Supplemental Figures

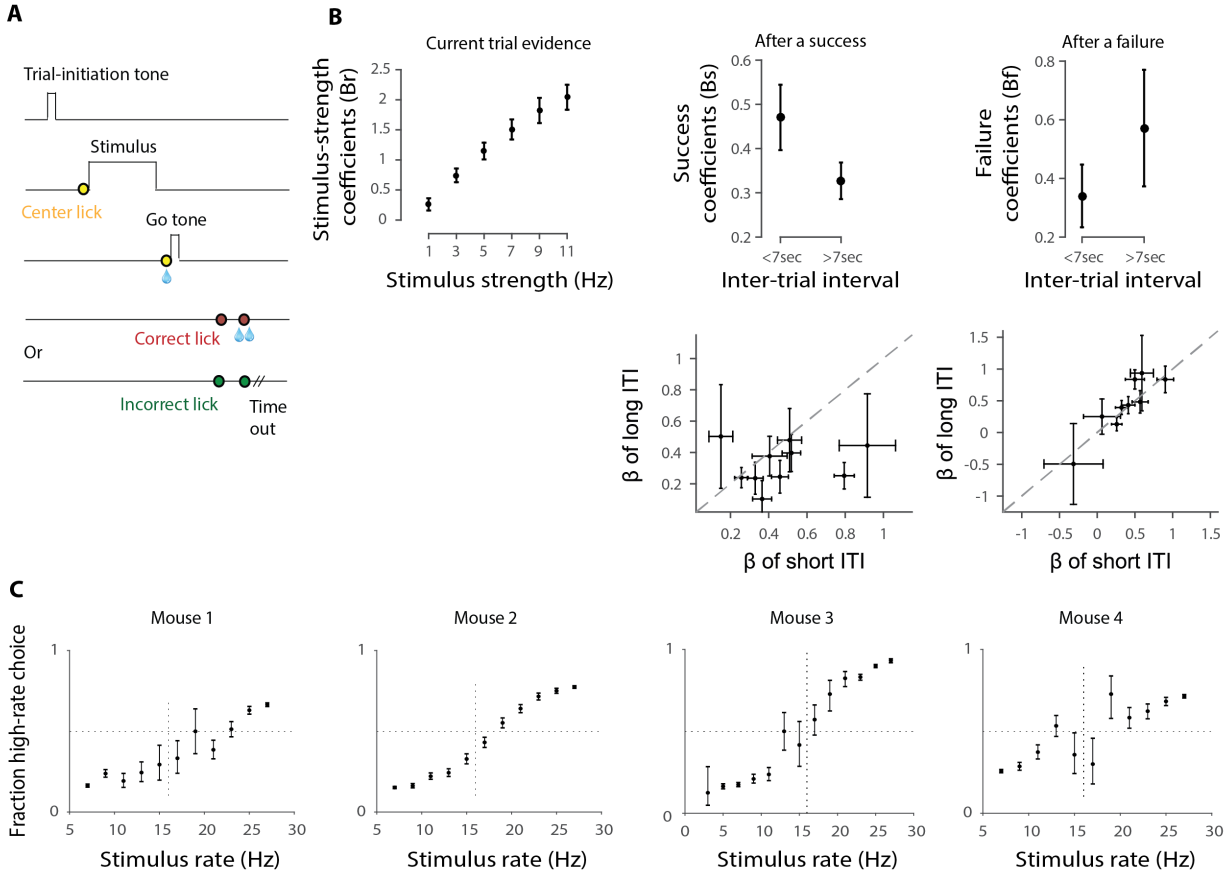

Figure S1. Related to Figure 1. Perceptual decisions about stimulus rate reflect current evidence, previous trial's outcome, and the time passed since the previous trial.

A**,** Trial structure. In each trial, first a brief tone is presented to indicate to the animal to initiate the trial ("trial-initiation tone"). Once the animal licks to the center waterspout (row 2: yellow circle), the stimulus is presented for 1 sec. At the end of the stimulus, the animal is required to lick again in the center (row 3: yellow circle). This will result in: 1) a small water reward in the center, 2) a "go tone" that indicates to the animal to make its choice. If the animal licks to the correct side (row 4, 1st red circle), and confirms this lick (row 4, 2nd red circle), it will receive water as a reward. If the animal licks to the wrong side (last row, 1st green circle), and confirms this lick (last row, 2nd green circle), there will be a time-out, i.e. longer time before the next trial can start. **B**, A logistic regression model was used to assess the extent to which the animal's choice depends on stimulus strength (how far the stimulus rate is from the category boundary at 16Hz), previous choice outcome, and the time interval since the previous choice. Stimulus strength was divided into 6 bins (**left**); previous success was divided into 2 bins: success after a long ITI and success after a short ITI (**middle**); previous failure was also divided into 2 bins: failure after a long ITI and failure after a short ITI (**right**). Plots in top row show β averaged across animals (same 10 animals as in Figure 1B). Error bars: S.E.M across subjects. **Top left**: strength of the sensory evidence affects the animal's choices: the stronger the evidence, the higher the impact. **Top middle**: Success of a previous trial also affects animal's decision; the effect is stronger when the previous success occurs after a short ITI (<7sec). **Top Right**: Same but for previous incorrect trials; the effect of ITI after a failure was not significant. Plots in **bottom** row show success (left) and failure (right) β for individual mice. Error bars: S.E.M returned from glmfit.m in Matlab. **C,** Behavioral performance of the four mice in which we imaged excitatory and inhibitory activity during decision-making. In mice 1, 2, and 4, imaging was performed throughout learning by tracking the same group of neurons. Plots reflect data from all sessions. Errors bars: Wilson Binomial Confidence Interval.

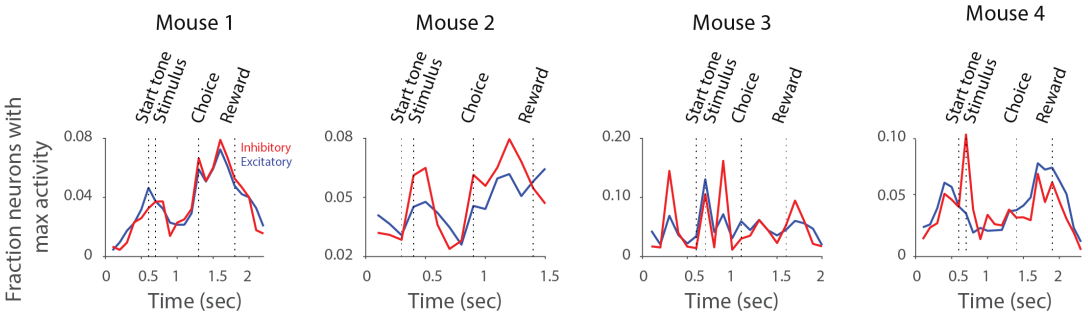

Figure S2. Related to Figure 1. Excitatory and inhibitory neurons have similar temporal dynamics.

For each session, the fraction of neurons with peak activity in each 100ms time window was computed. This quantity is an estimate of the temporal-epoch tuning of neurons. Curves show mean across sessions, for excitatory (blue) and inhibitory (red) neurons, for each mouse. Similar to Figure 1E, traces were aligned for each trial event (start tone, stimulus, choice, reward), and then concatenated (see Figure 1E, legend).

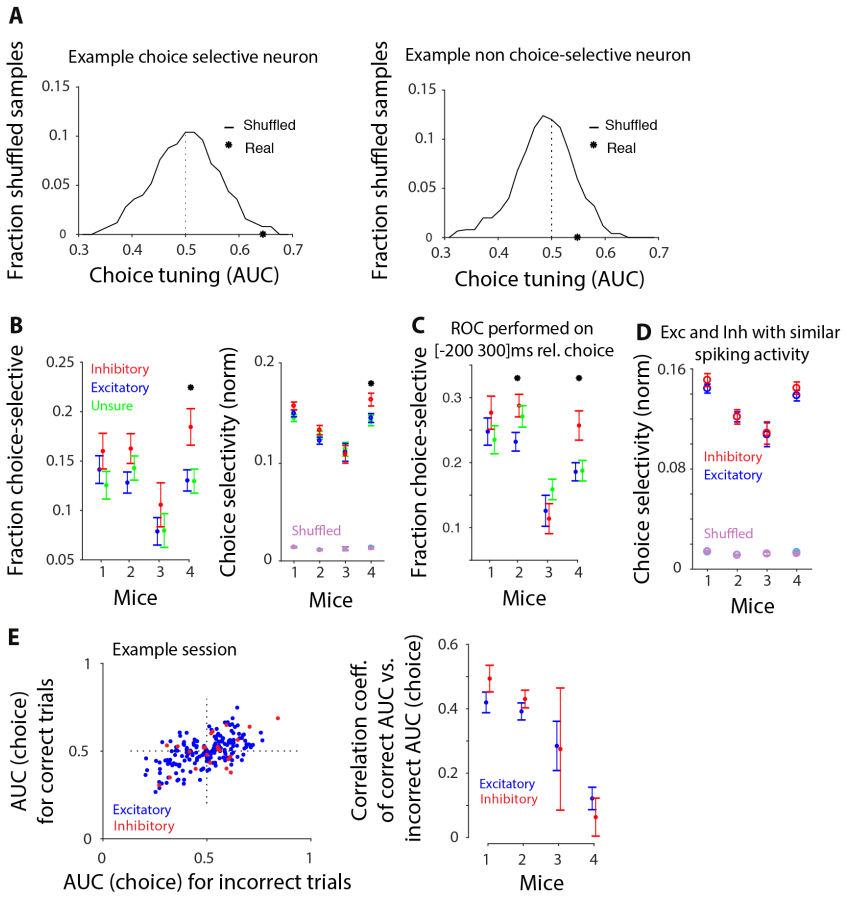

Figure S3. Related to Figure 2. Single neuron measures reveal similar choice selectivity in excitatory and inhibitory neurons.

A, Example neurons to illustrate the method for assessing significant choice selectivity in individual neurons. In both panels, the solid line shows the distribution of values for the area under the ROC curve (AUC) generated by 50 different trial shuffles in which trials were randomly assigned to a left vs. right choice. Star indicates the actual AUC value of the neuron. Significance was assessed from the probability of occurrence of the actual AUC value in the shuffle distribution. When probabilities were <0.05, neurons were considered choice selective. Only the neuron on the left has significant choice selectivity. **B,** Fraction (left) and magnitude (right) of choice selectivity are shown for the unsure neurons (i.e. neurons classified as neither excitatory nor inhibitory; green), as well as excitatory (blue) and inhibitory (red) neurons. Data for each mouse show mean +/- standard error across sessions. **C,** Fraction of choice-selective neurons based on ROC analysis on [-200 300]ms relative to the choice. Fraction selective neurons at this time window (median across mice): excitatory: 21%; inhibitory: 27%, resulting in approximately 11 inhibitory and 69 excitatory neurons with significant choice selectivity per session. There is a considerable increase in the fraction of selective neurons when using this time window rather than 0-97ms window (see Figure 2C for comparison). **D,** ROC analysis restricted to those excitatory and inhibitory neurons that had the same spiking activity. Choice selectivity is still similar between the two cell types. Note that the significant difference observed for mouse 4 in Figure 2C is absent after controlling for the difference in spiking activity of inhibitory and excitatory neurons. Mean +/- standard error across sessions. **E, left:** Choice selectivity was computed on correct trials (vertical axis) as well as error trials (horizontal axis), and was correlated between the two conditions. Data is from a single session, each point shows an individual neuron whose cell type is indicated by its color. The positive correlation indicates that choice selectivity was overall similar on correct and error trials (Pearsons' correlation coefficient, excitatory neurons: r=0.58; p<0.001; inhibitory neurons: r=0.55, p=0.007). The small number of points in quadrants 2 and 4 indicate less frequent neurons that showed opposite selectivity on correct vs. error trials. **Right,** Summary of correlation coefficient of AUC on correct trials vs. AUC on incorrect trials, mean across sessions for each animal. Error bars: S.E.M. across sessions. The weaker correlation in mouse 4 indicates that this animal had a mixture of cells selective for the stimulus and cells selective for the choice. Note that although the center of the imaging window was identical in all animals, the imaging location within the window of this animal was slightly posterior to the others. The enrichment of cells selective for the stimulus, in this mouse compared to other mice, may reflect that the region we imaged in mouse 4 was closer to primary visual cortex.

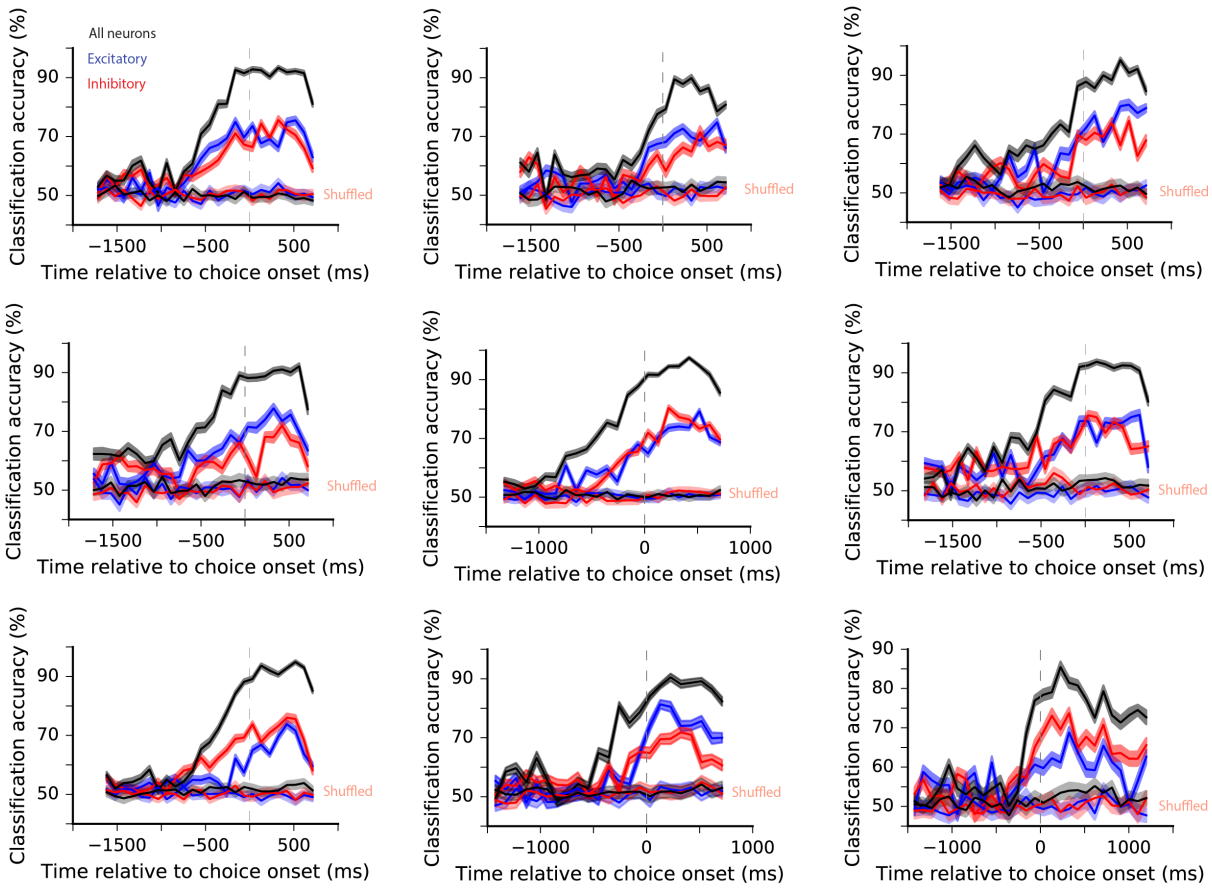

Figure S4. Related to Figure 3. Population activity is highly selective for the animal's choice; excitatory and inhibitory neurons are similarly selective.

Classification accuracy of the choice decoder at each moment in the trial for 9 additional example sessions. Dashed lines: choice onset. Black: all neurons included in the decoder; blue: subsampled excitatory neurons; red: inhibitory neurons; dim colors: shuffled control. In most sessions, inhibitory and subsampled excitatory populations have comparable classification accuracy.

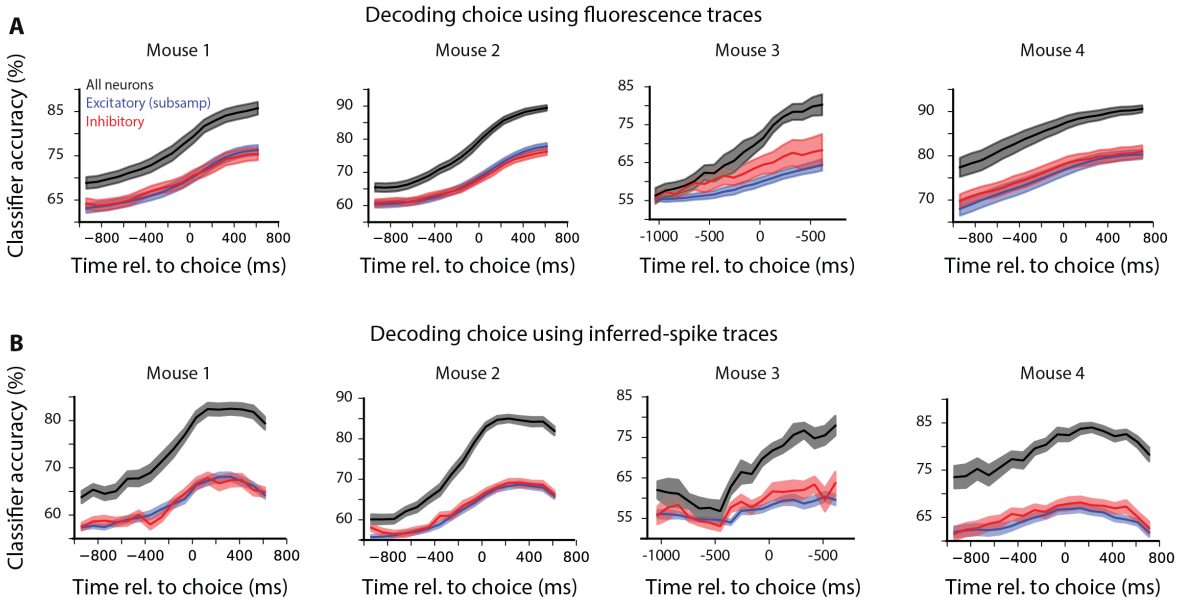

Figure S5. Related to Figure 3. Classification accuracy is similar for excitatory and inhibitory populations, whether the choice decoder is trained/tested on fluorescence traces or on inferred spikes.
SVM classifiers were trained to decode choice from the population activity of all neurons (black), inhibitory neurons (red), or subsampled excitatory neurons (blue). In **(A)** fluorescence traces (Figure 1D middle) were used, and in **(B)** inferred spikes (Figure 1D right) were used. In both cases, decoder accuracy is similarly high for excitatory and inhibitory neurons.

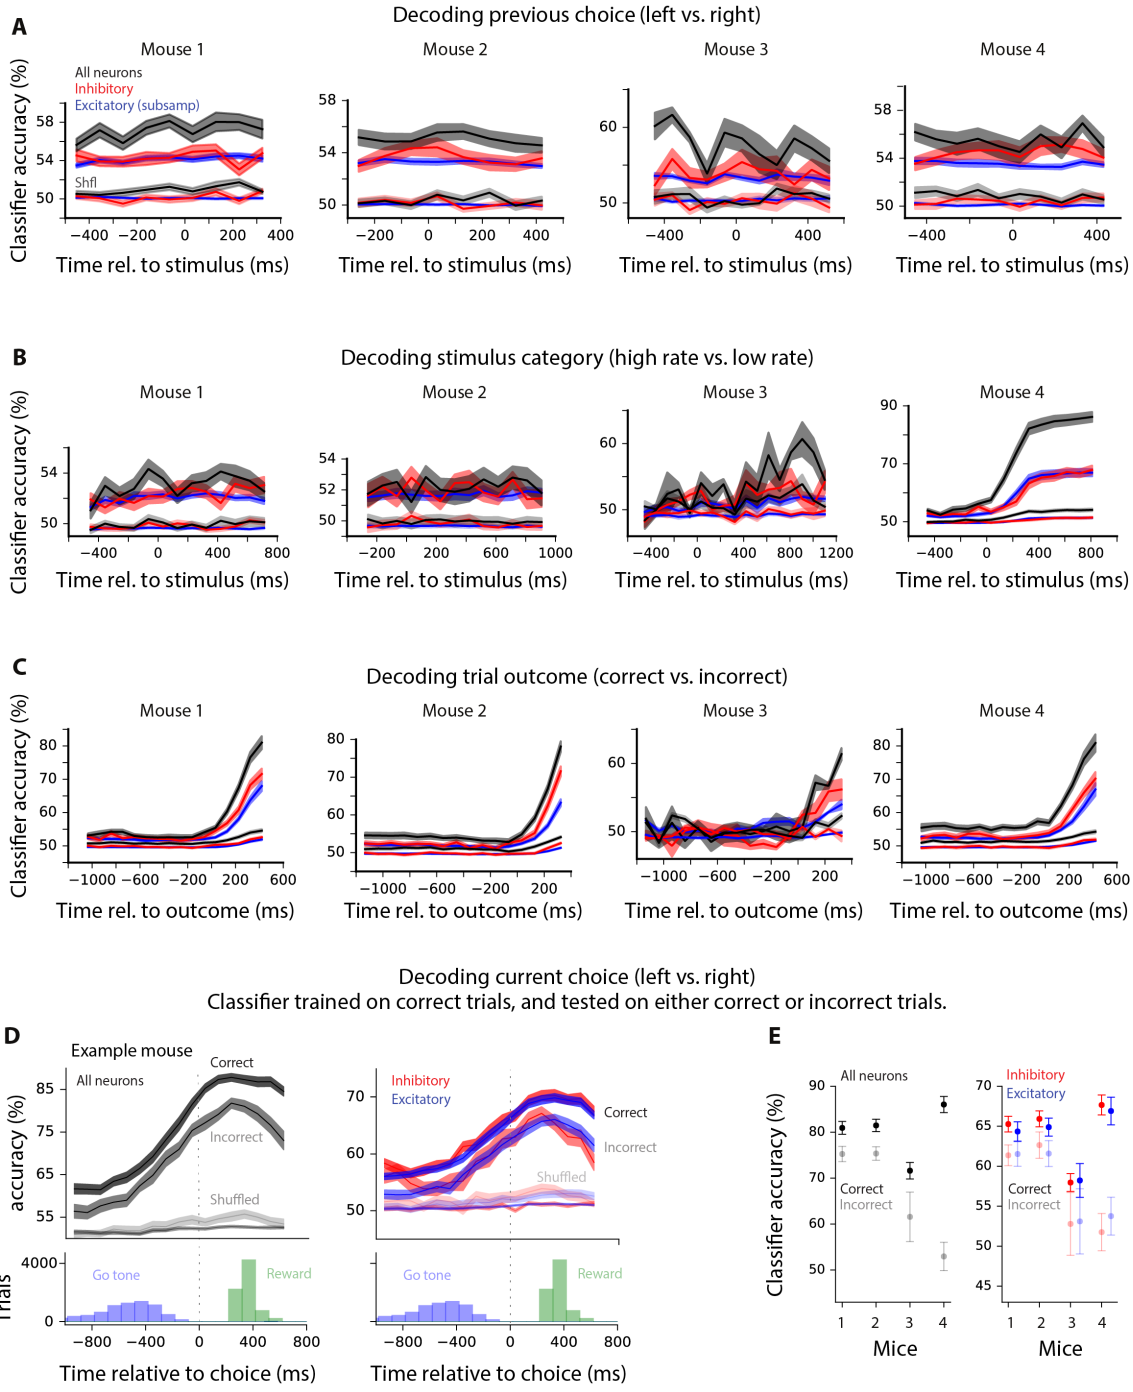
A
Decoding previous choice (left vs. right)
Mouse 1
Mouse 2
Mouse 3
Mouse 4
All neurons
Inhibitory
Excitatory (subsamp)
Shfl
Classifier accuracy (%)
Time rel. to stimulus (ms)
B
Decoding stimulus category (high rate vs. low rate)
C
Decoding trial outcome (correct vs. incorrect)
Time rel. to outcome (ms)
Decoding current choice (left vs. right)
Classifier trained on correct trials, and tested on either correct or incorrect trials.
D
Example mouse
Correct
Incorrect
Shuffled
Trials
Go tone
Reward
Time relative to choice (ms)
Inhibitory
Excitatory
E
Mice

Figure S6. Related to Figure 3. Population activity is strongly selective for the trial outcome, and to a lesser degree to the stimulus category and previous choice.

A, SVM classifier trained to decode previous choice from the activity of all neurons (black), inhibitory neurons (red), or subsampled excitatory (blue) neurons. "shfl" indicates classifier accuracy trained using shuffled trial labels. Previous choice is reflected, though weakly, in the population activity of the current trial. **B,** SVM classifier trained to decode the stimulus category, i.e. whether the stimulus is high rate (above 16Hz) or low rate (below 16Hz). Except for mouse 4, in which the imaging location was slightly more posterior (see Figure S3E, legend), stimulus category is weakly reflected in the population activity. **C,** SVM classifier trained to decode the trial outcome (i.e. correct vs. incorrect). Classification accuracy gradually increases and reaches 80% (median across mice) approximately 400ms after the animal confirms his choice (Figure S1A). Inhibitory neurons showed slightly higher selectivity for the outcome. Unsaturated lines in B and C: performance on shuffled trials. **D,** SVM classifier trained on correct trials to decode choice and tested on correct as well as incorrect trials. Data from an example animal (48 sessions). **Top**: Classification accuracy of decoders trained on all neurons (left), subsampled excitatory neurons (right, blue trace), and inhibitory neurons (right, red trace). In all cases, classifiers were trained on correct trials; however they were tested either on correct (dark lines: "Correct") or incorrect (dim lines: "Incorrect") trials. Classification accuracy on incorrect trials was high; indicating that population activity primarily reflects the animal's choice, yet it differs at least slightly for correct and incorrect trials. This reduction was similar for excitatory and inhibitory neurons (blue are red traces are overlapping in the right panel). **Bottom**: Across-trial distribution of go tones and reward delivery (See Fig. 3B bottom). Left and right panels are the same plots and are duplicated to facilitate alignment to each corresponding plot above. **E,** Summary across all mice for all neurons (left) and excitatory and inhibitory neurons separately (right). Classifier performance on correct (dark colors) and incorrect (dim colors) trials is shown. Mouse 4 had the largest difference in classification accuracy for correct vs. error trials. As with the single-neuron analysis (Figure S3E) and decoding of stimulus category (Figure S6B), this difference likely reflects that the imaging region was slightly posterior within the window for this animal. Importantly, for all mice, the change in classification accuracy was quite similar for excitatory and inhibitory neurons (right), indicating that both populations reflect choice vs. stimulus to a comparable degree.

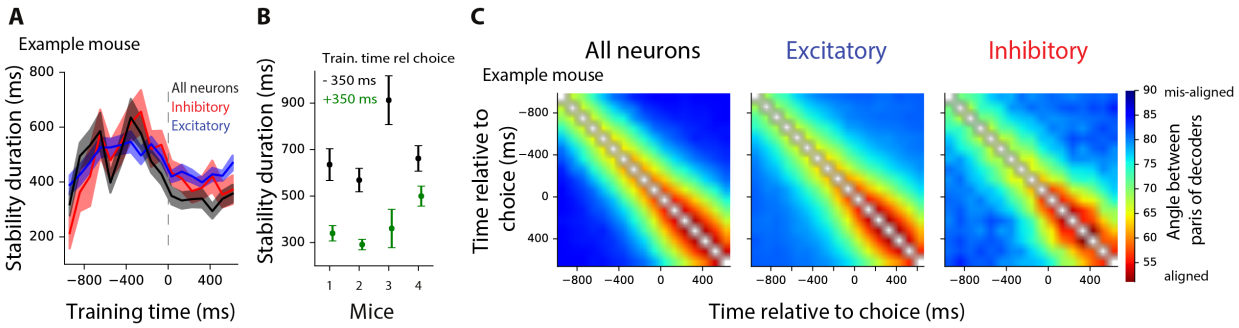

Figure S7. Related to Figure 4. Additional analyses provide more evidence for similar temporal stability of the choice decoder in excitatory and inhibitory populations.

A, In an example mouse, population activity that predicts the animal's choice is similarly stable for excitatory and inhibitory neurons during the course of a trial. The vertical axis shows the stability duration for decoders trained at different times during the trial. Stability duration is defined as the width of the testing window over which decoder accuracy does not statistically differ from that within the training window (red regions of Figure 4C) from that obtained by using the same training and testing times (diagonal of Figure 4A). Error bars: S.E.M. across sessions. Summary data for all mice at training time 0-97 ms before choice (dashed line) are shown in Figure 4D. **B,** Stability duration of the all-neuron decoder (black in panel A) is compared for decoders trained 350ms before the choice (black), and 350ms after the choice (green). Population stability was lower after the choice than before the choice. This may be due to additional events, e.g. reward delivery and repeated licking, which follow the choice. **C**, Another measurement of stability likewise suggests similar temporal stability for excitatory and inhibitory populations. Stability was assessed by measuring the angle between pairs of decoders trained at different time points in the trial. If a similar pattern of population activity represents choice from moment t_1 to moment t_2, the choice classifiers trained at these times will be aligned, i.e. the angle between the two classifiers will be small. The colors indicate the angle between pairs of decoders trained at different moments in the trial. Small angles (hot colors) indicate alignment of choice decoders; hence stable activity patterns, related to choice, across neurons. Left: all neurons; middle: excitatory neurons (subsampled to match the number of inhibitory neurons); right: inhibitory neurons. As with our other method (Figure 4), the time course of stability was similar for excitatory and inhibitory neurons.

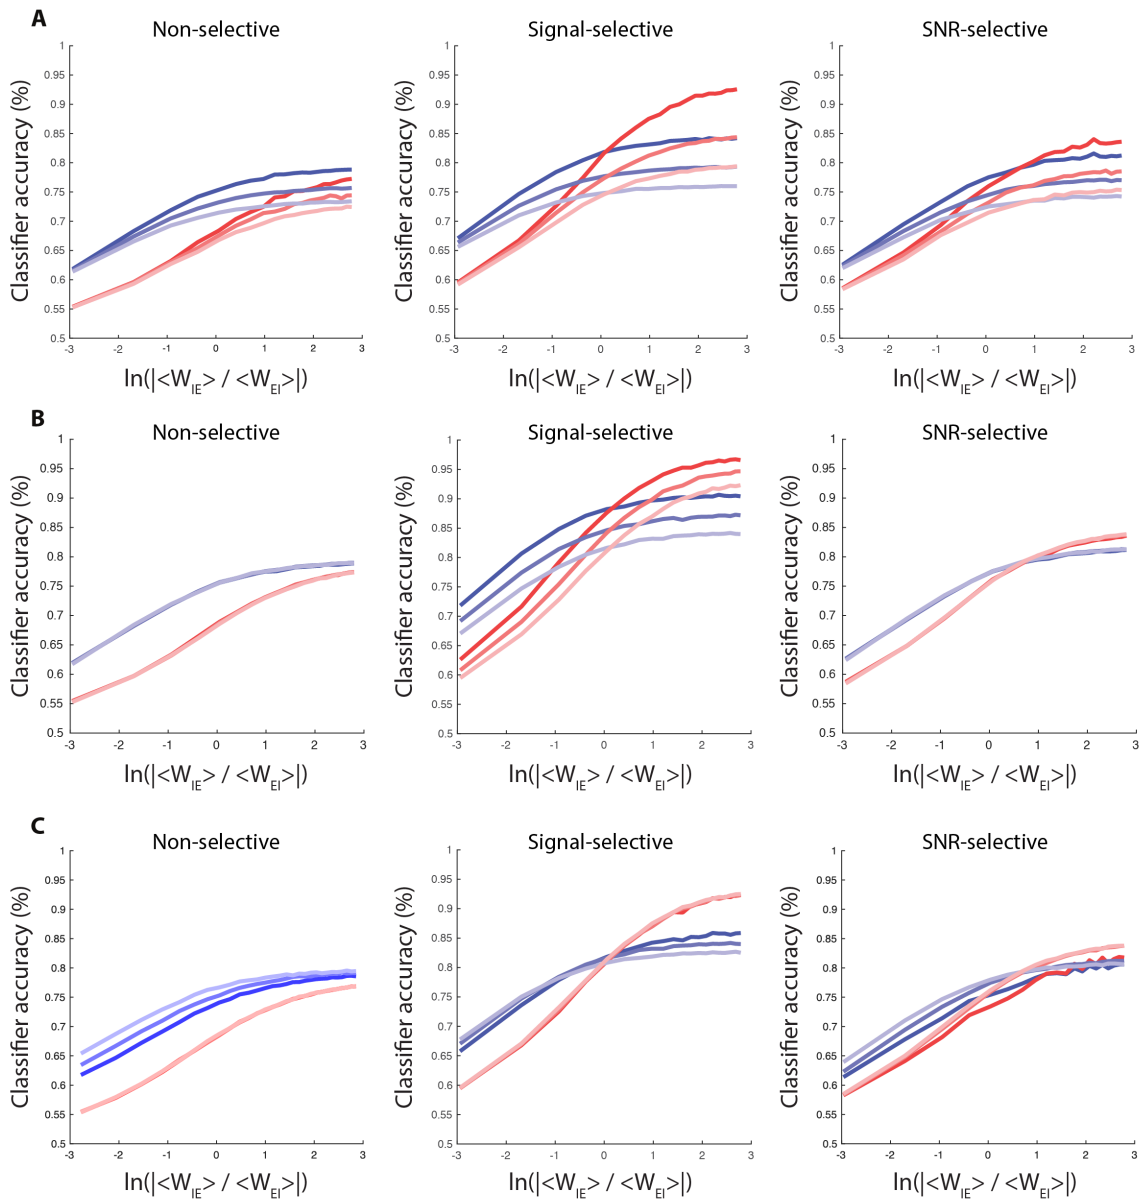

Figure S8. Related to Figure 5. Selective connectivity between excitatory and inhibitory neurons allows for matched classification accuracy in the two populations.

Decoding accuracy versus three parameters. **A,** Differential correlations, which are known to be present in any realistic network (Moreno-Bote et al., 2014). $\mathbf{\Sigma}_{EE} \rightarrow \mathbf{\Sigma}_{EE} + \epsilon \Delta\mathbf{h}\Delta\mathbf{h}$. Dark to light hues: $\epsilon = 0, 0.25, 0.50$. **B,** Excitatory to excitatory connections. Dark to light hues: $w_{EE} =$ 0.35, 0.3, 0.25 (default). **C,** Inhibitory to inhibitory connections. Dark to light hues: $w_{II} =$ -2.4, -2.0 (default), and -1.6.

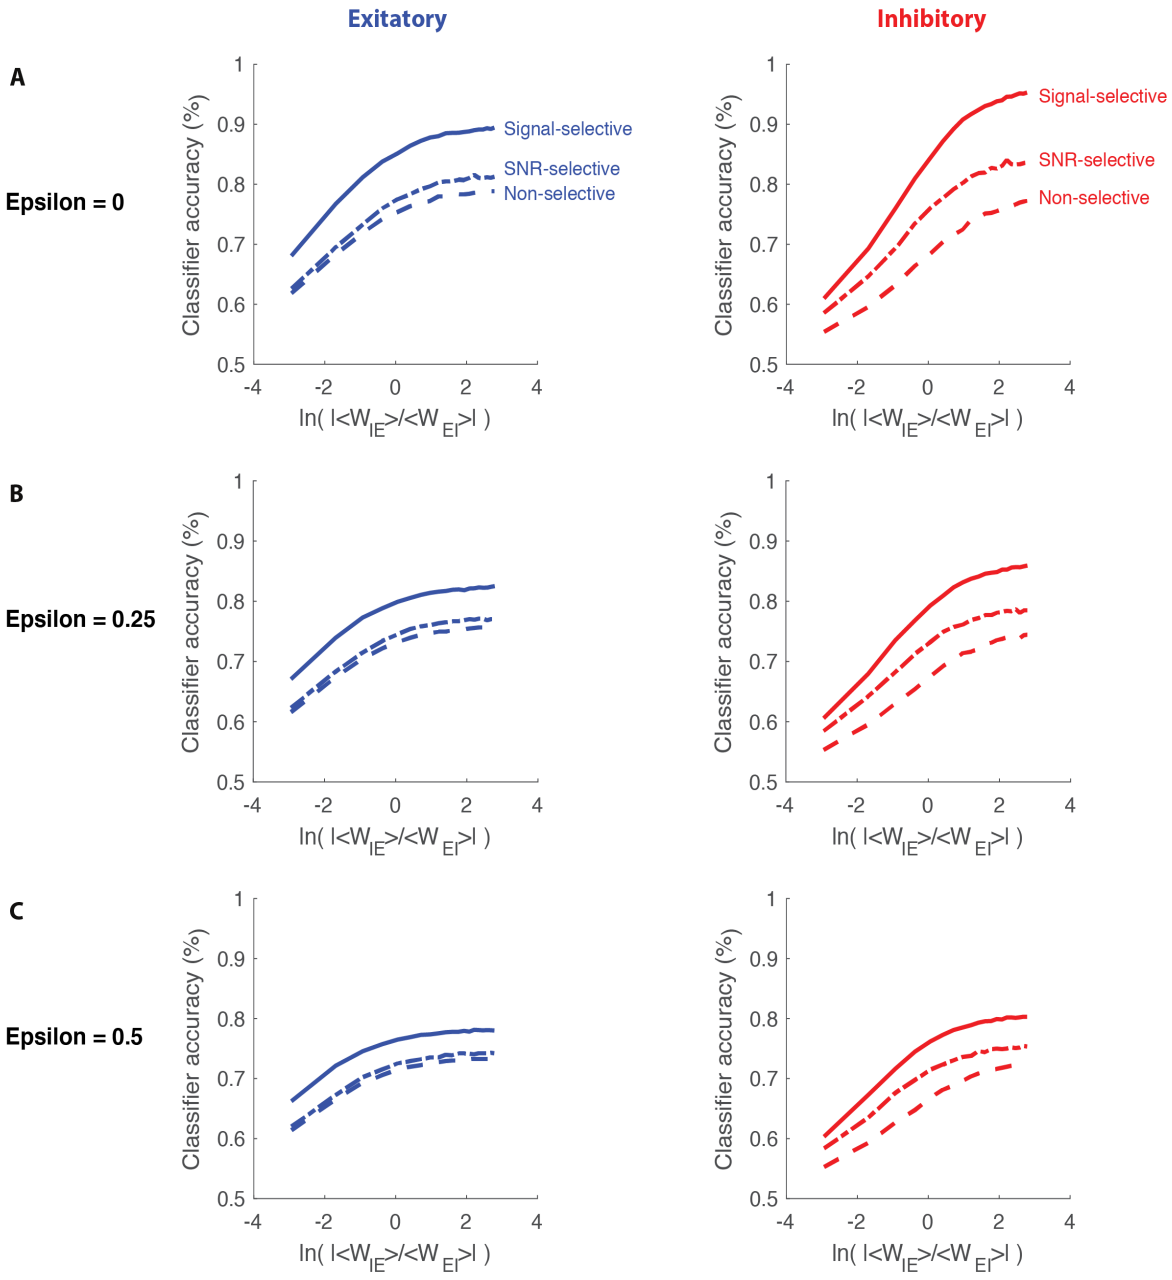

Figure S9. Related to Figure 5. Signal-selective models of decision-making have significantly higher classification accuracy compared to non-selective and SNR-selective models.

Classification accuracy of excitatory (blue) and inhibitory (red) neurons is compared for 3 models of decision-making with different levels of differential correlations, as controlled by ε (see Fig. S8 caption). We focused on differential correlations because they are known to be present in any realistic network (Moreno-Bote et al., 2014). Solid: signal-selective; Dashed-dotted: SNR-selective; Dashed: non-selective. **A, B, C:** $\epsilon = 0, 0.25, 0.50$, respectively. Horizontal axis shows the relative strength of excitatory-to-inhibitory vs. inhibitory-to-excitatory connections. For all values of this parameter and in both cell types, signal-selective model has the highest classification accuracy for decoding the animal's choice.

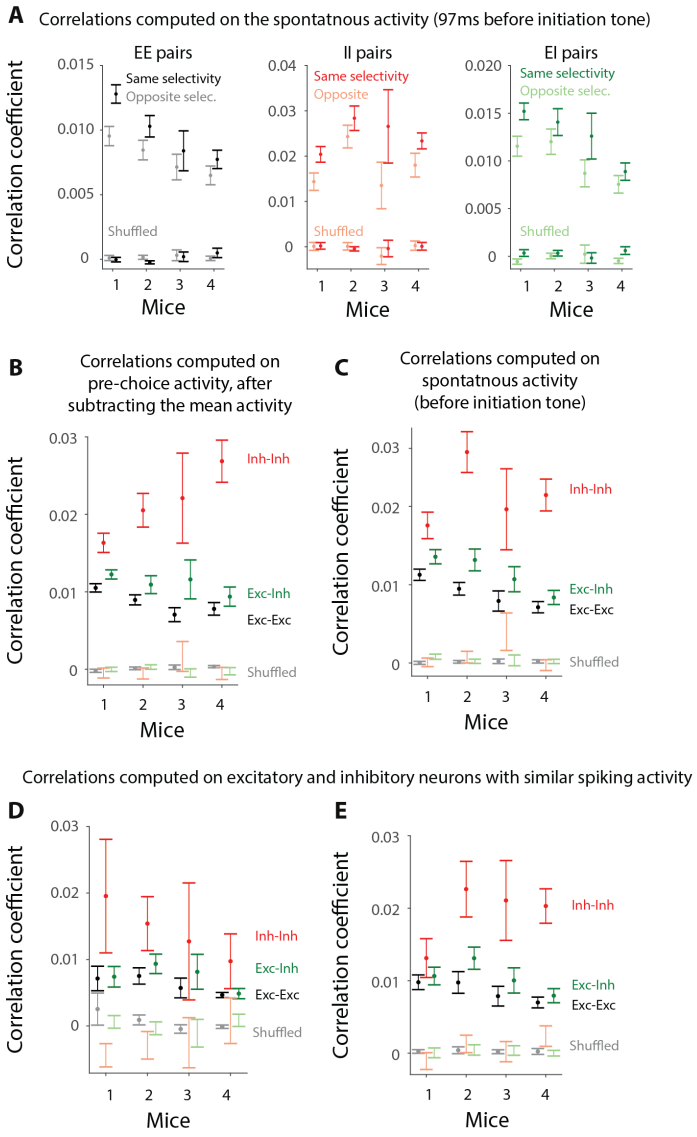

Figure S10. Related to Figure 6. Higher noise correlations between neurons with similar choice selectivity. Also, inhibitory neurons are more strongly correlated.

A, Noise correlations between neurons with the same choice selectivity (dark colors) vs. those with opposite choice selectivity (dim colors), for pairs of excitatory neurons (left), pairs of inhibitory neurons (middle) or excitatory-inhibitory pairs (right). Signal correlations were not present because correlations were computed 0-97 ms before the trial initiation tone, when the stimulus is not present, and the activity is spontaneous. **B,** Noise correlations were much stronger for inhibitory-inhibitory pairs (red) than excitatory-excitatory pairs (black), and had intermediate values for excitatory-inhibitory pairs (green). Correlations are computed on 0-97ms before the choice after subtracting off the mean choice activity, hence removing the signal correlations. **C,** Same as B but for the time period 0-97 ms before the trial initiation tone (i.e. the spontaneous activity). **D,E,** same as in B,C, except correlations were computed only on those excitatory and inhibitory neurons with the same median spiking activity (Methods).

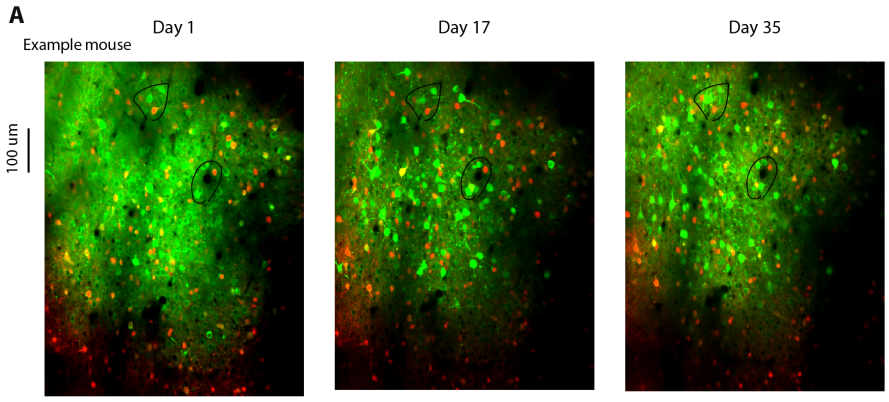

Figure S11. Related to Figure 8. The same field of view was imaged during learning.
A, Field of view from three example sessions of a mouse: 1st days of imaging (left), a middle imaging session (middle), and last day of imaging (right). Left to right panels span 60 days, out of which 35 days were experimental days. Black circles mark example areas that can be easily matched among the sessions. Each panel is an average image of all the frames imaged in the session. Green and red (bleedthrough corrected) images were merged.

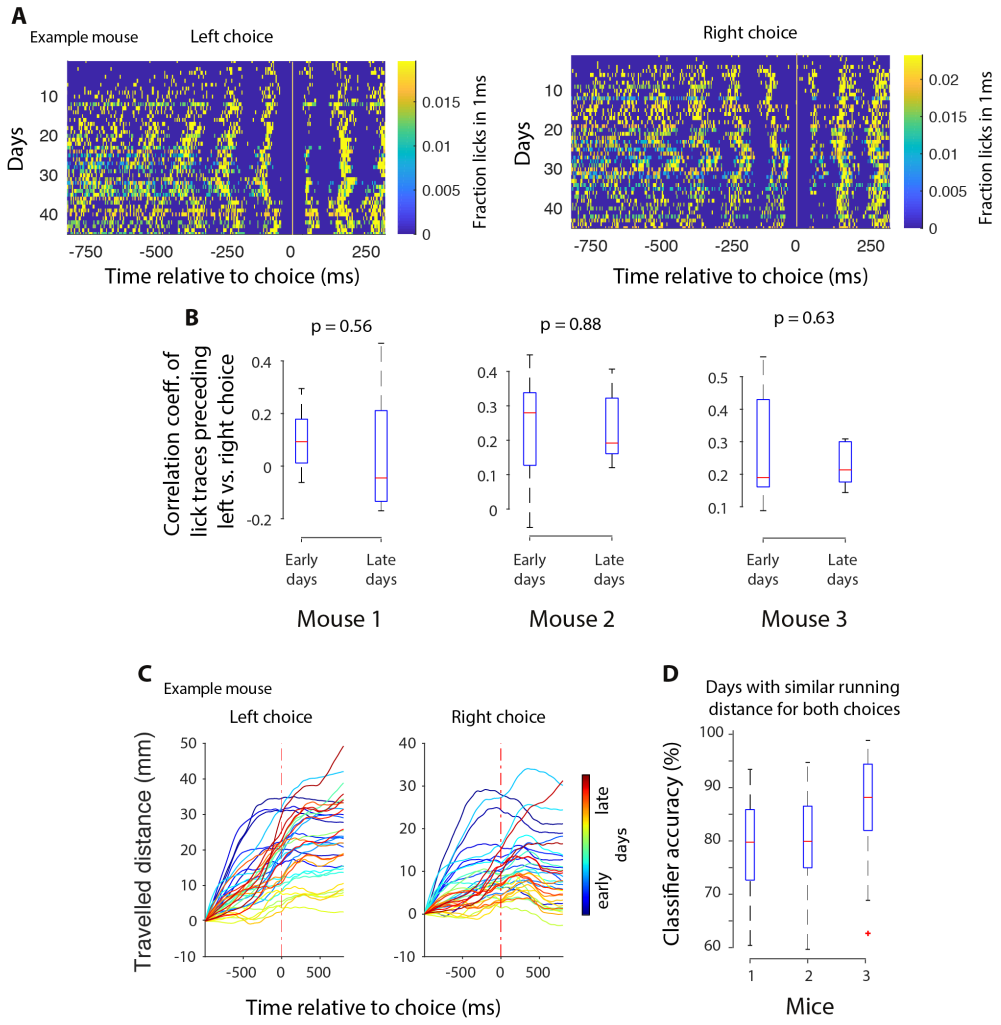

Figure S12. Related to Figure 8. Further analysis of learning-induced changes in the population activity: changes in licking and running movements are unlikely to account for improved classifier accuracy during learning.

A, Licking was similar in advance of high rate vs. low rate choices, both early and late in training. Licks that occur before the choice (vertical line at 0) are to the center waterspout, and licks that occur after the choice are to the side waterspouts; example mouse. **B,** Each plot shows the Pearson's correlation coefficient between licking patterns, to the center waterspout, preceding left and right choices, calculated 250ms before the choice. These correlations were typically similar for early vs. late training days, indicating that animal's licking pattern preceding left vs. right choices did not change drastically over the course of learning. **C,** Distance that the animal travelled during the decision (as measured by the rotary encoder on the running wheel) was similar in advance of left vs. right choices; example mouse; each line represents a session (cold colors: early sessions; hot colors: late sessions). **D,** Classifier accuracy (0-97 ms before the choice) of the full population was high even when the analysis was restricted to sessions in which the distance travelled was not significantly different (t-test, P>0.05; time 0-97 ms before the choice) for left vs. right choices. This analysis was necessary because for some mice in some sessions, there were idiosyncratic differences between the distances travelled in advance of left vs. right choices. In (B) and (D), median (red horizontal line), inter-quartile range (blue box), and the entire range of data (dashed black lines) are shown. There is a single red '+' at the bottom of mouse 3. What is the story there?

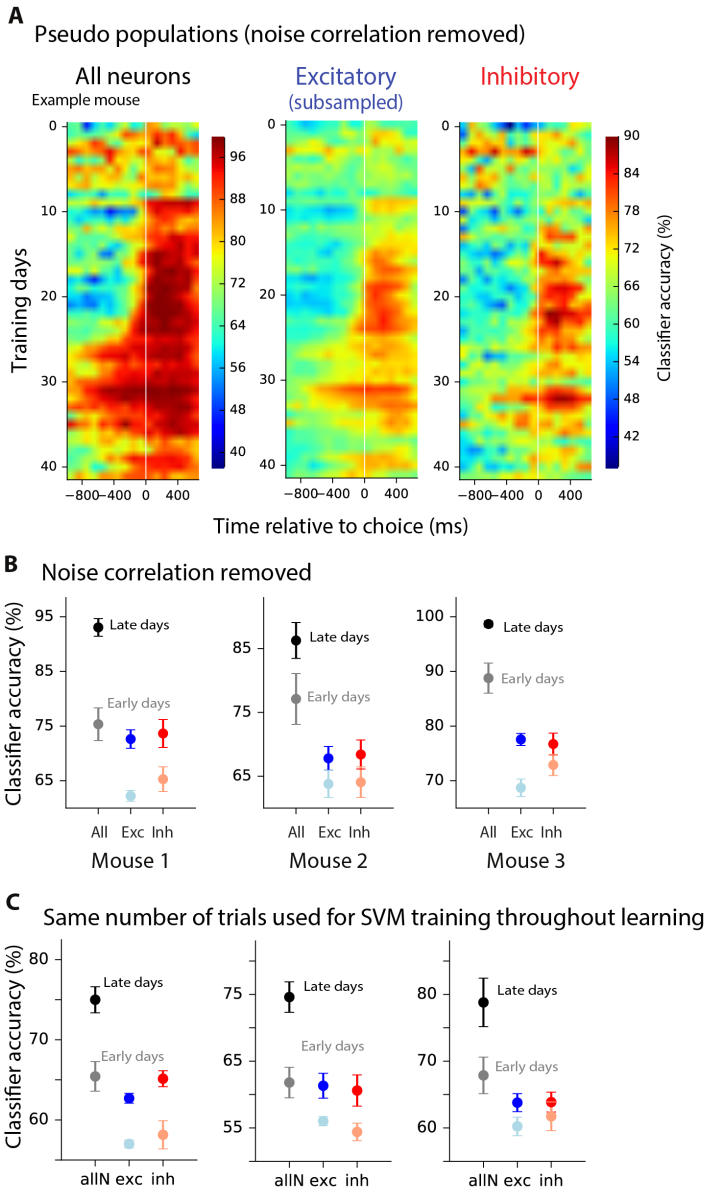

Figure S13. Related to Figure 8. Further analysis of learning-induced changes in the population activity: the reduction in noise correlations is insufficient to account for the improved classification accuracy during learning. Instead, the improvement can be explained by an increase in the fraction of significantly choice-selective neurons. A, Classification accuracy for each training session (average of cross-validation samples), for all neurons (left), subsampled excitatory (middle), and inhibitory neurons (right); example mouse. White vertical line: choice onset. This format is the same as Figure 8A, but here the noise correlations are removed by making pseudo populations (similar procedure as in Figure 7). **B,** Summary of each mouse, showing classification accuracy averaged across early (unsaturated colors) vs. late (saturated colors) training days, at 0-97ms before the choice. As in (A), data are based on pseudo-populations in which the noise correlations are removed. The learning-induced improvement in the classifier accuracy in pseudo populations indicates that reduced noise correlations (Figure 8F) cannot solely account for the enhanced classifier accuracy in the population during learning (Figure 8A). **C**, Equal trial numbers were used to train the choice classifier in every session to control for any effects of trial numbers on classifier accuracy. An increase in classifier accuracy is still observed as a result of learning. Classifiers were trained only on correct trials.

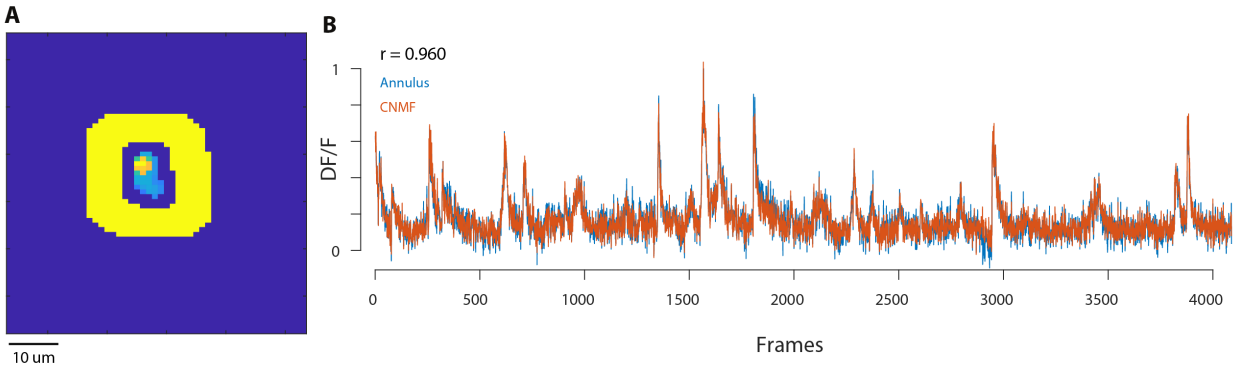

Figure S14. Related to Methods section "Neuropil Contamination removal". Removing neuropil contamination with CNMF or manually using an annulus leads to the same results.

A, An example spatial component in the FOV and its surrounding annulus (yellow). **B,** ΔF/F trace for the same component obtained by manually subtracting the neuropil activity averaged over the annulus region (blue trace) or by using the output of the CNMF processing pipeline (red trace). The two traces look nearly identical as also demonstrated by their high correlation coefficient (r = 0.96; the traces are not denoised). These results demonstrate the ability of the CNMF framework to properly capture neuropil contamination and remove it from the detected calcium traces.

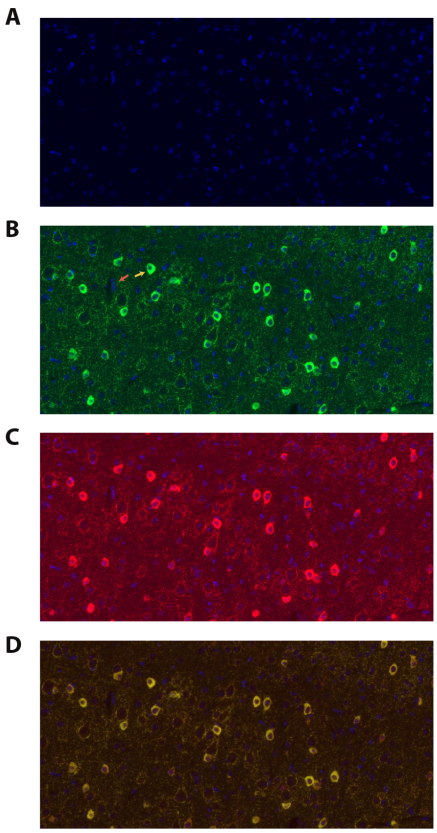

Figure S15. Related to Methods section "Immunofluorescence staining for TdTomato and GABA". GABA and tdTomato are highly colocalized.

A, DAPI staining shows nuclei in blue. **B,** Overlap of GABA antibody (green) and DAPI (blue). Neurons with the green stain surrounding the nucleus are inhibitory (example neuron: yellow arrow). Neurons that lack the green stain are excitatory (example: red arrow). **C,** Overlap of tdTomato antibody (red) and DAPI (blue). Note that GABA (B) and tdTomato (C) are almost completely colocalized. **D,** Overlap of DAPI (blue), GABA (green), and tdTomato (red). Yellow indicates colocalization of GABA (green) and tdTomato (red). 98.2% of neurons that express tdTomato also express GABA. All neurons that express GABA also express tdTomato.