficial layers and then "closes the loop" by projecting back to L4 (Binzegger et al., 2004; Briggs and Callaway, 2001; Douglas and Martin, 2004; Gilbert, 1983). In rodents, L6 CT neurons project to L5a and, to a lesser extent, L4 (Kim et al., 2014) and receive strong long-range inputs (Kinnischtzke et al., 2016). There are putative discrepancies in the basic cortical circuitry between different mammalian species, implying that a unifying "canonical cortical circuit" might not exist across mammals. However, despite their heterogeneity, the existing data still indicate that cortical circuits across brain areas and species share some common functional principles that are key to understanding their function.

Layer-Specific Perturbations: Insights and Challenges

Understanding any neural circuit requires perturbing it and observing changes in the computations it performs. Examples of putative computations cortical circuits implement are schematized in Figure 2. These include summations that give rise to oriented edge detectors (Figures 2A and 2B), signal amplification through recurrent excitation (Figure 2C), coincidence detection (Figure 2D), generation of a sparse code (Figure 2E), lateral integration that might facilitate contour or boundary detection (Figure 2F), and coding through synchronization (Figure 2G). Although the anatomy and physiology of the neurons in any layer can help us build theories and propose hypotheses for how layers contribute to each of these computations, only manipulating the activity of specific layers or subsets of layers can test and validate these theoretical hypotheses. Prior to the advent of cell type-specific perturbations (via opto- or chemogenetics), the primary tools for perturbation were chemical lesions (reversible or irreversible), cortical cooling, and electrical microstimulation. A common weakness of all of these tools is that precisely calibrating the spatial extent of the perturbation is extremely challenging, and they cannot be absolutely layer-specific because cortical neurons' dendrites and axons often stretch across laminar boundaries. This last fact even muddies the concept of what a layer is in the cortex because some deep-layer pyramidal neurons derive much of their synaptic input from their dendrites, which occupy different layers than the one in which their cell body resides (Larkum et al., 2018).

Layer-specific optogenetic manipulation overcomes this last problem, but the results of both chemical and optogenetic perturbations of cortical layers have often challenged the canonical model of information flow across the cortical layers. For instance, reversibly blocking L4 activity (by chemically silencing
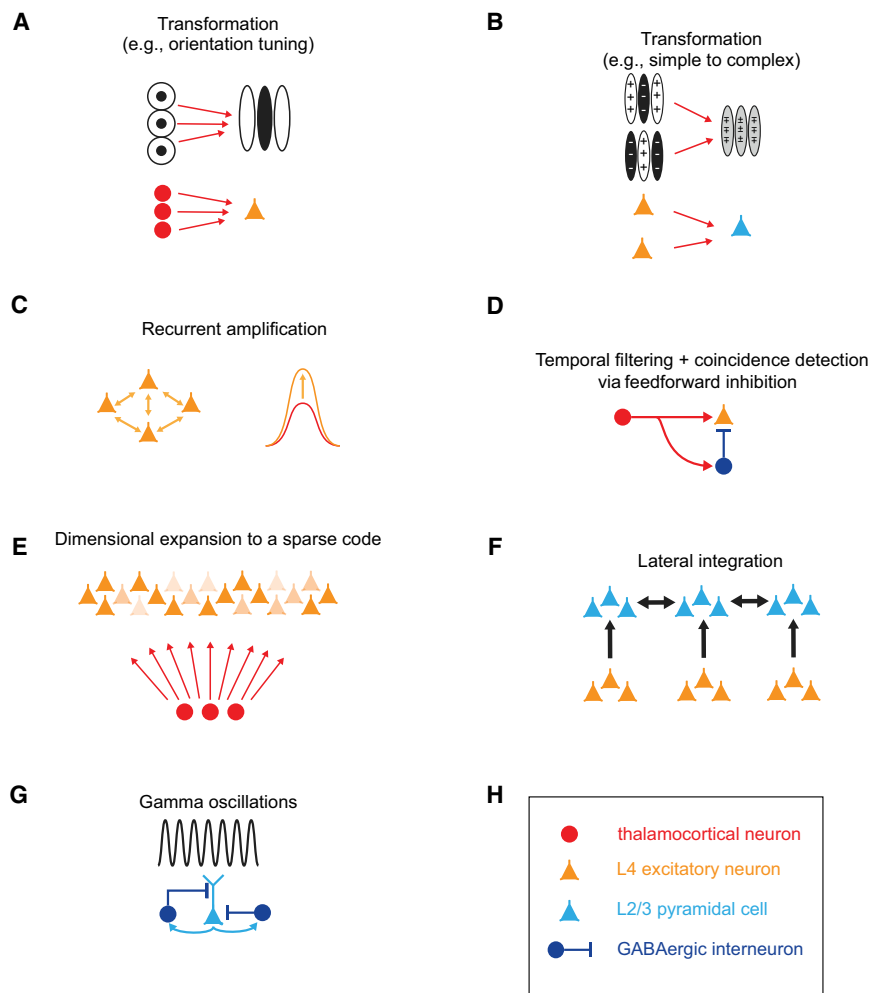
A Transformation
(e.g., orientation tuning)

B Transformation
(e.g., simple to complex)

C Recurrent amplification

D Temporal filtering + coincidence detection
via feedforward inhibition

E Dimensional expansion to a sparse code

F Lateral integration

G Gamma oscillations

H
- ● thalamocortical neuron
- ▲ L4 excitatory neuron
- ▲ L2/3 pyramidal cell
- ●— GABAergic interneuron

Figure 2. Examples of Circuit Motifs for Basic Sensory Transformations and Computations in L4 and L2/3
(A) Diagram of the thalamocortical circuit that generates orientation selectivity between the visual thalamus and primary visual cortex. Top: structure of receptive fields of the corresponding neurons.
(B) Diagram of a circuit between L4 and L2/3 that could generate complex cells from simple cells in V1. Top: structure of receptive fields of the corresponding neurons.
(C) Recurrent excitatory circuitry in L4 that linearly amplifies thalamocortical input.
(D) A simple feedforward inhibitory circuit that enforces coincidence detection between the thalamus and L4.
(E) Dimensional expansion of the sensory code between thalamus and cortical L4.
(F) Diagram of horizontal circuits in L2/3 that might contribute to contextual modulation.
(G) Schematic of the minimal recurrent excitatory and inhibitory circuit for generating gamma frequency oscillations in L2/3.
(H) Legend.

could facilitate activity in downstream layers by generating a more favorable excitation to inhibition ratio (Adesnik, 2018), a feature previously only identified in L2/3 of the barrel cortex (Adesnik and Scanziani, 2010).

The outcomes of these experiments argue that the traditional feedforward layer-centric model of cortical activity can account for some but not all aspects of cortical processing. However, a major problem persists with both chemical and layer-specific optogenetic manipulations: because of the highly inter-connected nature of neurons across layers and cortical areas, perturbation of cells in one layer will almost necessarily affect activity in other layers and areas. Therefore, it is difficult to attribute any measurable physiological or behavioral effect to the action of the pharmacologically or optogenetically targeted layer. Although this is a criticism that can be more generally leveled at any sort of brain perturbation (Otchy et al., 2015), it stands as one of the key challenges in interpreting data gained from perturbation studies.

Layers Acting Alone
Perhaps the most fundamental, first-order question concerning cortical layers is whether the lamination is functionally relevant. Species such as birds exhibit high cognitive abilities but have no such layering in their pallium; instead, computation seems to be organized around cells clustered into ''nuclei'' (Calabrese and Woolley, 2015). Strikingly, a mutant mouse with completely disorganized cortical layers has no immediately apparent sensory physiological deficits, suggesting that layers, per se, may not be critical for many aspects of cortical function (Guy and Staiger, 2017). This raises the alternative hypothesis that the layered structure of the neocortex is largely a consequence of

specific layers of the visual thalamus) does not block much of the sensory evoked responses in L2/3 of the visual cortex of anesthetized cats (Malpeli, 1983; but see Martinez and Alonso, 2001) but simultaneously suppressing higher visual cortex areas does (Mignard and Malpeli, 1991). In a similar vein, direct application of the action potential blocker lidocaine to the superficial layers of the somatosensory cortex has essentially no effect on whisker-evoked activity in L5 pyramidal cells (PCs) in sedated and paralyzed rats (Constantinople and Bruno, 2013; but see Wright and Fox, 2010). Direct optogenetic suppression of L4 in awake, locomoting mice leads to a modest reduction in sensory evoked activity in L2/3 of V1 or S1. However, it simultaneously leads to a potent disinhibition of activity in L5 because of a disynaptic translaminar inhibitory circuit between L4 and L5 (Pluta et al., 2015). Optogenetic suppression of L6 in awake mice also has a largely disinhibitory effect across most layers of the cortex, an effect attributed to the deactivation of a broadly inhibiting translaminar inhibitory neuron (Bortone et al., 2014; Olsen et al., 2012). In one study, optogenetic activation of L4, L2/3, or L5 in brain slices from the mouse primary visual cortex revealed that activation of each of these layers suppressed activity within their own layer by recruiting potent recurrent inhibition but

early cortical development and that layers have no intrinsic function in the adult brain. Instead, they just happen to harbor unique cell types that, through their translaminar circuits, constitute the functional substrate of cortical computation. However, lamination is a conserved feature of brain circuits from insects to fish to mammals, particularly evident in structures such as the retina, cerebellum, tectum, and hippocampus (Striedter, 2005). In these structures, lamination organizes many aspects of synaptic connectivity and function, such as the segregation of ON and OFF pathways in the retina. Thus, one is tempted to conclude that many circuits evolved lamination to achieve specific computational goals, even when it is not strictly necessary for computation in general.

We suspect that neural ensembles composed of co-active neurons distributed across layers are the major substrate of cortical computation. However, the selective effects of the perturbation experiments described above imply that, under certain circumstances, layers can have specific and identifiable effects on their local cortical circuit and even on behavioral output. These studies raise the possibility that, for particular computations or behaviors, only one cortical layer, or even a specific subtype within a single cortical layer, might be sufficient to execute the task based on its afferent input and its long-range output. This is surely possible from an anatomical perspective. Principal neurons across all layers receive bottom-up thalamic inputs, most receive top-down cortical input from higher cortical areas, and many project a long-range axon out of the cortex. If any of their long-range input pathways are sufficiently strong (or strong when provided in combination), many cortical neurons could operate independently of their local circuits. There is evidence to support this notion. A recent study identified a compact long-range circuit from the whisker to the barrel cortex in which a subset of L5 PCs projects directly back to pre-motor spinal trigeminal neurons that control whisker retraction (Matyas et al., 2010). Because L5 neurons can be driven directly by the thalamus (Constantinople and Bruno, 2013) or through the conjunction of bottom-up and top-down motor cortical input (Manita et al., 2015), under specific conditions, motor control of certain features of whisking (e.g., touch-induced pumps; Deutsch et al., 2012) might only require activation of these specific L5 PCs. If so, it would represent a complete sensorimotor behavior that could be cortically dependent but does not require activity of superficial layers at all. In a similar vein, because other projection subtypes of L5 PCs target various sub-cortical nuclei, we suspect that there might exist a suite of reflex-like sensorimotor behaviors that require minimal local processing within the primary sensory cortex and that, instead, rely exclusively on the integration of various long-range pathways to specific subtypes of L5 PCs. Conversely, because L2/3 PCs across species receive direct thalamic input and project to downstream cortical areas (Fitzpatrick et al., 1983; Meyer et al., 2010; Petreanu et al., 2009), various types of sensory computations in L2/3 could also be executed with limited involvement of L4–L6. Taken together, this suggests that there may be specific behaviors and cortical computations that can depend on a single layer or even a single cell type within one layer.

A Holistic Attack on Cortical Layer Function

Why has the extensive physiological data collected over the last few decades provided few concrete examples of new transformations that occur across layers? This absence is problematic because it has precluded the generation and validation of hypotheses concerning the differential roles of layers in sensory computation. We can consider several explanations: first, higher-order feature selectivity might not emerge across layers within a single area and only across cortical areas; second, physiologists have, until recently, lacked the proper tools to observe these higher-order features; third, most measurements have been made under too impoverished behavioral conditions; or fourth, data have been analyzed in the wrong conceptual framework (i.e., synthetic versus inferential).

Cortical circuits in any given brain area can process sensory input from multiple modalities, encode behavior or movement-related information, and directly contribute to motor control (Ibrahim et al., 2016; Iurilli et al., 2012; Matyas et al., 2010; Vinck et al., 2015). Accordingly, investigating the function of microcircuits in primary sensory cortex function by only focusing on sensory transformations is sure to lead to a partial understanding of cortical computation at best. Identifying "canonical computation" will therefore require investigating how the canonical circuit operates as but one functional part of a much larger architecture. Although the corticocortical network is hierarchically organized, it is also massively recurrent and fundamentally sensorimotor (D'Souza and Burkhalter, 2017; Gämänuţ et al., 2018). It could be that certain computational principles of cortical microcircuits will be comprehensible only in the context of naturalistic, sensorimotor behaviors (Juavinett et al., 2018; Krakauer et al., 2017).

All of these issues enumerated above have likely hindered progress toward a fuller understanding of cortical layers and cortical computation in general. Based on the available data from primates, from the synthetic perspective it appears that major new transformations primarily emerge across cortical areas and not between layers of any individual area (Liu et al., 2016; Ziemba et al., 2016). Alternatively, rather than further refining the encoding of sensory stimuli, different layers of the sensory cortex might instead be specialized to perform different computations with respect to non-sensory factors, such as the movement, motivation, or stored memories of an animal. Such computations might only become evident when the appropriate non-sensory factors are taken into account (Keller et al., 2012; Roelfsema et al., 1998; Saleem et al., 2013).

We propose that differentiating between these alternative hypotheses for the computational role of cortical layers requires a holistic approach to monitoring and manipulating activity along the cortical axis. Rarely does one have enough *a priori* knowledge for the specific path information will take under a specific condition so that focused, small scale recording will be sufficient. Furthermore, key types of transformations might only become apparent at the population level and across neurons that are spatially intermingled with other ensembles and distributed across multiple cortical areas (Ma et al., 2006; Pouget et al., 2000; Wohrer et al., 2013). Thus, the appropriate holistic approach includes recording densely across all layers, identifying the cell types of the recorded cells within this volume,

recording neural activity simultaneously in multiple connected brain regions (e.g., thalamus, secondary sensory areas, and relevant motor cortices), and perturbing individual layers and cell types with high spatial and temporal precision. Just as importantly, behavioral data on motor actions, brain state, attentional focus, and potential goals should be acquired to inform all analysis. Where possible, experiments should target ethological behaviors in as natural a context as the experimental design will permit. Applying this approach comparatively across different cortical areas and species will help reveal the circuit elements that are conserved in different cortical systems, which will be instrumental in identifying simplifying motifs. Whether one considers cortical computation from the synthetic or inferential vantage point, such an all-inclusive experimental approach can provide the comprehensive data needed to determine whether cortical circuits implement canonical computations and, if so, how specific layers and cell types contribute to specific computations or behaviors.

In particular, from such a dataset, one can execute several key critical tests. The first is an analysis of the neural activity across layers during a specific computation. One must determine which layers (and which cell types within each layer) show activity that can explain how the computation is mediated and then track the transformation of the encoded signal across the layers, from cortical input to cortical output, to motivate clear hypotheses for which layers are causally involved. A simple example of this process is the lateral geniculate nucleus (LGN) to V1 transformation for the synthesis of oriented receptive fields, which was apparent even under conditions of sparse sampling because it probably dominates the first stage of V1 processing.

After observation, the second test requires a causal manipulation of the layer(s) and cell type(s) thought to be involved in a computational process based on the results of step one. To achieve this, first one would want to abolish the computation by eliminating the activity of the putative contributing layers. However, for the reasons mentioned above, loss-of-function or gain-of-function experiments cannot, on their own, entirely discriminate between various circuit models that could account for the same effect. To overcome this challenge, more sophisticated and more integrated approaches are required. Consider two examples. L6 activity has been proposed to control both the size of upper layer receptive fields and their gain by recruiting intracortical inhibition in cats and mice (Bolz and Gilbert, 1986; Olsen et al., 2012). Silencing L6 activity enhances upper layer responses and expands their receptive fields, but this could be due to direct intracortical effects, or it could operate indirectly through L6's effect on the LGN. One study in awake mice, therefore, employed simultaneous recording from the LGN and V1 layers and demonstrated that the disinhibition of the cortex during optogenetic stimulation of L6 preceded effects on the thalamus, arguing for a direct effect (Olsen et al., 2012).

Despite the elegance of this example, temporal offsets in the effects of perturbations might not always be able to distinguish between competing models. Consider the possible explanations for the synthesis of complex receptive fields in L2/3. One model argues for a simple summation over phase-offset iso-oriented simple cells in L4 (Hubel and Wiesel, 1962; Martinez and Alonso, 2001; Figure 2G), whereas another argues that this computation depends on V2 feedback (Mignard and Malpeli, 1991), and a third proposes that appropriate summation over geniculocortical input directly to L2/3 is sufficient to account for it (Hoffman and Stone, 1971). One study showed that inactivating L4 suppresses L2/3 complex cell activity in anesthetized cats (Martinez and Alonso, 2001). This manipulation, however, only proves that L4 ascending input is required for L2/3 visually evoked drive but does not unequivocally demonstrate that complex cells emerge from the summation of L4 simple cells with spatially offset ON and OFF subfields. Instead, if one could photo-activate a small subset of iso-oriented L4 neurons with appropriately arranged subfields and recapitulate the non-phasic response of the L2/3 complex cell, then one could distinguish between these models. Thus, a more general approach for revealing the mechanism of a transformation across layers with great certainty is to entirely reconstitute the computation by activating the appropriate ensemble of presynaptic neurons with extreme spatial and temporal precision. An alternative scheme is retrogradely and transsynaptically labeling an ensemble of cortical neurons from one cortical starter cell with a GCaMP6-expressing rabies virus (Wertz et al., 2015). This permits measuring the functional responses of a large subset of putative presynaptic cells to a single cortical neuron. Employing a virus that drives a microbial opsin in addition to GCaMP6 could further enable photo-activation of the presynaptic ensemble for causal, functional tests.

Although we have focused on how the above approach would apply to sensory transformations, a holistic attack on cortical layer function will also require a holistic perspective on how cortical layers and cell types contribute to behavior, including perceptual learning, decision-making, and motor control. In the motor cortex, a key conceptual advance has been to investigate how cortical dynamics robustly control movements instead of asking how cortical populations represent movement parameters (Churchland et al., 2012; Gallego et al., 2017; Michaels et al., 2016). By the same token, in the sensory cortex, it may be fruitful to shift the focus from asking how sensory information is represented and transformed per se and instead ask how sensory cortices and their component layers and cell types subserve the broader function of allowing an animal to flexibly adapt its behavior to changing environmental conditions in its environment.

Cracking Layers: New Tools from the Technical Revolution in Neuroscience
The holistic approach that understanding the cortex calls for is no longer beyond reach. Several emerging technologies will soon make combining large-scale recording with extremely precise perturbations of neural activity experimentally possible. We review several key advances and propose new experimental paradigms combining these advances that could address the function of cortical layers. The reader is referred to other reviews on this subject for more in-depth consideration of some of these approaches (Emiliani et al., 2015; Luo et al., 2018).

Massively Parallel Electrophysiology
Unquestionably, measuring neural activity simultaneously across cortical layers and cortical input and output structures will be highly advantageous for gauging the computational role and effect of each cortical layer. For more than a decade, laminar

multi-electrode arrays have represented the key tool in the cortical physiologist's arsenal for addressing this challenge (Csicsvari et al., 2003), but their relatively low channel count and large size have precluded achieving the needed scale. Over the same period, multiphoton calcium imaging has begun to provide an incredibly dense view of local neural computation, albeit in circumscribed regions and specific layers of the cortex (Stosiek et al., 2003). However, major advances in both silicon probe technology and large-scale calcium imaging place us at the threshold of probing neural activity at the scale and speed needed to address the problem at hand.

Leveraging new fabrication techniques, extremely compact multi-electrode arrays with up to 1,000 contacts now permit simultaneous physiological access to dozens of brain areas and, just as importantly, densely within and across layers (Jun et al., 2017; Shobe et al., 2015). Clever mechanical systems for inserting multiple such probes into the same head-fixed animal should allow tracking the activity across all layers of each area and across a series of connected areas in the cortical hierarchy simultaneously (Figure 3A). This possibility is fundamental; it will allow investigators to track the evolution of neural activity, with the resolution of single neurons and single spikes, from massive populations of neurons representing the key nodes of computation. Although these probes do net yet offer integrated micro-optics, structured illumination from a separate optical system could provide simultaneous high-resolution optogenetic manipulation, at least for superficial cortical areas. Adding micro-light emitting diodes (LEDs) (Scharf et al., 2016) or optical waveguides (Wu et al., 2013) directly onto these probes could provide optical control over their full extent so that investigators could optogenetically control specific layers within specific cortical areas (Figure 3B). Furthermore, the integration of nano- or microfluidic ports into the same probes could additionally provide spatiotemporally resolved pharmacological manipulations that are needed for determining how various neuromodulators influence laminar interactions and computations (Canales et al., 2015).

Large-Scale Multiphoton Imaging

Major strides are also being made in scaling up the field of view, depth, and speed of multiphoton calcium imaging. New approaches for engineering the optical wavefront or point spread function of the ultrafast laser provides volumetric imaging of neural activity across depth, with limited compromise in the overall data acquisition speed (Lu et al., 2017; Prevedel et al., 2016; Thériault et al., 2014). Random access imaging via ultra-fast scanning or volumetric access with spatial light modulators represent alternative approaches (Katona et al., 2012; Duemani Reddy et al., 2008; Yang et al., 2016). These tools open up the possibility of capturing most if not all of the neural activity of an entire layer in a single cortical area or possibly even of two layers or more at a time. Further advances, such as splitting the scanning beam and multiplexing the volumetric imaging path, could, in theory, allow coverage of all neurons across all layers in a single area (Figure 3C; Cheng et al., 2011). Obtaining such comprehensive datasets on neurons imaged simultaneously across layers will potentially reveal types of transformations or forms of computations not readily apparent with sparser sampling techniques.

Just as sampling all cortical layers in a single cortical area extremely densely is surely advantageous, sampling a series of connected cortical areas simultaneously and at the same level of density could be even more revealing. For instance, dense sampling of L2/3 of V1 and multiple layers in V2 at the same time should unequivocally aid in better understanding how the output of L2/3 leads to new feature selectivity in V2. To achieve this, several groups have designed low-magnification optical pathways that permit mesoscopic views of the cortex with single-neuron resolution (Sofroniew et al., 2016; Stirman et al., 2016; Tsai et al., 2015). In the case of the mouse visual cortex, this places V1 and many of its associated V2-like downstream structures within the same accessible imaging field. For S1, one should be able to capture data in S1, S2, M1, and more frontal pre-motor areas all at the same time.

A major outstanding challenge for two-photon microscopy, however, is imaging the infragranular layers. Two-photon imaging quality degrades with depth, and in species such as rats and primates that have thicker cortices than mice, L5 and L6 may be out of reach. Even in mice, imaging down to L6 is typically only possible when selectively labeling the deep layers because out-of-focus multiphoton excitation of fluorophores in superficial layers leads to poor signal contrast. Approaches to overcome this challenge include implanting optics such as gradient index (GRIN) lenses and micro-prisms (Jung et al., 2004), the latter of which offers edge-on views of the cortical lamina (Andermann et al., 2013). These approaches, however, necessarily cause substantial tissue damage, a problem that may be avoided by employing three-photon imaging, which can reach more than 1.2 mm into the rodent brain (Ouzounov et al., 2017). Taken together, some combination of these optical and technical advances will soon permit the measurement of neural activity at the requisite spatial and temporal scale to make major new discoveries in understanding the function of each layer in sensory coding.

Multiphoton Optogenetics

Although conventional optogenetic approaches readily provide potent and reversible control of genetically defined subtypes, one-photon photo-stimulation with visible light offers little spatial precision because of limits on focusing light along the axial dimension and the scattering properties of cortical tissue. Multiphoton excitation, however, can provide tight spatial and temporal control of the photo-activated ensemble (Mardinly et al., 2018; Packer et al., 2012; Paluch-Siegler et al., 2015; Papagiakoumou et al., 2010; Pégard et al., 2017). Under appropriate conditions, resolution can reach to the cellular level, permitting individual neurons to be photo-activated or photo-suppressed based not only on their genetic subtype but also on their precise location and even based on their functional properties, such as their tuning to a specific stimulus feature (Packer et al., 2015; Rickgauer et al., 2014). With the introduction of computer-generated holography, the optical wavefront can be shaped in three dimensions, allowing for simultaneous targeting of neurons across layers. Over the past decade, major strides have been made in engineering the multiphoton optical system and the opsin proteins themselves to offer potent control over large neural ensembles in the intact brain (Hernandez et al., 2016; Mardinly et al., 2018; Pégard et al., 2017; Prakash et al., 2012;

Figure 3. New Methods for the Dissection of Cortical Layer Function

(A) Schematic of a mechanical setup for the simultaneous insertion of multiple ultra-high-density multielectrode arrays (courtesy of New Scale Technologies).

(B) Schematic of a multi-layer silicon probe integrating electrodes and optical waveguides for recoding and optogenetics (courtesy of V. Lanzio and S. Cabrini).

(C) Schematic of a very large field of view high-speed volumetric two-photon microscope for densely imaging across multiple layers and multiple connected cortical areas simultaneously.

(D) Schematic of using a holographic multiphoton microscope to selectively perturb individual layers in the same animal on different trials.

(E) Schematic of a paradigm for mapping the physiological responses of L2/3 and L4 cortical neurons with calcium imaging and then optogenetically stimulating a precise ensemble of presynaptic L4 neurons while recording from a L2/3 neuron in an attempt to recapitulate the sensory response properties of the L2/3 neuron.

(F) Sequence of proposed experiments for correlating the sensory physiology, synaptic connectivity, and transcriptionally defined cell types of densely imaged cortical neurons across layers.

Yang et al., 2018). Using this new technology makes it possible to execute a broad array of experiments targeting cortical layers, their cell types, or groups of cells across layers based on a common physiological property (Figure 3D). For instance, it should now be possible to execute the experiment suggested by Hubel and Wiesel (1962) to address the mechanism of the synthesis of complex receptive fields in V1; one could map the receptive fields and orientation tuning of L4 neurons first with multiphoton calcium imaging and then photo-stimulate an identified set of these neurons while monitoring activity in L2/3 neurons (Figure 3E). This would allow investigators to literally execute the type of experiments called for by Hubel and Wiesel (1962) in their pioneering study.

More generally, multiphoton holographic optogenetics can be used to manipulate the activity of individual layers independently and all in the same animal. Although prior studies have used one-photon optogenetics and layer-specific Cre lines to activate or silence neurons specifically in a single layer, this does not allow direct comparison of perturbations to each of the cortical layers in the same animal and under the same condition. It should also be noted that most layer-specific Cre lines that are available invariably label a subset of neurons in a given layer (Adesnik, 2018), leading to some ambiguity in how the data should be interpreted. With multiphoton optogenetics, one can target the opsin to all cortical neurons in all layers and use its superior spatial resolution to choose, moment to moment, which layer to perturb.

Multiphoton Voltage Imaging

Although multiphoton calcium imaging has revolutionized population recording in the brain, it suffers several drawbacks, principally because of its reliance on calcium flux as an indirect measure of neural spiking. Calcium indicators such as GCaMP6 provide only a highly temporally filtered measurement of spike rates changes (Chen et al., 2013), precluding precise assignment of spike times, and, under many conditions, may not report spike rates with a linear transfer function. One way to overcome these issues is to image voltage (and, thus, action potentials and potentially sub-threshold activity) directly. Achieving cellular resolution, millisecond precision voltage imaging of cortical layers would allow us to address major questions about cortical computation as it might be implemented by each layer. For example, correlated activity in cortical ensembles at fast timescales has been proposed to be a key means of encoding and computing aspects of sensory stimuli (Salinas and Sejnowski, 2001). Calcium imaging cannot properly report fine timescale correlations, and multi-electrode physiology cannot sample neuronal activity densely enough and does not as easily allow one to identify the genetic identity of the recorded cells (except via "opto-tagging," which has various inherent drawbacks) (Lima et al., 2009). Multiphoton voltage imaging could, in principle, overcome these challenges and address the long-standing debate of the importance of spike timing in sensory coding in the cortex. Extending this type of analysis to data acquired simultaneously from neurons in multiple layers has the power to address whether and how spike timing, synchronization, and neuronal correlations contribute to the differential roles of cortical layers during sensory perception and behavior.

There are at least two major technical hurdles that must be overcome to achieve large-scale, fast population voltage imaging. First, we require genetically targetable, fast voltage sensors with sufficient sensitivity to robustly report action potentials from many neurons simultaneously in vivo. Second, we require an optical system that can simultaneously sample the fluorescence of such a sensor from large neuronal populations at a high enough rate to capture all emitted spikes. Recently, there has been dramatic progress in developing new and ever more sensitive voltage sensors, including genetically encoded voltage indicators (GEVIs) or synthetic dyes that can be genetically targeted by ligation to a genetically expressed receptor. These have been reviewed elsewhere (Canepari et al., 2015; Luo et al., 2018; Perron et al., 2009; Xu et al., 2017), so we focus on how existing voltage sensors meet the key attributes of the ideal voltage sensor for cracking cortical layers. The ideal sensor should be genetically targeted, show high voltage sensitivity, have fast kinetics, and be bright, photo-stable, and efficiently excited by multiphoton excitation. Of the available voltage sensors, none meets all of these requirements, but the rate of progress has been substantial. The accelerated sensor of action potentials (ASAP) and Arclight GEVIs show usable two-photon cross-sections but have comparatively less voltage sensitivity, slower kinetics, and faster photo-bleaching than the Archaerhodopsin (Arch)-based GEVIs (Akemann et al., 2013; Chamberland et al., 2017; Xu et al., 2017). However, the latter show extremely low light sensitivity, precluding excitation in the multiphoton regime (Hochbaum et al., 2014). The Arch-based GEVIs can capture sub-threshold responses with reasonable signal to noise (Piatkevich et al., 2018) and, under restricted circumstances, can be used in vivo with one-photon illumination; however, their requirements for extremely high laser excitation powers have so far limited simultaneous in vivo imaging to less than a dozen neurons (Adam et al., 2018). Synthetic voltage dyes, which could be more rapidly customized to achieve the features of the ideal voltage sensor, represent a promising alternative path (Miller, 2016). They can be genetically targeted in a variety of ways, although the downside is that the dyes must be micro-injected into the brain, probably repeatedly for chronic imaging over several days. In light of these recent advances, we believe that a suitable voltage sensor for in vivo population voltage imaging is on the near horizon.

The second challenge, that of sampling the fluorescence of the voltage sensor across large populations of neurons at extremely high speeds, presents perhaps an even larger obstacle. Conventional raster-scanning multiphoton imaging used with calcium indicators is 1–2 orders of magnitude too slow to sample voltage for capturing action potentials with millisecond precision. Random access multiphoton imaging could address this issue in part because of its much higher speed of beam steering (Bullen and Saggau, 1999) but, in its current form, might still not provide the sampling rates of the large populations ideally required and will provide a relatively low signal because of the very short dwell time of the laser on each neuron. A second alternative is to use temporally focused multiphoton holography, as described above for optogenetic stimulation, to illuminate a targeted ensemble of neurons and to capture their fluorescence emission on a high-speed camera (Bovetti et al., 2017; Tanese

et al., 2017). The major advantage of such a scheme is that the target cells can be illuminated simultaneously, continuously, and across their entire somata, putatively providing large signals at a temporal rate only limited by the frame rate of the imaging sensor, which can reach into to 5–10 kHz range. The two challenges for such a system would be the limit of laser energy that can be delivered to the brain without causing heating or phototoxic damage and computational challenges with de-mixing the scattered emission from the illuminated cells on the imaging sensor, particularly from densely labeled tissue where the somatic membranes of adjacent neurons can physically touch.

For any of these voltage imaging approaches, it seems likely that imaging a complete volume of tissue comparable with what can be achieved with two-photon calcium imaging will be unachievable without truly dramatic improvements in voltage dyes. Instead, one can conceive a hybrid approach wherein cortical cells co-express a calcium indicator and a voltage sensor. First, the response properties of large populations of cortical neurons can be probed via calcium flux, and then a restricted subset that shows physiological features relevant to the experiment can be probed in a much more targeted fashion for voltage imaging. In the case of the sparse population codes in most cortical layers, such a hybrid approach should be particularly advantageous because, under many circumstances, only a small fraction of the neurons in the imaging volume show significant responses to the stimulus set or behavioral task under study.

Transsynaptic Viral Tracing
Despite the power of the aforementioned approaches, none provide direct access to populations of cortical neurons that are synaptically coupled. Enriched synaptic connectivity is arguably a crucial hallmark of the neural ensembles, within or across layers, which represent the fundamental units of cortical computation. Genetic labeling strategies that can drive fluorescent proteins, activity indicators, or actuators in synaptically coupled ensembles thus represent a key tool for dissecting the laminar basis of cortical computation. Despite the existing knowledge of interlaminar connectivity, transsynaptic labeling strategies can become a key tool in the cortical physiologists' arsenal for correlating the structure and function of cortical layers. The primary strategy to achieve this is retrograde transsynaptic labeling with deletion mutant rabies viruses that can only jump one synapse (Wickersham et al., 2007). This approach has gained wide adoption for mapping brain inputs to specific cell types and, when the tracing is initiated from just one neuron, permits direct imaging of the spatial distribution and even response properties of large numbers of cortical neurons that converge onto single cells (Marshel et al., 2010; Rancz et al., 2011; Wertz et al., 2015; Wickersham et al., 2007). More recent studies have described anterograde transsynaptic viral tracers that can complement the retrograde rabies strategy (Lo and Anderson, 2011; Zeng et al., 2017; Zingg et al., 2017). These approaches have at least two caveats that have largely remained unresolved: to what extent are the labeled neurons truly synaptically connected and what bias exists in the labeling because clearly not all presynaptic cells are labeled; in other words, what is the false positive and false negative rate of the tracing?

This issue was originally addressed by ground truth paired whole-cell recordings in organotypic slices (Wickersham et al., 2007); however, cultured slices harbor highly aberrant wiring patterns. Similar validation in brain slices or intact brain tissue, although perhaps vastly more difficult to achieve, is essentially lacking, implying that those relying on this tool are doing so somewhat blindly. These are challenging questions to address but would be worthwhile to answer so that studies employing transsynaptic tracers could be interpreted in the proper light. As an approach, transsynaptic tracing has the potential to directly reveal and permit the perturbation of the neural ensembles at the core of cortical function with relatively simple tools. Long-range projection patterns can also be probed with a combination of single-cell labeling, whole-brain fluorescence imaging, and multiplexed analysis of projections by sequencing ("MAPseq"), which identifies projections by sequencing axon-targeted DNA bar codes (Han et al., 2018).

Genetically Labeling Correlated Neuronal Ensembles
An attractive alternative strategy for labeling neurons that exhibit correlated activity *in vivo* is through genetic approaches that drive reporter genes under the control of activity-dependent promoters, principally the *c-fos* locus (Barth et al., 2004). This general strategy allows one to express a fluorescent protein, activity dye, or actuator in a large, brain-wide ensemble of neurons that exhibit *c-fos* promoter activity in some restricted time frame. This tool could be particularly useful for labeling *trans*-laminar ensembles united by a common *in vivo* function for further functional imaging or perturbation. Originally this strategy was done with GFP (Barth et al., 2004) but was then extended to driver systems, including the tetracycline transactivator (tTa), Cre recombinase, inducible Cre, and, most recently, the avian receptor for pseudo-typed viruses (Guenthner et al., 2013; Reijmers et al., 2007; Sakurai et al., 2016). It has been used to great effect to label neuronal ensembles in various brain nuclei that might encode memories because they were co-active during memory encoding, retrieval, or consolidation (Liu et al., 2012; Reijmers et al., 2007). The main caveats to this approach are that the *c-fos* promoter shows variable basal levels of activity according to cell type and cortical layer and that the relationships between neuronal activity, *c-fos* promoter activity, and downstream indicator or actuator expression remain largely undefined and may depend on cell type and other contingencies. Therefore, like viral transsynaptic labeling, identifying neural ensembles with activity-dependent transcription represents a particularly powerful tool in the neuroscientists' toolbox that requires more thorough validation. When further validated or suitably refined, however, it could likewise lead to transformative advances in our understanding of cortical layer function.

An alternative to *c-fos*-based labeling is turning on reporter transgenes with light. Such strategies include photo-activatable recombinases (including Cre and Cas9) that could allow expression of genetically encoded indicators or actuators with exquisite spatial precision when combined with multiphoton activation of the recombinase (Schindler et al., 2015). This could permit a range of new types of experiments not possible with existing approaches. For example, an experimenter can identify a functionally correlated ensemble of neurons with two-photon imaging

and then photo-label these neurons to drive the expression of ChR2. Subsequently, these neurons can be stimulated with simple, fiber-optic-based one-photon photo-stimulation under the freely behaving condition. One could selectively photo-label neurons in just one cortical layer and compare the effects of perturbing their activity with a similar ensemble composed of neurons distributed across several layers. Although the pattern of optical stimulation would not have the specificity of multiphoton optogenetics, the target population would be just as specific, and this approach would obviate the need for head fixation or the complex hardware required for holographic optogenetics.

Cell Typing by *In Situ* *RNA Profiling*

All cortical layers are composed of multiple cell types that have been differentiated by their molecular markers, morphology, intrinsic physiology, and local and long-range connectivity. To understand how layers, by virtue of containing a discrete set of identifiable circuit elements, differentially contribute to cortical computation, we need to determine how each cell type in each layer differentially encodes sensory stimuli or correlates with behavior. Indeed, abundant evidence indicates that distinct cell classes exhibit strikingly different responses to varying stimulus features, brain states, or behavioral choices. Furthermore, we would also like to determine the complete synaptic wiring diagram between all of these cell types in each layer, which is necessary for both conceptual and quantitative models of cortical circuit dynamics. Although an absolute parsing of the number and distinction of cell types is a matter of ongoing and fruitful debate, the more important question is how one can identify the myriad cell types of neurons in the same tissue that can be probed *in vivo*. The most popular approach so far has been to genetically label a specific cardinal type using transgenic lines and identify them *in vivo* with two-photon imaging or electrophysiologically with opto-tagging. The downside is that the number of cell types that can be identified in the same animal is limited by the number of independent reporter transgenes (with two or three being a typical practical limit). An alternative approach is to register *in vivo*-imaged tissue with *post hoc* immunostaining for multiple cell markers (typically up to three). This has been used to correlate sensory responses with three of the major markers for cortical GABAergic neurons (Kerlin et al., 2010). However, immunostaining is limited by both the number of available antibodies and the number of separable color channels. *In situ* RNA profiling is an exciting new technology for overcoming this limitation. This emergent technology (sequential fluorescence in situ hybridization, termed "seqFISH" or multiplexed error-robust fluorescence in situ hybridization, termed "MerFISH") uses *in situ* temporal multiplexed hybridization of a user-defined set of DNA bar codes to quantitatively identify the set of mRNAs each neuron in a tissue section expresses (Chen et al., 2015; Lubeck et al., 2014). Importantly, the appropriate set of bar codes to differentiate cortical neurons can be drawn from recent single-cell RNA sequencing datasets of cortical neurons (Tasic et al., 2016). This approach was recently validated in hippocampal brain slices (Shah et al., 2016).

Based on this technology, we can propose two types of experiments to obtain powerful new datasets on the *in vivo* physiology and synaptic connectivity of all of the transcriptionally defined sub-classes that compose each cortical layer. First, one could sequentially image the activity of a volume of cortex *in vivo* and then assign each imaged neuron to a specific transcriptionally defined cell class by *post hoc* seqFISH and careful image registration (Figure 3F). In one fell swoop, one could thus begin to define the physiological responses of more than 40 cell classes for highly quantitative, within-tissue comparison. Second, one could map the monosynaptic connectivity among all the cortical cell types composing each layer by combining singe cell-resolution multiphoton optogenetic mapping with patch-clamp electrophysiology, ideally *in vivo* to avoid cutting any projections. In this latter scheme, one labels all cortical neurons with a soma-targeted opsin and then systematically maps presynaptic neurons to a single target patched neuron (Baker et al., 2016). Subsequently, one probes the mapped tissue with seqFISH. This then provides the transcriptional identity of both the target neurons and all of its positively identified presynaptic partners. Perhaps a particularly exciting possibility would be to combine all of these steps into a single pipeline: imaging a volume of cortical tissue *in vivo* to obtain physiological responses across layers, then mapping the monosynaptic inputs to a small number of these neurons with patch-clamp and multiphoton stimulation, and finally assigning cell types with seqFISH. Such a scheme would allow physiologists to begin to directly correlate synaptic connectivity with the *in vivo* physiology of all the major subtypes of cortical neurons. It could provide, for example, a nearly complete picture for how complex cell receptive fields might emerge in L2/3 by identifying all of a complex cell's presynaptic neurons and quantifying their synaptic strengths. More generally, it can take a type of transformation that is only observable across cortical layers and track it back to the presynaptic pool of neurons that should generate it and provide much of the needed data to build a highly realistic model of the computation.

Neural Data Science

Continued advances in data analysis, experimental design, modeling, and theory are critical for harnessing and fully exploiting the new experimental techniques available for observing and perturbing cortical circuits. These topics have been addressed elsewhere (Paninski and Cunningham, 2018), so we will summarize the most relevant points. As the number of neurons that can be simultaneously recorded continues to grow exponentially, neuroscientists must adapt their analytical techniques. In addition to the data engineering challenges associated with collecting large datasets, neural data typically require extensive preprocessing; multi-electrode array data require spike sorting, and calcium imaging datasets require denoising and motion correction, demixing to isolate signals from individual neurons, and temporal deconvolution to recover estimates of spiking activity. These processes have historically been expensive in terms of time and labor because they have typically involved some level of manual involvement of the experimenter for quality control. An issue that must soon be addressed is that the preprocessing steps used by different groups vary considerably; as datasets grow, broadly accepted standards will become increasingly important for rigorously comparing the outcomes of different studies (Harris et al., 2016). One potential solution is the increasing development of fully automated pre-processing (such as spike sorting and calcium source separation) that requires minimal human input (Chung et al., 2017; Dhawale et al.,

2017; Pnevmatikakis et al., 2016). Not only is such automation necessary for handling the greatly expanding datasets, but they could ensure that different groups would extract the same neural signals given the same data, promoting reproducibility. The challenge is that these tools, just like the experimental approaches listed above, must be rigorously and extensively validated with technically challenging ground truth experiments involving simultaneous single-cell electrophysiology (Harris et al., 2000). Thankfully, the development and adoption of new analysis tools has increasingly automated these steps, which eases a key bottleneck for data acquisition and should improve the quality and reproducibility of preprocessing. The scalability of these methods to ever larger populations and to longer timescales is also a priority.

Closed-Loop Experimental Design

As increasingly sophisticated manipulations become available to experimenters, the question of how to optimally design these manipulations to definitively test key hypotheses becomes increasingly important and challenging. For instance, the high spatial and temporal resolution of multiphoton optogenetics, where different ensembles of neurons can be photo-activated every few milliseconds within a large volume of thousands of neurons, presents a combinatorial explosion of different possible patterns with which to perturb the system. How the experimenter chooses among this nearly limitless repertoire of stimulation patterns to gain insight into the circuit is far from trivial. One useful direction is generating stimulus patterns that mimic endogenous dynamics and then parametrically altering key parameters of these patterns, such as the distribution of firing rates, spike times, and neuronal correlations.

A companion problem, however, is that perturbations of highly recurrent networks, such as cortical circuits, are just as difficult to interpret. Indeed, the canonical model has largely failed to predict the results of even simple, relatively non-specific manipulations of individual layers. Theoretically, highly recurrent networks can be dissected via simultaneous manipulations of multiple components. However, even in very simple cases, this is extremely challenging because of the combinatorial explosion of perturbations that might need to be performed (Kumar et al., 2013). Tools like multiphoton optogenetics provide a potential solution to this problem, but, arguably, an experimenter, no matter how knowledgeable, cannot predict by intuition alone which perturbations to implement to overcome these challenges.

Fortunately, new tools are making it increasingly possible to extract meaningful data in real time (Chung et al., 2017; Friedrich et al., 2017), enabling closed-loop experiments based on the online analysis of neural activity. This has two major implications. First, closed-loop design could address many of the challenges that arise as experiments become increasingly sophisticated. Closed-loop design will allow the exploration of otherwise intractable spaces, such as mapping thousands of synaptic connections or estimating higher-order terms of nonlinear stimulus-response functions (Bolus et al., 2018; DiMattina and Zhang, 2013; Grosenick et al., 2015). Second, closed-loop design should enable whole new classes of experiments that could provide insight into cortical dynamics not obtainable with more conventional approaches. For instance, many recent studies of cortical computation have used dimensionality reduction- and dynamical systems-based approaches to identify latent structure in the activity of cortical populations. A particularly exciting direction for future studies will be to identify intrinsic manifolds in this space online and then perform optogenetic perturbations with respect to this structure (Jazayeri and Afraz, 2017) (e.g., compare the effects of perturbations that allow the population to remain within the intrinsic manifold versus ones that push the population outside of it). More generally, although these analyses have begun to point to interesting differences between cortical areas (Russo et al., 2018; Seely et al., 2016), they have so far been used in a manner that is largely agnostic to the layer and cell type of the recorded neurons. An important goal for future work will be to identify how the structure of cortical population dynamics relates to the anatomical structure of the underlying microcircuits. Ultimately, we predict that closed-loop design will allow the field to shift from the hammer toward the scalpel, forgoing manipulations that dramatically alter the activity of many neurons for ones that precisely perturb smaller ensembles while leaving the circuit as a whole within its normal regime of activity.

Missing Tools in the Cortical Physiologists' Arsenal

Despite the advent of new technical approaches for monitoring and manipulating neural activity, there is one key tool that is noticeably lacking in the neurophysiologist's toolbox: the ability of synapse-specific silencing between user-defined pre- and post-synaptic neurons. Such a tool would be revolutionary for cracking the function of layers and their component cell types. Consider the challenge of understanding how L4 affects other layers during sensory coding. Because L4 neurons target excitatory neurons in all other layers (including in L4 itself), the data obtained from simply suppressing L4 spiking does not allow one to disambiguate which of these synaptic connections was causally involved in the observed effect, whether it be on neural activity or behavior. If, instead, one could selectively silence just one of its output pathways (e.g., L4 to L2/3 or L4 to L5), then one could obtain a more detailed understanding of its role in translaminar computation. Some investigators have employed synaptic terminal silencing with inhibitory opsins; the efficacy of this approach appears to be spotty and is difficult to confirm, and it is not useful for cracking cortical circuits in any case; intracortical circuits are too spatially intermingled. Instead, what is needed is an approach that is conceptually analogous to the GFP reconstitution across synaptic partners ("GRASP") technology, which allows one to target a split fluorescent reporter to two genetically defined classes of neurons (Feinberg et al., 2008); if they contact one another, the fluorescent protein is reconstituted. Because splitting and reconstituting an integral membrane protein like an opsin seems implausible, alternative schemes must be developed. If this technical obstacle can be overcome, then a whole new range of critical optogenetic experiments could be performed to dissect the function of each intracortical pathway.

Conclusion

With the advent of new technologies for large-scale recording and manipulation of neural activity across cortical layers, coupled with the possibility of identifying all of the cardinal cell types that compose each layer, we are on the threshold of obtaining the new data that could finally address how cortical

layers differentially contribute to cortical computation. Mounting a holistic attack on cortical circuits with these new approaches, particularly in the appropriate behavioral paradigms, promises to provide fundamental new mechanistic insight into cortical circuit structure and function. Ultimately, the proximal goal should be to distinguish between the competing conceptual frameworks that attempt to explain how cortical circuits generate perception. Probing the synthetic model requires us to reveal new types of transformations of sensory coding across and between the layers of each cortical area within the sensory hierarchy, if and when they exist. Ultimately, it is necessary to provide a complete picture for how complex feature selectivity (such as for faces) is synthesized across this hierarchy and how this depends on the activity and the connectivity among the cell types at each stage. Probing the inferential model requires the direct observation of neural activity that encodes predictions, prediction errors, and the sensory or motor memories that collectively lead to the moment-to-moment comparison between a generative model of the world and the sensory data that arrive via the sense organs.

DECLARATION OF INTERESTS

H.A. has a patent related to this work: 3D Sparse Holographic Temporal Focusing, 2016, L. Waller, N. Pegard, and H. Adesnik, Provisional Patent Application #62-429,017.

REFERENCES

Adam, Y., Kim, J.J., Lou, S., Zhao, Y., Brinks, D., Wu, H., Mostajo-Radji, M.A., Kheifets, S., Parot, V., Chettih, S., et al. (2018). All-optical electrophysiology reveals brain-state dependent changes in hippocampal subthreshold dynamics and excitability. bioRxiv. https://doi.org/10.1101/281618.

Adesnik, H. (2018). Layer-specific excitation/inhibition balances during neuronal synchronization in the visual cortex. J. Physiol. 596, 1639–1657.

Adesnik, H., and Scanziani, M. (2010). Lateral competition for cortical space by layer-specific horizontal circuits. Nature 464, 1155–1160.

Akemann, W., Sasaki, M., Mutoh, H., Imamura, T., Honkura, N., and Knopfel, T. (2013). Two-photon voltage imaging using a genetically encoded voltage indicator. Sci. Rep. 3, 2231.

Alonso, J.M., and Martinez, L.M. (1998). Functional connectivity between simple cells and complex cells in cat striate cortex. Nat. Neurosci. 1, 395–403.

Andermann, M.L., Gilfoy, N.B., Goldey, G.J., Sachdev, R.N.S., Wölfel, M., McCormick, D.A., Reid, R.C., and Levene, M.J. (2013). Chronic cellular imaging of entire cortical columns in awake mice using microprisms. Neuron 80, 900–913.

Armstrong-James, M., Fox, K., and Das-Gupta, A. (1992). Flow of excitation within rat barrel cortex on striking a single vibrissa. J. Neurophysiol. 68, 1345–1358.

Baker, C.A., Elyada, Y.M., Parra, A., and Bolton, M.M. (2016). Cellular resolution circuit mapping with temporal-focused excitation of soma-targeted channelrhodopsin. eLife 5, e14193.

Barth, A.L., Gerkin, R.C., and Dean, K.L. (2004). Alteration of neuronal firing properties after in vivo experience in a FosGFP transgenic mouse. J. Neurosci. 24, 6466–6475.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., and Friston, K.J. (2012). Canonical microcircuits for predictive coding. Neuron 76, 695–711.

Binzegger, T., Douglas, R.J., and Martin, K.A. (2004). A quantitative map of the circuit of cat primary visual cortex. J. Neurosci. 24, 8441–8453.

Bolus, M.F., Willats, A.A., Whitmire, C.J., Rozell, C.J., and Stanley, G.B. (2018). Design strategies for dynamic closed-loop optogenetic neurocontrol in vivo. J. Neural Eng. 15, 026011.

Bolz, J., and Gilbert, C.D. (1986). Generation of end-inhibition in the visual cortex via interlaminar connections. Nature 320, 362–365.

Bortone, D.S., Olsen, S.R., and Scanziani, M. (2014). Translaminar inhibitory cells recruited by layer 6 corticothalamic neurons suppress visual cortex. Neuron 82, 474–485.

Bovetti, S., Moretti, C., Zucca, S., Dal Maschio, M., Bonifazi, P., and Fellin, T. (2017). Simultaneous high-speed imaging and optogenetic inhibition in the intact mouse brain. Sci. Rep. 7, 40041.

Brecht, M. (2017). The body model theory of somatosensory cortex. Neuron 94, 985–992.

Brecht, M., Roth, A., and Sakmann, B. (2003). Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. J. Physiol. 553, 243–265.

Briggs, F., and Callaway, E.M. (2001). Layer-specific input to distinct cell types in layer 6 of monkey primary visual cortex. J. Neurosci. 21, 3600–3608.

Bruce, C., Desimone, R., and Gross, C.G. (1981). Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. J. Neurophysiol. 46, 369–384.

Bullen, A., and Saggau, P. (1999). High-speed, random-access fluorescence microscopy: II. Fast quantitative measurements with voltage-sensitive dyes. Biophys. J. 76, 2272–2287.

Calabrese, A., and Woolley, S.M.N. (2015). Coding principles of the canonical cortical microcircuit in the avian brain. Proc. Natl. Acad. Sci. USA 112, 3517–3522.

Callaway, E.M. (1998). Local circuits in primary visual cortex of the macaque monkey. Annu. Rev. Neurosci. 21, 47–74.

Canales, A., Jia, X., Froriep, U.P., Koppes, R.A., Tringides, C.M., Selvidge, J., Lu, C., Hou, C., Wei, L., Fink, Y., and Anikeeva, P. (2015). Multifunctional fibers for simultaneous optical, electrical and chemical interrogation of neural circuits in vivo. Nat. Biotechnol. 33, 277–284.

Canepari, M., Bernus, O., and Zecevic, D. (2015). Preface. In Membrane Potential Imaging in the Nervous System and Heart, M. Canepari, O. Bernus, and D. Zecevic, eds. (Springer).

Chamberland, S., Yang, H.H., Pan, M.M., Evans, S.W., Guan, S., Chavarha, M., Yang, Y., Salesse, C., Wu, H., Wu, J.C., et al. (2017). Fast two-photon imaging of subcellular voltage dynamics in neuronal tissue with genetically encoded indicators. eLife 6, 6.

Chang, L., and Tsao, D.Y. (2017). The code for facial identity in the primate brain. Cell 169, 1013–1028.e14.

Chen, T.W., Wardill, T.J., Sun, Y., Pulver, S.R., Renninger, S.L., Baohan, A., Schreiter, E.R., Kerr, R.A., Orger, M.B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. Nature 499, 295–300.

Chen, K.H., Boettiger, A.N., Moffitt, J.R., Wang, S., and Zhuang, X. (2015). RNA imaging. Spatially resolved, highly multiplexed RNA profiling in single cells. Science 348, aaa6090.

Cheng, A., Gonçalves, J.T., Golshani, P., Arisaka, K., and Portera-Cailliau, C. (2011). Simultaneous two-photon calcium imaging at different depths with spatiotemporal multiplexing. Nat. Methods 8, 139–142.

Chung, J.E., Magland, J.F., Barnett, A.H., Tolosa, V.M., Tooker, A.C., Lee, K.Y., Shah, K.G., Felix, S.H., Frank, L.M., and Greengard, L.F. (2017). A fully automated approach to spike sorting. Neuron 95, 1381–1394.e6.

Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Foster, J.D., Nuyuju-kian, P., Ryu, S.I., and Shenoy, K.V. (2012). Neural population dynamics during reaching. Nature 487, 51–56.

Constantinople, C.M., and Bruno, R.M. (2013). Deep cortical layers are activated directly by thalamus. Science 340, 1591–1594.

Cruikshank, S.J., Urabe, H., Nurmikko, A.V., and Connors, B.W. (2010). Pathway-specific feedforward circuits between thalamus and neocortex revealed by selective optical stimulation of axons. Neuron 65, 230–245.

Csicsvari, J., Henze, D.A., Jamieson, B., Harris, K.D., Sirota, A., Barthó, P., Wise, K.D., and Buzsáki, G. (2003). Massively parallel recording of unit and local field potentials with silicon-based electrodes. J. Neurophysiol. 90, 1314–1323.

D'Souza, R.D., and Burkhalter, A. (2017). A laminar organization for selective cortico-cortical communication. Front. Neuroanat. 11, 71.

Deutsch, D., Pietr, M., Knutsen, P.M., Ahissar, E., and Schneidman, E. (2012). Fast feedback in active sensing: touch-induced changes to whisker-object interaction. PLoS ONE 7, e44272.

Dhawale, A.K., Poddar, R., Wolff, S.B.E., Normand, V.A., Kopelowitz, E., and Ölveczky, B.P. (2017). Automated long-term recording and analysis of neural activity in behaving animals. eLife 6, 6.

DiMattina, C., and Zhang, K. (2013). Adaptive stimulus optimization for sensory systems neuroscience. Front. Neural Circuits 7, 101.

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

Douglas, R.J., and Martin, K.A. (2004). Neuronal circuits of the neocortex. Annu. Rev. Neurosci. 27, 419–451.

Duemani Reddy, G., Kelleher, K., Fink, R., and Saggau, P. (2008). Three-dimensional random access multiphoton microscopy for functional imaging of neuronal activity. Nat. Neurosci. 11, 713–720.

Emiliani, V., Cohen, A.E., Deisseroth, K., and Häusser, M. (2015). All-optical interrogation of neural circuits. J. Neurosci. 35, 13917–13926.

Feinberg, E.H., Vanhoven, M.K., Bendesky, A., Wang, G., Fetter, R.D., Shen, K., and Bargmann, C.I. (2008). GFP Reconstitution Across Synaptic Partners (GRASP) defines cell contacts and synapses in living nervous systems. Neuron 57, 353–363.

Feldmeyer, D. (2012). Excitatory neuronal connectivity in the barrel cortex. Front. Neuroanat. 6, 24.

Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex 1, 1–47.

Fitzpatrick, D., Itoh, K., and Diamond, I.T. (1983). The laminar organization of the lateral geniculate body and the striate cortex in the squirrel monkey (Saimiri sciureus). J. Neurosci. 3, 673–702.

Friedrich, J., Zhou, P., and Paninski, L. (2017). Fast online deconvolution of calcium imaging data. PLoS Comput. Biol. 13, e1005423.

Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural manifolds for the control of movement. Neuron 94, 978–984.

Gămănuț, R., Kennedy, H., Toroczkai, Z., Ercsey-Ravasz, M., Van Essen, D.C., Knoblauch, K., and Burkhalter, A. (2018). The mouse cortical connectome, characterized by an ultra-dense cortical graph, maintains specificity by distinct connectivity profiles. Neuron 97, 698–715.e10.

Gilbert, C.D. (1983). Microcircuitry of the visual cortex. Annu. Rev. Neurosci. 6, 217–247.

Grosenick, L., Marshel, J.H., and Deisseroth, K. (2015). Closed-loop and activity-guided optogenetic control. Neuron 86, 106–139.

Gross, C.G., Rocha-Miranda, C.E., and Bender, D.B. (1972). Visual properties of neurons in inferotemporal cortex of the Macaque. J. Neurophysiol. 35, 96–111.

Guenthner, C.J., Miyamichi, K., Yang, H.H., Heller, H.C., and Luo, L. (2013). Permanent genetic access to transiently active neurons via TRAP: targeted recombination in active populations. Neuron 78, 773–784.

Guy, J., and Staiger, J.F. (2017). The functioning of a cortex without layers. Front. Neuroanat. 11, 54.

Han, Y., Kebschull, J.M., Campbell, R.A.A., Cowan, D., Imhof, F., Zador, A.M., and Mrsic-Flogel, T.D. (2018). The logic of single-cell projections from visual cortex. Nature 556, 51–56.

Harris, K.D., and Shepherd, G.M. (2015). The neocortical circuit: themes and variations. Nat. Neurosci. 18, 170–181.

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

Harris, K.D., Quiroga, R.Q., Freeman, J., and Smith, S.L. (2016). Improving data quality in neuronal population recordings. Nat. Neurosci. 19, 1165–1174.

Hernandez, O., Papagiakoumou, E., Tanese, D., Fidelin, K., Wyart, C., and Emiliani, V. (2016). Three-dimensional spatiotemporal focusing of holographic patterns. Nat. Commun. 7, 11928.

Hochbaum, D.R., Zhao, Y., Farhi, S.L., Klapoetke, N., Werley, C.A., Kapoor, V., Zou, P., Kralj, J.M., Maclaurin, D., Smedemark-Margulies, N., et al. (2014). All-optical electrophysiology in mammalian neurons using engineered microbial rhodopsins. Nat. Methods 11, 825–833.

Hoffman, K.P., and Stone, J. (1971). Conduction velocity of afferents to cat visual cortex: a correlation with cortical receptive field properties. Brain Res. 32, 460–466.

Homann, J., Koay, S.A., Glidden, A.M., Tank, D.W., and Berry, M.J. (2017). Predictive coding of novel versus familiar stimuli in the primary visual cortex. bioRxiv. https://doi.org/10.1101/197608.

Hubel, D.H., and Wiesel, T.N. (1959). Receptive fields of single neurones in the cat's striate cortex. J. Physiol. 148, 574–591.

Hubel, D.H., and Wiesel, T.N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. J. Physiol. 160, 106–154.

Ibrahim, L.A., Mesik, L., Ji, X.Y., Fang, Q., Li, H.F., Li, Y.T., Zingg, B., Zhang, L.I., and Tao, H.W. (2016). Cross-modality sharpening of visual cortical processing through layer-1-mediated inhibition and disinhibition. Neuron 89, 1031–1045.

Iurilli, G., Ghezzi, D., Olcese, U., Lassi, G., Nazzaro, C., Tonini, R., Tucci, V., Benfenati, F., and Medini, P. (2012). Sound-driven synaptic inhibition in primary visual cortex. Neuron 73, 814–828.

Jazayeri, M., and Afraz, A. (2017). Navigating the neural space in search of the neural code. Neuron 93, 1003–1014.

Juavinett, A.L., Erlich, J.C., and Churchland, A.K. (2018). Decision-making behaviors: weighing ethology, complexity, and sensorimotor compatibility. Curr. Opin. Neurobiol. 49, 42–50.

Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydın, Ç., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature 551, 232–236.

Jung, J.C., Mehta, A.D., Aksay, E., Stepnoski, R., and Schnitzer, M.J. (2004). In vivo mammalian brain imaging using one- and two-photon fluorescence microendoscopy. J. Neurophysiol. 92, 3121–3133.

Katona, G., Szalay, G., Maák, P., Kaszás, A., Veress, M., Hillier, D., Chiovini, B., Vizi, E.S., Roska, B., and Rózsa, B. (2012). Fast two-photon in vivo imaging with three-dimensional random-access scanning in large tissue volumes. Nat. Methods 9, 201–208.

Keller, G.B., Bonhoeffer, T., and Hübener, M. (2012). Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. Neuron 74, 809–815.

Kerlin, A.M., Andermann, M.L., Berezovskii, V.K., and Reid, R.C. (2010). Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. Neuron 67, 858–871.

Kim, J., Matney, C.J., Blankenship, A., Hestrin, S., and Brown, S.P. (2014). Layer 6 corticothalamic neurons activate a cortical output layer, layer 5a. J. Neurosci. 34, 9656–9664.

Kinnischtzke, A.K., Fanselow, E.E., and Simons, D.J. (2016). Target-specific M1 inputs to infragranular S1 pyramidal neurons. J. Neurophysiol. *116*, 1261–1274.

Krakauer, J.W., Ghazanfar, A.A., Gomez-Marin, A., MacIver, M.A., and Poeppel, D. (2017). Neuroscience needs behavior: correcting a reductionist bias. Neuron *93*, 480–490.

Kumar, A., Vlachos, I., Aertsen, A., and Boucsein, C. (2013). Challenges of understanding brain function by selective modulation of neuronal subpopulations. Trends Neurosci. *36*, 579–586.

Larkum, M.E., Petro, L.S., Sachdev, R.N.S., and Muckli, L. (2018). A perspective on cortical layering and layer-spanning neuronal elements. Front. Neuroanat. *12*, 56.

Lefort, S., Tomm, C., Floyd Sarria, J.C., and Petersen, C.C. (2009). The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. Neuron *61*, 301–316.

Lima, S.Q., Hromádka, T., Znamenskiy, P., and Zador, A.M. (2009). PINP: a new method of tagging neuronal populations for identification during in vivo electrophysiological recording. PLoS ONE *4*, e6099.

Liu, X., Ramirez, S., Pang, P.T., Puryear, C.B., Govindarajan, A., Deisseroth, K., and Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. Nature *484*, 381–385.

Liu, L., She, L., Chen, M., Liu, T., Lu, H.D.D., Dan, Y., and Poo, M.M. (2016). Spatial structure of neuronal receptive field in awake monkey secondary visual cortex (V2). Proc. Natl. Acad. Sci. USA *113*, 1913–1918.

Lo, L., and Anderson, D.J. (2011). A Cre-dependent, anterograde transsynaptic viral tracer for mapping output pathways of genetically marked neurons. Neuron *72*, 938–950.

Lu, R., Sun, W., Liang, Y., Kerlin, A., Bierfeld, J., Seelig, J.D., Wilson, D.E., Scholl, B., Mohar, B., Tanimoto, M., et al. (2017). Video-rate volumetric functional imaging of the brain at synaptic resolution. Nat. Neurosci. *20*, 620–628.

Lubeck, E., Coskun, A.F., Zhiyentayev, T., Ahmad, M., and Cai, L. (2014). Single-cell in situ RNA profiling by sequential hybridization. Nat. Methods *11*, 360–361.

Luo, L., Callaway, E.M., and Svoboda, K. (2018). Genetic dissection of neural circuits: a decade of progress. Neuron *98*, 256–281.

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

Malpeli, J.G. (1983). Activity of cells in area 17 of the cat in absence of input from layer a of lateral geniculate nucleus. J. Neurophysiol. *49*, 595–610.

Manita, S., Suzuki, T., Homma, C., Matsumoto, T., Odagawa, M., Yamada, K., Ota, K., Matsubara, C., Inutsuka, A., Sato, M., et al. (2015). A top-down cortical circuit for accurate sensory perception. Neuron *86*, 1304–1316.

Mardinly, A.R., Oldenburg, I.A., Pégard, N.C., Sridharan, S., Lyall, E.H., Chesnov, K., Brohawn, S.G., Waller, L., and Adesnik, H. (2018). Precise multimodal optical control of neural ensemble activity. Nat. Neurosci. *21*, 881–893.

Marr, D. (2010). Vision: A Computational Investigation Into The Human Representation And Processing of Visual Information (MIT Press).

Marshel, J.H., Mori, T., Nielsen, K.J., and Callaway, E.M. (2010). Targeting single neuronal networks for gene expression and cell labeling in vivo. Neuron *67*, 562–574.

Martinez, L.M., and Alonso, J.M. (2001). Construction of complex receptive fields in cat primary visual cortex. Neuron *32*, 515–525.

Martinez, L.M., Wang, Q., Reid, R.C., Pillai, C., Alonso, J.M., Sommer, F.T., and Hirsch, J.A. (2005). Receptive field structure varies with layer in the primary visual cortex. Nat. Neurosci. *8*, 372–379.

Matyas, F., Sreenivasan, V., Marbach, F., Wacongne, C., Barsy, B., Mateo, C., Aronoff, R., and Petersen, C.C.H. (2010). Motor control by sensory cortex. Science *330*, 1240–1243.

Meyer, H.S., Wimmer, V.C., Hemberger, M., Bruno, R.M., de Kock, C.P.J., Frick, A., Sakmann, B., and Helmstaedter, M. (2010). Cell type-specific thalamic innervation in a column of rat vibrissal cortex. Cereb. Cortex *20*, 2287–2303.

Michaels, J.A., Dann, B., and Scherberger, H. (2016). Neural population dynamics during reaching are better explained by a dynamical system than representational tuning. PLoS Comput. Biol. *12*, e1005175.

Mignard, M., and Malpeli, J.G. (1991). Paths of information flow through visual cortex. Science *251*, 1249–1251.

Miller, E.W. (2016). Small molecule fluorescent voltage indicators for studying membrane potential. Curr. Opin. Chem. Biol. *33*, 74–80.

Olsen, S.R., Bortone, D.S., Adesnik, H., and Scanziani, M. (2012). Gain control by layer six in cortical circuits of vision. Nature *483*, 47–52.

Otchy, T.M., Wolff, S.B.E., Rhee, J.Y., Pehlevan, C., Kawai, R., Kempf, A., Gobes, S.M.H., and Ölveczky, B.P. (2015). Acute off-target effects of neural circuit manipulations. Nature *528*, 358–363.

Ouzounov, D.G., Wang, T., Wang, M., Feng, D.D., Horton, N.G., Cruz-Hernández, J.C., Cheng, Y.T., Reimer, J., Tolias, A.S., Nishimura, N., and Xu, C. (2017). In vivo three-photon imaging of activity of GCaMP6-labeled neurons deep in intact mouse brain. Nat. Methods *14*, 388–390.

Packer, A.M., Peterka, D.S., Hirtz, J.J., Prakash, R., Deisseroth, K., and Yuste, R. (2012). Two-photon optogenetics of dendritic spines and neural circuits. Nat. Methods *9*, 1202–1205.

Packer, A.M., Russell, L.E., Dalgleish, H.W.P., and Häusser, M. (2015). Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. Nat. Methods *12*, 140–146.

Paluch-Siegler, S., Mayblum, T., Dana, H., Brosh, I., Gefen, I., and Shoham, S. (2015). All-optical bidirectional neural interfacing using hybrid multiphoton holographic optogenetic stimulation. Neurophotonics *2*, 031208.

Paninski, L., and Cunningham, J.P. (2018). Neural data science: accelerating the experiment-analysis-theory cycle in large-scale neuroscience. Curr. Opin. Neurobiol. *50*, 232–241.

Papagiakoumou, E., Anselmi, F., Bègue, A., de Sars, V., Glückstad, J., Isacoff, E.Y., and Emiliani, V. (2010). Scanless two-photon excitation of channelrhodopsin-2. Nat. Methods *7*, 848–854.

Pégard, N.C., Mardinly, A.R., Oldenburg, I.A., Sridharan, S., Waller, L., and Adesnik, H. (2017). Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT). Nat. Commun. *8*, 1228.

Perron, A., Mutoh, H., Akemann, W., Gautam, S.G., Dimitrov, D., Iwamoto, Y., and Knöpfel, T. (2009). Second and third generation voltage-sensitive fluorescent proteins for monitoring membrane potential. Front. Mol. Neurosci. *2*, 5.

Petreanu, L., Mao, T., Sternson, S.M., and Svoboda, K. (2009). The subcellular organization of neocortical excitatory connections. Nature *457*, 1142–1145.

Piatkevich, K.D., Jung, E.E., Straub, C., Linghu, C., Park, D., Suk, H.J., Hochbaum, D.R., Goodwin, D., Pnevmatikakis, E., Pak, N., et al. (2018). A robotic multidimensional directed evolution approach applied to fluorescent voltage reporters. Nat. Chem. Biol. *14*, 352–360.

Pluta, S., Naka, A., Veit, J., Telian, G., Yao, L., Hakim, R., Taylor, D., and Adesnik, H. (2015). A direct translaminar inhibitory circuit tunes cortical output. Nat. Neurosci. *18*, 1631–1640.

Pnevmatikakis, E.A., Soudry, D., Gao, Y., Machado, T.A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. Neuron *89*, 285–299.

Pouget, A., Dayan, P., and Zemel, R. (2000). Information processing with population codes. Nat. Rev. Neurosci. *1*, 125–132.

Prakash, R., Yizhar, O., Grewe, B., Ramakrishnan, C., Wang, N., Goshen, I., Packer, A.M., Peterka, D.S., Yuste, R., Schnitzer, M.J., and Deisseroth, K. (2012). Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. Nat. Methods *9*, 1171–1179.

Prevedel, R., Verhoef, A.J., Pernía-Andrade, A.J., Weisenburger, S., Huang, B.S., Nöbauer, T., Fernández, A., Delcour, J.E., Golshani, P., Baltuska, A., and Vaziri, A. (2016). Fast volumetric calcium imaging across multiple cortical layers using sculpted light. Nat. Methods *13*, 1021–1028.

Rancz, E.A., Franks, K.M., Schwarz, M.K., Pichler, B., Schaefer, A.T., and Margrie, T.W. (2011). Transfection via whole-cell recording in vivo: bridging single-cell physiology, genetics and connectomics. Nat. Neurosci. *14*, 527–532.

Rao, R.P.N., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nat. Neurosci. *2*, 79–87.

Reijmers, L.G., Perkins, B.L., Matsuo, N., and Mayford, M. (2007). Localization of a stable neural correlate of associative memory. Science *317*, 1230–1233.

Rickgauer, J.P., Deisseroth, K., and Tank, D.W. (2014). Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. Nat. Neurosci. *17*, 1816–1824.

Roelfsema, P.R., Lamme, V.A.F., and Spekreijse, H. (1998). Object-based attention in the primary visual cortex of the macaque monkey. Nature *395*, 376–381.

Russo, A.A., Bittner, S.R., Perkins, S.M., Seely, J.S., London, B.M., Lara, A.H., Miri, A., Marshall, N.J., Kohn, A., Jessell, T.M., et al. (2018). Motor cortex embeds muscle-like commands in an untangled population response. Neuron *97*, 953–966.e8.

Sakurai, K., Zhao, S., Takatoh, J., Rodriguez, E., Lu, J., Leavitt, A.D., Fu, M., Han, B.X., and Wang, F. (2016). Capturing and manipulating activated neuronal ensembles with CANE delineates a hypothalamic social-fear circuit. Neuron *92*, 739–753.

Saleem, A.B., Ayaz, A., Jeffery, K.J., Harris, K.D., and Carandini, M. (2013). Integration of visual motion and locomotion in mouse visual cortex. Nat. Neurosci. *16*, 1864–1869.

Salinas, E., and Sejnowski, T.J. (2001). Correlated neuronal activity and the flow of neural information. Nat. Rev. Neurosci. *2*, 539–550.

Scharf, R., Tsunematsu, T., McAlinden, N., Dawson, M.D., Sakata, S., and Mathieson, K. (2016). Depth-specific optogenetic control in vivo with a scalable, high-density μLED neural probe. Sci. Rep. *6*, 28381.

Schindler, S.E., McCall, J.G., Yan, P., Hyrc, K.L., Li, M.J., Tucker, C.L., Lee, J.M., Bruchas, M.R., and Diamond, M.I. (2015). Photo-activatable Cre recombinase regulates gene expression in vivo. Sci. Rep. *5*, 13627.

Seely, J.S., Kaufman, M.T., Ryu, S.I., Shenoy, K.V., Cunningham, J.P., and Churchland, M.M. (2016). Tensor analysis reveals distinct population structure that parallels the different computational roles of areas M1 and V1. PLoS Comput. Biol. *12*, e1005164.

Shah, S., Lubeck, E., Zhou, W., and Cai, L. (2016). In situ transcription profiling of single cells reveals spatial organization of cells in the mouse hippocampus. Neuron *92*, 342–357.

Shobe, J.L., Claar, L.D., Parhami, S., Bakhurin, K.I., and Masmanidis, S.C. (2015). Brain activity mapping at multiple scales with silicon microprobes containing 1,024 electrodes. J. Neurophysiol. *114*, 2043–2052.

Sofroniew, N.J., Flickinger, D., King, J., and Svoboda, K. (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. eLife *5*, e14472.

Stirman, J.N., Smith, I.T., Kudenov, M.W., and Smith, S.L. (2016). Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the mammalian brain. Nat. Biotechnol. *34*, 857–862.

Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). In vivo two-photon calcium imaging of neuronal networks. Proc. Natl. Acad. Sci. USA *100*, 7319–7324.

Striedter, G.F. (2005). Principles of Brain Evolution (Sinauer Associates).

Sun, W., Tan, Z., Mensh, B.D., and Ji, N. (2016). Thalamus provides layer 4 of primary visual cortex with orientation- and direction-tuned inputs. Nat. Neurosci. *19*, 308–315.

Tanese, D., Weng, J.Y., Zampini, V., De Sars, V., Canepari, M., Rozsa, B., Emiliani, V., and Zecevic, D. (2017). Imaging membrane potential changes from dendritic spines using computer-generated holography. Neurophotonics *4*, 031211.

Tasic, B., Menon, V., Nguyen, T.N., Kim, T.K., Jarsky, T., Yao, Z., Levi, B., Gray, L.T., Sorensen, S.A., Dolbeare, T., et al. (2016). Adult mouse cortical cell taxonomy revealed by single cell transcriptomics. Nat. Neurosci. *19*, 335–346.

Thériault, G., Cottet, M., Castonguay, A., McCarthy, N., and De Koninck, Y. (2014). Extended two-photon microscopy in live samples with Bessel beams: steadier focus, faster volume scans, and simpler stereoscopic imaging. Front. Cell. Neurosci. *8*, 139.

Thomson, A.M., and Lamy, C. (2007). Functional maps of neocortical local circuitry. Front. Neurosci. *1*, 19–42.

Thorpe, S., Fize, D., and Marlot, C. (1996). Speed of processing in the human visual system. Nature *381*, 520–522.

Tsai, P.S., Mateo, C., Field, J.J., Schaffer, C.B., Anderson, M.E., and Kleinfeld, D. (2015). Ultra-large field-of-view two-photon microscopy. Opt. Express *23*, 13833–13847.

Vinck, M., Batista-Brito, R., Knoblich, U., and Cardin, J.A. (2015). Arousal and locomotion make distinct contributions to cortical activity patterns and visual encoding. Neuron *86*, 740–754.

Wertz, A., Trenholm, S., Yonehara, K., Hillier, D., Raics, Z., Leinweber, M., Szalay, G., Ghanem, A., Keller, G., Rózsa, B., et al. (2015). PRESYNAPTIC NETWORKS. Single-cell-initiated monosynaptic tracing reveals layer-specific cortical network modules. Science *349*, 70–74.

Wickersham, I.R., Finke, S., Conzelmann, K.K., and Callaway, E.M. (2007). Retrograde neuronal tracing with a deletion-mutant rabies virus. Nat. Methods *4*, 47–49.

Wohrer, A., Humphries, M.D., and Machens, C.K. (2013). Population-wide distributions of neural activity during perceptual decision-making. Prog. Neurobiol. *103*, 156–193.

Wright, N., and Fox, K. (2010). Origins of cortical layer V surround receptive fields in the rat barrel cortex. J. Neurophysiol. *103*, 709–724.

Wu, F., Stark, E., Im, M., Cho, I.J., Yoon, E.S., Buzsáki, G., Wise, K.D., and Yoon, E. (2013). An implantable neural probe with monolithically integrated dielectric waveguide and recording electrodes for optogenetics applications. J. Neural Eng. *10*, 056012.

Xu, Y., Zou, P., and Cohen, A.E. (2017). Voltage imaging with genetically encoded indicators. Curr. Opin. Chem. Biol. *39*, 1–10.

Yamawaki, N., and Shepherd, G.M.G. (2015). Synaptic circuit organization of motor corticothalamic neurons. J. Neurosci. *35*, 2293–2307.

Yang, W., Miller, J.E.K., Carrillo-Reid, L., Pnevmatikakis, E., Paninski, L., Yuste, R., and Peterka, D.S. (2016). Simultaneous multi-plane imaging of neural circuits. Neuron *89*, 269–284.

Yang, W., Carrillo-Reid, L., Bando, Y., Peterka, D.S., and Yuste, R. (2018). Simultaneous two-photon imaging and two-photon optogenetics of cortical circuits in three dimensions. eLife *7*, e32671.

Zeng, W.B., Jiang, H.F., Gang, Y.D., Song, Y.G., Shen, Z.Z., Yang, H., Dong, X., Tian, Y.L., Ni, R.J., Liu, Y., et al. (2017). Anterograde monosynaptic transneuronal tracers derived from herpes simplex virus 1 strain H129. Mol. Neurodegener. *12*, 38.

Ziemba, C.M., Freeman, J., Movshon, J.A., and Simoncelli, E.P. (2016). Selectivity and tolerance for visual texture in macaque V2. Proc. Natl. Acad. Sci. USA *113*, E3140–E3149.

Zingg, B., Chou, X.L., Zhang, Z.G., Mesik, L., Liang, F., Tao, H.W., and Zhang, L.I. (2017). AAV-mediated anterograde transsynaptic tagging: mapping corticocollicular input-defined neural pathways for defense behaviors. Neuron *93*, 33–47.