Contextual processing in mouse visual cortex

Jorge Aurelio Menéndez

CoMPLEX Summer Project Thesis

Supervisors: Matteo Carandini & Marius Pachitariu

15th August, 2016

ABSTRACT

A classical finding of visual cortical responses in anesthetized cat is that they are sensitive to the temporal context in which a stimulus is presented. Here, we investigate the classic tilt after-effect in awake mice, using 2-photon calcium imaging to simultaneously record the responses of 1000+ neurons in transgenic mice with labeled paravalbumin- (PV) or somatostatin- (SOM) expressing interneurons. We found that both pyramidal and PV cells showed highly diverse responses, with a minority of cells in both populations showing stable orientation tuning. Very few PV cells had high orientation tuning, but SOM cells seemed to be equally orientation tuned as pyramidal cells. We observed classical adaptation effects in putative pyramidal and SOM cells, with reduced responses of cells with preferred orientations/directions near that of the adaptor. PV cells, on the other hand, did not seem to undergo any adapation effects. Notably, however, we did not observe previously reported repulsive tuning curve shifts during an adaptation-inducing stimulus, in any of our populations. We conclude that adaptation effects in mouse V1 likely arise from single-cell adaptation, and propose future modeling work that may clarify whether other cortical contributions may play a role.

Contents

1. Introduction

One of the most fascinating and eminent challenges of modern neuroscience is understanding the connection between brain and behavior. From visual perception to social exchanges, mammals demand that their nervous system accomplish astonishing computational feats on a second-by-second basis. A thorough understanding of such nervous systems and the behaviors they produce thus necessitates understanding these computations and how the underlying biological processes give rise to them. [10, 46]

A practical starting point in this venture is sensory perception, where both ends of the brain-behavior spectrum can be observed. From the seminal work of Hubel & Weisel a half-century ago [32], neuroscientists have had much success mapping cortical activity to perceptual phenomena. Furthermore, the abstract mathematical problem that perception poses provides a framework through which to understand such mappings in terms of computation. [10, 11, 47]

An illustrative example of this approach is the study of the tilt after-effect, whereby being exposed to a visual stimulus of a particular orientation for an extended period of time induces a perceptual bias in subsequently presented orientations. [63, 27, 38, 16] The resulting biases are systematic and can be interpreted as serving a functional computational role. [63, 16] Such *adaptation* to the temporal (or spatial [63]) context in which a stimulus is encountered can also be observed in the brain, where visual cortical neurons show corresponding systematic changes in their response properties. [48, 39, 12, 19, 9, 62, 17] Hence, adaptation provides an elegant example where we can neatly map changes in neural activity to changes in behavior (in this case, perception). [38, 63] In a certain sense, it allows to *observe* neural computation.

Here, we present a series of findings from recording visual cortical neurons responding to such oriented stimuli. By presenting the stimuli in different temporal contexts, we can ask how neurons change their responses under different settings and how these changes may underlie functional computations.

Importantly, it is clear that the mammalian cortex employs a diverse arsenal of cells that may each play distinct computational roles. [29, 30] In understanding the biological substrate of cortical computation it will thus be essential to disentangle the response properties of and interactions between these different types of neurons. One particularly crucial distinction in this respect is that of inhibitory (GABAergic) and excitatory (glutamatergic) neurons. [33] GABAergic cells have drastically different firing dynamics [42, 41] as well as differential connectivity patterns than those of typical excitatory pyramidal cells. [31, 55, 23, 22, 56] These differences must in some sense underlie differential functional roles played by these two broad classes of cortical cells. [33, 31, 30, 29]

In this vein, we used 2-photon calcium imaging with transgenic mice to record from many cells simultaneously and selectively record parvalbumin-positive and somatostatin-positive interneurons. With this technique, we were able to observe

many different kinds of activity in these different classes of cortical cells. What follows is mainly a descriptive study of how these cells respond to stimuli in different temporal contexts. This lays the groundwork for future analysis to dig into the mechanisms giving rise to the neural responses intrinsic to contextual processing in visual cortex.[1]

I start by reviewing the literature surrounding orientation tuning and the influence of context on visual cortical neuron responses, followed by a discussion of current hypotheses about the computational roles of parvalbumin and somatostatin interneurons in visual cortex. I then outline the approach taken here before reporting the method details and results.

1.1. **Contextual processing in visual cortex.** The crux of the challenge faced by any visual system is that vision is inherently an *inverse* problem: the sensory data available to an organism severely underdetermines what is actually out there in the world. [47] It is because of this that context is critical to perceptual processing. The context an object is in can provide crucial information for inferring what that object is or what it may be doing, particularly because the space of possible contexts encountered on planet earth turns out to have a significant amount of structure. [63, 21] Intuitively, it would be thus be advantageous for a perceptual system's functioning to be sensitive to changes in context. [63, 38, 16]

Indeed, this strategy seems to be prominent in mammalian sensory cortices. [63, 39, 48, 50, 25, 45, 17] In visual cortex, in particular, one domain where such adaptation effects are seen robustly is in neural responses to orientation. Orientation tuning is a robust phenomenon in visual cortex whereby neurons preferentially respond to orientations within a certain range. [32] However, the tuning of these neurons (i.e. the orientations they prefer) is strikingly modulated by the spatial and temporal context in which oriented stimuli are encountered. [63]

For example, in a classic finding termed the tilt after-effect, neural responses to an oriented stimulus are modulated in systematic ways when a stimulus is shown in a temporal context favoring one orientation. A cell's orientation tuning can be measured precisely by presenting an animal with a sequence of different orientation and measuring the cell's responses to each one. This will yield a tuning curve reflecting the orientations that cell responds to most strongly. But when the animal is presented repeatedly with stimuli of one single orientation (e.g. by showing a grating of a particular orientation for an extended period of time spanning minutes [19], or by showing a sequence of oriented gratings where one orientation is significantly more frequent than any other [9]), that cell's tuning begins to change. Primarily, two effects are observed: [63, 38, 19, 9, 34]

(1) Responses to the "adaptor" orientation are strongly reduced. If the animal is presented a 90° oriented stimulus repeatedly or for a duration of

[1]At least in the temporal domain, although there are reasons to think that results would generalize to spatial contextual processing as well [63]

many seconds, orientation tuned neurons in visual cortex will subsequently respond much more weakly to 90° oriented stimuli, as well as to other orientations within $\sim 20°$. Responses to orientations more than 30°+ away from 90°, however, will be largely unaffected.

(2) The cell's preferred orientation shifts away from the adaptor orientation. Again, this effect is strongest for cells with preferred orientations near the adaptor orientation. So if a cell prefers 95° oriented stimuli when they are presented in isolation, it will respond preferentially to 100° orientations after being stimulated by a 90° stimulus over an extended period of time. However, a cell that prefers 175° will maintain its tuning.

Such changes in response properties arising from changing the context in which a stimulus is presented are termed *adaptation* effects.

In fact, adaptation effects can be observed across processing stages in the visual system hierarchy, from cortical responses to spatial frequency [48], direction of motion [39], spatial location [17], and contrast [48, 12], down to the lateral geniculate nucleus cells [17] and retinal ganglion cells. [4] The effects of adaptation can be observed at the level of behavior as well, where preceptual biases arise after prolonged exposure to a particular stimulus like an orientation or even a face. [27, 16] Similar effects can be found in the auditory domain as well, where repeated presentation to the same frequency leads to changes in the frequency tuning of individual cells and a similiar stimulus-specific attenuation of responses. [50] These adaptation effects found across sensory cortices suggest that it might be a general strategy employed by mammalian perceptual systems, or even by the mammalian cortex in general.

Why would such perceptual biases and response changes be useful? Indeed, a change in the orientation tuning of cell should lead to the so-called *coding catastrophe*, where a population code of orientation is suddenly changed from one moment to the next, leading to downstream decoding errors and incorrect inferences. [63] However, it turns out that if one assumes classical population decoding schemes (e.g. winner-take-all, maximum-likelihood, population vector), this problem is in fact dealt with via the interplay of the the two effects mentioned above: whereas the reduction of responses to stimuli near the adaptor amplifies the decoding error, the repulsive shifts in preferred orientation ameliorate it. [34] A plausible hypothesis is thus that biological constraints impose that neurons can't respond at the same strength when repeatedly stimulated, leading to a kind of neural fatigue enhanced in the cells than usually repsond most strongly to the repeatedly presented adaptor stimulus. Indeed such effects are in fact seen *in vitro*, where stimulation of individual neurons over a long period of time leads to an attenuation in spiking responses. [43] These *spike-frequency adaptation* effects seem to be a consequence of the cellular mechanisms inherent to neural spiking, explained by the dynamics of membrane ion channels. [43, 59] Thus, it could be that visual cortex might is wired up in such a way that repulsive tuning curve shifts occur under certain temporal

contexts that would otherwise lead to the coding catastrophe via spike-frequency adaptation. [68]

Alternatively, these adaptation effects may be desirable despite the coding catastrophe. Selectively reducing responses to stimuli appearing frequently in the sensory environment would lead to improved deviance detection, which may be advantageous for survival even under erroneous perceptual biases. [50, 28, 63, 38] Relatedly, it would improve coding efficiency by reducing redundant signals and decorrelating neural responses. [63, 38, 9] Adapting to the statistics of the sensory context may also help maintain stimulus discriminability across different environments while staying within a healthy dynamic firing range [38], although the psychophysical evidence for this is dubious. [63]

In summary, it is clear that sensory adaptation is ubiquitous throughout sensory cortices. It is unclear, however, whether it is a property brains are designed to have or whether it is simply a byproduct of biological constraints that the brain must deal with. Elucidating the mechanisms by which it arises will be crucial for disentangling these two explanations. Here, we focus on orientation adaptation because adaptation effects in orientation tuned V1 cells are well studied and robust. Furthermore they provide a probe for understanding cortical computation, since they seem to arise from cortical activity: the sharp orientation tuning obseved in V1 seems to necessitate recurrent cortical input to begin with [65, 6, 13], and lateral geniculate nuclues (LGN) neurons are very weakly orientation tuned [17], meaning that they could not underlie the stimulus-specific adaptation effects observed (i.e. the fact that cells tuned to orientations near the adaptor's are more affected than those tuned to orientations orthogonal to the adaptor). Importantly, however, all previous experiments mentioned here were performed on anesthetized cats. It remains an open question whether they will be replicated in awake mice, whose visual cortex is differently organized. [29]

Lastly, inhibitory neurons have also been shown to be somewhat orientation selective [52, 67, 61, 3, 36, 14, 42] and seem to play an important role in establishing and sharpening the orientation tuning of pyramidal cells [65, 70, 44, 35, 6, 2]. These observations pose the question of whether they play a role in driving cortical adaptation effects, or whether they undergo adaptation effects themselves. [68, 50] Thus, orientation tuning and tilt after-effects provide a paradigm well-suited to probing the cortical circuitry underlying contextual processing in V1.

1.2. **Cortical inhibition in visual cortex.** One of the questions currently at the forefront of research on mammalian neocortex is the question of what computational role is played by GABAergic inhibitory neurons. In general, inhibitory neurons have very distinct firing and wiring patters that clearly distinguish them from cortical pyramidal cells [30, 22, 29], and even within the class of GABAergic interneurons, many subtypes exist with different properties. A full taxonomy remains a matter of debate, but it is largely agreed upon that three main classes can be distinguished by their gene expresssion patterns, response properties, and

connectivity: parvalbumin expressing (PV), somatostatin expressing (SOM), and vasoactive intestinal peptide expressing interneurons (VIP). [56, 71] VIP interneurons seem to provide very little inhibition to pyramidal cells [56], so we focus on PV and SOM interneurons.

Parvalbumin-expressing interneurons seem to be the most abundant in visual cortex, as well as the most responsive to visual stimulation [36] and the ones that evoke the strongest inhibitory post-synaptic conductances in pyramidal cells. [56] They synapse onto perisomatic regions of pyramidal cells [22, 18], with highly dense connections to and from the local population (input and output connectivity rates of $\sim 80\%$). [31, 55] In V1, excitatory inputs to PV cells tend to be unselective and local [31], leading to low orientation selectivity and a tuning bias toward that of the surrounding local population. [36] These observations seem to suggest that PV cell activity simply reflects the average activity of pyramidal cells in the local vicinity. [36, 31]

A hypothesis that arises out of these findings is that the computational role of PV cells may be to control the response gain of pyramidal cells. [3] Indeed, optogenetic activation or suppression of PV cells in V1 leads to changes in pyramidal cell response gain while leaving orientation tuning and contrast gain intact [3, 70] (although see [40] for a contradictory result[2]). Similarly, pharmacological blocking of GABA receptors increases pyramidal cell response gain but not contrast gain. [35] The fact that PV cells synapse onto areas near or on the somata of pyramidal cells makes them particularly well-suited to regulating pyramidal cell output. They also show short response latencies similar to pyramidal cells [41, 42] and seem to be particularly efficient at transforming synaptic input into spikes [41], putting them in an ideal position to coactivate with and thus efficiently regulate pyramidal cell responses to cortical or thalamic input. Indeed, PV cells receive synaptic input from thalamus [58] and respond to visual stimulation on the same timescale as pyramidal cells, allowing them to drive feed-forward inhibition. [42] Controlling pyramidal cell response gain in this way may be crucial for allowing them to operate within a reasonable dynamic range in the face of massive excitatory synaptic input. [64]

On the other hand, somatostatin-expressing interneurons have complementary properties that seem to make them better suited for modulating the input to pyramidal cells. [70, 26] They target pyramidal cell dendrites rather than cell bodies [22, 18], allowing them to directly interfere with the driving ability of specific synaptic inputs to a cell. Rather than providing strong and fast inhibition like PV cells, SOM cells tend to show much longer response latencies [41, 42] and produce inhibitory post-synaptic conductances half as strong as those produced by PV cells. [56] They also receive significantly weaker and delayed excitatory drive. [41, 42] Together, these observations suggest that SOM interneurons implement

[2]Note, however, that the discrepancy may be explained by the fact that the authors in [40] used 4 seconds of laser stimulation, which could have led to confounding network effects [20]

weak modulatory feedback inhibition in cortex. Indeed, the inhibition stemming from SOM cells seems to modulate the excitatory drive to pyramidal cells, affecting orientation tuning [70, 44], contrast sensitivity [70], and distal dendritic activity in L1 [26]. Interestingly, somatostatin-expressing interneurons seem to be orientation tuned [42] (although see [36]), suggesting they receive highly specific synaptic input from pyramidal cells.

What role could each of these play in contextual processing? The short-term dynamics of the synapses between pyramidal cells and each of these types of interneuron suggest they may each play fundamentally different roles. Specifically, PC-SOM synapses are typically facilitating [5, 66, 60] whereas PC-PV synapses are depressing [5, 60, 24]. Indeed, in auditory cortex, optogenetic suppression of PV vs SOM cells leads to differential changes in pyramidal cell responses to deviant and regular tones. [50] SOM-mediated inhibition seems to selectively suppress responses to repeated stimuli, whereas PV inhibition seems to affect responses to all stimuli unselectively. Again, this is reminiscent of the hypothesis that PV cells implement unselective response gain control whereas SOM cells implement selective feedback modulation of dendritic input. Other studies have also pointed out the importance of SOM cell activity for deviance detection in mice. [28]

In fact, both PV and SOM interneurons seem to undergo adaptation effects themselves in anesthetized mouse A1. [50] However, anaesthesia seems to strongly modulate SOM cell responses [1], so it remains an open question whether this is the case in awake mouse V1. Indeed, as reported below, we found no adaptation effects in PV and SOM responses to oriented gratings.

1.3. **Measuring neural responses.** To be able to probe the neural circuitry underlying contextual processing in visual cortex, we need a way to (1) record many neurons simultanesouly and (2) identify different types of neurons, *in vivo*. To do so, we use 2-photon calcium imaging, a technique that allows the recording of action potentials of many neurons simultaneously. [36, 31, 15, 37] Importantly, we take advantage of the fact that different genetically expressed fluorescent proteins can be imaged at different wavelengths, allowing us to use transgenic mice to label different classes of neurons. We can then selectively record PV or SOM interneurons along with the rest of the cells in their surrounding local population. This paradigm also provides the huge bonus of allowing us to image *awake* mice by headfixing them to the microscope after surgically implanting a transparent cover window over visual cortex.

However, like so many things in life, the huge advantages provided by 2-photon calcium imaging don't come without their costs. The raw data obtained from the microscope consists of an array of pixels with intensity values that vary over time. Determining which pixels belong to cells, which don't, and which belong to the same cell is by no means a trivial problem, and even pixels that we can be sure belong to a certain cell will inevitably have signals contaminated by irrelevant neural activity in the surrounding tissue (so-called "neuropil" contamination). Such

issues are critical to take into account when trying to infer neural activity from raw calcium imaging data. [54, 15, 57] We are not interested in how fluorescence intensity values measured by the microscope change with different stimulus presentations - we are interested in how a given neuron's pattern of firing changes. As such, need a way to accurately infer cell spiking activity from the recorded calcium signal. [54, 15, 57]

In addressing all of these challenges, we utilized a recently developed software for pre-processing 2-photon calcium imaging data, called *Suite2P*. [54] This software performs the whole pipeline, from cell detection to neuropil subtraction to spike deconvolution, and outperforms all current algorithms on ground truth data (http://neurofinder.codeneuro.org/). The result is a spike train for each neuron in the imaged field-of-view, which we can then analyze with respect to the stimuli presented to the mouse during recording. In the following methods section I describe in detail the processing steps implemented by *Suite2P*.

Using two different transgenic mouse lines, we imaged awake mice with labelled parvalbumin or somatostatin neurons while they were presented oriented gratings. The sequence of gratings were desigend to induce adaptation effects, which we measured and characterized in putative pyramidal cells, parvalbumin cells, and somatostatin cells.

2. Methods

2.1. **Mice & Imaging.** Experiment 1 was performed on three PV-Cre×tdTomato transgenic mice, imaging at 30Hz in L2/3. Experiment 2 was performed on a single SOM-Cre×tdTomato mouse, imaging at 3Hz across 10 planes spanning L2/3 and L4. These particular transgenic lines are such that parvalbumin- (in PV-Cre) or somatostatin- (in SOM-Cre) positive cells emit red light when stimulated with a laser at the correct frequency. Calcium activity was recorded by injecting an AAV1 virus at the field of view that led to all neurons near the injection site to express GCaMP6m/f, a protein that emits green light when stimulated by a laser whilst in the presence of calcium. Importantly, this particular calcium indicator has fast enough dynamics to detect spiking activity. [15] This allowed us to later infer spike times from the resulting recorded fluorescence signal.

An experimental session consisted of several blocks. First, retinotopic mapping was performed to find the receptive fields of the cells in the field of view. This was necessary to then be able to present the stimulus at the optimal location to stimulate the cells being imaged. Then, imaging was performed with the laser at a wavelength that would maximally activate the red fluoorescent protein in the PV+ or SOM+ cells to be able to later identify which recorded cells were PV/SOM and which were not (termed putative pyramidal cells). A different photodiode was used for this, with a filter to ensure only red light was being detected. The rest of the blocks consisted of presenting the head-fixed mouse with a series of visual stimuli while recording with the laser at a wavelength chosen to elicit emission from the

calcium indicator GCaMP. These provided us with fluorescence timecourses for each pixel in the field of view, which we then grouped into cells (see below) that we later classified as PV/SOM or putative pyramidal based on the red images.

In all experiments, mice were surgically implanted with a cover window over visual cortex that allowed us to image visual cortical neurons without exposing the brain. By head-fixing the mice during stimulus presentation, we were able to perform all experiments *in vivo* without anaesthesia. This is a crucial point, as the majority of previous studies on adaptation and orientation selectivity have been performed on anesthetized animals [36, 61, 20, 35, 2], and stimulus-evoked activity may not be the same in an awake animal as in an anaesthetized one (although see [53]), particularly that of somatostatin cells. [1]

2.2. **Stimuli.** In experiment 1, stimuli consisted of circular patches of oriented bars (circular gratings). They were presented at the optimal receptive field for the cells in the field of view. Stimuli were shown under three different conditions. In the 1Hz condition, the experimental block consisted of showing a sequence of gratings at a frequency of 1Hz for a duration of 100ms each. The 10Hz condition accordingly consisted of gratings presented at a frequency of 10Hz, also for a duration of 100ms each. In both of these cases, 12 different orientations ($0°, 15°, 30°, 45°, 60°, 75°, 90°, 105°, 120°, 135°, 150°, 165°$) were shown, presented in a random order such that any orientation was equally likely to succeed any other orienation. In addition, blank stimuli were interleaved randomly between oriented gratings. The third condition was the biased stimulus condition, which was just like the 10Hz stimulus except that one orientation ($90°$) was more frequent than any other. This type of stimulus provides the advantage of allowing us to stmultaneously induce and probe the effects of adaptation. [9] To ensure that all other aspects were kept constant relative to the unbiased 10Hz stimulus sequence, we equated the number of presentations of all other orientations, reducing the number of blank trials to increase the number of $90°$ stimulus presentations while keeping the total experiment length constant.

In experiment 2, oriented bars were presented drifting across the screen at different angles. The sequence of simuli consisted of a 30 second adaptation stimulus followed be a sequence of 2 second alternating "top-up" and test stimuli, separated by 1 second inter-stimulus blank intervals. In the control condition, the adaptation stimulus and top-up stimuli consisted of blanks. In the adaptation condition, these consisted of oriented bars drifting at $180°$ in both the 30s adaptor period and in each of the top-ups.

Stimuli were presented using Psychtoolbox, using a photodiode to precisely detect stimulus onset times relative to the calcium recording.

2.3. **Imaging data analysis.**

2.3.1. *Cell detection and neuropil subtraction.* As mentioned above, the raw recorded imaging data consists of a timecourse of fluorescence intensity for each pixel in the

imaged field of view. However, what we seek is *cell* spiking activity. To obtain this, we modeled each pixel's fluorescence activity as a series calcium transients arising from the spiking activity of a cell the pixel is a part of (if it is part of one). By fitting this model to the measured pixel intensity values, we were able to infer the groups of pixels that comprised single cells, along with those cells' spike trains.

Importantly, however, it is well known that the point spread function of the 2-photon microscope is elongated in the z axis by an order of magnitude more than in the x or y axes, meaning that the measured fluorescence from even a cell perfectly in focus will inevitably contain signal from adjacent somata or dendrites and axons behind and in front of it. [54, 15, 36] Such sources of fluorescence activity irrelevant to the actual spiking of the cell are termed *neuropil*, and are crucial to account for in inferring individual cell activity from the fluorescence signal. To do this, we incorporate into our generative model of pixel fluorescence instensities a term for neuropil contamination that models the spatial structure of neuropil activity. Previous studies have shown that this activity is correlated in space [54], so the model is designed such that the neuropil signal is forced to vary smoothly over space. Specifically, we model the neuropil signal across the image as a series of raised cosine functions uniformly tiling the field of view. By fitting weights and timecourses to each of these components, we obtain the neuropil signal at each pixel, enforced to vary smoothly between neighboring pixels. [54]

The generative model of the raw recorded signal r_k of the kth pixel is thus expressed as follows:

$$r_k = \sum_n \Lambda_{kn}\mathbf{f}_n + \alpha_k \sum_i B_{ki}\mathbf{n}_i + \eta$$

where

- Λ is a sparse matrix designating pixel membership to individual cells in the image (ROIs). $\Lambda_{kn} > 0$ if pixel k is part of the nth cell and 0 otherwise.
- $\mathbf{f}_n$ is the fluorescence signal reflecting the calcium activity arising from action potentials of the nth cell.
- $\{B_i\}$ are the set of basis functions modelling the spatial structure of the neuropil signal, each B_i consisting of a two-dimensional raised cosine. Each of these components has an associated neuropil signal timecourse $\mathbf{n}_i$. We thus denote the resulting contribution of the ith neuropil component to the kth pixel as B_{ki}, which decays as a function of the pixel's distance from the spatial center of B_i (the peak of the cosine).
- α_k is a scaling constant that estimates the level of neuropil contamination in each pixel. This is important because the amount of neuropil contamination in any given part of the image will depend on whether there is a cell there and whether it is in focus. The presence of a cell will displace the neuropil along the z-axis, reducing its contribution. If, however, that cell is out of focus, then the neuropil contribution will be greater. On the other hand, if

there is no cell there at all the signal will be dominated by neuropil activity and α_k will be high for pixels in that area.

- $\eta \sim \mathcal{N}(0, \sigma)$ is the random measurement noise

By jointly modelling the neuropil contamination across the field of view and the individual cell fluorescence timecourses, our fitting procedure simultaneously performs cell identification and neuropil subtraction.

The model is fit via a modified expectation-maximization (EM) algorithm that minimizes the squared error between the model's estimate of each pixel's recorded signal and that from the actual recording. [54] In the E step, the algorithm optimizes the matrix of pixel membership Λ and the neuropil scaling parameter α. Given these, the M step estimates the cell and neuropil timecourses $\mathbf{f}$ and $\mathbf{n}$ by linear regression. As mentioned above, the neuropil basis functions B_i are fixed to be two-dimensional raised cosines tiling the field of view. This algorithm converges on the optimal set of parameters relatively quickly [54], which together provide us with clusters of pixels (ROIs) corresponding to individual cells and the neuropil-free fluorescence timecourses for each of these.

To ensure that each of the identified ROIs corresponds to an actual cell, we impose certain spatial constraints (e.g. member pixels should lie close to the centroid). We then manually checked the remaining ROIs using a custom GUI [54], removing any that resembled axons or dendrites such that our resulting data set consisted only of cell somata. In this manual curation step, we also removed any cells with an aberrant fluorescent signal due to saturation with GCaMP.

In our data set, this resulted in a total of 1012 cells across 4 data sets (from 3 mice) in experiment 1, and 4953 cells across 10 recording planes in experiment 2. We then classified each cell as putative pyramidal cell or parvalbumin (experiment 1) or somatostatin (experiment 2) using the images obtained from the red-pass filtered photodiode to find the cells transgenically labelled with tdTomato.

2.3.2. *Spike detection.* One component that the above model is missing is the calcium dynamics of spiking cells. In extracting spikes from fluorescence data, we used the neuropil signal and ROIs obtained above to refit a new generative model of cell fluorescence traces that incorporated cell spiking, as has been done in previous calcium imaging data processing pipelines. [57, 69] Here, we model the fluorescence timecourse $\mathbf{f}_n$ of the nth cell as a spike train $\mathbf{s}_n$ convolved with a kernel k, with a contribution from neuropil:

$$\mathbf{f}_n = \mathbf{s}_n * \mathbf{k} + \beta_n \mathbf{p}_n + b_n$$

In the analyses performed here, we used a decaying exponential kernel with a decay constant of 100ms. $\mathbf{p}_n$ is determined by averaging the neuropil signal obtained in the pervious step (above) for each pixel over all the pixels in the ROI corresponding to the nth cell. b_n denotes baseline activity. We now fit this model to the average raw fluorescence signal in each ROI, giving us the spike trains $\mathbf{s}_n$ of each neuron and

a new estimate of level neuropil contamination β_n, presumably now more accurate because the model incorporates calcium dyanmics arising from cell spiking. [54]

Importantly, each spike train consists of a series of spikes at different times and with different amplitudes that are fit to each individual cell. Because different cells may show different levels of GCaMP expression (or different calcium dynamics if they are different types of cells [36]), the resulting spike amplitudes provide only a relative measure of spike rate for that cell itself. In other words, an amplitude of 2 indicates twice as many spikes as a previous spike of amplitude 1 by the same cell. But it does not necessarily indicate twice as many spikes as a spike of amplitude 1 in a differnt cell. This caveat is important to take into account when interpreting spiking activity in different neurons, as we do at times below. It is for this reason that we refer to "spiking activity" rather than spike rate.

The whole pipeline has been shown to accurately infer ground truth spiking activity quite accurately [54], currently leading other prominent 2-photon calcium imaging data processing pipelines (http://neurofinder.codeneuro.org/). The resulting spike trains are thus a good estimate of actual cell spiking activity.

2.4. **Characterizing neural responses.**

2.4.1. *Classifying neural responses.* Individual cell responses were classified by the difference between mean activity during the 333ms (10 imaging samples) prior to stimulus onset and the mean activity during the 333ms just after stimulus onset. We term this difference the *evoked response*. For each cell and orientation, we computed the evoked response at each stimulus presentation and then performed a sign test at $\alpha = .05$ to determine statistical significane. If a cell had a positive and statistically significant evoked response to at least one orientation, it was classified as *responsive*. If no orientations provoked a statistically significant evoked response, the cell was classified as *non-responsive*. On the other hand, if no orientations provoked a statistically significant positive evoked response and at least one orientation provoked a statistically significant *negative* evoked response (i.e. pre-stimulus onset activity was higher than post-onset activity), it was classified as *suppressive*.

2.4.2. *Estimating and quantifying orientation tuning.* Tuning curves were estimated by fitting circular Gaussian functions to the mean responses to each orientation. Unless otherwise noted, the preferred orientation of a cell was taken to be the orientation at which the Gaussian tuning curve peaked. OSI was computed as 1 - circular variance:

$$global\, OSI = \frac{|\sum_k R_k e^{i2\theta_k}|}{\sum_k R_k}$$

where R_k is the mean response to the orientation θ_k. Tuning amplitude and half-width at half-height were comptued as usual.

We quantified tuning stability as 1 minus the circular variance of the bootstrapped distribution of preferred orientation. For a given cell, this was found by randomly sampling half of all trials and estimating the preferred orientation based on these, using the vector averaging method[3]:

$$\theta_{pref} = \frac{1}{2} angle\left(\sum_k R_k e^{i2\theta_k}\right) = \frac{1}{2}\arctan\left(\frac{\sum_k R_k \sin 2\theta_k}{\sum_k R_k \cos 2\theta_k}\right)$$

By iterating this procedure 1000 times, a distribution of preferred orientations could be obtained for any given cell. The tuning stability ψ was then computed as:

$$\psi = \frac{|\sum_k e^{i2\theta_{pref}^{(k)}}|}{1000}$$

where $\theta_{pref}^{(k)}$ denotes the preferred orientation estimated from the trials sampled in the kth iteration of the sampling procedure. 1-ψ is called the circular variance of the distribution. Because an orientation is a circular quantity, circular statistics must be used to characterize a distribution of preferred orientations. ψ varies between 0 and 1, 1 denoting the most narrow distribution of preferred orientations across random samples of trials (i.e. a distribution with 0 circular variance).

To facilitate interpretation of this measure, figure 1 plots the value of ψ for orientation distributions of different widths. These are obtained by drawing 10^6 random samples from a normal distribution centered at 180° with standard deviation designated by the x-axis and measuring the ψ value of the resulting samples:

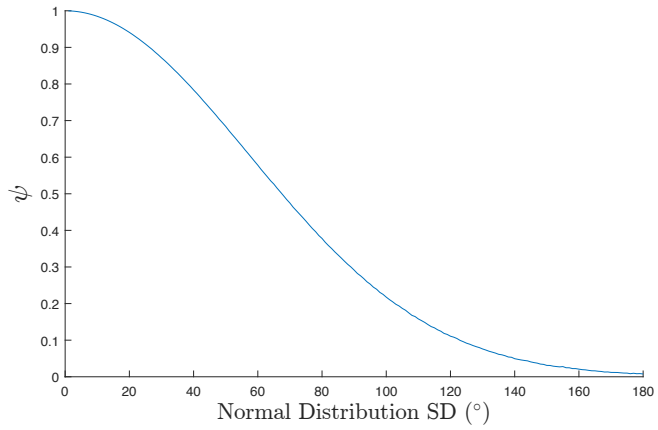

FIGURE 1. *ψ values for normal distributions of orientation with different standard deviations.* See text for details.

We thus designated $\psi = .90$ as an appropriate cut-off for deciding whether or not a cell was appropriately tuned. This value of ψ corresponds to a normal

[3]We chose this estimate of preferred orientation as opposed to that of a Gaussian fit for speed purposes

distribution of preferred orientations with a standard deviation of approximately 25°. Such a distribution would yield a 95% confidence interval 100° wide. In other words, you can be 95% sure that the true preferred orientation of a cell with $\psi = .90$ lies within 50° of your estimate, in either direction.

3. Experiment 1 Results

To investigate neural responses in different temporal contexts, we compared responses between unbiased 1hz and unbiased 10hz and then separately compared responses between unbiased and biased 10hz. We did this for putative pyramidal cells and labelled parvalbumin expressing cells.

3.1. **Putative pyramidal cells.**

3.1.1. *1Hz and 10Hz stimulus.* We found that putative pyramidal cells in our population were strikingly heterogenous in their visual evoked responses. We begin by focusing on responses to the 1Hz stimulus, where each presented grating was preceded by a 900ms inter-stimulus interval consisting of a blank screen. Figure 2 depicts the stimulus-triggered average fluorescence and spiking responses of four representative cells with different response properties. Figure 2**a** shows a particularly "well-behaved" cell, with a strong evoked response precisely time-locked to stimulus onset. Cell **b**, however, shows a similarly temporally precie response but in the opposite direction. We found several cells in our population that responded like this, seemingly suppressed by visual stimulation. Other cells did not seem to be at all responsive to oriented gratings, like that shown in figure 2**c**.

To better probe the underlying differences between such cells, we quantitatively classified all the cells in our population by their response, using the difference in activity before and after the stimulus (see methods, section 2.4.1) This resulted in classifying 424 (45%) cells as *responsive* (e.g. cell **a**), 267 (29%) as *suppressive* (cell **b**), and 243 (26%) as *non-responsive* (cell **c**). Indeed, as we show below, each of these three classes of cells showed substantial differences in their response dynamics and tuning.

We first looked at the temporal properties of the neural response by computing the cross-correlation of each cell's spike train with the stimulus timecourse[4] at different time lags. The time lag at which this cross-correlation is maximized provides an estimate of the latency at which a given cell responds strongest to visual stimulation. Figure 3 shows the relative frequency of these peak latencies in each response class, within a range of 0 to 1 second after stimulus onset (discretized by the 30Hz sampling rate of our imaging). It is readily apparent that whereas the positive responding neurons consistently responded strongest at $\sim 120 - 150$ms after stimulus onset, the suppressive and non-responsive cells were much more variable in the timing of their response.

[4]A series of 0's with a 1 at each onset time

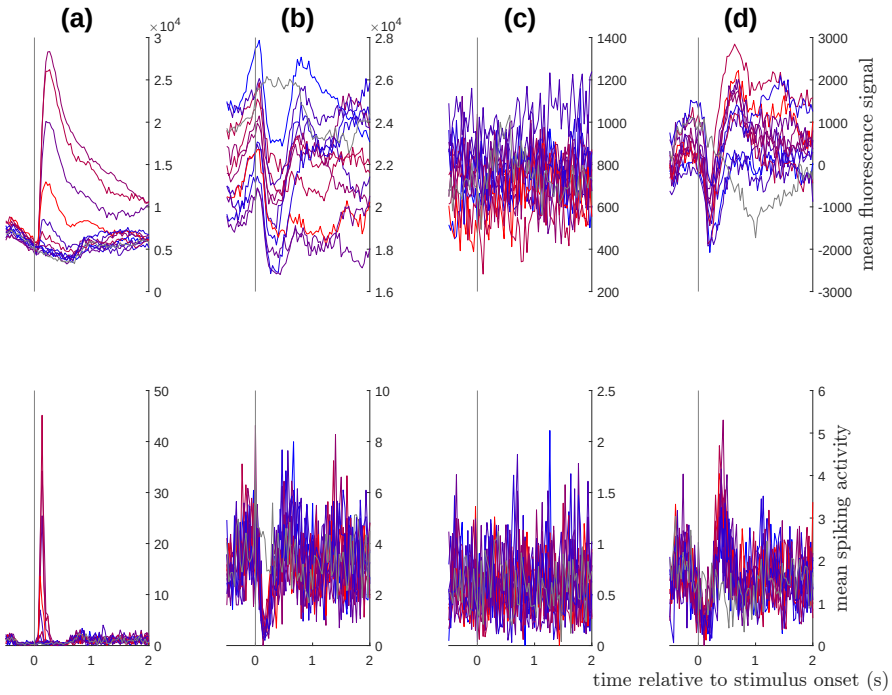

FIGURE 2. *Examples of differential responses in the putative pyramidal population.* Top panels show stimulus-triggered average of the fluorescence response, bottom panels show the same for the inferred spike trains.

Some further subtleties can be appreciated by looking at the average peri-stimulus histogram (PSTH) for each class (fig.4). While the response peaks of the *suppressive* cells indeed tend to be noisy and variable, the timing of the peak of suppression in these cells is highly precise, tightly coupled with the response peak of the *responsive* cells. The class averages also reveal that the suppressive cells tended to have much higher baseline activity than all the others, even during the blank stimulus (gray line). In fact, the average spiking actvity over the whole experiment timecourse was significantly higher in *suppressive* cells than *responsive* ones (two-sample t-test, $p < .001$), despite their substantially weaker stimulus-evoked response.[5]

Another interesting observation about the *suppressive* cells is that they often showed a rebound in activity overshooting the baseline. A representative example is shown in figure 2**d**. 40 out of the 267 *suppressive* cells showed responses like this, where the mean activity 330-660ms after stimulus onset was significantly higher

[5]Although it is important to remember here that spiking activity is not necessarily comparable between different cells, see methods section 2.3.2.

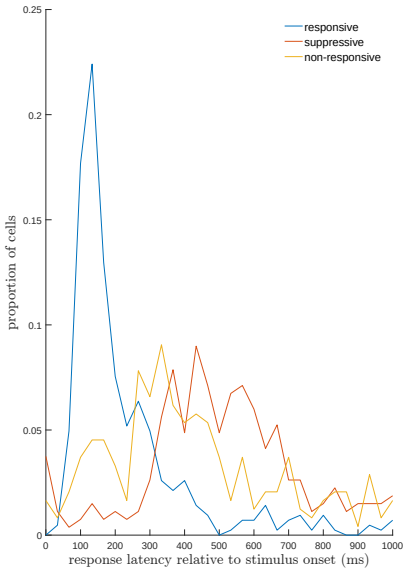

FIGURE 3. *Relative frequency histogram of peak stimulus - spike-train cross-correlation time lags.* y-axis is the proportion of cells in the given class that had their highest cross-correlation at the time lag indicated by the x-axis.

than mean acivity during the 330ms just prior to the stimulus (sign test, using FDR correction for mutiple comparisons [7] at $\alpha = .05$). Each stimulus in this condition was presented for 100ms, meaning that these cells had a strong response $\sim 200 - 400$ms after stimulus offset. Indeed, such an OFF-type response [32] can also be glimpsed in the mean PSTH (fig.4), where suppressive cells consistently peaked at around 500ms. The responses at this peak tended to be unselective for orientation, however, with a median global OSI of .05.

On the other hand, stimulus-evoked responses in the responsive classes were highly orientation tuned, as can be glimpsed from the smooth color gradient in the zoom-in of the response peak (fig.4, inset), color-coded by distance from preferred orientation (see figure caption for details). Indeed, the median OSI of *responsive* cell responses at the time of this peak (133ms) was .32. The differential orientation selectivity of these three classes is made evident by cross-validating tuning curves measured at each response latency for each cell class (fig.5). Using a cross-validation measure of tuning stability (see methods, section 2.4.2), we found that *suppressive* and *non-responsive* cell responses did not seem to differentiate between orientations at any timepoint after stimulus onset, whereas *responsive* cell

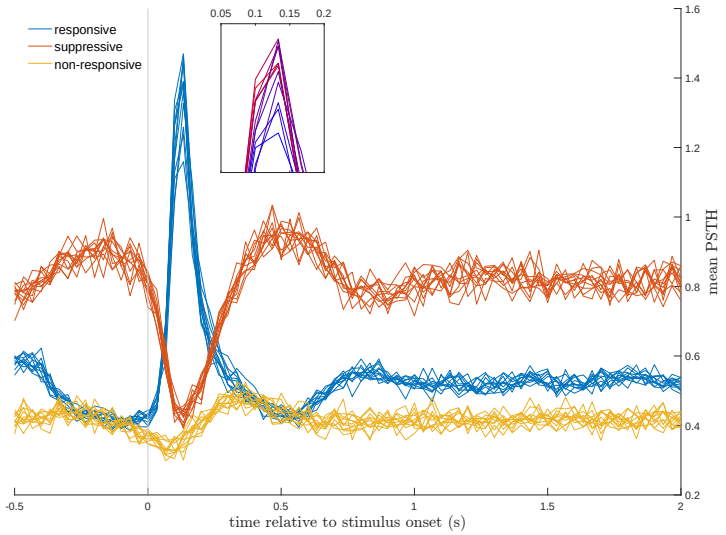

FIGURE 4. *Mean PSTH for each class.* Each trace of a given color was constructed by taking the stimulus-triggered average spike train for each cell and then averaging over all cells in each response class. This was done separately for each orientation, yielding 12 traces for each class. Importantly, we grouped orientations by their distance from each cell's preferred orientation (determined by the stimulus that produced the highest mean response during the 333ms immediately succeeding stimulus onset). For example, the responses to the 60° stimulus of a cell with orientation preference 15° were averaged together with the responses to the 150° simulus of a cell with orientation preference 105° (i.e. +45°). The inset provides a zoom-in of the peak of the mean PSTH for the *responsive* class. Here, each trace is color coded by distance from preferred orientation, where red indicates the preferred orientation (±0°) and blue indcates the orthogonal one (±90°)

responses were highly selective at 100-150ms. Interestingly, however, the magnitude of suppression in *suppressive* cells was weakly orientation tuned, with median OSI .15[6].

These sharp distinctions in response properties raise the question of whether we are distinguishing between different types of neurons. Indeed, L2/3 of visual cortex contains many types of inhibitory GABAergic as well as excitatory glutamatergic cells. Current estimates [30, 33] put the proportion of each at about 80 and 20%, respectively, with ∼ 30-40% of GABAergic cells being parvalbumin positive. [56]

[6]This was computed by taking the difference between mean response to a visual stimulus and to the blank (averaged over 333ms period after stimulus onset). This gave us a measure of the magnitude of suppression evoked by each orientation, for each cell. Taking only the *suppressive* cells that were suppressed at all 12 orientations (208/267), we computed the OSI on the magnitude of suppression at each orientation for each cell and took the median over this population.

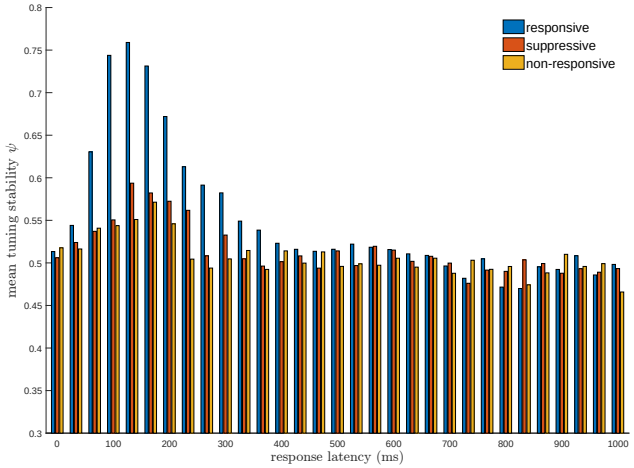

FIGURE 5. *Tuning stability at different response latencies.* y-axis is mean tuning stability computed across cells in each population, by cross-validating tuning curves obtained from responses at each latency on the x-axis. Tuning stability was measured as 1 - circular variance of the boostrapped distribution of preferred orientations, obtained by iteratively sampling a random half of trials and reestimating the preferred orientation (see methods, section 2.4.2

Because our transgenic mice have PV cells labelled with tdTomato, we can be sure our population of putative pyramidal cells did not contain any, leaving a rough estimate of $(1 - .4) \times .2)/(.8 + (1 - .4) \times .2)) \approx 13\%$ of our putative pyramidal cells likely to be non-PV interneurons. This is far below the 29% of cells we classified as *suppressive*, indicating that the cells in this class are likely pyramidal cells being inhibited. However, it is possible that the delayed OFF-type response pattern observed (fig.2*d*, 4% of all putative pyramidal cells) could be that of a certain subtype of interneuron. For example, it is known that somatostatin positive interneurons have a particularly delayed and variable reponse latency relative to pyramidal and parvalbumin cells [42, 41], which could explain the late peak in activity.

Focusing on the *responsive* class of cells, we then asked if and how stimulus-evoked responses changed in the 10Hz condition. In this condition, each stimulus was preceded by 100ms of another stimulus, with no inter-stimulus interval (except, effectively, when a blank stimulus was presented). Such a dynamic setting more closely resembles the kind of constant stream of visual stimulation the visual cortex would receive in nature, and can lead to altered single cell responses. [51, 49, 8] We found a very simple pattern of changes, where cells typically saw reduced activity in the higher temporal frequency setting. Furthermore, the cells with strongest stimulus-evoked responses in the 1Hz condition saw the largest reduction (fig.6).

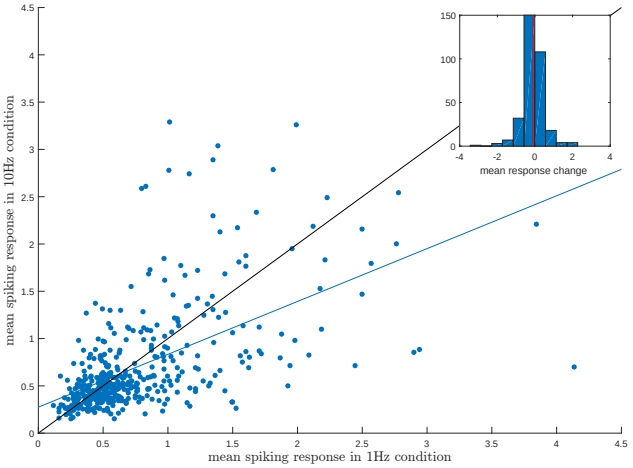

FIGURE 6. *Mean stimulus-evoked spiking responses in 1hz vs 10hz stimulus.* Stimulus-evoked response computed by taking the mean spiking response over the 333ms period immediately succeeding stimulus onset and averaging across stimulus presentations. Linear fit shows a highly significant linear relationship ($p < .001$), with slope significantly below 1 (t-test, $p < .001$). Inset shows histogram of difference between the two conditions, only for those cells with statistically significant changes (two-sample t-test, unequal variance, using FDR correction for multiple comparisons [7] with $\alpha = .05$). The median of these differences was significantly below 0 (sign test, $p < .01$), indicating that cells that saw significant differences in their responses to the two stimulus sequences on average produced weaker resposnes in the 10Hz stimulus.

Such changes are evocative of homeostatic response gain control. An example of this is plotted in fig.7**a**. That said, we observed large variability in how responses differed between the two conditions (fig.7). Indeed, some cells saw an increase in activity (fig.7**b**), and others were completely silenced (fig.7**c**). Furthermore, median time-averaged spiking activity in the population was significantly higher in the 10Hz than in the 1Hz condition (sign test, $p < .001$), suggesting that homeostatic mechanisms in fact may not be at play here.

Did these changes in response translate to changes in tuning? If so, this would have important implications for the population coding of orientation, which has previously been investigated in dynamic high temporal frequency stimuli such as our 10Hz stimulus, but in anesthetized cat. [8] We fit tuning curves to each cell's responses at 133ms after stimulus onset to estimate their preferred orientations, and computed a cross-validation measure of the stability of these tuning curves (see methods, section 2.4.2). Restricting our analysis to only those cells with stable orientation tuning in the 10Hz condition[7] ($\psi > 0.9$, see methods), we found

[7]We computed tuning stability from responses in the 10Hz condition rather than in the 1Hz condition because of the greater number of trials.

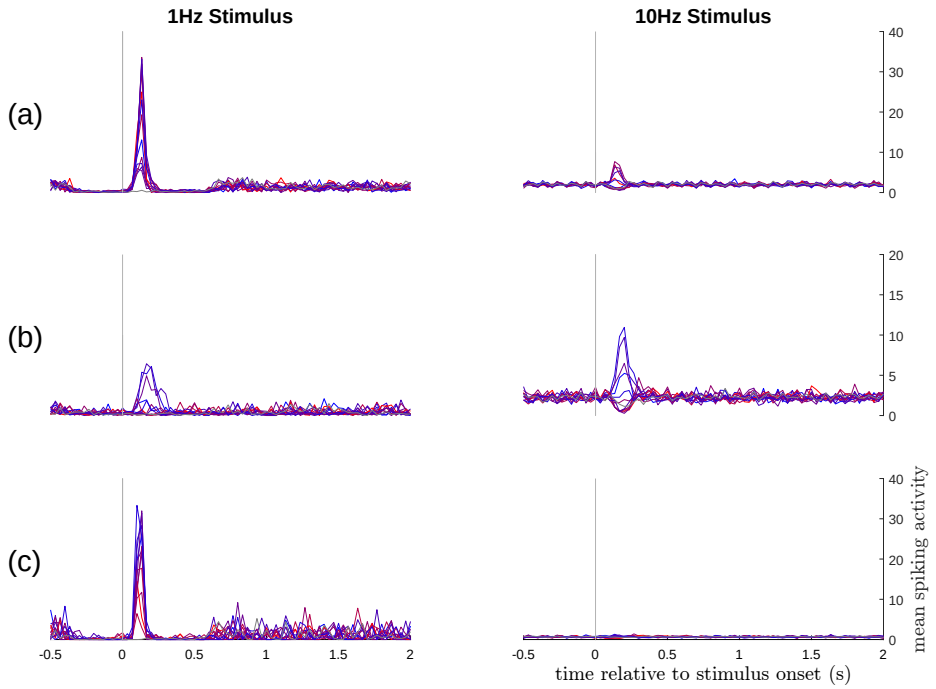

FIGURE 7. *Individual examples of differences in stimulus-evoked response between 1hz and 10Hz conditions.* Left panels show stimulus-triggered mean spike trains during the 1Hz stimulus, right panels show the same for the 10Hz stimulus. Different colored traces represent responses to different orientations, where red corresponds to 90° and blue is 0°. Row (a) shows classic response gain control type of modulation. Rows (b) and (c) show cells that underwent different changes.

that cells largely maintained their tuning, with only slight (10-20°) deviations of the preferred orientation (fig.8, top left panel). However, global OSI and amplitude were significantly reduced in the 10Hz stimulus (sign test, $p < .05$, .001, respsectively). That said, we again saw great variability in 1Hz vs 10Hz tuning differences, with some cells seeing a significant increase in tuning curve amplitude (fig.9**c**). A few representative examples are plotted in figure 9.

3.1.2. *Unbiased vs biased stimulus.* In comparing responses in the 10Hz unbiased and biased stimulus conditions we looked only at cells identified as *responsive* in the 1Hz stimulus and with a tuning stability $\psi > .90$ in the 10Hz unbiased stimulus, to constrain our analysis to cells that we were confident were orientation tuned. This resulted in a total of 270 cells, with median global OSI 0.39 and HWHH 20.2°, in line with previous reports [53, 31, 67, 3, 36] regarding orientation tuning in mouse visual cortex (although substantially lower than that found in [61]). Note, however, that this population consists of about 29% of all the putative pyramidal

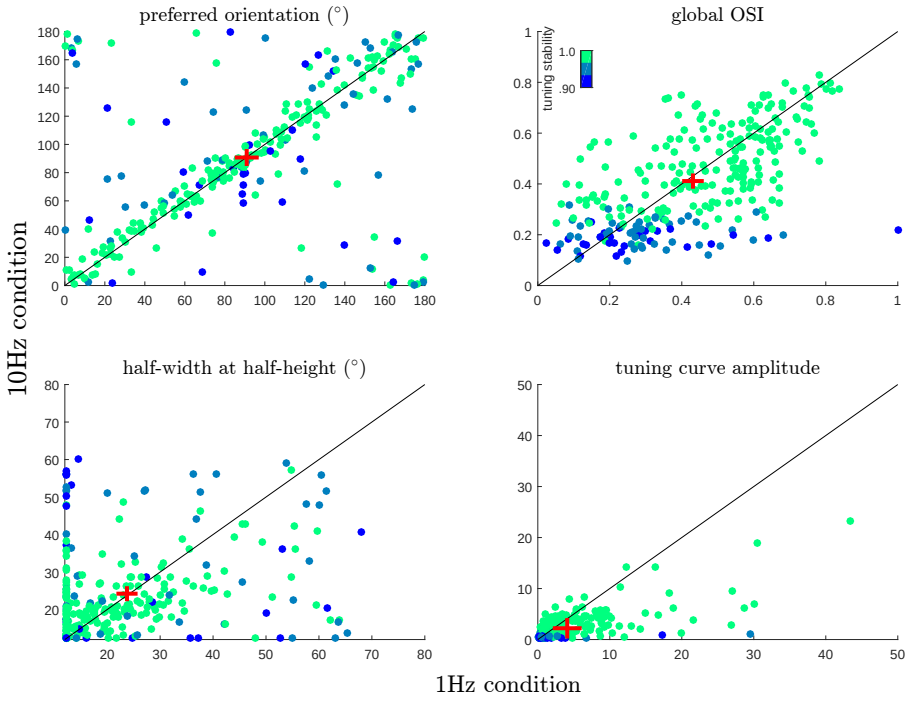

FIGURE 8. *Differences in tuning properties between 1hz and 10Hz conditions.* Only showing cells judged as having stable tuning (see text). Color designates our measure of tuning stability ψ (see section 2.4.2). Red crosses designate means.

cells we recorded. The large number of untuned cells we found may relate to the fact that we recorded from awake mice, as opposed to all the studies cited (except [53], who found largely unchanged tuning in awake mouse V1).

A few examples of the kinds of response differences observed in the biased condition are shown in figure 10. In general, the strength of the evoked response was strongly reduced in the biased stimulus, with tuning curve amplitudes reduced by 20-40%. To look at tuning changes across the population, we focused on the well-tuned cells with global OSI $> 0.25, (n = 201)$ and looked to see if classical adaptation effects were replicated in our experiment. Validating our approach, we found that responses to the adaptor stimulus (90°) were significantly lower in the biased condition - the hallmark of adaptation (fig.11**A**; sign test, $p < .001$). This can be seen as well in the cell-specific reduction in tuning curve amplitude, strongest for cells tuned to orientations near the adaptor orientation (fig.11**B**). The magnitude of this reduction at its strongest (40%) is consistent with that found in anesthetized cats. [9] However, our results did not replicate previous findings regarding adaptation-induced shifts in preferred orientation in anesthetized cat V1.

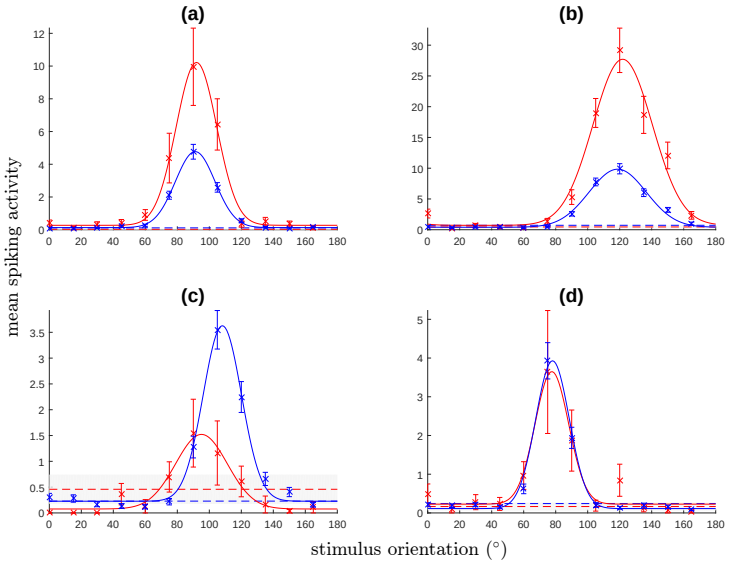

FIGURE 9. *Individual examples of tuning curves in 1Hz (red) and 10Hz (blue) conditions.* These cells were selected based on their high global OSI across both conditions and on their differential tuning differences between conditions. Error bars designate SEM across stimulus presentations and dotted lines are mean response to the blank stimulus, with shaded area showing SEM in the 1Hz condition. The same follows for all subsequently shown tuning curve plots.

[19, 9] We found that cells tuned to orientations far from the adaptor suffered tuning shifts of the same magnitude as those tuned to orientations close to the adaptor, and there was no systematicity in the direction of these shifts (fig.11**C**). Looking at the population tuning curves (obtained by averaging over cells with similar preferred orientations, see figure caption for details), the significant effect on tuning curve amplitude reduction and null effect on preferred orientation become obvious (contrast with figures 1c,d in [9]). The overall effect of these tuning changes did not lead to a homeostatic equalization of time-averaged responses (fig.13), as opposed to that seen in anesthetized cats. [9]

3.2. **Parvalbumin cell responses.**

3.2.1. *1Hz and 10Hz stimulus.* Our whole recorded cell population contained 78 identified parvalbumin-positive (PV) cells. We found that responses in this population were by and large quite similar to those of the putative pyramidal population. Indeed, we encountered the same three types of responses in the 1Hz stimulus (fig.14), leading us to classify 28 (36%) cells as *responsive* (fig.14**a**), 40 (51%) as *suppressive* (fig.14**b**), and 10 (13%) as non-responsive (fig.14**c**). 12 *suppressive* cells showed significant OFF-type responses (fig.14**d**). The finding that only 36%

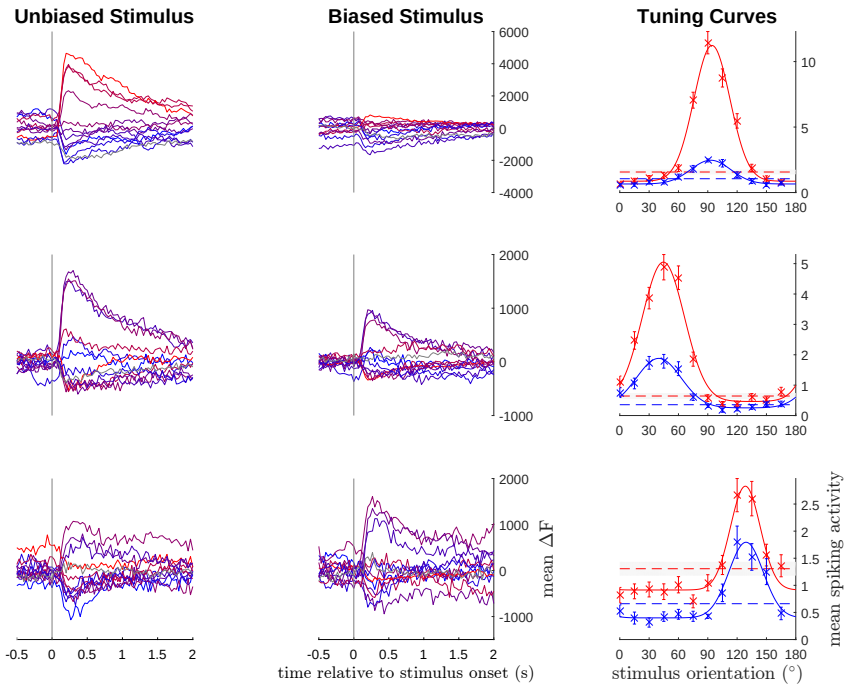

FIGURE 10. *Examples of differential responses in the biased and unbiased stimulus conditions.* First two panels in each row plot stimulus-triggered mean fluorescence traces in the unbiased (left) and biased (middle) stimulus conditions. Mean activity prior to stimulus onset is subtracted from each trace (Δ F) to facilitate comparison of the evoked response. Rightmost panel in each row shows resulting tuning curves in unbiased (red) and biased (blue) conditions. Note the stability of the preferred orientation - orientation tuning shifts were not typically observed. Adaptor orientation was 90°, so a noticeable shift right would have been expected in the top cell, from previous findings [19, 9]. Tuning curves plotted as in fig.9

of PV cells reponded strongly to visual stimulation is surprising in light of previous reports yielding proportions as high as 70% in anesthetized mouse V1. [36] Again, this could be a result of our mice being awake during recording.

In this case, however, we found that very few cells were tuned for orientation. Performing the same analysis as in figure 5 above, we found that PV cell activity did not seem to be selective for orientation at any response latency (fig.15). The mean tuning stability was maximal for *responsive* cells at 133ms after stimulus onset, just like for the *responsive* putative pyramidal cells, but with $\psi = 0.6$, ∼20% lower. Importantly, a ψ value of 0.6 corresponds to highly variable tuning across stimulus presentations, approximately equivalent to a normally distributed preferred orientation with standard deviation of 58° (see methods, section 2.4.2). This implies that the 95% confidence interval of the preferred orientation of a

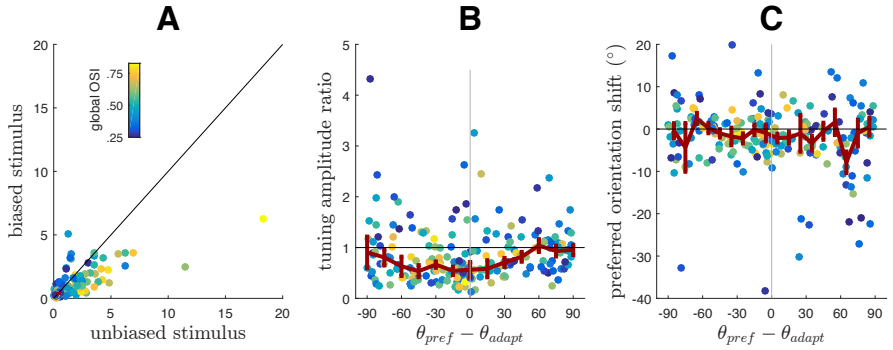

FIGURE 11. *Changes in tuning properties in biased stimulus condition.* Color of each point denotes OSI. Only cells with stable tuning ($\psi > 0.90$) and high OSI (global OSI > 0.25) are included ($n = 201$). Panel A plots mean responses to the adapator stimulus (90°) in each condition. Dark red cross denotes median across cells. Panel B plots the biased to unbiased condition ratio of tuning amplitudes. Dark red bars denote bin medians and SEM. Cells were binned into 15° bins by distance of preferred orientation from adaptor orientation ($\theta_{pref} - \theta_{adapt}$), same bins as in fig.12. Panel C depicts changes in preferred orientation between unbiased and biased conditions. Again, dark red data points denote bin medians and SEM, here using 10° bins.

cell with $\psi = 0.6$ is about 116° wide. We found that only 12 PV cells (15% of PV population) satisfied our criterion for sufficiently stable tuning ($\psi > .90$ in the 10Hz condition), and these had a median OSI of .12 and median half-width at half-height of 32.0. Four representative cell tuning curves are plotted in fig.16). This is in line with the .1 to .26 range found in previous studies in anesthetized mouse V1 [31, 36, 3] (although substantially lower than others [61]). The population of PV cells as a whole, however, had a substantially lower median OSI of .08.

An interesting difference between the *responsive* PV cells and the *responsive* putative pyramidal cells was that the 10Hz stimulus elicited an increase in activity in the PV cells, rather than a decrease (fig.17, vs fig.6). These changes seemed to largely consist of an increase in baseline spiking activity, leading to a reduction of the amplitude of the stimulus-evoked response. Two representative examples are shown in figure 18. Indeed, the time-averaged spiking activity for this group of cells was significantly higher in the 10Hz condition than in the 1Hz (sign test, $p < .001, n = 28$), as was the response to the blank stimulus (dotted lines in fig.16), suggesting that spontaneous activity was much higher in this condition. This is consistent with the hypothesis that PV cell activity simply reflects the average activity of the local pyramidal population [31, 36], as many pyramidal cells would have been actively responding to the high abundance and frequency of stimuli in that condition. This would also explain our observation that *responsive* PV cell OSIs were significantly lower in the 10Hz condition (sign test, $p < .001, n = 28$), and preferred orientations seemed to vary between the two conditions (see fig.16 for a few examples). In summary, it seems that PV cells in the *responsive*

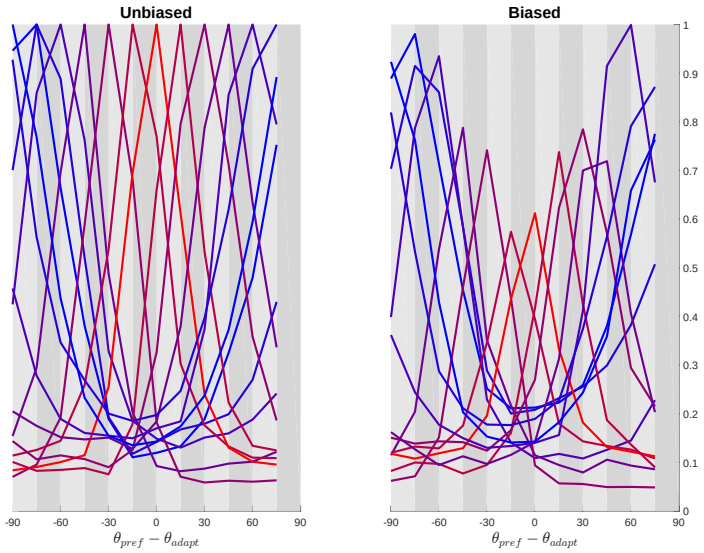

FIGURE 12. *Mean population tuning curves in unbiased and biased stimulus conditions.* Cells were binned into 15° bins by their preferred orientations in the unbiased stimulus condition, and their responses to each orientation were then averaged within each bin to produce each of the plotted population tuning curves. Population tuning curves were then scaled by their maximum in the unbiased condition. Color coded by bin center, with red = 90° (adaptor orientation) and blue = 0°. Shaded vertical bars indicate space between bin centers, such that each bin center lies at the border between two bars with different shading. These make it evident that, in the biased stimulus condition, the population tuning curves do not stray from their preferred orientation bin centers - the classic tuning curve repulsion effect of adaptation [19, 9] is not replicated.

class unselectively increase their spontaneous activity when presented with higher frequency stimuli.

3.2.2. *Unbiased vs biased stimulus.* As we did with the putative pyramidal cells, we restricted our analysis here to *responsive* PV cells with $\psi > 0.9$ ($n = 12$). Some examples of the kinds of response changes induced by the biased stimulus are shown in figure 19. Responses did not seem to modulate in the same manner as in the putative pyramidal cell population. The amplitude of the stimulus-evoked fluorescence response remained largely unchanged, and responses to the adaptor orientation were, on average, maintained across both conditions (sign test, $p = .77, n = 12$). Looking at individual tuning curves of the 12 tuned PV cells (fig.20) reveals the lack of orientation tuning of these cells and the variability in tuning changes between the unbiased and biased stimulus conditions. Overall, PV cells did not seem to be orientation tuned, leading to what looks like largely noise in the tuning differences between the unbiased and biased stimulus.

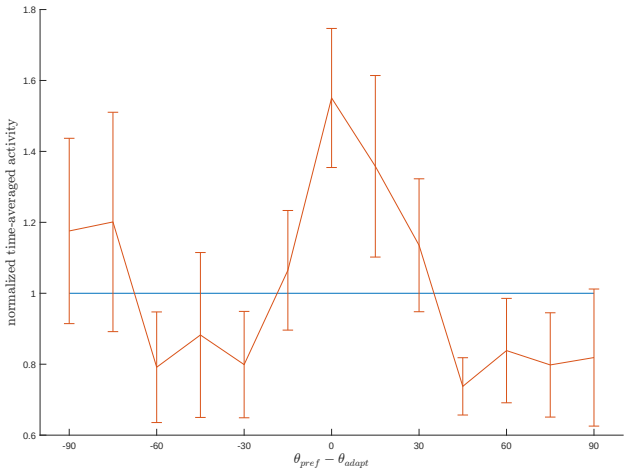

FIGURE 13. *Mean time-averaged activity of each preferred orientation bin.* Spiking activity over the whole timecourse in each condition (unbaised, biased) was averaged for each cell, and then means were computed for each bin from figure 12. The mean for each bin was then scaled to that in the unbiased condition. Red line indicates bin means in biased condition, with error bars showing SEM for each bin.

3.3. **Noise correlations between different response classes.** To furhter investigate the possibility that the different response classes we observed formed differential subtypes of neurons, we computed noise correlations between pairs of cells in each class (Pyr/PV, *responsive*/*suppressive*/*OFF-type*/*non-responsive*). Different patterns of noise correlations for each class would suggest differential connectivity patterns, providing a distinction between them beyond the their responses that we classified them by.

Due to time constraints, we were not able to fully delve into this analysis, but I briefly mention two points. Firstly, parvalbumin-positive cells had substantially more and larger noise correlations between each other than did putative pyramidal cells (the yellowish bottom section of the triangle in each matrix), replicating previous results in anesthetized mouse V1. [31] Secondly, putative pyramidal cells with OFF-type responses (oPC) had significantly larger noise correlations with PV cells than any other class of putative pyramidal cell, particularly with PV cells with likewise OFF-type responses. It is difficult to exactly interpret this result, but it indeed agrees with the idea that these cells comprised a subtype of cell distinct from the rest of the putative pyramidal population.

The above results were computed from responses to the 1Hz stimulus. This correlation structure remained qualitatively similar across the 1Hz, 10Hz, and 10Hz biased stimulus conditions, with somewhat more and stronger noise correlations

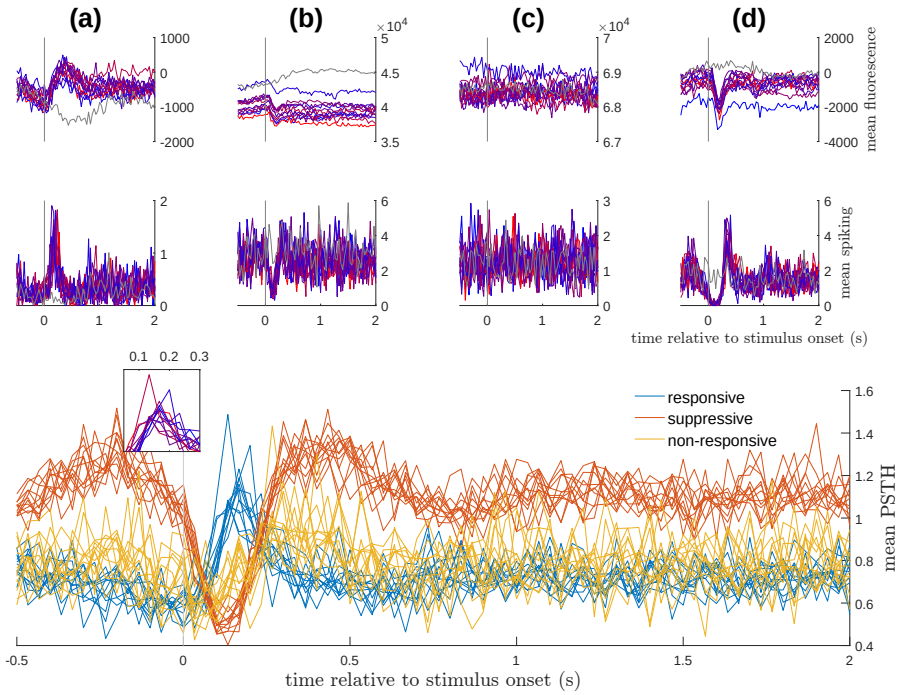

FIGURE 14. *Examples of PV cell responses in the 1Hz stimulus and mean PSTHs for each response class.* Analysis performed exactly as in figures 2 and 4 above.

in the 10Hz stimuli than in the 1Hz stimulus. Unfortunately we did not have time to look into this carefully.

4. Experiment 2 Results

Due to lack of time, we were not able to apply the same battery of analyses to our somatostatin cell labelled recordings, but I will briefly report two findings.

Firstly, we found that somatostatin cells are highly direction tuned, with a median global DSI[8] of .26, compared to .25 for putative pyramidal cells. Indeed the distribution of global DSI across each of the populations were statistically indistinguishable (two-sample Kolmogorov-Smirnov goodness-of-fit test, $p = .12$, fig.22). We were not able to classify different responses within each of these populations so we can't say at this moment whether there were underyling subpopulations with differntially lower or higher levels of tuning. Having collapsed across all response classes, this number is still significantly higher than the .18 median OSI of PV

[8]This was computed exactly as global OSI, but adapted for 360° (i.e. the same equation for global OSI in section 2.4.2 but without multiplying θ_k by 2)

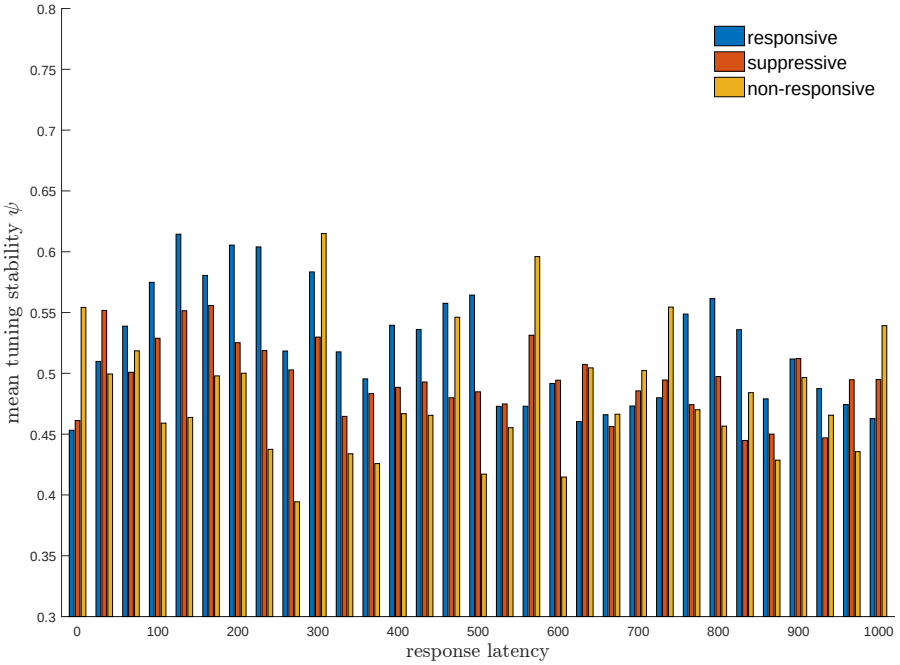

FIGURE 15. *PV cell tuning stability at different response latencies.* Analaysis and plot exactly as in fig.5

cells in the 1Hz condition of experiment 1. That said, we did not get to quantifying tuning stability of any cells in this experiment, leaving open the possibility that some of the highly tuned SOM cells were simply showing spurious selective responses.

Interestingly, by comparing responses in the control and adaptation conditions, we found that somatostatin-expressing and putative pyramidal cells both showed the hallmarks of adapatation. Narrowing down our population to only those with global DSI > 0.25, we found that responses to the adaptor orientation (180°) were significantly lower in the adaptor condition than in the control condition (see methods, section 2.2), for both putative pyramidal cells (sign test, $p < .001$) and SOM cells (sign test, $p < .05$). In line with previous findings in A1 [50], this suggests that SOM cells in awake mouse V1 undergo adaptation effects alike those seen in pyramidal cells. However, the effecs were markedly weaker for SOM cells, with a mean difference between conditions of -.07, versus -.14 in the putative pyramidal cells.

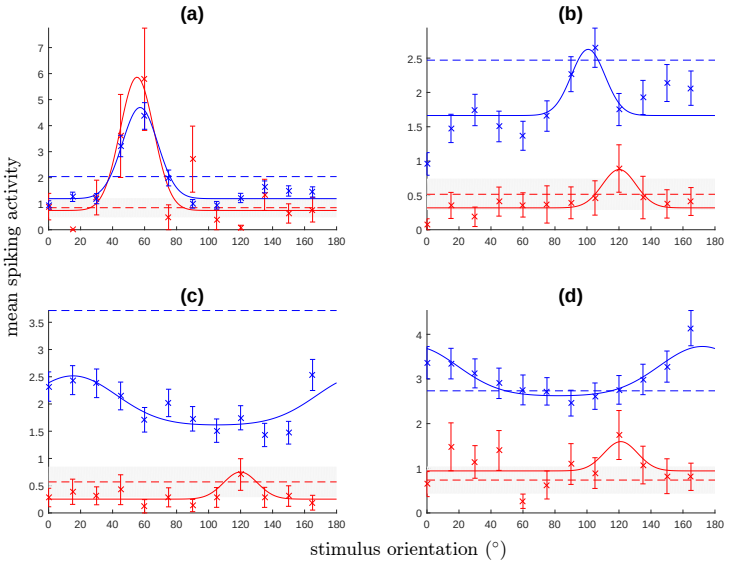

FIGURE 16. *Individual PV cell tuning curves in 1Hz stimulus and 10Hz stimulus.* Cells were selected based on satisfying our criterion for tuning stability ($\psi > 0.9$ in the 10Hz condition) and on having high gobal OSI relative to the rest of the population. Red indicates responses in the 1Hz stimulus, blue indicates 10Hz stimulus. Almost all 12 PV cells with stable tuning showed this pattern of increased activity across the board in the 10Hz stimulus, including responses to the blank stimulus (dotted lines). Most also saw shifts in preferred orientation, as in cells (b), (c), and (d).

5. Discussion

One of the main novelties of our experiment was that we measured orientation tuning in the visual cortex of *awake* mice. Most previous studies regarding orientation tuning in V1 have focused on experiments with anesthetized animals, where neural responses may be affected by the anaesthetic. We found that responses were extremely diverse (figs.2 & 14), with a minority of cells being orientation tuned. Specifically, we found that 29% of our putative pyramidal cell population and 15% of our PV cell population showed stable orientation tuning. While V1 cells may be responding solely to visual input when in the ansethetized state, our results suggest that recurrent cortical activity arising from other ongoing processes in the awake mouse may be influencing their responses in the awake state.

A related outstanding question from our results is the origins of the different types of responses observed. Particularly, we found that a large proportion of both parvalbumin-expressing interneurons (51%) and putative pyramidal cells (29%) were suppressed by visual stimulation (fig.2,14 b). The origin of this suppression is not clarified by the analyses performed here, but our finding that the magnitude of the suppression was mildly orientation tuned suggests that it may arise from

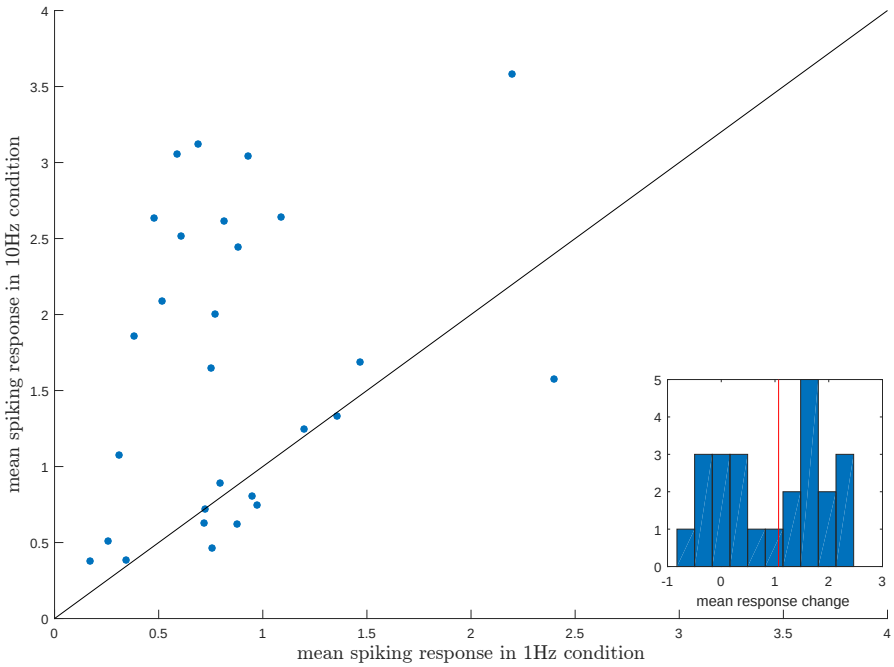

FIGURE 17. *Mean stimulus-evoked spiking responses in 1hz vs 10hz stimulus.* Analysis performed just as in fig.6. Histogram shows cells with statistically significant changes (24/28), median of this group of cells is significantly greater than 0 (sign test, $p < .05, n = 24$).

recurrent cortical inputs. Again, this may be the result of cortical processing ongoing in the awake mouse that may not be present in the anesthetized state.

A second interesting type of response we encountered was a delayed response following suppression. These responses were not orientation tuned but were relatively time-locked to stimulus onset/offset. Furthermore, cells with this type of response showed a distinct noise correlation pattern (fig.21), suggesting they were indeed a separate subpopulation of cells. One plausible hypothesis might be that these are a different type of cell from the rest. However, the fact that this type of response was found in both labelled PV cells as well as putative pyramidal cells suggests that, alternatively, it might be a response pattern arising from differential cortical and/or feed-forward interactions, rather than the result of differential physiological properties of a distinct type of neuron.

In seeking to answer questions about contextual processing in visual cortex, we also investigated responses in several different temporal contexts. Firstly, we looked to see how responses were modulated by an increase in stimulus frequency. Unsurprisingly, we found that putative pyramidal cells largely maintained their

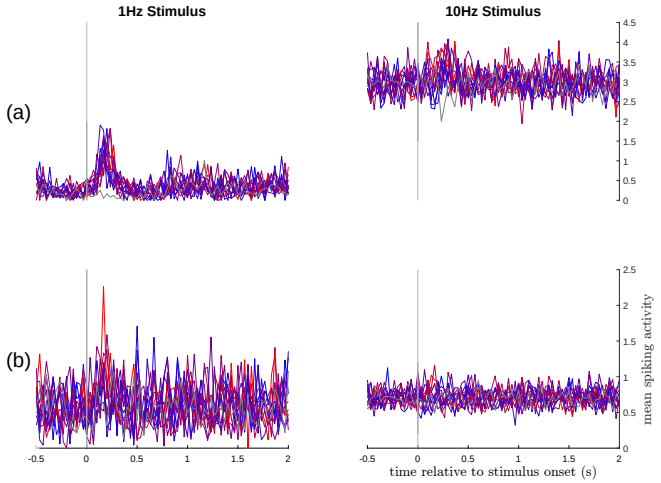

FIGURE 18. *Individual examples of differences in PV cell stimulus-evoked response between 1hz and 10Hz conditions.* These were picked based on showing large and significant differences between the two conditions. Most *responsive* PV cells showed differences like that seen in cell (a). However, a significant number of cells showed differences similar to cell (b), where baseline activity was largely unchanged and the evoked response was completely eliminated.

orientation tuning properties, maintaining their population code. [8] However, PV cells showed a wide variety of tuning changes, supporting the idea that thay are laregly untuned for orientation. [31, 36, 67]

We also found classic signatures of response gain control in the putative pyramidal cell population, whereby response strength was diminished when responses became more frequent, as in the 10Hz stimulus. Furthermore, neurons with the strongest responses in the 1Hz stimulus saw the largest reductions in their responses. Interestingly, PV cells showed the opposite pattern: their evoked responses were stronger in the 10Hz stimulus than in the 1Hz stimulus. These complementary differences could potentially be explained by increased pyramidal cell firing from the more frequent visual stimulation during the 10Hz stimulus triggering an increase in PV cell firing. Via the resultive IPSPs in their pyramidal cell targets, this would then lead to attenuation of pyramidal cell responses. This explanation is in line with claims that PV cells implement response gain control in V1. [3, 70] However, it is worth noting here that synapses from pyramidal cells to PV cells have been shown to be depressing [5, 60], which could interefere with this putative causal mechanism. Alternatively, the reduction in putative pyramidal cell responses may simply be the result of spike-frequency adaptation [43, 59] or other short-term adaptation effects elicited by presenting stimuli in quick succession with no inter-stimulus interval. [51, 49]

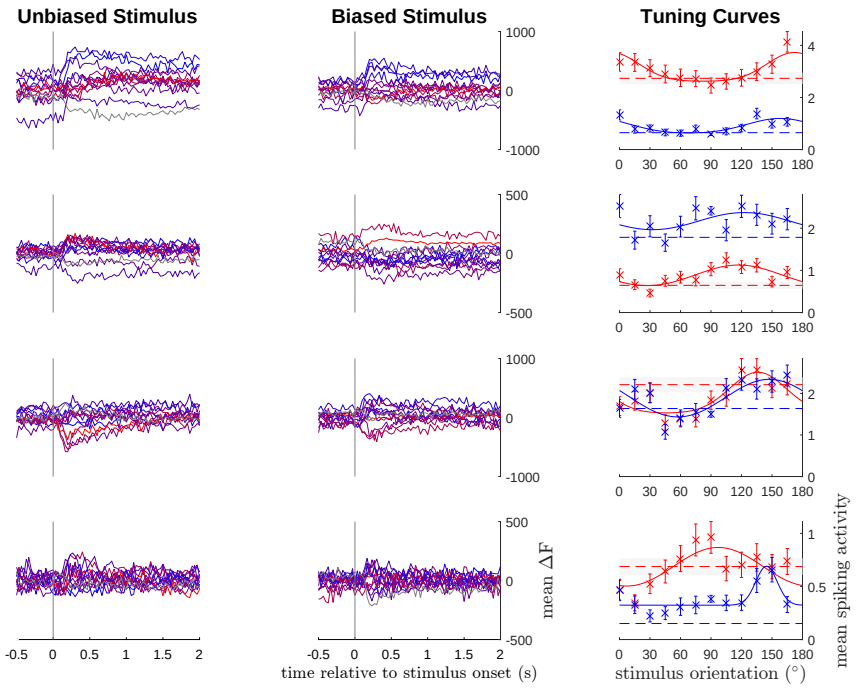

FIGURE 19. *Individual PV cell stimulus-trigerred responses and tuning curves in the unbiased and biased conditions.* Figure organized exactly as fig.10. Note that, in contrast to the putative pyramidal cell population, amplitude of PV cell evoked fluorescence responses tended to remain unchanged between the two conditions. Furthemore, tuning curve differences between the two conditions had little systematicity, with a wide mix of vertical shifts (in activity) and horizontal shifts (in preferred orientation).

Although not reported in the text above, we also found that 70% of all *responsive* putative pyramidal - *responsive* PV cell pairs had statistically significant noise correlations ($p < .05$, FDR correction for multiple comparisons [7]) in the 10Hz condition, although with a relatively low mean Pearson correlation of .012. Importantly, these two numbers dropped to 32% and .008 in the 1Hz condition, suggesting that the higher stimulus frequency elicited more pyramidal-PV cell coactivation. This possibility supports the hypothesis that stronger PV cell activity is recruited in the 10Hz condition by the more frequent visual stimulation of pyramidal cells, possibly leading to response gain control. Further analyses in this vein may better clarify the cortical dynamics underlying the observed differences in responses to stimuli with different temporal frequencies.

A more behaviorally relevant type of contextual modulation is that elicited by changes in the statistics of the stimuli. This is what we tested with the biased and

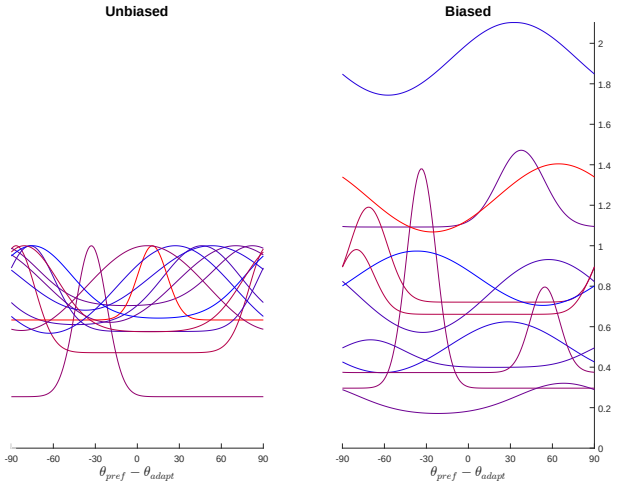

FIGURE 20. *Tuning curves tuned PV cells.* Tuning curves plotted for all 12 PV cells satisfying our criterion for stable tuning. As in fig.12, tuning curves were scaled to the peak in the unbiased condition. Cells showed a wide variety of differences between the two conditions. Note the cell with preferred orientation closest to the adaptor orientation (the red curve), which saw an increase in responses to all orientations.

unbiased stimulus conditions. We replicated the classical finding that V1 pyramidal cell responses to an adaptor stimulus are significantly reduced (fig.11A), leading to larger reduction of the tuning curve amplitude of cells tuned to orientations near that of the adaptor stimulus (fig.11B). These results suggest that adaptation effects are generated cortically, as upstream lateral geniculate nucleus neurons are not orientation tuned and thus shouldn't be able to elicit orientation-specific adaptation effects. [17]

However, we did not replicate the classical tuning curve repulsion effects observed in anesthetized cats. [19, 9] It is hard to see how anesthesia would lead to such an effect, so it is more likely that these effects may arise from cortical interactions inherent to the orientation column organization found in cat V1. Mouse V1, on the other hand, has a "salt and pepper" organization [29] such that the distance between any two pyramidal cells has no relationship to the difference in their preferred orientations. However, pyramidal cells in mouse V1 nonetheless preferentially connect to other pyramidal cells with similar orientation preference. The crucial difference between cat and mouse V1 may thus be that interneurons, which make dense local connections [22], will have post-synaptic targets with more similar orientation tuning in cat V1 than in mouse V1, since in cat V1 the locality constraint implies connnections onto cells within the same cortical column. Indeed, modelling work shows that if inhibitory neuron responses are slightly tuned (which would arise naturally in cat V1 from dense inputs within an orientation

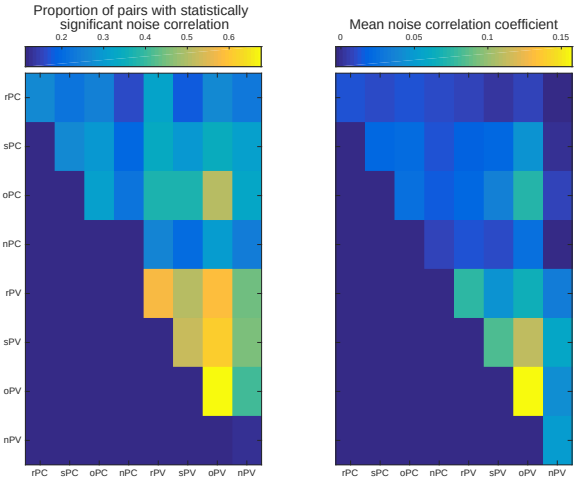

FIGURE 21. *Noise correlations between pairs of neurons of different classes.* Left panel shows the proportion of cell pairs with statistically significant noise correlations in each combination of response classes, computed by Pearson correlation between each cell's responses across stimulus presentations (averaged over 333ms after stimulus onset) and correcting for multiple comparisons with FDR correction [7] at $\alpha = .05$. Right panel shows the mean Pearson correlation coefficient between all cell pairs in each combination of response classes. These matrices would usually be symmetric, so we simply set all values below the diagonal to 0. PC = putative pyramidal cell, PV = parvalbumin-positive cell. r = *responsive* (e.g. fig.2/14a), s = *suppressive* (fig.2/14b), o = *OFF-type* (fig.2/14d), and n = *non-responsive* (fig.2/14c).

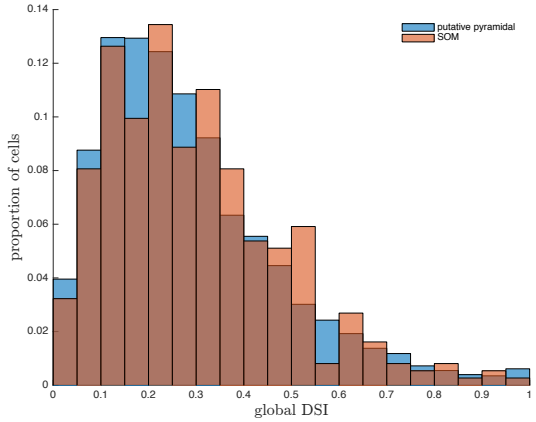

FIGURE 22. *Distribution of direction selectivity in putative pyramidal and somatostatin populations.* Labelled somatostatin-expressing interneurons showed equally high direction tuning as putative pyramidal cells, with statistically indistinguishable distributions (two-sample Kolmogorov-Smirnov goodness-of-fit test, $p = .12$).

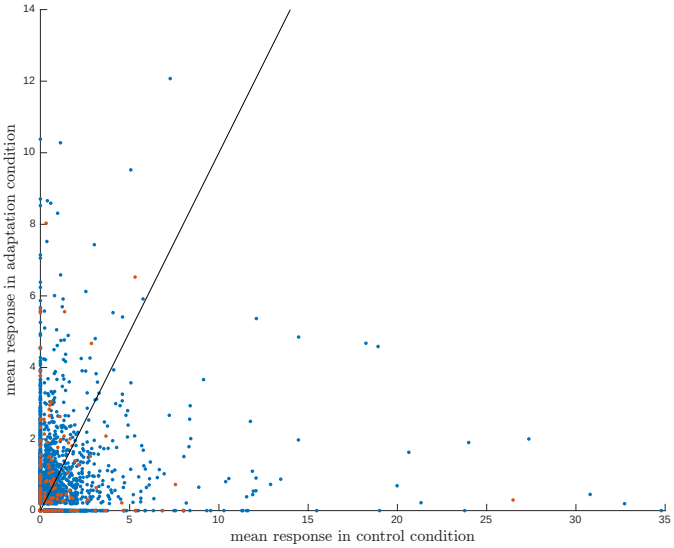

FIGURE 23. *Individual cell mean responses to adaptor orientation in control condition and adaptor condition.* Black line is the unity line - note the different scales on the x- and y- axes.

column) and are reduced by adaptation (or, alternatively, by depressing synapses from and/or onto PV cells [24, 5]), repulsive tuning curve shifts may arise from the resulting recurrent dynamics. [68]

In mouse visual cortex, however, we found no reduction in parvalbumin-expressing interneuron responses by adaptation and no repulsive tuning curve effects. In fact, our results our explained simply by single cell spike-frequency adaptation [43, 59], leading to response reductions proportional to the difference between cells' preferred orientations and the adaptor orientation. In principle, this simple explanation should account for our population-level adaptation effects (fig.12). Future modeling work should confirm this.

As already mentioned, we found PV cells not to show any adaptation effects, their responses to the adaptor stimulus unaffected by its increased frequency in the biased stimulus condition. This contrasts previous findings in anesthetized mouse A1, where stimulus-specific adaptation is observed in PV cells. [50] Our data are consistent with PV cells being largely untuned for orientation, with only 15% of our whole population showing stable tuning curves. PV cells in A1, on the other hand, may be more tuned to frequency. Overall, our results seem to support the hypothesis that PV cell activity in V1 simply reflects average local activity via unselective excitatory pyramidal inputs. [31, 36]

Interestingly, however, we did find that somatostatin-expressing interneurons were strongly direction tuned and showed the hallmark effects of adaptation. SOMs have been shown to be important for stimulus-specific adaptation in A1 [50] and seem to undergo stimulus-specific adaptation themselves. [50] Indeed, SOMs are well-suited to drive adaptation in pyramidal cells by via their facilitating synapses. [5, 50] It thus remains an open possibility that adaptation in mouse V1 pyramidal cells is driven by cortical connections from somatostatin-expressing interneurons. One could imagine a simple mechanism whereby highly tuned SOM cells selectively inhibit similarly tuned pyramidal cells. Under a stimulus context heavily favoring one orientation or direction, those SOM cells tuned to the adaptor orientation/direction would be preferentially activated, thus leading to the selective inhibition of pyramidal cells with orientation preferences also near the adaptor. Recent evidence [70], however, suggests that SOM cells do not selectively inhibit similarly tuned pyramidal cells. Regardless, adaptation in SOM cells could not be mediated by SOM cell inhibition, since SOM cells do not typically inhibit each other. [56] It is thus likely that spike-frequency adaptation does play a prominent role in adaptation effects in awake mouse visual cortex.

To summarize, contextual processing in mouse V1 seems to be mediated by single cell adapatation effects, rather than inhibitory cortical circuitry. The results presented here, however, do not settle the question. We emphasize that future modelling work may be able to elucidate cortical contributions to the kinds of dynamics we observed. These models should be constrained by the fact that adaptation effects seem to involve modulation of the activity, and not the selectivity, of cells with stimulus selectivity matching the context. Understanding the responses and dynamics described above within a framework of temporal contextual processing may lead to a better understanding of cortical circuitry and computation.

References

[1] Adesnik, H., Bruns, W., Taniguchi, H., Huang, Z. J., and Scanziani, M. A neural circuit for spatial summation in visual cortex. *Nature 490*, 7419 (2012), 226–231.

[2] Anderson, J. S., Carandini, M., and Ferster, D. Orientation tuning of input conductance, excitation, and inhibition in cat primary visual cortex. *Journal of neurophysiology 84*, 2 (2000), 909–926.

[3] Atallah, B. V., Bruns, W., Carandini, M., and Scanziani, M. Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron 73*, 1 (2012), 159–170.

[4] Baccus, S. A., and Meister, M. Fast and slow contrast adaptation in retinal circuitry. *Neuron 36*, 5 (2002), 909–919.

[5] Beierlein, M., Gibson, J. R., and Connors, B. W. Two dynamically distinct inhibitory networks in layer 4 of the neocortex. *Journal of neurophysiology 90*, 5 (2003), 2987–3000.

[6] Ben-Yishai, R., Bar-Or, R. L., and Sompolinsky, H. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences 92*, 9 (1995), 3844–3848.

[7] Benjamini, Y., and Hochberg, Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)* (1995), 289–300.

[8] Benucci, A., Ringach, D. L., and Carandini, M. Coding of stimulus sequences by population responses in visual cortex. *Nature neuroscience 12*, 10 (2009), 1317–1324.

[9] Benucci, A., Saleem, A. B., and Carandini, M. Adaptation maintains population homeostasis in primary visual cortex. *Nature neuroscience 16*, 6 (2013), 724–729.

[10] Carandini, M. From circuits to behavior: a bridge too far? *Nature neuroscience 15*, 4 (2012), 507–509.

[11] Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., Gallant, J. L., and Rust, N. C. Do we know what the early visual system does? *The Journal of neuroscience 25*, 46 (2005), 10577–10597.

[12] Carandini, M., and Ferster, D. A tonic hyperpolarization underlying contrast adaptation in cat visual cortex. *Science 276*, 5314 (1997), 949–952.

[13] Carandini, M., and Ringach, D. L. Predictions of a recurrent model of orientation selectivity. *Vision research 37*, 21 (1997), 3061–3071.

[14] Cardin, J. A., Palmer, L. A., and Contreras, D. Stimulus feature selectivity in excitatory and inhibitory neurons in primary visual cortex. *The Journal of Neuroscience 27*, 39 (2007), 10333–10344.

[15] 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. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature 499*, 7458 (2013), 295–300.

[16] Clifford, C. W., Webster, M. A., Stanley, G. B., Stocker, A. A., Kohn, A., Sharpee, T. O., and Schwartz, O. Visual adaptation: Neural, psychological and computational aspects. *Vision research 47*, 25 (2007), 3125–3131.

[17] Dhruv, N. T., and Carandini, M. Cascaded effects of spatial adaptation in the early visual system. *Neuron 81*, 3 (2014), 529–535.

[18] Di Cristo, G., Wu, C., Chattopadhyaya, B., Ango, F., Knott, G., Welker, E., Svoboda, K., and Huang, Z. J. Subcellular domain-restricted gabaergic innervation in primary visual cortex in the absence of sensory and thalamic inputs. *Nature neuroscience 7*, 11 (2004), 1184–1186.

[19] Dragoi, V., Sharma, J., and Sur, M. Adaptation-induced plasticity of orientation tuning in adult visual cortex. *Neuron 28*, 1 (2000), 287–298.

[20] El-Boustani, S., Wilson, N. R., Runyan, C. A., and Sur, M. El-boustani et al. reply. *Nature 508*, 7494 (2014), E3–E4.

[21] Felsen, G., Touryan, J., and Dan, Y. Contextual modulation of orientation tuning contributes to efficient processing of natural stimuli. *Network: Computation in Neural Systems 16*, 2-3 (2005), 139–149.

[22] Fino, E., Packer, A. M., and Yuste, R. The logic of inhibitory connectivity in the neocortex. *The Neuroscientist 19*, 3 (2013), 228–237.

[23] Fino, E., and Yuste, R. Dense inhibitory connectivity in neocortex. *Neuron 69*, 6 (2011), 1188–1203.

[24] Gabernet, L., Jadhav, S. P., Feldman, D. E., Carandini, M., and Scanziani, M. Somatosensory integration controlled by dynamic thalamocortical feed-forward inhibition. *Neuron 48*, 2 (2005), 315–327.

[25] Garcia-Lazaro, J., Ho, S., Nair, A., and Schnupp, J. Shifting and scaling adaptation to dynamic stimuli in somatosensory cortex. *European Journal of Neuroscience 26*, 8 (2007), 2359–2368.

[26] Gentet, L. J., Kremer, Y., Taniguchi, H., Huang, Z. J., Staiger, J. F., and Petersen, C. C. Unique functional properties of somatostatin-expressing gabaergic neurons in mouse barrel cortex. *Nature neuroscience 15*, 4 (2012), 607–612.

[27] Gibson, J. J., and Radner, M. Adaptation, after-effect and contrast in the perception of tilted lines. i. quantitative studies. *Journal of Experimental Psychology 20*, 5 (1937), 453.

[28] Hamm, J. P., and Yuste, R. Somatostatin interneurons control a key component of mismatch negativity in mouse visual cortex. *Cell reports 16*, 3 (2016), 597–604.

[29] Harris, K. D., and Mrsic-Flogel, T. D. Cortical connectivity and sensory coding. *Nature 503*, 7474 (2013), 51–58.

[30] Harris, K. D., and Shepherd, G. M. The neocortical circuit: themes and variations. *Nature neuroscience 18*, 2 (2015), 170–181.

[31] Hofer, S. B., Ko, H., Pichler, B., Vogelstein, J., Ros, H., Zeng, H., Lein, E., Lesica, N. A., and Mrsic-Flogel, T. D. Differential connectivity and response dynamics of excitatory and inhibitory neurons in visual cortex. *Nature neuroscience 14*, 8 (2011), 1045–1052.

[32] Hubel, D. H., and Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of physiology 160*, 1 (1962), 106–154.

[33] Isaacson, J. S., and Scanziani, M. How inhibition shapes cortical activity. *Neuron 72*, 2 (2011), 231–243.

[34] Jin, D. Z., Dragoi, V., Sur, M., and Seung, H. S. Tilt aftereffect and adaptation-induced changes in orientation tuning in visual cortex. *Journal of Neurophysiology 94*, 6 (2005), 4038–4050.

[35] Katzner, S., Busse, L., and Carandini, M. Gabaa inhibition controls response gain in visual cortex. *The Journal of neuroscience 31*, 16 (2011), 5931–5941.

[36] Kerlin, A. M., Andermann, M. L., Berezovskii, V. K., and Reid, R. C. Broadly tuned response properties of diverse inhibitory neuron subtypes in mouse visual cortex. *Neuron 67*, 5 (2010), 858–871.

[37] Ko, H., Hofer, S. B., Pichler, B., Buchanan, K. A., Sjöström, P. J., and Mrsic-Flogel, T. D. Functional specificity of local synaptic connections in neocortical networks. *Nature 473*, 7345 (2011), 87–91.

[38] Kohn, A. Visual adaptation: physiology, mechanisms, and functional benefits. *Journal of neurophysiology 97*, 5 (2007), 3155–3164.

[39] Kohn, A., and Movshon, J. A. Adaptation changes the direction tuning of macaque mt neurons. *Nature neuroscience 7*, 7 (2004), 764–772.

[40] Lee, S.-H., Kwan, A. C., Zhang, S., Phoumthipphavong, V., Flannery, J. G., Masmanidis, S. C., Taniguchi, H., Huang, Z. J., Zhang, F., Boyden, E. S., et al. Activation of specific interneurons improves v1 feature selectivity and visual perception. *Nature 488*, 7411 (2012), 379–383.

[41] Li, L.-y., Xiong, X. R., Ibrahim, L. A., Yuan, W., Tao, H. W., and Zhang, L. I. Differential receptive field properties of parvalbumin and somatostatin inhibitory neurons in mouse auditory cortex. *Cerebral Cortex* (2014), bht417.

[42] Ma, W.-p., Liu, B.-h., Li, Y.-t., Huang, Z. J., Zhang, L. I., and Tao, H. W. Visual representations by cortical somatostatin inhibitory neuronsselective but with weak and delayed responses. *The Journal of Neuroscience 30*, 43 (2010), 14371–14379.

[43] Madison, D., and Nicoll, R. Control of the repetitive discharge of rat ca 1 pyramidal neurones in vitro. *The Journal of Physiology 354* (1984), 319.

[44] Mao, R., Schummers, J., Knoblich, U., Lacey, C. J., Van Wart, A., Cobos, I., Kim, C., Huguenard, J. R., Rubenstein, J. L., and Sur, M. Influence of a subtype of inhibitory interneuron on stimulus-specific responses in visual cortex. *Cerebral Cortex 22*, 3 (2012), 493–508.

[45] Maravall, M., Petersen, R. S., Fairhall, A. L., Arabzadeh, E., and Diamond, M. E. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol 5*, 2 (2007), e19.

[46] Marcus, G., Marblestone, A., and Dean, T. The atoms of neural computation. *Science 346*, 6209 (2014), 551–552.

[47] Marr, D. Vision.

[48] Movshon, J. A., and Lennie, P. Pattern-selective adaptation in visual cortical neurones. *Nature* (1979).

[49] Müller, J. R., Metha, A. B., Krauskopf, J., and Lennie, P. Rapid adaptation in visual cortex to the structure of images. *Science 285*, 5432 (1999), 1405–1408.

[50] Natan, R. G., Briguglio, J. J., Mwilambwe-Tshilobo, L., Jones, S. I., Aizenberg, M., Goldberg, E. M., and Geffen, M. N. Complementary control of sensory adaptation by two types of cortical interneurons. *Elife 4* (2015), e09868.

[51] Nelson, S. B. Temporal interactions in the cat visual system. i. orientation-selective suppression in the visual cortex. *The Journal of neuroscience 11*, 2 (1991), 344–356.

[52] Niell, C. M., and Stryker, M. P. Highly selective receptive fields in mouse visual cortex. *The Journal of Neuroscience 28*, 30 (2008), 7520–7536.

[53] Niell, C. M., and Stryker, M. P. Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron 65*, 4 (2010), 472–479.

[54] Pachitariu, M., Stringer, C., Schröder, S., Dipoppa, M., Rossi, L. F., Carandini, M., and Harris, K. D. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* (2016), 061507.

[55] Packer, A. M., and Yuste, R. Dense, unspecific connectivity of neocortical parvalbumin-positive interneurons: a canonical microcircuit for inhibition? *The Journal of Neuroscience 31*, 37 (2011), 13260–13271.

[56] Pfeffer, C. K., Xue, M., He, M., Huang, Z. J., and Scanziani, M. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature neuroscience 16*, 8 (2013), 1068–1076.

[57] Pnevmatikakis, E. A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron 89*, 2 (2016), 285–299.

[58] Porter, J. T., Johnson, C. K., and Agmon, A. Diverse types of interneurons generate thalamus-evoked feedforward inhibition in the mouse barrel cortex. *The Journal of Neuroscience 21*, 8 (2001), 2699–2710.

[59] Prescott, S. A., and Sejnowski, T. J. Spike-rate coding and spike-time coding are affected oppositely by different adaptation mechanisms. *The Journal of Neuroscience 28*, 50 (2008), 13649–13661.

[60] Reyes, A., Lujan, R., Rozov, A., Burnashev, N., Somogyi, P., and Sakmann, B. Target-cell-specific facilitation and depression in neocortical circuits. *Nature neuroscience 1*, 4 (1998), 279–285.

[61] Runyan, C. A., Schummers, J., Van Wart, A., Kuhlman, S. J., Wilson, N. R., Huang, Z. J., and Sur, M. Response features of parvalbumin-expressing interneurons suggest precise roles for subtypes of inhibition in visual cortex. *Neuron 67*, 5 (2010), 847–857.

[62] Sanchez-Vives, M. V., Nowak, L. G., and McCormick, D. A. Membrane mechanisms underlying contrast adaptation in cat area 17in vivo. *The Journal of Neuroscience 20*, 11 (2000), 4267–4285.

[63] Schwartz, O., Hsu, A., and Dayan, P. Space and time in visual context. *Nature Reviews Neuroscience 8*, 7 (2007), 522–535.

[64] Shadlen, M. N., and Newsome, W. T. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of neuroscience 18*, 10 (1998), 3870–3896.

[65] Shapley, R., Hawken, M., and Ringach, D. L. Dynamics of orientation selectivity in the primary visual cortex and the importance of cortical inhibition. *Neuron 38*, 5 (2003), 689–699.

[66] Silberberg, G., and Markram, H. Disynaptic inhibition between neocortical pyramidal cells mediated by martinotti cells. *Neuron 53*, 5 (2007), 735–746.

[67] Sohya, K., Kameyama, K., Yanagawa, Y., Obata, K., and Tsumoto, T. Gabaergic neurons are less selective to stimulus orientation than excitatory neurons in layer ii/iii of visual cortex, as revealed by in vivo functional ca2+ imaging in transgenic mice. *The Journal of neuroscience 27*, 8 (2007), 2145–2149.

[68] Teich, A. F., and Qian, N. Learning and adaptation in a recurrent model of v1 orientation selectivity. *Journal of Neurophysiology 89*, 4 (2003), 2086–2100.

[69] Theis, L., Berens, P., Froudarakis, E., Reimer, J., Rosón, M. R., Baden, T., Euler, T., Tolias, A. S., and Bethge, M. Benchmarking spike rate inference in population calcium imaging. *Neuron 90*, 3 (2016), 471–482.

[70] Wilson, N. R., Runyan, C. A., Wang, F. L., and Sur, M. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature 488*, 7411 (2012), 343–348.

[71] Xu, X., Roby, K. D., and Callaway, E. M. Immunochemical characterization of inhibitory mouse cortical neurons: three chemically distinct classes of inhibitory cells. *Journal of Comparative Neurology 518*, 3 (2010), 389–404.