A

Muscimol, 3.4 licks/s

Saline, 3.1 licks/s

5 s

B

Field of view | Activity map | ROIs

PYR (CaMKIIα)

L P
A M

100 µm

Activity (a.u.)
1
0

SST⁺

A L
M P

PV⁺

M
P A
L

VIP⁺

A L
M P

C

Task-modulated neurons (%)

100

80

60

PYR
SST⁺
PV⁺
VIP⁺

1

Figure S1. **Involvement of dmPFC in the go/no-go sensory discrimination task, related to Figure 1**

(A) Licking of a mouse with free access to water after bilateral injections of muscimol (top, green) or saline (bottom, gray). Muscimol caused no significant motor deficits as measured by rate of licking. Each vertical line represents a lick.

(B) Left: example fields of view for each of the neuron subtypes (labels on the left). Middle: pixel-wise activity maps for the recordings on the left (see Extended Experimental Procedures). Right: regions of interest (ROIs) were selected based on the activity map (manually drawn) for further analysis.

(C) Percentage of task-modulated neurons ($p < 0.01$, 3-way ANOVA) for each cell type (out of 816 PYR, 255 PV$^+$, 423 SST$^+$ and 477 VIP$^+$ neurons).
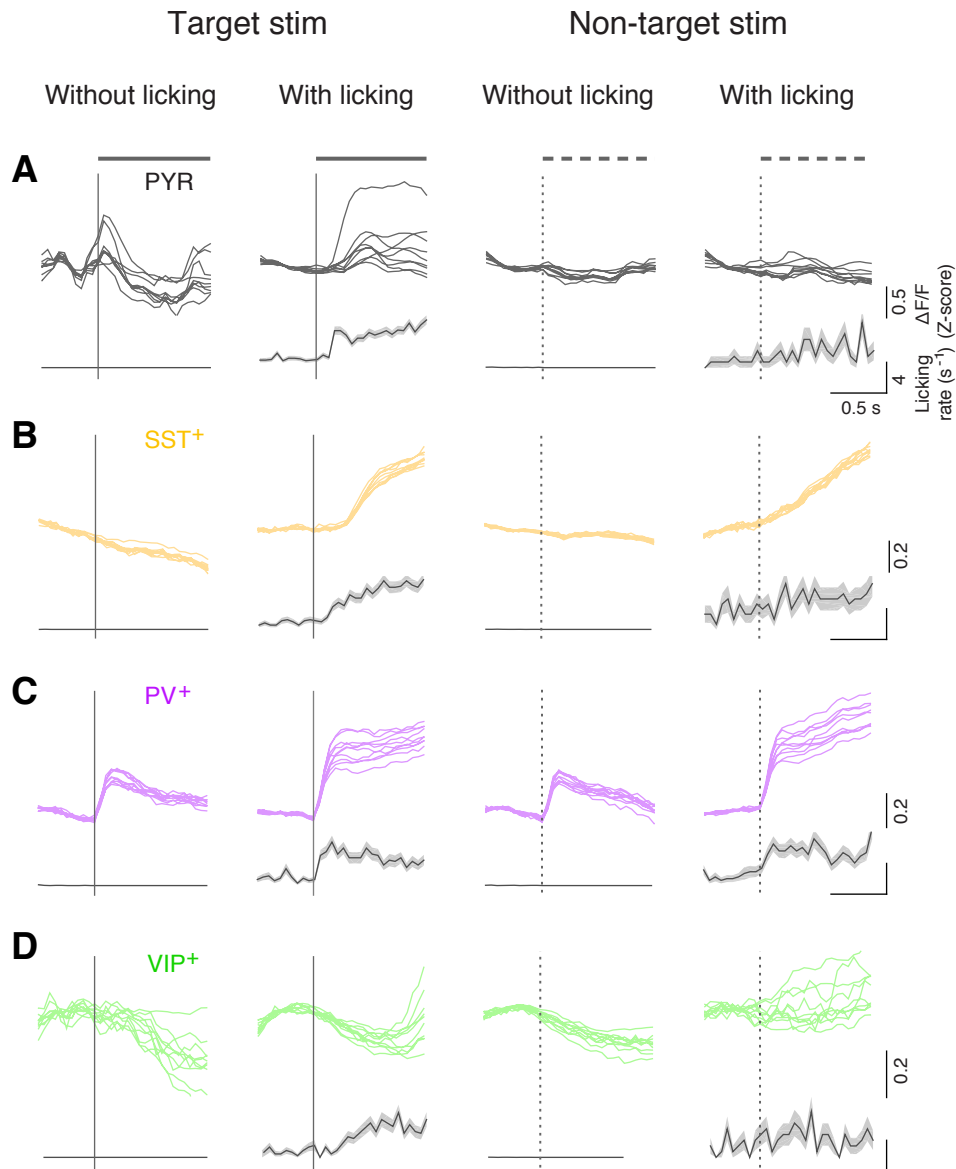
Figure S2. **Apparent responses of PYR and SST$^+$ neurons to target stimuli were in fact associated with licking, related to Figure 3**

(A – D) Trial-averaged ΔF/F traces at presentation of the target or non-target stimulus, in trials with or without licking within 1 s of stimulus presentation. Each plot contains the trial-averaged responses of 10 example neurons recorded simultaneously in the same field of view (top, thin colored lines) with corresponding licking histograms (bottom). Neurons are the same as those shown in Figures 3F-3I. Shaded areas, ± SEM.
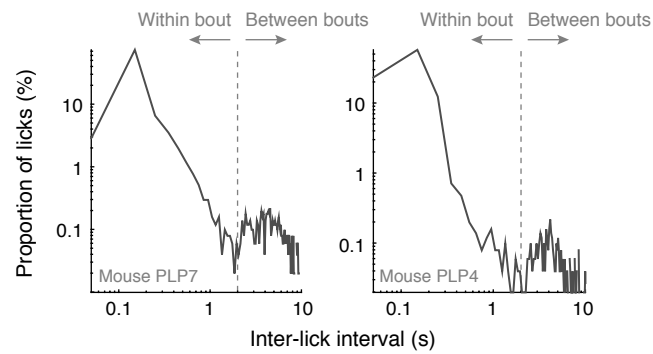
Figure S3. **Licking is organized in bouts, related to Figure 4**

Inter-lick interval distribution from two example sessions of two different mice.

The clear bimodal distribution indicates that licking is organized in bouts, with

short inter-lick intervals within each bout (larger peak, < 2 s) and long intervals

between bouts (smaller peak, > 2 s). Dashed line, criterion interval used to
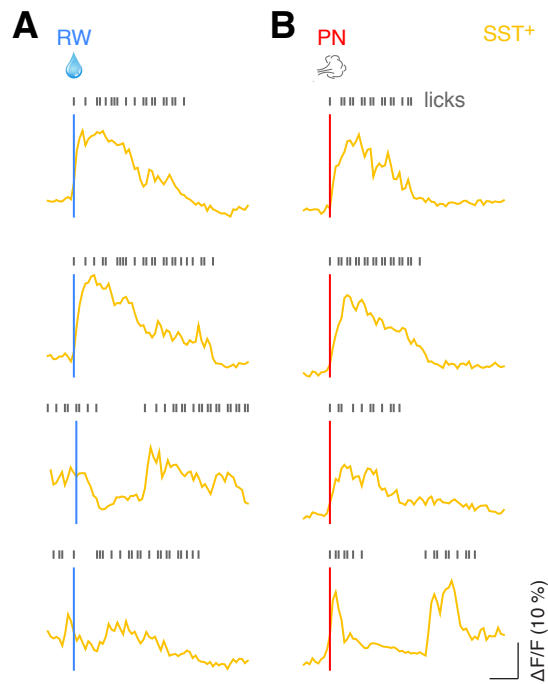
identify licking bouts.

Figure S4. **SST+ neuron activity after trial outcome is associated with licking, related to Figure 5**

(A – B) Shown are ΔF/F traces of a representative SST+ neuron at several example RW (A) and PN (B) trials. Note that the activity was strongly correlated with licking rather than time-locked to RW or PN.

Figure S5. Differences in post-PN and post-RW prep cue responses are not caused by differences in baseline, related to Figure 6

Responses to prep cue averaged across all task-modulated neurons of each cell type, in post-RW (blue) and post-PN (red) trials. For each neuron, only a subset of the trials was included such that the baseline in post-PN and post-RW trials was matched (see Supplemental Experimental Procedures).

A

Field of view

PYR

SST+

M
P ← → A
L

A
L
M ← → P

100 µm

Activity map

1
0

Activity (a.u.)

1
0

Activity ratio

Activity ratio > 3
Activity ratio > 6

11
1

Activity ratio

13
1

B

Neurons in panel C

PYR

40

Neurons in panel D

SST+

Number of neurons

PV+

VIP+

10   20   30   >30

Activity ratio

C

Activity ratio > 3

0.8

0.6

0.4

0.2

CC (ΔF/F)

100   200   300

D

Activity ratio > 6

0.8

0.6

0.4

0.2

CC (ΔF/F)

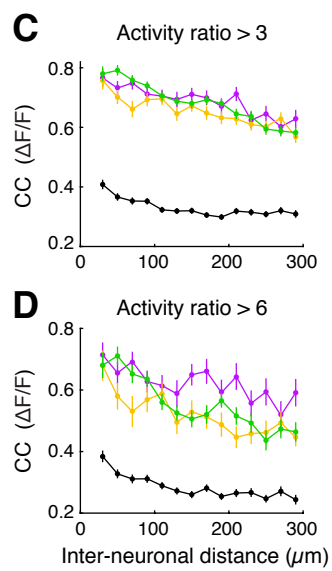0   100   200   300

Inter-neuronal distance (µm)

Figure S6. Higher correlations between inhibitory neurons were not caused by contamination from out-of-focus neuropil fluorescence, related to Figure 7

(A) Example fields of view (top) with corresponding activity maps (middle) for a PYR (left) and PV$^+$ (right) recording. Bottom: the luminance of each ROI represents its activity ratio, defined as the ratio between the mean activity of within-ROI pixels and that of pixels outside of all ROIs. ROIs outlined in red have activity ratios between 3 and 6 and those outline in yellow have ratios higher than 6 (panel D).

(B) Distribution of activity ratios for all neurons of each type. The distribution of ratios was comparable across cell types and was not higher for PYR neurons (PYR: 9.1 ± 0.3, SST$^+$: 7.5 ± 0.5, PV$^+$: 18.0 ± 2.9, VIP$^+$: 11.1 ± 0.8, mean ± SEM). Dashed lines indicate the threshold used for the analysis in panel C (activity ratio > 3) and D (activity ratio > 6).

(C and D) CC between ΔF/F traces of each cell pair vs. distance of the pair, averaged across all recordings for each cell type. (C) Only neurons with activity ratios > 3 were included in this analysis (PYR: n = 9,630 neuron pairs; SST$^+$: n = 3,309; PV$^+$: n = 1,431; VIP$^+$: n = 4,683). (D) Only neurons with activity ratios > 6 were included in this analysis (PYR: n = 5,429 neuron pairs; SST$^+$: n = 1,060; PV$^+$: n = 668; VIP$^+$: n = 1,714). Error bars, ± SEM.
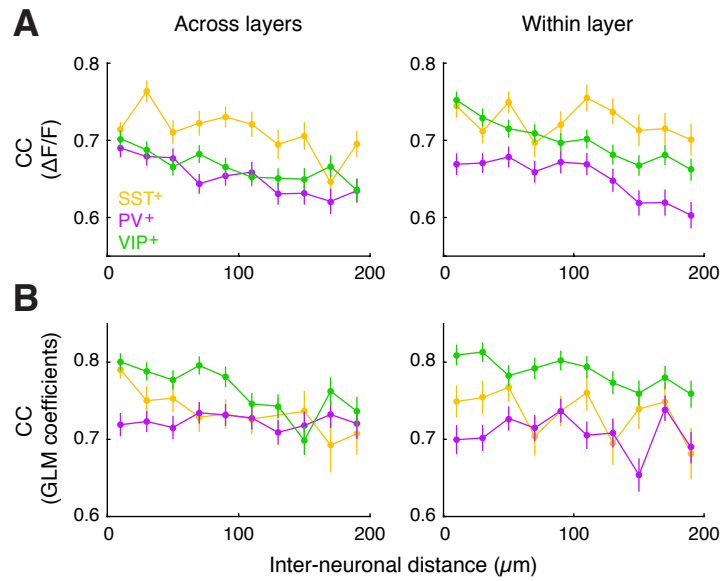
Figure S7. **Between-layer and within-layer CCs for inhibitory neurons, related to Figure 8**

(A) CC between ΔF/F traces of each neuron pair vs. distance across (perpendicular to) layers (left) or distance within (parallel to) each layer (right). Across layers, SST[+]: n = 3,647 neuron pairs; PV[+]: n = 1,747; VIP[+]: n = 4,811; within layer, SST[+]: n = 3,101 neuron pairs; PV[+]: n = 1,328; VIP[+]: n = 4,077. Error bars, ± SEM.

(B) CC between GLM coefficients of each neuron pair vs. distance across (perpendicular to) layers (left) or distance within (parallel to) each layer (right). Across layers, SST[+]: n = 3,535 neuron pairs; PV[+]: n = 1,730; VIP[+]: n = 4,196; within layer, SST[+]: n = 2,998 neuron pairs; PV[+]: n = 1,311; VIP[+]: n = 3,627. Error bars, ± SEM.

SUPPLEMENTAL MOVIE LEGENDS

Movie S1. Microendoscopic Ca^{2+} imaging from the mouse dmPFC during goal-directed behavior, related to Figure 1

For each of the cell types, shown are a 20-s segment of an example recording along with ΔF/F traces for 4 ROIs. The position of each ROI is indicated by the circle of corresponding color in the fields of view (~ 500 x 500 μm for all examples).

SUPPLEMENTAL EXPERIMENTAL PROCEDURES

Mouse lines

The following mouse lines were used in the present study: *CAMKIIα-Cre* [Jackson Laboratories, *B6.Cg-Tg(Camk2a-cre)T29-1Stl/J*, stock number 005359], *PV-Cre* (Jackson Laboratories, *B6;129p2-pvalb*^{tm1(cre)Arbr}*/J*, stock number 008069), *SST-Cre* (Jackson Laboratories, *Sst*^{tm2.1(cre)Zjh}*/J*, stock number 013044) and *VIP-Cre* (Jackson Laboratories, *Vip*^{tm1(cre)Zjh}*/J*, stock number 010908). Some of the *PV-Cre*, *SST-Cre* and *VIP-Cre* mice were crossed with loxP-flanked tdTomato mice [Jackson Laboratories, *B6.Cg-Gt(ROSA)26Sortm14(CAG-tdTomato)Hze/J*, stock number 007914] and therefore expressed tdTomato in the cell types expressing Cre.

Surgical procedures

For all invasive procedures the mice were anesthetized with isofluorane (5% induction and 1.5% maintenance) and placed on a stereotaxic frame (David Kopf Instruments). Temperature was kept at 37 $^{\circ}$C throughout the procedure using a heating pad. The mice received 1 – 2 doses of buprenorphine (0.05 mg/kg, the first before surgery and the other 6 – 8 h later if necessary) and supplementary analgesia (meloxicam, 5 mg/kg) if needed.

Animals used in imaging experiments underwent two surgical procedures, the first for headplate implantation and injection of the genetically encoded Ca^{2+} indicator GCaMP6f (Chen et al., 2013b), and the second to implant the GRIN lens. For the first procedure, after asepsis we exposed the

11

skull and removed the overlying connective tissue. A craniotomy of ~500 µm was then performed to target the dorsomedial prefrontal cortex in the right hemisphere (dmPFC, primarily prelimbic area, AP +2.1 mm, ML 0.3 mm, DV 1.6 mm), where we injected 300-500 nL of either AAV1-*flex-GCaMP6f* or AAV1-*syn-GCaMP6f* (Penn Vector Core) through a borosilicate pipette using a microinjector (Nanoject, Drummond Scientific). We waited for at least 10 min after the conclusion of the injection before removing the pipette. The headplate was then affixed to the skull using machine screws and dental cement mixed with carbon powder (Sigma) to make it opaque. The region surrounding the craniotomy was protected with a silicone elastomer (Kwik-Cast – World Precision Instruments).

After a week of recovery from this procedure, the mice underwent the initial stages of behavioral training for approximately 2 weeks (see below). We then interrupted water restriction and performed a second surgical procedure to implant the gradient refractive index (GRIN) lens (Inscopix). Briefly, after performing asepsis and removing the silicone cap we expanded the craniotomy to ~1 mm and slowly aspirated the superficial brain tissue (corresponding to part of area M2 and of the anterior cingulate cortex) using a blunt 27G needle connected to a vacuum pump, avoiding blood clotting by constantly irrigating the tissue with sterile Ringer's solution. We ensured that the aspirated area was just wide enough to fit the lens (i.e. 1 mm in diameter). We stopped aspirating upon visual identification of white matter in the lateral portion of the craniotomy, corresponding to the forceps minor of the corpus callosum, which at that anteroposterior (AP) coordinate is at the same depth

as the dorsal most portion of the prelimbic area (~ 1.5 mm)(Paxinos and Franklin, 2004) (although in some cases post-hoc histology revealed it to be at the border between the prelimbic area and the anterior cingulate cortex). After ensuring there was no active bleeding, the GRIN lens was gently placed upon the tissue and the space between the lens and the skull was filled with 1.5% agar. The lens was then cemented to the rest of the implant and covered with a plastic cap. The mice recovered from this procedure for at least 3 days and water restriction was reinstated for training.

Once the mice reached criterion performance (see below), we implanted the baseplate for chronic imaging (Ghosh et al., 2011; Ziv et al., 2013). Since this was not an invasive procedure, the mice were head-fixed and lightly anesthetized with 0.5 – 1% isofluorane. The protective plastic cap was removed to expose the surface of the lens and the miniature microscope attached to the baseplate was lowered with a micromanipulator to the desired focal plane. The baseplate was then cemented to the rest of the implant, and covered with a protective cap after the microscope had been retracted.

For mice used in pharmacological inactivation experiments, we implanted a headplate as described above along with a bilateral cannula (26G, Plastics One) targeting dmPFC using the same AP and mediolateral (ML) coordinates as described above, and at a depth of 1.1 mm (instead of 1.6), since the internal cannula used for injection protruded 0.5 mm beyond the external one.

Histology and immunohistochemistry

13

We performed histology to confirm the location of the implanted GRIN lens and cannulas. The mice were transcardially perfused under deep anesthesia with ~5 mL of phosphate-buffered saline (PBS) followed by ~5 mL of 4% w/v paraformaldehyde (PFA). The brain was removed from the skull and post-fixed in 4% w/v PFA at 4 °C for 4 – 12 h, followed by cryoprotection with a 30% w/v sucrose solution for 3 – 7 days. The brain was then sectioned in 30-µm thick sagittal slices using a cryostat (Leica Microsystems). The slides were mounted with VECTASHIELD® mounting medium with DAPI (Vector Laboratories) and examined with a Microphot-SA fluorescence microscope (Nikon Corp.).

To confirm cell-type specificity of GCaMP6f expression we performed immunohistochemistry for CaMKIIα, PV, SST and VIP. Initial sample processing and brain slicing were carried out as above. After slicing, samples were thoroughly washed with 0.5% v/v TritonX-100 in PBS and incubated for 2 h with blocking buffer (2% w/v normal goat serum in PBS). The buffer was washed out with PBS and the samples were incubated overnight at 4 °C with the following primary antibodies: anti-PV (PVG-214, Swant; 1:1,000), anti-SST (MAB353, Millipore; 1:200), anti-VIP (20077, ImmunoStar; 1:500) or anti-CAMKIIα (sc-13141, Santa Cruz Biotechnology; 1:50). After incubation slices were washed with PBS and incubated for 2 h at room temperature with the following secondary antibodies: for PV, VIP, and SST, Alexa594-conjugated anti-rabbit IgG (Invitrogen, 1:1,000 for PV and VIP and 1:200 for SST); for CAMKIIα, Alexa568-conjugated anti-mouse IgG (Invitrogen, 1:200). The protocol and antibody concentrations are compatible with previous reports

(Lee et al., 2012). Images were acquired with a 20x water immersion lens or a 40x oil immersion lens in a Zeiss LSM 780 Confocal Microscope (Carl Zeiss Inc.) and processed in ImageJ (http://rsbweb.nih.gov/ij/). Cells were manually counted in ImageJ.

Pharmacological inactivation experiments

We implanted cannulas for bilateral mPFC pharmacological inactivation in 5 *PV-Cre* animals. Thirty minutes before each behavioral session, we inserted an internal cannula to target the dmPFC and injected 0.375 μL of either the GABA$_A$ agonist muscimol (Sigma, 1 μg/μL) or saline at a rate of 0.25 μL/min. We waited for 5 min after the end of the injection to allow diffusion of the drug before removing the cannula. The dose and injection rate are within the range reported in other studies (Harvey et al., 2012; Muller et al., 1997; Shah et al., 2004). To ensure that the behavioral deficits resulting from the infusion were not of a motor nature, at the end of the session we measured the animals' lick rate by allowing the mice to drink water that was delivered whenever licks were detected.

Selection criteria for imaging sessions

We imaged from a total of 6 PV$^+$, 4 SST$^+$, 4 VIP$^+$ and 6 PYR animals. Three of the PYR animals were *PV-tdTomato* mice injected with AAV1-*syn-GCaMP6f*. We consider those to be putative PYR cells, since this class comprises at least 80% of cortical neurons (Markram et al., 2004). Moreover, we observed that the *syn-GCaMP6f* virus had little tropism for the inhibitory interneurons,

since there was very little overlap between GCaMP6f and tdTomato expression when we injected AAV1-*syn-GCaMP6f* into either *PV-tdTomato* [n = 2 mice, 9 fields of view, 1.6% (19/1180) of GCaMP6f-expressing cells were PV⁺, 3.6% (19/521) of PV⁺ cells were infected] or *SST-tdTomato* mice [n = 2 mice, 9 fields of view, 4.2% (44/1041) of GCaMP6f-expressing cells were SST⁺, 16.8% (44/252) of SST⁺ cells were infected]. Given that PV⁺ and SST⁺ cells comprise 35% and 30% of all inhibitory interneurons (Pfeffer et al., 2013), respectively, we estimate the proportion of GCaMP6f-expressing cells that were PYR neurons in *syn-GCaMP6f*-injected animals to be > 90% (assuming similarly low infection efficacy for other inhibitory interneuron subclasses). Thus, we included these recordings in the PYR neuron group. Our major findings (diversity of PYR neurons and the laminar organization of their response properties) are not affected by inclusion of these animals. We excluded sessions in which the mice did not perform above 65% correct over at least 100 consecutive trials (analyzed with a sliding window). We also included sessions in which the false alarm rate was below 10% and the hit rate was at least 3 times as high. Using these criteria, we had a total of 104 imaging sessions from 20 mice (1 – 10/mouse, mean ± SEM: 5.2 ± 0.5).

Image processing

We corrected for lateral motion using an algorithm based on a previous study (Ziv et al., 2013). We created a new image stack by subtracting from each frame a low-pass filtered version of itself (disk filter, radius 20 pixels) to decrease the influence of large fluorescence transients when calculating

correlations. Each frame was then moved by the x-y shift that maximized its cross-correlation with the average of all frames in the stack.

Because the focusing mechanism rotates the objective, it was necessary to register each field of view with a reference image taken on the day of the baseplate implant, for which the anatomical coordinates were known. This was done by searching for the x-y shift, image rotation and dilation that maximized the correlation between the field of view and the reference image, followed by manual verification.

After these initial pre-processing stages we selected regions of interest (ROIs) for further data analysis. We first computed an activity map as follows (Ahrens et al., 2012):

$$m_{x,y} = \left\langle \left\langle \left(\left(f_{x,y}(t) - f_{x,y} \right) \middle/ \left(f_{x,y} + f \right) \right)_w^3 \right\rangle \right\rangle_t$$

Where $m_{x,y}$ is the activity at pixel (x,y), brackets indicate averaging, $f_{x,y}(t)$ is the fluorescence value at frame t, $f_{x,y}$ is the average of $f_{x,y}(t)$ over time, f is the average of $f_{x,y}$ over w, and w is a sliding window of 2 x 2 pixels ($\sim$ 5 x 5 μm). ROIs were then manually selected on the activity map by drawing polygons. For each ROI we extracted a raw fluorescence trace (F_{raw}) as the average across pixels belonging to that ROI for each frame.

To correct for potential contamination from out-of-focus neuropil fluorescence, for each ROI we subtracted from the raw fluorescence (F_{raw}) the average fluorescence over a 20-μm ring surrounding that ROI (F_{np}): $F_{subt}(t) = F_{raw}(t) - cf \times F_{np}(t)$, where cf is a correction factor estimated as the ratio between the average fluorescence over a blood vessel and the neuropil

adjacent to that blood vessel, both subtracted by a DC offset given by the intensity of off-lens pixels. The values thus obtained for *cf* ranged from 0.15 to 0.87 (mean ± SEM: 0.61 ± 0.02), compatible with other reports that used similar methods (Chen et al., 2013b; Kerlin et al., 2010). Our finding that interneurons of the same type were functionally much more similar to each other compared to the PYR population was robust over a range of *cf* values (0 – 1 in 0.2 steps, i.e. from no subtraction to full subtraction of neuropil fluorescence).

Imaging sessions lasted for up to an hour, which typically resulted in some bleaching of the Ca^{2+} indicator. To correct for such decreases in baseline fluorescence we subtracted slow fluctuations in baseline according to the expression: $F_{corrected}(t) = F_{subt}(t) - G(t) + <F_{subt}(t)>_t$, where brackets indicate time average over the entire recording session and $G(t)$ is the average of $F_{subt}(t)$ over a 300-s sliding window.

GLM

We used a generalized linear model (GLM) to parametrically describe task-related activity of each cell. Each Z-scored ΔF/F trace was binned at 100 ms (2 frames) and modeled as the linear combination of various task events at several delays with respect to event onset:

$$\widehat{\Delta F / F(t)} = \beta_0 + \sum_{n=-2}^{7} \beta_n^{pc} F_{t-n}^{pc} + \sum_{n=-2}^{5} \beta_n^{ts} F_{t-n}^{ts} + \sum_{n=-2}^{5} \beta_n^{nts} F_{t-n}^{nts} + \ldots$$

$$\sum_{n=-6}^{29} \beta_n^{lon} F_{t-n}^{lon} + \sum_{n=-6}^{79} \beta_n^{loff} F_{t-n}^{loff} + \sum_{n=-6}^{14} \beta_n^{lk} F_{t-n}^{lk} + \ldots$$

$$\sum_{n=-2}^{47} \beta_n^{rw} F_{t-n}^{rw} + \sum_{n=-2}^{107} \beta_n^{pn} F_{t-n}^{pn} + \sum_{n=-2}^{27} \beta_n^{cr} F_{t-n}^{cr} + \sum_{n=-2}^{27} \beta_n^{ms} F_{t-n}^{ms}$$

Where β_0 is the bias term, β_n^x indicates model coefficient for task event x at temporal delay n, F_t^x indicates occurrence of task event x at time t (1 if the event occurred, 0 otherwise), *pc* is prep cue, *ts* is target stimulus, *nts* is non-target stimulus, *lon* is lick-bout onset, defined as a lick occurring at least 2 s after the previous one, *loff* is licking offset, defined as the last lick in a bout not followed by another lick within 2 s, *lk* are mid-bout licks, *rw* is reward, *pn* is punishment, *cr* is correct rejection and *ms* is miss (note that for correct rejections and misses the event timing is defined as the time of auditory stimulus offset). Modeling licking onset and offset separately yielded significantly better predictions, since for many neurons responses to licking onset and offset were much stronger than for licks occurring in the middle of a bout (Figures 4A and S3). For each regressor, F is 1 only at onset of the corresponding event (duration, 1 frame). We chose the number of time delays for each regressor so that different regressors are non-overlapping (except for licks) while spanning the entire trial and inter-trial periods. The baseline for each regressor was defined as the coefficients at $n < 0$. For lick regressors (*lon*, *loff* and *lk*), we considered n = -5 and -6 as baseline, since activity often

started increasing a few hundred ms before lick onset. The final model had a total of 389 coefficients plus the general bias term.

GLMs were fit using ridge regression following a procedure based on a previous study (Huth et al., 2012). We first set aside data from 20% of the trials for testing. The first step was to estimate the regularization parameter λ. The remainder 80% of the data for model fitting were further divided into a training set (80%) and a validation set (20%). The model was fit using a range of λ values. Model performance was measured using the validation set by computing the correlation coefficient (CC) between the predicted and measured activity. We repeated this procedure 10 times, selected the value of λ that maximized CC for the validation set, and used it to fit the combined training and validation data set. The final model performance was evaluated by the CC between the measured and predicted responses in the test data set. To assess statistical significance of model predictions, we bootstrapped the test data set 5,000 times to obtain a distribution of CC values and calculated a p value as the proportion of iterations with CC $\leq$ 0. Neurons with $p < 0.05$ were considered to have significant GLM fits.

Spatial organization analysis

To analyze the spatial distribution of responses, we calculated two types of correlation coefficients (CC) between pairs of neurons: between their whole $\Delta F/F$ traces binned at 100 ms or between the baseline-subtracted GLM coefficients of neurons with statistically significant fits. We then plotted the CCs against the inter-neuronal distance (distance between the ROI centroids),

binned at 20 μm. For this analysis we only included neuron pairs with distances up to 300 μm, since only a small minority of fields of view had pairs separated by longer distances. The first bin (0 – 20 μm) was excluded because there were very few neuron pairs with distances < 20 μm. For the analysis in Figures 8 and S7, we estimated the absolute laminar position (i.e. distance from dmPFC pial surface) of each neuron by adding its relative mediolateral position in the field of view to 1) the distance between the medial border of the field of view and the medial border of the lens (determined by registration with a reference image collected on the day of the baseplate implant, see above); and 2) a fixed offset for each animal, corresponding to the distance between the medial border of the lens and the cortical surface, estimated histologically after the imaging experiments. Amplitude of the response to the prep cue was measured as the average GLM coefficient for prep cue between 0 and 0.8 s after baseline subtraction.

Analysis of activity modulation by action outcome

To analyze the dependence of the response to prep cues on previous trial outcome (Figure 6), we first divided trials into post-PN and post-RW groups, excluding trials with licking between 0 and 1 s from prep cue onset to avoid contamination by activity associated with licking. To compare the responses in these two groups of trials, it was necessary to perform a baseline correction. Due to the slow dynamics of Ca^{2+} activity, the baseline often exhibited a significant slope. We estimated this slope by fitting a straight line to the baseline period (0 - 500 ms before prep cue onset) and extrapolated it to 1 s after the cue onset. The baseline-corrected response was computed as the

difference between the measured response and the extrapolated baseline. To control for the possibility that the different responses to prep cue in post-RW and post-PN trials were caused by differences in the pre-trial baseline activity, for each cell we selected a subset of post-PN and post-RW trials in order to match the average baseline in the two conditions. The result was qualitatively similar to that based on the entire set of trials (Figure S5).

SUPPLEMENTAL REFERENCES

Ahrens, M.B., Li, J.M., Orger, M.B., Robson, D.N., Schier, A.F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. Nature *485*, 471–477.

Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature *484*, 62–68.

Huth, A.G., Nishimoto, S., Vu, A.T., and Gallant, J.L. (2012). A continuous semantic space describes the representation of thousands of object and action categories across the human brain. Neuron *76*, 1210–1224.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Interneurons of the neocortical inhibitory system. Nat. Rev. Neurosci. *5*, 793–807.

Muller, J., Corodimas, K.P., Fridel, Z., and LeDoux, J.E. (1997). Functional inactivation of the lateral and basal nuclei of the amygdala by muscimol infusion prevents fear conditioning to an explicit conditioned stimulus and to contextual stimuli. Behav. Neurosci. *111*, 683–691.

Paxinos, G., and Franklin, K.B.J. (2004). The Mouse Brain in Stereotaxic Coordinates (San Diego: Academic Press).

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

Shah, A.A., Sjovold, T., and Treit, D. (2004). Inactivation of the medial prefrontal cortex with the GABAA receptor agonist muscimol increases open-arm activity in the elevated plus-maze and attenuates shock-probe burying in rats. Brain Res. *1028*, 112–115.

Ziv, Y., Burns, L.D., Cocker, E.D., Hamel, E.O., Ghosh, K.K., Kitch, L.J., Gamal, El, A., and Schnitzer, M.J. (2013). Long-term dynamics of CA1 hippocampal place codes. Nat. Neurosci. *16*, 264–266.