THE UNIVERSITY OF CHICAGO

PERCEPTUAL WEIGHTING OF MOUSE V1 SPIKES REVEALED BY OPTOGENETIC WHITE NOISE STIMULATION

A DISSERTATION SUBMITTED TO

THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES

AND THE PRITZKER SCHOOL OF MEDICINE

IN CANDIDACY FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

COMMITTEE ON NEUROBIOLOGY

BY

JULIAN DAY-COONEY

CHICAGO, ILLINOIS

MARCH 2021

TABLE OF CONTENTS

List of figures……iii

Abstract……iv

Introduction……1

Results……7

Discussion……20

Methods……23

Supplemental Figures……...34

References……..38

LIST OF FIGURES

Figure 1. White noise optogenetic stimulation of V1 during a visual detection task……..6

Figure 2. The optogenetic-behavioral kernel…….……10

Figure 3. Alignment of optogenetic-behavioral kernel with V1 neuronal activity……..13

Figure 4. The optogenetic-behavioral kernel and electrophysiology for a ramping visual stimulus.……...17

Supplemental Figure 1. Optogenetic-behavioral kernels from individual mice...…….34

Supplemental Figure 2. Optogenetic-behavioral kernel requires alignment of stimulation to be applied to cortical neurons representing the visual stimulus…….35

Supplemental Figure 3. Alternative alignments for the optogenetic-behavioral kernels.36

Supplemental Figure 4. Covariation of spike weighting and reaction time……..…….37

Abstract

After photons enter the eye, neurons transduce them into patterns of spikes that propagate through the brain until we visually perceive them. We and other animals do this quickly and reliably, but the specific qualities of the spike trains that lead to perception remain elusive. Using optogenetics to perturb these patterns of spikes as they flow through the visual system offers experimenters the unique opportunity to understand which neurons contribute to perception and when their activity is most critical. In mice trained in a visual detection task, we optogenetically excited inhibitory interneurons in primary visual cortex to test which periods contained the most important spikes for performing the task. We employed white noise optogenetic stimulation—a random Bernoulli process of optogenetic stimulation—that allowed us to attain an unbiased, high-resolution picture of when optogenetic stimulation affected behavior. A reverse correlation analysis was performed on the optogenetic time courses that led to a successful detection of the stimulus, yielding an optogenetic-behavioral kernel. We performed acute electrophysiological recordings to understand how the optogenetic-behavioral kernel relates to patterns of sensory evoked spiking in V1. With this optogenetic-behavioral kernel derived from white noise optogenetic stimulation in the context of two different visual stimuli, we determined the time course of V1 readout is biased toward the earliest stimulus-evoked spikes. These data demonstrate the power of white noise stimulation to uncover the readout weighting of spikes from a given population.

Introduction

How signals originating in our sensory organs are transformed into percepts to guide behavior is one of the central questions in neuroscience. Relating the features of neuronal activity in visual areas to visually-guided behavior has been an important approach for revealing how the activity of neurons is decoded to drive perceptual decisions. A number of metrics have been developed that quantify the covariation of features (e.g. firing rate) with behavioral outcomes. While these correlational approaches such as choice probability (Britten et al. 1996; Nienborg and Cumming 2009; Bosking and Maunsell 2011) and the covariation of neuronal response latency and reaction time (DiCarlo and Maunsell 2005; Lee et al. 2016; Seal et al. 1983) have proved useful, these approaches are not able to determine if variation observed in neural activity is actually driving the variation in the behavior. For example, choice probability has been used to quantify the relationship between different time windows of neural activity and behavior (Cook and Maunsell 2002; Williams et al. 2003), but the observations that significant choice-correlated activity can be found during or even after the choice is made (e.g. Kim and Basso 2008; Cohen and Newsome 2009; Nienborg and Cumming 2009) draws into question causal relationships between neural activity and perceptual decisions.

To overcome these limitations, neuronal activity needs to be directly inserted or removed from the population of neurons putatively involved in the generation of

behavior. This type of causal approach facilitates direct comparisons between neuronal activity and behavioral outcomes. Causal perturbations in this vein have been applied to the activity of visual cortical neurons to test which neurons most contribute to a behavior and when. Lesion studies (T. T. Liu and Behrmann 2017), electrical microstimulation (Mark H. Histed, Ni, and Maunsell 2013; Seidemann, Zohary, and Newsome 1998), pharmacological (Schiller, Sandell, and Maunsell 1986), and chemogenetic (Hamm and Yuste 2016) approaches have all been used to better understand which neurons contribute to specific perceptions and behaviors. However, these methods each have limitations in either temporal precision or specificity in targeting populations of neurons in the area of interest. Optogenetics offers the ability to perturb spiking with millisecond resolution and in genetically-defined cell classes (Bernstein and Boyden 2011), yet most optogenetic studies in the context of a behavioral task probe circuits for hundreds of milliseconds to seconds in order to generate a robust and reliable behavioral effect. Practically speaking, it is difficult to probe different epochs that are very brief (<20 ms) due to the number of measurements that need to be collected across behavioral sessions. Combining data across behavioral sessions can be distorted by daily fluctuations in behavioral performance, and more epochs requires more days of behavioral data to be combined. Estimating the effect of optogenetic perturbation on behavioral performance could take hundreds of both stimulated and un-stimulated trials to yield a precise measurement of effect size. However, the amount of trials that mice can perform in a single session is limited to a few hundred trials so only one or two timepoints can be perturbed per session and have sufficient statistical power. Other

animal models, such as non-human primates, can perform more trials within individual sessions, but these models are much less genetically tractable and optogenetic perturbation with cell-type specificity is not as developed as is the mouse. To overcome this problem of having to rely on many days of stable behavioral performance, experimenters have reduced the number of perturbed periods of neural activity to gain more precision while losing temporal resolution. As a result, more is known about which populations of neurons contribute to perceptual behaviors than when spikes from these populations are readout and used. Knowing the periods of spiking that are used to perform a behavior will shed light on how neural activity propagates across populations of neurons and will ultimately lead to better-constrained models of perceptual decision making.

Evidence from a tactile discrimination task in mice demonstrated that optogenetic inhibition of primary somatosensory cortex impaired behavior only if it was delivered during the stimulus presentation period and not the behavioral response period (Guo et al. 2014). This dissociated the timing of behaviorally relevant spikes as being in certain areas at specific periods. However, the duration of inhibition spanned over a second, and evidence from a reaction time version of this task demonstrated that mice can behaviorally respond to such a stimulus in less than 300 ms (Sachidhanandam et al. 2013). Studies that applied optogenetic inhibition to primary visual cortex during visual detection tasks observed that the first 80-100 ms of stimulus-evoked activity seemed to be preferentially used to correctly detect the stimulus (Resulaj et al. 2018). Despite the

brevity of these perturbations, it has been proposed that behaviorally-relevant neuronal activity may be in the sub-ten millisecond range (VanRullen and Thorpe 2002; Shriki, Kohn, and Shamir 2012; Kirchner and Thorpe 2006; H. Liu et al. 2009; Tchumatchenko et al. 2011). Resulaj and colleagues binned their inhibition into 40 ms increments, but it remains to be seen whether or not visual cortical activity was read out at an even briefer temporal window. For example, Histed and Maunsell found that mice could detect the optogenetic activation of excitatory neurons in V1 with as brief as a 1 ms pulse (M. H. Histed and Maunsell 2014), implying that readout of V1 activity could operate at much faster timescales than previously thought.

Sampling neuronal activity at a millisecond resolution and getting a behavioral readout of the efficacy of different timesteps requires pooling data across many days, but logistical constraints make collecting such data nontrivial. Although opsins can operate at a millisecond resolution, the spike failure rate increases dramatically and more trials become necessary (Mattis et al. 2011). In addition, even small changes in behavioral performance and opsin expression can make comparing effect sizes across days difficult. To accommodate this variation, we used a powerful optogenetic stimulation protocol that uses randomly delivered white noise pulse trains on each trial. With each period in the trial receiving optogenetic inhibition, at random, the effective periods can be estimated across many trials. However, unlike stimulating different periods on different days, the fluctuations in behavior and opsin efficacy across days will be shared

across all periods of stimulation, thus giving an estimate of relative temporal efficacy of perturbation with minimal bias.

Using white noise patterns of stimulation have been used to probe the dynamics of biological systems for over half a century (Marmarelis and Marmarelis 1978; Sakai, Naka, and Korenberg 1988). It has been a highly efficient tool for mapping spatiotemporal receptive fields of individual neurons (Eggermont, Johannesma, and Aertsen 1983; Schwartz et al. 2006; Neri and Heeger 2002) as well as determining aspects of a stimulus that are preferentially used to perform a task by using a behavioral response-triggered average (Okazawa et al. 2018; Nienborg and Cumming 2009). The use of white noise optogenetic stimulation in combination with behavior has been applied to invertebrates as they engage in naturalistic behaviors (Hernandez-Nunez et al. 2015; Porto et al. 2019). Both of these studies used behavioral state changes as the readout of the effect of their perturbation. Because the timescales of these chemosensory-driven behaviors were in the hundreds of milliseconds to seconds, the white noise stimuli used were low-pass filtered at 2 Hz (Porto et al. 2019) and 4 Hz (Hernandez-Nunez et al. 2015) because higher temporal resolution was not necessary. However, using the right opsin and behavioral readout, white noise optogenetics has the potential to characterize dynamics with orders of magnitude faster precision.

White noise stimulation can probe a system at high temporal resolution over long time scales. When applied to behavioral effects, results from different sessions can be

normalized and combined across days and animals, creating a precise picture of how a stimulus affects behavior. We show that a white noise optogenetic stimulus applied to the primary visual cortex (V1) of mice during a visual detection task yields an opto-behavioral kernel that represents the temporal weighting of V1 spiking.

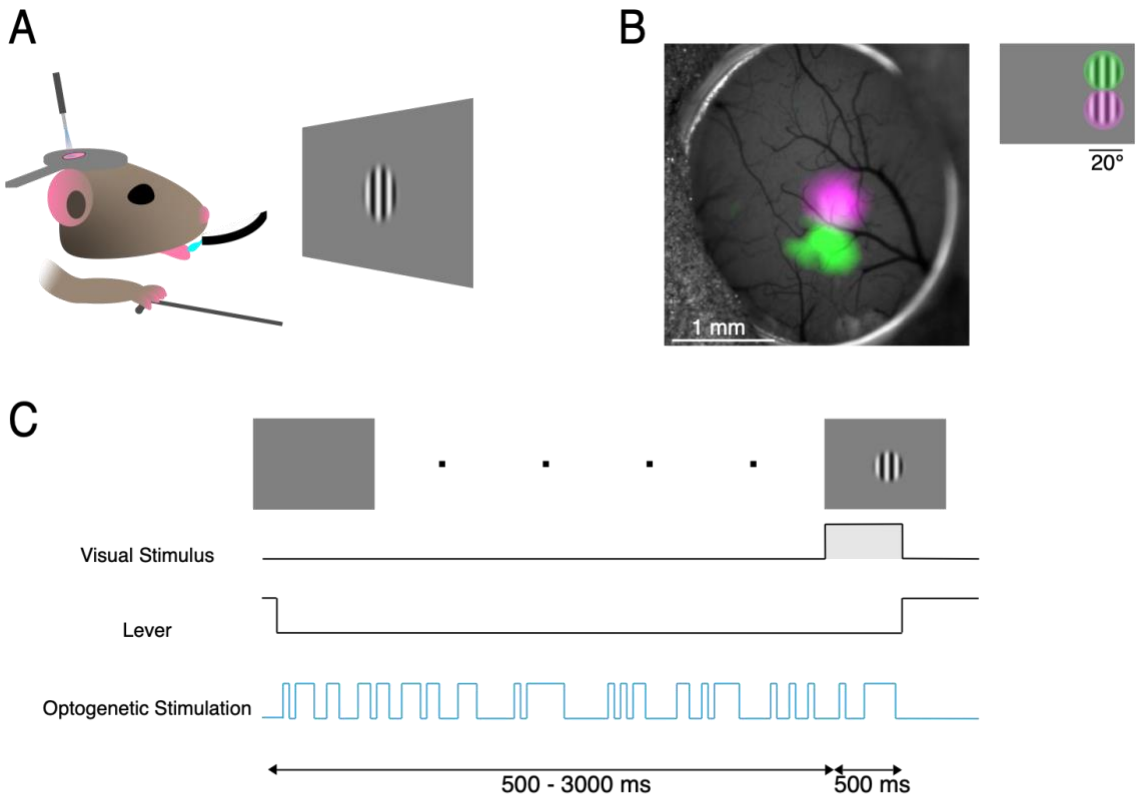

Figure 1. White noise optogenetic stimulation of V1 during a visual detection task. **A**) Cartoon of the behavioral task apparatus. Water scheduled mice earned drops of water sweetened with saccharin delivered from a syringe pump. Mice earned drops by releasing the lever positioned under their forepaw within a 700 ms response window following the onset of a Gabor stimulus on a screen in front of them. **B**) Intrinsic signals through an example cranial window. Intrinsic signals were evoked using 20 degree diameter drifting Gabor stimuli. Two of the five retinotopic locations (corresponding stimuli on the schematic of the screen to the right) are shown for clarity. **C**) Task design with example white noise optogenetic time course. Mice depressed a lever to initiate a trial, and after a 500-3000 ms prestimulus period the Gabor stimulus appeared for 500 ms and the mice had 700 ms to respond in order to receive a liquid reward. If the trial is a stimulation trial, a trial-unique white noise optogenetic pulse train was delivered to V1 during the prestimulus and stimulus periods, exciting parvalbumin-expressing interneurons and thereby delivering fast and local inhibition to V1.

Results

White Noise Optogenetic Stimulation Affects Behavior

We took advantage of the parvalbumin(PV)-expressing subclass of interneurons since they exert fast and potent inhibition on local pyramidal neurons (Packer and Yuste 2011). Delivery of fast optogenetic excitation to PV interneurons has been shone to potently inhibit V1 neurons in the mouse and impair visual detection of stimuli (Glickfeld, Histed, and Maunsell 2013; Cone et al. 2019, 2020). PV-Cre mice were implanted with a head post and a cranial window over V1 of one hemisphere. Mice were then trained to perform a visual detection task in which they had to lift their forepaw to release a lever rapidly after a Gabor stimulus appeared on the video display in front of them (Figure 1A; see Methods). The positioning of the window was confirmed by retinotopically mapping V1 using intrinsic signal imaging (Figure 1B). After the animal reached stable behavioral performance detecting a small Gabor (SD <10°, median: 61 days, range: 32-74 days), we injected a Cre-dependent channelrhodopsin-2/tdTomato AAV virus into the part of V1 corresponding to the location of the visual stimulus. After allowing 3 weeks for viral expression, an optic fiber was mounted above the cranial window directed at the area of

viral expression to deliver consistent optogenetic stimulation in daily sessions over a period of several weeks.

Once behavioral performance had stabilized (>60% hit rate for >300 trials a day), white noise optogenetic stimulation was delivered to V1 throughout one third to one half of trials, using a power that was set for each animal to achieve a modest decrease in behavioral performance.

Optogenetic-Behavioral Kernel

Behavioral outcomes (hits or misses) were used to construct an optogenetic-behavior kernel (OBK; see Methods). The OBK is essentially the average optogenetic V1 stimulus preceding a successful detection. Figure 2A shows the OBK derived from one animal's performance in a single testing session. A power of zero corresponds to optogenetic stimulation of V1 having no net effect on the animal's success in detecting the visual stimulus. The negative peak 52 ms after the appearance of the visual stimulus indicates that reduced power around this time was associated with a greater chance of detecting that stimulus. This negative deflection corresponds to less activation of PV inhibitory neurons, which would remove inhibition from primary neurons in V1. Optogenetic stimulation 50 ms earlier or later than this peak was far less effective, although the stimulus remained on until the animal responded. The brief OBK indicates that optogenetic stimulation did not have appreciable persistent effects on

neuronal circuits in V1 or other structures contributing to the behavior. The critical period ended hundreds of milliseconds before the animal's lever release (Figure 2B).

Observing a reliable OBK within an individual session depends on many factors, including the number of trials, the ratio of hits to misses, false alarm rate and efficacy of the optogenetic stimulus. Significant OBK were not observed in all individual sessions, but trials from different sessions and animals can be combined. Figure 2C-D shows the combined OBK from 216 sessions from 10 mice (individual OBKs are plotted in Supplemental Figure 1). The resulting OBK is only somewhat broader than the OBK from a single session, suggesting a remarkable consistency in the timing of effects across animals and sessions. The peak of the average OBK occurs 52 ms after the onset of the visual stimulus and there were no obvious effects of stimulation before 0 ms and after 120 ms post-visual stimulus.

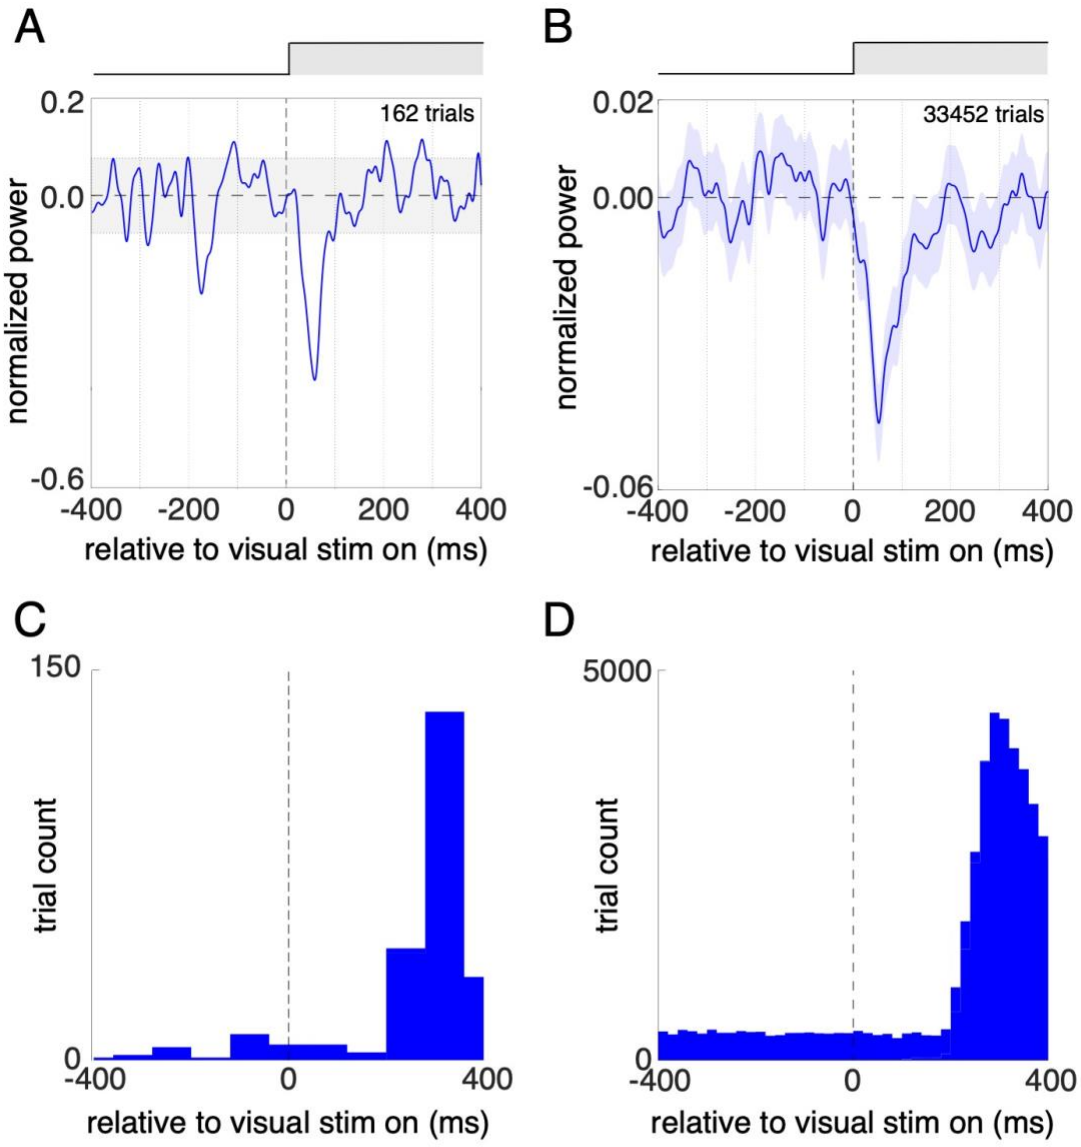

Figure 2. The Optogenetic-Behavioral Kernel. **A**) The optogenetic-behavioral kernel from an example session. The optogenetic time courses that allowed for correct detection of the visual stimulus were aligned to the onset of the stimulus and averaged. The shaded area are the 95% confidence intervals expected from chance. **B**) Same as in (A) but across 216 sessions from 10 animals total (14,612 hit trials, 15,811 miss trials). The error bars are 95% bootstrapped confidence intervals. The peak of the kernel occurs at 52 ms after stimulus onset. **C**) The reaction time distribution from the example day in (A). **D**) the reaction time distribution of all responses used in (B). The median reaction time is 350 ms.

The OBKs in Figure 2 were constructed using optogenetic stimuli that had been aligned to the onset of the visual stimuli. We did this because we expected the V1 neuronal response (and therefore the V1 behavioral contribution) to be aligned with the visual stimuli. Other alignments are possible. In particular, an OBK can be calculated after aligning optogenetic stimuli from hit trials with the response time for each trial. A response-time-aligned OBK would be an obvious choice if optogenetic perturbation were delivered to motoneurons. Although a response-time-aligned V1 OBK could conceivably reveal more temporal precision than the visual-stimulus-aligned OBKs in Figure 2, that was not the case (Supplemental Figure 3). Other alignments, such as stretching time on individual trials to align both stimulus onset and response time across trials also did not produce OBKs that were more limited in time (Supplemental Figure 3). This suggests that the spikes in V1 that contribute most of the detection of a visual stimulus are those that occur shortly after the onset of that stimulus, and that most of the variability in response time occurs in downstream brain structures that lie between V1 and the motoneurons that generate the response (Lee et al. 2016).

If a V1 OBK was constructed after aligning the optogenetic stimulus to lever releases on trials with a false alarm, no significant peaks were observed (Supplemental Figure 3). This suggests that false alarms were not driven by fluctuations in V1 activity that led the animals to perceive a fictive visual stimulus, but rather by other brain regions. This result is consistent with previous behavioral data showing that mice cannot detect activation of their V1 PV neurons (Cone et al. 2019, 2020). Control measurements

made with optogenetic stimulation that was misaligned with the opsin-expressing neurons (which moves the inhibitory perturbation away from the patch of cortex the mouse is using to detect the stimulus) showed that the behavioral effects did not arise from cortical heating, direct retina stimulation by the optogenetic stimulus, or non-specific behavioral effects (Supplemental Figure 2).

The peak of the opto-behavioral kernel occurs at 52 ms whereas the median response times taken from the same sessions is 350 ms (Figure 2A, C). This demonstrates that the width of the opto-behavioral kernel is not simply spanning from the visual stimulus onset to the motor response, but rather early stimulus-evoked activity is preferentially used to perform this reaction time task. This early activity seems to be sufficient to drive the perceptual decision, rendering all activity moot after about 100 ms. However, to establish what aspects of the neuronal responses underlie this preferential weighting, we performed acute electrophysiology outside the context of the task.

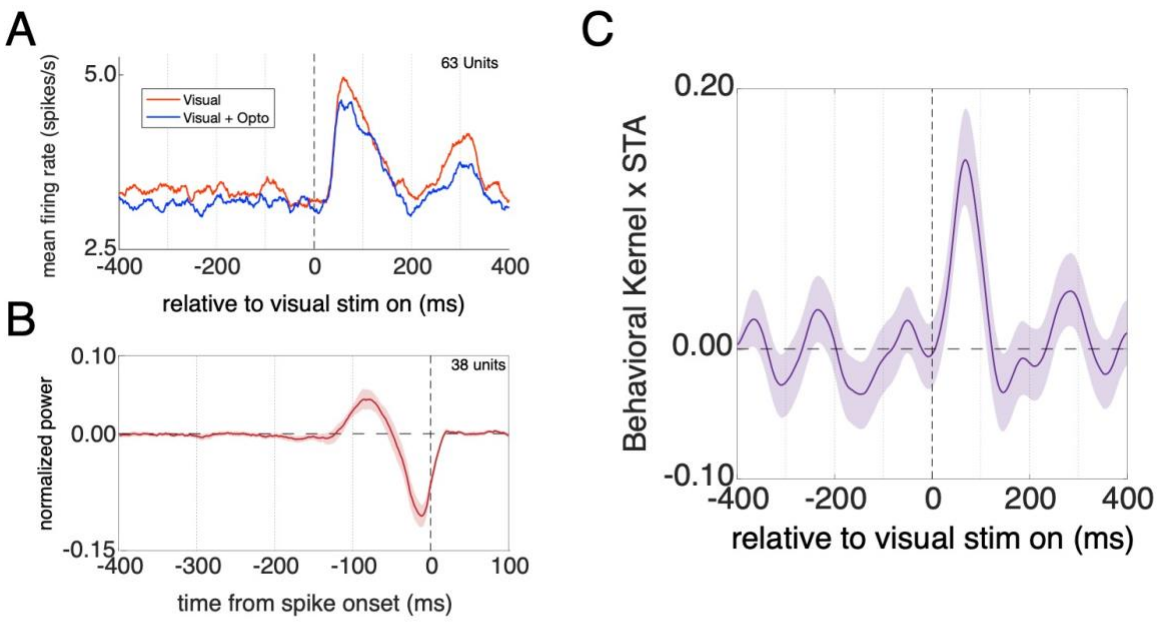

Figure 3. Alignment of optogenetic-behavioral kernel with V1 neuronal activity. **A**) Peristimulus time histograms of V1 neurons (63 units) presented a Gabor stimulus only (red trace) or a Gabor stimulus paired with optogenetic white noise stimulation (blue trace). The peak response during the unstimulated condition occurs at 60 ms. **B**) The spike-triggered average of the optogenetic time courses. **C**) The convolution of the spike-triggered average and the optogenetic-behavioral kernel. This convolution captures the effective periods of optogenetic stimulation of V1 after taking into account the delays inherent to the perturbation. The peak of the convolution occurs at 68 ms.

Comparing Opto-Behavioral Kernel and Neuronal Activity in V1

The early peak in the V1 OBK (Figure 2A, C) suggests that the earliest portion of the V1 response is critical for detecting the visual stimulus. To examine this directly, we made acute recordings from V1 using multielectrode arrays in awake mice expressing ChR2 in PV interneurons. On half the trials a visual stimulus was presented alone, while on the other half white noise optogenetic stimulation was presented before, during and after the visual stimulus. Because we recorded simultaneously from many neurons, the visual stimulus was not optimized for most units. Overall, 38% of units (24 of 63) were significantly excited by visual stimuli relative to baseline, and 3% (2 of 63) were

significantly inhibited (at least $p < 0.05$, Wilcoxon signed-rank test, median response = 0.6 spikes/s, IQR 0.0 -1.3; $p < 10^{-6}$, Wilcoxon signed-rank test).

Using an optogenetic power that was consistent with that used in the behavioral studies (mean = 0.25 mW; range = 0.01 - 2.00 mW), we found that optogenetic stimulation of V1 PV neurons caused only a slight reduction in V1 spiking (Figure 3A, median = -0.14 spikes/s, IQR -0.42 - 0.10; $p < 0.05$, Wilcoxon signed-rank test). Thirty-three units were significantly modulated by optogenetic stimulation. Most (67%, 22/33) were inhibited by optogenetic stimulation, but some were excited (34%, 11/33). Overall, these data show that whole-trial white noise optogenetic inhibition, when averaged across many trials, scales both baseline and visual stimulus evoked firing rates.

The peak of the OBK at 52 ms seems to align with the absolute earliest V1 stimulus-evoked spiking. However, the OBK should not be compared directly with the V1 PSTH. The effects of the optogenetic stimulus on V1 spiking will be delayed by the kinetics of the opsin ChR2 offset time constant of 11 ms (Mattis et al. 2011), the latency of PV cell spiking, and transmission and synaptic delays in the delivery of inhibition to V1 neurons and postsynaptic response times (PV to pyramidal neuron decay constant is also about 12 ms (Packer and Yuste 2011). There will also be polysynaptic effects that will unfold over a longer time course.

Fortunately, the white noise optogenetic stimulation allows us to estimate all the relevant delays. A spike-triggered average (STA) of a white noise stimulus provides the impulse response of a neuron (Bryant and Segundo 1976). The aggregated STA of V1 neurons provides the overall response of V1 spiking to an impulse of optogenetic PV stimulation. The average STA for V1 units is plotted in Figure 3B. The population STA shows that spikes in V1 (at time 0) tend to occur when optogenetic PV has dropped about 12 ms earlier). There appears to be a rebound effect such that spiking is more likely if optogenetic PV stimulation had been increased about 80 ms earlier. Optogenetic white noise stimulation occurring after a spike (positive time) should have no influence on that spike, so the STA is expected to have a normalized power of 0.50 from 0 ms on. It is worth noting that a non-zero average power after the spike occurs is most likely an artifact of the minimum 25 ms pulse width used for the optogenetic stimulation, as it is impossible that stimulation delivered after the spike occurs affects the generation of the spike.

The STA reveals the average delays between optogenetic stimulation delivery and changes in spiking (Figure 3B). In particular, the timing of the negative peak of the STA indicates that most of the effects of optogenetic PV stimulation on overall V1 spiking are delayed by 12 ms. That is, the OBK kernel peak at 52 ms will have its greatest effect at 52+12 ms. A complete estimate of the effects of the OBK on V1 spikes can be obtained by convolving the STA (the impulse response function) with the OBK. This convolution captures all the delayed responses (positive and negative) to optogenetic stimulation.

That convolution is shown in Figure 3C, superimposed on a scaled version of the unstimulated PSTH from Figure 3A.

As expected, this convolution peaks later than the OBK (68 ms), a time that corresponds closely with the peak in the PSTH. However, the convolution returns to zero while the spike rate is still well above baseline, indicating that it is the earlier half of the transient response that contains the spikes that contribute to the behavioral response. With enough trials and statistical power, later spikes may show a non-zero contribution to the behavior so these periods may still contain behaviorally-relevant spiking, but on average these spikes are given far less weight than the initial stimulus-evoked spikes.

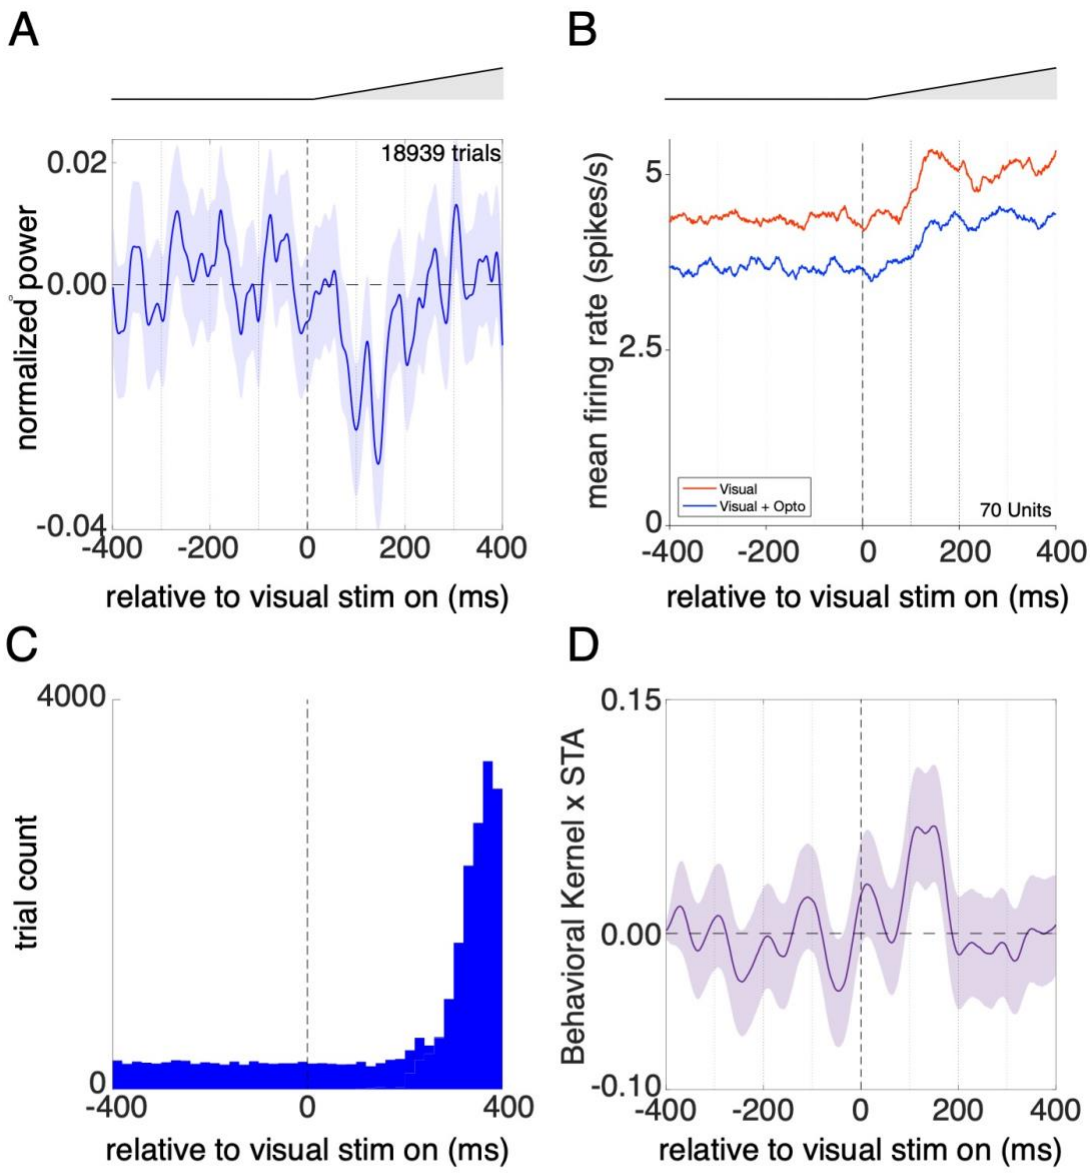

Figure 4. The optogenetic-behavioral kernel and electrophysiology for a ramping visual stimulus. **A**) The optogenetic-behavioral kernel produced from a Gabor stimulus that ramped up from 0 to 100% contrast over 500 ms (6 mice, 143 sessions, 8,984 hit trials, 10,076 miss trials). The error bars are 95% bootstrapped confidence intervals. The peak response occurs at 145 ms. **B**) The peristimulus time histograms for the ramping stimulus. The blue trace were trials where optogenetic white noise was delivered and the red trace was without any stimulation. The peak response of the unstimulated response occurs at 142 ms. **C**) The reaction time distribution from the same trials used to generate the kernel in (A). The median of the reaction time distribution was 413 ms. **D**) The convolution of the optogenetic-behavioral kernel (shown in (A)) with the STA (shown in Figure 3B). The peak of the convolution occurs at 149 ms.

Optogenetic-Behavioral Kernel for Ramping Stimuli

These results suggest that optogenetic white noise perturbations can be a powerful tool for exploring how specific neuronal populations contribute to specific behaviors. For example, OBKs should differ in characteristic ways between various neuronal populations (e.g., sensory versus motor neurons) or different tasks (e.g., respond to vision stimuli versus respond to auditory stimuli). To examine whether differences in OBKs can be readily detected, we modified the visual stimulus detection task to use stimuli that cause visual information to be protracted over longer periods as opposed to a sudden burst of activity.

We encouraged mice to use more time to integrate the stimulus by presenting a Gabor stimulus that had contrast that ramped from 0 to 100% over 500 ms, while keeping all other task parameters the same (Figure 4). The ramping stimulus was chosen because it is known to decrease the prominence of the initial transient burst in firing rate observed when stimuli immediately appear, which we observed in our data (compare earliest stimulus-evoked spiking in 4D and 3A) and in previous characterizations of the contrast response functions of V1 neurons (Albrecht et al. 2002). The median reaction time for the ramping stimulus was 63 ms longer than for the step onset (413 ms versus 350 ms; Figure 4B). Correspondingly, the peaks of the V1 PV neuron OBK (Figure 4A) and the V1 PSTH (Figure 4C) were also shifted about 100 ms later relative to the onset of the stimulus (step kernel peak: 52 ms, step PSTH peak: 60 ms; ramp kernel peak:

145 ms, ramp PSTH peak: 142 ms). When comparing the delay corrected cross-correlogram (Figure 4D) to the neuronal responses, the peak and breadth of the corrected opto-behavioral kernel corresponds to the earliest stimulus-evoked neuronal activity. As was the case with the stepped onset visual stimulus, the opto-behavioral kernel drops back to 0 before the peak of spiking activity and hundreds of milliseconds before the animals begin to respond to the stimulus. These data reinforce the notion that the earliest stimulus-evoked spiking in V1 are preferentially used to perform a visual detection task.

Because the ramping stimulus required more time to detect the stimulus, we tested if faster reaction times were related to briefer windows of integrating V1 spikes. When a median split was performed of all reaction times, fast and slow RT trials were used to generate separate OBKs (Supplemental Figure 4). For the ramp data, we observed a later shift in the peak of the OBK generated from the later RTs compared to faster RTs (100 ms versus 146 ms). We did not observe such a difference for the step stimulus (Supplemental Figure 4, left column), but there is evidence of a longer tail of the kernel in the slow RT kernel condition. These data point toward the importance of variation in the integration of V1 spikes in driving differences in RT.

Discussion

We took advantage of the fast and potent inhibition PV interneurons exert on local neurons (Packer and Yuste 2011) to test which V1 spikes are most important for driving detection of a visual stimulus. We found that spikes occurring within the first 60 ms of stimulus-evoked V1 activity affect behavioral performance most as indicated by the peak in the STA-convolved OBK. This finding is consistent with previous experimental work tested at coarser temporal resolutions (Resulaj et al. 2018). Unlike previous work, our convolution of the STA with the OBK took into account of opsin and synaptic delays (Figure 3D, 4C) yielding an estimate of not only the most effective moments to deliver optogenetic stimulation, but also the most effective moments to inhibit putative V1 pyramidal neurons.

Using white noise optogenetic stimulation, the resulting opto-behavioral kernels represent a high-resolution estimate of the dynamics of the readout of V1 activity. The outsized importance of the earliest stimulus-evoked spikes in V1 is not a new idea (VanRullen and Thorpe 2002; Berens et al. 2012; Tchumatchenko et al. 2011; Wilson et al. 2017; Resulaj et al. 2018), but we have demonstrated causally that these early spikes, if inhibited, greatly impair visual awareness of the stimulus. The OBK drops to zero while the largest peak (the transient) in firing rates is still occurring. These transient bursts in activity have been hypothesized to contain the information needed for animals

to decode the stimulus (Müller et al. 2001), but these results suggest that only a subset of the transient spiking period can be used to detect the stimulus.

We determined that the effect on spiking that inhibition of pre-stimulus spikes did not affect behavior as indicated by the kernel beginning to rise after 0 ms. This observation also reinforces that our inhibition did not persist for hundreds of milliseconds, but was relatively limited to the periods of illumination. In addition to the decay constants of ChR2 and the IPSCs of PV interneurons each being less than the pulse duration of our stimulation (Lin 2012; Packer and Yuste 2011), the convolution of the STA of optogenetic stimulation directly addresses this possibility. The delay between the release from inhibition and the average spike timing was 12 ms, indicating that the influence of stimulation on V1 spiking was short-lived and that our estimate of large temporal weighting of earlier spikes was accurate. The convolution between the STA and the OBK further indicated that the 60 ms period following the onset of stimulus-evoked activity was the most efficacious in causing a perceptual deficit.

It is important to note that the 25 ms minimum pulse widths of our optogenetic time courses artificially broadens the OBK and the STA. Deconvolution methods could be of value to extract an even more precise estimate of the effects optogenetic stimulation on V1 spiking and behavior. In addition, future experiments could reduce this broadening by using faster opsins that can accommodate briefer pulse widths.

Using a relatively small proportion of spikes in V1 to perform a visual detection task is in accordance with the speed of some visually-guided behaviors being as fast at 120 ms (Kirchner and Thorpe 2006). Other studies have pointed toward the earliest spikes being preferentially used in a variety of tasks and areas (Oram and Perrett 1992; Müller et al. 2001; Chen, Geisler, and Seidemann 2008; Shriki, Kohn, and Shamir 2012). Shriki, et al., found that using only the relative first spike latencies of a small number of visual cortical neurons was sufficient to decode the orientation of the visual stimulus. Alternative theories have proposed that feedback mechanisms to early visual areas may be necessary for perception (Vugt et al. 2018). Our data show a brief time window during which feedback from higher areas can affect perceptually-relevant activity, but we cannot exclude the role of feedback in our task.

White noise optogenetic stimulation is a powerful approach for increasing temporal resolution when using optogenetics in combination with behavior. Although used at longer timescales (Hernandez-Nunez et al. 2015; Porto et al. 2019), we've demonstrated that it can be applied to a system with a time resolution orders or magnitude faster while also taking into account the kinetics of the circuitry and opsins. In addition, the opto-behavioral kernels can easily be combined across sessions and animals, despite the varying levels of performance and opsin expression. Optogenetic-behavioral kernels allow experimenters to probe the windows of neuronal activity that contribute to behavior in an unbiased fashion.

Methods

Animals

All animal procedures followed NIH guidelines and were approved by the Institutional Animal Care and Use Committee of the University of Chicago. The transgenic mice used in this study were heterozygous for Cre-driver expressed in Parvalbumin-expressing cells.

Cranial Window Surgery

Three-month-old mice were implanted with a cranial window over their left primary visual cortex (V1). They were anesthetized with isoflurane gas in 100% O_2 (induction, 3%; maintenance 1-1.5%) and i.p. ketamine (40 mg/kg) and xylazine (2 mg/kg). Once anesthetized, they were put in ear bars and their corneas were protected with silicone gel. Body temperature was maintained with a regulated heating pad. Using aseptic technique, the scalp was retracted and a custom aluminum headpost was placed above the skull and attached to the skull with acrylic cement.

A 3 mm craniotomy was drilled over V1 (3.1 mm lateral to lambda) using a 0.25 mm diameter spherical drill bit. A sterile glass cranial window (3 mm diameter, 0.8 mm thick)

and then placed over the craniotomy and sealed to the head post is acrylic cement. Mice were given postoperative analgesics (buprenorphine, 0.1 mg/kg and meloxicam, 2 mg/kg).

Retinotopic Mapping with Optical Imaging of Intrinsic Signals

After recovery from the surgery the retinotopy of V1 was measured through the cranial windows using the autofluorescence signals in response to visual stimulation. This mapping was used to precisely target viral injections in V1. Drifting Gabors were displayed on a screen in front of the head-fixed mice, and the fluorescence changes were measured using widefield imaging through the cranial window. Autofluorescence was excited with blue light (470 ± 40 nm, Chroma) and the fluorescence was collected with a green long-pass filter (500-nm cutoff) and a 1.0x air objective (Zeiss; StereoDiscovery V8 microscope; ~0.11 NA). Fluorescence was captured with a CCD camera (AxioCam MRm, Zeiss; 460 x 344 pixels; 4 x 3mm field of view). The Gabors (10° SD, 0.1 cycles/deg, 30 deg/s) were centered on two locations along the vertical meridian and 3 locations in the right visual hemifield (25° azimuth) at various elevations (±10° elevation) and the averaged changes in fluorescence for each visual location were used to determine the retinotopy of V1.

Viral Injections and Optic Fiber Placement

For virus injections, mice were anesthetized with isoflurane (1-2% in 100% O_2) and held in a stereotaxic apparatus. Using aseptic technique, the cranial window was removed.

Beveled glass pipettes with a 20 µm opening were used to inject virus (AAV9-Flex-ChR2-tdTomato; 10^{11} viral particles; Penn Vector Core) using surface vasculature in combination with the retinotopic map to target a region of monocular V1 (~20° azimuth, ±10 degrees elevation). Injections were made 250 and 400 µm below the cortical surface (50 nL/min). The craniotomy was sealed using a new sterilized cranial window when the injections were completed.

After allowing 2-3 weeks for the virus to express, the location and extent of expression was confirmed using the fluorescence of the tdTomato fluorophore. An optic fiber (400-µm diameter; 0.48 NA; Doric Lenses) was then mounted using acrylic cement above the cranial window, directed at the viral expression and illuminating ~1 mm^2 of cortex through the window.

Behavioral Task

During behavioral training and testing, mice were given scheduled access to water. Water scheduling began no earlier than 3 weeks following surgery. For training, the mice were placed in a custom sled that allowed them to be comfortably head-fixed in front of a gamma corrected 54.6 cm monitor 8.4 cm away. Mice were trained to depress the lever to initiate a trial. This triggered a neutral tone to indicate the trial start and a pre-stimulus period which had a duration drawn from a uniform distribution between 500 and 3000 ms. During the pre-stimulus period, the monitor displayed a uniform gray after which a full contrast, isoluminant Gabor stimulus (5° SD, 0.1 cycles/deg) appeared.

Mice then had to release the lever within a response window of 200 to 600 or 700 ms after stimulus onset to receive a reward (1.5-4 µL). Early responses were punished with a timeout (1500-2500 ms) before another trial would start. The parameters of the task were slowly adjusted over a period of many weeks until mice responded reliably for hundreds of trials in a day to a Gabor stimulus in the part of visual space corresponding to the location in V1 map where optogenetic stimulation would be delivered.

On trials with ramping contrast, a Gabor (5° S.D., 0.1 cycles/deg) appeared in the same location used for the stepped contrast stimuli, but with contrast increased from 0% to 100% over 500 ms. The response window and stimulus duration were the same between the stepped and ramped conditions for each mouse.

White Noise Optogenetic Stimulation

Once mice reached stable behavioral performance and viral expression, optogenetic stimulation was delivered through the optic fiber using an LED source (455 nm; Thorlabs). A photodiode trigger presented at the beginning of a trial in a corner of the screen was used to trigger optogenetic stimulation, ensuring precise alignment between the optogenetic and visual stimuli. On all trials with optogenetic stimulation, that stimulus ramped up from zero to average mean power over the first 250 ms of stimulation. In preliminary sessions the optogenetic power was adjusted by presenting a constant optogenetic stimulus throughout some trials. These trials were used to find an LED power that produced a reliable decrease in the proportion of correct trials, which

was used as the mean power in subsequent testing. This power varied between mice, presumably owing to differences in the strength and spatial distribution of virus expression, optic fiber alignment, and behavioral capacity (mean power delivered to the fiber optic stub: 0.25 mW, range 0.01 - 2.00 mW).

A binary white noise optogenetic stimulus was generated by randomly assigning zero or maximum (2x mean) power to a series of 25 ms bins with equal probability. Briefer bins could not accommodate the kinetics of the ChR2 (~10 ms decay time; Mattis et al. 2011) as well as the kinetics of the PV-to-pyramidal neuron synapse (~15 ms rise and decay time; Packer and Yuste 2011). The phase of the 25 ms bin sequence was randomly offset to the nearest millisecond on each trial. The resulting optogenetic stimulus produced equal power across all frequencies represented and is, therefore, a quasi-white noise stimulus by definition (Marmarelis and Marmarelis 1978).

The white noise stimulation was delivered on a randomly subset of trials (33-50%). Stimulation was only delivered on days where mice had a low false alarm rate (< ~40%) and a large enough proportion of trials correct (> ~60%). Days were only included in further analysis that showed a decrease in percent of correct trials on simulation trials as compared to unstimulated trials.

Optogenetic Behavioral Kernel

A first-order Weiner kernel was calculated from the optogenetic pulse trains from trials that ended with a hit or a miss, excluding trials that were ignored or ended with a false alarm. Data from different animals were integrated by normalizing the maximum optogenetic power to 1. For the primary analysis, the optogenetic stimulus profiles from hit trials and miss trials were aligned to the onset of the visual stimulus onset, profiles from miss trials were inverted, and then all the profiles were averaged (14,612 hit trials, 15,811 miss trials for the stepped contrast stimulus; 8,984 hit trials, 10,076 miss trials for the ramped contrast stimulus). The resulting optogenetic kernel represents the linear temporal weighting of optogenetic activation of V1 PV-expressing neurons in perturbing behavioral performance.

The OBKs for both the step and ramp visual stimuli were aligned to the reaction time, false alarms, and a fixed interval stimulus-RT stretched alignment (Supplemental Figure 3). These alternative alignments demonstrated that the OBKs could be informative about the relationship between V1 spiking and the timing of the motor response. However, the lack of significant deflections in the false alarm-aligned case shows that the optogenetic stimulation, itself, did not compel the mice to respond (Supplemental Figure 3, right column).

For the control stimulation (Supplemental Figure 2), 3 mice were trained on the stepped contrast and white noise optogenetic stimulation was delivered as described. OBKs were generated from 5 consecutive training days when the visual stimulus was aligned

to the area of V1 being stimulated (25 degrees azimuth, 0 degrees elevation). Then, mice were trained for a few days to detect the same stimulus presented in the opposite hemifield (-25 degrees azimuth, 0 degrees elevation). Once mice were proficiently detecting the stimulus misaligned to the area of V1 being perturbed, white noise optogenetic stimulation was delivered for 10 consecutive days (5 days in one of the mice). From these control stimulus days, the OBK was calculated across all mice (Supplemental Figure 2A) without any exclusion criteria applied. Two of the mice were then re-trained at the original, aligned stimulus location and were stimulated for 5 more consecutive days. The aligned stimulus OBKs were then combined across pre- and post-control days such that the session number was matched between aligned and misaligned stimulation days for each mouse (Supplemental Figure 2B). All control data did not have any exclusion criteria applied.

Electrophysiological recordings

We recorded extracellularly from V1 in awake, passively viewing, head fixed mice using multisite silicon probes (Neuronexus, Inc.; 32-site model 4 × 8-100–200-177). Data were collected from 10 total recording locations across the 5 mice (1 female). One mouse was first used for data collection in behavioral experiments, while the rest were untrained but injected to express opsins following the same procedures. Electrodes were electroplated with a gold solution mixed with carbon nanotubes (Keefer et al. 2008; Ferguson, Boldt, and Redish 2009) to impedances between 200-500 kΩ.

At the start of recording sessions, mice were anesthetized with isoflurane (1.2–2% in 100% O_2), placed in a sled and the headpost mounted as in the behavioral rig. A gamma corrected video display was positioned in the visual hemifield opposite of the recording site (~10 cm viewing distance). The eyes were kept moist with 0.9% saline throughout the session. We visualized ChR2-expressing areas of monocular V1 by imaging tdTomato fluorescence with a fluorescence microscope and camera (Zeiss, Inc). The cranial window was then removed and the electrodes lowered through a slit in the dura. We then positioned an optic fiber above the cortex at a distance comparable to that used during behavioral experiments (1.0-1.5 mm). The craniotomy was then covered with 3% agarose dissolved in aCSF (MilliporeSigma Inc., TOCRIS respectively). The agar was prevented from drying out over the course of the recording session by covering it with silicone oil (60,000 centistokes; MilliporeSigma Inc). Following the recovery period of 1 hour, anesthetic was removed, and we waited at least an additional hour for recovery from anesthesia before recording.

The electrode was advanced to locate responsive units and was allowed to settle for 30 minutes before collecting data. Delivery of visual and optogenetic stimuli and data acquisition was computer controlled. Concurrent visual and optogenetic stimuli were similar to those used during behavioral experiments (stepped contrast stimuli: a Gabor patch with SD= 13°; ramped contrast stimuli: full screen grating) except that visual stimuli were presented for 500 ms. Except when stimuli were presented the video display was held a uniform neutral gray. Optogenetic stimuli were triggered at the start

of each trial using a photodiode positioned on the stimulus display to align the timing of visual and optogenetic stimuli and spikes.

We recorded at least 100 repetitions of each stimulus condition (visual stimulus, visual stimulus + white noise optogenetic stimuli). White noise optogenetic stimuli were binary (zero to fixed power) with (25 ms frames). A different optogenetic stimulus (frame-phase randomized to 1 ms) was presented on each trial. To avoid artifacts at the onset of stimulation, the mean optogenetic power ramped up over 250 ms at the beginning of each trial to a mean power 0.25 mW.

Electrode signals were amplified, bandpass filtered (750 Hz to 7.5 kHz) sampled around threshold crossings (Blackrock, Inc.) and spikes were sorted offline (OfflineSorter, Plexon, Inc.). Data analyses were conducted in MATLAB (MathWorks Inc.) using custom code. We recorded both single and multi-units but did not differentiate between them because our primary interest was how optogenetic manipulations affected the V1 population.

Statistical Analyses

Mice were excluded from analyses that produced less than 8 days of behavioral data while being optogenetically stimulated. Based on previous work from the lab showing that PV neuron stimulation can only produce behavioral deficits (Cone et al. 2019), we also excluded individual sessions that did not produce a decrease in the percent of

correct trials during stimulation as compared to unstimulated trials. While the stimulation power being delivered was modest and did not produce a prominent decrease in spike rate (Figure 3B; Figure 4C), a decrease in performance was indicative of correct alignment between optogenetic stimulation and visually-evoked spiking as well as sufficient viral expression to produce a behavioral effect.

The confidence intervals for the opto-behavioral kernels, the spike-triggered averages, and the cross-correlograms were generated using a bootstrap procedure. In each case, the average of 10,000 draws without replacement were made from the distribution in question, and the 95% confidence intervals were generated from the variance of these distributions.

To produce the OBKs differentiated by fast and slow reaction times (Supplemental Figure 4), a median split was performed on all hit trials on each individual session (step stimulus: 216 sessions, ramp stimulus: 143 sessions). The fast and slow hit kernels was then computed from the different sets of optogenetic time courses before being combined with the inverted miss kernel to produce the resulting fast and slow RT OBKs.

For the electrophysiology, visually responsive units to the stepped contrast stimulus were taken as those with a significant change in the average firing rate ($p < 0.05$; Wilcoxon signed-rank test) during the 50-200 ms after stimulus onset (stimulus period) relative to the average firing rate during the -200 to -50 ms before stimulus onset

(baseline). For the ramped contrast stimulus, the analysis window extended to the end of visual stimulus presentation (50-500 ms), compared to the -500 to -50 ms before visual stimulus onset. For analysis of visual responsivity, only trials without optogenetic stimuli were considered. Units were defined as optogenetically responsive if there was a significant difference ($p < 0.05$; Wilcoxon signed-rank test) in the average firing rate measured during the visual stimulus window (50-500 ms) for trials compared to without optogenetic stimulation.

Supplemental Figures

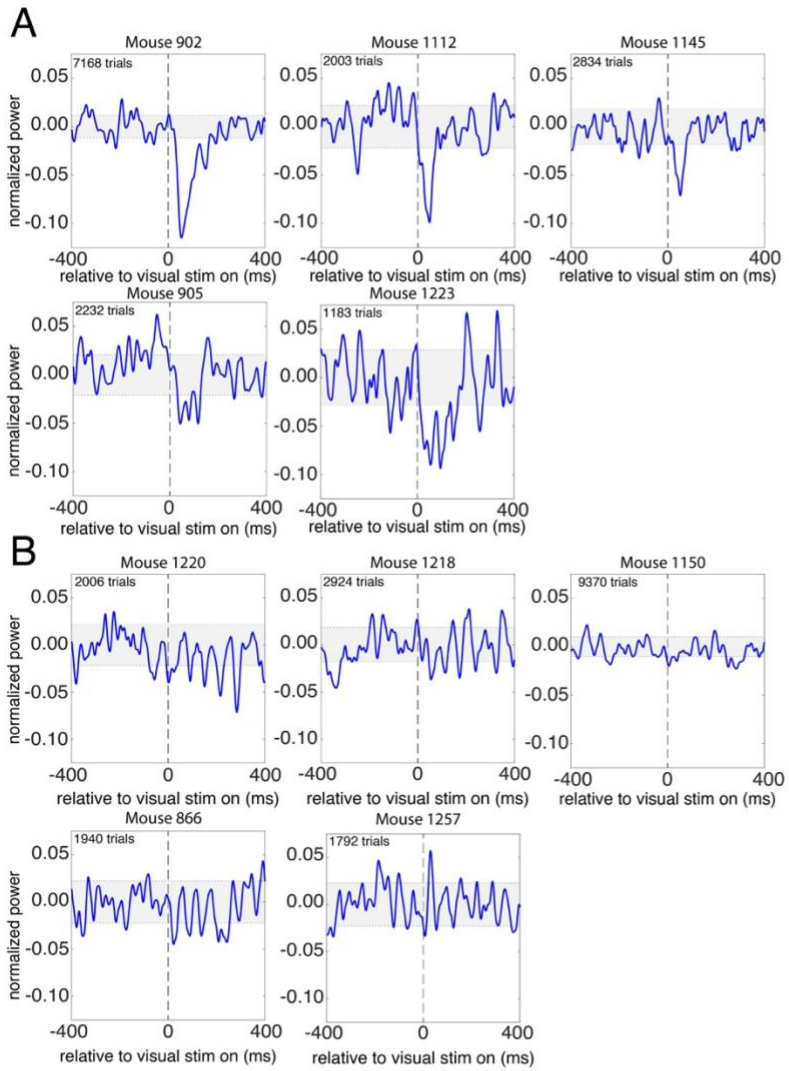

Supplemental Figure 1. Optogenetic-behavioral kernels from individual mice. A) Five mice that showed a reliable kernel that peaked around 50 ms after onset of a stepped contrast Gabor. B) The other five mice showed little or no peak at this time, but showed more negative deflections after the onset of the visual stimulus. The shaded areas are the 95% confidence intervals.

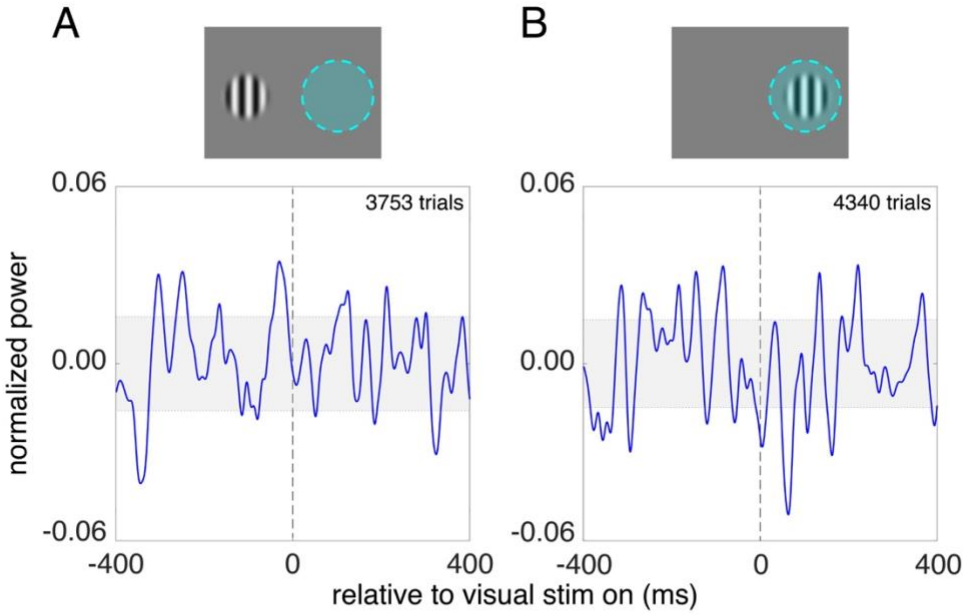

Supplemental Figure 2. Optogenetic-behavioral kernel requires stimulation to be applied to cortical neurons representing the visual stimulus. **A**) OBK resulting from 3 mice that were trained to detect a Gabor stimulus in the hemifield ipsilateral to the V1 being stimulated. On the top is a schematic of the area of visual space being perturbed by optogenetic stimulation (blue circle). If the optogenetic stimulation was misaligned from the area of cortex that is processing the visual stimulus, no kernel appears. **B**) OBK from the same 3 mice taken from sessions before and after the sessions used in (A). The same optogenetic power was task parameters were kept the same, except that the visual stimulus was aligned with the area of V1 being perturbed.

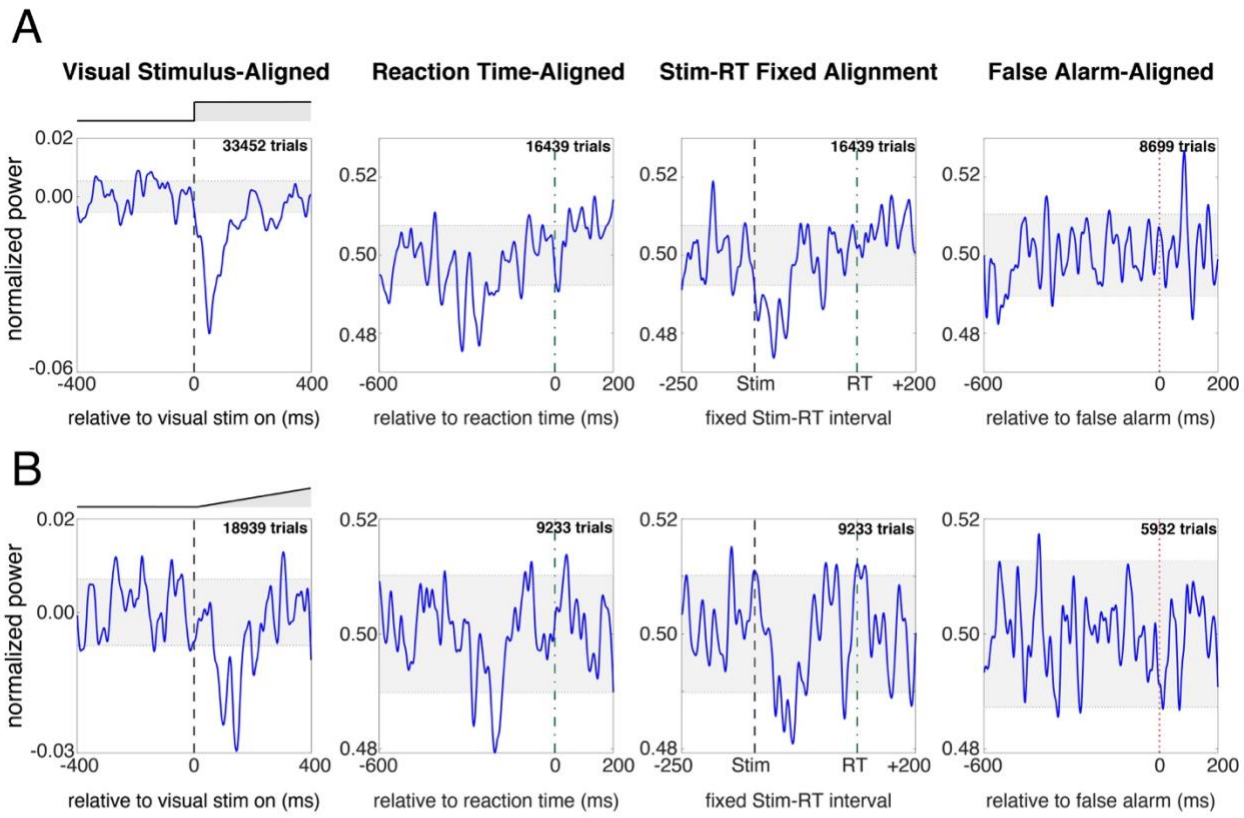

Supplemental Figure 3. Alternative alignments for the optogenetic-behavioral kernels. **A**) OBK for the stepped contrast stimulus. From left to right: aligned the onset of the visual stimulus, aligned to the reaction time, aligned to both the stimulus and the reaction time with stretching or compressing the kernel to a fixed interval, and aligned to false alarms. **B**) The same as in (A) but for the ramped contrast stimulus.

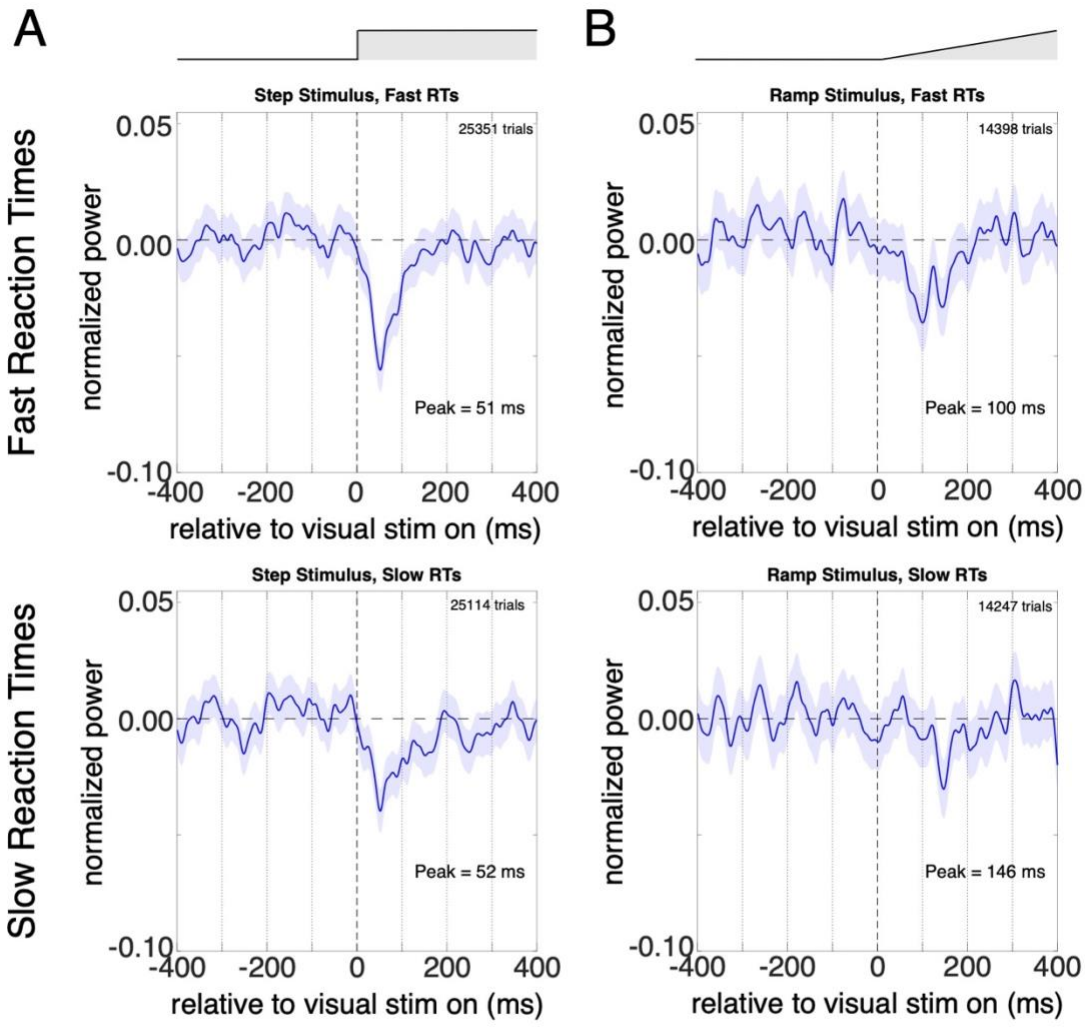

Supplemental Figure 4. Covariation of spike weighting and reaction time. **A**) The OBKs for the stepped contrast visual stimulus. A median split was performed on the reaction times within each session and an OBK was separately calculated from the fast and slow RT trials (top panel: faster RT OBK; bottom panel: slower RT OBK). **B**) The same as in (A) but for the ramping contrast stimulus.

References

Albrecht, Duane G, Wilson S Geisler, Robert A Frazor, and Alison M Crane. 2002. "Visual Cortex Neurons of Monkeys and Cats: Temporal Dynamics of the Contrast Response Function." *J Neurophysiol* 88: 888–913.

Berens, Philipp, Alexander S. Ecker, R. James Cotton, Wei Ji Ma, Matthias Bethge, and Andreas S. Tolias. 2012. "A Fast and Simple Population Code for Orientation in Primate V1." *Journal of Neuroscience* 32 (31): 10618–26. https://doi.org/10.1523/JNEUROSCI.1335-12.2012.

Bernstein, Jacob G., and Edward S. Boyden. 2011. "Optogenetic Tools for Analyzing the Neural Circuits of Behavior." *Trends in Cognitive Sciences* 15 (12): 592–600. https://doi.org/10.1016/j.tics.2011.10.003.Optogenetic.

Bosking, William H, and John H R Maunsell. 2011. "Effects of Stimulus Direction on the Correlation between Behavior and Single Units in Area MT during a Motion Detection Task" 31 (22): 8230–38. https://doi.org/10.1523/JNEUROSCI.0126-11.2011.

Britten, KH, WT Newsome, Michael N. Shadlen, S Celebrini, and JA Movshon. 1996. "A Relationship between Behavoiral Choice and the Visual Responses of Neurons in Macaque MT." *VIsual Neuroscience* 13 (May): 87–100. https://doi.org/10.1017/S095252380000715X.

Bryant, Hugh L, and José P Segundo. 1976. "Spike Initiation by Transmembrane Current: A White-noise Analysis." *The Journal of Physiology* 260 (2): 279–314. https://doi.org/10.1113/jphysiol.1976.sp011516.

Chen, Yuzhi, Wilson S. Geisler, and Eyal Seidemann. 2008. "Optimal Temporal Decoding of Neural Population Responses in a Reaction-Time Visual Detection Task." *Journal of Neurophysiology* 99 (3): 1366–79. https://doi.org/10.1152/jn.00698.2007.

Cohen, Marlene R, and William T Newsome. 2009. "Estimates of the Contribution of Single Neurons to Perception Depend on Timescale and Noise Correlation" 29 (20): 6635–48. https://doi.org/10.1523/JNEUROSCI.5179-08.2009.

Cone, Jackson J., Morgan L. Bade, Nicolas Y. Masse, Elizabeth A. Page, David J. Freedman, and John H.R. Maunsell. 2020. "Mice Preferentially Use Increases in Cerebral Cortex Spiking to Detect Changes in Visual Stimuli." *The Journal of Neuroscience*, no. August: JN-RM-1124-20. https://doi.org/10.1523/jneurosci.1124-20.2020.

Cone, Jackson J., Megan D. Scantlen, Mark H. Histed, and John H.R. Maunsell. 2019. “Different Inhibitory Interneuron Cell Classes Make Distinct Contributions to Visual Contrast Perception.” *ENeuro* 6 (1): 1–12. https://doi.org/10.1523/ENEURO.0337-18.2019.

Cook, Erik P., and John H.R. Maunsell. 2002. “Dynamics of Neuronal Responses in Macaque MT and VIP during Motion Detection.” *Nature Neuroscience* 5 (10): 985–94. https://doi.org/10.1038/nn924.

DiCarlo, James J., and John H.R. Maunsell. 2005. “Using Neuronal Latency to Determine Sensory-Motor Processing Pathways in Reaction Time Tasks.” *Journal of Neurophysiology* 93 (5): 2974–86. https://doi.org/10.1152/jn.00508.2004.

Eggermont, J. J., P. I. M. Johannesma, and A. M. H. J. Aertsen. 1983. “Reverse-Correlation Methods in Auditory Research.” *Quarterly Reviews of Biophysics* 16 (03): 341. https://doi.org/10.1017/s0033583500005126.

Ferguson, John E., Chris Boldt, and A. David Redish. 2009. “Creating Low-Impedance Tetrodes by Electroplating with Additives.” *Sensors and Actuators, A: Physical* 156 (2): 388–93. https://doi.org/10.1016/j.sna.2009.10.001.

Glickfeld, Lindsey L, Mark H Histed, and John H R Maunsell. 2013. “Mouse Primary Visual Cortex Is Used to Detect Both Orientation and Contrast Changes.” *The Journal of Neuroscience* 33 (50): 19416–22. https://doi.org/10.1523/JNEUROSCI.3560-13.2013.

Guo, ZengcaiV, Nuo Li, Daniel Huber, Eran Ophir, Diego Gutnisky, JonathanT Ting, Guoping Feng, and Karel Svoboda. 2014. “Flow of Cortical Activity Underlying a Tactile Decision in Mice.” *Neuron* 81 (1): 179–94. https://doi.org/10.1016/j.neuron.2013.10.020.

Hamm, Jordan P, and Rafael Yuste. 2016. “Somatostatin Interneurons Control a Key Component of Mismatch Negativity in Mouse Visual Cortex.” *Cell Reports* 16: 597–604. https://doi.org/10.1016/j.celrep.2016.06.037.

Hernandez-Nunez, Luis, Jonas Belina, Mason Klein, Guangwei Si, Lindsey Claus, John R. Carlson, and Aravinthan D.T. Samuel. 2015. “Reverse-Correlation Analysis of Navigation Dynamics in Drosophila Larva Using Optogenetics.” *ELife* 4 (MAY): 1–16. https://doi.org/10.7554/eLife.06225.

Histed, M. H., and J. H. R. Maunsell. 2014. “Cortical Neural Populations Can Guide Behavior by Integrating Inputs Linearly, Independent of Synchrony.” *Proceedings of the National Academy of Sciences* 111 (1): E178–87. https://doi.org/10.1073/pnas.1318750111.

Histed, Mark H., Amy M. Ni, and John H R Maunsell. 2013. “Insights into Cortical Mechanisms of Behavior from Microstimulation Experiments.” *Progress in Neurobiology*. https://doi.org/10.1016/j.pneurobio.2012.01.006.

Keefer, Edward W., Barry R. Botterman, Mario I. Romero, Andrew F. Rossi, and Guenter W. Gross. 2008. “Carbon Nanotube Coating Improves Neuronal Recordings.” *Nature Nanotechnology* 3 (7): 434–39. https://doi.org/10.1038/nnano.2008.174.

Kim, Byounghoon, and Michele A. Basso. 2008. “Saccade Target Selection in the Superior Colliculus: A Signal Detection Theory Approach.” *Journal of Neuroscience* 28 (12): 2991–3007. https://doi.org/10.1523/JNEUROSCI.5424-07.2008.

Kirchner, Holle, and Simon J. Thorpe. 2006. “Ultra-Rapid Object Detection with Saccadic Eye Movements: Visual Processing Speed Revisited.” *Vision Research* 46 (11): 1762–76. https://doi.org/10.1016/j.visres.2005.10.002.

Lee, Joonyeol, Mati Joshua, Javier F F. Medina, and Stephen G G. Lisberger. 2016. “Signal, Noise, and Variation in Neural and Sensory-Motor Latency.” *Neuron* 90 (1): 165–76. https://doi.org/10.1016/j.neuron.2016.02.012.

Lin, John Y. 2012. “A User’s Guide to Channelrhodopsin Variants: Features, Limitations and Future Developments.” *Experiments in Physiology* 96 (1): 19–25. https://doi.org/10.1113/expphysiol.2009.051961.A.

Liu, Hesheng, Yigal Agam, Joseph R. Madsen, and Gabriel Kreiman. 2009. “Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex.” *Neuron* 62 (2): 281–90. https://doi.org/10.1016/j.neuron.2009.02.025.

Liu, Tina T., and Marlene Behrmann. 2017. “Functional Outcomes Following Lesions in Visual Cortex: Implications for Plasticity of High-Level Vision.” *Neuropsychologia* 105 (June): 197–214. https://doi.org/10.1016/j.neuropsychologia.2017.06.030.

Marmarelis, Panos Z, and Vasilis Z Marmarelis. 1978. *Analysis of Physiological Systems: The White-Noise Approach*. *Springer Science & Business Media*. https://doi.org/10.1007/978-1-4613-3970-0.

Mattis, Joanna, Kay M Tye, Emily a Ferenczi, Charu Ramakrishnan, Daniel J O’Shea, Rohit Prakash, Lisa a Gunaydin, et al. 2011. “Principles for Applying Optogenetic Tools Derived from Direct Comparative Analysis of Microbial Opsins.” *Nature Methods* 9 (2): 159–72. https://doi.org/10.1038/nmeth.1808.

Müller, J R, Andrew B Metha, John Krauskopf, and Peter Lennie. 2001. “Information Conveyed by Onset Transients in Responses of Striate Cortical Neurons.” *The*

Journal of Neuroscience : The Official Journal of the Society for Neuroscience 21 (17): 6978–90. https://doi.org/21/17/6978 [pii].

Neri, Peter, and David J. Heeger. 2002. “Spatiotemporal Mechanisms for Detecting and Identifying Image Features in Human Vision.” *Nature Neuroscience* 5 (8): 812–16. https://doi.org/10.1038/nn886.

Nienborg, Hendrikje, and Bruce G. Cumming. 2009. “Decision-Related Activity in Sensory Neurons Reflects More than a Neuron’s Causal Effect.” *Nature* 459 (7243): 89–92. https://doi.org/10.1038/nature07821.

Okazawa, Gouki, Long Sha, Braden A Purcell, and Roozbeh Kiani. 2018. “Psychophysical Reverse Correlation Reflects Both Sensory and Decision-Making Processes.” *Nature Communications* 9 (1). https://doi.org/10.1038/s41467-018-05797-y.

Oram, M W, and D I Perrett. 1992. “Time Course of Neural Responses Discriminating Different Views of the Face and Head.” *Journal of Neurophysiology* 68 (1): 70–84.

Packer, A. M., and R. Yuste. 2011. “Dense, Unspecific Connectivity of Neocortical Parvalbumin-Positive Interneurons: A Canonical Microcircuit for Inhibition?” *Journal of Neuroscience* 31 (37): 13260–71. https://doi.org/10.1523/JNEUROSCI.3131-11.2011.

Porto, Daniel A., John Giblin, Yiran Zhao, and Hang Lu. 2019. “Reverse-Correlation Analysis of the Mechanosensation Circuit and Behavior in C. Elegans Reveals Temporal and Spatial Encoding.” *Scientific Reports* 9 (1): 1–14. https://doi.org/10.1038/s41598-019-41349-0.

Resulaj, Arbora, Sarah Ruediger, Shawn R Olsen, and Massimo Scanziani. 2018. “First Spikes in Visual Cortex Enable Perceptual Discrimination.” *ELife* 7: e34044. https://doi.org/10.7554/eLife.34044.

Sachidhanandam, Shankar, Varun Sreenivasan, Alexandros Kyriakatos, Yves Kremer, and Carl C.H. Petersen. 2013. “Membrane Potential Correlates of Sensory Perception in Mouse Barrel Cortex.” *Nature Neuroscience* 16 (11): 1671–77. https://doi.org/10.1038/nn.3532.

Sakai, Hiroko M., Ken-Ichi Naka, and Michael J. Korenberg. 1988. “White-Noise Analysis in Visual Neuroscience.” *Visual Neuroscience* 1 (3): 287–96. https://doi.org/10.1017/S0952523800001942.

Schiller, Peter H, Julie H Sandell, and John H R Maunsell. 1986. “Functions of the ON and OFF Channels of the Visual System.” *Nature* 322: 824–25.

Schwartz, Odelia, Jonathan W. Pillow, Nicole C. Rust, and Eero P. Simoncelli. 2006. “Spike-Triggered Neural Characterization.” *Journal of Vision* 6 (4): 484–507. https://doi.org/10.1167/6.4.13.

Seal, John, Daniel Commenges, Roger Salamon, and Bernard Bioulac. 1983. “A Statistical Method for the Estimation of Neuronal Response Latency and Its Functional Interpretation.” *Brain Research* 278 (1–2): 382–86. https://doi.org/10.1016/0006-8993(83)90279-2.

Seidemann, Eyal, Ehud Zohary, and William T. Newsome. 1998. “Temporal Gating of Neural Signals Duringperformance of a Visual Discrimination Task.” *Nature* 394: 72–75. https://doi.org/10.1038/255243a0.

Shriki, Oren, Adam Kohn, and Maoz Shamir. 2012. “Fast Coding of Orientation in Primary Visual Cortex.” *PLoS Computational Biology* 8 (6). https://doi.org/10.1371/journal.pcbi.1002536.

Tchumatchenko, Tatjana, Aleksey Malyshev, Fred Wolf, and Maxim Volgushev. 2011. “Ultrafast Population Encoding by Cortical Neurons.” *Journal of Neuroscience* 31 (34): 12171–79. https://doi.org/10.1523/JNEUROSCI.2182-11.2011.

VanRullen, Rufin, and Simon J. Thorpe. 2002. “Surfing a Spike Wave down the Ventral Stream.” *Vision Research*. https://doi.org/10.1016/S0042-6989(02)00298-5.

Vugt, Bram Van, Bruno Dagnino, Devavrat Vartak, and Houman Safaai. 2018. “The Threshold for Conscious Report: Signal L” 7186 (March): 537–42.

Williams, Ziv M., John C. Elfar, Emad N. Eskandar, Louis J. Toth, and John A. Assad. 2003. “Parietal Activity and the Perceived Direction of Ambiguous Apparent Motion.” *Nature Neuroscience* 6 (6): 616–23. https://doi.org/10.1038/nn1055.

Wilson, Christopher D., Gabriela O. Serrano, Alexei A. Koulakov, and Dmitry Rinberg. 2017. “A Primacy Code for Odor Identity.” *Nature Communications* 8 (1). https://doi.org/10.1038/s41467-017-01432-4.