Title: A parameter-free statistical test that improves the detection of neuronal responsiveness

Manuscript Length

Number of figures: 8

Number of supplementary figures: 2

Authors

Jorrit S. Montijn[1], Koen Seignette[1], Marcus H. Howlett[1], J. Leonie Cazemier[1], Maarten Kamermans[1], Christiaan N. Levelt[1], J. Alexander Heimel[1]

Affiliations

1 Netherlands Institute for Neuroscience, Royal Netherlands Academy of Arts and Sciences, Amsterdam, 1105 BA, the Netherlands

Contact Information

Correspondence and requests for materials should be addressed to J.S.M. (jsmontijn@gmail.com) or J.A.H. (heimel@nin.knaw.nl).

Abstract

Neurophysiological studies depend on a reliable quantification of whether and when a neuron responds to stimulation. Current methods to determine responsiveness require arbitrary parameter choices, such as binning size. These choices can change the results, which invites bad statistical practice and reduces the replicability. Moreover, many methods only detect mean-rate modulated cells. New recording techniques that yield increasingly large numbers of

cells would benefit from a test for cell-inclusion that requires no manual curation. Here, we present the parameter-free ZETA-test, which outperforms t-tests and ANOVAs by including more cells at a similar false-positive rate. We show that our procedure works across brain regions and recording techniques, including calcium imaging and Neuropixels data. Furthermore, in illustration of the method, we show in mouse visual cortex that 1) visuomotor-mismatch and spatial location are encoded by different neuronal subpopulations; and 2) optogenetic stimulation of VIP cells leads to early inhibition and subsequent disinhibition.

Keywords

Neural data analysis, Statistics, Neurons, Responsiveness, Response Latency, Electrophysiology, Calcium Imaging, Open Source, Visual Cortex, VIP disinhibition

Introduction

Many neuroscience studies rely on the analysis and visualization of neuronal spiking signals. Classical studies used manual curation during experiments to select cells for analysis (Hubel and Wiesel, 1959; Mountcastle, 1957), but this method cannot provide a statistically unbiased sample. Moreover, such manual curation is unsuitable for state-of-the-art large-scale recording techniques, such as Neuropixels and high-density multi-electrode arrays (Bartolo et al., 2020; Jun et al., 2017; Semedo et al., 2019; Steinmetz et al., 2019).

Despite the widespread application of neuronal responsiveness analyses, neuroscience currently lacks a standard practice for determining whether a neuron is responsive to an experimental stimulus or treatment (Mesa et al., 2020). Common approaches, such as comparing a neuron's average spike rate during the presence and absence of a stimulus, can only detect mean-rate modulated cells (Mazurek et al., 2014; Ringach et al., 2002). On the other hand, approaches such as computing a peri-stimulus time histogram (PSTH) and applying an ANOVA, require the *a priori* selection of an arbitrary binning window size (Palm et al., 1988). Choosing the wrong bin size reduces the test's sensitivity, while optimizing from a range of window sizes creates a multiple-comparison problem. This lowers the approach's statistical power when corrections are applied, or can even lead to (unintentional) "p-hacking" if the results are not corrected (Head et al., 2015).

To solve these problems, we developed a method that detects whether a cell is responsive to stimulation in a statistically robust way and avoids binning altogether. This method, which we call ZETA (Zenith of Event-based Time-locked Anomalies), either outperformed or matched that of t-tests/ANOVAs in all conditions tested. Building upon this framework, we also present a procedure to visualize instantaneous spiking rates without the need for binning, and show how this can be used to estimate peak-activity latencies with sub-millisecond accuracy. We apply these methods to transient-detected two-photon calcium imaging data from the visual cortex of mice traversing a virtual linear track and find that

visuomotor mismatch signals and spatial location are encoded by different V1 neuronal subpopulations. Finally, we apply our approach to Allen Brain Institute Neuropixels data and show that optogenetic stimulation of VIP-expressing cells in mouse visual cortex has a separable early inhibitory and late disinhibitory effect on the local neural circuit. We anticipate that the ZETA-test will be a useful resource for a wide range of applications across various disciplines.

Results

ZETA: Zenith of Event-based Time-locked Anomalies

A common procedure in pre-processing neural data is removing cells that are not responsive to an experimental stimulus. Many experimenters determine the "stimulus responsiveness" of a cell by comparing its average spiking rate before and during the presentation of a stimulus (fig. 1A-C). This procedure will therefore remove neurons that show no response, but has the risk of also removing neurons that show a strong, but complex time-locked response to stimuli. To remedy this shortcoming, we developed a binning-free method for determining whether a neuron shows any time-locked modulation. We call this statistical test ZETA for Zenith of Event-based Time-locked Anomalies (fig. 1D-F). It represents whether a neuron's spike train could be observed by chance, if it were not responding to an experimenter's event of interest: e.g., the presentation of a visual stimulus, the onset of optogenetic stimulation, or a self-generated variable, such as an animal's location on a track.

ZETA is calculated on a single cell by performing the following steps. First, we align all spikes to stimulus onsets, i.e. building a raster plot (fig. 1A). Pooling all spikes across trials, we obtain a single vector of spike times relative to stimulus onset, and calculate the cumulative distribution as a function of time (fig. 1D). The deviation of this curve from a linear baseline represents whether the neuron has a higher or lower spiking density relative to a non-modulated

spiking rate (fig. 1E, blue curve). We compare this pattern to the likelihood of observing it by chance by running multiple jittered bootstraps to generate a null hypothesis distribution (fig. 1E, grey curves). After scaling the experimentally observed curve to the variation in the null hypothesis distribution, we use it to obtain a p-value corresponding to the Zenith of Event-based Time-locked Anomalies. Low ZETA-test p-values indicate that the neuron's firing pattern is statistically unlikely to be observed if the neuron is not modulated by the event of interest.

ZETA's main difference from other approaches used to infer a neuron's stimulus responsiveness, such as mean-rate or PSTH-based methods, is that this metric makes no *a priori* assumptions about the underlying distribution of temporal modulations and is binning-free. It can therefore detect both long-timescale changes in mean firing rate, as well as short-timescale stimulus-locked bursts or lapses of activity at any point in time relative to stimulus onset.

Benchmarking the ZETA-test

To investigate whether the ZETA-test includes more cells recorded in mouse visual cortex in response to a drifting grating, while still retaining a false-alarm rate at the significance level $\alpha = 0.05$, we used a benchmarking test comparing jittered and non-jittered data (fig. 2). In the non-jittered case, we compared the inclusion rate of the ZETA-test, as described above, to that of a mean-rate t-test. For the t-test, we calculated the average spiking rate of a cell during stimulation (0s – 1.0s after onset) and before stimulation (-0.5s – 0s), and performed a paired t-test over trial repetitions (fig. 1C). This showed that cells included with a t-test were almost exclusively a subset of the cells detected with ZETA (see fig. 2B for all V1 cells recorded with Neuropixels). In other words, if a cell is detected as being visually responsive with a t-test, it is almost guaranteed to also be detected by the ZETA-test. In addition to these cells, the ZETA-test also includes cells that were not registered by a t-test. Although many varieties exist, fig.

2C shows an example cell that is detected by both t-tests and ZETA (top, sustained change in firing rate), and an example only detected by ZETA (bottom, balanced on/off peaks).

Next, we ran the same tests again, but now on data where we randomly jittered the stimulus onset times between $-\tau$ and $+\tau$, where τ is median onset-to-onset duration. This procedure preserves the properties of a neuron's spike train, but removes locking of the responses to the stimulus (fig. 2D). If the ZETA-test simply always gave low p-values, then this would result in a high false-alarm rate as many neurons would still be included. In contrast, the false-alarm rate of the ZETA-test was generally low, and consistent with that expected for a significance level of $\alpha=0.05$ (fig. 2E-F).

Sensitivity of ZETA is superior to mean-rate t-tests

We performed this benchmark for single-cell activity obtained from n=12 combinations of various visual regions (V1, AM, PM, LGN, SC, LP, NOT, APN, Retina) using multiple techniques (Neuropixels, n=8; Neuronexus, n=2; GCaMP6, n=1; pMEA, n=1), in response to light flashes (retina, n=1) or drifting gratings (all others, n=11). Under all conditions, the inclusion rate using ZETA-tests was higher than using t-tests, while the false alarm rate (FA) did not differ significantly: average inclusion rate ZETA-tests = 77%, t-tests = 60%, $p=1.0 \times 10^{-7}$; ZETA-tests FA = 4.0%, t-tests FA = 4.1%, p=0.886; $\alpha=0.05$, n=12 data sets (fig. 1G, fig. 3A). This means that ZETA includes 28% more cells that are truly stimulus-modulated at the same level of false positives.

The percentage of visually responsive cells detected by the ZETA-test is higher than typically reported. For example, using gratings that only differed by orientation, we found 93.3% of all V1 cells to be visually-modulated. Even more striking was that the lower bound of the binomial 95%-confidence interval was at 89.7%. This lower-bound is higher than the responsiveness previously reported in many studies (Montijn et al., 2016a; Niell and Stryker,

2008; Shuler and Bear, 2006; Steinmetz et al., 2019), exemplifying the advantage of the ZETA-test.

The data used so far for benchmarking contains more stimulus repetitions than commonly used in neuroscience. This raises the possibility that the ZETA-test is only advantageous when large numbers of trial are used. To test this, we randomly subsampled the number of trials included in the analysis, and repeated our benchmark procedure. As the results for V1 in fig. 3B show, the ZETA-test consistently included more than the t-test, regardless of the number of trials.

We hypothesized that the t-test's worse performance might result from pooling the responses to different orientations into one group. Therefore, we repeated the t-test's benchmark after first splitting the trials into 24 groups corresponding to the 24 directions we presented. A neuron was included if the spiking rate during stimulus presentation was significantly different from its pre-stimulus baseline rate in any group, after applying a Bonferroni correction. However, this procedure reduced the t-tests' performance (fig. 3B middle panel).

We noticed that the t-test's false alarms were rather low after Bonferroni corrections. To test whether we over-corrected the t-test, we removed the multiple-comparison correction (fig. 3B, right-hand panel). In this case, the t-test false alarm rate increased to >50%, while its inclusion rate (89.9%) remained lower than that of the ZETA-test (95.0%). Finally, we investigated whether the t-test's performance was hampered by including the immediate off-response after stimulus offset in the baseline period. We reran the above analyses, but now limited the baseline to the 300 ms preceding the stimulus onset. As can be seen in Supplementary Figure 1, this did not improve the t-test's performance.

Taken together, these results show that the t-test is most sensitive when using the full 500-ms epoch preceding the stimulus onset as baseline period and when pooling data across all orientations. Moreover, they show that the ZETA-test is superior to the mean-rate t-test in all

cases, regardless of number of trial repetitions, brain region where the data were recorded, or specifics of the data preparation.

ZETA-test in the absence of short peaks of activity

Having established that the ZETA-test performs well in real neural data, we searched for conditions under which the ZETA-test breaks down. First, we asked whether the ZETA-test's performance depended on short bursts of activity, as often seen in the initial onset response of neurons to visual stimulation. We therefore ran a simulation with artificial neurons that followed a Poisson distribution for inter-spike intervals. We generated neurons with a range of background activities (see methods, mean μ = 5 Hz) and a range of stimulus-driven activities (μ = 25 Hz). Unlike many real neurons, these artificial neurons do not show a short onset peak of activity (fig. 4, top row). Running the above described benchmark for ZETA and mean-rate t-tests, we found that the ZETA-test performed better than a t-test in discriminating stimulus-responsive neurons (α=0.05, inclusion rate ZETA = 92.5%, t-test = 89.1%, two-sample binomial test, n=10000, p=2.2 x 10^{-16}, FA rate ZETA = 4.8%, t-test = 3.8%, p=7.2 x 10^{-4}). However, note that the t-test's FA rate was low. Therefore, we raised the t-test's α to 0.06 to investigate if it could be made more sensitive by raising its rejection threshold. At α=0.06, the t-test's inclusion rate grew to 89.6% at an FA rate of 4.8%, but this was still significantly lower than the ZETA-test at α=0.05 (two-sample binomial test, 92.5% vs 89.6%, n=10000, p=6.4 x 10^{-13}). In conclusion, even under these conditions, which are ideally suited to the properties of a mean-rate t-test, the ZETA-test still exceeds its performance.

We asked whether this generalized to other recording techniques as well (fig. 4). We repeated the benchmarking procedure described in the previous paragraph on several data sets and found that the ZETA-test included more cells than a t-test in all cases; mouse V1 responses to drifting gratings recorded with a Neuronexus probe (N=40 neurons, inclusion ZETA=67.5%,

t-test=52.5%) and GCaMP6f (N=183, Z=41.5%, t=18.0%), mouse superior colliculus responses to drifting gratings recorded with a Neuronexus probe (N=89, Z=64.0%, t=55.1%), and zebrafish retinal ganglion cell responses to on/off light flashes recorded with a perforated micro-electrode array (N=280, Z=95.7%, t=85.4%).

Neuronal responsiveness to natural movies

Next, we asked how the performance of the ZETA-test compares to that of an ANOVA, in a case where there is no *a priori* knowledge regarding the neuronal response profile. We determined the responsiveness of neurons to natural movies, using either the ZETA-test, or a one-way ANOVA across bins, repeated for different bin sizes. Single-cell data were recorded using Neuropixels in 7 visual brain areas of 3 mice, while the animals were presented with repetitions of 20 s long natural movies (fig. 5A).

To ensure that the ANOVA approach could detect short bursts of activity, as well as long timescale whole-scene modulations of firing rates, we chose a wide range of bin sizes. We picked a single movie frame duration (0.0167s) as the centre point, and used bins from 1/512th up to 512 movie frames, spaced equidistantly on a base-2 logarithmic scale in 18 steps. For each area, we pooled all cells, and calculated the area-level inclusion rate with either the ZETA-test or Bonferroni-corrected ANOVAs (fig. 5B). Bin sizes of 33 – 267 ms did not differ significantly from the ZETA-test's inclusion rate (FDR-corrected paired t-tests, $p>0.05$, n=7 areas), while the ZETA-test's inclusion rate was higher than with an ANOVA for all short (< 33 ms) and long (> 267 ms) bin durations ($p<0.05$).

The above approaches give some insight into which bin sizes best capture the dominant temporal components in neuronal responses in our data, but do not rigorously compare the sensitivity of ANOVAs to the ZETA-test. To better address this issue, we compared the ZETA-test to the ANOVAs' performance when classifying a neuron as "included" whenever any of

the 19 ANOVAs reached significance (i.e., $p < \alpha$). Repeating this procedure for various significance levels α on the interval (0,1) produces an ROC curve (fig. 5C). The ROC's summary statistic is the area under the curve (AUC); 1 being a perfect discriminator. The AUC is conveniently similar to a Wilcoxon-Mann-Whitney statistic and can be used to determine which procedure is more sensitive (Calders and Jaroszewicz, 2007). Using this approach, we found there was no significant difference in performance between a set of ANOVAs and the ZETA-test (z-test, p=0.554): ANOVAs AUC=0.798 ± 0.010, ZETA-test AUC=0.792 ± 0.011 (mean ± sd). Overall, the above results show that, under these conditions, the ZETA-test performs as well as an aggregate set of ANOVAs.

Instantaneous firing rates (IFRs) for visualization and onset latency detection

The above paragraphs have shown that the ZETA-test is a sensitive statistical tool to detect whether neurons respond to a stimulus. However, it cannot be used to determine when exactly the strongest response of a neuron occurs. Therefore, we developed a method that determines the instantaneous firing rate (IFR) using the temporal deviations upon which ZETA is based. Like the ZETA-test, it avoids the bin size selection issue of peri-stimulus time histograms (PSTHs). Another advantage of this IFR is that its temporal resolution is limited only by the neuron's spike density. It is therefore a useful tool for determining spike train features with high precision, such as a neuron's onset latency.

Fig. 6A,B show two example V1 neurons with a relatively high (fig. 6A) and low (fig. 6B) firing rate. Here we define "onset" as the time the half-maximal response of the peak is first crossed, a metric that is heavily influenced by the chosen bin size when using PSTH-based analyses. Moreover, this bin-width-dependent estimation makes PSTH-based comparisons across cell classes problematic, as the choice of optimal bin width depends on spiking properties such as the firing rate and peak firing duration, which are heterogeneous across neuronal types

(fig. 6A4, B4). Hence, the main advantage of our method is that it avoids having to tailor the bin size to each neuron individually; allowing for better comparisons across varying cell types and brain regions.

To test performance on real data, we compared the estimated onset latency using our IFR to what we obtain from a PSTH analysis with varying bin sizes. Across all V1 neurons (fig. 6C, n=119 neurons) as well as all neurons recorded in visual areas (fig. 6D, n=1403 neurons), our metric showed the strongest agreement in latency estimation with PSTH bin sizes between 1 – 10 ms.

As no ground truth is known for the real latencies of experimentally recorded neurons, we performed a benchmark test with artificially generated spike trains. Like before, we used Poisson neurons with a constant background spiking rate, where each neuron was assigned a background rate ranging from 0.5 to 32 Hz (fig. 6E-G). We used 100 trials, each 2 seconds long, and created a peak response on top of this baseline rate by adding a single spike in 50% of all trials (i.e., 50 spikes in total). We also varied the peak-response width by jittering the time each spike was added according to a normal distribution, σ ranging from 1 – 10 ms. We found that the error in the peak estimate was independent of background spiking rate (fig. 6H), indicating robust performance even when the spikes contributing to the peak were only 0.78% of the total (i.e., for 32 Hz). The mean error was also independent of the jitter, while the standard deviation of the error estimate grew in a theoretically optimal fashion as $O(\sigma)$ (fig. 6I).

We next tested the estimator accuracy when using a binning-based PSTH approach (bin width 1 – 58 ms) and found it depends on both the bin width and the width of the peak response (fig. 6J). This means that accurate latency determination using PSTHs requires the use of multiple bin sizes. More importantly, when comparing the estimation error between these bin-based methods and our IFR, we found that the binning-less IFR-based latencies were consistently as accurate as, or more accurate than, the best possible bin-width for any given

peak-width (fig. 6K). In other words, the IFR-based accuracy supersedes the PSTH-based accuracy without the need to hand-pick the optimal binning width per neuron.

Visuomotor mismatch and spatial location are mediated by different neuronal subpopulations

Having established that the ZETA-test and IFR are statistically robust and have clear advantages over mean-rate approaches and PSTHs, we applied these tools to a GCaMP6 data set. Many theories, such as predictive coding (Friston, 2005; Gregory et al., 1980; Rao and Ballard, 1999), biologically realistic error backpropagation (Ooyen and Roelfsema, 2003; Whittington and Bogacz, 2019), and canonical cortical microcircuit operation (Bastos et al., 2012; Douglas et al., 1989), define a "top-down" signal representing an expectation, error or surprise signal as distinct from a bottom-up sensory drive. In mouse V1, such top-down visuomotor mismatch signals have been reported previously (Attinger et al., 2017; Keller et al., 2012; Leinweber et al., 2017; Saleem et al., 2013). However, whether individual V1 neurons can be classified into different groups based on their encoding of top-down visuomotor mismatch or bottom-up sensory-driven spatial location signals has not been studied.

To examine this issue, and to provide an example of how one could use ZETA in a neurophysiology study, we used neuronal calcium data recorded in L2/3 V1 of 4 mice running on a virtual-reality linear track (fig. 7A). In 87% of all corridor runs (N= 622/713 trials), the track was rendered normally, and the mice received visual feedback matching its running speed. In the remaining 13% of runs (N=91), rendering was halted at a random location for 500 ms before resuming (fig. 7B). After performing calcium-transient detection to obtain putative spike times (Montijn et al., 2016b), we calculated ZETA-scores for all neurons in three different ways. We aligned the spikes to mismatch-onsets, to trial starts, or converted the spike times into locations on the track, and aligned these spike locations to the start. For each recording

(n=7), we calculated the Pearson correlation for each pair of these three ZETA-scores (fig. 7C-F). Across recordings, time- and location-modulation were positively correlated (mean r=0.23, one-sample t-test, n=7 recordings, p=0.016); time- and mismatch-modulation were not significant (r=0.13, p=0.28); and location- and mismatch-modulation were negatively correlated (r=-0.22, p=0.04).

We next used a kernel-density estimate (KDE) to directly test whether the joint encoding of two features within single neurons was different from chance (fig. 7G-I; Supplemental Figure 2). Indeed, we found that it was less likely that a neuron showed high modulation values for both spatial location and visuomotor mismatch than if the two features were encoded independently (z=-4.3, p=2.0 x 10^{-5}). We also found time and location to be more likely to be encoded by the same neurons (z=3.8, p=1.1 x 10^{-4}), while time- and mismatch-modulation show no effect (z=-0.06, p=0.95). This suggests a functional specialization for many neurons to encode either visuomotor mismatch signals or spatial location, but not both. While it is possible that location- and mismatch-encoding are also mediated by specific genetic subtypes of interneurons (Attinger et al., 2017), our analysis of putative pyramidal cells demonstrates that principal cells also show encoding specialization.

Optogenetic stimulation of VIP cells disinhibits visual cortex

Finally, we applied the IFR and ZETA-test to data recorded at the Allen Brain Institute (Siegle et al., 2019). In this case, we investigate whether optogenetic stimulation of VIP cells in visual cortex disinhibits the local circuit, as has been shown previously for auditory cortex and mPFC (Pi et al., 2013). The analysis of optogenetic stimulation in visually-responsive areas is complicated by the fact that mice can see the blue light used for optogenetic stimulation. In other words, if a neuron is active after a laser pulse, it could be caused by direct stimulation, indirect circuit disinhibition, or simply be a sensory-driven response. Using the methods

described in this paper to overcome these issues, we show that optogenetic stimulation of VIP-expressing cells in mouse visual cortex causes short-latency inhibition and longer-latency disinhibition in separate neuronal subpopulations (fig. 8).

Data used for this analysis were recorded from 5 Vip-IRES-Cre; Ai32 mice at the Allen Brain Institute. Cells were only included for analysis if the clustering quality was sufficient, and they were recorded in a visual cortex area (1706 cells total in AL, AM, PM, L, V1, RL, MMP). We performed a ZETA-test to discard cells that were not modulated within the interval (-0.5, +0.5 s) after optogenetic stimulation (N=1144 cells remaining). Instantaneous firing rate peak- and trough-latencies were computed and cells were discarded if their peak occurred earlier than +1 ms after optogenetic stimulation onset. The remaining cells (N=909) were classified as VIP (N=13) if their peak latencies occurred within the 10ms duration of the optogenetic stimulation, as Inhibited (N=59) or Activated (N=137) if they displayed respectively decreased or increased firing within 20ms from the stimulus offset, or Other (N=700). We limited our classification window to the first 30ms after the onset of the optogenetic stimulus as it takes about 30-50ms for retinal light responses to emerge in the visual cortex. Hence, the majority of 'Other' cells likely show sensory-driven responses.

Based on these classifications, we compared the single-cell latencies of Inhibited and Activated neurons, restricting our analysis to peaks within the 10-30ms post-optogenetic stimulation window. The latency of Inhibited cells was significantly shorter after optogenetic stimulation than it was for Activated cells (median latency Inh = 16.5 ms, Act = 23.6, Mann-Whitney U-test, $p=5.1 \times 10^{-7}$). These results show a VIP-mediated disinhibition mechanism operates in vivo in visual cortex, confirm slice connectivity studies (Pfeffer et al., 2013), and are in general agreement with results from the auditory cortex and mPFC (Pi et al., 2013).

Discussion

We developed the ZETA-test, a statistical method for determining whether neuronal spiking responses are modulated by the occurrence of events, such as the onset of sensory stimulation. The ZETA-test is widely applicable: we have shown that it can accurately detect neuronal responsiveness in a wide range of settings, across various brain regions, stimuli, and recording techniques. In some cases, its performance was equal to established and powerful statistical techniques, such as t-tests and ANOVAs. In all other cases, the ZETA-test showed markedly improved statistical sensitivity. For example, the ZETA-test detected 28% more visually modulated cells in response to drifting gratings than did the stimulus-epoch based t-tests. In addition to its improved statistical performance, the ZETA-test avoids arbitrary parameter choices, as it does not require the selection of a temporal bin size. This makes the ZETA-test even easier to use than established methods, as it can be applied directly to raw spike times and stimulus onsets, and the lack of a parameter selection naturally lends itself to the bulk-analysis of large numbers of cells.

Secondly, we developed an instantaneous firing rate for analysis and visualization of neuronal firing patterns. Similar to the ZETA-test, it has two main advantages over alternative common approaches, such as PSTHs: 1) it does not require binning, and therefore removes the need to manually tune binning widths for different cell types; and 2) its temporal resolution is only limited by the spike density, which allows more accurate determination of spike train events such as peak response latencies. When benchmarking its performance for latency-estimation, it outperformed PSTH-based approaches without the need to fine-tune bin sizes for each neuron individually (fig. 6K).

We investigated the performance of the ZETA-test in a variety of cases and observed that the ZETA-test sometimes performed as well as other techniques and sometimes better. So when exactly does the ZETA-test provide extra statistical power? We hypothesize that this is in cases where the complexity and heterogeneity of the neuronal spike trains exceeds the

complexity of the statistical test that would naturally follow from the stimulation protocol. For example, we observed that the ZETA-test performed similarly to an aggregate set of ANOVAs with various bin sizes (fig. 5). In this case, there are many degrees of freedom for the set of ANOVAs. Inherently complex neuronal spike trains, even if they were only in response to a single movie epoch, can therefore already be accurately captured within the aggregate ANOVA analysis. In contrast, using a t-test on "stimulus" versus "no-stimulus" activity averages out any time-varying modulations in neuronal activity (fig. 3).

An alternative to using the ZETA-test could be to perform a set of ANOVAs as in fig. 5. While this would certainly be an improvement upon doing a paired t-test, multiple ANOVAs are more complex to implement and the procedure needs multiple-comparison correction (Head et al., 2015). More importantly, it still requires choosing an arbitrary set of bin sizes. The ZETA-test therefore has the advantage of not having to choose any parameter at all, while still providing the same statistical sensitivity as the more complex procedure of a set of ANOVAs. For example, investigating the effect of optogenetic stimulation of VIP neurons on the local neural circuit requires analysing heterogeneous latencies and response profiles of neurons. Using the ZETA-test and IFR, we could significantly simplify various analysis steps (fig. 8).

However, like all statistical tools, the ZETA-test also has its limitations. The most obvious is that it only works when comparing point-events (e.g., spike times) to other point-events (e.g., stimulus onsets). It is therefore a powerful tool for state-of-the-art electrophysiological techniques, like Neuropixels recordings, but does not apply to calcium imaging dF/F0 traces. However, as we show in fig. 7, the ZETA-test performs well when applied to putative spiking events extracted from calcium imaging data. Therefore, the ZETA-test would still be useful in calcium imaging data sets where the temporal response profile of neurons to the experimental treatment is unknown. Moreover, optical recording techniques continue to improve, allowing better single-spike extractions with calcium imaging (Packer et

al., 2015) and genetically encoded voltage indicators (Knöpfel and Song, 2019). As such, in all likelihood near-future neural data will remain spike-based regardless of the underlying recording technique used.

The main aim of this paper is to present the IFR and ZETA-test and show their improvements over commonly used analysis tools, such as PSTHs, t-tests and ANOVAs. However, we also show some results that by themselves are scientifically noteworthy. For example, our results indicate that almost all V1 neurons (>93%) are responsive to drifting gratings, even if the spatial frequency (SF) and temporal frequency (TF) parameters are not optimal to drive these cells (fig. 3). This calls into question the idea that V1 cells are generally narrowly tuned and can only be driven by specific stimulus feature parameters (Xing et al., 2004). Our analysis indicates that neuronal responsiveness rarely drops to 0 for a combination of stimulus features, as otherwise we would not have found such a large number of responsive V1 cells using a feature-sparse stimulus (24 directions, 1 SF: 0.05cpd, 1 TF: 2Hz).

This suggests that the response of V1 neurons to features such as SF and TF may drop off rather slowly with distance to their preferred stimulus properties. Conceptually, this would mean that the V1 neural code may in fact be more "dense" than proposed by some (Ohiorhenuan et al., 2010; Olshausen and Field, 1997; Vinje and Gallant, 2000). That said, our results show that only a relatively small group of neurons is strongly driven by any one stimulus: sparse coding may therefore still operate in V1, but it acts on top of a dense, but weak code.

Two other noteworthy findings are that VIP cell activation drives disinhibition in visual cortex (fig. 8) and that the strength of visuomotor signals and modulation by spatial location are negatively correlated (fig. 7). This suggests that bottom-up and top-down processing in visual cortex may be mediated by distinct neuronal subnetworks. While this functional segregation has previously been shown based on a macroanatomical (e.g., laminar) analysis (Kok et al., 2016; Markov et al., 2014; Poort et al., 2012; Self et al., 2019), our results suggest

that this functional segregation also holds at a microanatomical level for neurons located within a single recording plane (fig. 7).

In conclusion, the IFR and ZETA-test are simpler, more statistically powerful, and less error-prone tools than bin-based PSTHs, t-tests and ANOVAs widely used in neuroscience today. Statistically underpowered studies are still common in neuroscience, which makes the development of statistically sensitive tools especially important (Button et al., 2013). Moreover, when power analyses are used to determine the necessary sample size, then increased statistical sensitivity, such as with the ZETA-test, can reduce the number of required experimental animals. To facilitate the adoption of the ZETA-test by the neuroscientific community, we provide easy-to-use and well-documented open-source implementations online in MATLAB (https://github.com/JorritMontijn/ZETA) and Python (https://github.com/JorritMontijn/zetapy).

Acknowledgments

JSM was funded by a grant awarded to JAH by the Stichting Vrienden van het Herseninstituut. Brain images in figures 3-5 were generated using Brain Explorer 2 (Lau et al., 2008). We thank the Allen Brain Institute for their openly accessible data sets. We also thank the engineers of the mechatronics workshop at the NIN.

Author contributions

JSM designed ZETA, performed the neuropixels recordings, and performed the analyses. KS and JSM performed the GCaMP6 recordings. JSM and JLC performed the neuronexus recordings. MHH performed the retinal recordings. JSM and JAH refined ZETA and wrote the manuscript. All authors discussed the results and contributed to the final manuscript.

Declaration of Interests

The authors declare no conflicts of interest exist.

Figures and legends

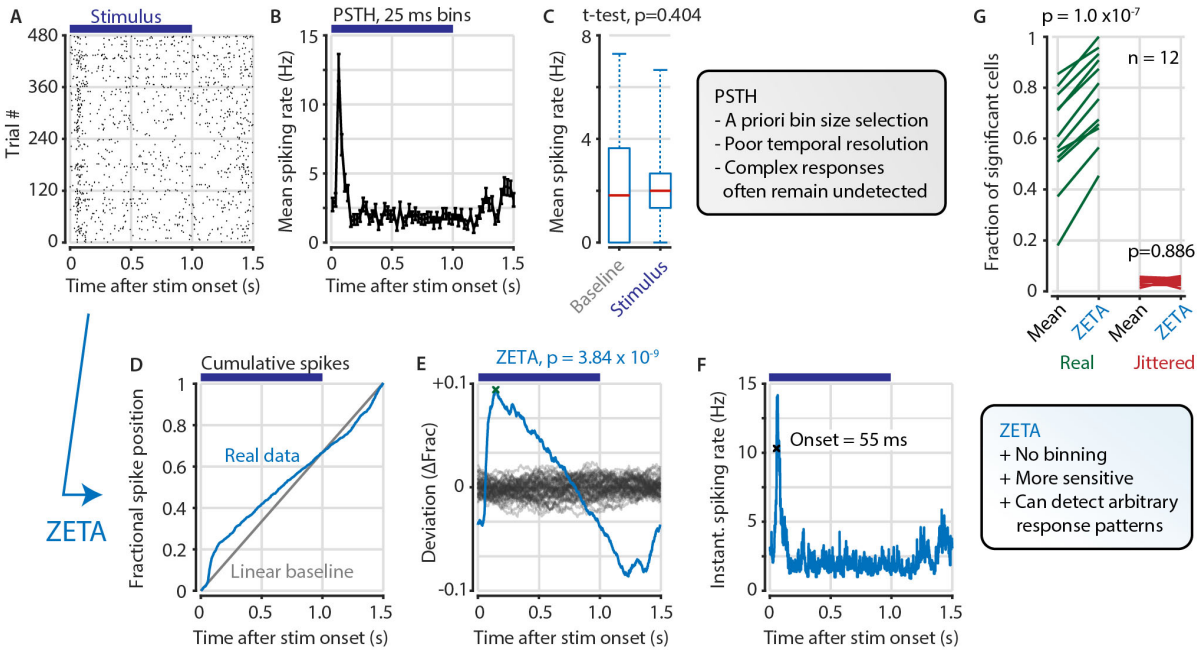

Figure 1. The ZETA-test detects visually responsive neurons. A,B) Raster plot (A) and PSTH (B) for a neuron that shows an onset peak and a reduced sustained spiking in response to a visual stimulus (purple bars). C) A common approach for computing a neuron's responsiveness is to perform a t-test on the average activity per trial during stimulus presence (0 – 1s) and absence (-0.5 – 0s), which fails to detect this neuron's response (paired t-test, p=0.404, n=480 trials). D) ZETA avoids binning, by using the spike times to construct a fractional position for each spike (blue) and compares this with a null-distribution of a stationary rate (grey). E) The difference between the real and null curves gives a deviation from expectation (blue), where the most extreme value is defined as the Zenith of Event-based Time-locked Anomalies (ZETA, green cross). To compute its statistical significance (here, $p=3.84 \times 10^{-9}$), we scale ZETA by the standard deviation over repeats of the procedure with jittered onsets (grey curves). F) A ZETA-derived instantaneous spiking rate allows a reliable estimation of response onset. G) ZETA detects more stimulus-responsive cells than a mean-rate t-test (paired t-test, $p = 1.0 \times 10^{-7}$, n=12 data sets) at an identical false-alarm rate (p = 0.886).

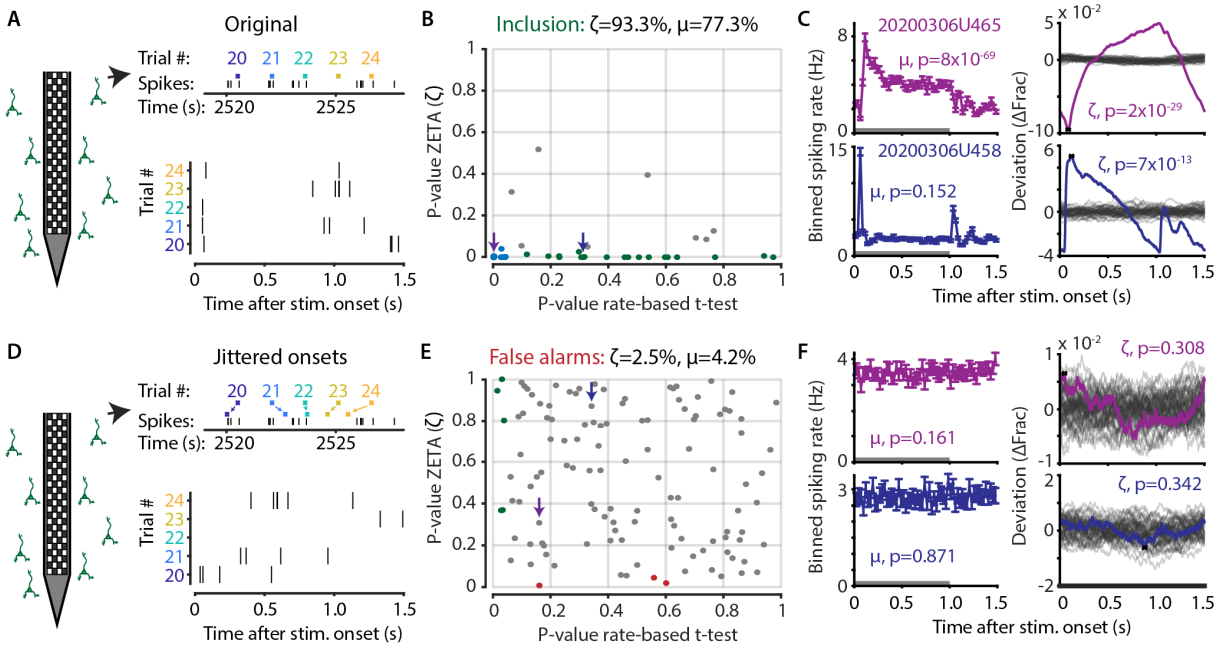

Figure 2. The ZETA-test outperforms the t-test for V1 neurons. A) Spikes were recorded with neuropixels in mouse V1 and aligned with the onset of square-wave drifting gratings. B) Applying the ZETA-test, we found that 93.3% of all V1 neurons showed significant firing rate modulations that were time-locked to stimulus onset. Using a rate-based t-test between stimulation (0 - 1s) and no stimulation (-0.5 - 0s) epochs across trials registered 77.3% of neurons to be visually responsive. Neurons detected by both methods are blue, by neither are grey, and by only ZETA are green. All neurons detected by the t-test were also detected by ZETA. Arrows indicate the neurons shown in C. C) Two example neurons that are 1) detected by a rate-based t-test as well as ZETA (top) and 2) not detected by a rate-based t-test, but only by ZETA (bottom). D) We investigated the false-alarm rate of both approaches by jittering the onsets of visual stimuli; this preserved the temporal structure of the spiking response, but destroys the time-locked modulations in activity. E-F) same cells and analyses as in B-C but for jittered onsets. Red indicates neurons included by ZETA, but not a t-test; green indicates neurons included by a t-test, but not ZETA. As expected, the false alarm rate was around 5% (α=0.05) for both approaches.

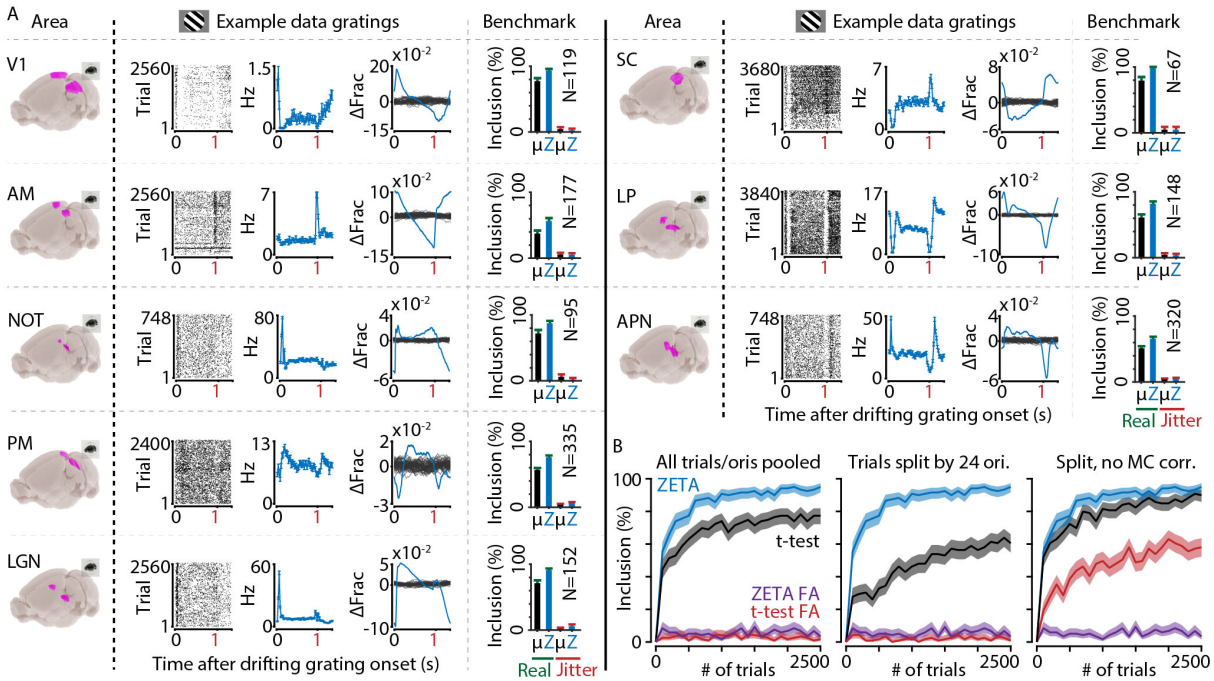

Figure 3. The ZETA-test's superior sensitivity is independent of brain area. A) We recorded neuronal responses to drifting gratings (1s with 500ms blank ITIs) from various visual brain areas: V1, SC, AM, LP, NOT, APN, PM, and LGN. For each area, we show an example neuron's raster plot (left), binned responses (middle) and temporal deviation from a static firing rate (right). All cells depicted here were significant using ZETA. Area-level benchmark summaries are presented in the histograms, comparing the inclusion rate ('Real', green) and false-alarm rate ('Jitter', red) for rate-based t-tests (μ, black) and ZETA (Z, blue). In all cases, ZETA included more neurons, while both approaches showed a false-alarm rate around 5%. Bar graphs show inclusion rate and 75th percentile confidence interval. B) The ZETA-test performs better than t-tests at low trial numbers and regardless of data preparation. Graphs show V1 inclusion percentage as a function of number of trials subsampled from the whole data set, using the ZETA-test (blue) and t-test (black), and false alarms after jittering using the ZETA-test (purple) and t-test (red). Left-hand panel shows results when trials of all orientations are pooled. Middle and right-hand panels show results when t-test p-values are calculated separately per orientation (24 groups) and Bonferroni-corrected (middle) or not corrected (right).

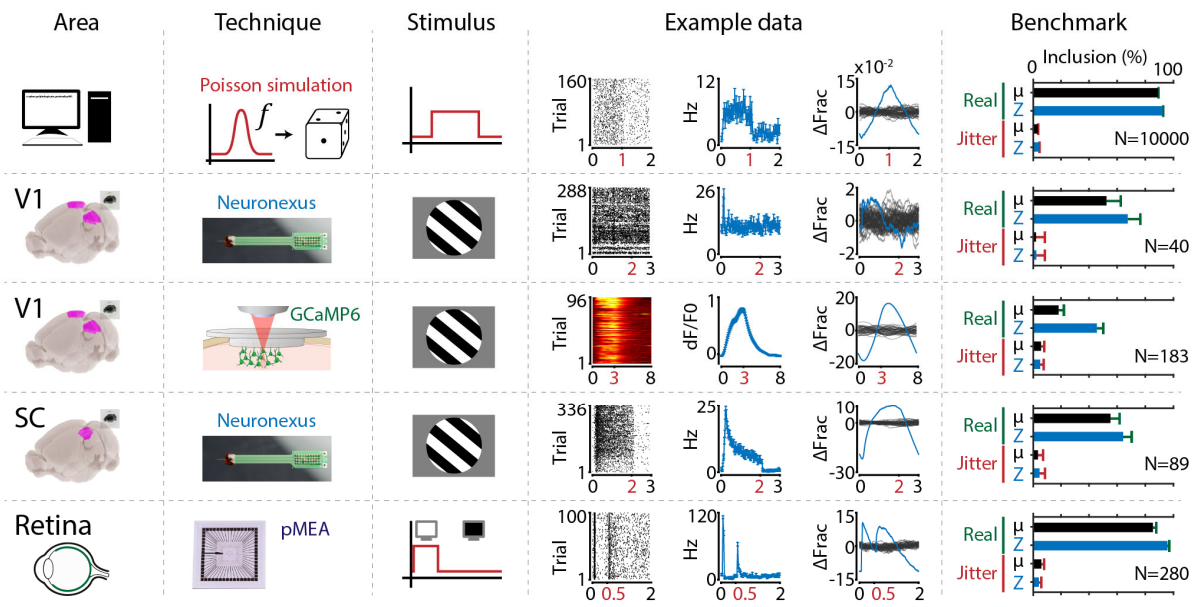

Figure 4. The ZETA-test is more sensitive than a mean-rate approach across a range of areas, recording techniques and stimuli. From left to right the figure shows the following. First column: the recorded area. Second column: recording technique. Third column: stimulus. Fourth column: example data including raster plot or heat map (lhs), neural activity in spiking rate (Hz) or dF/F0 (middle) and activation deviation (rhs, ΔFrac). All cells depicted here were significant using ZETA. Fifth column: benchmark performance for neuronal inclusion percentage ('Real', green) and false alarm rate after jittering ('Jitter', red) for tests using mean-rate (black, μ) and ZETA (blue, Z). From top to bottom, we investigated several data sets: simulated Poisson neuron responses to on/off stimuli (top row), mouse V1 responses to drifting gratings recorded with a Neuronexus probe (second row) and GCaMP6f (third row), mouse superior colliculus responses to drifting gratings recorded with a Neuronexus probe (fourth row), and zebrafish retinal ganglion cell responses to on/off light flashes recorded with a perforated micro-electrode array (fifth row). Bar graphs show inclusion rate and 75th percentile confidence interval.

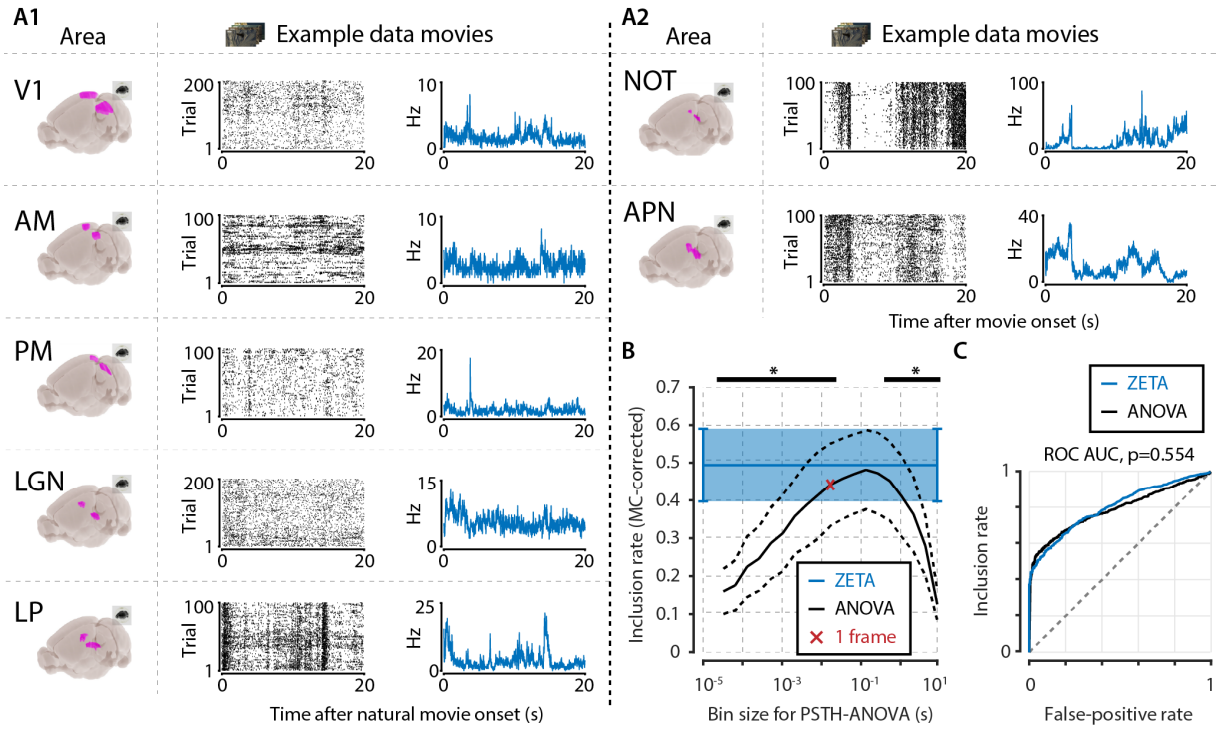

Figure 5. Determining neuronal responsiveness to natural movies. A) We recorded neuronal responses to natural movies (4 scenes that repeated every 20 seconds) from various visual brain areas: V1, AM, PM, LGN, LP, NOT, and APN. For each area, we show an example neuron's raster plot (left) and binned responses (right). All cells depicted here were significant using ZETA. B) We determined a neuron's responsiveness using a 1-way ANOVA over all binned responses (i.e., PSTH-level ANOVA) for various bin sizes and MC-corrected (black) and using the ZETA-test (blue). Curves show mean ± SEM over brain regions (n=7). ZETA shows an inclusion rate similar to the most optimal bin sizes (0.0333s – 0.2667s), and significantly higher than bin sizes of 0.0167s or shorter, as well as 0.533s or longer (*, FDR-corrected paired t-tests, $p<0.05$). C) An ROC analysis on all n=977 cells showed that the ZETA-test and a combined set of ANOVAs were similarly sensitive (ZETA-test AUC: 0.792 ± 0.011; ANOVAs AUC: 0.798 ± 0.010; z-test, p=0.554).

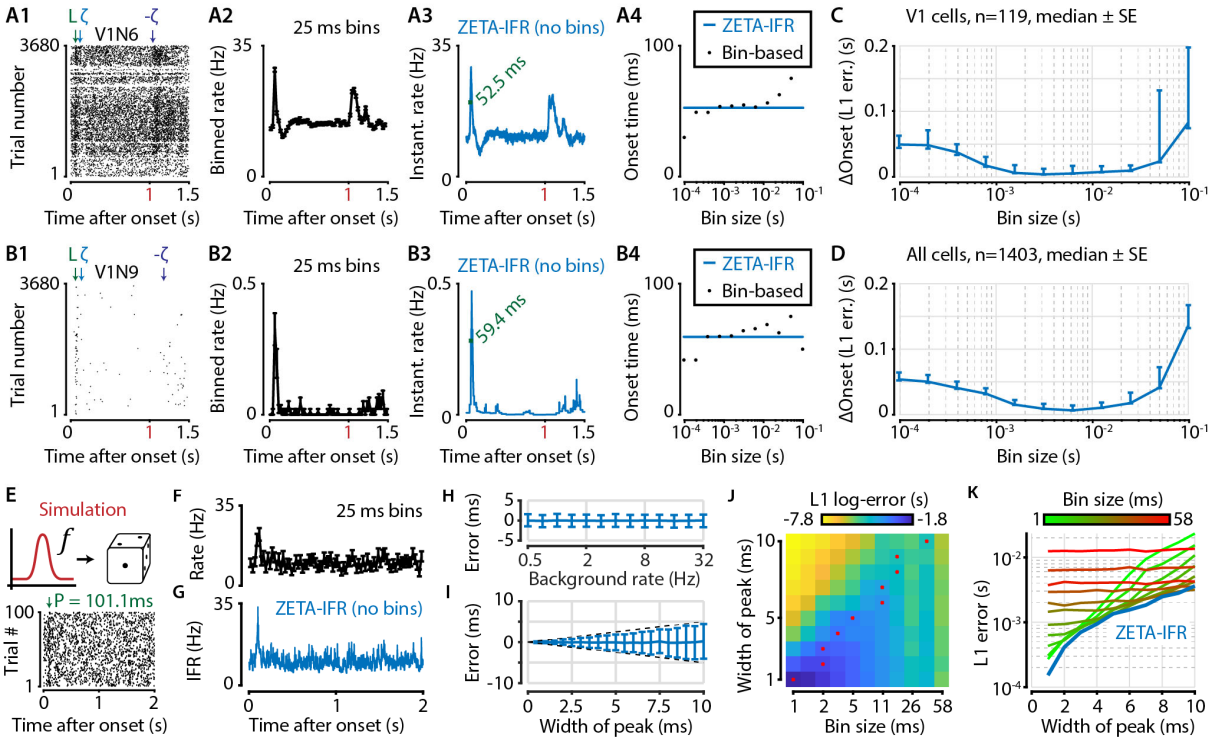

Figure 6. A ZETA-derived measure for instantaneous firing rates (IFR) avoids binning and allows more accurate latency determination than with PSTHs. A,B) Responses of two example V1 cells to drifting gratings. From left to right: 1) Raster plots showing the estimated onset latency ('L', green), and times of ZETA (blue) and –ZETA (purple). 2) Spiking rates using 25 ms bins. 3) Binning-free instantaneous spiking rates provide a much higher temporal resolution. Using the first crossing of half-peak firing rates, we determined the onset latencies of these cells to be 52.5 and 59.4 ms. 4) Estimated onset times using instantaneous firing rates (blue) or bins of durations from 0.1 to 100 ms (black). Onset latency estimates depend on the chosen bin size, and the optimal size varies across cells. C-D) The median (±SE) difference in onsets estimated by our instantaneous firing rates compared to that of various sized bins for V1 cells (C, n=119) and for cells from all brain regions (D, n=1403). Both C and D show the onsets estimated by the two methods were most similar for bin sizes between 1 – 10 ms. E-K) Benchmarking of peak detection using artificial Poisson neurons that show a transient peak. E) Example Poisson cell for a background rate of 10Hz and a peak-width of 10 ms. With the true peak at 100 ms, the estimation error here was 1.1 ms. F) Binning the cell's spiking response in 25-ms bins

reduces the peak height and temporal precision. G) Our instantaneous spiking rate preserves a sharper peak response and allows for a temporally accurate latency estimation. H) The detection of peak latencies is insensitive to realistic levels of a stationary Poisson background firing rate (13 base rates, 0.5 – 32 Hz). I) The mean error is unbiased, and the standard deviation in the onset peak latency estimate scales linearly with the width of the peak. Dotted lines show the real peak width. Graphs in H-I show mean ± sd. J) The error in peak latency estimation depends on both the bin width and the width of the neuron's peak response. Red crosses indicate the bin-width with the lowest error. K) Plotting the latency estimation error shows that different bin sizes are optimal for different peak-widths. The error based on our binning-less IFR (blue) is as low as the most optimal bin size, for any peak width.

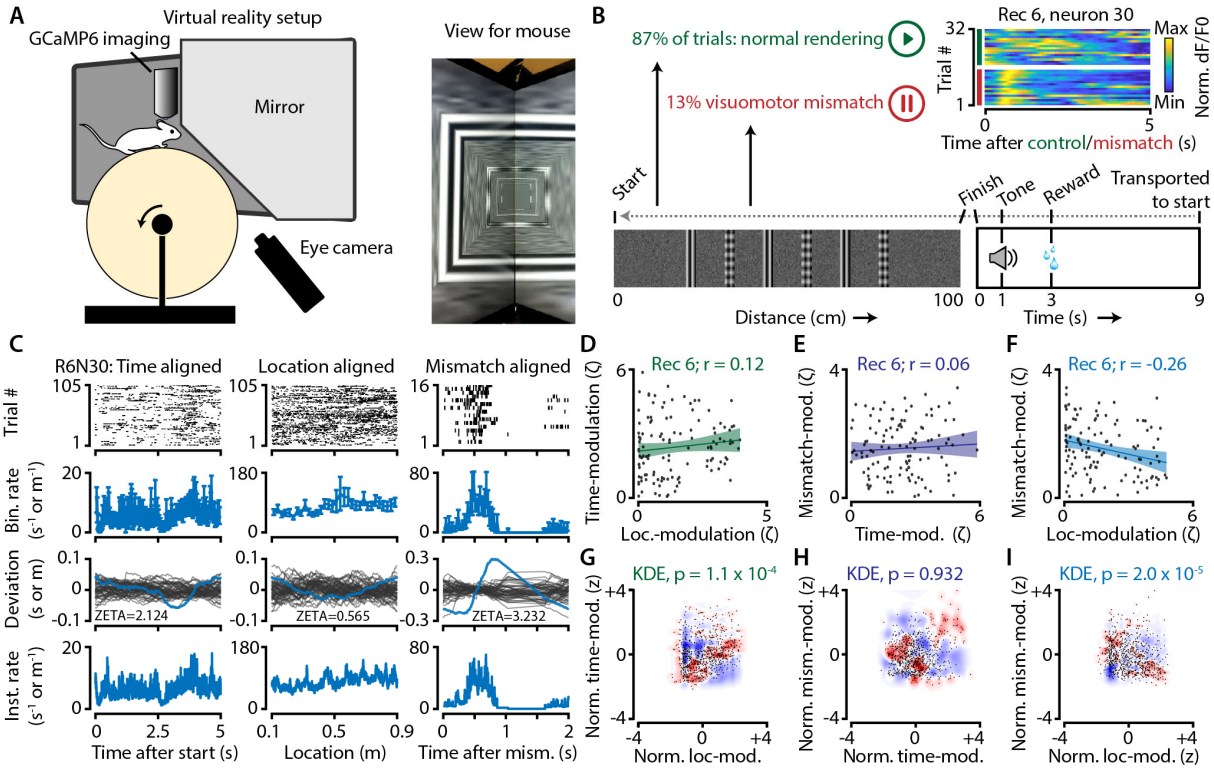

Figure 7. Neurons in V1 encode either visuomotor mismatch signals or spatial location. A) Schematic of setup showing mouse on running wheel (lhs) viewing a virtual tunnel (rhs). B) Trials consist of a 100cm linear track. One second after the mice ran to the end of the tunnel, an auditory stimulus signaled that a water reward would be delivered two seconds later. 6s after reward delivery, mice were transported to the start of the virtual tunnel. In a subset of trials, the rendering of the tunnel was paused at a random location, eliciting a visuomotor mismatch signal. Top right shows calcium imaging data for an example "mismatch neuron" during 16 control and 16 mismatch trials. C) Spiking data for example neuron obtained from exponential fits of the dF/F0 signals. Putative spikes were aligned to start (left), location of the animal on the track (middle), or mismatch onset (right). From top to bottom: raster plot of putative spike times; mean +/- SEM of firing rates over trials (n=105 trials, of which n=16 mismatch trials); spiking deviation underlying ZETA; instantaneous firing rate. D-I) Relationship between time-, location-, and mismatch-modulation. One point is one neuron. D-F) ZETA-scores for example recording 6 (N=120 neurons). G-I) Analysis using a kernel-density estimate (KDE) to test

whether joint-encoding of two features is more common than expected by chance (see supplementary figure 2). G) More neurons showed joint-encoding of both spatial and temporal location than expected by chance ($p=1.1 \times 10^{-4}$). H) Joint-encoding of spatial location and mismatch was not significantly different from chance ($p=0.932$). I) Location on the virtual track and visuomotor mismatch are less likely to be encoded by the same neuron than expected from chance ($p=2.0 \times 10^{-5}$).

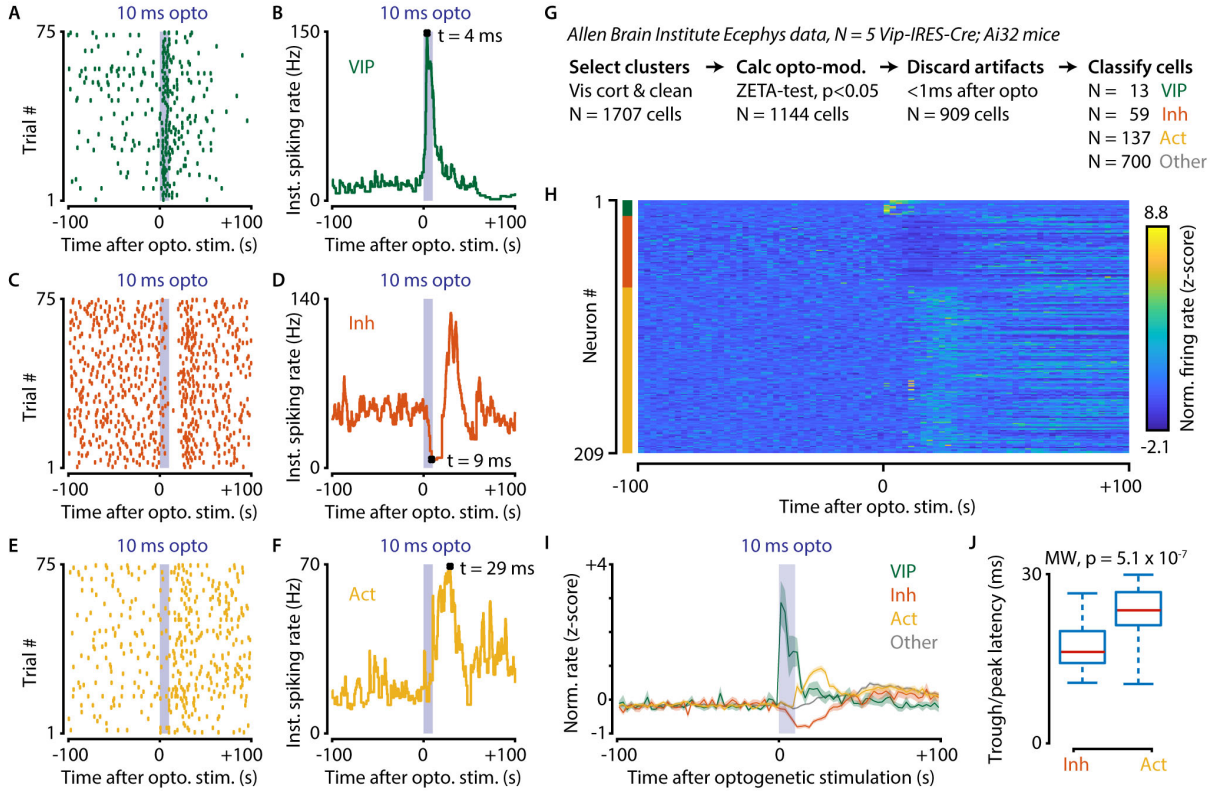

Figure 8. Optogenetic stimulation of VIP-expressing cells in mouse visual cortex causes short-latency inhibition and longer-latency disinhibition of the local neural circuit. A-F) Response of example cells classified as VIP (A,B), Inhibited (C,D) and Activated (E,F). G) Data were recorded in visual cortex from 5 Vip-Cre mice at the Allen Brain Institute. Cells were only included if the clustering quality was sufficient (1707 cells total). A ZETA-test included cells that were modulated within (-0.5, +0.5 s) after optogenetic stimulation (N=1144 cells). IFR peak- and trough-latency was computed and cells were discarded if their peak was earlier than 1 ms after optogenetic stimulation onset. Remaining cells were classified as VIP (N=13), Inhibited (N=59), Activated (N=137), or Other (N=700). H) Heat map showing normalized firing rate of VIP (top), Inhibited (middle), and Activated (bottom) cells. I) Mean ± SEM of PSTH (2.5 ms bin size) over all VIP (green), Inhibited (orange), Activated (yellow), and Other (grey) cells. J) Inhibited cells showed significantly lower mean IFR-peak latencies after optogenetic stimulation than Activated cells (Inh: 17.4 ms; Act: 23.0 ms; Mann-Whitney U-test, $p=5.1 \times 10^{-7}$).

Methods

ZETA

Well documented and easy-to-use Matlab and python code performing the procedures described in the following paragraphs can be found here: https://github.com/JorritMontijn/ZETA and https://github.com/JorritMontijn/zetapy.

We developed a timescale-free, binning-less statistical test for determining whether a neuron shows a time-locked modulation of spiking activity. It is derived from a metric that represents the reliability, as number of standard deviations away from chance, that the temporal density of spikes is non-random across trial repetitions. This metric, we call ZETA (ζ), can be computed on a vector of $i = [1 \ldots N]$ spike times $\boldsymbol{v}$, and a vector of $k = [1 \ldots q]$ event times (e.g., stimulus onsets), $\boldsymbol{w}$, using the following steps.

First, we make the spike times in $\boldsymbol{v}$ relative to the most recent stimulus onset, as when making a raster plot of spike times:

$$v_i = v_i - w_k, \text{ if } w_k < v_i < w_{k+1} \qquad 1$$

Next, we remove all spike times that are larger than a cut-off value τ, for example the trial duration, and add two artificial spikes at t=0 and t=τ to ensure coverage of the full epoch. We sort the n spike times in $\boldsymbol{v}$ such that $v_i < v_{i+1}$, and calculate the fractional position g_i, ranging from $1/n$ to 1, of each spike time in $\boldsymbol{v}$:

$$g_i = i/n \qquad 2$$

Therefore, another interpretation is that $\boldsymbol{g}$ represents a neuron's cumulative density function sampled at the spike times in $\boldsymbol{v}$. In order to quantify whether this distribution is different from our null hypothesis – i.e., that the neuron's firing rate is not modulated with respect to the stimulus onset – we compare this vector to a linear baseline density vector $\boldsymbol{b}$. If a neuron's spiking rate is constant, the cumulative density function is linear over time, and therefore the

expected fractional position of spike i at time v_i converges to the spike time divided by the trial duration τ as the number of events q increases:

$$\lim_{q \to \infty} b_i = v_i / \tau \qquad 3$$

The difference δ_i between g_i and b_i therefore gives a neuron's deviation from a temporally non-modulated spiking rate at time point v_i:

$$\delta_i = g_i - b_i \qquad 4$$

Finally, we create $\boldsymbol{d}$, a mean-normalized version of $\boldsymbol{\delta}$.

$$d_i = \delta_i - \bar{\delta} \qquad 5$$

Where

$$\bar{\delta} = \frac{1}{n} \sum_{i=1}^{n} \delta_i \qquad 6$$

We then define the Zenith of Event-based Time-locked Anomalies (ZETA, or ζ) as the most extreme value, i.e. the maximum of the absolute values:

$$\zeta \equiv \max(|\boldsymbol{d}|) \qquad 7$$

Null hypothesis for ZETA

Having calculated ZETA from the temporal deviation vector $\boldsymbol{d}$, we wish to quantify its statistical significance. First, we scale it such that its value is interpretable as a z-score. We therefore construct a null hypothesis distribution by repeating the above procedure P times with jittered event-times $\boldsymbol{w'}$, where we move each event time by a random sample drawn from the interval [-τ, τ]. This way, we calculate the chance of observing randomly high values in $\boldsymbol{d}$ without having to make assumptions about the underlying distribution of $\boldsymbol{d}$. However, a naïve approach would lead to difficulties here, as jittering $\boldsymbol{w}$ also changes the corresponding values of $\boldsymbol{v}$; and any jittered vector $\boldsymbol{d'}$ we obtain would be sampled at different times than the original vector $\boldsymbol{d}$. Therefore, we instead linearly interpolate the values of jittered fractional position vector $\boldsymbol{g'}$ at

the original spike times of $\boldsymbol{v}$. First, we construct a vector $\boldsymbol{f}$ of fractional spiking positions analogously to $\boldsymbol{g}$, but based on jittered event-times $\boldsymbol{w'}$:

$$f_i = i/n' \qquad 8$$

Note that we cannot simply take $\boldsymbol{g}$, as the total number of spikes n' in this jittered version is likely different from the original number of spikes n, because we only consider the spike times in the interval [0, τ] after the (jittered) event times. Next, we interpolate the values of $\boldsymbol{f}$ at sample times $\boldsymbol{v'}$ to the original sample times $\boldsymbol{v}$:

$$\boldsymbol{g}' = interp1(\boldsymbol{v}', \boldsymbol{f}, \boldsymbol{v}) \qquad 9$$

Here, *interp1* corresponds to a 1-dimensional piece-wise linear interpolation function. We repeat this process P times; where for each jitter iteration j, and calculate $\boldsymbol{\delta}'(j)$:

$$\boldsymbol{\delta}'(j) = \boldsymbol{g}'(j) - \boldsymbol{b} \qquad 10$$

Note that $\boldsymbol{b}$ is invariant with respect to the jitter iteration, as it is simply the n-element linear vector from $1/n$ to n. As before, we mean-normalize $\boldsymbol{\delta}'(j)$ to obtain a temporal deviation vector $\boldsymbol{d}'(j)$.

$$\boldsymbol{d}'(j) = \boldsymbol{\delta}'(j) - \overline{\delta}'(j) \qquad 11$$

Now we can define a null-hypothesis ZETA sample j as:

$$\zeta'(j) \equiv \max(|\boldsymbol{d}'(j)|) \qquad 12$$

Statistical significance of ZETA

Having constructed a way to generate samples from a null-hypothesis distribution, we are left with the task of using it in calculating the statistical significance of ZETA. If we had infinite samples, we could directly calculate the percentile of the empirical ζ from the null-distribution. However, as we have to work with a finite number of null-hypothesis samples, we need to estimate the true distribution. From extreme value theory we know that the distribution of

maximum values is known as a Gumbel distribution (Gumbel, 1941). Its cumulative density is given by:

$$F(x;\ m, \beta) = e^{-e^{-(x-m)/\beta}} \qquad 13$$

Here, x is the sample maximum (i.e., ζ), m is the mode, and β is the scale parameter. Therefore, we need to find m and β, which can be derived from the estimated sample mean and variance over jittered ZETAs of $\boldsymbol{\zeta'}$. The scale parameter β depends on the variance:

$$\beta = \frac{\sqrt{6 \cdot \mathrm{Var}(\boldsymbol{\zeta'})}}{\pi} \qquad 14$$

Then the mode can be computed from β and the mean $\overline{\zeta'}$:

$$m = \overline{\zeta'} - \beta\gamma \qquad 15$$

Here, γ is the Euler–Mascheroni constant ($\gamma \approx 0.577$). Now we can define the p-value by reading out the cumulative Gumbel distribution at ζ:

$$p = 1 - F(\zeta;\ m, \beta) \qquad 16$$

Finally, we can use p to obtain a corrected ZETA ζ_c that is interpretable as a z-score:

$$\zeta_c = \Phi^{-1}\left(1 - \frac{p}{2}\right) \qquad 17$$

Note that when we refer to ZETA or ζ in the rest of the manuscript, we mean the corrected version and its p-value as defined above.

Multi-scale derivatives of ZETA for latency detection

The ZETA-test indicates whether a neuron shows reliable deviations in spiking rate with respect to a particular series of events. The time of this maximum deviation, however, is not necessarily when the neuron shows its strongest firing rate modulation, but rather when the cumulative distribution reaches peak statistical significance. Therefore, in order to use the ZETA procedure to calculate the time of peaks in modulations (e.g., onset latencies), we should take the

derivative of the temporal deviation vector $\boldsymbol{d}$ underlying ZETA. A naïve approach with a simple spike-to-spike derivative unfortunately yields a curve with many spurious peaks. One solution would be to calculate the derivative over a larger time interval, but this comes at the expense of temporal resolution. Moreover, many different cell types exist with different dominant time constants. To balance temporal resolution and robustness, we therefore developed a multi-scale derivative procedure. First, we define a vector $\boldsymbol{t}$ of S timescales at which to compute derivatives. By default, we define the timescales to lie on a logarithmic scale with base 1.5, as this gave a reasonable trade-off between computational speed and accuracy. Base values closer to 1 will give more accurate results at the cost of computational speed. For base b and a trial duration of τ:

$$\boldsymbol{t} = b^{\boldsymbol{p}} \tag{18}$$

Where

$$\boldsymbol{p} = \{x \mid x \in \mathbb{Z}, \log_b 10^{-3} < x < \log_b \frac{\tau}{10}\} \tag{19}$$

The derivative $\dot{d}$ at spike i can then be defined for d_i and timescale t_k as:

$$\dot{d}_{i,k} = \frac{d_b - d_a}{v_b - v_a} \tag{20}$$

Where

$$a = \arg\max v_a \{v_a \mid v_a \in \boldsymbol{v}, v_a < v_i - \frac{t_k}{2}\} \tag{21}$$

$$b = \arg\min v_b \{v_b \mid v_b \in \boldsymbol{v}, v_b > v_i + \frac{t_k}{2}\}$$

Here, $\boldsymbol{v}$ are spike times, following the definition above. To avoid undefined edges, we set a and b to 1 and n respectively, iff $v_i \pm t_k/2$ falls outside the interval $[0, \tau]$. Taking the mean over all S timescales, we obtain an average of multi-scale derivatives $\boldsymbol{m}$,

$$m_i = \frac{1}{S}\sum_{k=1}^{S} \dot{d}_{i,k} \tag{22}$$

which has two important properties. First, long-timescale derivatives tend to 0, so there is a bias of $\boldsymbol{m}$ to more strongly follow shorter timescales. Secondly, random noise at the shortest timescales averages out over multiple short-timescale derivatives. Therefore, these two properties combined lead $\boldsymbol{m}$ to reflect the shortest timescales at which a real signal starts to emerge from random noise.

Calculation of high-resolution instantaneous firing rates

Another interesting property of the mean multi-scale derivative $\boldsymbol{m}$ is that it scales with the actual firing rate. In other words, it produces a time-locked neural activation curve, similar to a peri-stimulus time histogram (PSTH). If we properly rescale $\boldsymbol{m}$, we can therefore create an instantaneous firing rate metric with a temporal resolution that is only limited by the spike density.

Remember that the temporal deviation vector $\boldsymbol{\delta}$ itself is scaled to lie between -1 and +1, as its value depends on the difference between the fractional position of a spike (from 0 to 1) and the linear interval from $x,y=[0,0]$ to $[\tau,1]$. The theoretical lower limit of the multi-scale derivative is therefore $-1/\tau$. This can be illustrated as follows. Imagine a hypothetical neuron where all spikes are fired in an arbitrarily short interval close the start of each trial. This means that $\boldsymbol{\delta}$ rises from 0 to 1 in a short interval and from then decays linearly from 1 to 0 over an interval of τ. As $\dot{d}$ is defined between two spikes (including window edges 0 and τ), this means that the lowest possible value in this extreme case is the point between the last spike n and τ, i.e.:

$$\dot{d}_{\min} = \frac{d_\tau - d_n}{v_\tau - v_n} = \frac{0-1}{\tau - 0} = -1/\tau \qquad 23$$

The lowest possible firing rate is 0 Hz, which therefore corresponds to $-1/\tau$. An upper bound for $\dot{d}$ does not exist, as an arbitrarily short interval with a finite number of n spikes would lead to arbitrarily high $\dot{d}$:

$$\dot{d}_{\max} = \frac{d_j - d_i}{v_i - v_j} = \frac{1/n}{t - (t + 1/\infty)} = \infty \qquad 24$$

This is a desirable property, as the maximum instantaneous firing rate of a neuron is not theoretically constrained. Finally, the average firing rate of our metric should correspond to the real average firing rate in Hz ($n/\tau q$), where q is the number of events as defined above. We therefore define our instantaneous firing rate metric r as:

$$\boldsymbol{r} = \frac{n}{\tau \cdot q} \cdot \left(\frac{\boldsymbol{m} + 1/\tau}{\overline{m} + 1/\tau}\right) \qquad 25$$

Here, $\overline{m}$ is the weighted average of $\boldsymbol{m}$ by the inter-spike interval, such that the averaging occurs in the time domain and not the spike-number domain:

$$\overline{m} = \frac{1}{\tau} \sum_{i=2}^{n} \frac{m_{i-1} + m_i}{2} (v_{i-1} - v_i) \qquad 26$$

Mean-rate artificial Poisson neurons

We tested whether ZETA required short bursts of activity to work by generating artificial spike trains that only varied in mean-rate between a 1-second stimulus presentation and a 1-second inter-stimulus interval, and did not show onset peak responses. We created spike trains for 100 neurons with an orientation preference θ randomly sampled from a uniform distribution on the interval (0,π). The shape of the tuning curve was defined as the sum of two von Mises distributions centered at preferred orientation θ and θ+π with a concentration parameter of κ=5+ε, where ε was randomly sampled from a uniform distribution on the interval (0,5)). The baseline mean spiking rate μ_{base} was defined by randomly sampling from an exponential distribution with a mean of λ_{base} = 5 Hz. μ_{base} defined the trough of the neuron's tuning curve (i.e., the activity at θ+π/2 and θ-π/2) as well as the activity of the neuron when no stimulus was present. The firing rate for the preferred stimulus μ_{stim} was determined by similarly sampling from an exponential distribution with a mean of $\lambda_{stim} = \mu_{base}$ + 20 Hz. We generated spiking

activity for 160 trials (20 repetitions of 8 stimulus orientations: θ_{stim}= [0, 45, …, 315]). The average baseline firing across all n=10000 artificial neurons was therefore 5 Hz and the average firing rate during the preferred stimulus was 25 Hz. Spike times were generated for each trial-epoch (stimulus/baseline) independently by consecutively drawing inter-spike intervals from a Poisson distribution with $\lambda = 1/\mu$.

Artificial Poisson neurons for peak-latency benchmarking

We also addressed the question whether our instantaneous firing rate metric was sufficiently robust to allow accurate peak-time detection over a range of background firing rates and a range of peak widths. The procedure here was similar as above, with the exception that the firing rate during baseline and stimulus periods was identical: $\mu_{base} = \mu_{stim}$. We tested 13 base rates (0.5 – 32 Hz) and 19 jitter widths (1 – 10 ms in steps of 0.5). For each combination of base rate and jitter width, we generated 100 neurons and 100 trials per neuron. Peaks were added to the background activity by adding a single spike to half of all trials. Spike times were chosen by random sampling from a normal distribution with the standard deviation equal to the above jitter width and centered at 100 ± 10 ms after stimulus onset.

Figure panels 6J,K used slightly different parameters. Instead, we used 10 jitter widths (1 – 10 ms in steps of 1), simulated only a base rate of 32 Hz, used 160 trials for 1000 neurons, and compared the peak-latency detection using the ZETA-IFR with binning windows ranging from 1.00 ms to 57.67 ms; a logarithmic scale of base 1.5 with the exponent ranging from 0 to 10 in steps of 1.

Acquisition and preprocessing of laminar probe data (Neuronexus and Neuropixels)

We performed silicon probe recordings in six C57BL/6 mice, 2–7 months of age. Mice were housed in a 12 h/12 h dark/light cycle with ad libitum access to food and water. All experiments

were approved by the animal ethics committee of the Royal Netherlands Academy of Arts and Sciences, in compliance with all relevant ethical regulations.

Mice were habituated for 1-4 weeks before being implanted with a cranial bar used for head-fixation. Mice were anesthetized with isoflurane (3% induction, 1-1.5% maintenance in 50% O2) and injected subcutaneously with an analgesic and anti-inflammatory compound (Metacam, 2 mg/kg). The eyes were protected from drying by Cavasan eye ointment. They were moved to a stereotact with a thermal mat to keep their core temperature at 37 degrees Celsius, and the fur on their heads was removed. Once anesthesia was sufficiently deep, as indicated by the absence of a toe-pinch reflex, we applied lidocaine locally on the skin of the head and sterilized it with 70%-ethanol or betadine. The skin was removed, the skull cleaned, and a small metal rod was fixed to the skull anterior of bregma with the use of blue-light curing dental cement. If necessary, we sutured the skin, and let the mice recover for 2-7 days.

After the mice recovered as indicated by a return to their pre-operative weight, mice were habituated to sitting head-fixed in the electrophysiology rig for 3-10 days. Once habituated, they underwent a craniotomy surgery, following the same preparatory steps as described above. Before performing the craniotomy (1.5-3 mm in diameter), we first constructed a small ring of dental cement so the brain could be bathed in saline during recordings to avoid tissue desiccation. Once the craniotomy was complete, this ring was filled with sterile silicone, and the animals were left to recover for at least 16 hours. Over the next 1-7 days, we performed repeated-insertion recordings using either NeuroNexus or Neuropixels silicon probes. For a subset of animals, we dipped the probe into DiI on the last day of recording, and perfused the animal to perform post-hoc tracing of the electrode recording locations. For the remaining recordings, the probe position was determined using anatomical landmarks (i.e., bregma and lambda).

NeuroNexus recordings were performed using either a Tucker-Davis Technologies digitizer and custom-written MATLAB code as described previously (Ahmadlou and Heimel, 2015). Neuropixels recordings were performed using a National Instruments I/O PXIe-6341 module and SpikeGLX (https://github.com/billkarsh/SpikeGLX). Visual stimulation was performed as described previously (Montijn et al., 2016a), and synchronized with high accuracy (<1 ms) using photodiode signals that recorded visual stimulus onsets. Spikes were sorted post-hoc using Kilosort2 (https://github.com/MouseLand/Kilosort2) and only clusters of sufficient quality, as defined by Kilosort2's default threshold, were included for further analysis. High-quality clusters (i.e., putative neurons) were assigned a brain region using the AllenCCF MATLAB toolbox (https://github.com/cortex-lab/allenCCF), which automatically calculates a neuron's anatomical position based on penetration location, angle, and depth of the silicon probe. Abbreviations for brain areas mostly follow the Allen Brain Atlas area codes. V1: primary visual cortex. AM: anteromedial visual cortex. PM: posteromedial visual cortex. LGN: lateral geniculate nucleus. LP: lateral posterior nucleus. NOT: nucleus of the optic tract. APN: anterior pretectal nucleus. SC: superior colliculus. All code used in laminar probe data acquisition and pre-processing is available online (https://github.com/JorritMontijn/Acquipix).

Visual stimulus parameters

Visual stimuli during Neuronexus and Neuropixels recordings were shown at 60 Hz on a 51 by 29 cm Dell screen at 17 cm distance from the animal's eyes, using Psychtoolbox 3 in Matlab. Drifting gratings were displayed within a 120 visual-degree diameter window with 2 visual-degree cosine edge that faded smoothly into a neutral-grey background. Drifting gratings were shown in 24 directions: [0, 15, … 345] degrees at a spatial frequency of 0.05 cycles per degree and a temporal frequency of 1 cycle per second. Natural movies were 20 s long and consisted of 4 distinct scenes taken from the BBC nature documentary Earthflight (Montijn et al., 2016a).

Acquisition and preprocessing of retinal multi-electrode array data

Adult (>1year) zebrafish (*Danio rerio*), were dark-adapted for at least 1 h. Then under IR illumination fish were euthanized by rapid immersion in ice-cold water, the eyes removed, the retina isolated and placed photoreceptor side up on a perforated 60 electrode array (60pMEA200/30iR-Ti using a MEA2100 system: Multichannel systems, Reutlingen, Germany) in a recording chamber mounted on an Nikon Optiphot-2 upright microscope and viewed under IR with an Olympus 2x objective and video camera (Abus TVCC 20530). Room temperature Ames' medium (Sigma-Aldrich) gassed with a mixture of O2 and CO2 at pH of 7.6 continuously superfused the MEA recording chamber.

Extracellular multiunit GC activity was recorded at 25 kHz in MC rack (Multichannel systems, Reutlingen, Germany), zero-phase bandpass filtered (250–6,250 Hz) with a fourth-order Butterworth filter in Matlab (MathWorks, Natick, MA, USA), and sorted into single-unit activity with "offline spike sorter" (Plexon, Dallas, TX, USA). Spikes were detected using an amplitude threshold > $4\sigma_n$ where σ_n is an estimation of the background noise

$$\sigma_n = median\left(\frac{|x|}{0.6745}\right) \qquad 27$$

with x being the bandpass-filtered signal (Quiroga et al., 2004). The detected spikes were manually sorted into single units based on the first two principal components versus time.

The light stimulus consisted of a 500 ms full field light flash, preceded and followed by a 500- and 1000- ms period of darkness, generated using Psychophysics Toolbox Version 3 (Brainard, 1997; Kleiner et al., 2007), and repeated either 50 or 100 times. Stimuli were projected onto the retina from the photoreceptor side by a DLP projector (Light Crafter 4500, Wintech, Carlsbad, CA, USA) using a custom-built 2x water immersion objective. Only white light stimuli were used. The "dark" light intensity was 6 μW/m2, and the maximal "light" intensity was 176.2 μW/m2.

Acquisition and preprocessing of calcium imaging data

Drifting grating responses were recorded as previously described (Montijn et al., 2016a) and putative spike times were extracted using an exponential fitting algorithm (Montijn et al., 2016b). All code used in pre-processing of drifting grating calcium imaging data is available online (https://github.com/JorritMontijn/Preprocessing_Toolbox).

Virtual corridor experiments were performed on adult male RCFL-tdTOM x PV FlpO x VIP2r-cre (n = 1) and VIP2r-cre mice (n = 3). During water restriction, the weight of the animals was carefully monitored and kept stable through powdered milk intake (10% diluted in drinking water) in the corridor or additionally fed solid drink. Mice were kept on a reversed day/night schedule (12h/12h) and experiments were performed in their active dark phase.

For virus injection and chronic window implantation procedures, animals were anesthetized with isoflurane at 5% for induction and 1.0-1.5% for maintenance. Metacam (1 mg/kg) was administered subcutaneously for systemic analgesia and dexamethasone (8 mg/kg) was administered to prevent brain swelling. The eyes were covered throughout the surgeries with Cavasan ointment. For virus injections, three small holes were drilled around the center of the right V1 (2.9 mm lateral from midline, 0.5 mm anterior from lambda). Animals were injected with saline-diluted AAV2/1.hSyn.GCaMP6f.WPRE.SV40 (Addgene) at a depth of 250 and 550 µm (final concentration, 10^{12} viral particles/ml; 36.8 nl/depth/location). Mice were allowed to recover for a minimum of two weeks in their home cage before window implantations. For chronic window implantations, a custom-made head bar was positioned above V1 and fixed on the skull with dental cement (Kerr). A circular 3 mm craniotomy was made on the area of injection while the dura was kept intact. A double glass window (3 + 4 mm diameter) was placed inside the craniotomy and fixed on the skull with dental cement (Kerr). Animals were allowed to recover for a minimum of one week before training began.

Before imaging, mice were trained to being head fixed and run through a virtual corridor on a custom-made treadmill (Fig. 7A). Running speed of the mice was measured with a rotary encoder and processed using an Arduino and Matlab. Absolute running speed was used to render the virtual corridor in real-time. The left half of the virtual corridor was rendered on a gamma-corrected monitor (Dell) placed under an angle of 45° with a mirror to create the perception of a symmetrical corridor. The virtual corridor was written in Matlab using OpenGL and Psychophysics Toolbox 3 and contained a 100 cm black and white Gaussian noise texture with overlying visual stimuli. We used three vertical gratings and three checkerboard stimuli positioned 11 cm apart, at locations 22-77 cm (Fig. 7B). After completing a run in the corridor, mice were immediately shown a luminance-matched gray screen followed by a 0.5 s auditory cue (8 kHz) 1 second later. Two seconds after this cue, the mice received a 5-10 µl milk reward. Mice were trained for ~10 sessions until they completed up to ~150 trials during imaging sessions.

Imaging was performed with a two-photon microscope (Neurolabware) equipped with a Ti-sapphire laser (Mai-Tai 'Deepsee', Spectraphysics; wavelength, 920 nm) and a 16x, 0.8 NA water immersion objective (Nikon). The microscope was controlled by Scanbox (Neurolabware) running on Matlab. One one mouse we performed dual-plane imaging at 15.5 Hz/plane using an electrically tunable lens (OptoTune). The other three mice were imaged in a single plane at 15.5 Hz. During the pre-processing stage, we discarded all interneurons and only included putative pyramidal cells for further analysis.

Analysis of virtual corridor joint-feature encoding

For each neuron, we calculated ZETA-scores after aligning the putative spike times to trial start ('time-aligned'), converting spike times to locations and aligning them to the beginning of the corridor ('location-aligned'), or aligning the spike times to a visuomotor mismatch event, where

rendering of the virtual corridor was paused for 500 ms ('mismatch-aligned') (fig. 7C). We call the resulting values time-modulation, location-modulation and mismatch-modulation respectively. In the case of the spatial "location-aligned" analysis, we furthermore discarded the first and last 10% of the track to avoid including the start and end box locations. We investigated whether there was a relationship between the modulation values that neurons showed for these three different features in two ways. First, we simply calculated the Pearson correlation at the level of a single recording between time/location, time/mismatch, and location/mismatch modulations (fig. 7D-F). A one-sample t-test on these correlation values showed there was a significant, positive correlation for time/location and a significant, negative correlation for location/mismatch. However, this analysis does not directly answer the question whether the joint encoding of two features at the level of single neurons is more or less likely than we would expect by chance.

We therefore z-scored the modulation values per recording and pooled all neurons. We then used a kernel-density estimator (KDE) to construct a probability density estimate for the distribution of modulation values for time, location, and mismatch (supplemental figure 2). Combining these distributions for a pair of features, we obtain a null-hypothesis distribution of what the joint-feature-encoding distribution would look like if the modulation-scores are independent of each other. Using the same bandwidths as for the single-feature KDE-distributions, we smoothed the real data so we could directly compare which regions are over- or under-represented in the real data, producing the heat maps in fig. 7G-I. To quantify this over/underrepresentation, we counted the number of neurons in the upper-right quadrant, where neurons lie that have high modulation scores for both features. We compared the real count to the distribution obtained from 10,000 random samples taken from the joint-feature KDE-derived null-hypothesis distribution, where there is no correlation between the two feature modulation scores. As shown in supplemental figure 2 and fig. 7G-I, time and location-

modulation scores were more often both high in single neurons than expected from chance. Moreover, there were fewer neurons that showed a joint encoding of both spatial location and visuomotor mismatch than expected from chance.

VIP-optogenetics analysis of Allen Brain Institute Ecephys data

Detailed information on experimental and data acquisition procedures can be found online at the Allen Brain Institute website: https://portal.brain-map.org/explore/circuits/visual-coding-neuropixels. We used data from 5 Vip-IRES-Cre; Ai32 mice that underwent laser-based optogenetic stimulation. We pre-selected 1706 clusters that were recorded in visually-responsive cortex (AL, AM, PM, L, V1, RL, or MMP) and were of sufficient quality, specifically: KiloSort2 tagged the cluster as "good", ISI violations were under 0.5, amplitude thresholds under 0.1, and presence ratios over 0.9. The next step was to cull this population to only cells that showed modulated spiking with respect to the onset of optogenetic stimulation. We therefore included N=1144 cells that showed $p<0.05$ with a ZETA-test on the window between (-0.5, +0.5 s) after optogenetic stimulation. We calculated the peak-response latency using our IFR method and discarded cells with peak responses earlier than 1 ms after optogenetic stimulation onset (N=909). The remaining cells were classified as VIP if their IFR peak latency was < 10 ms (N=13), as Inhibited if their firing rate was significantly lower during 10-30 ms after optogenetic stimulation than during the 50 ms preceding optogenetic onsets (N=59), as Activated if their rate was significantly higher during 10-30 ms after optogenetic onset than during the pre-stimulus baseline (N=137), and otherwise as Other (N=700). The large size of this latter group could be explained by many cells being visually-responsive to the blue laser light. We chose a 10-30 ms window to compare the IFR peak/trough latencies of Activated and Inhibited cells, as 10 ms was the time of laser offset, and visual responses start to emerge in visual cortex after about 30-50 ms.

References

Ahmadlou, M., and Heimel, J.A. (2015). Preference for concentric orientations in the mouse superior colliculus. Nat. Commun. *6*, ncomms7773.

Attinger, A., Wang, B., and Keller, G.B. (2017). Visuomotor Coupling Shapes the Functional Development of Mouse Visual Cortex. Cell *169*, 1291-1302.e14.

Bartolo, R., Saunders, R.C., Mitz, A.R., and Averbeck, B.B. (2020). Information-Limiting Correlations in Large Neural Populations. J. Neurosci. *40*, 1668–1678.

Bastos, A.M., Usrey, W.M., Adams, R.A., Mangun, G.R., Fries, P., and Friston, K.J. (2012). Canonical Microcircuits for Predictive Coding. Neuron *76*, 695–711.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. *10*, 433–436.

Button, K.S., Ioannidis, J.P.A., Mokrysz, C., Nosek, B.A., Flint, J., Robinson, E.S.J., and Munafò, M.R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. Nat. Rev. Neurosci. *14*, 365–376.

Calders, T., and Jaroszewicz, S. (2007). Efficient AUC Optimization for Classification. In Knowledge Discovery in Databases: PKDD 2007, J.N. Kok, J. Koronacki, R. Lopez de Mantaras, S. Matwin, D. Mladenič, and A. Skowron, eds. (Berlin, Heidelberg: Springer Berlin Heidelberg), pp. 42–53.

Douglas, R.J., Martin, K.A., and Whitteridge, D. (1989). A canonical microcircuit for neocortex. Neural Comput. *1*, 480–488.

Friston, K. (2005). A theory of cortical responses. Philos. Trans. R. Soc. B Biol. Sci. *360*, 815–836.

Gregory, R.L., Longuet-Higgins, H.C., and Sutherland, N.S. (1980). Perceptions as hypotheses. Philos. Trans. R. Soc. Lond. B Biol. Sci. *290*, 181–197.

Gumbel, E.J. (1941). The Return Period of Flood Flows. Ann. Math. Stat. *12*, 163–190.

Head, M.L., Holman, L., Lanfear, R., Kahn, A.T., and Jennions, M.D. (2015). The Extent and Consequences of P-Hacking in Science. PLoS Biol. *13*.

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

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

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

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., and Broussard, C. (2007). What's new in psychtoolbox-3. Perception *36*, 1–16.

Knöpfel, T., and Song, C. (2019). Optical voltage imaging in neurons: moving from technology development to practical tool. Nat. Rev. Neurosci. *20*, 719–727.

Kok, P., Bains, L.J., van Mourik, T., Norris, D.G., and de Lange, F.P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. Curr. Biol. *26*, 371–376.

Lau, C., Ng, L., Thompson, C., Pathak, S., Kuan, L., Jones, A., and Hawrylycz, M. (2008). Exploration and visualization of gene expression with neuroanatomy in the adult mouse brain. BMC Bioinformatics *9*, 153.

Leinweber, M., Ward, D.R., Sobczak, J.M., Attinger, A., and Keller, G.B. (2017). A Sensorimotor Circuit in Mouse Cortex for Visual Flow Predictions. Neuron *95*, 1420-1432.e5.

Markov, N.T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., Lamy, C., Misery, P., Giroud, P., Ullman, S., et al. (2014). Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. J. Comp. Neurol. *522*, 225–259.

Mazurek, M., Kager, M., and Van Hooser, S.D. (2014). Robust quantification of orientation selectivity and direction selectivity. Front. Neural Circuits *8*.

Mesa, N., Waters, J., and Vries, S.E.J. de (2020). The effect of inclusion criteria on the functional properties reported in mouse visual cortex. BioRxiv 2020.05.12.091595.

Montijn, J.S., Meijer, G.T., Lansink, C.S., and Pennartz, C.M.A. (2016a). Population-Level Neural Codes Are Robust to Single-Neuron Variability from a Multidimensional Coding Perspective. Cell Rep. *16*, 2486–2498.

Montijn, J.S., Olcese, U., and Pennartz, C.M.A. (2016b). Visual Stimulus Detection Correlates with the Consistency of Temporal Sequences within Stereotyped Events of V1 Neuronal Population Activity. J. Neurosci. *36*, 8624–8640.

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

Niell, C.M., and Stryker, M.P. (2008). Highly Selective Receptive Fields in Mouse Visual Cortex. J. Neurosci. *28*, 7520–7536.

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

Olshausen, B.A., and Field, D.J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by VI? Vision Res. *37*, 3311–3326.

Ooyen, A., and Roelfsema, P. (2003). A biologically plausible implementation of error-backpropagation for classification tasks.

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

Palm, G., Aertsen, A.M.H.J., and Gerstein, G.L. (1988). On the significance of correlations among neuronal spike trains. Biol. Cybern. *59*, 1–11.

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

Pi, H.-J., Hangya, B., Kvitsiani, D., Sanders, J.I., Huang, Z.J., and Kepecs, A. (2013). Cortical interneurons that specialize in disinhibitory control. Nature *503*, 521–524.

Poort, J., Raudies, F., Wannig, A., Lamme, V.A.F., Neumann, H., and Roelfsema, P.R. (2012). The Role of Attention in Figure-Ground Segregation in Areas V1 and V4 of the Visual Cortex. Neuron *75*, 143–156.

Quiroga, R.Q., Nadasdy, Z., and Ben-Shaul, Y. (2004). Unsupervised spike detection and sorting with wavelets and superparamagnetic clustering. Neural Comput. *16*, 1661–1687.

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

Ringach, D.L., Shapley, R.M., and Hawken, M.J. (2002). Orientation Selectivity in Macaque V1: Diversity and Laminar Dependence. J. Neurosci. *22*, 5639–5651.

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

Self, M.W., van Kerkoerle, T., Goebel, R., and Roelfsema, P.R. (2019). Benchmarking laminar fMRI: Neuronal spiking and synaptic activity during top-down and bottom-up processing in the different layers of cortex. NeuroImage *197*, 806–817.

Semedo, J.D., Zandvakili, A., Machens, C.K., Yu, B.M., and Kohn, A. (2019). Cortical Areas Interact through a Communication Subspace. Neuron *102*, 249-259.e4.

Shuler, M.G., and Bear, M.F. (2006). Reward timing in the primary visual cortex. Science *311*, 1606–1609.

Siegle, J.H., Jia, X., Durand, S., Gale, S., Bennett, C., Graddis, N., Heller, G., Ramirez, T.K., Choi, H., Luviano, J.A., et al. (2019). A survey of spiking activity reveals a functional hierarchy of mouse corticothalamic visual areas. BioRxiv 805010.

Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273.

Vinje, W.E., and Gallant, J.L. (2000). Sparse Coding and Decorrelation in Primary Visual Cortex During Natural Vision. Science *287*, 1273–1276.

Whittington, J.C.R., and Bogacz, R. (2019). Theories of Error Back-Propagation in the Brain. Trends Cogn. Sci. *23*, 235–250.

Xing, D., Ringach, D.L., Shapley, R., and Hawken, M.J. (2004). Correlation of local and global orientation and spatial frequency tuning in macaque V1. J. Physiol. *557*, 923–933.