

1 ***Recurrent interactions can explain the variance in single trial responses***

2

3 Subhodh Kotekal<sup>1</sup> and Jason N. MacLean<sup>1,2</sup>

4 1. Department of Neurobiology, University of Chicago

5 2. Committee on Computational Neuroscience

6

7 **1. Abstract**

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

21

22 **2. Introduction**

23           The earliest single unit recordings in awake primary visual cortex demonstrated that  
24 many of the recorded action potentials could not be readily attributed to visual stimulus (Hubel  
25 1959). Across visual cortex, single neuron responses to repeated presentations of a given  
26 stimulus exhibit high variability from presentation to presentation (Charles, Park, Weller,  
27 Horwitz, & Pillow, 2018; Dean, 1981; Dechery & MacLean, 2018; Goris, Movshon, &  
28 Simoncelli, 2014; Heggelund & Albus, 1978; Shadlen & Newsome, 1998; Tolhurst, Movshon, &  
29 Dean, 1983). Response variability is largely restricted to neocortex (Scholvinck, Saleem,  
30 Benucci, Harris, & Carandini, 2015) and is shared across the neocortical neuronal population  
31 (Cohen & Kohn, 2011; Lin, Okun, Carandini, & Harris, 2015). However, neurons from a  
32 physiological perspective are capable of being highly reliable (Deweese & Zador, 2004; Mainen  
33 & Sejnowski, 1995) suggesting that variance arises primarily from synaptic inputs (Carandini,  
34 2004; Softky & Koch, 1993) and from extraretinal factors that are not visual stimulus such as  
35 arousal and locomotion (Goris et al., 2014; Niell & Stryker, 2010).

36           Neuronal variability can be taken to be purely noise (Faisal, Selen, & Wolpert, 2008),  
37 insofar as it is detrimental to the stability of sensory representation. Alternatively, single trial  
38 variability may reflect ongoing cortical dynamics associated with sensory processing and  
39 representation (Buonomano & Maass, 2009; Harris, 2005). Shared variability may similarly  
40 reflect ongoing dynamics (Shimaoka, Steinmetz, Harris, & Carandini, 2019) as pairwise  
41 correlations in a population of neurons can affect sensory representation (Abbott & Dayan, 1999;  
42 Moreno-Bote et al., 2014; Schneidman, Berry II, Segev, & Bialek, 2006). Neurons are highly  
43 interconnected and thus complex network interactions likely shape the activity of neurons  
44 (Dechery & MacLean, 2018; Song, Sjöström, Reigl, Nelson, & Chklovskii, 2005). In visual  
45 cortex, visually tuned neurons with similar stimulus selectivity are more likely to be synaptically

46 interconnected (Ko, Mrsic-Flogel, & Hofer, 2014) and these connections manifest as specific  
47 motifs (Song et al., 2005) which coordinate synaptic integration (Chambers & MacLean, 2016).  
48 Moreover, populations of neurons have been shown to exhibit higher-order state correlations  
49 (Ohiorhenuan et al., 2010; Schneidman et al., 2006), and topological network features have been  
50 shown to shape spike propagation and information transfer (Chambers & MacLean, 2016; Hu et  
51 al., 2018; Womelsdorf, Valiante, Sahin, Miller, & Tiesinga, 2014). However, the relative role of  
52 local synaptic connectivity, as compared to stimulus-related input or more global variables such  
53 as locomotion, in the generation of a sensory representation, particularly in real time, remains  
54 unclear. A comprehensive theory of primary visual cortex must encompass local correlations  
55 instead of treating neurons as independent units.

56         The cell assembly hypothesis, laid out by Donald Hebb, proposes that neurons are  
57 organized into mutually excitable groups called “assemblies” which are strongly coactive  
58 (Harris, 2005; Hebb, 1949). Under this hypothesis, sensory representations and cortical state  
59 dynamics manifest as sequences of assembly activations and internal state dynamics need not be  
60 deterministically tied to stimulus from trial-to-trial (Buonomano & Maass, 2009; Harris, 2005;  
61 Hebb, 1949; Scholvinck et al., 2015). Importantly, the assembly hypothesis suggests that single  
62 trial dynamics are strongly influenced by the group of coactive neurons rather than entirely by  
63 the external stimulus.

64         In this work, we investigate whether the variance in V1 neuronal activity over the time  
65 course of single trials is best explained by visual stimuli, locomotion, or by coactive groups of  
66 V1 neurons. We define coactive groups of neurons, which we term as “functional groups”, by  
67 local pairwise correlations of activity within short time intervals after accounting for stimulus  
68 and running effects. We used a generalized linear model (GLM) on calcium imaging data of

69 visually-evoked activity recorded from V1 in freely running mice. In the model, we relate the  
70 target neuron's activity during single trials to the functional group's activity, stimulus condition,  
71 and mouse running speed. Functional group coactivity is understood to be informative of a target  
72 neuron's activity if inclusion of coupling in the GLM improves predictions of the target's neuron  
73 activity, which is suggestive of the importance of internal circuit dynamics in shaping single  
74 trials (Harris, 2005; Harris, Csicsvari, Hirase, Dragoi, & Buzsáki, 2003). With the functional  
75 group as an operational definition of an assembly, we provide concrete descriptions of the  
76 timescales, numerical scale, correlation structure, and computational capabilities of assemblies.

77

### 78 **3. Results**

79         Trial-to-trial, the activity of individual neurons in visual cortex is highly variable despite  
80 ostensibly identical conditions (Shadlen & Newsome, 1998). To evaluate local and global effects  
81 on single trial variability, we built an encoding model using a generalized linear model (GLM),  
82 comprised of three terms: neuron coupling, stimulus and running speed. To determine each  
83 term's contribution, we constructed multiple restricted GLMs. Each excluded a term of interest  
84 and we compared test set predictions to those of the unrestricted GLM (i.e. the model containing  
85 all three terms). Relative change in a restricted model's test set performance indicated the  
86 excluded term's importance in modelling single trial activity. For example, a large increase in a  
87 restricted model's test set error indicated that the excluded variable was informative to single  
88 trial activity. A marginal change indicated that the excluded variable was uninformative.

89         Notably, we used prediction performance rather than the GLM coefficients to determine  
90 the importance of model terms. The full set of variables that truly explain V1 single trial activity  
91 is not *a priori* known, nor experimentally observable. Coefficient estimates would be different if

92 the omitted variables were included in the model, and may lead to erroneous conclusions; this is  
93 omitted variable bias (Stevenson, 2018). Consequently, we used the unrestricted GLM's mean  
94 squared error (MSE) on the test set data as a benchmark against which we compared each  
95 restricted model, and thus evaluated the importance each variable.

96

### 97 **3.1. Mouse Visual Cortical Data**

98 The imaging data used throughout this manuscript was described in a previous study  
99 (Dechery and MacLean 2018). Briefly, the activity of L2/3 excitatory neurons expressing  
100 GCaMP6s in mouse visual cortex (73–347 neurons; n = 8 animals; 21 distinct fields of view; Fig.  
101 1A) were imaged using two photon laser scanning microscopy (25–33 Hz; Dechery and  
102 MacLean 2018; Sadvovsky et al 2011). Mice were awake and allowed to freely run on a linear  
103 treadmill while viewing drifting square-wave gratings presented in 12 directions in pseudo-  
104 random order interleaved with mean-luminance matched grey screen; each trial was a 5-minute  
105 block of stimulus presentation in this format.

106

### 107 **3.2. Experimentally measuring coupling coefficients for use in a GLM**

108 Visual stimuli (Dechery & MacLean, 2018) were repeatedly presented in identical 5-  
109 minute blocks. Partial correlations, i.e. *functional weights*, were computed for each pair of  
110 neurons (Fig. 1B), where each functional weight captured the reliability of coactivity after  
111 accounting for shared changes in activity due to stimulus and running speed. This measure is  
112 analogous to “noise correlations” (Cohen & Kohn, 2011), but also allowed us to account for the  
113 running-induced changes in circuit activity (Niell & Stryker, 2010). To summarize the temporal  
114 component of coactivity, we computed the mean traces of each neuron and constructed a cross-

115 correlogram, the maximum value of which denoted the *functional lag* between the neurons.  
116 Functional connections were often not symmetric, and the direction of the lag gives a direction of  
117 the connection. The associated functional lags most frequently were 0 frames (Fig. 1Biii) and  
118 frequency decreased rapidly with lag amplitude; most coupling was bidirectional (lag 0) (mean =  
119 50.1%, sem = 2.8%, mean over n = 21 datasets). For each neuron, the set of other neurons with  
120 functional connections directed towards the given neuron was called its *functional group*. The  
121 average number of incoming functional connections for a target neuron (i.e. functional group  
122 size) was  $43.18 \pm 3.76$  (mean  $\pm$  s.e.m.). We employed these experimentally measured functional  
123 connections as coupling coefficients in a GLM rather than fitted coefficients to mitigate omitted  
124 variable bias (Stevenson, 2018).

125

### 126 **3.2.1. Performance of the unrestricted GLM**

127 To assess the accuracy of a simple GLM model (consisting of coupling coefficients, a  
128 stimulus term, and a running term), we evaluated our ability to predict single trial neuron  
129 fluorescence. To model the stimulus, we first determined neuronal response properties by  
130 computing the average stimulus-dependent response (averaged across all stimulus presentation  
131 time bins and all blocks). Neurons significantly tuned to orientation or direction were labeled as  
132 *tuned* with the procedure described in (Dechery & MacLean, 2018); all others were labeled  
133 *untuned*. For tuned and untuned neurons, the stimulus term in our GLM was the experimentally  
134 observed average stimulus-dependent response and overall average response respectively. The  
135 running term for both types of neurons was given by the mouse's running speed as measured by  
136 a rotary encoder attached to the axle of the mouse treadmill. Since most functional weights  
137 ( $70.62\% \pm 2.76\%$ ) (mean  $\pm$  s.e.m where mean is over datasets) had associated lags with values 0

138 or 1 (in units of roughly 30 *ms* time bins), we only included these coupling lags in our model.  
139 We note that each lag 0 connection was counted twice to indicate bidirectionality. Pooling across  
140 neurons, the unrestricted GLM exhibited accurate predictions of single trial neuronal  
141 fluorescence (fluorescence traces normalized for each neuron) as indicated by a small test set  
142 mean squared error (MSE) (median = 0.0402, IQR = 0.0616, [25<sup>th</sup> percentile = 0.0233, 75<sup>th</sup>  
143 percentile = 0.0849] Fig. 2A). The test set MSE was stable over the entirety of stimulus  
144 presentation since the first second after stimulus onset exhibited similar test MSE relative to the  
145 entire trial (median = -0.80%, IQR = 17.54%).

146 The accurate predictions of the unrestricted GLM allowed us to next evaluate which  
147 individual terms best contribute to capturing single trial neuronal dynamics. To isolate the  
148 relative contribution of each term to the prediction accuracy of single trial fluorescence, we  
149 excluded the stimulus or running term (*stim-restricted GLM* and *run-restricted GLM*  
150 respectively) and compared these models against the prediction accuracy of the *unrestricted*  
151 *GLM* (Fig. 2C). Each of the three GLMs were individually fit on the same training data and  
152 compared on the same test data (roughly a 70/30 data split). Additionally, we fit GLMs in which  
153 all neuronal couplings were excluded (i.e. *coupling-restricted GLM*) to determine the extent to  
154 which the functional group contributed to accurate prediction of single trial responses.

155

### 156 **3.2.2. Performance of the stimulus-restricted GLM**

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

159 First, as in the unrestricted GLM above, we chose to model the visual stimuli as trial  
160 averaged responses given that the transform(s) of visual information as it transits from retina to

161 LGN to layer 4 to layer 2/3 is not fully known. In particular, we averaged the empirically  
162 measured stimulus-dependent fluorescence across all time bins per stimulus condition across all  
163 trials (each trial being a single 5-minute block) (Fig. 3A), as described above. We then fit a GLM  
164 containing all terms (i.e. unrestricted GLM) and a GLM with all terms except the stimulus term  
165 (stimulus-restricted GLM). When comparing between our two models of the stimulus, we  
166 restricted our analysis of the results to only the drifting grating (i.e. non-grey stimulus) imaging  
167 frames.

168 Pooling across neurons, the stimulus-restricted GLM exhibited marginal test set MSE  
169 increase (all frames: median = +0.20%, IQR = 1.43%, stimulus frames: median = +0.22%, IQR =  
170 1.79%, Fig. 3B). The stim-restricted GLM's marginal MSE increase suggests that knowledge of  
171 the average response does not improve single trial predictions when conditioning on the  
172 functional group's activity and the running speed. It must be noted, however, that neurons in the  
173 95<sup>th</sup> percentile of stimulus-restricted GLM test set MSE increase exhibit a large increase (median  
174 = +29.73%, IQR = 29.11% across stimulus frames, Fig. 3B), indicating that average response  
175 captured a large component of the single trial dynamics for a small subset of neurons.

176 Our second model of the stimulus attempted to account for the temporal dynamics of a  
177 stimulus-dependent response. In this model, tuned and untuned neurons were treated the same.  
178 For each neuron, we averaged its fluorescence trace in response to a given stimulus across all  
179 presentations and across all trials to obtain the "block averaged trace" for a given stimulus (Fig.  
180 3A). The block averaged trace was thus obtained after averaging over 9-15 traces depending on  
181 the dataset; as noted above, we tested this model only on stimulus frames (i.e. non-grey frames).  
182 In order to avoid averaging over different numbers of presentations, we took a 50/50 split of the  
183 data for training and testing the models. Notably, by considering fewer trials, we necessarily



184 were less able to completely isolate and account for the other two variables, coupling and  
185 running speed using partial correlation. As before, we fit an unrestricted GLM and the stimulus-  
186 restricted GLM, and we restricted analysis to the drifting grating (i.e. non-grey stimulus) frames.

187         Across the population, this variant of the stimulus-restricted GLM again exhibited  
188 marginal increase in test set MSE (median = +6.64%, IQR = 5.02% across stimulus frames, Fig.  
189 3B). For the majority of neurons, averaged summaries of stimulus-dependent responses are ill-  
190 equipped to explain single trial activity even when accounting for the average temporal  
191 component of dynamics. Again, neurons in the 95<sup>th</sup> percentile of stimulus-restricted test MSE  
192 increase exhibited large MSE increase (median = +50.59%, IQR = 36.87% across stimulus  
193 frames, Fig. 3B). The remainder of our analysis in this paper used the average response model  
194 for the stimulus term unless otherwise specified.

195

### 196 **3.2.3. Performance of the run-restricted GLM**

197         We evaluated the contribution of the running term (running speed measured by rotary  
198 encoder) to accurate predictions of single neuron fluorescence transients. The run-restricted  
199 GLM exhibited marginal changes in test set MSE compared to the unrestricted GLM (median = -  
200 0.01%, IQR = 0.20%, Fig. 3C), indicating that inclusion of running speed failed to improve  
201 prediction accuracy when conditioning on functional group activity and the average response. In  
202 numerous cases, we obtained a negative coefficient for the running term as we did not constrain  
203 model coefficients to be nonnegative. While we do not focus on the coefficients in the majority  
204 of the paper, it was necessary to address these particular values because it is well known that  
205 running speed strongly modulates spike rates (Niell & Stryker, 2010), and our data similarly  
206 showed enhanced population activity during periods of running (Dechery & MacLean, 2018).

207 Given that (1) we are inevitably subsampling visual cortical circuits and thus cannot include all  
208 relevant variables in the GLM, and (2) we had found that run speed coefficients are mostly small  
209 in magnitude (median = -0.0001, IQR = 0.0062), it was possible that sign errors were due to  
210 omitted variable bias (Stevenson, 2018). Consequently, we refit the unrestricted GLM and the  
211 run restricted GLM with the additional constraint that coefficients of all model terms (i.e.  
212 coupling, stimulus, and running terms) be nonnegative. Again the change in test set mean  
213 squared error (MSE) relative to the unrestricted GLM was small (median = 0.0060%, IQR =  
214 0.49%, Fig. 3C). Although we observed global increases in neuronal activity across all imaged  
215 neurons during running, the inclusion of a running term did not substantially contribute to  
216 accurate prediction of neuronal fluorescence in held out data.

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

225

#### 226 **3.2.4. Performance of the coupling-restricted GLM**

227 The above results suggested that inclusion of the stimulus condition and running speed  
228 were largely uninformative to generating accurate predictions of single trial neuronal dynamics  
229 in the full GLM. However, it was unclear whether these terms enable accurate predictions when

230 *not* conditioning on functional group activity. We fit a GLM with all terms (with the average  
231 response stimulus model and rotary encoder running model) except coupling on the training set.  
232 The coupling-restricted GLM exhibited a large test set MSE increase relative to the unrestricted  
233 GLM (median = 28.93%, IQR = 47.03%, Fig. 5A top right; across all frames), demonstrating  
234 that the functional group captured variance unexplainable by the average stimulus response and  
235 running speed. This result held when the block-averaged trace was used as the stimulus term  
236 (with the rotary encoder running model), with the coupling-restricted GLM exhibiting a large  
237 increase in test set MSE (median = 22.35%, IQR = 39.80%; across stimulus frames).  
238 Collectively these results indicate that the couplings between neurons contribute substantially to  
239 accurate model predictions.

240

### 241 **3.3. GLM prediction accuracy is sensitive to the specifics of topology and large weights**

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

248 Previous work has shown that weak functional weights are uninformative of single trial  
249 dynamics (Dechery & MacLean, 2018). To examine the importance of specific topological  
250 features amongst the strongest weighted edges in the functional group, we first identified the  
251 edges with weights in the top quartile of magnitudes (termed the “strong edges” and “strong  
252 weights”) and permuted the weights (along with the corresponding lag) among these edges. This

253 permutation scheme preserved the topology of the functional groups but shuffled the associated  
254 weights (Fig. 4A). The new weights of the functional group were then substituted into the  
255 unrestricted GLM. We did not refit the unrestricted GLM, but rather computed the change of  
256 MSE on the training data (here, we revert to analyzing all frames). Refitting the GLM would  
257 have optimized model coefficients to minimize model error, thus obscuring the effect of shuffled  
258 weights. This shuffling procedure was repeated 1000 times, and resulted in a large median  
259 increase of training MSE (median of median increase = +23.57%, IQR of median increase =  
260 57.99%, Fig. 4B top). These results demonstrate that the predictive abilities of functional groups  
261 are highly sensitive to the precise value and association of each strong functional weight to a  
262 specific edge.

263         Second, we again identified the edges with strong weights (termed “strong edges”). Now,  
264 rather than shuffle only the weights, we shuffled the strongest edges and thereby generated new  
265 topologies (Fig. 4A). This allowed us to determine the sensitivity of prediction ability to the  
266 underlying strong edges. To elaborate, this shuffling procedure shuffled *all* of the strong edges  
267 across *all* functional groups. After shuffling, each neuron’s functional group was changed since  
268 the composition of a functional group is defined by the specific incoming edges that each neuron  
269 receives. Again the GLM was not refit for the same reasons articulated above, and this procedure  
270 was repeated 1000 times. Shuffled edges resulted in a large median increase of training set MSE  
271 (median of median increase = +20.00%, median of IQR = 31.43%, Fig. 4B right).

272         Thus, both the specific weights and the specific connections in the upper quartile of edges  
273 define the functional group which in turn underlies the impressive performance of the coupled  
274 model. Large weights thus explain the importance of specific edges in the functional group for  
275 accurate predictions and are a signature of the important timescales and weights (regardless of

276 sign) of the functional group; consequently, the timescale and weights of Hebbian assemblies can  
277 be further characterized by observing which timescales and weight sign correspond to the edges  
278 with large weights in the functional group.

279

### 280 **3.4. Recurrent coupling coefficients contribute to accurate predictions**

281 Given the importance of the functional group, we next asked how its temporal structure  
282 related to prediction. To do so we constructed two models,  $GLM_0$  and  $GLM_1$ , which respectively  
283 contained only a lag 0 or a lag 1 coupling term, and compared prediction performance to an  
284 unrestricted GLM. We additionally examined the large weight distributions for lag 0 and lag 1  
285 coupling.

286 We found that the  $GLM_0$  exhibited a marginally increased test set MSE while  $GLM_1$   
287 exhibited a large test set MSE increase as compared to the unrestricted model that contained both  
288 coupling terms ( $GLM_0$  median = +0.63%, IQR = 2.16%,  $GLM_1$  median = +9.81%, IQR =  
289 21.10%, Fig. 5B). This result demonstrates that lag 0 coupling is necessary to generate accurate  
290 single trial predictions and suggests that that prediction-relevant coactivity between neuronal  
291 pairs occurs within a window less than or equal to  $\sim 30$  ms. Hence, the predictive ability of  
292 functional group coactivity is driven by the functional group's bidirectional edges.

293 The sufficiency of lag 0 weights and insufficiency of lag 1 weights to predict single trials  
294 is informed by the prevalence of strong weights associated with lag 0 edges (“strong” defined as  
295 belonging to the top quartile) (Fig. 5C). Furthermore, we observe that strong coupling exhibits  
296 more lag 0 weights than lag 1, which is consistent with our previous results regarding large  
297 weights. While it is the case that lag 0 and lag 1 functional input are linearly correlated (median  $r$   
298 = 0.56, IQR = 0.40, Fig. 5D), collinearity cannot totally explain the sufficiency of lag 0. If

299 activity from lag 0 and lag 1 coupling are collinear, then  $GLM_1$  should have performed  
300 comparably to  $GLM_0$ , as the lag 1 coupling GLM coefficient would be modified linearly to  
301 account for the absence of lag 0 coupling; yet,  $GLM_1$  performed worse than  $GLM_0$ .  
302

### 303 **3.5. Necessity of positive coupling and redundancy of negative coupling**

304 We next asked to what extent prediction was attributable to positive or negative coupling  
305 coefficients. Negative correlations have been shown theoretically to increase coding capacity  
306 (Sompolinsky, Yoon, Kang, & Shamir, 2001) and consequently were of particular interest here.  
307 Two models,  $GLM^+$  and  $GLM^-$ , were constructed where only the positive weight term or  
308 negative weight term was included respectively; additionally, an unrestricted GLM was fit.  
309 Stimulus and running speed terms were included in all models; however, we did not restrict the  
310 set of lags under consideration.  $GLM^+$  exhibited a marginal increase in test set MSE (median =  
311 +0.24%, IQR = 1.43%, Fig. 6A) whereas  $GLM^-$  exhibited a large increase in test set MSE  
312 (median = +25.38%, IQR = 36.37%, Fig. 6A). These results indicate that positive weights are  
313 sufficient for accurate prediction while negatively correlated neurons contributed minimally.  
314 This result further implies that the positive and negative coupling model terms are not collinear  
315 as  $GLM^+$  and  $GLM^-$  did not exhibit comparable performance. The performance difference of  
316  $GLM^+$  and  $GLM^-$  is notably marked by the prevalence of strong positive as compared to negative  
317 weights. Across datasets, the top quartiles of weight magnitudes are comprised of many positive  
318 weights and few negative weights (Fig. 6B).

319

### 320 **3.6. The size of the informative functional group saturates**

321           Previously we had found that the accuracy of an encoding model increased with the  
322 number of neurons imaged in a dataset (Dechery and MacLean 2018). We set out to establish  
323 whether this positive relationship was due to more of each neuron's respective functional group  
324 being sampled, or whether we had imaged a greater number of complete functional groups. As  
325 expected the total number of incoming edges increased with the number of neurons imaged per  
326 dataset (Fig. 6C – slope=0.21). However, this was not the case when isolating the positive  
327 bidirectional edges in the upper quartile of weights. Rather the number of edges in this subset  
328 remained relatively stable regardless of the number of neurons imaged (slope was 0.028; Fig.  
329 6C). These results suggest that the increase in model accuracy with increasing numbers of  
330 neurons sampled is the consequence of capturing a greater number of complete functional  
331 groups. Correspondingly, more neurons are accurately modeled instead of obtaining a better  
332 delineation of neurons' functional groups. This result suggests a numerical size for a functional  
333 group.

334

### 335   **3.7. Functional weights enhance single trial decoding**

336           While the functional group enables accurate single trial encoding predictions, it remained  
337 unclear if functional groups are computationally relevant to decoding. More concretely, we asked  
338 if using the unrestricted GLM, in which the structure of the functional group is known, to decode  
339 the stimulus would result in better performance than decoding when neurons are assumed to be  
340 uncoupled since the decoding performance is suggestive of the computational relevance of the  
341 functional group (Pillow et al., 2008).

342           We constructed coupled and uncoupled decoders under the Bayesian decoding  
343 framework. In particular, a uniform prior over all stimulus conditions was adopted and the

344 stimulus was decoded via the *maximum a posteriori* estimate. All frames (i.e. both training and  
345 test frames) were decoded, and both decoders performed better than chance (Fig. 7A), where  
346 chance performance is given by  $\frac{1}{13} \approx 7.69\%$  (12 stimulus conditions and 1 grey condition). The  
347 coupled GLM decoder decoded more accurately than the uncoupled decoder (coupled GLM  
348 mean accuracy = 64.77%, sem = 2.30%, uncoupled mean accuracy = 40.90%, sem = 2.74%,  
349 accuracy pooled over all frames, mean over datasets), indicating that knowledge of coupling aids  
350 in the extraction of sensory information from the population response. We further computed the  
351 mutual information between the true stimulus and the predicted stimulus (Quiari Quiroga &  
352 Panzeri, 2009), determining the coupled GLM decoder extracts roughly 64% more sensory  
353 information (mean = +63.94%, sem = 14.69%, mean over datasets, Fig. 7B). These results  
354 indicate, supposing downstream elements have the theoretical capability to read out the visual  
355 stimulus, that the functional group dramatically enhances the capability of decoding of the  
356 stimulus.

357

#### 358 **4. Discussion**

359 Knowledge of the coactivity of a V1 neuron's functional group enables accurate  
360 predictions of short term dynamics within single trials. Global descriptions such as averaged  
361 stimulus-dependent responses and running fail to meaningfully capture single trial activity in the  
362 vast majority of L2/3 neurons in visual cortex. With the functional group operationalizing the  
363 notion of a Hebbian assembly, our results give concrete descriptions of the timescales, numerical  
364 scale, correlation structure, and computational capabilities of assemblies.

365

#### 366 **4.1. Modelling fluorescence with GLMs**



367           The assembly hypothesis posits that internal state dynamics of the underlying circuit  
368 shape single trial dynamics; consequently, assembly coactivity is proposed to capture single  
369 trials to a greater extent than global variables. To test this proposition, we used various GLMs as  
370 encoding models (Paninski, Pillow, & Lewi, 2007; Park, Meister, Huk, & Pillow, 2014; Pillow et  
371 al., 2008; Runyan, Piasini, Panzeri, & Harvey, 2017). We used partial correlations to capture  
372 shared variability and act as coupling coefficients in the GLM rather than relying on fitted  
373 coefficients (Park et al., 2014; Pillow et al., 2008; Runyan et al., 2017). Using empirically  
374 measured correlations for coupling coefficients is an attractive modeling choice; fitted coupling  
375 coefficients necessarily depend on the model and variable specification, and are thus at risk of  
376 omitted variable bias (Stevenson, 2018). Since we use the functional group as an operational  
377 notion of an assembly, coupling must not change with specifications of the linear model.

378

#### 379 **4.2. Functional group structure and prediction**

380           We found that the unrestricted GLM accurately modeled the dynamics of single trial  
381 activity, thus justifying the decision to adopt a linear model, use partial correlations, and restrict  
382 to lags 0 and 1. The target neuron's variability captured by the model is a deterministic function  
383 of the variabilities of the functional group, averaged stimulus-dependent response, and running  
384 speed and did not require an explicit stochastic term. Our model results showed functional group  
385 coactivity was the main predictor of single trial activity of V1 neurons, in support of the  
386 assembly perspective (Harris et al., 2003) as well as previous work showing that pairwise  
387 correlations can explain activity patterns (Schneidman et al., 2006; Shlens et al., 2006).

388           Investigating the robustness of prediction to modulations of the functional group, we  
389 found that modulating the topology of large weights severely degrades the model's ability to

390 predict single trial responses. Hence, the precise instantiation of the strong weights and edges in  
391 the functional group is essential for prediction. While global population coupling and  
392 fluctuations have been shown to be informative to predicting activity (Clancy, Orsolich, & Mrcsic-  
393 Fogel, 2019; Lin et al., 2015; Okun et al., 2015; Scholvinck et al., 2015; Stringer et al., 2019),  
394 somewhat consistent with our block averaged stimulus and population average response  
395 locomotion models, our results point to the inclusion of a local description when predicting  
396 neuronal dynamics within single trial responses. Our results largely suggest that strong  
397 functional group coupling (a local description) disproportionately captures dynamics, in  
398 congruence with the assembly hypothesis and other previous work (Ohiorhenuan et al., 2010).

399         With the functional group established as the primary predictor of single trials, we  
400 investigated its features. Functional group edges with functional lag 0 are sufficient for accurate  
401 predictions, suggesting that assembly coactivity largely occurs on timescales of  $\sim 30ms$ , as found  
402 in other regions (Harris et al., 2003). Furthermore, while it is known that correlations between  
403 neurons are predominantly positive and stronger between similarly tuned neurons (Abbott &  
404 Dayan, 1999; Kohn & Smith, 2005; Shadlen & Newsome, 1998), we additionally found that  
405 precisely the large, positive weights most enabled accurate predictions.

406         Using a Bayesian decoding framework, we found that knowledge of the functional  
407 improves decoding over treating neurons as independent units, corroborating previous work  
408 (Ecker, Berens, Tolias, & Bethge, 2011; Maynard et al., 1999) and suggesting that shared  
409 variability is computationally relevant. While the Bayesian decoder does not reveal whether  
410 downstream circuits actually have access to or use the functional group structure, these results  
411 nonetheless demonstrate the theoretical relevance of functional groups (Quian Quiroga &  
412 Panzeri, 2009).

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

418

## 419 **5. Methods**

### 420 **5.1. Data**

421           A subset of imaging data ( $n = 21$  datasets) of mouse visual cortical neurons was taken  
422 from (Dechery & MacLean, 2018). As described in (Dechery & MacLean, 2018), data was  
423 collected from  $n = 4$  male and 4 female C57BL/6J mice expressing transgene Tg(Thy1-  
424 GCaMP6s)GP4.12Dkim (Jackson Laboratory). Activity of L2/3 excitatory neurons in mouse  
425 visual cortex (73–347 neurons; 25–33 Hz; 21 distinct fields of view) were imaged using high  
426 speed two photon laser scanning microscopy (Dechery and MacLean 2018; Sadvovsky et al  
427 2011). Mice were awake, head-fixed, and allowed to freely run on a linear treadmill while  
428 viewing drifting square-wave gratings presented in 12 directions in pseudo-random order  
429 interleaved with mean-luminance matched grey screen.

430

### 431 **5.2. Functional Weights**

432           Partial correlations were used as coupling coefficients in our GLMs; we called these  
433 *functional weights*. The functional weight between a pair of neurons is the partial correlation  
434 between the corresponding fluorescence traces accounting for the stimulus and population-wide  
435 co-activity driven by running. More precisely, for each 5-minute block of oriented drifting

436 grating stimulus presentations (one trial), a trial-specific partial correlation between neuron  $x$  and  
437 neuron  $y$  was computed while accounting for the mean response of neuron  $x$  in all other trials,  
438 the mean response of neuron  $y$  in all other trials, and the mean response of the population  
439 excluding neurons  $x$  and  $y$  in the current trial. Then the trial-specific partial correlations were  
440 averaged across trials to obtain the final partial correlation between neuron  $x$  and neuron  $y$ ; this  
441 was taken to be the functional weight between neurons  $x$  and  $y$ . The mean activity of neurons  $x$   
442 and  $y$  across all other trials was accounted for in order to control for stimulus-dependent effects  
443 resulting in a measure analogous to noise correlations (Dechery & MacLean, 2018). The mean  
444 population activity excluding neurons  $x$  and  $y$  in the current trial is accounted for in order to  
445 control for population co-activity driven by running (Dechery & MacLean, 2018). We used the  
446 MATLAB 2016a function `parcor.m` to compute partial correlations.

447         In order to capture the temporal structure of neuron coactivity, functional weights were  
448 given directionality in the following manner. The mean fluorescence trace (averaged across  
449 trials) for neuron  $x$  and neuron  $y$  was computed, and the cross-correlogram was computed. Due  
450 to the nature of two-photon imaging, neuronal fluorescence time series are given by a series of  
451 time bins with roughly 30 *ms* width; consequently, lags are given in units of time bins. The lag  
452 corresponding to the maximum of the cross-correlogram determined the lag of the functional  
453 weight and the sign of the lag determined the direction of the functional weight (Dechery &  
454 MacLean, 2018), and this lag was termed the *functional lag*. More concretely, if neurons  $x$  and  $y$   
455 had a functional weight with a positive functional lag, then the corresponding functional weight  
456 has direction corresponding to initial point  $x$  and terminal point  $y$ . If the lag is negative, then the  
457 functional weight has direction corresponding to initial point  $y$  and terminal point  $x$ . Lag 0  
458 corresponds to a bidirectional functional weight. The *functional edge* refers to this directed

459 functional connection. Cross-correlograms were calculated using the MATLAB 2016a functions  
460 `xcorr.m`. For each neuron, the set of neurons with outgoing functional edges to the given  
461 neuron (i.e. neurons with functional edges with which the given neuron was the terminal neuron)  
462 was called the *functional group* of the given neuron.

463

## 464 **5.3. Generalized Linear Model**

### 465 **5.3.1. Predicting Fluorescence**

466 A generalized linear model (GLM) was used as an encoding model to model calcium  
467 imaging data from V1 neurons (Kass, Eden, & Brown, 2014; Park et al., 2014; Pillow et al.,  
468 2008; Runyan et al., 2017). The GLM framework allows us to determine which variables are  
469 important in predicting the target neuron's activity by iteratively excluding variables of interest  
470 and comparing prediction accuracy.

471 Stimulus effects were modeled in two ways in order to determine the extent to which  
472 averaged summaries of stimulus-dependent responses contribute to accurate predictions of single  
473 trial dynamics. First, we modeled the stimulus by the stimulus-dependent average response  
474 (which we refer to as “average response”), in which response to a given stimulus was averaged  
475 across all time bins of the given stimulus presentation across all trials. Neurons significantly  
476 tuned to orientation or direction were labeled as *tuned* with the procedure described in (Dechery  
477 & MacLean, 2018), and all remaining neurons were labeled as *untuned*. The stimulus term in our  
478 GLM was the average stimulus-dependent response for tuned neurons. The stimulus term for  
479 untuned neurons was given by the response given by averaging the stimulus-dependent average  
480 response across all stimulus conditions. During grey frames, the stimulus term was set to zero in  
481 this model. With this model, we used a 70/30 split of the data for training and testing in order to

482 fit the GLM coefficients and test performance. We modeled the stimulus a second way by the  
483 “block averaged trace” in order to reflect dynamics associated to neuronal response. In this  
484 model, tuned and untuned neurons were treated exactly the same. The fluorescence traces in each  
485 presentation of a given stimulus were averaged across blocks (i.e. trials) but not time bins,  
486 producing a block-averaged trace which preserved the dynamics. Since the block-averaged trace  
487 preserves dynamics, we only take averages over the training set when fitting the GLM to avoid  
488 overfitting (when testing, averages are taken over the testing set). In order to ensure that averages  
489 occurred over similar numbers of stimulus presentations between the training and testing sets, we  
490 used a 50/50 split of the data. In our analyses, the block averaged trace was obtained after  
491 averaging over 9-15 traces depending on the dataset. The stimulus term is represented by  $s(t)$  in  
492 our GLM equation. Running effects were modeled without neuron-specific modulation, i.e. we  
493 included a term  $v(t)$  (denoting the running velocity at time  $t$ ) that was constant across all  
494 neurons.

495 The functional group was incorporated into the GLM by computing a linear combination  
496 (with weights given by the functional weights) of the lagged (lag given by the functional lag)  
497 fluorescence traces of all neurons in the functional group (i.e. neurons with outgoing functional  
498 edges to the given neuron). For compact notation, we let  $\mathbf{r}(t)$  be the vector of population  
499 fluorescence at a given time,  $T$  be the total duration of the event, and  $N$  be the total number of  
500 neurons in the population. Then we defined  $W_i^k \in \mathbb{R}^N$  to be the vector of functional weights  
501 corresponding to incoming edges to neuron  $i$  with lag  $k$ . In the entire investigation, we restricted  
502 attention to  $k \in \{0, 1\}$  as the frequency of lagged functional relationships with  $k > 1$  decreased  
503 rapidly.

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

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

506 Here,  $c, \beta_0, \beta_1, \Psi, \Gamma \in \mathbb{R}$  are coefficients and  $\varepsilon_t$  is mean zero noise. The linear offset  $c$  is needed  
507 as different neurons have different baseline fluorescence. The gain terms  $\beta_0, \beta_1$  are needed to  
508 account for the range of incoming edges to each neuron. The  $\Psi$  and  $\Gamma$  coefficients are the  
509 modulation associated with stimulus and running respectively. The notation  $W_i^0$  and  $W_i^1$  refers to  
510 the vector of lag 0 and lag 1 incoming functional weights to neuron  $i$ . More specifically,  
511  $W_i^0, W_i^1 \in \mathbb{R}^N$ , and

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

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

517

## 518 **5.4. GLM Variants**

### 519 **5.4.1. Stimulus-restricted, Run-restricted, and Coupling-restricted GLMs**

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

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

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

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

528 and for the *coupling-restricted GLM*, we had the model

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

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

531

#### 532 **5.4.2. GLM<sub>0</sub> and GLM<sub>1</sub>**

533 To determine the contributions of lag 0 and lag 1 coupling to predictions of single trials,  
534 we similarly excluded these terms iteratively and compared predictive performance to the  
535 unrestricted model. In particular, *GLM<sub>0</sub>* is given by

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

536 This model was then fit on the training data and compared against the unrestricted model. Note  
537 that  $W_i^0$  is not affected by refitting as it determined by the functional weights and functional lags,  
538 which are empirically measured. The process of fitting only affects the model coefficients, i.e.  
539  $c, \beta_0, \Psi, \Gamma$ . Similarly, *GLM<sub>1</sub>* is given by

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

540 and the model is fitted and evaluated with exactly the same procedure as *GLM<sub>0</sub>*.



541

### 542 **5.4.3. GLM<sup>+</sup> and GLM<sup>-</sup>**

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

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

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

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

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

552

### 553 **5.5. Permutations of Functional Group Topology**

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

557 We shuffled the strong weights to determine the importance of strong weights for  
558 prediction. We shuffled the top quartile of weights while preserving the corresponding  
559 underlying edges. Formally, we enumerated the edges with strong weights in our entire  
560 population  $\{(i_1, j_1, w_1), \dots, (i_K, j_K, w_K)\}$ , where a directed edge from neuron  $i_\alpha$  to neuron  $j_\alpha$   
561 exists and  $w_\alpha$  is the associated weight for  $1 \leq \alpha \leq K$ . This enumeration contains all tuples

562  $(i_\alpha, j_\alpha, w_\alpha)$  such that  $|w_\alpha|$  is in the top quartile of all weight magnitudes. We then applied a  
563 uniformly random permutation  $\sigma \in S_K$  (where  $S_K$  is the finite symmetric group on  $K$  points) to  
564 act on the weights, furnishing the enumeration  $\{(i_1, j_1, w_{\sigma(1)}), \dots, (i_K, j_K, w_{\sigma(K)})\}$ . We then  
565 assigned these permuted weights to the corresponding edges (i.e. the directed edge from neuron  
566  $i_\alpha$  to neuron  $j_\alpha$  was assigned the weight  $w_{\sigma(\alpha)}$ ) in the functional group. Furthermore, we shuffled  
567 the functional lags in this way (thus changing the directionality of the underlying edges). This  
568 construction preserves the existence of edges, preserves which edges are in the top quartile of  
569 weight magnitudes, and preserves structure between functional lags and functional weights.  
570 Hence, changes in prediction accuracy are directly related to the importance of the functional  
571 weight to prediction.

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

578

## 579 **5.6. Decoding**

### 580 **5.6.1. Constructing the decoder**

581 To determine whether the functional group is computationally relevant, we used the  
582 functional group and GLM framework to decode the stimulus. We constructed a Bayesian  
583 decoder which estimates the stimulus at time  $t$  given the population response and the running  
584 speed at time  $t$ . The unrestricted GLM (using average response for the stimulus term) was

585 incorporated into the decoder (Paninski et al., 2007) in order to directly link the functional  
586 structure of the population to decoding; this was termed the coupled decoder. The coupled  
587 decoder was then compared to an uncoupled decoder: a Bayesian decoder which treats all units  
588 as conditionally independent (Pillow et al., 2008; Runyan et al., 2017).

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

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

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

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

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

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

597 This is standard in Bayesian decoders (Paninski et al., 2007; Quian Quiroga & Panzeri, 2009).  
598 The uncoupled decoder is constructed similarly. The prior and conditional distributions are  
599 specified as

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

600 Here,  $s(t) = (s_1(t), \dots, s_N(t))$  is the vector of stimulus-dependent average response. Further,  
601  $\text{diag}(\sigma_1^2, \dots, \sigma_N^2)$  is the diagonal matrix with  $\sigma_i^2$  denoting the variance of neuron  $i$ 's fluorescence  
602 trace during all frames when the stimulus  $s(t)$  was presented. Note that since the conditional

603 distribution posits a diagonal covariance matrix in a multivariate Gaussian, the neurons are  
604 independent, and we have the factorization

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

605 Here,  $p_i$  is the marginal density of neuron  $i$ 's trace. Similarly, the decoded stimulus is the MAP  
606 estimate. Using the 70/30 data split for training and testing from previous analyses, we fit the  
607 GLM on the training set, then decoded all of the frames (both training and testing frames) in our  
608 analysis.

609

### 610 **5.6.2. Estimation of information about stimulus**

611 The extent to which sensory information was extracted from the population response was  
612 investigated via an information theoretic approach. Specifically, we computed the mutual  
613 information between the decoded stimulus and the true stimulus (Quiari Quiroga & Panzeri,  
614 2009), given by

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

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

621

622 **Acknowledgments**

623           This work was supported by NIH grant R01EY022338. The authors wholeheartedly  
624 thank Joseph Dechery for collecting the mouse data sets used in the analysis. Additionally, the  
625 authors thank John Maunsell, Jackson Cone, Maayan Levy, and Elizabeth de Laittre for their  
626 helpful and deliberate comments on the manuscript.

627 **References**

- 628 Abbott, L. F., & Dayan, P. (1999). The Effect of Correlated Variability on the Accuracy of a  
629 Population Code. *Neural Computation*, *11*(1), 91–101.  
630 <https://doi.org/10.1162/089976699300016827>
- 631 Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal  
632 processing in cortical networks. *Nature Reviews Neuroscience*, *10*(2), 113–125.  
633 <https://doi.org/10.1038/nrn2558>
- 634 Carandini, M. (2004). Amplification of Trial-to-Trial Response Variability by Neurons in Visual  
635 Cortex. *PLOS Biology*, *2*(9), e264. <https://doi.org/10.1371/journal.pbio.0020264>
- 636 Chambers, B., & MacLean, J. N. (2016). Higher-Order Synaptic Interactions Coordinate  
637 Dynamics in Recurrent Networks. *PLOS Computational Biology*, *12*(8), e1005078.  
638 <https://doi.org/10.1371/journal.pcbi.1005078>
- 639 Charles, A. S., Park, M., Weller, J. P., Horwitz, G. D., & Pillow, J. W. (2018). Dethroning the  
640 Fano Factor: A Flexible, Model-Based Approach to Partitioning Neural Variability.  
641 *Neural Computation*, *30*(4), 1012–1045. [https://doi.org/10.1162/neco\\_a\\_01062](https://doi.org/10.1162/neco_a_01062)
- 642 Clancy, K. B., Orsolic, I., & Mrsic-Flogel, T. D. (2019). Locomotion-dependent remapping of  
643 distributed cortical networks. *Nature Neuroscience*, *1*. [https://doi.org/10.1038/s41593-](https://doi.org/10.1038/s41593-019-0357-8)  
644 [019-0357-8](https://doi.org/10.1038/s41593-019-0357-8)
- 645 Cohen, M. R., & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature*  
646 *Neuroscience*, *14*, 811 EP-.
- 647 Dean, A. F. (1981). The variability of discharge of simple cells in the cat striate cortex.  
648 *Experimental Brain Research*, *44*(4), 437–440. <https://doi.org/10.1007/BF00238837>

- 649 Dechery, J. B., & MacLean, J. N. (2018). Functional triplet motifs underlie accurate predictions  
650 of single-trial responses in populations of tuned and untuned V1 neurons. *PLOS*  
651 *Computational Biology*, *14*(5), 1–23. <https://doi.org/10.1371/journal.pcbi.1006153>
- 652 Deweese, M. R., & Zador, A. M. (2004). Shared and Private Variability in the Auditory Cortex.  
653 *Journal of Neurophysiology*, *92*(3), 1840–1855. <https://doi.org/10.1152/jn.00197.2004>
- 654 Ecker, A. S., Berens, P., Tolias, A. S., & Bethge, M. (2011). The Effect of Noise Correlations in  
655 Populations of Diversely Tuned Neurons. *Journal of Neuroscience*, *31*(40), 14272–  
656 14283. <https://doi.org/10.1523/JNEUROSCI.2539-11.2011>
- 657 Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature*  
658 *Reviews Neuroscience*, *9*(4), 292–303. <https://doi.org/10.1038/nrn2258>
- 659 Goris, R. L. T., Movshon, J. A., & Simoncelli, E. P. (2014). Partitioning neuronal variability.  
660 *Nature Neuroscience*, *17*(6), 858–865. <https://doi.org/10.1038/nn.3711>
- 661 Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nature Reviews*  
662 *Neuroscience*, *6*(5), 399–407. <https://doi.org/10.1038/nrn1669>
- 663 Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., & Buzsáki, G. (2003). Organization of cell  
664 assemblies in the hippocampus. *Nature*, *424*(6948), 552–556.  
665 <https://doi.org/10.1038/nature01834>
- 666 Hebb, D. O. (1949). *The organization of behavior: a neuropsychological theory*. New York:  
667 Wiley.
- 668 Heggelund, P., & Albus, K. (1978). Orientation selectivity of single cells in striate cortex of cat:  
669 The shape of orientation tuning curves. *Vision Research*, *18*(8), 1067–1071.  
670 [https://doi.org/10.1016/0042-6989\(78\)90037-8](https://doi.org/10.1016/0042-6989(78)90037-8)

- 671 Kass, R. E., Eden, U. T., & Brown, E. N. (2014). Generalized Linear and Nonlinear Regression.  
672 In *Analysis of Neural Data* (pp. 391–412). [https://doi.org/10.1007/978-1-4614-9602-1\\_14](https://doi.org/10.1007/978-1-4614-9602-1_14)
- 673 Ko, H., Mrsic-Flogel, T. D., & Hofer, S. B. (2014). Emergence of Feature-Specific Connectivity  
674 in Cortical Microcircuits in the Absence of Visual Experience. *Journal of Neuroscience*,  
675 *34*(29), 9812–9816. <https://doi.org/10.1523/JNEUROSCI.0875-14.2014>
- 676 Kohn, A., & Smith, M. A. (2005). Stimulus Dependence of Neuronal Correlation in Primary  
677 Visual Cortex of the Macaque. *Journal of Neuroscience*, *25*(14), 3661–3673.  
678 <https://doi.org/10.1523/JNEUROSCI.5106-04.2005>
- 679 Lin, I.-C., Okun, M., Carandini, M., & Harris, K. D. (2015). The Nature of Shared Cortical  
680 Variability. *Neuron*, *87*(3), 644–656. <https://doi.org/10.1016/j.neuron.2015.06.035>
- 681 Mainen, Z. F., & Sejnowski, T. J. (1995). Reliability of spike timing in neocortical neurons.  
682 *Science*, *268*(5216), 1503–1506. <https://doi.org/10.1126/science.7770778>
- 683 Maynard, E. M., Hatsopoulos, N. G., Ojakangas, C. L., Acuna, B. D., Sanes, J. N., Normann, R.  
684 A., & Donoghue, J. P. (1999). Neuronal Interactions Improve Cortical Population Coding  
685 of Movement Direction. *Journal of Neuroscience*, *19*(18), 8083–8093.  
686 <https://doi.org/10.1523/JNEUROSCI.19-18-08083.1999>
- 687 Moreno-Bote, R., Beck, J., Kanitscheider, I., Pitkow, X., Latham, P., & Pouget, A. (2014).  
688 Information-limiting correlations. *Nature Neuroscience*, *17*(10), 1410–1417.  
689 <https://doi.org/10.1038/nn.3807>
- 690 Niell, C. M., & Stryker, M. P. (2010). Modulation of Visual Responses by Behavioral State in  
691 Mouse Visual Cortex. *Neuron*, *65*(4), 472–479.  
692 <https://doi.org/10.1016/j.neuron.2010.01.033>



- 693 Ohiorhenuan, I. E., Mechler, F., Purpura, K. P., Schmid, A. M., Hu, Q., & Victor, J. D. (2010).  
694 Sparse coding and high-order correlations in fine-scale cortical networks. *Nature*,  
695 *466*(7306), 617–621. <https://doi.org/10.1038/nature09178>
- 696 Okun, M., Steinmetz, N. A., Cossell, L., Iacaruso, M. F., Ko, H., Barthó, P., ... Harris, K. D.  
697 (2015). Diverse coupling of neurons to populations in sensory cortex. *Nature*, *521*(7553),  
698 511–515. <https://doi.org/10.1038/nature14273>
- 699 Paninski, L., Pillow, J., & Lewi, J. (2007). Statistical models for neural encoding, decoding, and  
700 optimal stimulus design. In P. Cisek, T. Drew, & J. F. Kalaska (Eds.), *Computational*  
701 *Neuroscience: Theoretical Insights into Brain Function* (pp. 493–507).  
702 [https://doi.org/10.1016/S0079-6123\(06\)65031-0](https://doi.org/10.1016/S0079-6123(06)65031-0)
- 703 Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in  
704 parietal cortex during sensorimotor decision-making. *Nature Neuroscience*, *17*, 1395 EP-.
- 705 Pillow, J. W., Shlens, J., Paninski, L., Sher, A., Litke, A. M., Chichilnisky, E. J., & Simoncelli,  
706 E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal  
707 population. *Nature*, *454*, 995 EP-.
- 708 Quian Quiroga, R., & Panzeri, S. (2009). Extracting information from neuronal populations:  
709 information theory and decoding approaches. *Nature Reviews Neuroscience*, *10*, 173 EP-.
- 710 Runyan, C. A., Piasini, E., Panzeri, S., & Harvey, C. D. (2017). Distinct timescales of population  
711 coding across cortex. *Nature*, *548*, 92 EP-.
- 712 Schneidman, E., Berry II, M. J., Segev, R., & Bialek, W. (2006). Weak pairwise correlations  
713 imply strongly correlated network states in a neural population. *Nature*, *440*, 1007 EP-.

- 714 Scholvinck, M. L., Saleem, A. B., Benucci, A., Harris, K. D., & Carandini, M. (2015). Cortical  
715 State Determines Global Variability and Correlations in Visual Cortex. *Journal of*  
716 *Neuroscience*, 35(1), 170–178. <https://doi.org/10.1523/JNEUROSCI.4994-13.2015>
- 717 Shadlen, M. N., & Newsome, W. T. (1998). The Variable Discharge of Cortical Neurons:  
718 Implications for Connectivity, Computation, and Information Coding. *Journal of*  
719 *Neuroscience*, 18(10), 3870–3896. [https://doi.org/10.1523/JNEUROSCI.18-10-](https://doi.org/10.1523/JNEUROSCI.18-10-03870.1998)  
720 03870.1998
- 721 Shimaoka, D., Steinmetz, N. A., Harris, K. D., & Carandini, M. (2019). The impact of bilateral  
722 ongoing activity on evoked responses in mouse cortex. *ELife*, 8, e43533.  
723 <https://doi.org/10.7554/eLife.43533>
- 724 Shlens, J., Field, G. D., Gauthier, J. L., Grivich, M. I., Petrusca, D., Sher, A., ... Chichilnisky, E.  
725 J. (2006). The Structure of Multi-Neuron Firing Patterns in Primate Retina. *Journal of*  
726 *Neuroscience*, 26(32), 8254–8266. <https://doi.org/10.1523/JNEUROSCI.1282-06.2006>
- 727 Softky, W. R., & Koch, C. (1993). The highly irregular firing of cortical cells is inconsistent with  
728 temporal integration of random EPSPs. *Journal of Neuroscience*, 13(1), 334–350.  
729 <https://doi.org/10.1523/JNEUROSCI.13-01-00334.1993>
- 730 Sompolinsky, H., Yoon, H., Kang, K., & Shamir, M. (2001). Population coding in neuronal  
731 systems with correlated noise. *Physical Review E*, 64(5), 051904.  
732 <https://doi.org/10.1103/PhysRevE.64.051904>
- 733 Song, S., Sjöström, P. J., Reigl, M., Nelson, S., & Chklovskii, D. B. (2005). Highly Nonrandom  
734 Features of Synaptic Connectivity in Local Cortical Circuits. *PLOS Biology*, 3(3), e68.  
735 <https://doi.org/10.1371/journal.pbio.0030068>

- 736 Stevenson, I. H. (2018). Omitted Variable Bias in GLMs of Neural Spiking Activity. *Neural*  
737 *Computation*, 30(12), 3227–3258. [https://doi.org/10.1162/neco\\_a\\_01138](https://doi.org/10.1162/neco_a_01138)
- 738 Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C. B., Carandini, M., & Harris, K. D. (2019).  
739 Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437),  
740 eaav7893. <https://doi.org/10.1126/science.aav7893>
- 741 Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in  
742 single neurons in cat and monkey visual cortex. *Vision Research*, 23(8), 775–785.  
743 [https://doi.org/10.1016/0042-6989\(83\)90200-6](https://doi.org/10.1016/0042-6989(83)90200-6)

Fig. 1: (A) Left, anecdotal example fluorescence traces of neuronal responses to oriented drifting gratings (neuron 260, dataset 3). 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 generalized linear models (GLMs); 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.

Fig. 2: (A) Distribution of the percent change in test set mean-squared error (MSE) for the unrestricted GLM across all neurons. The unrestricted GLM 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 GLM models the single trial responses of many neurons accurately with a right skewed distribution. (B) Top, time-varying model terms and subsequent prediction for anecdotal neuron (neuron 125, dataset 3). Bottom, diagram for construction of restricted models. The relative changes in test set MSE of the restricted GLMs compared to the unrestricted GLM indicate the importance of the excluded term in generating accurate predictions of single trial responses. For example, if the stimulus-restricted GLM exhibits a minimal test set MSE increase relative to the test set MSE of the unrestricted GLM, then the stimulus is understood to be uninformative to predicting single trial responses.

Fig. 3: (A) Top, stimulus-dependent average response profile for an anecdotal example neuron (neuron 2, dataset 3). 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 GLMs 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 GLM using the average response stimulus model and the stimulus-restricted GLM using the block averaged trace stimulus model. Right, zoom in on 95<sup>th</sup> percentile. (C) Cumulative distribution functions of the percent change of test set MSE across all frames for the run-restricted GLMs using the rotary encoder running model, nonnegative model coefficients, and average population response running model.

Fig. 4: (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  $\beta$ ,  $\delta$ , and  $\xi$  are all permuted, and the bidirectional coupling associated with  $\delta$  follows the permuted weight  $\delta$ . 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

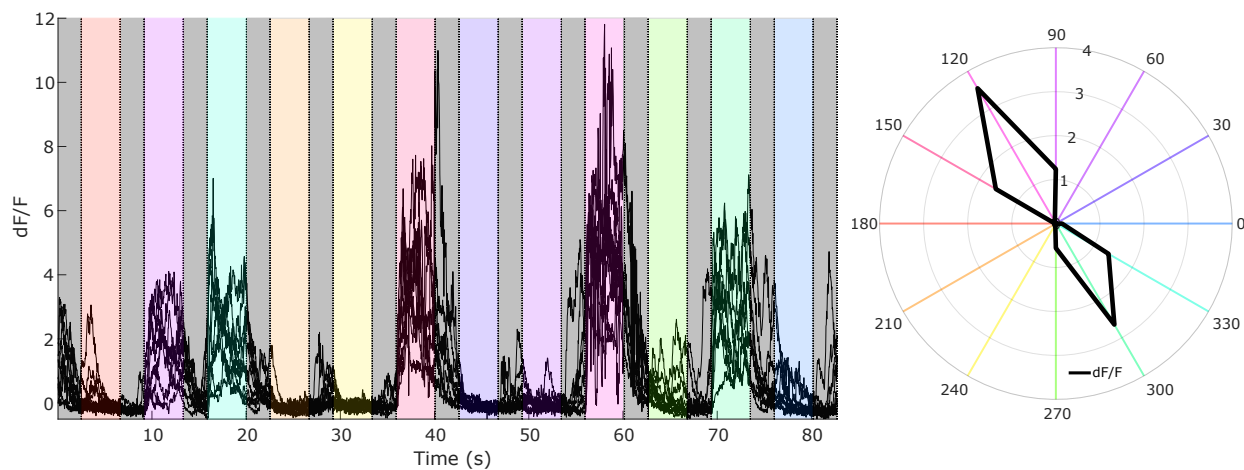
permuting edges, new functional group topologies (in the sense of neuron membership) are instantiated. (B) Top, distribution of the median percent change of training set MSE across all 1000 permutations of strong weights. Bottom, distribution of the median percent change of training set MSE across all 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.

Fig. 5: (A) Left, example of functional group of an anecdotal neuron (neuron 37, dataset 3). Diagram shows all neurons with nonzero weight (across all lags) that have directed coupling outgoing to neuron 37 in dataset 3; these neurons constitute the functional group for neuron 37 in dataset 3. Directionality is indicated by arrows and edge weight is denoted by thickness of the edge and arrow head size, with thickness being a monotonically increasing function of weight. Top right, cumulative distribution functions for the percent change in test set MSE across all frames for stimulus-restricted, run-restricted, and coupling-restricted GLMs using the average response stimulus model and the rotary encoder running model. (B) Cumulative distribution function of percent change in test set MSE of  $GLM_0$  and  $GLM_1$  with respect to the unrestricted GLM. 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.

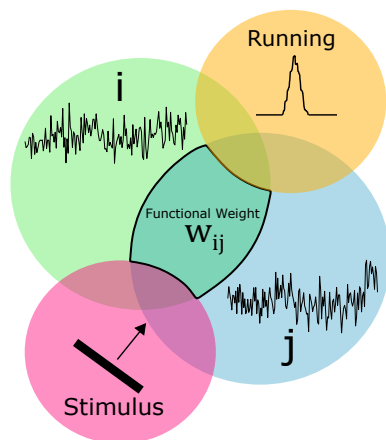
Fig. 6: (A) Cumulative distribution function of percent change in test set MSE of  $GLM_+$  and  $GLM_-$  with respect to the unrestricted GLM. The exclusion of edges with positive weights results in a large increase of test set MSE while the exclusion of edges with negative weights results in only minimal 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).

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


**A**




**Bi**




**Bii**

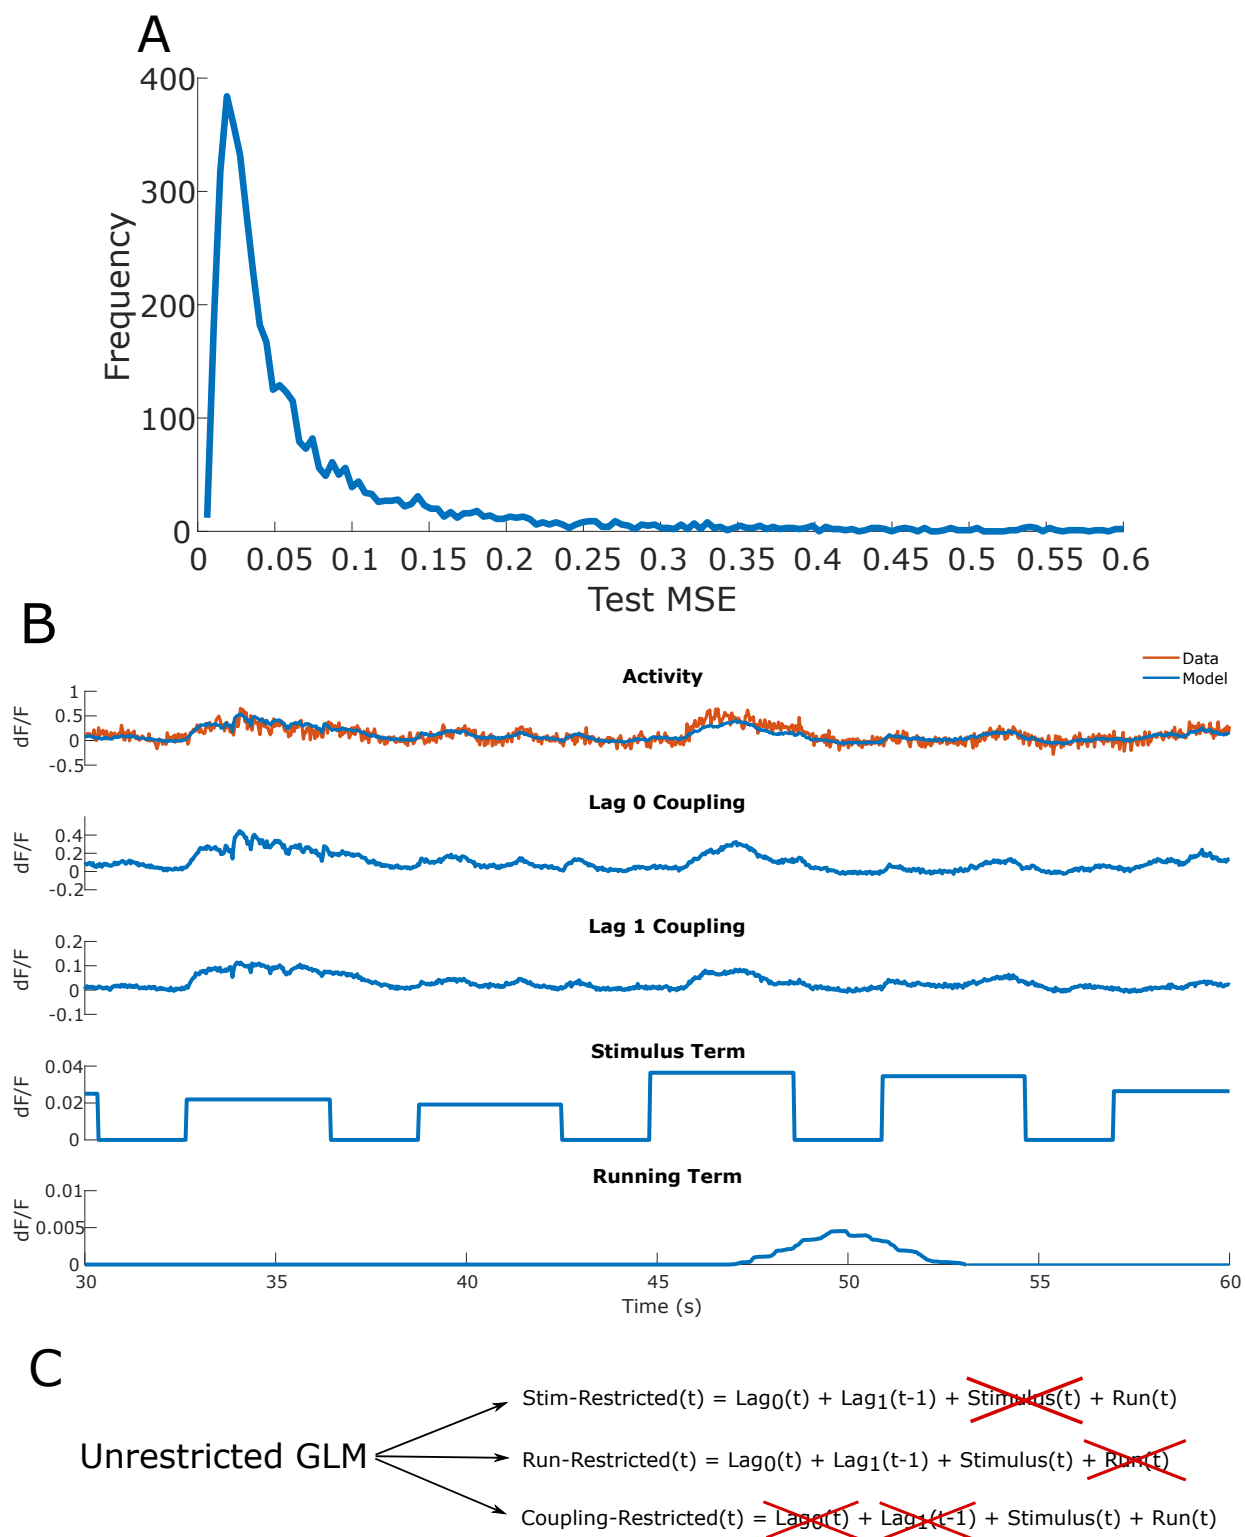


**Biii**

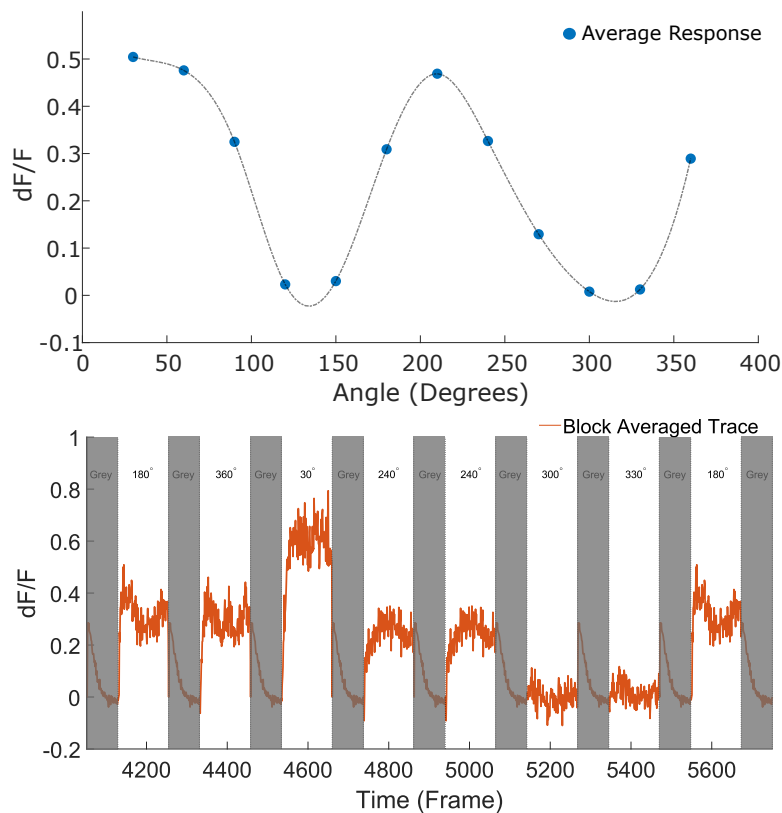


**Biv**



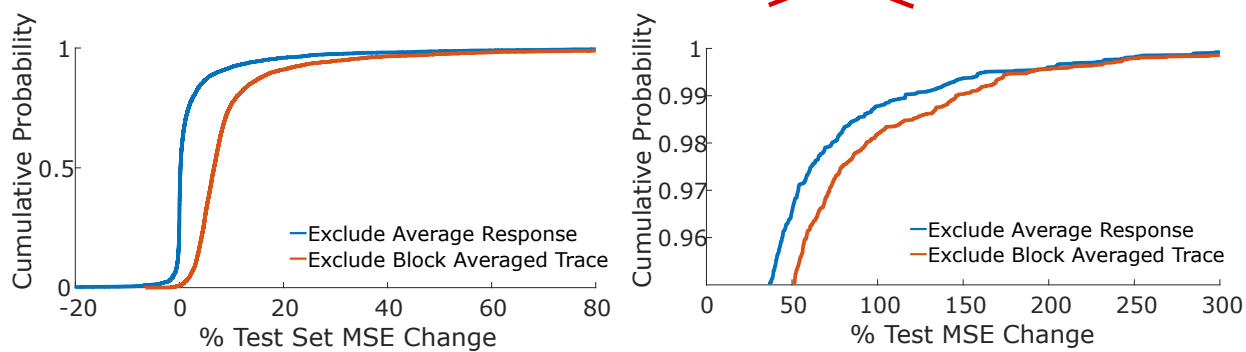


**A**



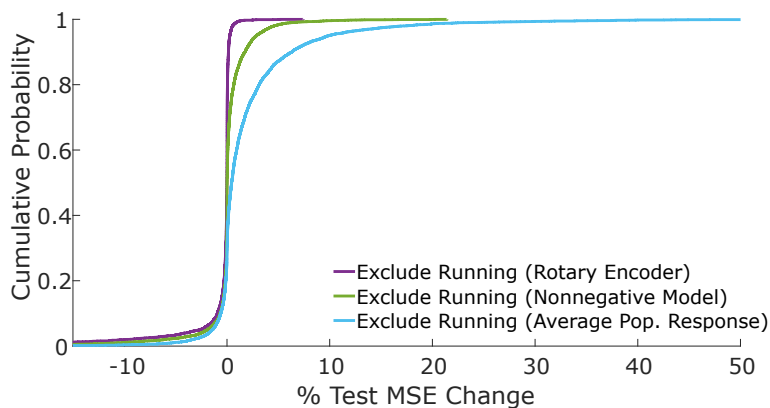
**B**

$$\text{Stim-Restricted}(t) = \text{Lag}_0(t) + \text{Lag}_1(t-1) + \text{Stimulus}(t) + \text{Run}(t)$$



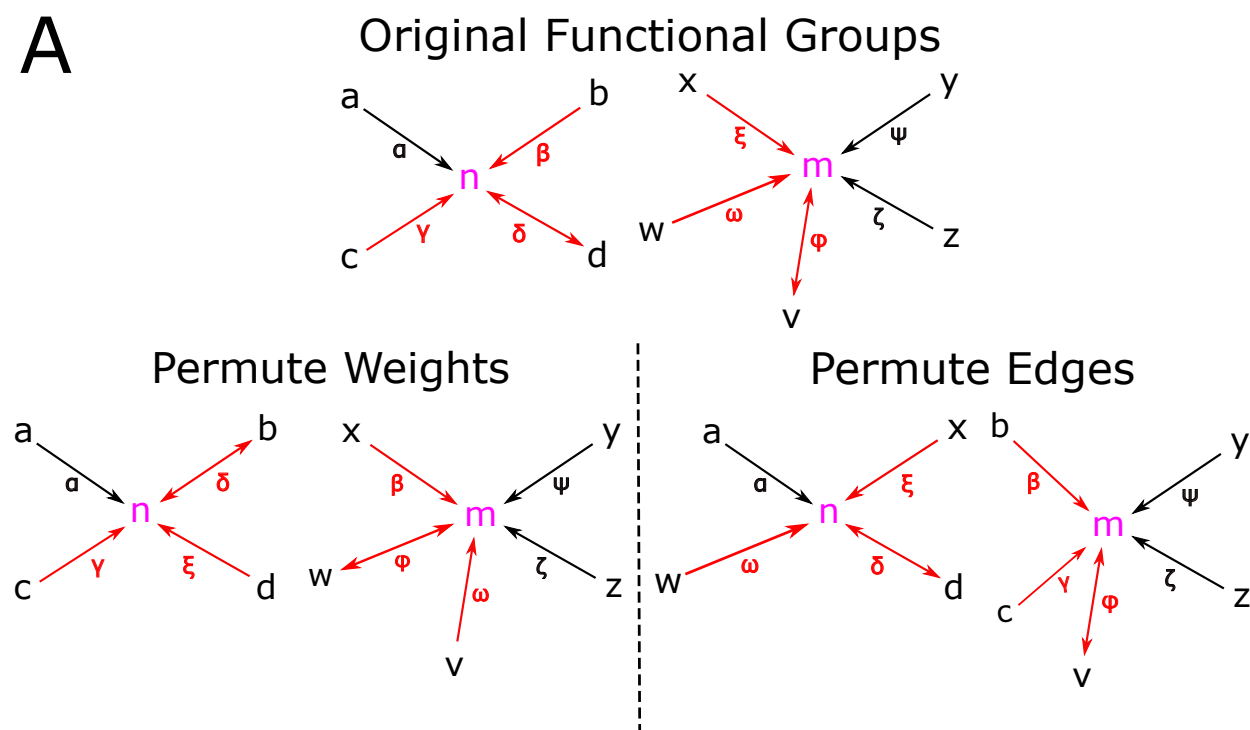
**C**

$$\text{Run-Restricted}(t) = \text{Lag}_0(t) + \text{Lag}_1(t-1) + \text{Stimulus}(t) + \text{Run}(t)$$

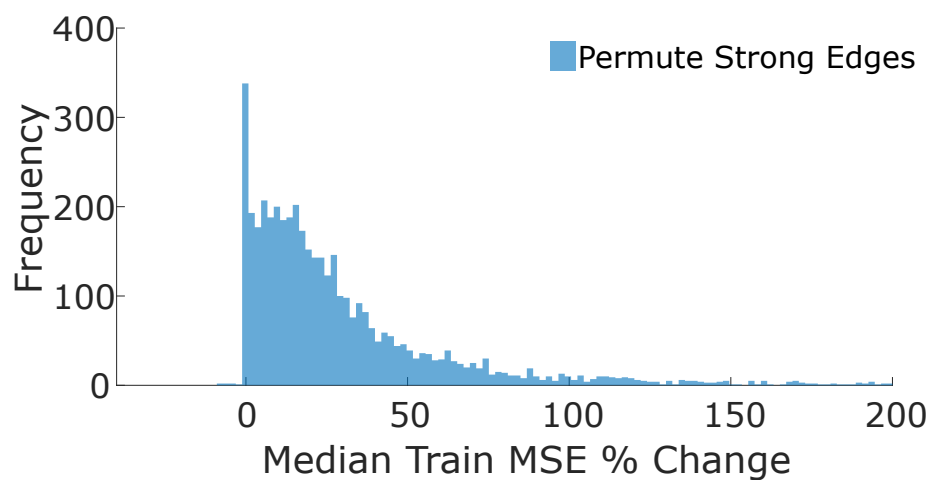
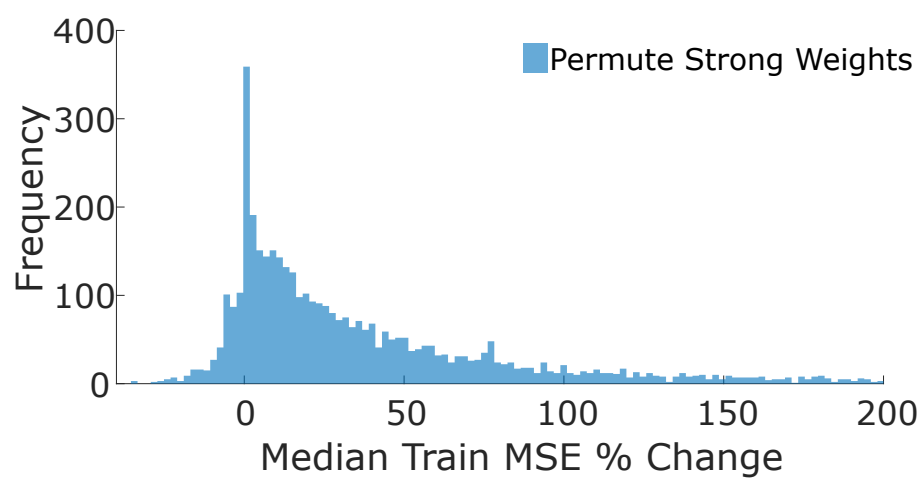




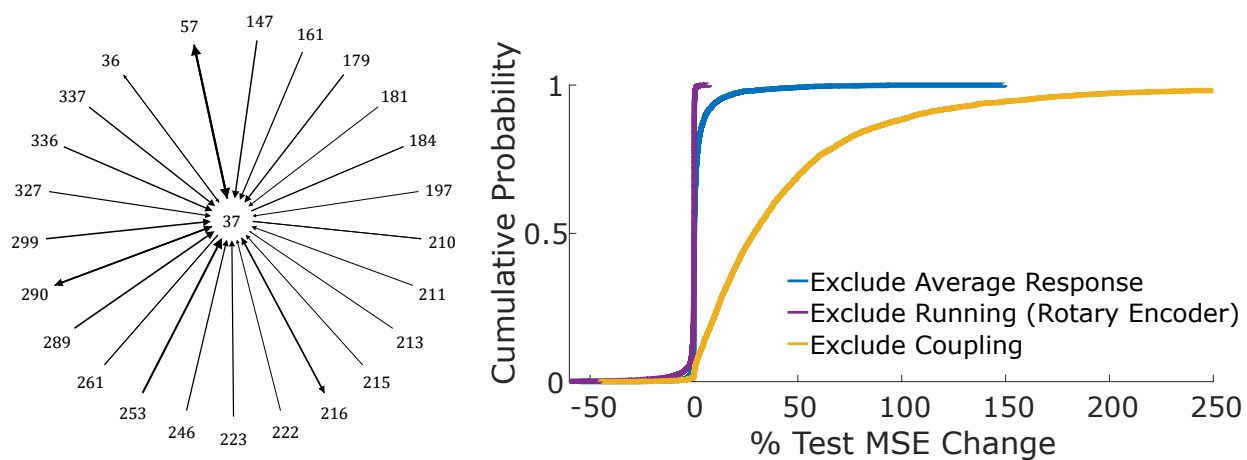
**A**



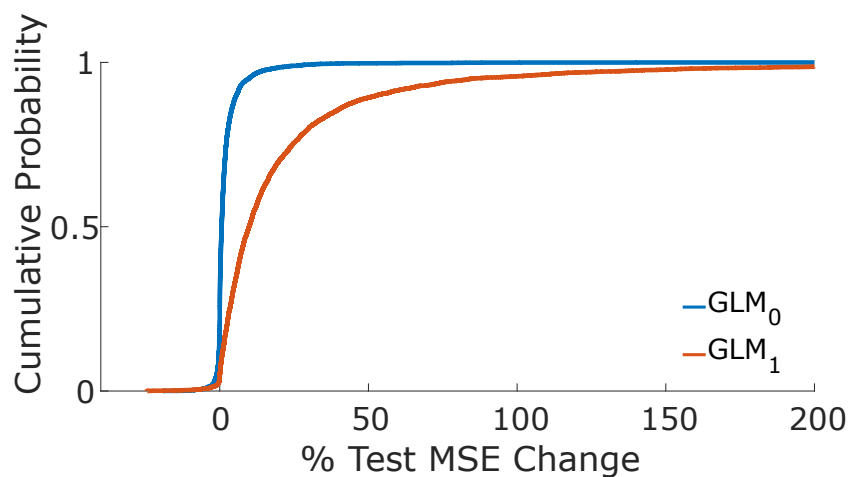
**B**



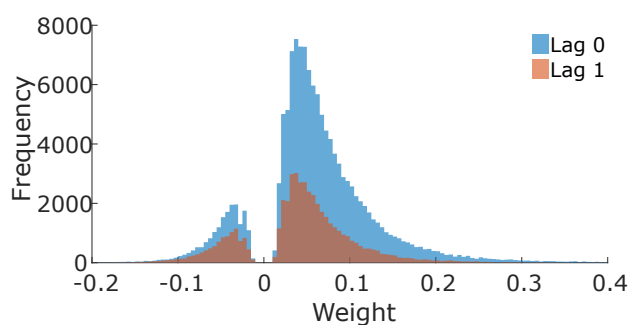
## A Coupling



## B



## C



## D

