

RESEARCH ARTICLE

Recurrent interactions can explain the variance in single trial responses

Subhodh Kotekal^{1☯✉}, Jason N. MacLean^{1,2,3☯✉*}

1 Department of Neurobiology, University of Chicago, Chicago, Illinois, United States of America, **2** Committee on Computational Neuroscience, University of Chicago, Chicago, Illinois, United States of America, **3** Grossman Institute for Neuroscience, Quantitative Biology and Human Behavior, Chicago, Illinois, United States of America

☯ These authors contributed equally to this work.

✉ Current address: Department of Statistics, University of Chicago, Chicago, Illinois, United States of America

* jmaclean@uchicago.edu



OPEN ACCESS

Citation: Kotekal S, MacLean JN (2020) Recurrent interactions can explain the variance in single trial responses. *PLoS Comput Biol* 16(1): e1007591. <https://doi.org/10.1371/journal.pcbi.1007591>

Editor: Leyla Isik, MIT, UNITED STATES

Received: June 3, 2019

Accepted: December 9, 2019

Published: January 30, 2020

Copyright: © 2020 Kotekal, MacLean. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The data underlying the results presented in this study were previously published in Dechery, J. B., & MacLean, J. N. (2018). "Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons." *PLOS Computational Biology*, 14(5), 1–23. <https://doi.org/10.1371/journal.pcbi.1006153>. Additional files pertaining to model output are available from figshare at: https://figshare.com/articles/models_and_partial_correlations_zip/11413932.

Funding: This work was supported by National Institutes of Health (NIH) grant R01EY022338 (JNM). The funders had no role in study design,

Abstract

To develop a complete description of sensory encoding, it is necessary to account for trial-to-trial variability in cortical neurons. Using a linear model with terms corresponding to the visual stimulus, mouse running speed, and experimentally measured neuronal correlations, we modeled short term dynamics of L2/3 murine visual cortical neurons to evaluate the relative importance of each factor to neuronal variability within single trials. We find single trial predictions improve most when conditioning on the experimentally measured local correlations in comparison to predictions based on the stimulus or running speed. Specifically, accurate predictions are driven by positively co-varying and synchronously active functional groups of neurons. Including functional groups in the model enhances decoding accuracy of sensory information compared to a model that assumes neuronal independence. Functional groups, in encoding and decoding frameworks, provide an operational definition of Hebbian assemblies in which local correlations largely explain neuronal responses on individual trials.

Author summary

V1 neurons exhibit substantial response variability while being massively interconnected, complicating the development of a comprehensive theory of visual encoding. Some work has pointed to a global description, in which stimuli and spontaneous behaviors are the main drivers of neuronal responses. However, other work concerning the retina and elsewhere suggest that local correlations between neighboring neurons shape a target neuron's single trial responses to a much greater extent than global variables. In this work, we compare and evaluate local and global descriptions of variable single trial responses in mouse V1. We compute pairwise correlations of neuronal activity after accounting for global variables, and define the "functional group" of a neuron to be those neurons exhibiting correlated activity over short time scales. We find that the knowledge of the activity of a target neuron's functional group enables much more accurate predictions of single

data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

trial responses as compared to knowledge of the stimulus condition or mouse locomotion. Moreover, we show that the informative correlations are strong, positive, and are exhibited on short time scales. Finally, we find that local correlations enable accurate decoding of the visual stimulus. These results point to a local description of the visual system over short time scales.

Introduction

The earliest single unit recordings in awake primary visual cortex demonstrated that many of the recorded action potentials could not be readily attributed to visual stimulus [1]. Across visual cortex, single neuron responses to repeated presentations of a given stimulus exhibit high variability from presentation to presentation [2–8]. Response variability is largely restricted to neocortex [9] and is shared across the neocortical neuronal population [10,11]. However, neurons from a physiological perspective are capable of being highly reliable [12,13] suggesting that variance arises primarily from synaptic inputs [14,15] and from extraretinal factors that are not visual stimulus such as arousal and locomotion [5,16].

Neuronal variability can be taken to be purely noise [17], insofar as it is detrimental to the stability of sensory representation. Alternatively, single trial variability may reflect ongoing cortical dynamics associated with sensory processing and representation [18,19]. Shared variability may similarly reflect ongoing dynamics [20] as pairwise correlations in a population of neurons can affect sensory representation [21–23]. Neurons are highly interconnected and thus complex network interactions likely shape the activity of neurons [4,24]. In visual cortex, visually tuned neurons with similar stimulus selectivity are more likely to be synaptically interconnected [25] and these connections manifest as specific motifs [24] which coordinate synaptic integration [26]. Moreover, populations of neurons have been shown to exhibit higher-order state correlations [23,27], and topological network features have been shown to shape spike propagation and information transfer [26,28,29]. However, the relative role of local synaptic connectivity, as compared to stimulus-related input or more global variables such as locomotion, in the generation of a sensory representation, particularly in real time, remains unclear. A comprehensive theory of primary visual cortex must encompass local correlations instead of treating neurons as independent units.

The cell assembly hypothesis, laid out by Donald Hebb, proposes that neurons are organized into mutually excitable groups called “assemblies” which are strongly coactive [19,30]. Under this hypothesis, sensory representations and cortical state dynamics manifest as sequences of assembly activations and internal state dynamics need not be deterministically tied to stimulus from trial-to-trial [9,18,19,30]. Importantly, the assembly hypothesis suggests that single trial dynamics are strongly influenced by the group of coactive neurons rather than entirely by the external stimulus.

In this work, we investigate whether the variance in V1 neuronal activity over the time course of single trials is best explained by visual stimuli, locomotion, or by coactive groups of V1 neurons. We define coactive groups of neurons, which we term as “functional groups”, by local pairwise correlations of activity within short time intervals after accounting for stimulus and running effects. We used a linear model (LM) on calcium imaging data of visually-evoked activity recorded from V1 in freely running mice. In the model, we relate the target neuron’s activity during single trials to the functional group’s activity, stimulus condition, and mouse running speed. Functional group coactivity is understood to be informative of a target neuron’s activity if inclusion of coupling in the LM improves predictions of the target’s neuron

activity, which is suggestive of the importance of internal circuit dynamics in shaping single trials [19,31]. With the functional group as an operational definition of an assembly, we provide concrete descriptions of the timescales, numerical scale, correlation structure, and computational capabilities of assemblies.

Results

Trial-to-trial, the activity of individual neurons in visual cortex is highly variable despite ostensibly identical conditions [7]. To evaluate local and global effects on single trial variability, we built an encoding model using a linear model (LM), comprised of three terms: neuron coupling, stimulus and running speed. To determine each term's contribution, we constructed multiple restricted LMs. Each excluded a term of interest and we compared test set predictions to those of the unrestricted LM (i.e. the model containing all three terms). Relative change in a restricted model's test set performance indicated the excluded term's importance in modelling single trial activity. For example, a large increase in a restricted model's test set error indicated that the excluded variable was informative to single trial activity. A marginal change indicated that the excluded variable was uninformative. These LMs are inspired by the large amount of existing work employing generalized linear models (GLMs).

Notably, we used prediction performance rather than the LM coefficients to determine the importance of model terms. The full set of variables that truly explain V1 single trial activity is not *a priori* known, nor experimentally observable. Coefficient estimates would be different if the omitted variables were included in the model, and may lead to erroneous conclusions; this is omitted variable bias [32]. Consequently, we used the unrestricted LM's mean squared error (MSE) on the test set data as a benchmark against which we compared each restricted model, and thus evaluated the importance each variable.

Mouse visual cortical data

The imaging data used throughout this manuscript was described in a previous study [4]. Briefly, the activity of L2/3 excitatory neurons expressing GCaMP6s in mouse visual cortex (73–347 neurons; $n = 8$ animals; 21 distinct fields of view; Fig 1A) were imaged using two-photon laser scanning microscopy (25–33 Hz; [4,33]). Mice were awake and allowed to freely run on a linear treadmill while viewing drifting square-wave gratings presented in 12 directions in pseudo-random order interleaved with mean-luminance matched grey screen; each trial was a 5-minute block of stimulus presentation in this format.

Experimentally measuring coupling coefficients for use in a linear model

Visual stimuli [4] were repeatedly presented in identical 5-minute blocks. Partial correlations, i.e. *functional weights*, were computed for each pair of neurons (Fig 1B), where each functional weight captured the reliability of coactivity after accounting for shared changes in activity due to stimulus and running speed. This measure is analogous to “noise correlations” [10], but also allowed us to account for the running-induced changes in circuit activity [16]. To summarize the temporal component of coactivity, we computed the mean traces of each neuron and constructed a cross-correlogram, the maximum value of which denoted the *functional lag* between the neurons. Functional connections were often not symmetric, and the direction of the lag gives a direction of the connection. The associated functional lags most frequently were 0 frames (Fig 1Biii) and frequency decreased rapidly with lag amplitude; most coupling was bidirectional (lag 0) (mean = 50.1%, sem = 2.8%, mean over $n = 21$ datasets). For each neuron, the set of other neurons with functional connections directed towards the given neuron was called its *functional group*. The average number of incoming functional connections for a target

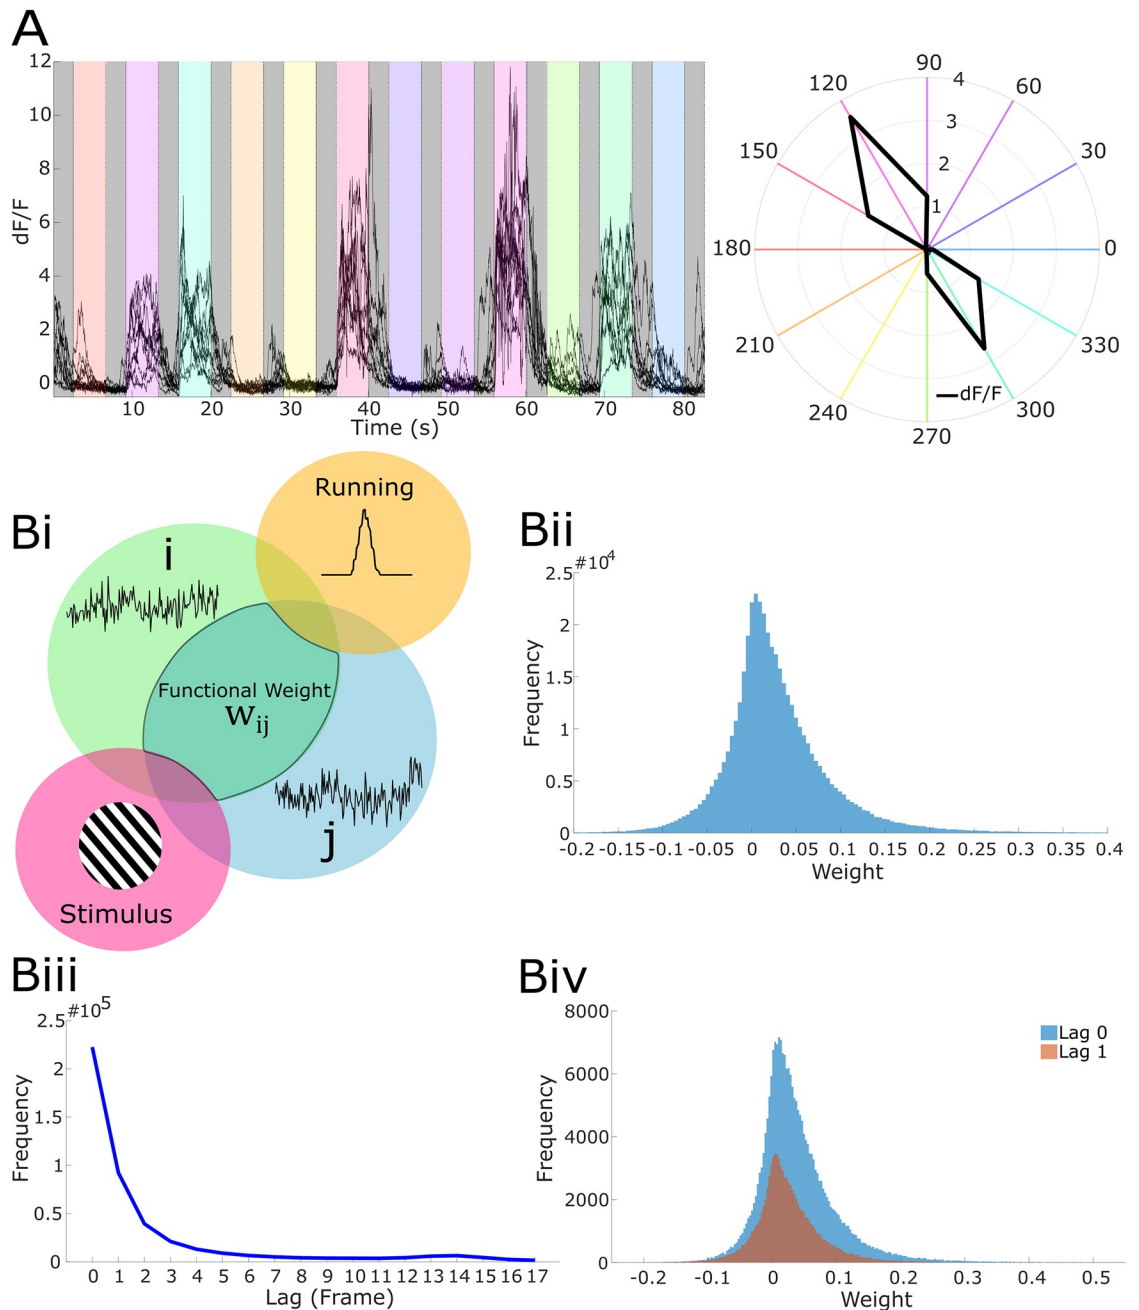


Fig 1. Characterizing local neuron correlations. (A) Left, anecdotal example fluorescence traces of neuronal responses to oriented drifting gratings. Right, polar plot of the average response of anecdotal neuron averaged across time bins and stimulus presentations. Colors in both panels correspond to different grating drift directions. (Bi) Diagram of partial correlation measure used for coupling in linear models (LMs); the partial correlation between neurons i and j is the correlation in the activities after accounting for stimulus and running induced coactivity. In the diagram, the partial correlation is denoted and outlined as the region of intersection between the regions of neurons i and j , but not including the regions of stimulus and running. (Bii) Functional weight (i.e. partial correlation magnitude) distribution across population. (Biii) Functional lag distribution with respect to lag length across population. (Biv) Functional weight distribution segregated by lag 0 and lag 1.

<https://doi.org/10.1371/journal.pcbi.1007591.g001>

neuron (i.e. functional group size) was 43.18 ± 3.76 (mean \pm s.e.m.). We employed these experimentally measured functional connections as coupling coefficients in a LM rather than fitted coefficients to mitigate omitted variable bias [32]. Moreover, using the measured functional connections, rather than learned couplings, negated the potential of these terms compensating for imperfect models of stimulus and running. In other words, this approach lessened the likelihood of imparting more importance to the couplings than warranted in our evaluations.

Performance of the unrestricted LM. To assess the accuracy of a simple LM (consisting of coupling coefficients, a stimulus term, and a running term), we evaluated our ability to predict single trial neuronal fluorescence. To model the stimulus, we first determined neuronal response properties by computing the average stimulus-dependent response (averaged across all stimulus presentation time bins and all blocks). Neurons significantly tuned to orientation or direction were labeled as *tuned* with the procedure described in [4]; all others were labeled *untuned*. For tuned and untuned neurons, the stimulus term in our LM was the experimentally observed average stimulus-dependent response and overall average response respectively. The running term for both types of neurons was given by the mouse's running speed as measured by a rotary encoder attached to the axle of the mouse treadmill. Since most functional weights ($70.62\% \pm 2.76\%$) (mean \pm s.e.m where mean is over datasets) had associated lags with values 0 or 1 (in units of roughly 30 ms time bins), we only included these coupling lags in our model. We note that each lag 0 connection was counted twice to indicate bidirectionality. Pooling across neurons, the unrestricted LM exhibited accurate predictions of single trial neuronal fluorescence (fluorescence traces normalized for each neuron) as indicated by a small test set mean squared error (MSE) (median = 0.0402, IQR = 0.0616, [25th percentile = 0.0233, 75th percentile = 0.0849] Fig 2A). The test set MSE was stable over the entirety of stimulus presentation since the first second after stimulus onset exhibited similar test MSE relative to the entire trial (median = -0.80%, IQR = 17.54%).

While most functional weights had associated lags with values 0 or 1, it remained possible that correlations on longer timescales, though small in proportion, impacted modeled neuronal dynamics. We constructed an all lag control model that included all lags up until the maximum lag (17 frames or ~510 ms) as well as stimulus and running terms as done in the unrestricted LM. After fitting the all lag control model on the training set, we computed the percent change in test set MSE of the all lag control model relative to the unrestricted LM (which contained only lags 0 and 1). The median percent change in test set MSE was -0.60% and the IQR was 2.45%, indicating that the inclusion of longer lags did not significantly improve predictions of single trial dynamics confirming our decision to restrict to lags 0 and 1.

While a linear model gives an accurate description of single trial activity, it was possible that a generalized linear model (GLM) with nonlinear link function may have been more suitable. In this case individual neuronal activity is described by a nonlinear function of coupling, stimulus, and running speed. We tested two standard nonlinear GLM link functions suitable for fluorescence data. First, we tested the reciprocal link function in which $\frac{1}{Y} = X\beta$, where Y denotes the vector of responses, X denotes the matrix of model terms (i.e. coupling, stimulus, and running speed), and β denotes the vector of model coefficients. The model was fit on the training data via MATLAB's `glmfit.m` function, and the percent change on test set MSE relative to the unrestricted GLM was computed, giving a median MSE of $1.97 \times 10^5\%$ and an IQR of $5.74 \times 10^5\%$. We also tested the link function in which $Y^3 = X\beta$. Fitting the model and computing percent change test set MSE, we obtained a median MSE of 59.01% and an IQR of 75.58%. These results indicate that these two typical nonlinear functions did not yield improved predictions over the linear model.

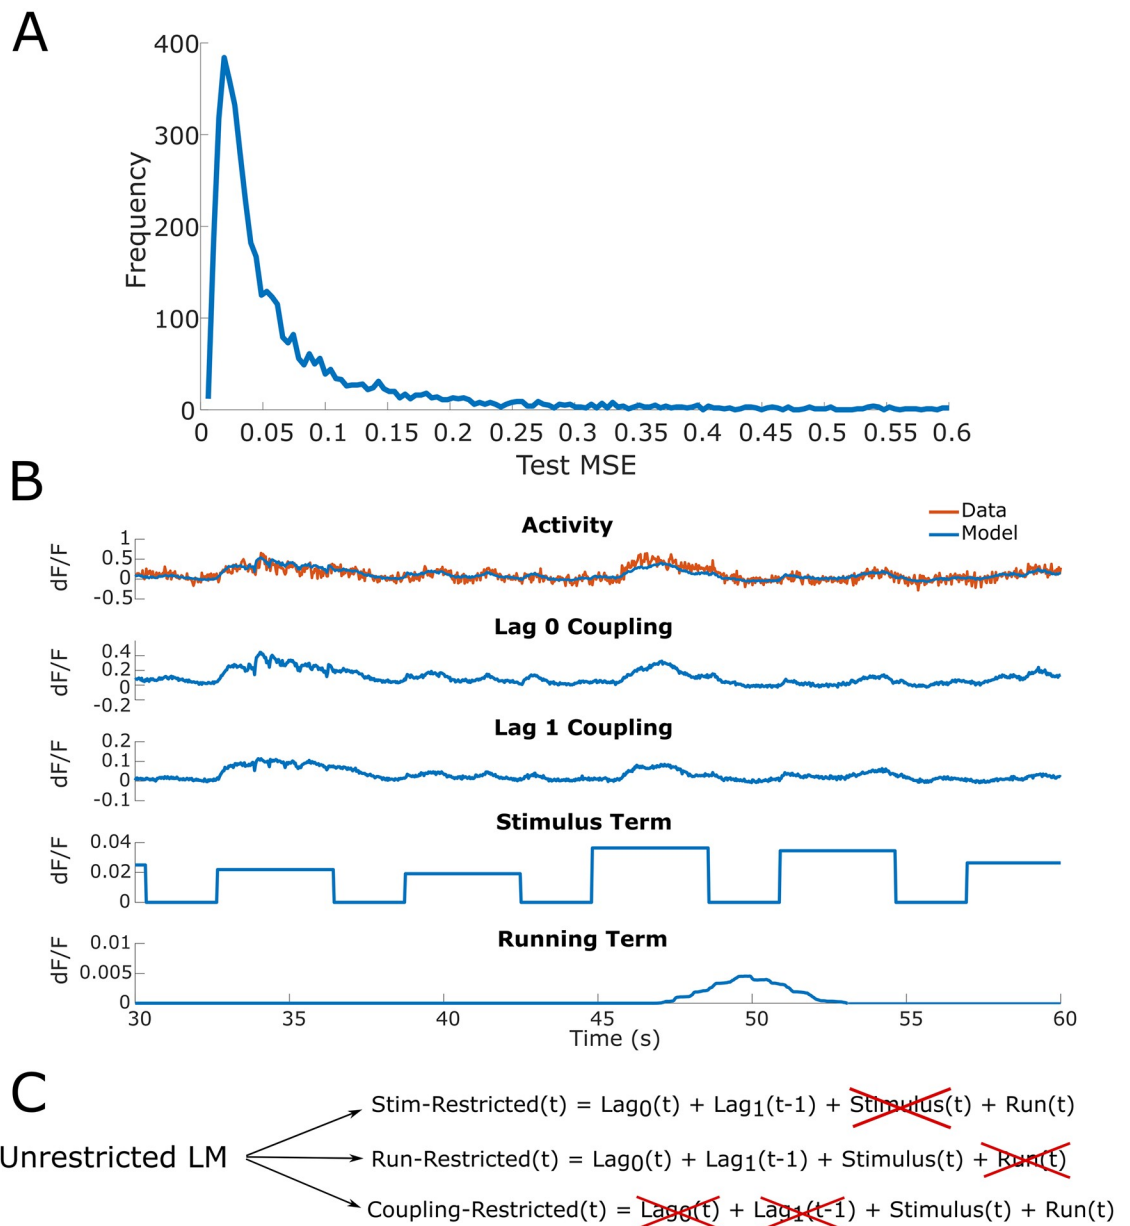


Fig 2. Modelling single trial responses with the unrestricted LM. (A) Distribution of the percent change in test set mean-squared error (MSE) for the unrestricted LM across all neurons. The unrestricted LM contains all three model terms, which are the coupling coefficients, stimulus term (average stimulus-dependent response), and running term (running speed given by rotary encoder). Across stimulus/grey conditions, the unrestricted LM models the single trial responses of many neurons accurately with a right skewed distribution. (B) Top, overlay of modeled (blue) and experimentally measured (red) fluorescence changes in an anecdotal neuron. Bottom, time-varying model terms that led to the prediction. The panels referring to model terms show the time-varying value of that model term; these are then linearly combined to give the model output, which is compared to data in the top panel. (C) Diagram for construction of restricted models. The relative changes in test set MSE of the restricted LMs compared to the unrestricted LM indicate the importance of the excluded term in generating accurate predictions of single trial responses. For example, if the stimulus-restricted LM exhibits a minimal test set MSE increase relative to the test set MSE of the unrestricted LM, then the stimulus is understood to be uninformative to predicting single trial responses.

<https://doi.org/10.1371/journal.pcbi.1007591.g002>

The accurate predictions of the unrestricted LM allowed us to next evaluate which individual terms best contribute to capturing single trial neuronal dynamics. To isolate the relative contribution of each term to the prediction accuracy of single trial fluorescence, we excluded the stimulus or running term (*stim-restricted LM* and *run-restricted LM* respectively) and

compared these models against the prediction accuracy of the *unrestricted LM* (Fig 2C). Each of the three LMs were individually fit on the same training data and compared on the same test data (roughly a 70/30 data split). Additionally, we fit LMs in which all neuronal couplings were excluded (i.e. *coupling-restricted LM*) to determine the extent to which the functional group contributed to accurate prediction of single trial responses.

Performance of the stimulus-restricted LM. We modeled the stimulus in two ways to examine the extent to which averaged summaries of stimulus-dependent responses can contribute to accurate predictions.

First, as in the unrestricted LM above, we chose to model the visual stimuli as trial averaged responses given that the transform(s) of visual information as it transits from retina to LGN to layer 4 to layer 2/3 is not fully known. In particular, we averaged the empirically measured stimulus-dependent fluorescence across all time bins per stimulus condition across all trials (each trial being a single 5-minute block) (Fig 3A), as described above. We then fit a LM containing all terms (i.e. unrestricted LM) and a LM with all terms except the stimulus term (stimulus-restricted LM). When comparing between our two models of the stimulus, we restricted our analysis of the results to only the drifting grating (i.e. non-grey stimulus) imaging frames.

Pooling across neurons, the stimulus-restricted LM exhibited marginal test set MSE increase (all frames: median = +0.20%, IQR = 1.43%, stimulus frames: median = +0.22%, IQR = 1.79%, Fig 3B). The stim-restricted LM's marginal MSE increase suggests that knowledge of the average response does not improve single trial predictions when conditioning on the functional group's activity and the running speed. It must be noted, however, that neurons in the 95th percentile of stimulus-restricted LM test set MSE increase exhibit a large increase (median = +29.70%, IQR = 29.11% across stimulus frames, Fig 3B), indicating that average response captured a large component of the single trial dynamics for a small subset of neurons.

Our second model of the stimulus attempted to account for the temporal dynamics of a stimulus-dependent response. In this model, tuned and untuned neurons were treated the same. For each neuron, we averaged its fluorescence trace in response to a given stimulus across all presentations and across all trials to obtain the "block averaged trace" for a given stimulus (Fig 3A). The block averaged trace was thus obtained after averaging over 9–15 traces depending on the dataset; as noted above, we tested this model only on stimulus frames (i.e. non-grey frames). In order to avoid averaging over different numbers of presentations, we took a 50/50 split of the data for training and testing the models. Notably, by considering fewer trials, we necessarily were less able to completely isolate and account for the other two variables, coupling and running speed using partial correlation. As before, we fit an unrestricted LM and the stimulus-restricted LM, and we restricted analysis to the drifting grating (i.e. non-grey stimulus) frames.

Across the population, this variant of the stimulus-restricted LM again exhibited marginal increase in test set MSE (median = +6.64%, IQR = 5.02% across stimulus frames, Fig 3B). For the majority of neurons, averaged summaries of stimulus-dependent responses are ill-equipped to explain single trial activity even when accounting for the average temporal component of dynamics. Again, neurons in the 95th percentile of stimulus-restricted test MSE increase exhibited large MSE increase (median = +50.59%, IQR = 36.87% across stimulus frames, Fig 3B). The remainder of our analysis in this paper used the average response model for the stimulus term unless otherwise specified.

Performance of the run-restricted LM. We evaluated the contribution of the running term (running speed measured by rotary encoder) to accurate predictions of single neuron fluorescence transients. The run-restricted LM exhibited marginal changes in test set MSE compared to the unrestricted LM (median = -0.01%, IQR = 0.20%, Fig 3C), indicating that

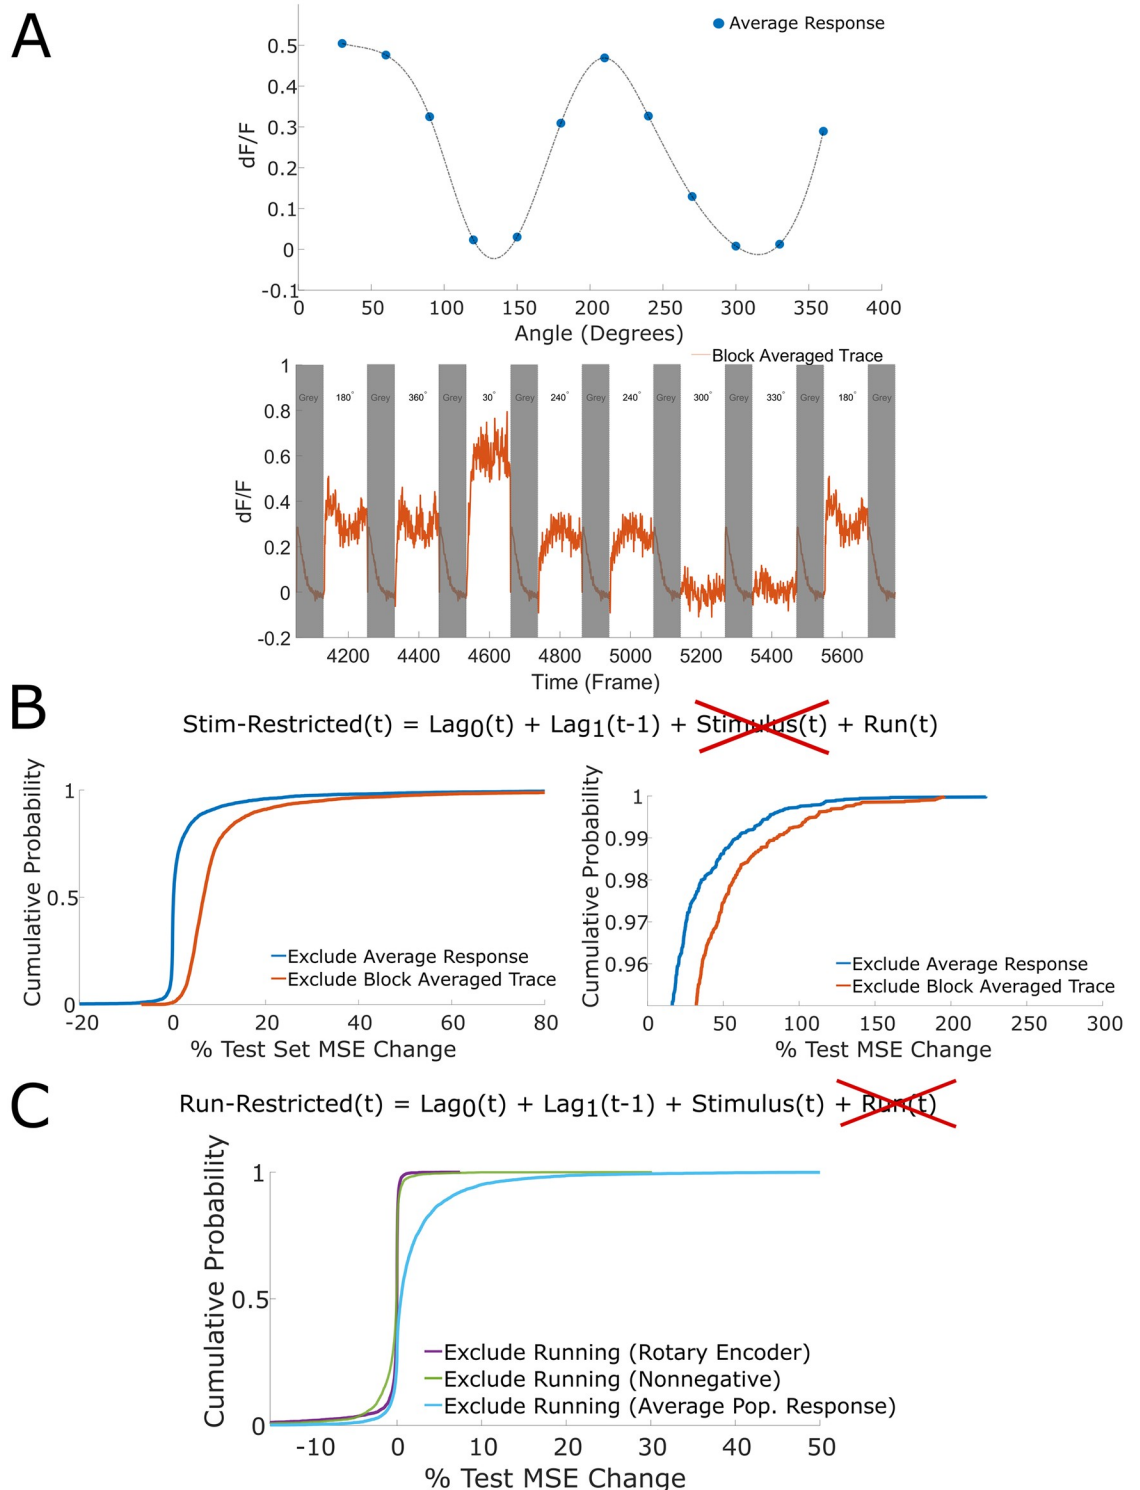


Fig 3. Performance of stimulus-restricted and run-restricted LMs in predicting single trials. (A) Top, stimulus-dependent average response profile for an anecdotal example neuron. The average response is a grand mean, i.e. the average response of a neuron to a given stimulus is the neuron's response to the stimulus averaged across all time bins of all presentations of the given stimulus. In effect, the average response is the average fluorescence change the neuron exhibits in response to a given stimulus. Bottom, an example of a block averaged trace for the same neuron (neuron 2, dataset 3). The block averaged trace of a neuron for a given stimulus is the trace obtained when averaging traces across all blocks (i.e. trials) and all presentations of a given stimulus within

each block. In effect, the block averaged trace represents the average trace of fluorescence changes in response to a given stimulus. For the LMs using the block averaged trace model, we restricted our analysis only to the stimulus frames and excluded the grey frames which are obscured here using grey bars. (B) Left, cumulative distribution functions of the percent change of test set MSE across stimulus frames for the stimulus-restricted LM using the average response stimulus model and the stimulus-restricted LM using the block averaged trace stimulus model. Right, zoom in on 95th percentile. (C) Cumulative distribution functions of the percent change of test set MSE across all frames for the run-restricted LMs using the rotary encoder running model, nonnegative model coefficients, and average population response running model.

<https://doi.org/10.1371/journal.pcbi.1007591.g003>

inclusion of running speed failed to improve prediction accuracy when conditioning on functional group activity and the average response. When restricting to frames where the animal runs, the run-restricted LM still exhibited marginal changes in test set MSE relative to the unrestricted LM (median = -0.024%, IQR = 0.3575%). In numerous cases, we obtained a negative coefficient for the running term as we did not constrain model coefficients to be nonnegative. While we do not focus on the coefficients in the majority of the paper, it was necessary to address these particular values because it is well known that running speed strongly modulates spike rates [16], and our data similarly showed enhanced population activity during periods of running [4]. Given that (1) we are inevitably subsampling visual cortical circuits and thus cannot include all relevant variables in the LM, and (2) we had found that run speed coefficients are mostly small in magnitude (median = -0.0001, IQR = 0.0062), it was possible that sign errors were due to omitted variable bias [32]. Consequently, we refit the unrestricted LM and the run restricted LM with the additional constraint that coefficients of all model terms (i.e. coupling, stimulus, and running terms) be nonnegative. Again the change in test set mean squared error (MSE) relative to the unrestricted LM was small (median = -0.0598%, IQR = 0.6575, Fig 3C). Although we observed global increases in neuronal activity across all imaged neurons during running, the inclusion of a running term did not substantially contribute to accurate prediction of neuronal fluorescence in held out data.

To account for the possibility that running speed does not linearly drive population activity, we constructed a second running model. We used the time-varying average population response as the running term since running is known to globally modulate the population. Fitting and testing the corresponding unrestricted and run-restricted LMs, we found that the average population response running model captured more of the dynamics than the nonnegative rotary encoder model, but was still relatively uninformative as changes in test set MSE were small (median = 0.45%, IQR = 2.32%, Fig 3C). In the remainder of our analysis, we used the rotary encoder running model when not otherwise specified.

Performance of the coupling-restricted LM. The above results suggested that inclusion of the stimulus condition and running speed were largely uninformative to generating accurate predictions of single trial neuronal dynamics in the full LM. However, it was unclear whether these terms enable accurate predictions when *not* conditioning on functional group activity. We fit a LM with all terms (with the average response stimulus model and rotary encoder running model) except coupling on the training set. The coupling-restricted LM exhibited a large test set MSE increase relative to the unrestricted LM (median = 28.93%, IQR = 47.03%, Fig 4A top right; across all frames), demonstrating that the functional group captured variance unexplainable by the average stimulus response and running speed. This result held when the block-averaged trace was used as the stimulus term (with the rotary encoder running model), with the coupling-restricted LM exhibiting a large increase in test set MSE (median = 22.35%, IQR = 39.80%; across stimulus frames). Collectively these results indicate that the couplings between neurons contribute substantially to accurate model predictions.

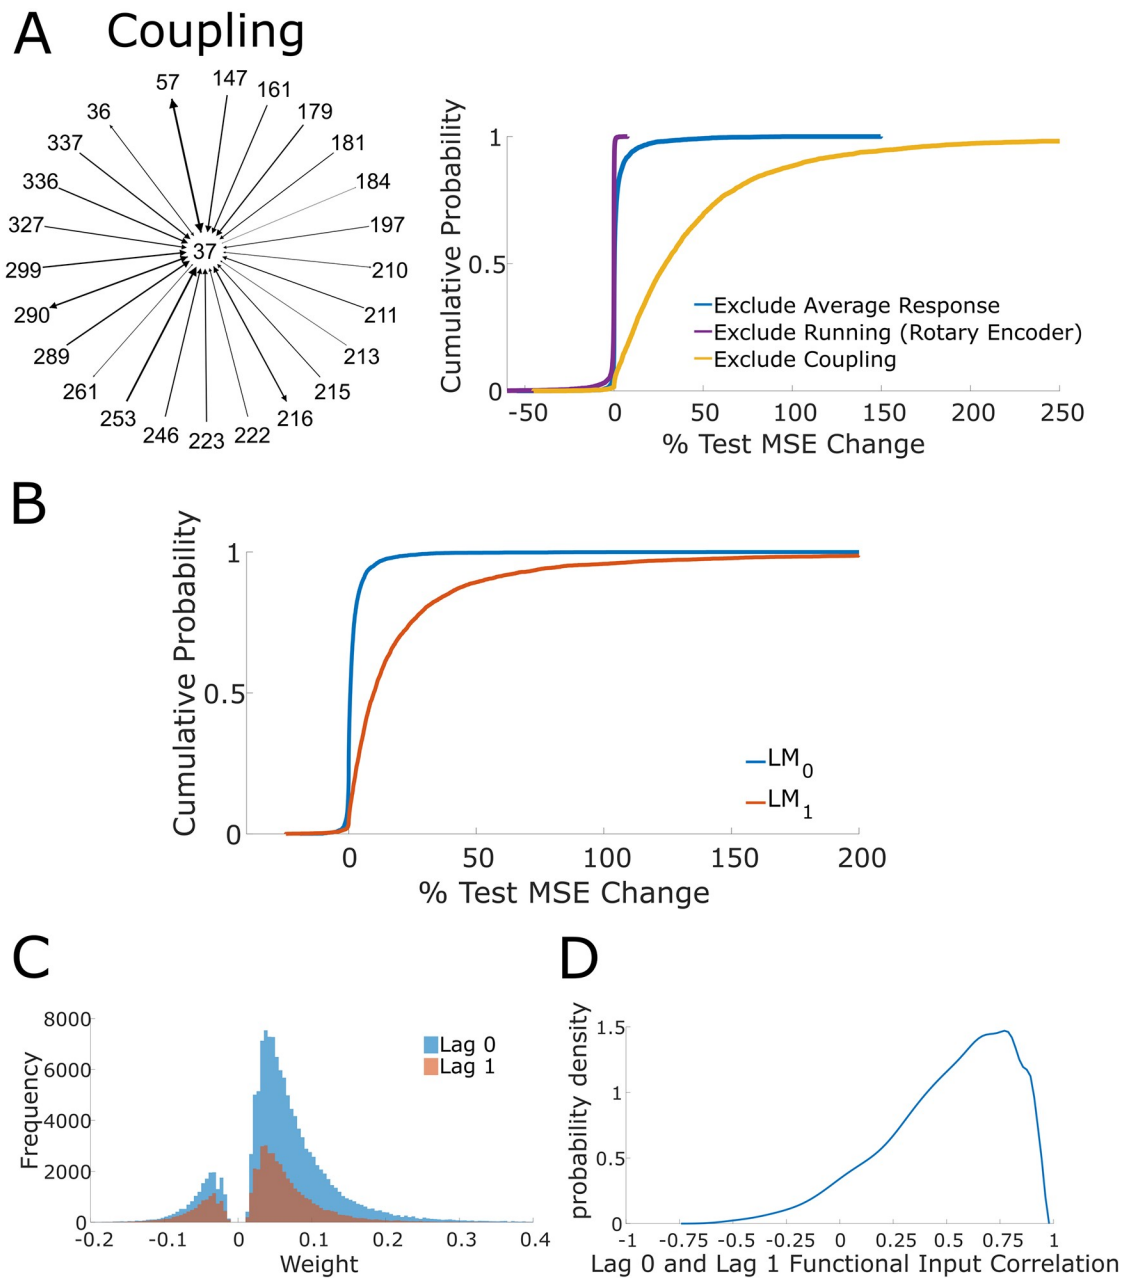


Fig 4. Coupling-restricted model performance and correlation time scales. (A) Left, example of functional group of an anecdotal neuron. Diagram shows all neurons with nonzero weight (across all lags) that have directed coupling outgoing to the anecdotal neuron; these neurons constitute the functional group for the anecdotal neuron. Directionality is indicated by arrows and edge weight is denoted by thickness of the edge and arrow head size, with increasing and decreasing thickness corresponding to increasing and decreasing weight respectively. Top right, cumulative distribution functions for the percent change in test set MSE across all frames for stimulus-restricted, run-restricted, and coupling-restricted LMs using the average response stimulus model and the rotary encoder running model. (B) Cumulative distribution function of percent change in test set MSE of LM₀ and LM₁ with respect to the unrestricted LM. The exclusion of lag 0 results in a large increase in test set MSE, indicating that lag 0 edges are informative to predicting single trial responses. The exclusion of lag 1 results in a marginal increases of test set MSE, indicating that lag 1 edges are mostly uninformative. (C) Strong weight distribution segregated by lag 0 and lag 1 ("strong" weight means functional weight in top quartile of magnitudes). Strong edges are more likely to be lag 0 edges, indicating why lag 0 edges are informative to accurate predictions. (D) Probability density estimate of Pearson correlation between lag 0 and lag 1 coupling across all neurons.

<https://doi.org/10.1371/journal.pcbi.1007591.g004>

LM prediction accuracy is sensitive to the specifics of topology and large weights

While it was clear that the functional group coupling enabled accurate predictions of single trial activity, the importance of specific topological features of the functional group was unclear. Is prediction more sensitive to the precise weight values or the underlying presence of specific connections? To explore these relationships, we employed graph theoretic methods to manipulate functional group weights and connections. In this framework, a connection is referred to as an edge.

Previous work has shown that weak functional weights are uninformative of single trial dynamics [4]. To examine the importance of specific topological features amongst the strongest weighted edges in the functional group, we first identified the edges with weights in the top quartile of magnitudes (termed the “strong edges” and “strong weights”) and permuted the weights (along with the corresponding lag) among these edges. This permutation scheme preserved the topology of the functional groups but shuffled the associated weights (Fig 5A). The new weights of the functional group were then substituted into the unrestricted LM. We did not refit the unrestricted LM, but rather computed the change of MSE on the training data (here, we revert to analyzing all frames). Refitting the LM would have optimized model coefficients to minimize model error, thus obscuring the effect of shuffled weights. This shuffling procedure was repeated 1000 times, and resulted in a large median increase of training MSE (median of median increase = +23.57%, IQR of median increase = 57.99%, Fig 5B top). These results demonstrate that the predictive abilities of functional groups are highly sensitive to the precise value and association of each strong functional weight to a specific edge.

Second, we again identified the edges with strong weights (termed “strong edges”). Now, rather than shuffle only the weights, we shuffled the strongest edges and thereby generated new topologies (Fig 5A). This allowed us to determine the sensitivity of prediction ability to the underlying strong edges. To elaborate, this shuffling procedure shuffled *all* of the strong edges across *all* functional groups. After shuffling, each neuron’s functional group was changed since the composition of a functional group is defined by the specific incoming edges that each neuron receives. Again the LM was not refit for the same reasons articulated above, and this procedure was repeated 1000 times. Shuffled edges resulted in a large median increase of training set MSE (median of median increase = +20.00%, median of IQR = 31.43%, Fig 5B right).

We additionally investigated the effect of permuting the strong weights and edges when refitting the LM coefficients after shuffling. We shuffled the strong weights and edges as described above, refit the LM coefficients on the training data, and then computed the MSE on the test data. Due to the additional computing time associated with refitting after every permutation, we reduced the number of permutations from 1000 to 100. After permuting the strong weights, the percent change in test set MSE relative to the original unrestricted LM exhibited a median change of 1.10% with an IQR of 14.15%, and after permuting the strong edges, the percent change in test set MSE exhibited a median increase of 9.68% and an IQR of 20.26%. These results are reported as a “grand median” (rather than a “median of medians”) after pooling across all neurons and all permutations to stay consistent with our restricted-model analyses which involved refitting LM coefficients. The LMs’ prediction accuracies are broadly robust to perturbations in the functional group’s strong weights, which is unsurprising as the coefficients are refit after shuffling. However, perturbations to the underlying topology of the functional group (i.e. shuffling strong edges) result in increased test set MSE, despite the fact that refitting LM coefficients reduces model error. This result thus provides further validation that specific strong edges are crucial for model prediction.

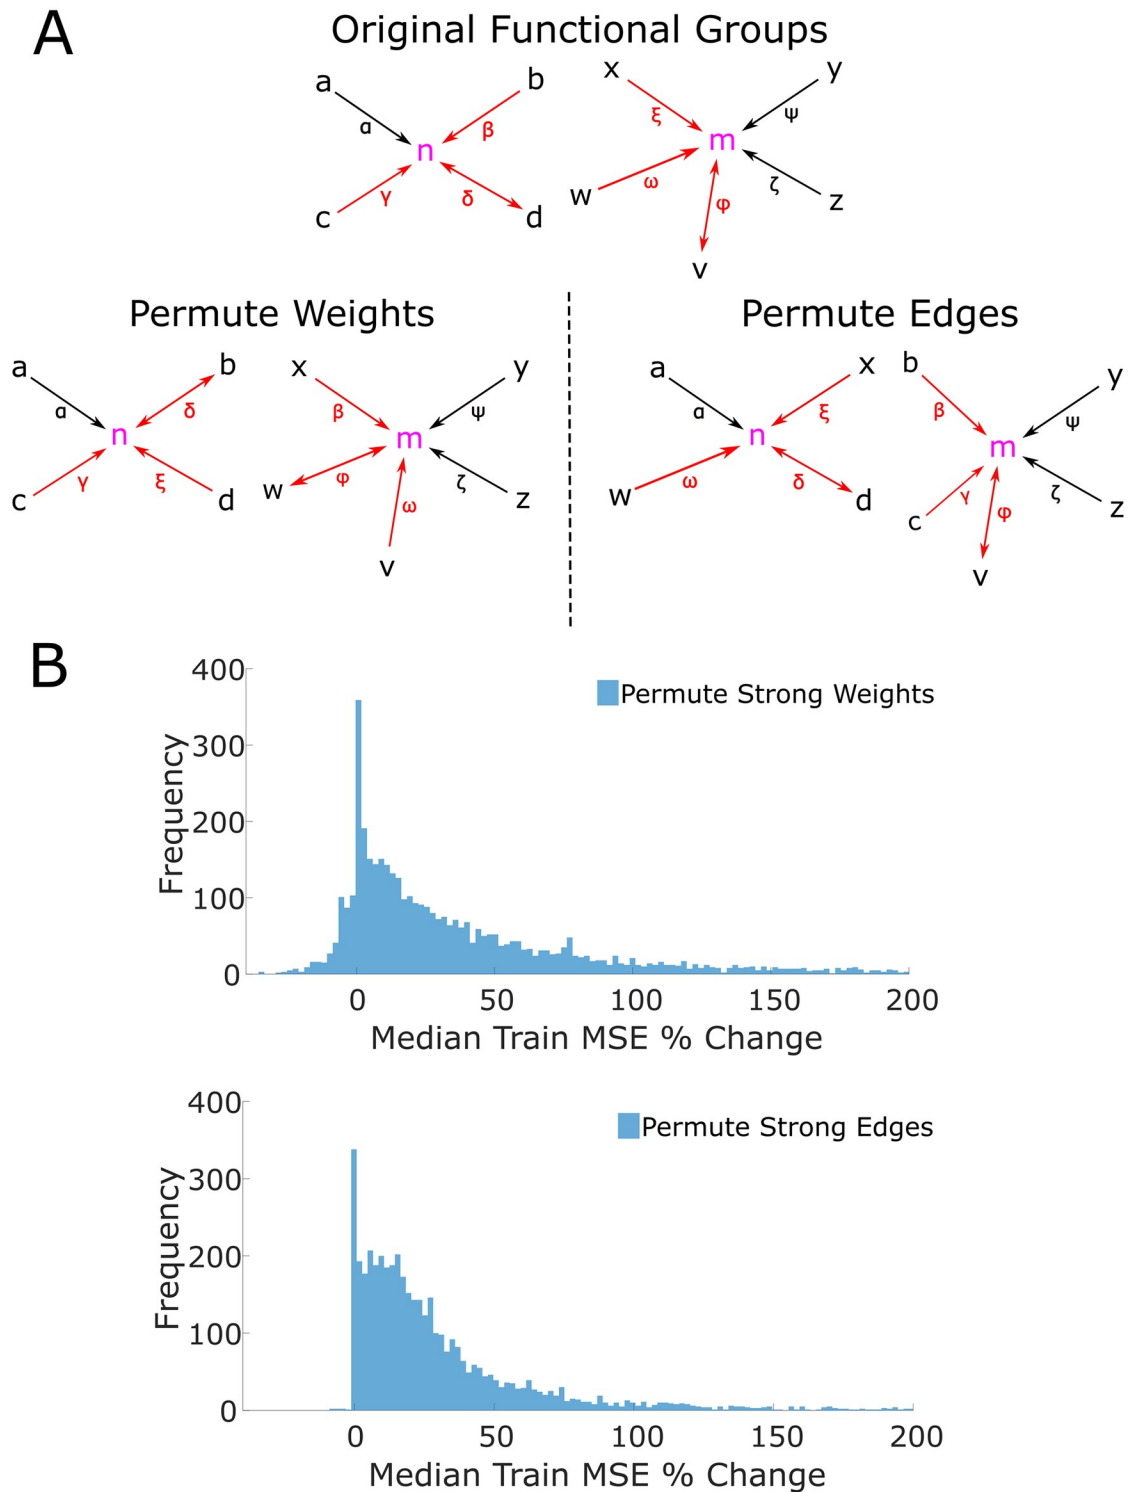


Fig 5. Role of strong weights and strong edges in the functional group. (A) Diagram of permutation methods. Top, example of the functional groups for neurons n and m . The roman letters denote neurons that are coupled to the target neurons. Greek letters denote the weight of the edge (i.e. the functional weight). The arrows show the directionality of the coupling, with bidirectional coupling indicating a lag 0 edge. Red coloring of edges and weights indicate that the edge is “strong” (i.e. has an edge in the top quartile of magnitudes). Only strong edges and weights are permuted in the two shuffle methods. Bottom left, an example permutation of the strong weights of the original functional groups. Note that each functional group retains the neurons in its group, but the strong weights (denoted by the red Greek letters) are freely permuted, even between functional groups. Note that the lag is

permuted along the weight, and so the directionality of the coupling follows the weight when it is permuted. For example, note that the weights β , δ , and ξ are all permuted, and the bidirectional coupling associated with δ follows the permuted weight δ . Bottom right, an example permutation of strong edges of the functional groups. When permuting edges, the neuron memberships of the functional groups change as strong edges are permuted in and out of each functional group. For example, neurons b and c are permuted into the functional group of m from the functional group of n . Similarly, neurons x and w are permuted into the functional group of n from the functional group of n . By permuting edges, new functional group topologies (in the sense of neuron membership) are instantiated. (B) Top, distribution of the median percent change of training set MSE across all 1000 permutations of strong weights. Bottom, distribution of the median percent change of training set MSE across all 1000 permutations of strong edges. These distributions indicate that accurate predictions from the model are sensitive to perturbations of both strong weights and edges of functional groups.

<https://doi.org/10.1371/journal.pcbi.1007591.g005>

Thus, the specific edges and specific weights in the upper quartile of weight magnitudes underlie the impressive performance of the coupled model. Large weights thus explain the importance of specific edges in the functional group for accurate predictions and are a signature of the important timescales and weights (regardless of sign) of the functional group; consequently, the timescale and weights of Hebbian assemblies can be further characterized by observing which timescales and weight sign correspond to the edges with large weights in the functional group. Having established that edges with large weights underlie accurate predictions, we next proceeded to describe various properties (such as timescale and weight and sign) of the functional group which were relevant to accurate prediction by comparing restricted models. More precisely, we did not attempt to establish causal effects since we had already established that large weights enabled accurate predictions. Instead by establishing the relationship between the specific features of strong weight edges and accurate predictions we generated a more comprehensive description of the functional group.

Magnitudes of recurrent coupling coefficients are sufficient for accurate predictions

Given the importance of the functional group, we next asked how its temporal structure related to prediction. To do so we constructed two models, LM_0 and LM_1 , which respectively contained only a lag 0 or a lag 1 coupling term, and compared prediction performance to an unrestricted LM. We additionally examined the large weight distributions for lag 0 and lag 1 coupling.

We found that the LM_0 exhibited a marginally increased test set MSE while LM_1 exhibited a large test set MSE increase as compared to the unrestricted model that contained both coupling terms (LM_0 median = +0.63%, IQR = 2.16%, LM_1 median = +9.84%, IQR = 21.10%, Fig 4B). This result demonstrates that lag 0 coupling is necessary to generate accurate single trial predictions and suggests that prediction-relevant coactivity between neuronal pairs occurs within a window less than or equal to ~ 30 ms. Hence, the predictive ability of functional group coactivity is driven by the functional group's bidirectional edges.

The sufficiency of lag 0 weights and insufficiency of lag 1 weights to predict single trials is informed by the prevalence of strong weights associated with lag 0 edges ("strong" defined as belonging to the top quartile) (Fig 4C). Furthermore, we observe that strong coupling exhibits more lag 0 weights than lag 1, which is consistent with our previous results regarding large weights. While it is the case that lag 0 and lag 1 functional input are linearly correlated (median $r = 0.56$, IQR = 0.40, Fig 4D), collinearity cannot totally explain the sufficiency of lag 0. If activity from lag 0 and lag 1 coupling are collinear, then LM_1 should have performed comparably to LM_0 , as the lag 1 coupling LM coefficient would be modified linearly to account for the absence of lag 0 coupling; yet, LM_1 performed worse than LM_0 .

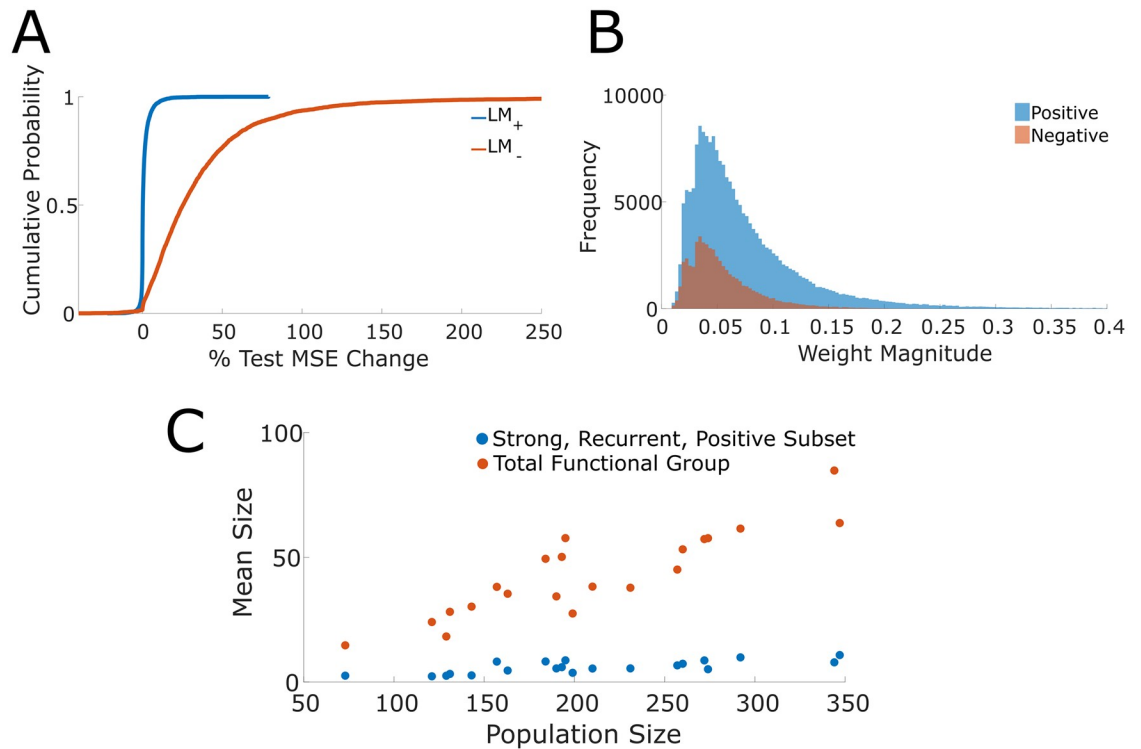


Fig 6. Correlation structure and size of functional groups. (A) Cumulative distribution function of percent change in test set MSE of LM⁺ and LM⁻ with respect to the unrestricted LM. The exclusion of edges with positive weights results in a large increase of test set MSE while the exclusion of edges with negative weights results in only minimal increases. These results suggest that edges with positive weights are informative while edges with negative weights are uninformative in generating accurate predictions of single trial responses. (B) Distribution of strong weight magnitudes segregated by positive and negative weights. Strong weights are more likely to be positive rather than negative, thus indicating why positive weighted edges are informative to accurate predictions. (C) Mean sizes of the total functional group and the strong, positive, recurrent subsets against the population size (i.e. the number of imaged neurons).

<https://doi.org/10.1371/journal.pcbi.1007591.g006>

Positively correlated coactivity sufficiently predicts single trial dynamics

We next asked to what extent prediction was attributable to positive or negative coupling coefficients. Negative correlations have been shown theoretically to increase coding capacity [34] and consequently were of particular interest here. Two models, LM⁺ and LM⁻, were constructed where only the positive weight term or negative weight term was included respectively; additionally, an unrestricted LM was fit. Stimulus and running speed terms were included in all models; however, we did not restrict the set of lags under consideration. LM⁺ exhibited a marginal increase in test set MSE (median = +0.24%, IQR = 1.43%, Fig 6A) whereas LM⁻ exhibited a large increase in test set MSE (median = +25.38%, IQR = 36.37%, Fig 6A). These results indicate that positive weights are sufficient for accurate prediction while negatively correlated neurons contributed minimally. This result further implies that the positive and negative coupling model terms are not collinear as LM⁺ and LM⁻ did not exhibit comparable performance. The performance difference of LM⁺ and LM⁻ is notably marked by the prevalence of strong positive as compared to negative weights. Across datasets, the top quartiles of weight magnitudes are comprised of many positive weights and few negative weights (Fig 6B).

The size of the informative functional group saturates

Previously we had found that the accuracy of an encoding model increased with the number of neurons imaged in a dataset [4]. We set out to establish whether this positive relationship was due to more of each neuron's respective functional group being sampled, or whether we had imaged a greater number of complete functional groups. As expected the total number of incoming edges increased with the number of neurons imaged per dataset (Fig 6C; slope = 0.21). However, this was not the case when isolating the positive bidirectional edges in the upper quartile of weights. Rather the number of edges in this subset remained relatively stable regardless of the number of neurons imaged (slope was 0.028; Fig 6C). These results suggest that the increase in model accuracy with increasing numbers of neurons sampled is the consequence of capturing a greater number of complete functional groups. Correspondingly, more neurons are accurately modeled instead of obtaining a better delineation of neurons' functional groups. This result suggests a numerical size for a functional group.

Functional weights enhance single trial decoding

While the functional group enables accurate single trial encoding predictions, it remained unclear if functional groups are computationally relevant to decoding. More concretely, we asked if using the unrestricted LM, in which the structure of the functional group is known, to decode the stimulus would result in better performance than decoding when neurons are assumed to be uncoupled since the decoding performance is suggestive of the computational relevance of the functional group [35].

We constructed coupled and uncoupled decoders under the Bayesian decoding framework. In particular, a uniform prior over all stimulus conditions was adopted and the stimulus was decoded via the *maximum a posteriori* estimate. All frames (i.e. both training and test frames) were decoded, and both decoders performed better than chance (Fig 7A), where chance performance is given by $\frac{1}{13} \approx 7.69\%$ (12 stimulus conditions and 1 grey condition). The coupled LM decoder decoded more accurately than the uncoupled decoder (coupled LM mean accuracy = 64.77%, sem = 2.30%, uncoupled mean accuracy = 40.90%, sem = 2.74%, accuracy pooled over all frames, mean over datasets), indicating that knowledge of coupling aids in the extraction of sensory information from the population response. On both confusion matrices (Fig 7A), the "bright bands" indicated that the decoder in some cases incorrectly decoded stimuli that were 180 degrees apart from the true stimulus (the drifting gratings are graded in increments of 30 degrees). Of course, the imaged population of neurons includes both orientation-tuned (tuned to two directions of drifting gratings that are 180 degrees apart) and direction-tuned neurons (tuned to primarily one direction of drifting gratings). The "bright band" phenomenon may be explained by the increased activity of orientation-tuned neurons pushing the decoder to select the incorrect stimulus. Nonetheless, the coupled LM decoder ultimately decodes more accurately than the uncoupled decoder.

We further computed the mutual information between the true stimulus and the predicted stimulus [36], determining the coupled LM decoder extracts roughly 64% more sensory information (mean = +63.94%, sem = 14.69%, mean over datasets, Fig 7B). These results indicate, supposing downstream elements have the theoretical capability to read out the visual stimulus, that the functional group dramatically enhances the capability of decoding of the stimulus.

Discussion

Knowledge of the coactivity of a V1 neuron's functional group enables accurate predictions of short term dynamics within single trials. Global descriptions such as averaged stimulus-

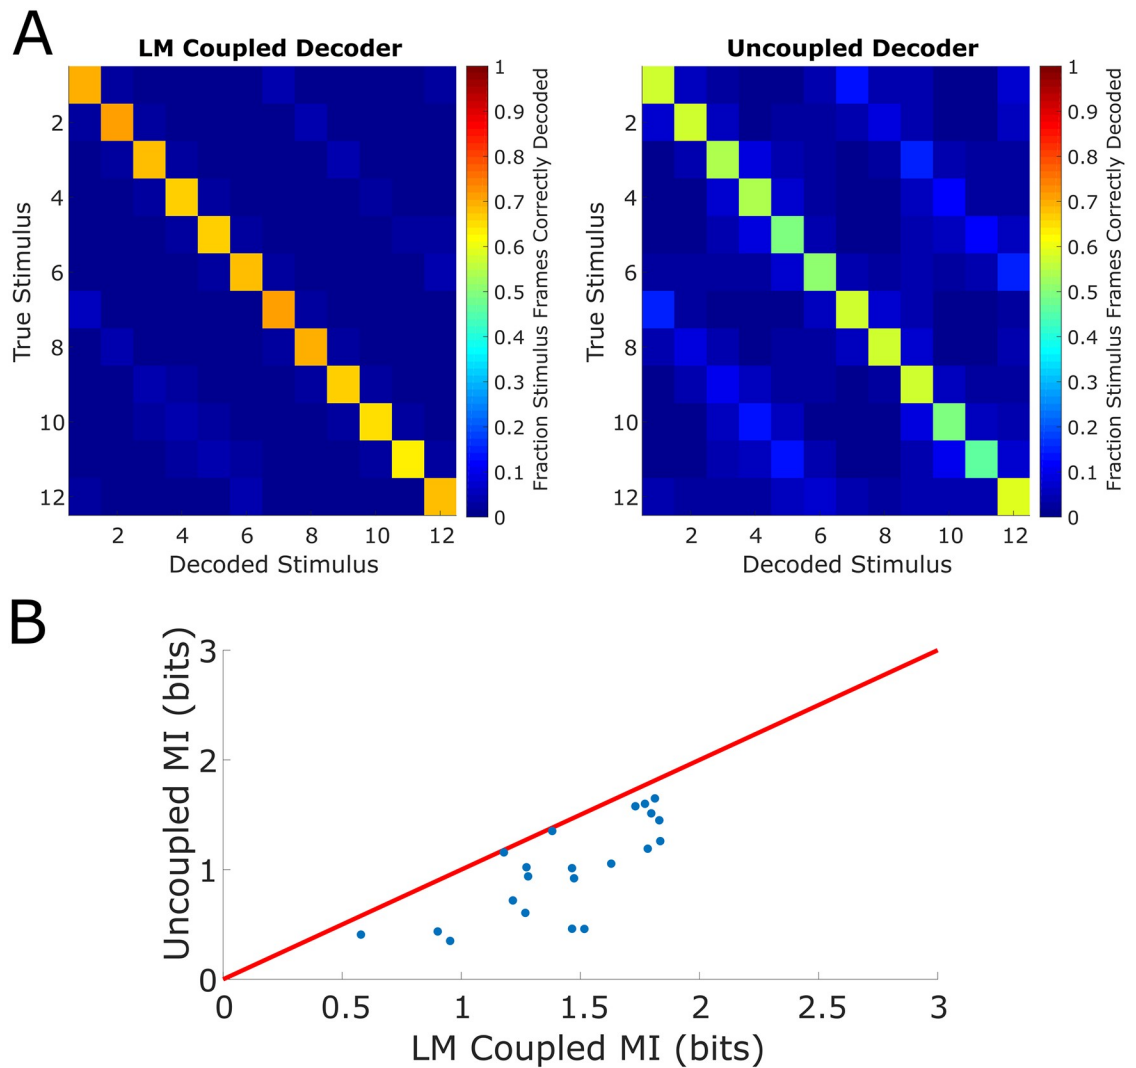


Fig 7. Decoding the visual stimulus with functional groups. (A) Left, confusion matrix (percent of corresponding stimulus frames correctly decoded) of LM Coupled Decoder. Right, confusion matrix of Uncoupled Decoder. The LM Coupled Decoder decodes at a higher accuracy than the Uncoupled Decoder across all stimulus conditions. (B) Comparison of mutual information (MI) between LM Coupled Decoder and Uncoupled Decoder (red line is unity).

<https://doi.org/10.1371/journal.pcbi.1007591.g007>

dependent responses and running fail to meaningfully capture single trial activity in the vast majority of L2/3 neurons in visual cortex. With the functional group operationalizing the notion of a Hebbian assembly, our results give concrete descriptions of the timescales, numerical scale, correlation structure, and computational capabilities of assemblies.

Modelling fluorescence with LMs

The assembly hypothesis posits that internal state dynamics of the underlying circuit shape single trial dynamics; consequently, assembly coactivity is proposed to capture single trials to a greater extent than global variables. To test this proposition, we used various LMs as encoding models [35,37–39]. We used partial correlations to capture shared variability and act as coupling coefficients in the LM rather than relying on fitted coefficients [35,38,39]. Using empirically measured correlations for coupling coefficients is an attractive modeling choice; fitted

coupling coefficients necessarily depend on the model and variable specification, and are thus at risk of omitted variable bias [32]. Our approach also minimized the potential for learned couplings correcting for any limitations of our stimulus or running models. Since we use the functional group as an operational notion of an assembly, coupling must not change with specifications of the linear model.

Functional group structure and prediction

We found that the unrestricted LM accurately modeled the dynamics of single trial activity, thus justifying the decision to adopt a linear model, use partial correlations, and restrict to lags 0 and 1. The target neuron's variability captured by the model is a deterministic function of the variabilities of the functional group, averaged stimulus-dependent response, and running speed and did not require an explicit stochastic term. Our model results showed functional group coactivity was the main predictor of single trial activity of V1 neurons, in support of the assembly perspective [31] as well as previous work showing that pairwise correlations can explain activity patterns [23,40].

Investigating the robustness of prediction to modulations of the functional group, we found that modulating the topology of large weights severely degrades the model's ability to predict single trial responses. Hence, the precise instantiation of the strong weights and edges in the functional group is essential for prediction. While global population coupling and fluctuations have been shown to be informative to predicting activity [9,11,41–43], somewhat consistent with our block averaged stimulus and population average response locomotion models, our results point to the inclusion of a local description when predicting neuronal dynamics within single trial responses. Our results largely suggest that strong functional group coupling (a local description) disproportionately captures dynamics, in congruence with the assembly hypothesis and other previous work [27].

With the functional group established as the primary predictor of single trials, we investigated its features. Functional group edges with functional lag 0 are sufficient for accurate predictions, suggesting that assembly coactivity largely occurs on timescales of $\sim 30ms$, as found in other regions [31]. Furthermore, while it is known that correlations between neurons are predominantly positive and stronger between similarly tuned neurons [7,21,44], we additionally found that precisely the large, positive weights most enabled accurate predictions. Moreover, we determined that imaging a larger population captured more functional groups rather than capturing finer details of existing ones. Hence, we obtained an estimate of the numerical size of a functional group, giving an indication of the dimensionality of informative local interactions in a 2D plane.

Using a Bayesian decoding framework, we found that knowledge of the functional group improves decoding over treating neurons as independent units, corroborating previous work [45,46] and suggesting that shared variability is computationally relevant. While the Bayesian decoder does not reveal whether downstream circuits actually have access to or use the functional group structure, these results nonetheless demonstrate the theoretical relevance of functional groups [36].

The use of the functional group as an operational definition of an assembly enables concrete descriptions of the timescale, correlation structure, and computational capability of a coactive population. The phenomenon that precise details of the functional group are critical to accurate predictions of dynamics over short time scales indicates that sensory representation and computation are comprised of local coupling in addition to global population-wide variables.

Methods

Ethics statement

The data underlying the results presented in this study were previously published in Dechery, J. B., & MacLean, J. N. (2018). “Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons.” *PLOS Computational Biology*, 14(5), 1–23. <https://doi.org/10.1371/journal.pcbi.1006153>. As indicated in that work, the data collection procedures were in accordance and approved by the Institutional Animal Care and Use Committee at the University of Chicago.

Data

A subset of imaging data ($n = 21$ datasets) of mouse visual cortical neurons was taken from [4]. As described in [4], data was collected from $n = 4$ male and 4 female C57BL/6J mice expressing transgene Tg(Thy1-GCaMP6s)GP4.12Dkim (Jackson Laboratory). Activity of L2/3 excitatory neurons in mouse visual cortex (73–347 neurons; 25–33 Hz; 21 distinct fields of view) were imaged using high speed two photon laser scanning microscopy [4,33]. Mice were awake, head-fixed, and allowed to freely run on a linear treadmill while viewing drifting square-wave gratings presented in 12 directions in pseudo-random order interleaved with mean-luminance matched grey screen.

Functional weights

Partial correlations were used as coupling coefficients in our LMs; we called these *functional weights*. The functional weight between a pair of neurons is the partial correlation between the corresponding fluorescence traces accounting for the stimulus and population-wide co-activity driven by running. More precisely, for each 5-minute block of oriented drifting grating stimulus presentations (one trial), a trial-specific partial correlation between neuron x and neuron y was computed while accounting for the mean response of neuron x in all other trials, the mean response of neuron y in all other trials, and the mean response of the population excluding neurons x and y in the current trial. Then the trial-specific partial correlations were averaged across trials to obtain the final partial correlation between neuron x and neuron y ; this was taken to be the functional weight between neurons x and y . The mean activity of neurons x and y across all other trials was accounted for in order to control for stimulus-dependent effects resulting in a measure analogous to noise correlations [4]. The mean population activity excluding neurons x and y in the current trial is accounted for in order to control for population co-activity driven by running [4]. We used the MATLAB 2016a function `parcor.m` to compute partial correlations. In mathematical notation, the partial correlation (within a particular trial of stimulus presentations) between neurons x and y after accounting for the stimulus and coactivity driven by running is given by

$$\rho_{xy} = - \frac{(\Omega^{-1})_{xy}}{\sqrt{(\Omega^{-1})_{xx}(\Omega^{-1})_{yy}}} \quad (1)$$

where Ω is the correlation matrix between the trial-specific fluorescence trace of neuron x , the trial-specific fluorescence trace of neuron y , the mean trace of neuron x averaged across all other trials, the mean trace of neuron y averaged across other trials, and the mean trace averaged across neurons x and y within the current trial. The indices x and y correspond to the indices of the trial-specific traces of neurons x and y in Ω . As noted above, the final partial correlation between neurons x and y is obtained by averaging the trial-specific ρ_{xy} across trials.

In order to capture the temporal structure of neuron coactivity, functional weights were given directionality in the following manner. The mean fluorescence trace (averaged across trials) for neuron x and neuron y was computed, and the cross-correlogram was computed. Due to the nature of two-photon imaging, neuronal fluorescence time series are given by a series of time bins with roughly 30 ms width; consequently, lags are given in units of time bins. The lag corresponding to the maximum of the cross-correlogram determined the lag of the functional weight and the sign of the lag determined the direction of the functional weight [4], and this lag was termed the *functional lag*. More concretely, if neurons x and y had a functional weight with a positive functional lag, then the corresponding functional weight has direction corresponding to initial point x and terminal point y . If the lag is negative, then the functional weight has direction corresponding to initial point y and terminal point x . Lag 0 corresponds to a bidirectional functional weight. The *functional edge* refers to this directed functional connection. Cross-correlograms were calculated using the MATLAB 2016a functions `xcorr.m`. For each neuron, the set of neurons with outgoing functional edges to the given neuron (i.e. neurons with functional edges with which the given neuron was the terminal neuron) was called the *functional group* of the given neuron.

Linear model

Predicting fluorescence. A linear model (LM) was used as an encoding model to model calcium imaging data from V1 neurons [35,38,39,47]. The linear model framework allows us to determine which variables are important in predicting the target neuron's activity by iteratively excluding variables of interest and comparing prediction accuracy.

Stimulus effects were modeled in two ways in order to determine the extent to which averaged summaries of stimulus-dependent responses contribute to accurate predictions of single trial dynamics. First, we modeled the stimulus by the stimulus-dependent average response (which we refer to as “average response”), in which response to a given stimulus was averaged across all time bins of the given stimulus presentation across all trials. Neurons significantly tuned to orientation or direction were labeled as *tuned* with the procedure described in [4], and all remaining neurons were labeled as *untuned*. The stimulus term in our LM was the average stimulus-dependent response for tuned neurons. The stimulus term for untuned neurons was given by the response given by averaging the stimulus-dependent average response across all stimulus conditions. During grey frames, the stimulus term was set to zero in this model. With this model, we used a 70/30 split of the data for training and testing in order to fit the LM coefficients and test performance. We modeled the stimulus a second way by the “block averaged trace” in order to reflect dynamics associated to neuronal response. In this model, tuned and untuned neurons were treated exactly the same. The fluorescence traces in each presentation of a given stimulus were averaged across blocks (i.e. trials) but not time bins, producing a block-averaged trace which preserved the dynamics. Since the block-averaged trace preserves dynamics, we only take averages over the training set when fitting the LM to avoid overfitting (when testing, averages are taken over the testing set). In order to ensure that averages occurred over similar numbers of stimulus presentations between the training and testing sets, we used a 50/50 split of the data. In our analyses, the block averaged trace was obtained after averaging over 9–15 traces depending on the dataset. The stimulus term is represented by $s(t)$ in our LM equation. Running effects were modeled without neuron-specific modulation, i.e. we included a term $v(t)$ (denoting the running velocity at time t) that was constant across all neurons.

The functional group was incorporated into the LM by computing a linear combination (with weights given by the functional weights) of the lagged (lag given by the functional lag)

fluorescence traces of all neurons in the functional group (i.e. neurons with outgoing functional edges to the given neuron). For compact notation, we let $\mathbf{r}(t)$ be the vector of population fluorescence at a given time, T be the total duration of the event, and N be the total number of neurons in the population. Then we defined $W_i^k \in \mathbb{R}^N$ to be the vector of functional weights corresponding to incoming edges to neuron i with lag k . In the entire investigation, we restricted attention to $k \in \{0, 1\}$ as the frequency of lagged functional relationships with $k > 1$ decreased rapidly.

With the coupling, stimulus term, and running term defined, for each neuron i , we constructed a LM to model its fluorescence,

$$\hat{r}_i(t) = c + \beta_0 \cdot W_i^{0^T} \mathbf{r}(t) + \beta_1 \cdot W_i^{1^T} \mathbf{r}(t - 1) + \Psi \cdot s(t) + \Gamma \cdot v(t) + \varepsilon_t \tag{2}$$

Here, $c, \beta_0, \beta_1, \Psi, \Gamma \in \mathbb{R}$ are free parameters that are fitted. The ε_t term is a mean zero noise term. Since we are not concerned with inferring the underlying coefficients, the true distribution of the noise is irrelevant and so a distribution does not have to be specified in our model. The linear offset c is needed as different neurons have different baseline fluorescence. The gain terms β_0, β_1 are needed to account for the range of incoming edges to each neuron. The Ψ and Γ coefficients are the modulation associated with stimulus and running respectively. The notation W_i^0 and W_i^1 refers to the vector of lag 0 and lag 1 incoming functional weights to neuron i . More specifically, $W_i^0, W_i^1 \in \mathbb{R}^N$, and

$$(W_i^\ell)_k = \begin{cases} \rho_{ki} & \text{if there exists a functional edge } k \rightarrow i \text{ with lag } \ell \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Here, ρ_{ki} is the functional weight associated to the directed edge $k \rightarrow i$. Since the vectors of functional weights W_i^0 and W_i^1 are computed directly from the data, the scalars $W_i^0 \mathbf{r}(t)$ and $W_i^1 \mathbf{r}(t)$ are features of the data and *not* free parameters. Hence, for a given neuron, there exist 5 parameters in the unrestricted LM. For each neuron, the model was fit on the training data by the method of least squares using a MATLAB R2016a function (`ls_cov.m`).

LM variants

Stimulus-restricted, run-restricted, and coupling-restricted LMs. To determine the extent to which the stimulus term, running term, and coupling terms contributed to predictions of single trials, we iteratively excluded terms of interest, refit the model, and compared predictive performance against the unrestricted model (i.e. containing all model terms described above). In particular, for the *stimulus-restricted LM*, we had the model

$$\hat{r}_i(t) = c + \beta_0 \cdot W_i^{0^T} \mathbf{r}(t) + \beta_1 \cdot W_i^{1^T} \mathbf{r}(t - 1) + \Gamma \cdot v(t) + \varepsilon_t \tag{4}$$

in which the stimulus term was excluded. The model coefficients all have the same interpretation as described in the previous section. We then fit this model on the training data, and compared its predictive performance on the test data with the unrestricted model. Similarly, for the *run-restricted LM*, we had the model

$$\hat{r}_i(t) = c + \beta_0 \cdot W_i^{0^T} \mathbf{r}(t) + \beta_1 \cdot W_i^{1^T} \mathbf{r}(t - 1) + \Psi \cdot s(t) + \varepsilon_t \tag{5}$$

and for the *coupling-restricted LM*, we had the model

$$\hat{r}_i(t) = c + \Psi \cdot s(t) + \Gamma \cdot v(t) + \varepsilon_t \tag{6}$$

These models were fit on the training data, and were compared against the unrestricted model in terms of predictive performance on the test data.

LM₀ and LM₁. To determine the contributions of lag 0 and lag 1 coupling to predictions of single trials, we similarly excluded these terms iteratively and compared predictive performance to the unrestricted model. In particular, LM_0 is given by

$$\hat{r}_i(t) = c + \beta_0 \cdot W_i^{0^T} \mathbf{r}(t) + \Psi \cdot s(t) + \Gamma \cdot v(t) + \varepsilon_t \tag{7}$$

This model was then fit on the training data and compared against the unrestricted model. Note that W_i^0 is not affected by refitting as it determined by the functional weights and functional lags, which are empirically measured. The process of fitting only affects the model coefficients, i.e. c, β_0, Ψ, Γ . Similarly, LM_1 is given by

$$\hat{r}_i(t) = c + \beta_1 \cdot W_i^{1^T} \mathbf{r}(t - 1) + \Psi \cdot s(t) + \Gamma \cdot v(t) + \varepsilon_t \tag{8}$$

and the model is fitted and evaluated with exactly the same procedure as LM_0 .

LM⁺ and LM⁻. To determine the contributions of positive and negative functional weights to predictions of single trial dynamics, we constructed a LM with explicit corresponding terms (again, with average response as the stimulus term). In particular, we had the following model

$$\hat{r}_i(t) = c + \beta_+ \cdot (W_i^+)^T \mathbf{r}(t) + \beta_- \cdot (W_i^-)^T \mathbf{r}(t) + \Psi \cdot s(t) + \Gamma \cdot v(t) + \varepsilon_t \tag{9}$$

The terms $\mathbf{r}(t), s(t), v(t)$ and ε_t as well as their corresponding coefficients are defined exactly as in the original model formulation. The notation W_i^+ and W_i^- refers respectively to the positive and negative functional weights in the functional group of neuron i . More specifically, $W_i^+, W_i^- \in \mathbb{R}^N$, and for neurons k and i and sign s (i.e. positive or negative), we had

$$(W_i^s)_k = \begin{cases} \rho_{ki} & \text{if there exists a functional edge } k \rightarrow i \text{ and } \rho_{ki} \text{ has sign } s \\ 0 & \text{otherwise} \end{cases} \tag{10}$$

In order to minimize over-fitting and keep model complexity comparable across different models, we used all lags in this model and treated them as lag 0.

Permutations of functional group topology

To determine if the topology of the functional group is essential to predictions of single trial activity, we shuffled the functional group in two ways; we shuffled the “strong weights” and the “strong edges”.

We shuffled the strong weights to determine the importance of strong weights for prediction. We shuffled the top quartile of weights while preserving the corresponding underlying edges. Formally, we enumerated the edges with strong weights in our entire population $\{(i_1, j_1, w_1), \dots, (i_K, j_K, w_K)\}$, where a directed edge from neuron i_α to neuron j_α exists and w_α is the associated weight for $1 \leq \alpha \leq K$. This enumeration contains all tuples $(i_\alpha, j_\alpha, w_\alpha)$ such that $|w_\alpha|$ is in the top quartile of all weight magnitudes. We then applied a uniformly random permutation $\sigma \in S_K$ (where S_K is the finite symmetric group on K points) to act on the weights, furnishing the enumeration $\{(i_1, j_1, w_{\sigma(1)}), \dots, (i_K, j_K, w_{\sigma(K)})\}$. We then assigned these permuted weights to the corresponding edges (i.e. the directed edge from neuron i_α to neuron j_α was assigned the weight $w_{\sigma(\alpha)}$) in the functional group. Furthermore, we shuffled the functional lags in this way (thus changing the directionality of the underlying edges). This construction preserves the existence of edges, preserves which edges are in the top quartile of weight magnitudes, and preserves structure between functional lags and functional weights. Hence, changes in prediction accuracy are directly related to the importance of the functional weight to prediction.

Similarly, to isolate the importance of the functional edges in prediction, we shuffled the strong edges of the functional group. We apply a uniformly random permutation $\sigma \in S_K$ to the list of terminal neurons. Formally, with the list $\{(i_1, j_1, w_1), \dots, (i_K, j_K, w_K)\}$ described above, we apply the permutation to obtain $\{(i_1, j_{\sigma(1)}, w_1), \dots, (i_K, j_{\sigma(K)}, w_K)\}$. This procedure permutes the strong edges while preserving the weight with respect to the source neuron. Hence, changes in prediction accuracy are directly related to the importance of functional edges to prediction.

Decoding

Constructing the decoder. To determine whether the functional group is computationally relevant, we used the functional group and LM framework to decode the stimulus. We constructed a Bayesian decoder which estimates the stimulus at time t given the population response and the running speed at time t . The unrestricted LM (using average response for the stimulus term) was incorporated into the decoder [37] in order to directly link the functional structure of the population to decoding; this was termed the coupled decoder. The coupled decoder was then compared to an uncoupled decoder: a Bayesian decoder which treats all units as conditionally independent [35,39].

The coupled decoder was constructed as follows. From Bayes' Theorem, the posterior probability of stimulus conditional on response is proportional to the product of the prior and the probability of response conditional on stimulus. Formally,

$$p(s(t)|\mathbf{r}(t)) \propto p(\mathbf{r}(t)|s(t)) \cdot p(s(t)) \tag{11}$$

Here, $\mathbf{r}(t)$ is the vector of population activity and $s(t)$ is the stimulus at time t . Estimating $p(s(t)|\mathbf{r}(t))$ requires specification of $p(s(t))$ and $p(\mathbf{r}(t)|s(t))$, which we gave as

$$\begin{aligned} s(t) &\sim \text{Uniform}(\{\text{Grey}, 30, 60, \dots, 360\}) \\ \mathbf{r}(t)|s(t) &\sim \mathcal{N}(\text{LM}(t, s(t)), \Sigma_{s(t)}) \end{aligned} \tag{12}$$

Here, $\text{LM}(t, s(t))$ is the vector of LM predicted responses at time t conditioned on the stimulus input $s(t)$. Further, $\Sigma_{s(t)}$ is the covariance of the neurons' traces during frames with $s(t)$ presentation. The decoded stimulus is then the *maximum a posteriori* estimate (MAP)

$$\hat{s}(t) = \text{argmax}_s p(\mathbf{r}(t)|s) \tag{13}$$

This is standard in Bayesian decoders [36,37]. The uncoupled decoder is constructed similarly. The prior and conditional distributions are specified as

$$\begin{aligned} s(t) &\sim \text{Uniform}(\{\text{Grey}, 30, 60, \dots, 360\}) \\ \mathbf{r}(t)|s(t) &\sim \mathcal{N}(\bar{s}(t), \text{diag}(\sigma_1^2, \dots, \sigma_N^2)) \end{aligned} \tag{14}$$

Here, $\bar{s}(t) = (\bar{s}_1(t), \dots, \bar{s}_N(t))$ is the vector of stimulus-dependent average response. Further, $\text{diag}(\sigma_1^2, \dots, \sigma_N^2)$ is the diagonal matrix with σ_i^2 denoting the variance of neuron i 's fluorescence trace during all frames when the stimulus $s(t)$ was presented. Note that since the conditional distribution posits a diagonal covariance matrix in a multivariate Gaussian, the neurons are independent, and we have the factorization

$$p(\mathbf{r}(t)|s(t)) = \prod_{i=1}^N p_i(\mathbf{r}_i(t)|s(t)) \tag{15}$$

Here, p_i is the marginal density of neuron i 's trace. Similarly, the decoded stimulus is the MAP estimate. Using the 70/30 data split for training and testing from previous analyses, we fit the

LM on the training set, then decoded all of the frames (both training and testing frames) in our analysis.

Estimation of information about stimulus. The extent to which sensory information was extracted from the population response was investigated via an information theoretic approach. Specifically, we computed the mutual information between the decoded stimulus and the true stimulus [36], given by

$$I(s(t), \hat{s}(t)) = \sum_{s(t), \hat{s}(t)} P(s(t), \hat{s}(t)) \log_2 \frac{P(s(t), \hat{s}(t))}{P(s(t))P(\hat{s}(t))} \quad (16)$$

Here, $s(t)$ is the true stimulus, $\hat{s}(t)$ is the decoded stimulus, and the mutual information $I(s(t), \hat{s}(t))$ has units in bits. The joint and marginal distributions were estimated directly from the decoding results. Given the large size of our data, subsampling the distributions was not a concern. This approach gave a meaningful quantification of information and enabled sensible comparisons between decoders. Further, it gave a measure for the computational relevance of the functional group beyond single trial predictive power.

Acknowledgments

The authors wholeheartedly thank Joseph Dechery for collecting the mouse data sets used in the analysis. Additionally, the authors thank John Maunsell, Jackson Cone, Maayan Levy, and Elizabeth de Laitre for their helpful and deliberate comments on the manuscript.

Author Contributions

Conceptualization: Subhodh Kotekal, Jason N. MacLean.

Formal analysis: Subhodh Kotekal, Jason N. MacLean.

Funding acquisition: Jason N. MacLean.

Investigation: Subhodh Kotekal, Jason N. MacLean.

Methodology: Subhodh Kotekal, Jason N. MacLean.

Project administration: Jason N. MacLean.

Resources: Jason N. MacLean.

Software: Subhodh Kotekal.

Supervision: Jason N. MacLean.

Validation: Subhodh Kotekal, Jason N. MacLean.

Visualization: Subhodh Kotekal, Jason N. MacLean.

Writing – original draft: Subhodh Kotekal.

Writing – review & editing: Subhodh Kotekal, Jason N. MacLean.

References

1. Hubel DH. Single unit activity in striate cortex of unrestrained cats. *J Physiol.* 1959 Sep 2; 147(2):226–238.2.
2. Charles AS, Park M, Weller JP, Horwitz GD, Pillow JW. Dethroning the Fano Factor: A Flexible, Model-Based Approach to Partitioning Neural Variability. *Neural Comput.* 2018 Jan 30; 30(4):1012–45. https://doi.org/10.1162/neco_a_01062 PMID: 29381442
3. Dean AF. The variability of discharge of simple cells in the cat striate cortex. *Exp Brain Res.* 1981 Dec 1; 44(4):437–40. <https://doi.org/10.1007/bf00238837> PMID: 7308358

4. Dechery JB, MacLean JN. Functional triplet motifs underlie accurate predictions of single-trial responses in populations of tuned and untuned V1 neurons. *PLOS Comput Biol.* 2018; 14(5):1–23.
5. Goris RLT, Movshon JA, Simoncelli EP. Partitioning neuronal variability. *Nat Neurosci.* 2014 Jun; 17(6):858–65. <https://doi.org/10.1038/nn.3711> PMID: 24777419
6. Heggelund P, Albus K. Orientation selectivity of single cells in striate cortex of cat: The shape of orientation tuning curves. *Vision Res.* 1978 Jan 1; 18(8):1067–71. [https://doi.org/10.1016/0042-6989\(78\)90037-8](https://doi.org/10.1016/0042-6989(78)90037-8) PMID: 706158
7. Shadlen MN, Newsome WT. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *J Neurosci.* 1998 May 15; 18(10):3870–96. <https://doi.org/10.1523/JNEUROSCI.18-10-03870.1998> PMID: 9570816
8. Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 1983 Jan 1; 23(8):775–85. [https://doi.org/10.1016/0042-6989\(83\)90200-6](https://doi.org/10.1016/0042-6989(83)90200-6) PMID: 6623937
9. Scholvinck ML, Saleem AB, Benucci A, Harris KD, Carandini M. Cortical State Determines Global Variability and Correlations in Visual Cortex. *J Neurosci.* 2015 Jan 7; 35(1):170–8. <https://doi.org/10.1523/JNEUROSCI.4994-13.2015> PMID: 25568112
10. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nat Neurosci.* 2011 Jun; 14:811 EP-. <https://doi.org/10.1038/nn.2842> PMID: 21709677
11. Lin I-C, Okun M, Carandini M, Harris KD. The Nature of Shared Cortical Variability. *Neuron.* 2015 Aug 5; 87(3):644–56. <https://doi.org/10.1016/j.neuron.2015.06.035> PMID: 26212710
12. Deweese MR, Zador AM. Shared and Private Variability in the Auditory Cortex. *J Neurophysiol.* 2004 Sep 1; 92(3):1840–55. <https://doi.org/10.1152/jn.00197.2004> PMID: 15115790
13. Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science.* 1995 Jun 9; 268(5216):1503–6. <https://doi.org/10.1126/science.7770778> PMID: 7770778
14. Carandini M. Amplification of Trial-to-Trial Response Variability by Neurons in Visual Cortex. *PLOS Biol.* 2004 Aug 24; 2(9):e264. <https://doi.org/10.1371/journal.pbio.0020264> PMID: 15328535
15. Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci.* 1993 Jan 1; 13(1):334–50. <https://doi.org/10.1523/JNEUROSCI.13-01-00334.1993> PMID: 8423479
16. Niell CM, Stryker MP. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. *Neuron.* 2010; 65(4):472–9. <https://doi.org/10.1016/j.neuron.2010.01.033> PMID: 20188652
17. Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci.* 2008 Apr; 9(4):292–303. <https://doi.org/10.1038/nrn2258> PMID: 18319728
18. Buonomano DV, Maass W. State-dependent computations: spatiotemporal processing in cortical networks. *Nat Rev Neurosci.* 2009 Feb; 10(2):113–25. <https://doi.org/10.1038/nrn2558> PMID: 19145235
19. Harris KD. Neural signatures of cell assembly organization. *Nat Rev Neurosci.* 2005 May; 6(5):399–407. <https://doi.org/10.1038/nrn1669> PMID: 15861182
20. Shimaoka D, Steinmetz NA, Harris KD, Carandini M. The impact of bilateral ongoing activity on evoked responses in mouse cortex. Palmer S, editor. *eLife.* 2019 Apr 30; 8:e43533. <https://doi.org/10.7554/eLife.43533> PMID: 31038456
21. Abbott LF, Dayan P. The Effect of Correlated Variability on the Accuracy of a Population Code. *Neural Comput.* 1999 Jan 1; 11(1):91–101. <https://doi.org/10.1162/089976699300016827> PMID: 9950724
22. Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A. Information-limiting correlations. *Nat Neurosci.* 2014 Oct; 17(10):1410–7. <https://doi.org/10.1038/nn.3807> PMID: 25195105
23. Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature.* 2006 Apr; 440:1007 EP-. PMID: 16625187
24. Song S, Sjöström PJ, Reigl M, Nelson S, Chklovskii DB. Highly Nonrandom Features of Synaptic Connectivity in Local Cortical Circuits. *PLOS Biol.* 2005 Mar 1; 3(3):e68. <https://doi.org/10.1371/journal.pbio.0030068> PMID: 15737062
25. Ko H, Mrsic-Flogel TD, Hofer SB. Emergence of Feature-Specific Connectivity in Cortical Microcircuits in the Absence of Visual Experience. *J Neurosci.* 2014 Jul 16; 34(29):9812–6. <https://doi.org/10.1523/JNEUROSCI.0875-14.2014> PMID: 25031418
26. Chambers B, MacLean JN. Higher-Order Synaptic Interactions Coordinate Dynamics in Recurrent Networks. *PLOS Comput Biol.* 2016 Aug 19; 12(8):e1005078. <https://doi.org/10.1371/journal.pcbi.1005078> PMID: 27542093
27. Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD. Sparse coding and high-order correlations in fine-scale cortical networks. *Nature.* 2010 Jul; 466(7306):617–21. <https://doi.org/10.1038/nature09178> PMID: 20601940

28. Hu Y, Brunton SL, Cain N, Mihalas S, Kutz JN, Shea-Brown E. Feedback through graph motifs relates structure and function in complex networks. *Phys Rev E*. 2018 Dec 14; 98(6):062312.
29. Womelsdorf T, Valiante TA, Sahin NT, Miller KJ, Tiesinga P. Dynamic circuit motifs underlying rhythmic gain control, gating and integration. *Nat Neurosci*. 2014 Aug; 17(8):1031–9. <https://doi.org/10.1038/nn.3764> PMID: 25065440
30. Hebb DO. *The organization of behavior: a neuropsychological theory*. New York: Wiley; 1949. xix, 335 p. (A Wiley book in clinical psychology).
31. Harris KD, Csicsvari J, Hirase H, Dragoi G, Buzsáki G. Organization of cell assemblies in the hippocampus. *Nature*. 2003 Jul; 424(6948):552–6. <https://doi.org/10.1038/nature01834> PMID: 12891358
32. Stevenson IH. Omitted Variable Bias in GLMs of Neural Spiking Activity. *Neural Comput*. 2018 Oct 12; 30(12):3227–58.
33. Sadovsky AJ, Kruskal PB, Kimmel JM, Ostmeier J, Neubauer FB, MacLean JN. Heuristically optimal path scanning for high-speed multiphoton circuit imaging. *J Neurophysiol*. 2011 Jun 29; 106(3):1591–8. <https://doi.org/10.1152/jn.00334.2011> PMID: 21715667
34. Sompolinsky H, Yoon H, Kang K, Shamir M. Population coding in neuronal systems with correlated noise. *Phys Rev E*. 2001 Oct 17; 64(5):051904.
35. Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, et al. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 2008 Jul; 454:995 EP-. <https://doi.org/10.1038/nature07140> PMID: 18650810
36. Quiñero Quiroga R, Panzeri S. Extracting information from neuronal populations: information theory and decoding approaches. *Nat Rev Neurosci*. 2009 Mar; 10:173 EP-. <https://doi.org/10.1038/nrn2578> PMID: 19229240
37. Paninski L, Pillow J, Lewi J. Statistical models for neural encoding, decoding, and optimal stimulus design. In: Cisek P, Drew T, Kalaska JF, editors. *Computational Neuroscience: Theoretical Insights into Brain Function* [Internet]. Elsevier; 2007. p. 493–507. (Progress in Brain Research; vol. 165). <http://www.sciencedirect.com/science/article/pii/S0079612306650310>
38. Park IM, Meister MLR, Huk AC, Pillow JW. Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci*. 2014 Aug; 17:1395 EP-. <https://doi.org/10.1038/nn.3800> PMID: 25174005
39. Runyan CA, Piasini E, Panzeri S, Harvey CD. Distinct timescales of population coding across cortex. *Nature*. 2017 Jul; 548:92 EP-. <https://doi.org/10.1038/nature23020> PMID: 28723889
40. Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, et al. The Structure of Multi-Neuron Firing Patterns in Primate Retina. *J Neurosci*. 2006; 26(32):8254–8266. <https://doi.org/10.1523/JNEUROSCI.1282-06.2006> PMID: 16899720
41. Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, et al. Diverse coupling of neurons to populations in sensory cortex. *Nature*. 2015 May; 521(7553):511–5. <https://doi.org/10.1038/nature14273> PMID: 25849776
42. Clancy KB, Orsolich I, Mrsic-Flogel TD. Locomotion-dependent remapping of distributed cortical networks. *Nat Neurosci*. 2019 Mar 11; 1.
43. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*. 2019 Apr 19; 364(6437):eaav7893.
44. Kohn A, Smith MA. Stimulus Dependence of Neuronal Correlation in Primary Visual Cortex of the Macaque. *J Neurosci*. 2005 Apr 6; 25(14):3661–73. <https://doi.org/10.1523/JNEUROSCI.5106-04.2005> PMID: 15814797
45. Ecker AS, Berens P, Tolias AS, Bethge M. The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. *J Neurosci*. 2011; 31(40):14272–14283. <https://doi.org/10.1523/JNEUROSCI.2539-11.2011> PMID: 21976512
46. Maynard EM, Hatsopoulos NG, Ojakangas CL, Acuna BD, Sanes JN, Normann RA, et al. Neuronal Interactions Improve Cortical Population Coding of Movement Direction. *J Neurosci*. 1999 Sep 15; 19(18):8083–93. <https://doi.org/10.1523/JNEUROSCI.19-18-08083.1999> PMID: 10479708
47. Kass RE, Eden UT, Brown EN. Generalized Linear and Nonlinear Regression. In: *Analysis of Neural Data* [Internet]. New York, NY: Springer New York; 2014. p. 391–412. https://doi.org/10.1007/978-1-4614-9602-1_14