Dynamic correlations help prefrontal ensembles transmit information about social behavior

Nicholas A. Frost[1,3,4,5], Anna Haggart[1,3,4,5], and Vikaas S. Sohal[2,3,4,5*]

[1]Department of Neurology, [2]Department of Psychiatry, [3]Center for Integrative Neuroscience, [4]Weill Institute for Neurosciences, [5]Kavli Institute for Fundamental Neuroscience
University of California, San Francisco
*Correspondence to: vikaas.sohal@ucsf.edu

ABSTRACT

How neurons encode behavior is a fundamental question. Neuronal ensembles increase or decrease activity during specific behaviors. However, it is unclear whether ensembles encode information solely via changes in activity levels, or whether changes in correlations between neurons carry additional information. We used microendoscopic GCaMP imaging to measure prefrontal activity while mice were either alone or engaged in social interaction. Using neural network classifiers to measure how well prefrontal neurons transmit information about social behavior to downstream neurons, we find that surrogate datasets which preserve dynamic correlations outperform those which preserve ensemble activity but not correlations. Notably, this ability of correlations to enhance the information transmitted by neuronal ensembles is lost in mice lacking the autism-associated gene Shank3. These results show that dynamically modulated correlations create patterns of coactive neurons which are behaviorally-specific and enhance the information transmitted by neuronal ensembles. Furthermore, this process may be disrupted in pathological states.

INTRODUCTION

During behavior, the activity of neurons is organized with precise temporal relationships (*1–3*). For example, during certain behaviors, subsets of neurons may exhibit correlated activity in which they become active at the same time or within short windows of time. However, it is unknown whether this sort of temporal organization is simply a byproduct of the interconnected nature of neuronal networks (*4*), or contributes in a meaningful way to information encoding (*5*). Groups of co-active neurons represent an attractive computational unit for information processing because they should optimize temporal summation in downstream targets. Thus, increases in correlations might further augment post-synaptic responses when pre-synaptic activity increases, or enhance post-synaptic responses even when the total level of pre-synaptic activity remains constant.

However, it is unclear whether behaviorally-driven changes in correlations actually encode additional behavioral information, beyond what is transmitted by changes in neuronal activity levels. In particular, with the advent of new technologies for simultaneously recording from large numbers of neurons in behaving animals, many studies have now shown that cortical ensembles encode behavioral information via increases or decreases in the activity of their constituent neurons. While correlations have been shown to contribute additional information for small groups (3-8 neurons) of cortical neurons (*6*), only a few studies have examined how correlations contribute to encoding within larger cortical ensembles. One study found that the identity of a conditioned stimulus was encoded in mean activity levels, but not in moment-to-moment patterns of co-activity (*7*). Another study found that in hippocampal region CA1, disrupting correlations impairs the decoding of position, head direction and speed, but did not directly examine whether correlations themselves are dynamically modulated to encode these behavioral variables (*8*). In particular, while multiple studies have shown that behavior can modulate correlations (*3*, *9*) the functional significance of this has remained unclear, because changes in correlations might simply reflect variation in activity levels (*10*) rather than contributing additional information.

To address these questions, we studied the mouse medial prefrontal cortex during simple social behaviors. The role of the medial prefrontal cortex in rodent social behavior is well-established (*11–14*). Many prefrontal neurons are recruited by social interaction (*2*, *13*, *14*) as well as social stimuli such as odors (*15*). These studies show that the activity levels of neuronal ensembles encode social behavior but have not examined whether changes in correlations between prefrontal neurons transmit additional information. Using microendoscopic GCAMP imaging in freely-moving mice, we identified prefrontal ensembles associated with social behavior. We used a neural network classifier to quantify how well these would transmit information about social behavior to downstream neurons. By examining the operation of this neural network and using surrogate datasets which preserve activity levels but either preserve or disrupt correlations, we find that changes in correlations enhance the information transmitted by neuronal ensembles. Notably, this was not the case in a mouse model of autism, demonstrating that this form of information transmission may be disrupted in pathological states.

RESULTS

Social interaction recruits prefrontal ensembles

We implanted microendoscopes (nVoke; Inscopix) into the medial prefrontal cortex (mPFC) of adult wildtype C57/B6 mice (WT) to image calcium transients using GCaMP6f expressed under control of the human synapsin promotor. We imaged freely moving mice during an assay which sequentially introduced 2 novel juvenile mice to the home cage of the subject mouse, first during an initial (novel) epoch and then again during a subsequent (familiar) epoch. These four epochs of social interaction were interleaved with epochs during which the subject mouse was alone in its home cage ('home cage' epochs). The first 5 minutes of each interaction epoch was scored by a blinded observer, and each wild-type mouse spent approximately 10 minutes interacting with the juvenile mice (393 +/- 25 s during the novel epochs and 235 +/- 18 s during the familiar epochs, $p = 0.00017$, paired t test, n = 10 WT mice).

We processed data using a modified PCA/ICA approach (*16*, *17*) to identify neurons which were active during the imaging session. To minimize the influence of the surrounding neuropil on neuronal signals, we calculated the mean signal within each ROI, then subtracted the mean signal calculated from a circular annulus surrounding each ROI (Supplementary Figure 1). Casual inspection of calcium traces revealed that some neurons were more active during epochs of social interaction (compared to periods of home cage exploration), whereas others exhibited the opposite pattern (**Figure 1A**). Correspondingly, aligning calcium traces to the onset of social interaction revealed many neurons that either increased or decreased activity at the onset of interaction (**Figure 1B**). Fluorescence traces were converted to binary event rasters (see Methods for details of event detection), in which most neurons were "active" in less than 5% of frames (**Figure 1C**). As a population, imaged neurons were more active during social interaction (Figure 1C, n = 663 neurons from 10 mice, percent time active in home cage: 1.8% +/- 0.1, percent time active during social interaction: 2.1 +/- 0.1%, $p = 0.00002$, paired t-test). There was a bimodal distribution of cells that were significantly more (>90th percentile, social: 152/663 neurons, home cage: 80/663 neurons; $p < 0.00001$, Chi-Squared Test) or less active (<10th percentile, social: 128/663 neurons, home cage: 119/663 neurons; $p = 0.5$) during either social interaction or matched periods when mice were alone in their home cage, as compared to circularly shuffled datasets (**Figure 1D**). These correspond to neuronal ensembles which are specifically recruited or inhibited during social interaction, respectively.

Using a neural network classifier to assess how well ensembles transmit information

Next, we sought to determine how well these prefrontal ensembles would transmit information about social behavior to downstream neurons, i.e., measure how well downstream neurons could decode whether a mouse was engaged in social behavior based on input from prefrontal neurons. For this we used a simple neural network classifier that received input from the recorded neurons. Our rationale for using this kind of neural network classifier was threefold. First, a simple neural network measures information that is immediately and readily available to downstream neurons. Second, for a neural network with only one hidden layer, it is straightforward to examine the weights to determine how the network performs the classification. This can provide insight into exactly how the neural network is able to decode behavior from the input activity. Third, examining how various parameters of the neural network

affect its performance can provide additional clues about how information is represented within the input population.

Figure 2A shows the design of the neural network classifier. The network consisted of a hidden layer containing 1000 units. We chose this number because it is both an order of magnitude larger than the number of input neurons and an order of magnitude smaller than the number of frames available for training (the latter helps ensure that there will be enough data to train the output weights). We simulated a different neural network for each mouse. Each hidden layer unit received input from a random subset of the prefrontal neurons from one mouse. I.e., each frame represents one timepoint and if neuron *i* is active in a frame then it provided an input of 1 to all the hidden units to which it is connected; otherwise it provides an input of 0. For each simulation, there was a fixed connection probability between each input neuron and each hidden layer unit. We tried different values for this connection probability in order to measure how classifier performance depends on the number of neurons that provide input to each hidden layer unit. Each hidden layer unit had an output weight which specifies how strongly that unit excites or inhibits a single output unit which classifies activity as belonging to periods in which a mouse was actively engaged in social interaction or alone in its home cage. E.g., output unit activity < 0 corresponds to the social condition, while output unit activity > 0 corresponds to the home cage condition. These output weights were adjusted during training (see Methods for details of the training rule) while the pattern of input connectivity was fixed. This models the situation in which prefrontal neurons transmit information to a downstream population of neurons (the hidden layer) that decode behavior via their output weights. We initially trained networks on 50% of the data (frames) and used the held-out data for testing. We trained and tested using intervals during which the mouse was actively engaged in social interaction or equivalent intervals when the mouse was alone in its home cage.

Classifier performance is optimal for intermediate connection probabilities

Classifier performance was strongly dependent on the probability that each input neuron was connected to each hidden unit. For the 8/10 datasets that could be classified above chance, classifier performance (measured on the 50% of data which was held-out during training) was near chance levels when the connection probability was < 0.1, but increased to a peak of 69 +/- 3% (**Figure 2B**) for a connection probability of 0.3. Accuracy decreased dramatically when the connection probability increased to 0.5 indicating that connection probabilities ~0.2 - 0.4 are optimal.

We also validated classifier performance by training and testing on surrogate datasets that were generated by 'swap shuffling' our original datasets. We created 'swap shuffled' surrogate datasets by randomly swapping blocks of activity between neurons (each block of activity = a set of consecutive frames during which the neuron was active). To understand this, think of the entire raster as a collection of blocks of activity. Each block occurs at a specific time, has a specific duration, and is associated with a particular neuron. Swap shuffling is equivalent to just shuffling the neurons associated with each block of activity (the start time and duration of each block do not change). For example, if neuron *i* originally became active at time *t1* for *n1* frames and neuron *j* was active at time *t2* for *n2* frames, then in the surrogate dataset neuron *i* might become active at *t2* (but not at *t1*) while neuron *j* might become active at *t1* (but not *t2*). Swap shuffling preserves the number of neurons active at each point in time (because the timing of

blocks of activity does not change). It also preserves the number of blocks of activity for each neuron, and this tends to preserve the overall level of activity of each neuron. Activity levels are not perfectly preserved, because blocks of activity can have different durations. Nevertheless, in practice, blocks of activity tend to have similar durations and the similarity between the mean activity level in each neuron before and after swap shuffling of entire datasets was 0.97 +/- 0.01. As expected, we found that neural network classifiers trained and tested on swap shuffled datasets performed near chance levels (Figure 2B).

Prefrontal neurons that drive classifier performance exhibit dramatic behaviorally-driven changes in their correlations

Next, we examined connections in trained networks to reveal factors which enable them to successfully classify social vs. home cage behavior (we analyzed networks with a connection probability = 0.3 since this maximized performance of the population). Each hidden layer unit has an output weight which measures how strongly it excites or inhibits the output unit that represents the 'decision' (social vs. home cage). Hidden units with output weights ~0 don't contribute to this decision. By contrast, hidden units with strong negative or positive weights promote the social or home cage decision, respectively (**Figure 3A**). Therefore, we hypothesized that there might be important differences in the pattern of input to hidden units, depending on whether those hidden units have large positive or negative output weights.

We arranged hidden layer units based on their output weights, i.e., the unit with the most negative weight was unit 1 and the unit with the most positive weight was unit 1000. Then we defined the 25 hidden layer units with the most negative weights as 'social units' and the 25 with the most positive weights as 'home cage units' (**Figure 3B**). For comparison we also defined the 25 hidden layer units with the weights closest to zero as 'neutral units.' For each pair of hidden units, we computed the similarity between their inputs (i.e., the correlation between their input vectors; **Figure 3C**). We then plotted the average input similarity of each hidden unit to either the social or home cage units (**Figure 3D**) or the neutral units (**Figure 3E**). Social and home cage units tended to receive input from the same prefrontal neurons as other hidden layer units with the same preference, i.e., which also had negative or positive output weights. By contrast neutral units did not exhibit any such relationship.

The preceding suggests that distinct ensembles of prefrontal neurons provide input to either social or home cage units. We hypothesized that there might be important features of activity in these ensembles that support the classification of social vs. home cage behavior. For example, one possibility is that prefrontal neurons which provide input to social units might tend to increase activity during social behavior, whereas prefrontal neurons which provide input to home cage units do the opposite. Surprisingly, this was not the case. In fact, both ensembles of prefrontal neurons significantly increased their activity when mice were engaged in social interaction (**Figure 3F**; social ensemble: mean activity level 1.4 +/- 0.3% in home cage vs. 1.8 +/- 0.3% during social interaction, $p < 0.05$, sign-rank test; home cage ensemble: mean activity level 1.5 +/- 0.3% in home cage vs. 1.9 +/- 0.3% during social interaction, $p < 0.001$, sign-rank test).. Next, we examined pairwise correlations between the activity of prefrontal neurons within each ensemble. Strikingly, mean correlations within the social ensemble increased during social interaction (**Figure 3G**; (mean correlation coefficient between neurons in the social ensemble: 0.009 +/- 0.002 in home cage vs. 0.012 +/- 0.002 during social interaction, $p < 0.05$). By

contrast, there was a non-significant decrease in correlations within the home cage ensemble (Figure 3G; home cage ensemble mean correlation coefficient 0.011 +/- 0.02 in home cage vs. 0.005 +/- 0.003 during social interaction, p=0.99, sign-rank).

Thus, the ensemble of prefrontal neurons which provide input to the social units actually form an assembly that collectively becomes more co-active during social behavior. In contrast, prefrontal neurons in the ensemble which provides input to the home cage units increase their activity, but not their co-activity, during social behavior. This suggests that changes in correlations associated with behavior may contribute to the encoding of social behavior.

Correlations enhance classifier performance

How can we quantitatively assess the contribution of these correlations, which are behaviorally-modulated, to classifier performance? Ideally we would first train a neural network on the original data. Then we would test this network's ability to classify data which maintained behaviorally-driven changes in activity levels, but either removed or preserved the correlations. Indeed, we have already developed methods for shuffling that achieve these goals. First, to shuffle the data in a manner that maintains behaviorally-driven changes in activity levels, but disrupts correlations, we can swap shuffle activity, but do so within each behavioral condition rather than across the entire testing dataset. In other words, we first divide up the raster into separate subrasters for each 5 minute behavior epoch (when the mouse was either engaged in social interaction or alone in its home cage). Then we performed swap shuffling (as described above) separately on each subraster, before recombining these swap shuffled subrasters to create the swap shuffled surrogate dataset for testing. Because swap shuffling tends to preserve activity levels, and because we swap shuffled activity within a behavioral condition, neurons that increase or decrease activity during periods of social interaction in the original dataset also tend to do so in the swap shuffled surrogate dataset.

To create surrogate datasets which preserve patterns of correlations as well as behaviorally-driven changes in activity, we used a method that we published previously: SHuffling Activity to Preserve Correlations, or SHARC (*18*). SHARC also re-assigns blocks of activity between neurons, but rather than doing so randomly, it instead follows an algorithm that achieves a target correlation matrix (in this case, the original correlation matrix) (**Figure 4B-C**). The full details of SHARC are presented in the Methods. Briefly: on each iteration, we randomly select one block of activity to be assigned to a new neuron. Instead of choosing the new neuron randomly, we first compute the difference between the target correlation matrix and the correlation matrix of the partially reconstructed surrogate dataset. Then we assign the block of activity to the neuron which will optimally reduce this difference. Finally, to maintain the mean activity level of each neuron, there is also an absolute limit on how many blocks of activity can be re-assigned to each neuron. We SHARC-shuffled each social or home cage subraster separately, then combined them to create a SHARC-shuffled surrogate dataset that preserves behaviorally-specific levels of activity and patterns of correlations.

We verified that both swap and SHARC shuffled surrogate datasets preserved levels of activity observed during both social interaction and periods when mice were alone in their home cages. Specifically, we computed the correlation between vectors in which each element represents the activity level of one neuron during one behavioral condition, and quantified the correlation

between each real and surrogate dataset. For swap shuffled surrogate datasets, the similarity of activity levels (compared to real data) was 0.89 +/- 0.02 in the home cage and 0.82 +/- 0.04 during social interaction. For SHARC shuffled surrogate datasets, the similarity of activity levels (compared to real data) was 0.88 +/- 0.03 in the home cage and 0.86 +/- 0.03 during social interaction (n = 10 mice/datasets). We also computed the similarity of the pattern of correlations between each surrogate dataset and the corresponding real dataset. In this case, only SHARC shuffled surrogate datasets preserved patterns of correlations. For swap shuffled surrogate datasets, the similarity of correlations to the real data was 0.01 +/- 0.01 in the home cage, and 0.03 +/- 0.01 during social interaction. For SHARC shuffled surrogate datasets, the similarity was 0.50 +/- 0.05 in home cage and 0.55 +/- 0.03 during social interaction.

We then trained classifiers on each dataset and tested each classifier using either swap or SHARC shuffled surrogate datasets generated from the same dataset using for training (**Figure 4C**). Classifiers performed better than chance when tested with either surrogate dataset indicating that changes in activity levels encode behavioral information. However, performance was significantly higher for SHARC shuffled surrogates datasets than for swap shuffled ones (**Figure 4D**; classifier accuracy for SHARC shuffled surrogate datasets = 68 +/- 4%, classifier accuracy for swap shuffled surrogate datasets = 61 +/- 4%, $p < 0.05$, sign-rank test). This demonstrates that behaviorally-modulated patterns of correlations transmit additional information, beyond what is readily decodable from activity levels alone.

Combinations of coactive neurons occur in a behaviorally-specific manner

The fact that neural networks perform classification better for connection probabilities ~0.2 – 0.4 than for connection probabilities < 0.1 indicates that the representations of social vs. home cage behavior are not linearly separable. (If the representations were linearly separable, then it should be possible to find a linear combination of single neuron activities which separate these behavioral conditions, i.e., a set of output weights associated with hidden units which each receive input from just one prefrontal neuron; this would correspond to a network that had a low connection probability and high classification accuracy). Together with the fact that classifier performance was higher for SHARC shuffled datasets than swap shuffled ones, this indicates that multineuron patterns of coactivity, rather than just levels of activity within neuronal ensembles, transmit information about social behavior. Therefore as a proof-of-concept, we directly examined whether 3-neuron patterns of coactivity occur in a behaviorally-specific manner.

First, we quantified how often each possible 3-neuron combination occurred in real datasets. Then we calculated how often each of these combinations in datasets that had been swap-shuffled (across the entirety of the dataset). For each real dataset we constructed 1000 swap-shuffled datasets, and identified 'enriched combinations,' which occurred more often in real datasets than in 95% of swap shuffled surrogate datasets. Enriched combinations are those which occur more often in real datasets than expected based on the chance overlap of activity between marginally independent neurons. Finally, we quantified how many of these enriched combinations were behaviorally-specific, i.e., occurred exclusively during social or home cage epochs. Combinations could appear to be behaviorally-specific simply because they only occurred at a single timepoint. Therefore we also restricted our analysis to enriched combinations which occurred during multiple distinct bouts of social interaction and/or matched sets of intervals during home cage epochs. Many of these repetitively-occurring enriched combinations

were behaviorally-specific: 43.5% occurred during social interaction, 26.5% during home cage epochs, and 30% during both conditions.

The selective occurrence of enriched combinations either during social interaction or when a mouse is alone in its home cage may reflect changes in single neuron activity (i.e., neurons that form a social combination are only active during the social condition), and/or changes in correlations (i.e., neurons are active in both conditions but only *co-active* during social behavior). To test the hypothesis that changes in correlations underlie the behavioral specificity of significant combinations, we examined the 3-neuron combinations that were specifically enriched during either periods of home cage exploration or social interaction (**Figure 5**). We defined specific enrichment as those combinations which occurred more often in real data than in 95% of swap-shuffled surrogate datasets for one behavioral context, and less in real data than in 50% of swap-shuffled surrogate datasets for the other behavioral context. Based on these criteria, 12,408 3-neuron combinations were specifically enriched during social interaction, and 9,572 were specifically enriched during home cage exploration. There were 55,696 instances in which a social and nonsocial 3-neuron combination overlapped in 2 out of 3 neurons. In 97.0% of these cases, the neuron which was part of a social 3-neuron combination (triplet) but left out of the overlapping home cage triplet was part of a different 3-neuron combination that was enriched during homecage exploration (Figure 5, top right). Conversely, the neuron which was part of a nonsocial triplet but left out of the overlapping social 3-neuron combination was part of a different socially-enriched 3-neuron combination in 99.1% of cases (Figure 5, bottom right). Overall, an average of 71 enriched homecage combinations contained the neuron missing from the social triplet, and 85 enriched social combinations contained the neuron missing from homecage triplets. Thus, the specificity of a combination of co-active neurons for social vs. nonsocial behavior does not occur simply because some neurons were only active during one condition, but rather reflects the dynamic reorganization of patterns formed by neurons which are active in both conditions, i.e., changes in correlations. This – the behaviorally-specific occurrence of multineuron patterns of coactivity – represents the substrate through which correlations can add to the behavioral information transmitted by neuronal ensembles.

Socially-enriched combinations are deficient in Shank3 KO mice

We were curious whether there might be conditions under which these phenomena – the occurrence of multineuron combinations of coactivity during social behavior, and the ability of correlations to enhance the transmission of information about social behavior – might be impaired. To explore this, we performed microendoscopic GCaMP imaging in mice lacking the autism-associated gene Shank3 (*19–21*). These mice have been extensively studied as models of Phelan-McDermid syndrome, which often includes autism as a clinical feature. Shank3$^{-/-}$ (KO) mice are known to have social deficits, and indeed, we found that compared to wild-type (WT) littermates, they spend significantly less time interacting with novel juvenile mice (**Figure 6A**).

We compared patterns of prefrontal activity in Shank3 KO mice and their WT littermates. As in WT mice, in Shank3 KO mice, many prefrontal neurons either increase or decrease activity during social interaction. However, compared to WT mice, the fraction of neurons whose activity increases during social interaction was significantly higher, whereas the fraction whose activity decreases was significantly lower (**Figure 6B-C;** 22% of 260 WT neurons vs. 39% of 290 KO neurons increased activity above the 90th percentile of shuffled data during social interaction, chi

squared = 17.7, $p < 0.0001$; 25% of WT vs. 15% of KO neurons decreased activity below the 10^{th} percentile of shuffled data during social interaction, chi squared = 8.2, $p < 0.0001$). Thus, Shank3 KO mice recruit abnormal neuronal ensembles during social behavior. We hypothesized that this might reflect a network-level disorganization that affects the normal patterning of activity during social behavior.

Indeed, we found that in KO mice a significantly smaller fraction of the 3-neuron combinations observed during social interaction were strongly enriched, i.e., occur more often in actual data than in 99.9% of swap-shuffled surrogate datasets (**Figure 6D**). This suggests that even though more neurons (i.e., larger ensembles), were recruited during social behavior in KO mice, these may have been less well-organized, such that the occurrence of socially-enriched patterns of activity is obscured by 'noise,' i.e., patterns formed by the chance overlap of activity between neurons that fire in a largely independent fashion. Notably, this deficiency was specific for social interaction. The fraction of 3-neuron combinations that were strongly enriched during home cage exploration (in comparison to swap-shuffled surrogate datasets) was similar in WT and KO mice (Figure 6D).

Correlations do not enhance the transmission of information about social behavior in Shank3 KO mice

The preceding shows that even though social behavior robustly recruits neuronal ensembles in Shank3 KO mice, the organization of these ensembles into multineuron combinations is disorganized. This suggests that the ability of patterns of co-activity to encode information about social behavior may be impaired in these mice. To test this, we directly examined whether correlations contribute to the transmission of information about social behavior in Shank3 KO mice. As before, we generated swap and SHARC shuffled surrogate datasets, then tested the ability of classifiers trained on the original datasets (from Shank3 KO mice) to classify activity associated with behavior during social interaction vs. in home cage. While we still observed above chance classification accuracy using a classifier with a connection probability of 0.3, there was no longer an increase in performance when correlations were preserved in SHARC shuffled surrogate datasets as compared to swap shuffled ones (**Figure 6E**; classifier accuracy: 62 +/ 4% for SHARC vs 63 +/- 2% for swap shuffled surrogate datasets, $p = 0.47$, sign-rank test). Thus, in Shank3 KO mice, multineuron patterns of coactivity during social behavior are disturbed, and correlations no longer add to the information about social behavior transmitted by prefrontal ensembles.

DISCUSSION

During complex behaviors, the brain can use many strategies to represent information about the external environment and internal state of the organism. The term 'ensemble' is often used to refer to a group of neurons whose activity is similarly modulated (either increased or decreased) during specific behaviors (*1*, *22–26*). It is generally accepted that ensembles transmit behavioral information via changes in the activity levels of their constituent neurons. On the other hand, many studies have also shown that correlations between neurons can change during specific behaviors (*3*, *9*) or behavioral states (*27–29*). Importantly, correlations reflect changes in coactivity which exceed those expected to occur simply because of changes in the activity levels of the individual neurons (*10*). I.e., when an ensemble becomes more active, its correlations could go up, down, or remain unchanged. By optimizing synaptic interactions such as temporal summation, changes in correlated activity could potentially enhance the behavioral information transmitted by changes in ensemble activity, or transmit entirely different types of information, e.g., about internal states. Correlations have been studied extensively for the isolated retina responding to visual stimuli (*30*). However, how correlations in recurrently connected cortical circuits such as the mPFC encode behavior has been more difficult to discern.

Here, we addressed this question using microendoscopic GCaMP imaging to measure activity from many (~80-100) prefrontal neurons during social behavior in mice. We used multiple approaches to disentangle the respective contributions of activity and correlations to the encoding of behavior. First, we used a simple neural network, in which prefrontal neurons provide input, there is one hidden layer, and a single output unit classifies social vs. nonsocial behavior, to quantify how well prefrontal ensembles would transmit information about social behavior to downstream neurons. We extended a method we previously published, (*18*), to non-randomly shuffle datasets in order to preserve both behaviorally-modulated correlations and ensemble activity. This enabled us to compare the amount of information about social behavior transmitted by either SHARC-shuffled surrogate datasets or randomly-shuffled surrogates which preserved ensemble activity but not correlations. In this way, we found that correlations enhance the amount of information that prefrontal ensembles transmit about social behavior. Indeed, when we examined connections within neural network classifiers, we found that prefrontal neurons which serve to detect social behavior increase their correlations during social behavior (whereas neurons which detect nonsocial behavior do not).

Correlations measure neuronal coactivity that occurs more often than expected based on the chance overlap of activity between neurons. Thus, in accordance with our finding that behavior modulates correlations, we found that multineuron patterns of coactivity which occur more often than expected by chance are behaviorally-specific. We then directly examined these behaviorally-specific and statistically-enriched combinations of coactive neurons. We found that they tend to be composed of neurons which are active in both conditions but only coactive in one, rather that neurons which are only active in one condition.

Interestingly, these statistically-enriched patterns of coactivity were specifically deficient during social behavior in mice lacking the autism-associated gene Shank3. Accordingly, in Shank3 KO mice, surrogate datasets which preserve behaviorally-modulated correlations failed to transmit more information about social behavior compared to randomly shuffled datasets which only

preserved ensemble activity. This shows that the ability of correlations to enhance the transmission of information about social behavior is not automatic, and can in fact be disrupted in pathological states.

What is the meaningful size of ensembles in the cortex?

Complex behavior is possible because the brain reliably encodes features pertaining to the external environment as well as the internal state of the organism. These features may be encoded by the modulation of activity in neuronal ensembles (*1*, *22–26*). How many neurons are needed to reliably encode an aspect of behavior? This is an important question because the capacity, robustness against noise, generalization ability, etc., of a network depend on how many neurons encode specific pieces of information.

We explored this question, not by measuring actual connections, but rather by asking what connection probability would optimize the ability of a downstream network to classify behavior based on input from prefrontal ensembles. Peak classifier performance occurred for connection probabilities ~0.2 - 0.3. Performance was markedly lower when the connection probability was 0.5. This is surprising because a connection probability of 0.5 would maximize the entropy of each connection; correspondingly, the number of distinct input combinations to a hidden unit is maximized when it receives connections input from half the input neurons. Thus, from the standpoint of encoding social behavior, combining activity from 20-30% of the input neurons must achieve some synergy that becomes degraded by including activity from additional neurons. This suggests that whatever mechanism normally generates behaviorally-meaningful patterns of coactivity in prefrontal neurons, the size of these patterns is limited to about 20-30% of the network. This may reflect nonrandom network connectivity (*31*, *32*) which produces correlated activity / coactivity within defined neuronal subgroups (*33*, *34*).

Combinatorial codes vs. sequential patterns of activity

Like many recent studies, we measured population-level activity in the mouse neocortex using genetically encoded calcium indicators. These indicators transduce neuronal spiking on timescales ~100 msec. Thus correlated activity / 'coactivity' imply that neurons jointly increase their activity within windows ~100 msec, and do not necessarily imply synchronous spiking on faster timescales (milliseconds or even tens of miliseconds). At the same time, correlated activity / coactivity on these timescales should be differentiated from sequential activity of neurons observed during the performance of sequential behaviors (i.e. spatial navigation or overtrained tasks) in which the activity of specific neurons corresponds to moving through a specific location or performing a specific portion of a complex task. As discussed above, in the neocortex correlations and coactivity likely reflect recurrent neural network connectivity (*33*). By contrast, sequential patterns of neuronal activation can occur simply as a byproduct of the arrangement of spatial locations along a trajectory, the stereotyped order in which cues are encountered during a task, etc.

Relevance to disease states

Interestingly, in *Shank3* KO mice, which exhibit social deficits, the mPFC successfully recruits specific neuronal ensembles during social interaction. However the organization of these ensembles into statistically-enriched patterns of coactivity is disrupted, and correlations fail to enhance the transmission of information by these ensembles. Thus, the computational units by

476 which information is processed in the mPFC appears to be inefficient, i.e., social behavioral
477 recruits an abnormally large number of neurons at the expense of the precise temporal patterning
478 of this activity. This central finding is similar to other findings in rodent models of autism at both
479 the single neuron and network levels (*15*, *19*, *35*). In particular, we found an increase in the
480 recruitment of prefrontal neurons during social interaction. This mirrors a recent study which
481 found hyperdynamic response to whisker stimulation in the same mice (*19*), possibly reflecting
482 GABAergic circuit dysfunction and/or homeostatic compensations (*36*). Here, we show how
483 such exaggerated responses and enlarged neural ensembles may disrupt information transmission
484 by degrading the ratio between signal (statistically meaningful patterns of coactivity) and noise
485 (the random overlap of activity between neurons).
486

MATERIALS AND METHODS

Behavior: C57/B6 mice were obtained from Jackson Laboratories. We utilized adult mice of either sex housed and bred in the UCSF animal facility. Adult mice were habituated to the room and observer for 3 days prior to test day. All videos were subsequently scored by a blinded observer. For imaging experiments, 5 WT and 4 KO littermates were generated through crosses between Shank3 heterozygous parents and injected with AAV5.Syn.GCaMP6f.WPRE.SV40. We included an additional 5 WT mice which were injected with AAV5.Syn.GCaMP6m.WPRE.SV40 (*37*). Viruses were obtained from Penn Viral Core. Injections and 500 um GRIN lens (Inscopix) implantations were carried out in 8-12 week old mice to express GCamp6f in prefrontal cortical neurons under control of the human Synapsin promotor. Mice were anesthetized with 2% isoflurane and mounted in a stereotactic frame. Craniotomies were made according to stereotaxic coordinates relative to Bregma. Coordinates for injection into mPFC were (in mm relative to Bregma): +1.7 anterior–posterior (AP), –0.3 mediolateral (ML) and –2.75 dorsoventral (DV), and GRIN lenses were implanted at the same AP and ML coordinates, to a depth of 2.25. We subsequently attached baseplates for attaching the microendoscope, ~4 weeks later depending on GCamp expression. Mice were habituated for three days with the scope attached, prior to test day. On test day, mice were habituated with the scope turned on, then imaged in alternating home cage and social epochs. During social epochs, one of 2 novel sex-matched juvenile mouse was introduced to the test mouse's homecage, in sequential order so that there were two 'novel' epochs, followed by two 'familiar' epochs interleaved with 'home cage' epochs during which the juvenile mice were removed and the test mouse was free to explore its home cage. The first and last home cage epoch were 10 minutes in length; the others were 5 minutes in length. Each social epoch lasted 10 minutes but only the first 5 minutes were recorded and scored. During each behavioral epoch, observer was not in the room. Interaction epochs were defined from the moment test mouse first sniffed the juvenile conspecific or object, until the test mouse turned away. Videos were recorded using Anymaze, and scored by a blinded observer. For the bulk of analysis we pooled data across 10 WT mice. Shank3 KO mice were compared only to recordings from WT littermates.

Image acquisition and segmentation: Images were acquired using an Inscopix nVoke microendoscope attached to a laptop computer and synced to a separate video acquisition computer running Anymaze. Frame rate was 20 Hz and the laser power was 0.2 mW. Acquisition was performed using 2x2 pixel binning, then subsequently downsampled again by 2.

We segmented neuronal signals using a modified PCA/ICA approach(*16*, *17*), modified so that each segment was expressed as a binary ROI consisting of pixels representing a single neuron. To deconvolve neuronal signals from background neuropil signals, we subtracted the mean signal from each identified segment from the mean value in pixels surrounding the edge of the segment (we excluded pixels that belonged to another ROI). Signals were subsequently lowpass filtered to remove high frequency noise using the Matlab command: designfilt('lowpassfir', 'PassbandFrequency', 0.5, 'StopbandFrequency', .65, 'PassbandRipple', 1, 'StopbandAttenuation', 25). All signal traces shown represent normalized versions of the dF/F_0 trace, where F_0 *is estimated by the median* value in the surround region. Threshold based event detection was performed on the traces by detecting increases in dF/F_0 exceeding 3σ over one second, then only keeping those events which exceeded a 15σ increase over two seconds, and a total area under the

curve of 250σ. As there were occasional downward deflections due to surround subtraction, we instituted a final parameter requiring that the peak cross an absolute value of $dF/F_0 = 0.0125$. σ is the standard deviation of dF/F_0, calculated over the least active 50% of the movie. In some cases these parameters were adjusted slightly to optimize event detection to > 95% sensitivity and specificity, based on visual inspection, for each movie. After identifying these events in the GCaMP signal from a cell, the cell was considered "active" during the entire period from the beginning of an event until the GCaMP signal decreased 30% from the peak of the event, up to a maximum of 2 seconds. The peak of the event was defined as the local maximum of the entire event, from the beginning of the event until dF/F_0 returned to the pre-event baseline value. Calcium traces from segmented neurons were visually inspected and a small number of segments were removed if they did not appear to represent a single, unduplicated neuron. We restricted further analysis to those mice with 25 or more active neurons. We then created 2-dimensional event rasters consisting of detected events for each neuron over the course of the experiment.

Detection of behaviorally modulated neurons: To determine the response of individual neurons to behavioral context, we averaged the activity of each neuron during frames corresponding to periods of social interaction, or to a temporally matched set of frames during the preceding home cage epoch. We then created a 'null distribution' for each neuron that represents the percent of time active expected in each condition based on chance, by circularly shuffling the data 10,000 times. We then compared the activity of each neuron during either social interaction or home cage exploration to this null distribution. Neurons were considered positively modulated if they exceeded the 90th percentile of that observed in circularly-shuffled datasets, and negatively modulated if the percent of frames that a neuron was active during a given context was below the 10th percentile of observations from circularly-shuffled data.

SHARC: SHARC (SHuffling Activity to Rearrange Correlations) is an iterative method for generating surrogate datasets. SHARC nonrandomly shuffles blocks of activity within a raster to generate a new (surrogate) raster in which the pairwise correlations between neurons match a target correlation matrix (*17*). Here we apply this previously-published method, with modifications to also preserve the activity level in each neuron (Figure 4B).

To begin, note that each raster is equivalent to a collection of blocks of activity. Each block of activity is defined by the time at which it begins, its duration, and the neuron which is active. On each iteration one block of activity is randomly chosen and assigned to a new neuron as follows. Suppose block i has been chosen to be re-assigned. First, we find all the blocks of activity that overlap with block i. Next, we selected the subset of these blocks for which new cell identities had already been assigned. Call this set X. Let r_j represent the number of timepoints over which block $j \in X$ overlaps with block i, and let n_j represent the identity of the cell assigned to block $j \in X$. L_i and L_j are the lengths of blocks i and j, respectively. Then we constructed a vector,

$$\vec{P}_i = \sum_{j \in X} \frac{r_j}{\sqrt{L_i L_j}} \left(\vec{C}_{n_j} - \vec{C'}_{n_j} \right)$$

where $\vec{C}_{n_j}$ represents row j of the target correlation matrix, i.e. the target correlations between neuron n_j and the other neurons, and $\overrightarrow{C'}_{n_j}$ contains the current values of the correlations between neuron n_j and the other neurons based on the blocks of activity that have already been re-assigned. This step can be thought of as "guessing" which cell *should* be assigned to a particular

block of activity by first figuring out what other cells are active at the same time, then choosing cells which are strongly correlated with these known active cells. Note that we assign values of $\vec{P}_i$ (i.e., construct "guesses" about which cell should be active), using the *difference* between the current correlation matrix ($\vec{C}'_{n_j}$) and the target correlation matrix ($\vec{C}_{n_j}$), in order to identify cell pairs for which the current correlation deviates from the target value, and force the new correlation matrix to progressively approximately the target correlation matrix.

We set elements of $\vec{P}_i$ to zero if the corresponding neuron had already been assigned to a block of activity that overlaps with block *i*, i.e. element n_j of $\vec{P}_i$ was set to zero $\forall\, j \in X$. Finally, we assigned block *i* to the neuron corresponding to the maximum value of $\vec{P}_i$. This can be thought of as choosing the cell that represents the "consensus" based on tallying up all of the "guesses" about which cells "should" be assigned to the block of activity being considered.

When all the elements of $\vec{P}_i$ were zero, e.g. because there no overlapping blocks of activity have had new cell identities assigned yet, then we chose a cell in order to match the originally observed level of activity. Specifically, after every iteration, we kept a log of the net number of blocks of activity that each neuron had donated or received. We used this vector to create a weighted probability whereby events from neurons which had received a net positive number of blocks were more likely to be chosen to be reassigned. To further ensure that the total number of active events for each neuron in the surrogate dataset was similar to the real dataset, if the difference between the number of blocks gained – lost in the reassignment process exceeded +10 for a particular neuron, then that neurons was no longer eligible to receive additional blocks of activity.

We extended this approach to generate surrogate datasets by shuffling data within shorter time windows (i.e., individual behavioral epochs). Here a discrete set of frames is chosen, corresponding to a subraster of the original raster. By repeating the process described above for each subrasters, then recombining the shuffled subrasters, we generate a complete shuffled dataset.

Classifier: We designed and trained a neural network to classify behavior (periods when a mouse was alone in its home cage vs. engaged in social interaction). This network contained 1000 units in a hidden layer, each of which received input from specific prefrontal neurons (from the real dataset). Thus, in each frame the activity of each hidden layer unit was just the summed activity of the connected prefrontal neurons. Each hidden layer unit had an output weight that represented the strength of its connection to a single output unit. On each frame the activity of the output unit was computed as:

$$y = \frac{1}{1 + e^{-\sum w_i x_i}}$$

where w_i is the output weight from hidden unit *i* and x_i is the activity of hidden unit *i*.

When we performed training and testing using the same dataset, we divided the dataset into alternating blocks of 500 frames for training vs. testing (in other cases we used the real dataset for training, then tested using a surrogate dataset). We restricted training or testing to frames in

which mice were scored as actively engaged in social interaction (or matched frames during periods when the mouse was alone in its home cage). We also limited training / testing to frames with at least 3 active neurons.

We trained the output weights by performing 500 passes through the training data (each pass visited all of the training frames in a random order). On each training timestep, we calculated y, the activity of the output unit, and then adjusted each output weight based on:

$$\Delta w_i = \varepsilon y(1-y)(y-z)x_i$$

where z is the correct classification of the frame (0 for social behavior, 1 for home cage) and ε, the learning rate, was set to 0.05.

Following training, we examined the pattern of input connections and output weights. The distribution of output weights was roughly gaussian and centered near 0. We identified the selection of prefrontal neurons most likely to be connected to hidden layer units with large positive or negative weights. Hidden layer units with large negative or positive output weights bias classification towards the social or home cage condition, respectively. Therefore, we refer to the 25 hidden units with the most negative or positive weights as 'social' or 'home cage' units respectively. We calculated the number of input connections between each prefrontal neuron and the 25 home cage units or 25 social units. We then defined 'home cage' or 'social' ensembles as the 20% of prefrontal neurons with the most input connections to home cage or social units, respectively. As described in the main text, we then analyzed properties of these two ensembles.

Quantification of multineuron combinations: Estimating chance overlap between activity of largely independent neurons requires accounting for two factors. First, neurons with higher activity are more likely to overlap by chance with other neurons. Second, overall network activity is dynamic over time, creating a tendency for otherwise independent neurons to be recruited at similar times. Thus, it is necessary to identify combinations which occur more often than expected based on 1) the activity levels of the constituent neurons, and 2) the fact that activity in a network is not constant over time. We can do this by quantifying the occurrence of combinations in datasets which have been shuffled to preserve 1) the overall level of activity in each neuron, and 2) the total level of activity in the network at each point in time.

3 neuron combinations were quantified by identifying each combination present in frames in which 2 or more neurons were active. The number of frames each combination was active in *real data* was stored in a n-dimensional matrix. Surrogate datasets were then generated from event rasters by *swapping* the identity of neurons associated with detected events (periods of activity). As the timing of events themselves is unchanged, and only the identity of the participating neurons are exchanged, this preserves both the number of events per frame and the number of events that each neuron participates in. Therefore, the total number of combinations in each frame and over the course of the experiment (i.e., the sum of occurrences across all combinations) is also preserved. The total number of combination occurrences in which a given neuron participates would also tend to be preserved in these swap-shuffled surrogate datasets.

We then quantified how often each combination occurred in real vs. swap-shuffled data. By comparing how often each combination occurred in real data vs. in 1,000 swap-shuffled surrogate dataset, we were able to quantify how 'enriched' each combination was, compared to the level of occurrence expected by chance based on the activity levels of its constituent neurons (and the overall temporal pattern of network activity). We expressed enrichment as a percentile, calculated relative to swap-shuffled surrogate data, e.g., the 100^{th} percentile indicates that a particular combination occurred more often in real data than in all 1,000 surrogate datasets. Further analysis was restricted to 'enriched combinations', i.e., combinations that occurred more often in real datasets than in 95% of surrogate datasets.

Statistical analysis: Neurons and significant combinations from all animals and groups were pooled and counted as single units. Proportions were compared using chi-squared test. Activity levels were compared using paired *t*-tests (2-sided), unless otherwise noted. Where applicable, error bars denote standard error.

REFERENCES

1. D. J. Cai, D. Aharoni, T. Shuman, J. Shobe, J. Biane, W. Song, B. Wei, M. Veshkini, M. La-Vu, J. Lou, S. E. Flores, I. Kim, Y. Sano, M. Zhou, K. Baumgaertel, A. Lavi, M. Kamata, M. Tuszynski, M. Mayford, P. Golshani, A. J. Silva, A shared neural ensemble links distinct contextual memories encoded close in time. *Nature* (2016), doi:10.1038/nature17955.

2. B. Liang, L. Zhang, G. Barbera, W. Fang, J. Zhang, X. Chen, R. Chen, Y. Li, D.-T. T. Lin, Distinct and Dynamic ON and OFF Neural Ensembles in the Prefrontal Cortex Code Social Exploration. *Neuron*. **100**, 700-714.e9 (2018).

3. R. C. deCharms, M. M. Merzenich, Primary cortical representation of sounds by the coordination of action-potential timing. *Nature*. **381**, 610–613 (1996).

4. D. O. Hebb, *The organization of behavior: A neuropsychological theory* (Wiley, New York, 1949).

5. G. Buzsáki, Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron*. **68**, 362–385 (2010).

6. B. B. Averbeck, P. E. Latham, A. Pouget, Neural correlations, population coding and computation, doi:10.1038/nrn1888.

7. M. S. Ahmed, J. B. Priestley, A. Castro, F. Stefanini, A. S. Solis Canales, E. M. Balough, E. Lavoie, L. Mazzucato, S. Fusi, A. Losonczy, Hippocampal Network Reorganization Underlies the Formation of a Temporal Association Memory. *Neuron*. **107**, 283-291.e6 (2020).

8. F. Stefanini, L. Kushnir, J. C. Jimenez, J. H. Jennings, N. I. Woods, G. D. Stuber, M. A. Kheirbek, R. Hen, S. Fusi, A Distributed Neural Code in the Dentate Gyrus and in CA1. *Neuron* (2020), doi:10.1016/j.neuron.2020.05.022.

9. E. Vaadia, I. Haalman, M. Abeles, H. Bergman, Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature*. **373**, 515–518 (1995).

10. J. De La Rocha, B. Doiron, E. Shea-Brown, K. Josić, A. Reyes, K. Josi, A. Reyes, Correlation between neural spike trains increases with firing rate. *Nature* (2007), doi:10.1038/nature06028.

11. O. Yizhar, L. E. Fenno, M. Prigge, F. Schneider, T. J. Davidson, D. J. O'Shea, V. S. Sohal, I. Goshen, J. Finkelstein, J. T. Paz, K. Stehfest, R. Fudim, C. Ramakrishnan, J. R. Huguenard, P. Hegemann, K. Deisseroth, Neocortical excitation/inhibition balance in information processing and social dysfunction. *Nature* (2011), doi:10.1038/nature10360.

12. A. Selimbeyoglu, C. K. Kim, M. Inoue, S. Y. Lee, A. S. O. Hong, I. Kauvar, C. Ramakrishnan, L. E. Fenno, T. J. Davidson, M. Wright, K. Deisseroth, Modulation of prefrontal cortex excitation/inhibition balance rescues social behavior in CNTNAP2-deficient mice. *Sci. Transl. Med.* **9**, eaah6733 (2017).

13. A. C. Brumback, I. T. Ellwood, C. Kjaerby, J. Iafrati, S. Robinson, A. T. Lee, T. Patel, S. Nagaraj, F. Davatolhagh, V. S. Sohal, Identifying specific prefrontal neurons that contribute to autism-associated abnormalities in physiology and social behavior. *Mol.*

Psychiatry (2017), doi:10.1038/mp.2017.213.

14. M. Murugan, H. J. Jang, M. Park, E. M. Miller, J. Cox, J. P. Taliaferro, N. F. Parker, V. Bhave, H. Hur, Y. Liang, A. R. Nectow, J. W. Pillow, I. B. Witten, Combined Social and Spatial Coding in a Descending Projection from the Prefrontal Cortex. *Cell*. **171**, 1663-1677.e16 (2017).

15. D. R. Levy, T. Tamir, M. Kaufman, A. Parabucki, A. Weissbrod, E. Schneidman, O. Yizhar, Dynamics of social representation in the mouse prefrontal cortex. *Nat. Neurosci.* (2019), doi:10.1038/s41593-019-0531-z.

16. E. A. Mukamel, A. Nimmerjahn, M. J. Schnitzer, Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*. **63**, 747–60 (2009).

17. F. J. Luongo, M. E. Horn, V. S. Sohal, Putative microcircuit-level substrates for attention are disrupted in mouse models of autism. *Biol. Psychiatry* (2016), doi:10.1016/j.biopsych.2015.04.014.

18. F. J. Luongo, C. A. Zimmerman, M. E. Horn, V. S. Sohal, Correlations between prefrontal neurons form a small-world network that optimizes the generation of multineuron sequences of activity. *J. Neurophysiol.* **115**, 2359–2375 (2016).

19. Q. Chen, C. A. Deister, X. Gao, B. Guo, T. Lynn-Jones, N. Chen, M. F. Wells, R. Liu, M. J. Goard, J. Dimidschstein, S. Feng, Y. Shi, W. Liao, Z. Lu, G. Fishell, C. I. Moore, G. Feng, Dysfunction of cortical GABAergic neurons leads to sensory hyper-reactivity in a Shank3 mouse model of ASD. *Nat. Neurosci.*, 1–13 (2020).

20. J. Peça, C. Feliciano, J. T. Ting, W. Wang, M. F. Wells, T. N. Venkatraman, C. D. Lascola, Z. Fu, G. Feng, Shank3 mutant mice display autistic-like behaviours and striatal dysfunction HHS Public Access. *Nature*. **472**, 437–442 (2011).

21. L. J. Duffney, P. Zhong, J. Wei, E. Matas, J. Cheng, L. Qin, K. Ma, D. M. Dietz, Y. Kajiwara, J. D. Buxbaum, Z. Yan, Autism-like Deficits in Shank3-Deficient Mice Are Rescued by Targeting Actin Regulators. *Cell Rep.* (2015), doi:10.1016/j.celrep.2015.04.064.

22. G. Corder, B. Ahanonu, B. F. Grewe, D. Wang, M. J. Schnitzer, G. Scherrer, An amygdalar neural ensemble that encodes the unpleasantness of pain. *Science (80-.).* **363**, 276–281 (2019).

23. J. Gründemann, Y. Bitterman, T. Lu, S. Krabbe, B. F. Grewe, M. J. Schnitzer, A. Lüthi, Amygdala ensembles encode behavioral states. *Science*. **364**, eaav8736 (2019).

24. K. Ghandour, N. Ohkawa, C. C. A. Fung, H. Asai, Y. Saitoh, T. Takekawa, R. Okubo-Suzuki, S. Soya, H. Nishizono, M. Matsuo, M. Osanai, M. Sato, M. Ohkura, J. Nakai, Y. Hayashi, T. Sakurai, T. Kitamura, T. Fukai, K. Inokuchi, Orchestrated ensemble activities constitute a hippocampal memory engram. *Nat. Commun.* **10**, 2637 (2019).

25. Y. Sakurai, How do cell assemblies encode information in the brain? *Neurosci. Biobehav. Rev.* (1999), , doi:10.1016/S0149-7634(99)00017-2.

26. Y. Sakurai, Y. Osako, Y. Tanisumi, E. Ishihara, J. Hirokawa, H. Manabe, Multiple approaches to the investigation of cell assembly in memory research—present and future. *Front. Syst. Neurosci.* **12** (2018), , doi:10.3389/fnsys.2018.00021.

27. L. Pinto, M. J. Goard, D. Estandian, M. Xu, A. C. Kwan, S.-H. Lee, T. C. Harrison, G. Feng, Y. Dan, Fast modulation of visual perception by basal forebrain cholinergic neurons. **16** (2013), doi:10.1038/nn.3552.

28. M. C. Dadarlat, M. P. Stryker, Locomotion enhances neural encoding of visual stimuli in mouse V1. *J. Neurosci.* **37**, 3764–3775 (2017).

29. M. Abeles, H. Bergman, I. Gat, I. Meilijson, E. Seidemann, N. Tishby, E. Vaadia, Cortical activity flips among quasi-stationary states. *Proc. Natl. Acad. Sci. U. S. A.* **92**, 8616–8620 (1995).

30. E. Schneidman, M. J. Berry, R. Segev, W. Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature*. **440**, 1007–1012 (2006).

31. X. Jiang, S. Shen, C. R. Cadwell, P. Berens, F. Sinz, A. S. Ecker, S. Patel, A. S. Tolias, Principles of connectivity among morphologically defined cell types in adult neocortex. *Science (80-.).* **350** (2015), doi:10.1126/science.aac9462.

32. A. P. B. Alex M. Thomson, David C. West, Yun Wang, Synaptic Connections and Small Circuits Involving Excitatory and Inhibitory Neurons in Layers 2–5 of Adult Rat and Cat Neocortex: Triple Intracellular Recordings and Biocytin Labelling In Vitro. *Cereb. Cortex*. **12**, 936–953 (2002).

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

34. A. Litwin-Kumar, B. Doiron, Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nat. Neurosci.* **15**, 1498–505 (2012).

35. J. P. Hamm, D. S. Peterka, J. A. Gogos, R. Yuste, Altered Cortical Ensembles in Mouse Models of Schizophrenia. *Neuron*. **94**, 153-167.e8 (2017).

36. S. B. Nelson, V. Valakh, Excitatory/Inhibitory Balance and Circuit Homeostasis in Autism Spectrum Disorders. *Neuron*. **87**, 684–698 (2015).

37. T.-W. Chen, T. J. Wardill, Y. Sun, S. R. Pulver, S. L. Renninger, A. Baohan, E. R. Schreiter, R. A. Kerr, M. B. Orger, V. Jayaraman, L. L. Looger, K. Svoboda, D. S. Kim, Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*. **499**, 295–300 (2013).

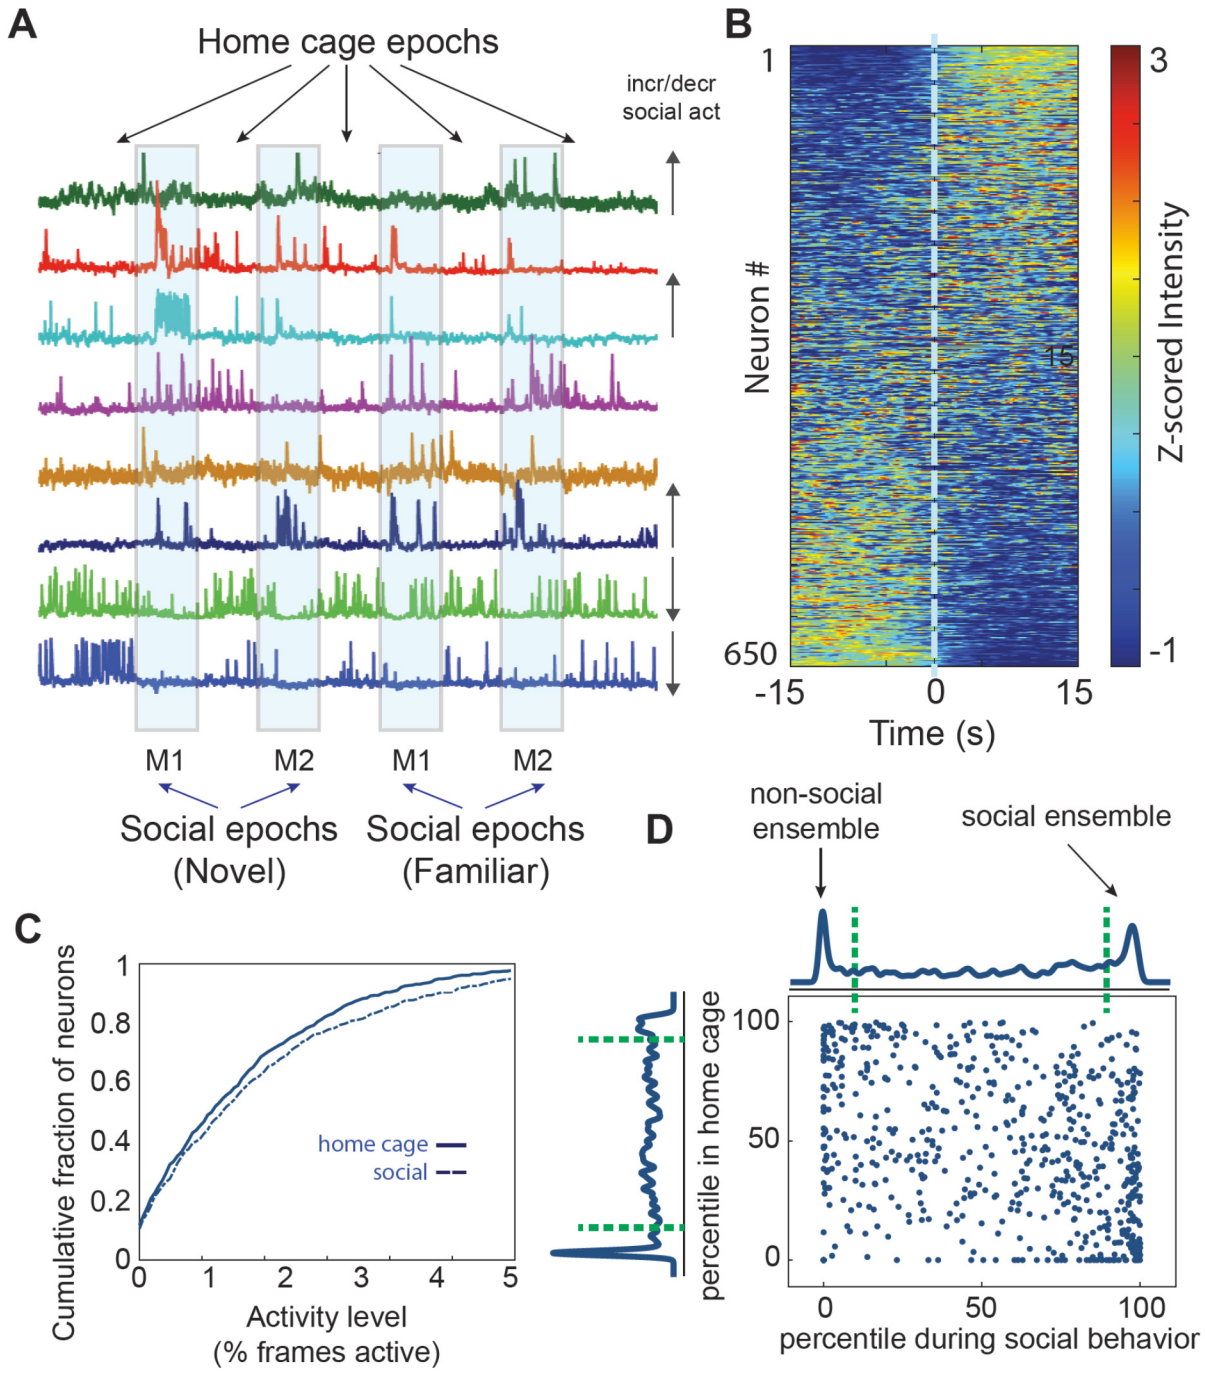

Figure 1. Social interaction modulates activity levels within prefrontal ensembles.

A. Mice were imaged across 9 consecutive behavioral epochs (each lasting 5 min) during which they were either alone in their homecage or interacted with one of two novel sex-matched juvenile mice introduced to the homecage ('M1' or '(M2'). Each novel mouse was subsequently re-introduced to the home cage during a familiar epoch. GCaMP traces during show examples of neurons that appear to increase or decrease activity during social epochs (see arrows at the right of each trace).

B. Mean z-scored GCaMP traces for all neurons recorded from wild-type mice (663 neurons from 10 mice) aligned to the onset of social interaction during the first bout of interaction within each social epoch.

C. Cumulative plot showing the distribution of activity levels for individual neurons during homecage epochs or periods of social interaction (percent time active in homecage: 1.8% +/- 0.1, percent time active during social interaction: 2.1 +/- 0.1%, $p = 0.00002$, paired t-test; n = 663 neurons from 10 WT mice).

D. Scatter-plot showing the activity of each neuron during each behavioral condition, expressed as a percentile relative to a null distribution generated by circularly shuffling that neuron's activity. Activity levels during social interaction or while the mouse was alone in its home cage are plotted on the x and y axis, respectively. Kernel density plots along the axes indicate the fraction of neurons whose activity was at a given percentile of the null distribution. Neurons with activity > 90^{th} percentile of shuffled datasets (green dotted line) were considered to be positively modulated, whereas neurons with activity < 10^{th} percentile (green dotted line) were considered to be negatively modulated during each behavior (>90^{th} percentile, social: 152/663 neurons, home cage: 80/663 neurons; $p < 0.00001$, chi-squared test; <10^{th} percentile, social: 128/663 neurons, home cage: 119/663 neurons; $p = 0.5$, chi-squared test).

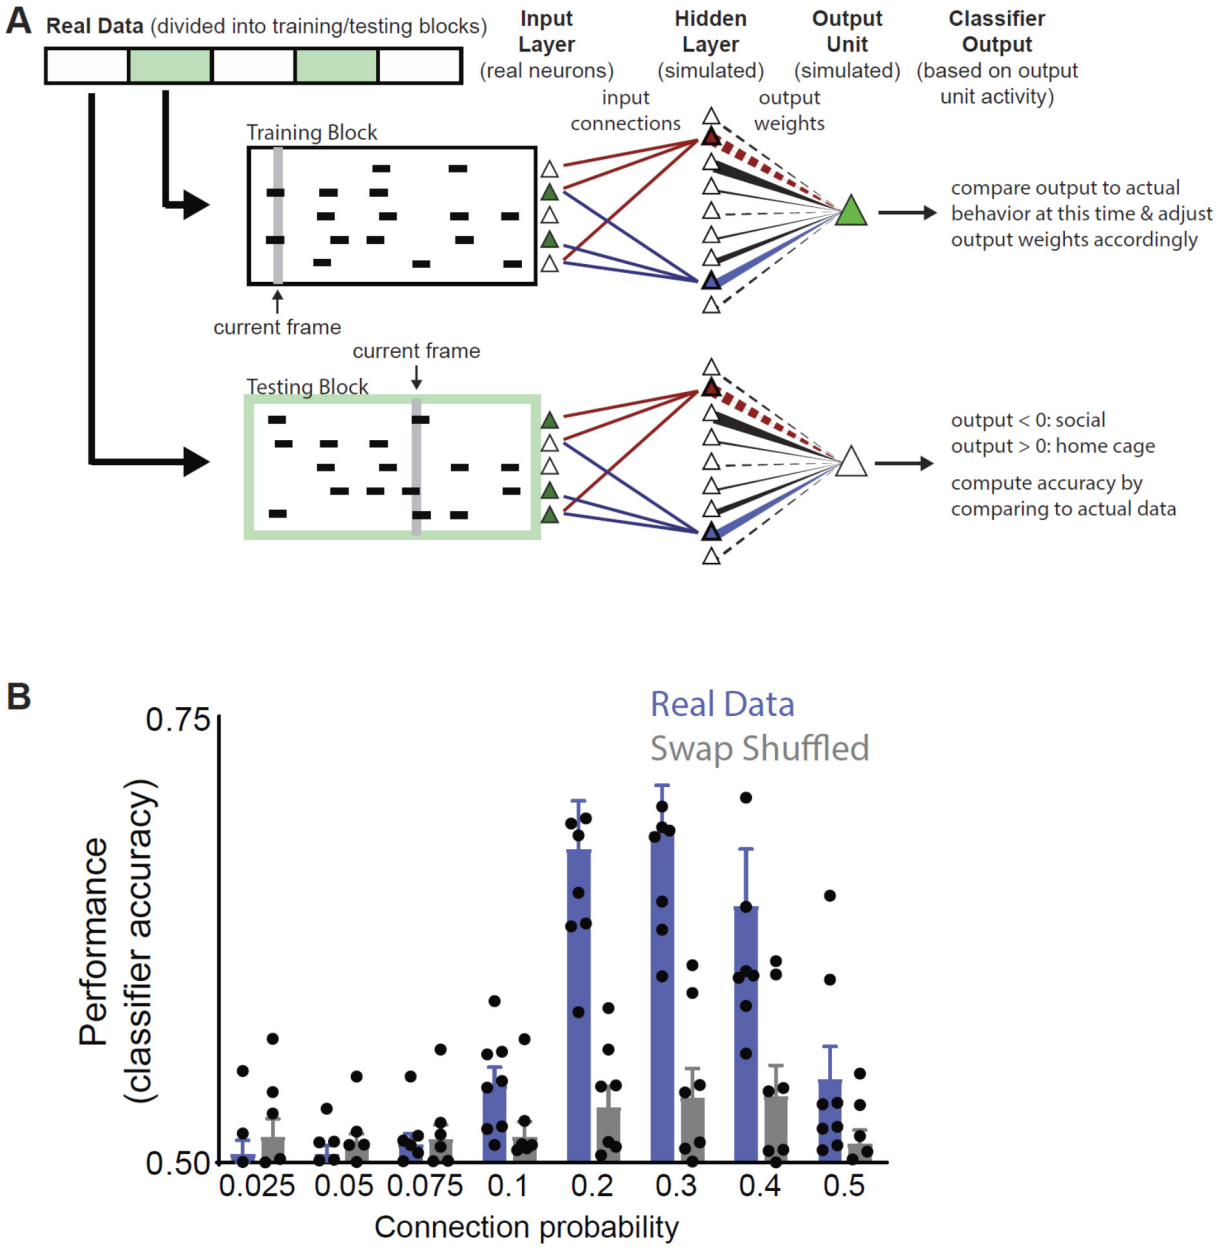

Figure 2. Classifying behavior from prefrontal ensembles using a simple neural network.

A. We constructed a neural network consisting of a single hidden layer (containing 1000 units) which were connected to a single output unit. The thickness of lines between each hidden layer unit and the output unit reflects the magnitude of the output weight. Positive and negative weights are indicated by solid and dashed lines, respectively. Each hidden layer unit received input from a random subset of prefrontal neurons from one real dataset. For clarity, we have only shown input connections to two hidden layer units (which are differentiated by their blue and red colors) – output weights from other hidden units are shown in black. The

output weight from each hidden layer neuron was iteratively updated during training. We trained the classifier to distinguish periods marked as home cage exploration or social interaction by dividing a dataset into 500-frame blocks, and then using alternating blocks for training or testing.

B. The classifier performed poorly (near chance) when the input connection probability (governing the number of prefrontal neurons that provided input to each hidden layer unit) was <10%. Classification accuracy was above chance in 8/10 mice and increased to a peak of 69 +/- 3% in these mice, before decreasing again for connection probabilities >30%. The classifier performed near chance levels when we trained and tested using data that had been randomly swap-shuffled.

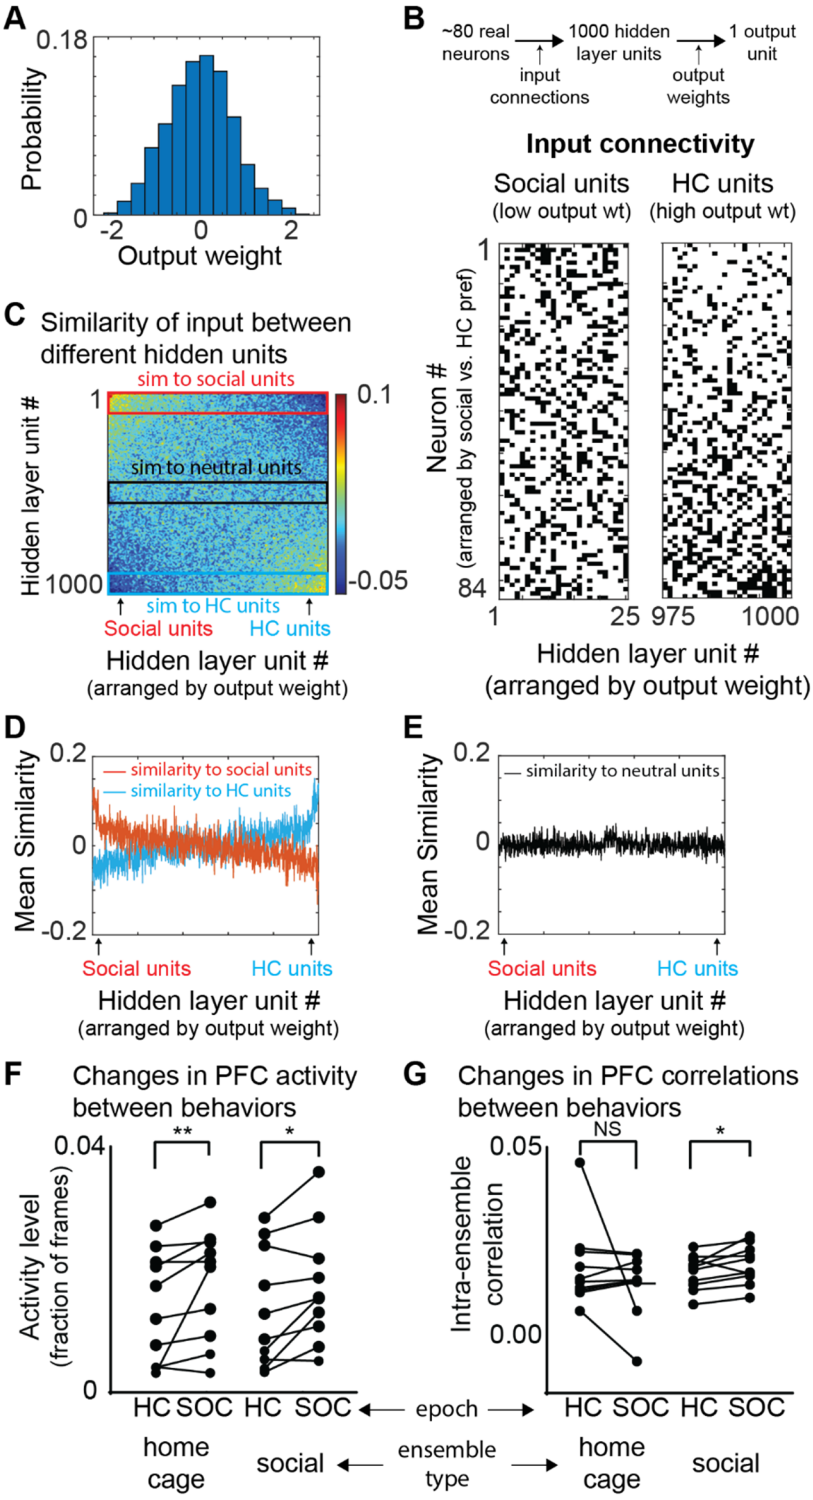

Figure 3. Classifier weights reveals an ensemble that increases correlations during social behavior and detects social behavior.

A. Example histogram depicting the distribution of output weights assigned to connections between hidden layer units and the output unit over the course of training.

B. Matrix of input connections for hidden units which detect the social (left) or home cage condition (right). The hidden layer units (x-axis) have been arranged in order of increasing output weights to identify 'social units' (25 most negative output weights) and 'home cage units' (25 largest positive output weights). Prefrontal neurons (y-axis) have been arranged in order of their preference for social interaction vs. home cage, i.e. the difference between their activity levels in the two conditions.

C. Correlation matrix showing the input similarity, i.e., the pairwise correlation between binary vectors representing the input connections to each pair of hidden layer units. Hidden layer units are arranged in order of increasing output weight. Red and blue rectangles indicate correlations with social or home cage units, respectively. A gaussian filter with a standard deviation of 3 was applied to the 1000x1000 matrix to improve visualization.

D. For each hidden layer unit, we plotted its average input similarity to either the 25 social units (red) or the 25 home cage units (blue). Hidden layer units (x-axis) are again arranged by output weight. Social units had similar patterns of input compared to each other but not to home cage units and vice-versa.

E. The average input similarity of each hidden layer unit to 25 hidden layer units with near-zero output weights ('neutral units'; black rectangle in C).

F. We defined social and home-cage (HC) ensembles as the 20% of prefrontal neurons most likely to provide input to the social or home cage units, respectively. The mean activity of both home cage and social ensembles increased during social interaction compared to the home cage condition (social ensemble: mean activity level 1.4 +/- 0.3% in home cage vs. 1.8 +/- 0.3% during social interaction, $p < 0.05$, sign-rank test; home cage ensemble: mean activity level 1.5 +/- 0.30% in home cage vs. 1.9 +/- 0.3% during interaction, $p < 0.001$, sign-rank test).

G. Correlations between neurons in the same ensemble increased during social interaction for the social ensemble but for the home cage ensemble (mean correlation coefficient between neurons in the social ensemble: 0.009 +/- 0.002 in home cage vs. 0.012 +/- 0.002 during social interaction, $p < 0.05$; home cage ensemble mean correlation coefficient 0.011 +/- 0.02 in home cage vs. 0.005+/- 0.003 during social interaction, $p=0.99$, sign-rank).

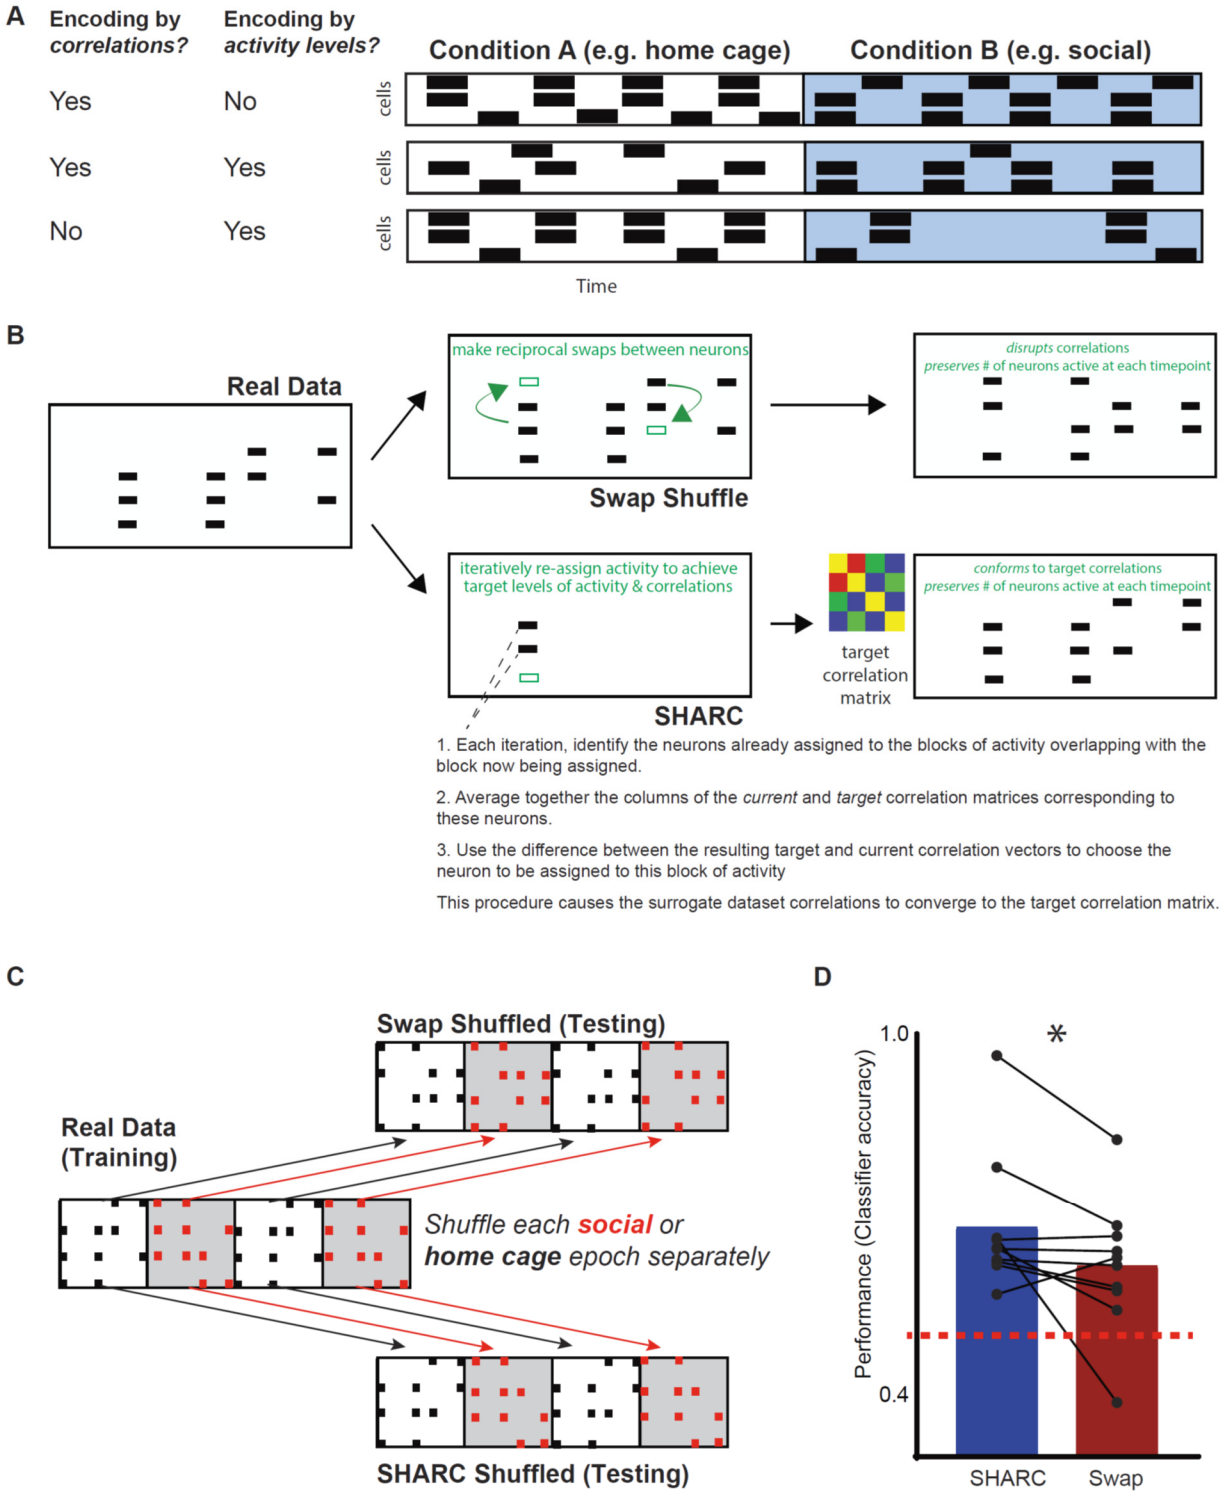

Figure 4. Correlations transmit additional information that is not efficiently conveyed by changes in activity levels alone.

A. Cartoon illustrating that information about behavior may be encoded through changes in activity levels, correlations between neurons, or both. When behavior modulates activity levels, correlations in two behavioral conditions may differ or be the same, and vice-versa.

B. To disentangle the roles of activity levels and correlations in transmitting information we used two different methods to create shuffled (surrogate) datasets which preserve changes in activity levels, but either do or do not preserve patterns of correlations. We made random, reciprocal swaps of activity between neurons to generate surrogate datasets which maintained network activity in each frame as well as the number of blocks of activity for each neuron. However, these datasets destroyed the correlation structure. In a second set of surrogate datasets we used SHARC to iteratively generate surrogates in which the correlation structure was also maintained.

C. To maintain dynamic changes in activity levels and correlations that are associated with the two behavioral conditions we swap-shuffled or performed SHARC separately for each behavioral epoch, then concatenated the 9 resulting surrogate subrasters to create each surrogate dataset.

D. We trained a classifier (with a connection probability = 0.3) on each real dataset, then tested that classifier on swap or SHARC-shuffled surrogate datasets generated from that real dataset. Accuracy was significantly higher for the SHARC-shuffled surrogates, which maintain the behaviorally-modulated correlations found in the original dataset (accuracy for SHARC shuffled surrogate datasets = 68 +/- 4%, classifier accuracy for swap shuffled surrogate datasets = 61 +/- 4%, $p < 0.05$, sign-rank test, n = 10 mice).

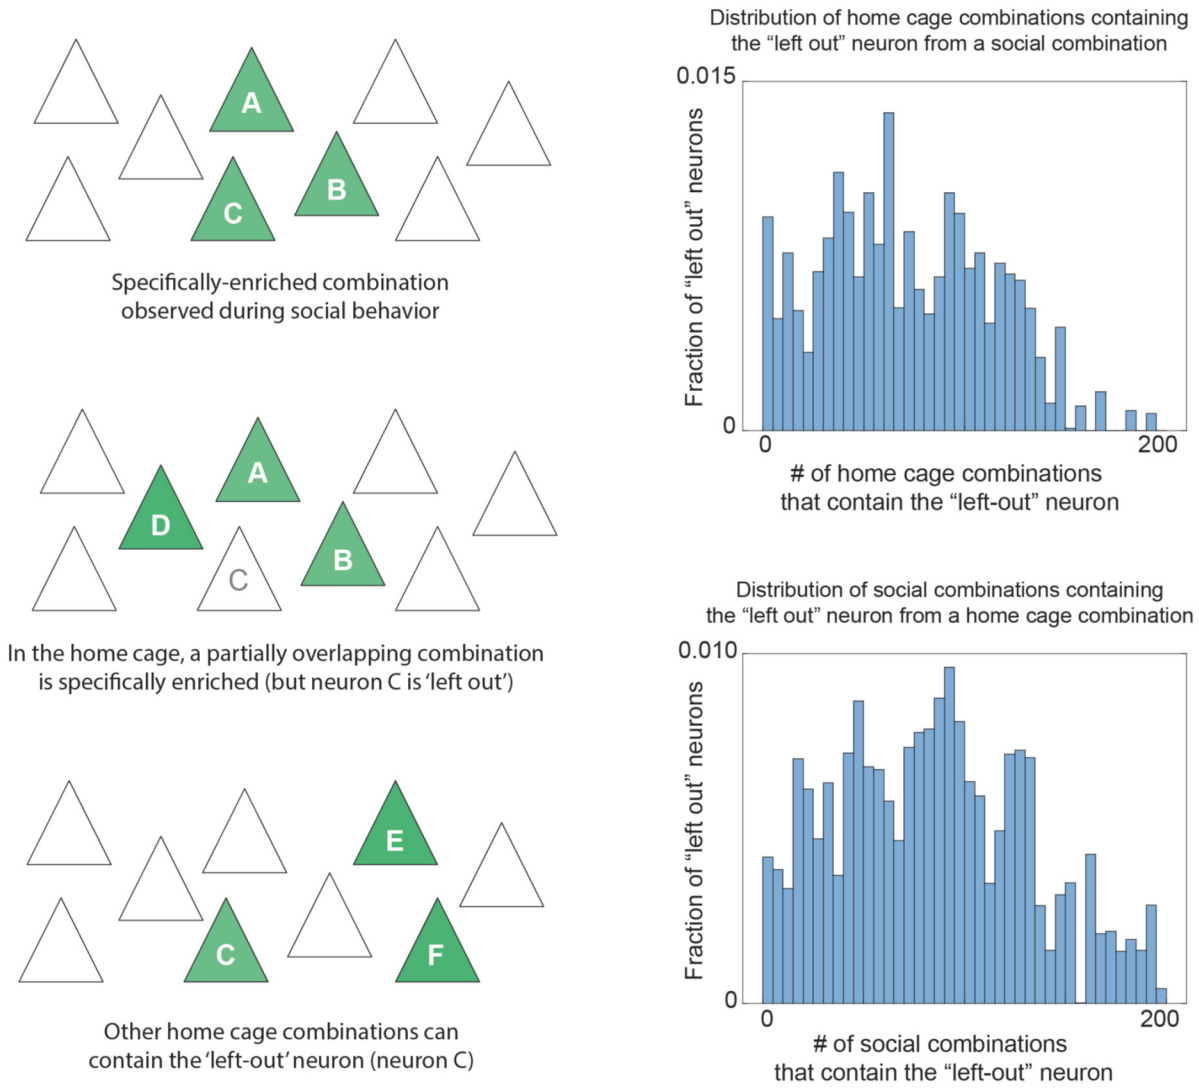

Figure 5. Behaviorally-specific patterns of coactivity are formed by neurons that are active in both conditions, but coactive with different partners in each condition.
Left: We identified combinations of 3 neurons that are specifically enriched during one behavioral condition (occurring more often during social interaction than in 95% of surrogate datasets, and occurring less often during home cage exploration than in 50% of surrogate datasets, or vice-versa). We then identified overlapping combinations occurring during the opposite behavioral condition in which a single neuron was 'left out.' In other words, we identified combinations from the two conditions that overlapped in exactly two neurons. **Right, top:** Histogram showing the number of distinct 3 neuron home cage combinations that contain the neuron which participates in a social combination but is 'left out' during home cage behavior. **Right, bottom:** Histogram showing the number of distinct 3 neuron social combinations that contain the neuron which participates in a home cage combination but is 'left

917 out' during social interaction. In the vast majority of cases, neurons that are 'left out' in one
918 condition are still active during that condition and participate in other combinations.

919

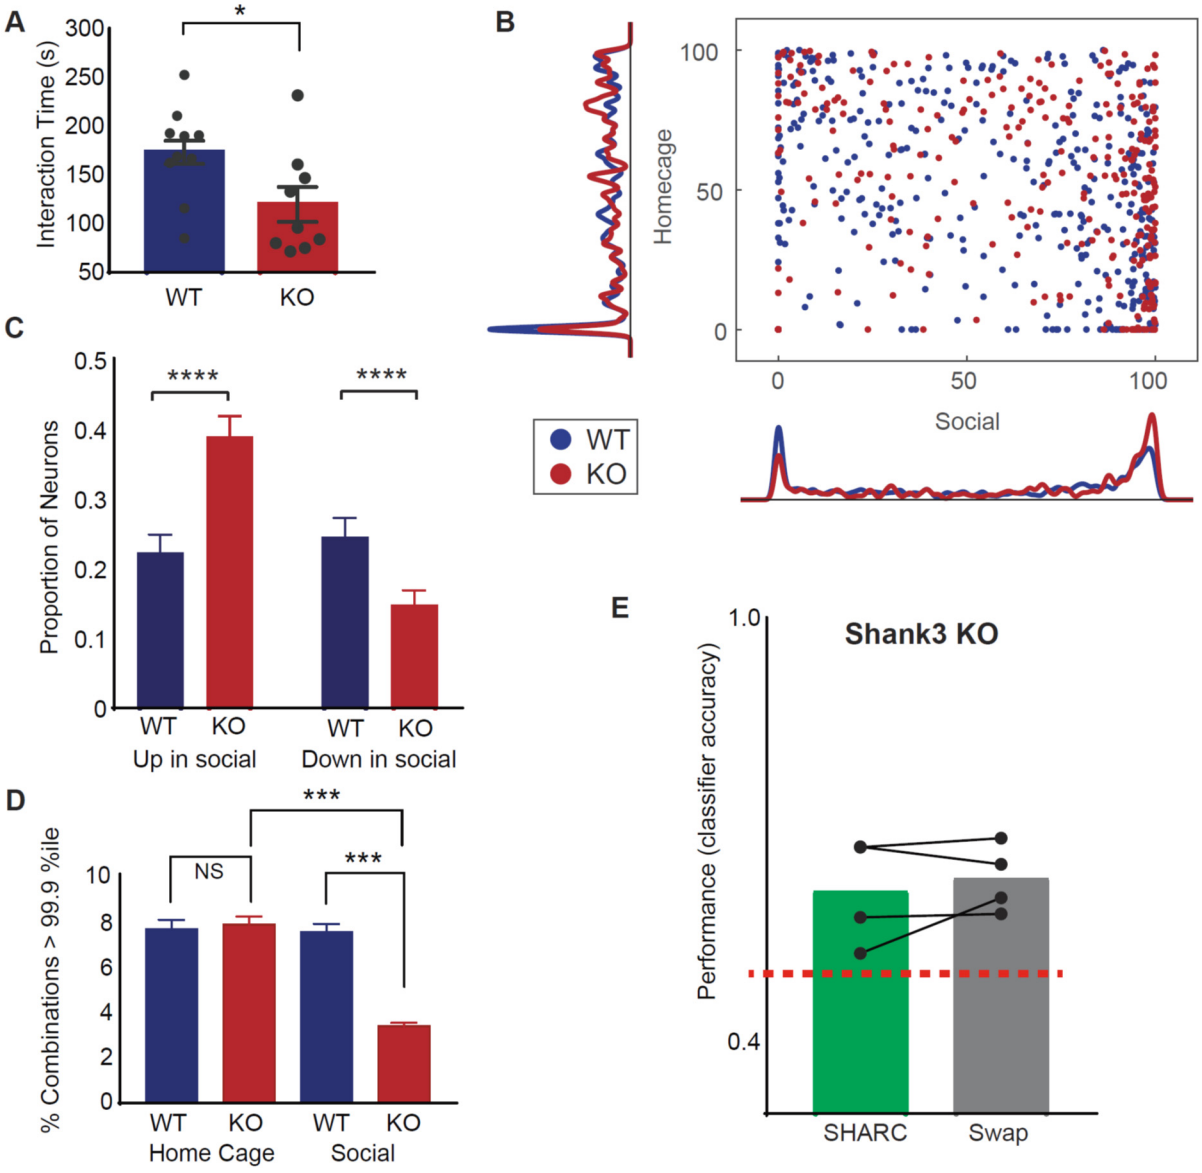

Figure 6. Shank3 KO mice have disorganized ensembles for which correlations fail to enhance the transmission of information about social behavior.

A. The mean time that *Shank3* KO mice or wild-type (WT) littermates spend interacting with a novel juvenile mouse of the same sex during a 5 min assay. Data has been pooled from 8 unimplanted WT mice as well as the 5 implanted WT mice used for microendoscopic imaging, and 5 unimplanted KO mice in addition to the 4 implanted mice used for imaging. For implanted mice we used the average of interaction time for the 2 novel mice. Pooled data showed decreased interaction in KO mice (173 +/- 12 s vs. 120 +/- 18 s for WT and KO respectively, $p < 0.05$, t-test). The un-implanted cohort alone shows a similar significant decrease in interaction time for KO mice (165 +/- 15 s vs 110 +/- 16 s for WT and KO respectively, $p < 0.05$). In the implanted cohort there was a similar trend toward decreased

interaction for KO mice (186 +/- 20 s vs 133 +/- 37 s, for WT and KO respectively, p = 0.21).

B. (Similar to Fig. 1D). Scatter-plot showing the activity of each neuron during each behavioral condition, expressed as a percentile relative to a null distribution generated by circularly shuffling that neuron's activity. Activity levels during social interaction or while the mouse was alone in its home cage are plotted on the x and y axis, respectively. Kernel density plots along the axes indicate the fraction of neurons whose activity was at a given percentile of the null distribution. Neurons with activity > 90th percentile of shuffled datasets (green dotted line) were considered to be positively modulated, whereas neurons with activity < 10th percentile (green dotted line) were considered to be negatively modulated during each behavior**.** Data is plotted for Shank3 KO mice (red) and WT littermates (blue) (mean percentile of activity during social interaction, WT: 50 +/- 2 percentile; KO: 64 +/- 2 percentile; p < 0.0001 by 2-sample *t*-test; mean percentile of activity during home cage, WT: 47 +/- 2 percentile; KO: 51 +/- 2 percentile; p = 0.1, *t*-test).

C. Bar graph showing the fraction of neurons whose activity was positively or negatively modulated (>90th percentile or <10th percentile) during social interaction. The proportion of neurons which increased activity during social interaction was significantly greater in KO mice (22% in WT vs. 39% in KO, chi-squared = 17.7, p < 0.0001), whereas the downregulated ensemble was significantly smaller in KO mice (25% in WT vs. 15% in KO, chi-squared test, 8.2, p < 0.005). Error bars denote the binomial S.E.M. algebraically derived from total number of neurons and the proportion that were modulated in the specified direction.

D. The proportion of 3 neuron combinations occurring during home cage exploration that are enriched > the 99.9th percentile compared to swap-shuffled datasets was similar across WT (Blue; 7.6%) and KO (Red; 7.8%) mice. By contrast, the proportion of 3 neuron combinations occurring during social interaction that are enriched > 99.9th percentile compared to swap-shuffled datasets was 7.5% in WT compared to only 3.4% in KO mice (total number of home cage combinations: 4,187 in 5 WT mice, and 5,878 in 4 KO mice; total number of social combinations: 5,487 in 5 WT mice, 16,326 in 4 KO mice). The top two plots show histograms of enrichment for the home cage (upper) or social conditions (middle); the lower panel is a bar graph showing the fraction of these combinations that were specifically enriched above the 99.9th percentile (chi-squared = 165, p < 0.0001). Error bars denote the S.E.M. algebraically derived from the binomial distribution, the number of 3 neuron combinations in each condition, and the proportion of those combinations that were enriched.

E. Performance of classifier trained on real datasets and tested on surrogate datasets. Performance was not better (and was non-significantly worse) when correlation structure was maintained using SHARC (classifier accuracy: 62 +/ 4% for SHARC vs 63 +/- 2% for swap shuffled surrogate datasets, p = 0.47, sign-rank test).

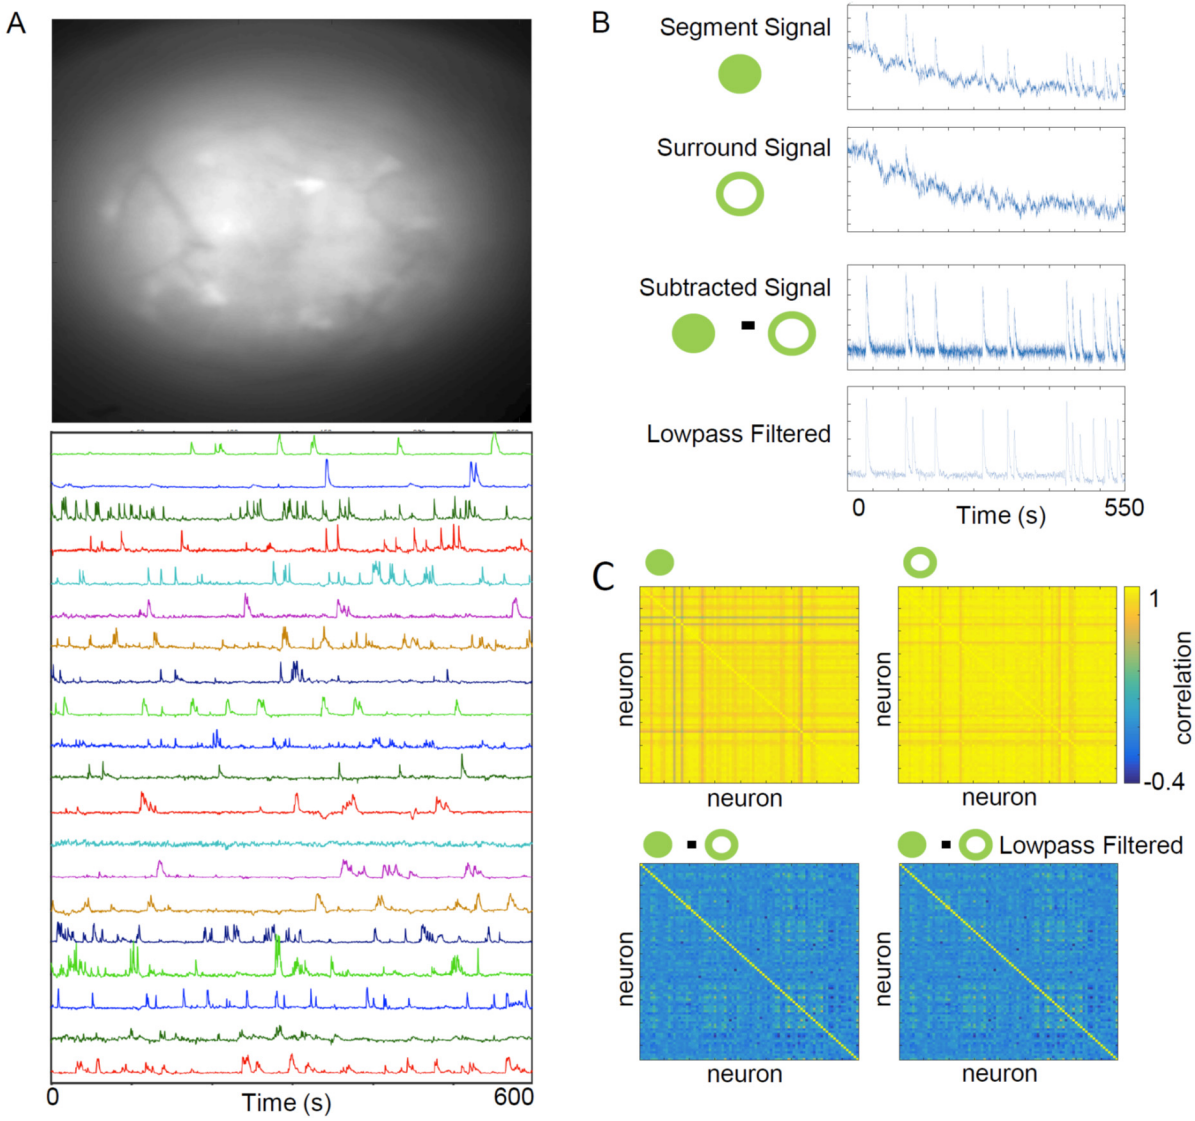

Supplementary Figure 1. Spatial Decorrelation of Neuronal Signals.

A. Example image (top) and individual neuron GCaMP traces (bottom) from prefrontal cortex imaged with implanted endoscope.

B. The average GCaMP signal from a region of interest (ROI), corresponding to one neuron, was corrected by subtracting the average GCaMP signal from the surrounding pixels, in order to spatially deconvolve signals from each ROI vs. the surrounding neuropil. Examples traces from a single neuron are shown.

C. The pairwise correlation matrix between signals from different neurons is shown (calculated from 550 seconds of activity from a single wildtype mouse), for the original GCaMP signals (top left), the surround signals (top right), the surround-subtracted signals (bottom left), and the surround-subtracted signals after lowpass filtering (bottom right).