

Supplementary note 1

Quantitative measure of optical tagging (Stimulus-Associated spike Latency Test, SALT)

We developed a statistically based, automated, unsupervised method for the detection of light-responsive neurons, which we termed Stimulus-Associated spike Latency Test (SALT).

We first describe the motivation to develop a robust, objective statistical test for the problem. Many naïve and simple, albeit arbitrary and not statistically grounded measures would identify most of the cells that appear light-activated on peri-stimulus time histograms (PSTHs). However, some cells would be inevitably misclassified, which is particularly problematic when identifying rare neuron-types that are outnumbered several fold by neighbouring pyramidal cells. We illustrate some traps of these approaches below.

Directly light-activated neurons could be identified based on a reliability threshold for spikes occurring within an arbitrary window of stimulation. High firing rate neurons may pass this reliability criterion without light-activation (see example in left panel of Supplementary Fig. 10c; compare its reliability value to those of Supplementary Fig. 10a and b), resulting in false positives. Moreover, light-related inhibition and quick rebound (see example in right panel of Supplementary Fig. 10c) can also result in false detection. This reliability-based approach can also lead to many false negatives in the

context of *in vivo* recordings because at high laser intensities spike shape can change due to the activation of other neural elements (e.g. cell bodies and fibres of passage)³⁶. For this reason, we applied relatively low intensities compared to *in vitro* studies, which still detectably activated nearby neurons, but did not cause a concurrent activation of a large number of units. Lower laser intensities, however, result in lower reliability (see Supplementary Fig. 10f), making it difficult to set a non-arbitrary threshold without the peril of false negatives. Supplementary Fig. 10a, b shows examples of neurons with low reliability that appear light-activated in PSTHs and can still be detected by statistical measures. (Please note that for this reason, we also employ a spike shape correlation measure to ensure that our spike sorting was not compromised due to high laser intensities. This is a complementary test as it pertains to spike sorting, whereas SALT tests light effects assuming that spike sorting is correct.)

Using a statistical comparison between baseline firing rate and light-induced firing can eliminate the arbitrary nature of the reliability threshold. This will ameliorate the problem of higher firing rate cells (Supplementary Fig. 10c, left panel) and false negatives with low reliability (Supplementary Fig. 10a). Nevertheless, the arbitrary choice of window becomes an issue for high firing tagged neurons that show an inhibited period after the light-induced spikes, a common type among fast spiking interneurons (Supplementary Fig. 10b).

Another straightforward idea is to identify neurons firing after light pulses at a low latency and jitter (Supplementary Fig. 10e). This method would also detect neurons that show fast light-induced inhibition (Supplementary Fig. 10d) that leads to low latency and jitter. Therefore while all tagged neurons show low latency and jitter (Supplementary

Fig. 2d), the reverse statement is not true, resulting in imperfect separation on the jitter versus latency scatter plot (Supplementary Fig. 10e).

What feature could best identify directly light-activated neurons? We believe the answer is significant firing with a specific latency following the light stimulus. Therefore the latencies of first spikes after light should significantly differ from such latencies relative to any arbitrary pre-stimulus time points, a criterion that differentiates inhibited neurons with (Supplementary Fig. 10b) or without light-evoked spikes (Supplementary Fig. 10d). Our method is a simple quantification of this property: a statistical test for the null hypothesis that light does not change firing latencies. This method identified all neurons that appeared activated based on visual inspection of the PSTHs and it did not identify units that were obviously non-activated (Supplementary Fig. 10c, d). This method has several advantages:

- 1) It eliminates the arbitrary choice of a reliability threshold.
- 2) It eliminates the arbitrary choice of a window for activation. (It should be noted that SALT also uses a window, but this is not critical as the only criterion it has to fulfill is to be long enough to include the first light-induced spike. We use a conservative window of 10 ms.)
- 3) It is a bona fide statistical test, lending objectivity and statistical basis for the identification.
- 4) It allows optimal use of data leading to lower false negative rates.

Note that reliable short latency activation may be indirect in some experimental conditions. For instance identifying densely labeled pyramidal neurons will require

additional criteria, such as reliable sustained light-activation for higher frequency stimulus trains³⁷.

Our implementation of SALT relying on an information theoretic measure is a straightforward mathematical realization of this statistical test but it is certainly not the only possible one. Below we provide a mathematical explanation for the full algorithm.

Stimulus-Associated spike Latency Test (SALT) is a statistical test to determine whether optogenetic activation caused a significant change in the timing of spikes after stimulation onset (Supplementary Fig. 11). Specifically, the distribution of first spike latencies relative to the light pulse, assessed in a 10 ms window after light-stimulation, was compared to epochs of the same duration in the stimulus-free baseline period. The choice of 10 ms window size provided sufficient statistical power without limiting the number of detected neurons (see below). For measuring distance between these distributions, we implemented a modified version of the information theory measure termed Jensen-Shannon divergence³⁸. Jensen-Shannon divergence, also known as information radius, is a symmetrical version of Kullback-Leibler divergence, denoted below $D(P||Q)$ between distributions P and Q ,

$$(1) \quad D_{JS}(P||Q) = \frac{1}{2}D(P||R) + \frac{1}{2}D(Q||R),$$

where

$$(2) \quad R = \frac{1}{2}(P + Q).$$

Let us now define D_{JSm}^2 by

$$(3) \quad D_{JSm}^2(P||Q) = D(P||R) + D(Q||R),$$

which can be expanded to

$$(4) \quad D_{JSm}^2(P||Q) = \sum_{i=1}^N \left(p_i \log \frac{2p_i}{p_i+q_i} + q_i \log \frac{2q_i}{p_i+q_i} \right),$$

where p_i and q_i are the probabilities of the N possible values of the discrete distributions P and Q . The advantages of the D_{JSm}^2 measure over the more generally used Kullback-Leibler divergence are that (i) it is bounded, resulting in meaningful values even when a probability estimate is near zero and (ii) it is a real metric, i.e. not only positive definite but also symmetric and subadditive³⁸. By using this metric, we tested the hypothesis that the post-stimulus spike-latency distribution is different from the set of baseline distributions for low frequency light stimulation (4 or 10 Hz). To this end, we compared the median distance of the post-stimulus distribution and the baseline distributions with the null-distribution of distances among the baseline distributions, computing a P value for significant short-latency light-activation for each neuron. The distribution of P values showed strong bimodality, with a well-separated cluster of near-zero values (Fig. 1d). We used a P value cutoff of 0.01, as it separated the two modes. Note that all directly activated neurons found with this algorithm displayed short latency spikes with low jitter (Supplementary Fig. 2d) and high waveform correlations ($r > 0.85$, Supplementary Fig. 2c). For high P values, i.e. in the absence of light-evoked action potentials, the correlation values provide a baseline for spike waveform variability of single units (top right corner of Supplementary Fig. 2c). We also note that the choice of the 10ms window did not result in the exclusion of neurons with similar features. We

provide Matlab code that implements this algorithm (<http://kepecslab.cshl.edu/software/>).

Supplementary note 2

Linear Discriminant Analysis (LDA)

LDA was used as a visual illustration to demonstrate differences in response profiles. This algorithm is designed to find the combination of features that maximally separates multiple predefined classes.

PETHs ($\pm 1s$) to relevant events were smoothed by a moving average. To visualize differences in response profiles of various cell-types, we embedded each cell in a 7 dimensional space in which each dimension quantified a separate aspect of a neuron's firing rate change around the reward-zone-exit and reward-zone-entry events. The dimensions were (1-4) similarity to average response profiles of *Pv* and *Som* neurons for the two events assessed by linear correlation of PETHs; (5-6) 'suppression time' around the two events, calculated as the temporal extent of below-threshold firing (threshold was determined as 25% of the difference between 1 and 99 percentile of population firing rates); (7) 'suppression index' for reward-zone-entry, difference between firing rate in the middle portion and second half of the PETH. These dimensions were designed to capture different aspects of the PETHs. For instance, suppression time quantifies marked differences in the duration of activated states, whereas the suppression index marks the direction of firing rate change after the relevant event. Neurons were grouped according to cell identity, i.e. *Pv*, *Som* and wide

spiking neurons recorded concurrently with identified cells (tetrode partners). The plane spanned by the centroids of these three groups in the 7-dimensional space was determined by linear discriminant analysis and all neurons of the three populations, as well as all non-tagged cells, were projected onto this plane (Supplementary Fig. 7a).

Supplementary note 3

Information theory analysis

The information gained by knowledge of cell-type (*Pv*, *Som*, narrow-spiking (NS), wide-spiking (WS)) can be computed using the Kullback-Leibler divergence between firing response distribution of a particular class and the full population of recorded neurons: $D_{KL}(P_{identified\ class}||P_{all})$, where $P_{identified\ class}$ denotes the response distribution of neurons with a particular label, whereas P_{all} stands for the prior distribution of all cells, spanning the space of possible responses.

Because the number of neurons per class was low compared to the full dimensionality of the response space, we used a dimension-reduced PETH. First, PETHs restricted to $\pm 0.5s$ windows around reward-zone-exit and reward-zone-entry events were smoothed with a moving average (window size, 110 ms). PETHs were projected to the corresponding group mean (e.g. mean *Pv* PETH) and discretized using the Freedman-Diaconis rule:

$$h = \frac{2 \cdot IQR}{n^{1/3}}$$

where h indicates bin width and IQR denotes interquartile range. Next, the Kullback-Leibler divergence was computed between PETHs of the selected group of neurons and all recorded neurons, quantifying the bit-wise information provided by knowing the group identity (posterior) over the prior distribution for all neurons. As a result of the dimensionality-reduction this information measure provides a lower bound on the total possible information. For calculating information provided by NS or WS neurons the spike width threshold value was optimized to yield the maximal information (NS < 230 μ s; WS > 330 μ s). Bootstrap standard errors were computed by resampling the PETH distributions 1000 times. We report the higher information value for each group, with the PETH triggered on either reward-zone-entry or reward-zone-exit (Supplementary Fig. 7b).

Note that this information theoretic analysis is a quantification of cell-type-specific differences and provides a lower bound on the information gained from cell type identity. Projecting the PETHs onto the group mean can only reduce the information available. This was a necessary step since the sample sizes were low compared to the full dimensionality of the PETH profiles. Because the same algorithm was used for the groups separately, it does not favor the information provided by any of the labels (e.g. *Pv*, *Som*, NS) included in this analysis.

Supplementary note 4

Gap statistic

Hierarchical clustering was performed on response profiles with respect to three behaviourally relevant events (reward zone entry, reward zone exit and cue presentation) using squared Euclidean distance measure, averaged over the three events. Number of clusters was iterated ranging from 1 to 100 and the gap statistic⁴⁷ was calculated to assess the number of clusters naturally present in the data set. Briefly, the gap statistic is calculated as a difference in the logarithm of within-cluster sum of squares between the experimental data and a simulated reference distribution, serving as an error measure. Maximum of the gap statistic in the function of number of clusters ('gap curve') shows the optimal number of clusters; an optimal number of 1 indicates the lack of any natural separation of the data points into distinct clusters. We used uniform reference distributions over a box aligned to the principal components of the data⁴⁷. Null-hypothesis estimates of within-cluster sum of squares were calculated via bootstrap resampling from the reference distributions (20 Monte Carlo replicates per number of clusters).

Supplementary note 5

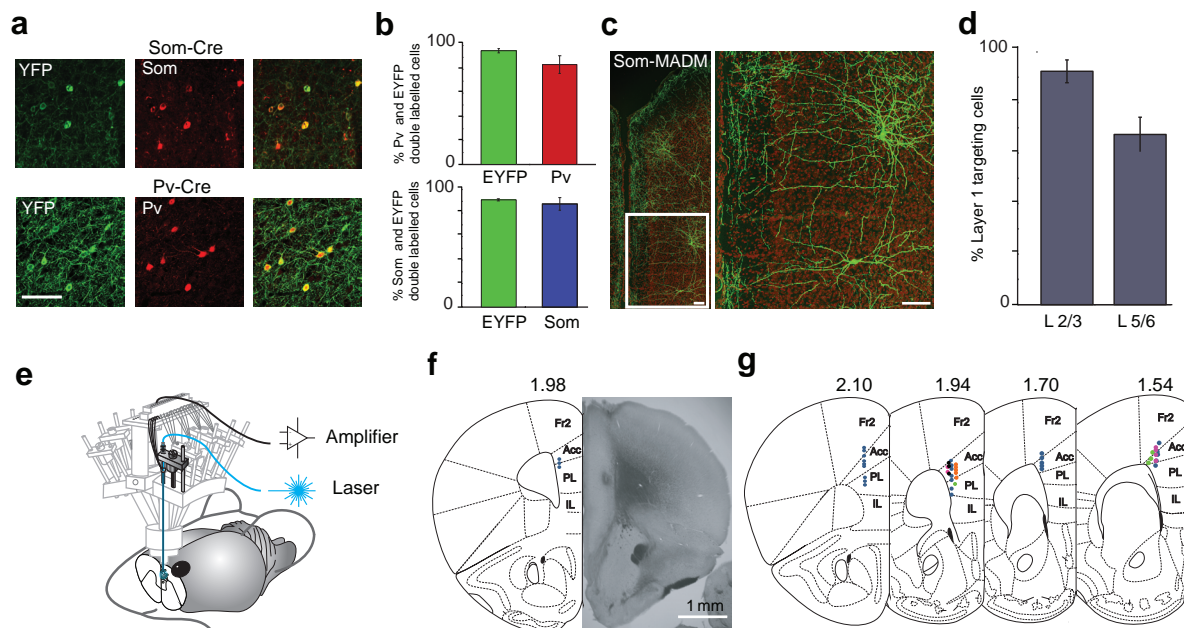
Immunohistochemistry, quantification of colocalization and single cell labeling

Upon conclusion of neuronal data collection, each mouse was anaesthetized with 500µl of 5% urethane and electrolytic lesions were made by passing 10 µA current with

a frequency of 1 kHz for 1-2 seconds through individual leads of each tetrode on which a tagged neuron was recorded. Lesions were confirmed by histological analyses of brain sections and track reconstruction using the Paxinos-Watson mouse brain atlas. We included optically tagged neurons that were mapped to the dorsomedial prefrontal cortex based on the cytoarchitectonic structure of the prefrontal cortex⁴⁹. Also note that according to this cytoarchitecture borders of ACC extends into more dorsal area (corresponding to M2 in Paxinos-atlas).

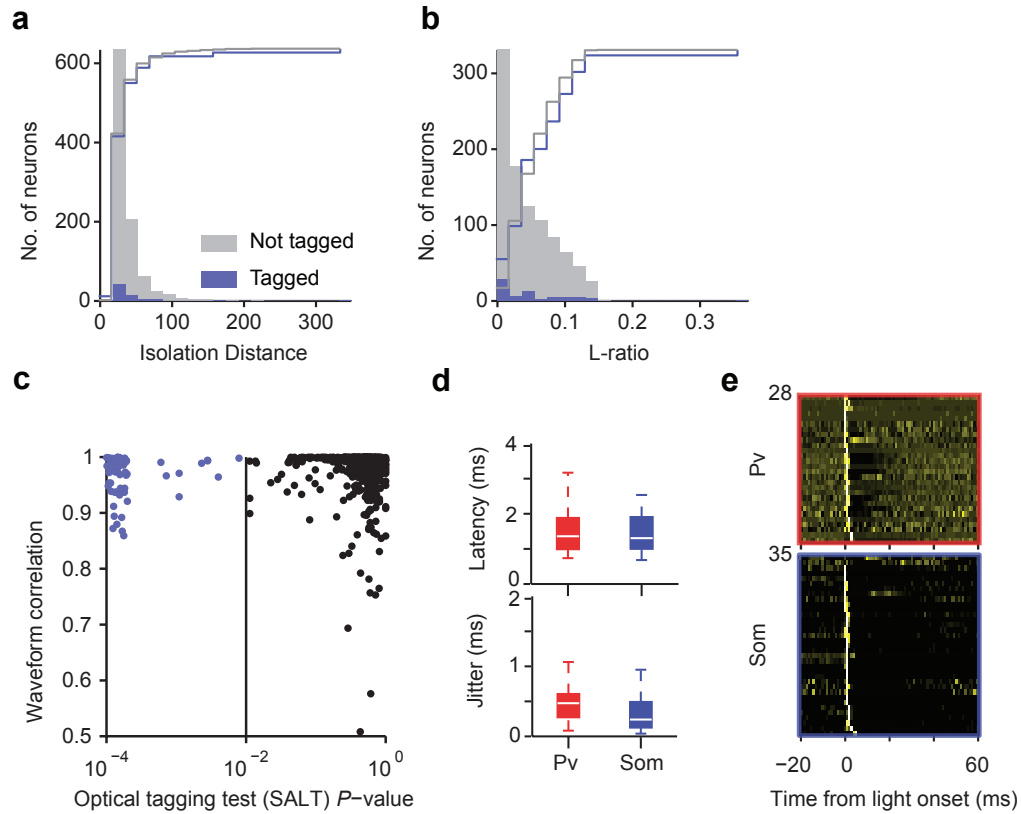
For immunostaining, mice were anaesthetized with 5% urethane and perfused transcardially with 4% paraformaldehyde (PFA) in phosphate buffer pH 7.4 (PBS). The brain was carefully extracted from the skull and incubated overnight at 4°C in 4% PFA, washed in PBS and stored in PBS at 4°C until sectioning. 50 µm coronal brain sections were made using Leica VT1000S vibratome. Sections were incubated with blocking solution (10% Normal Goat Serum and 1% Triton), washed in PBS (3 washes, 5 min each) and incubated overnight at 4°C with primary antibodies diluted in blocking solution. The following primary antibodies were used; anti-*Pv* (mouse monoclonal, 1:1000; Sigma), anti-*Som* (rat polyclonal, 1:300, Chemicon, Temecula), anti-c-fos (goat polyclonal, 1:500, Santa Cruz Biotechnology). Next day, sections were washed in PBS (3 washes, 5 min. each) and incubated for 2hr at room temperature with secondary antibodies at 1:400 dilution (alexa-633 goat anti rabbit, alexa-633 goat anti rat, and alexa-633 donkey anti goat, Molecular Probes). Sections were washed in PBS for 5 min and mounted using Vectashield mounting medium with DAPI (Vector Laboratories). Confocal images were acquired using LSM710 Zeiss microscope at X20 and X65 magnifications. The overlap between *Pv* and ChR2-EYFP, *Pv* and c-fos, *Som* and

ChR2-EYFP markers was quantified around the viral injection site (1mm² - 2mm² area, 50 μm section thickness). The density of EYFP expressing neurons was estimated in *Pv-Cre* and *Som-Cre* animals (3 sections per genotype) around the injection site by counting EYFP positive cells in a 40 μm² area (25 μm section thickness). For sparse *Som* labeling, MADM-11GT / MADM-11TG; *Som-IRES-Cre* line was used, in which GFP was sparsely expressed in *Som* interneurons throughout the cortex⁵⁰. Percentages of layer 2/3 and layer 5/6 dmPFC cells with processes to layer 1 was counted using epifluorescent microscope (Olympus BX51WI).




Supplementary Figure 1 | Genetic targeting of Pv and Som interneurons and localization of recording sites.

a, Virally induced ChR2-EYFP expression (left panel, green) co-localizes with interneuron markers (middle panel, Som or Pv); right panel shows merge. Scale bar, 100 μ m. **b**, Percentage of ChR2-EYFP cells expressing Pv (specificity, 208/220, n=3 sections from 3 animals) and percentage of Pv cells expressing ChR2-EYFP (efficiency, 208/260, n=3 sections from 3 animals) in ACC (top panel). Specificity (170/186, n=3 sections) and efficiency (170/201, n=3 sections) of ChR2-EYFP expression in Som interneurons in ACC (bottom panel). Bar height denotes mean and error bar standard error of the mean. **c**, Sparse labeling using Som-MADM shows projections ascending to the superficial portions of layer 1, which is marked by a sharp decrease in nuclear density. Right panel shows higher magnification image of a region in white rectangle. **d**, Percentage of layer 2/3 (34/38) and layer 5/6 (28/43) Som neurons sending processes to layer 1 (n=12 sections). Bar height denotes mean and error bar standard error of the mean. **e**, Cartoon showing mouse with implanted microdrive in ACC. **f**, Lesion and reconstruction of the position of optically tagged neurons in a representative brain section. **g**, Reconstruction of all anatomical sites where optically tagged Pv or Som neurons were recorded. Each schematic section contains corresponding lesion sites verified from 1-4 mouse brains. Recording sites from different animals are shown in different colors. Anterior - Posterior position (from bregma) of each brain schematics is indicated above in **(f)** and **(g)**.

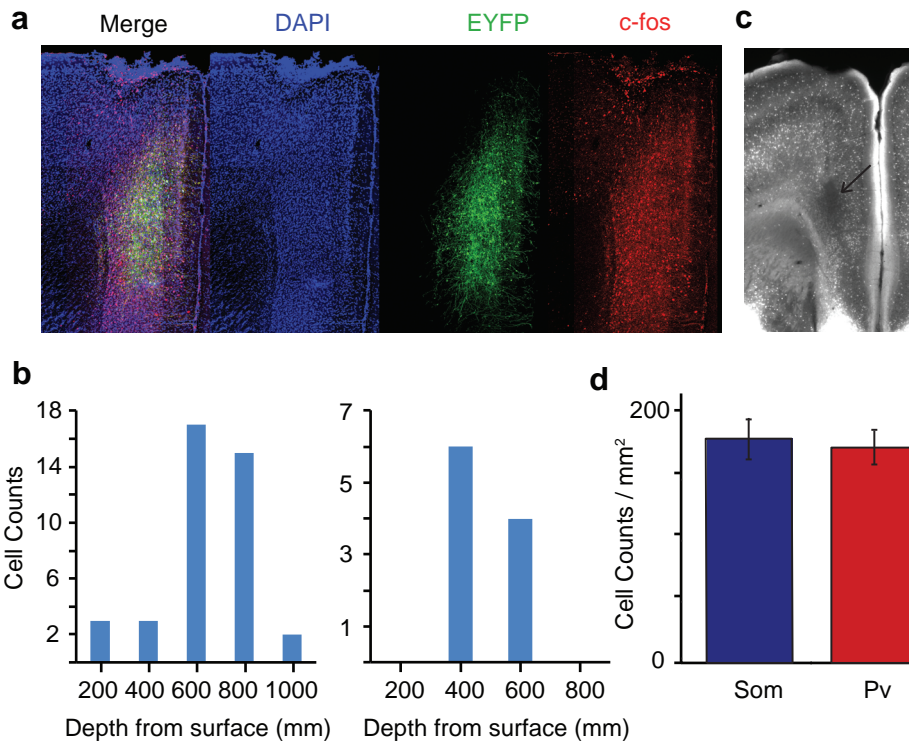


Supplementary Figure 2 | Single unit isolation quality for optically tagged neurons.

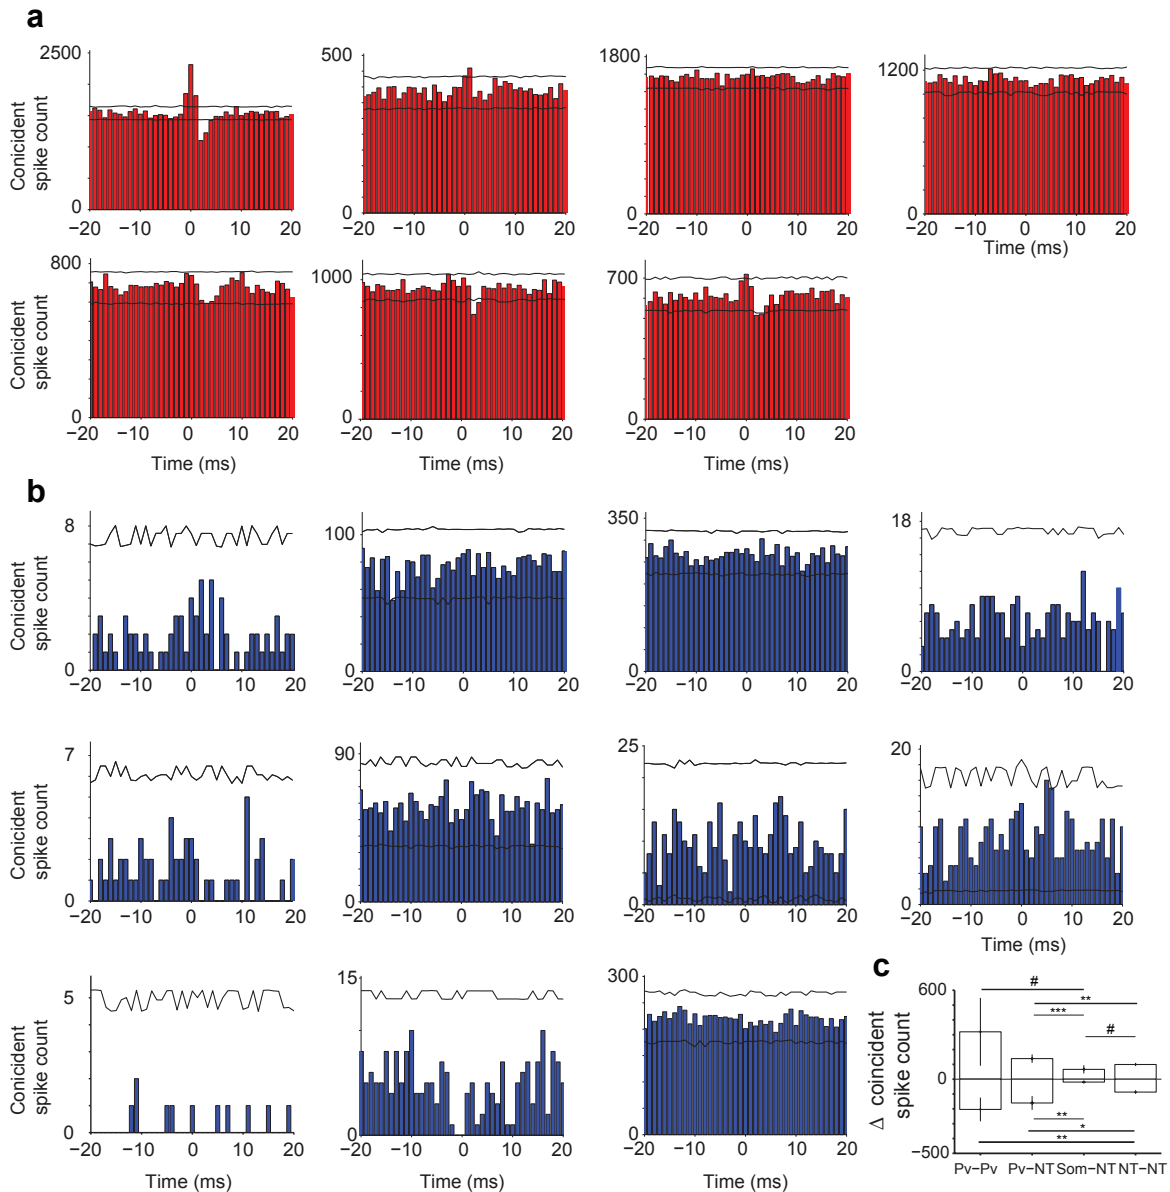
a, Histogram and cumulative histogram of isolation distance and **b**, L-ratio for all optically tagged neurons (blue) and not tagged neurons (gray). **c**, Scatter plot of waveform correlation (Pearson correlation coefficient between light evoked and spontaneous spikes) versus optical tagging test (SALT) for all recorded neurons. Neurons with $P < 0.01$ (by SALT) were considered optically tagged (blue). **d**, Box plots for all tagged Pv (red n=23) and Som (blue, n=35) neurons show low light evoked first spike latency (top panel) and small jitter (bottom panel). **e**, Z-scored PSTH of all tagged Pv (top panel) and Som (bottom panel) interneurons in response to 1 ms blue light stimulation.



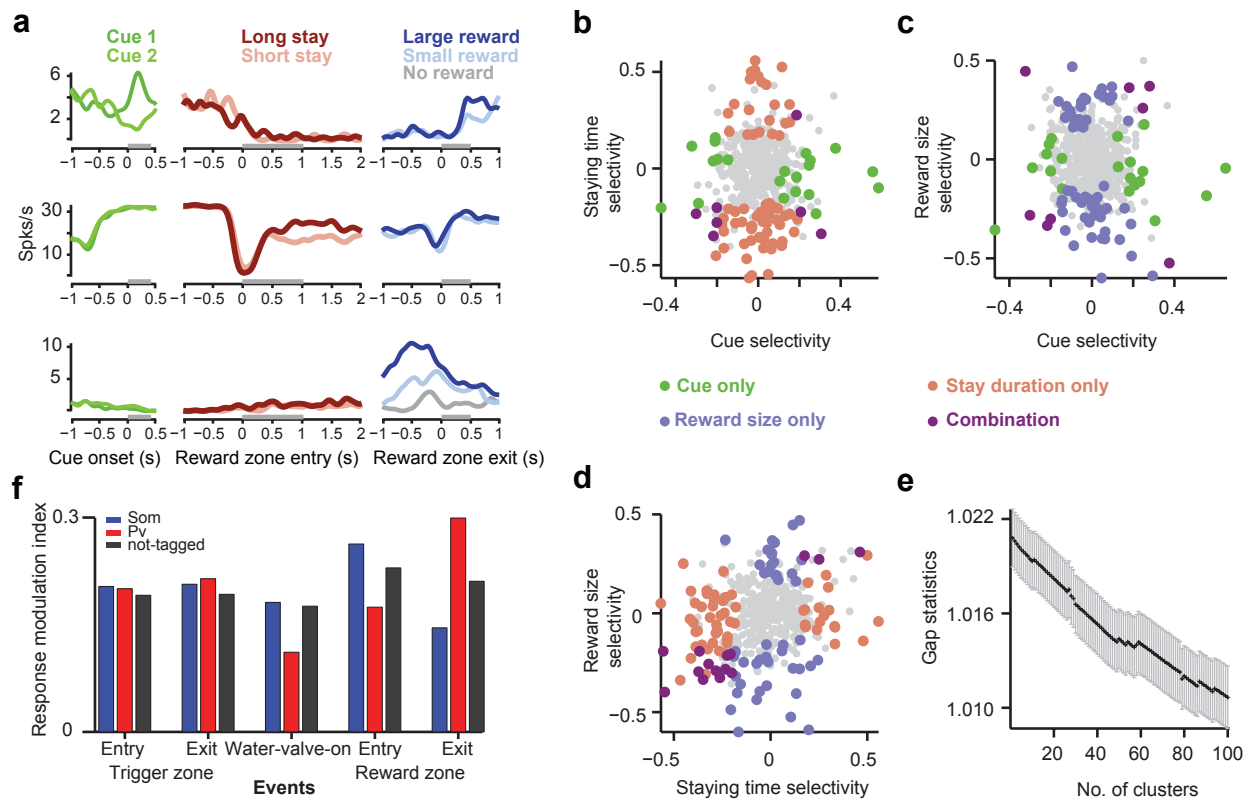
Supplementary Figure 5 | Individual cross-correlograms and mean counts for Pv, Som and NT population. **a**, Individual cross-correlograms for Pv-Pv pairs (red) and **b**, Som-Som pairs (blue). 99.9% confidence intervals are shown with black lines (see Methods). Note that number of significant short-timescale interactions among Pv-Pv pairs (5/7) compared to Som-Som pairs (0/11) were markedly different (two-sided Fisher's exact test $P < 0.01$). **c**, Average effect size for excitation and inhibition among all Pv-Pv, Pv-NT, Som-NT and NT-NT pairs. Effect size was quantified by the difference of peak (positive or negative) and baseline spike co-occurrence counts (non-normalized dimension of the CCGs) for all significant CCG effects. Bars above zero show average effect size for excitation whereas bars below zero show mean inhibitory effects. Error bars indicate SEM. Significantly different groups for inhibitory and excitatory effects are shown with horizontal bars (# $P = 0.05-0.06$, * $P=0.05$, ** $P<0.05$ and *** $P<0.01$ using two sided Mann-Whitney U-test).



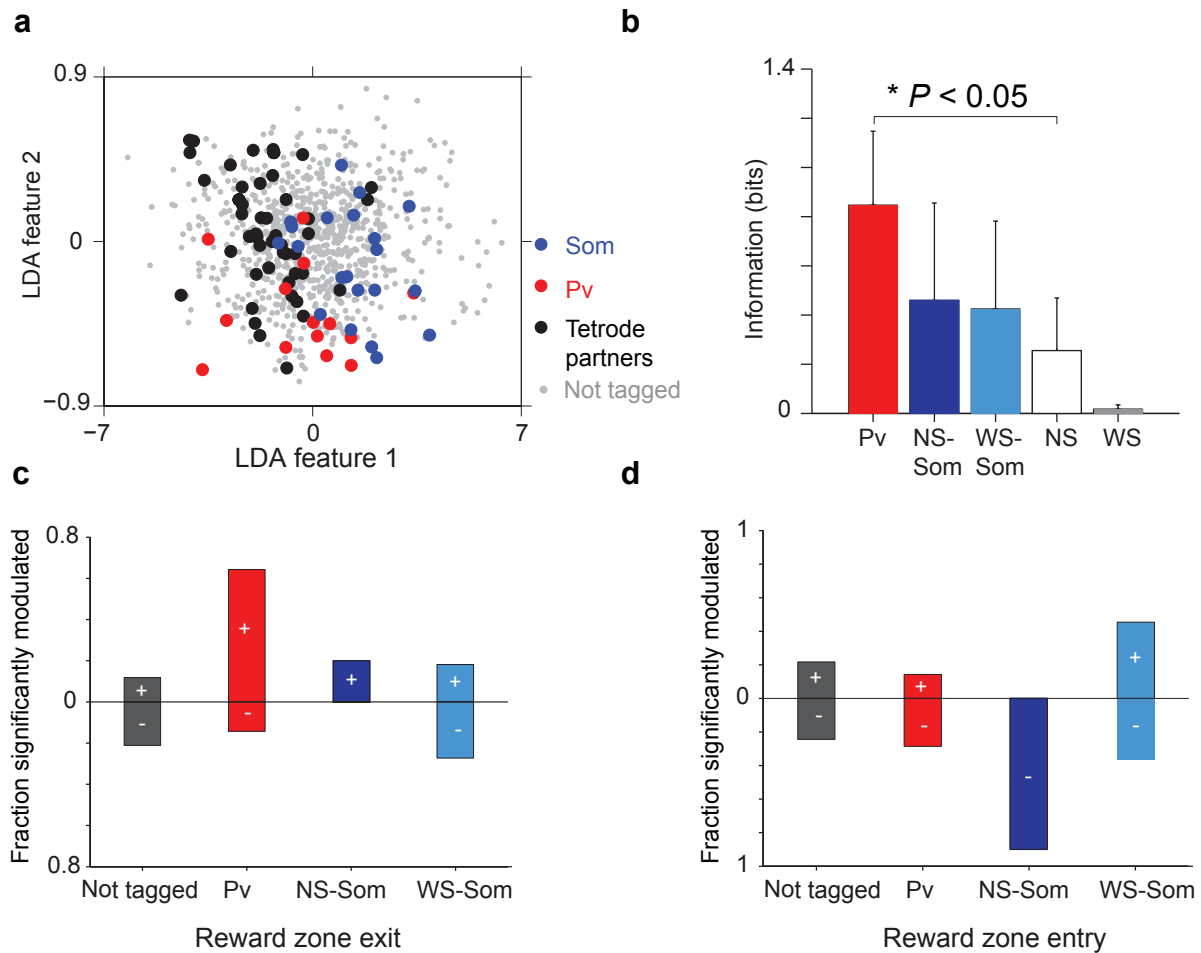
Supplementary Figure 4 | Estimation of blue light spread in brain tissue. **a**, Light induced c-fos expression in ACC of a *Pv-Cre* animal injected with AAV Ef1a-DIO-ChR2-EYFP virus. **b**, Quantification of Pv/c-fos double positive cells in the vicinity of ChR2 expression after the same light delivery protocol. We used representative sections from 2 *Pv-Cre* mice with the highest expression of ChR2-EYFP and counted EYFP positive neurons in consecutive and non-overlapping 0.25mm² patches (3-4 patches) starting from the dorsal surface. **c**, Photobleaching of the tdTomato fluorophore in a *Som-Cre::CAG-LSL-tdTomato* (Ai9) mouse line. Black arrow indicates bleached area. **d**, Density of ChR2-EYFP expression after transfection with AAV Ef1a-DIO-ChR2-EYFP viruses around the recorded sites for *Som-Cre* and *Pv-Cre* animals. Cells were counted within a 40 μm² area surrounding the lesion. Note that the density of ChR2 expression was similar in *Som-Cre* and *Pv-Cre* mice (183±17 neurons/mm², n=3 for *Som* and 176±14 neurons/mm², n=3 for *Pv*, as was the frequency of finding optically tagged neurons in electrophysiological recordings (5.4%; n=35/649 in *Som* and 4.7%; n=28/595 in *Pv* mice). Bar height denotes mean and error bar standard error of the mean.



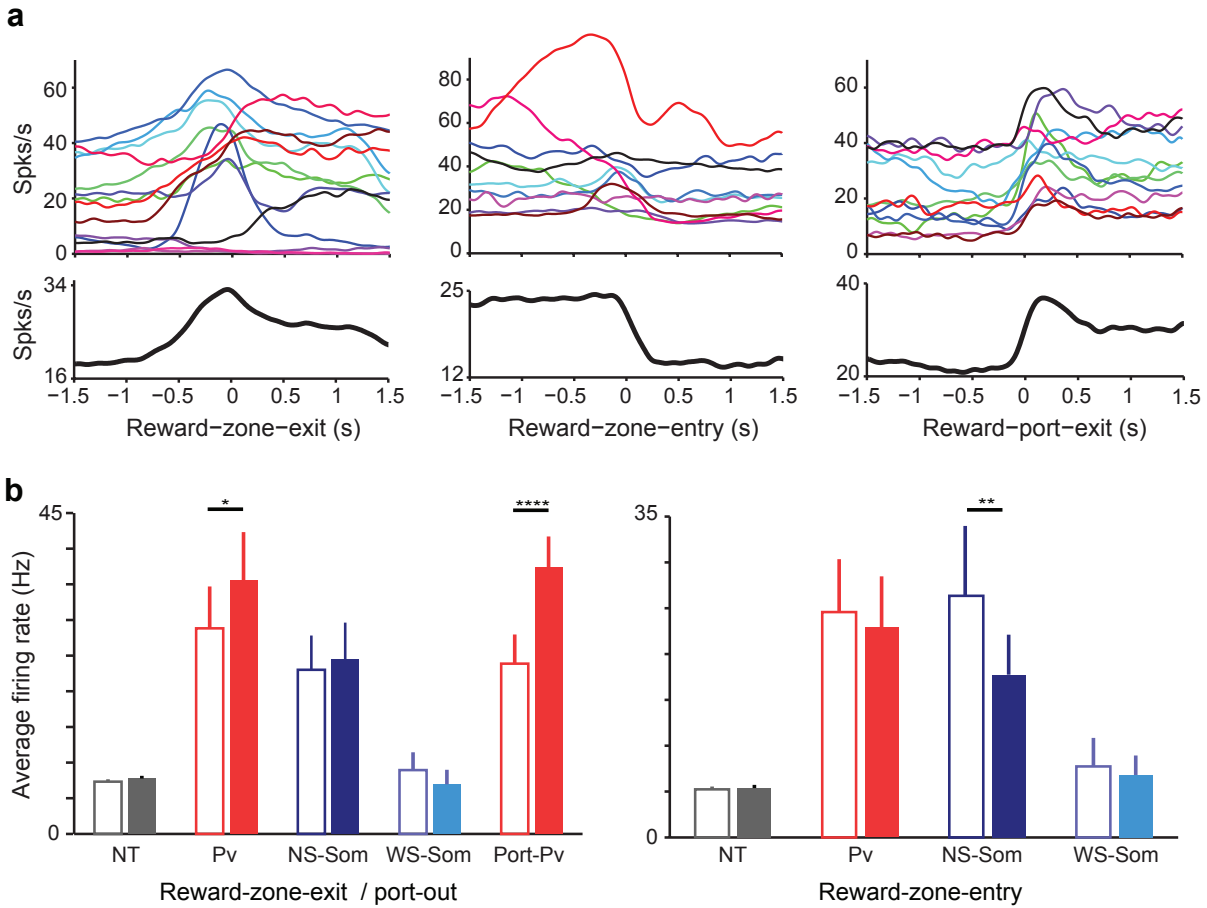
Supplementary Figure 5 | Individual cross-correlograms and mean counts for Pv, Som and NT population. **a**, Individual cross-correlograms for Pv-Pv pairs (red) and **b**, Som-Som pairs (blue). 99.9% confidence intervals are shown with black lines (see Methods). Note that number of significant short-timescale interactions among Pv-Pv pairs (5/7) compared to Som-Som pairs (0/11) were markedly different (two-sided Fisher's exact test $P < 0.01$). **c**, Average effect size for excitation and inhibition among all Pv-Pv, Pv-NT, Som-NT and NT-NT pairs. Effect size was quantified by the difference of peak (positive or negative) and baseline spike co-occurrence counts (non-normalized dimension of the CCGs) for all significant CCG effects. Bars above zero show average effect size for excitation whereas bars below zero show mean inhibitory effects. Error bars indicate SEM. Significantly different groups for inhibitory and excitatory effects are shown with horizontal bars (# $P = 0.05$ - 0.06 , * $P=0.05$, ** $P<0.05$ and *** $P<0.01$ using two sided Mann-Whitney U-test).



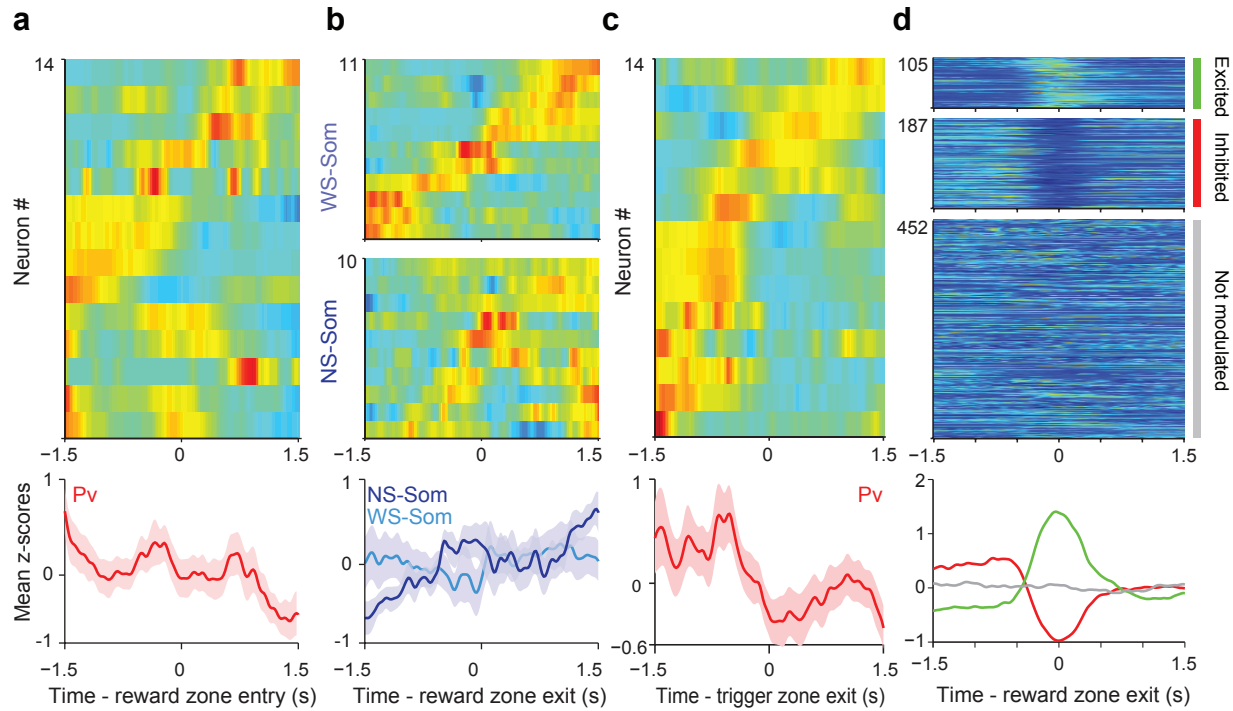
Supplementary Figure 6 | Heterogeneity of ACC neuronal responses in a reward foraging task. **a**, Peri-event time histograms (PETH) aligned to three events. Responses to cues predicting high and low reward (left column); responses to entry into reward zone separated by long versus short staying time in the reward zone (middle column) and responses to exit from reward zone separated by large, small and no reward i.e. reward omission trials (right column). Rows correspond to responses of three individual non-tagged neurons. Neurons show differential firing to auditory cues (top row, left column), staying time (middle row, center column) and outcome (bottom row, right column). Horizontal grey bars indicate time window used for selectivity analyses in **b-d**. **b-d**, ACC neurons showed response selectivity at multiple task events. Selectivity indices were calculated using ROC analysis (rescaled AUC between -1 and 1) comparing firing rates within epochs indicated by gray bars in **a**, for the two trial types. Statistically non-significant ($P > 0.05$, permutation test) values are indicated in grey, neurons with significant cue selectivity in green, staying time selectivity in orange and reward selectivity in blue; neurons showing selectivity to combinations of 2 events are shown in purple. **e**, Gap statistics plotted as a function of number of clusters does not reveal a preferential functional clustering in the behavioral correlates of ACC neurons. **f**, Response modulation index was computed for spike trains aligned to five major task events. Unidentified neurons show similar average modulation for all events without a preference for a particular event. *Pv* neurons show selectively higher average modulation at reward-zone exit event, while *Som* neurons are most modulated for reward-zone entry.



Supplementary Figure 7 | Distinct behavioral correlates of Pv and Som neurons. **a**, Dimension reduced response space (see Methods) illustrates that identified Pv (red) and Som (blue) interneurons, as well as their wide-spiking tetrode pairs (black) occupy distinct subspaces within the area spanned by the entire population of unidentified neurons (gray). **b**, A lower bound on the information gained about responses from knowledge of cell-type identity. Identified interneurons (Pv, NS-Som, WS-Som) are compared to unidentified narrow- (NS) and wide-spiking (WS) neurons. Error bars denote standard error of the mean. **c**, Fraction of significantly ($P < 0.01$, permutation test) activated (+) and suppressed neurons (-) for a given neuron-type during reward zone exit and **d**, reward zone entry event.

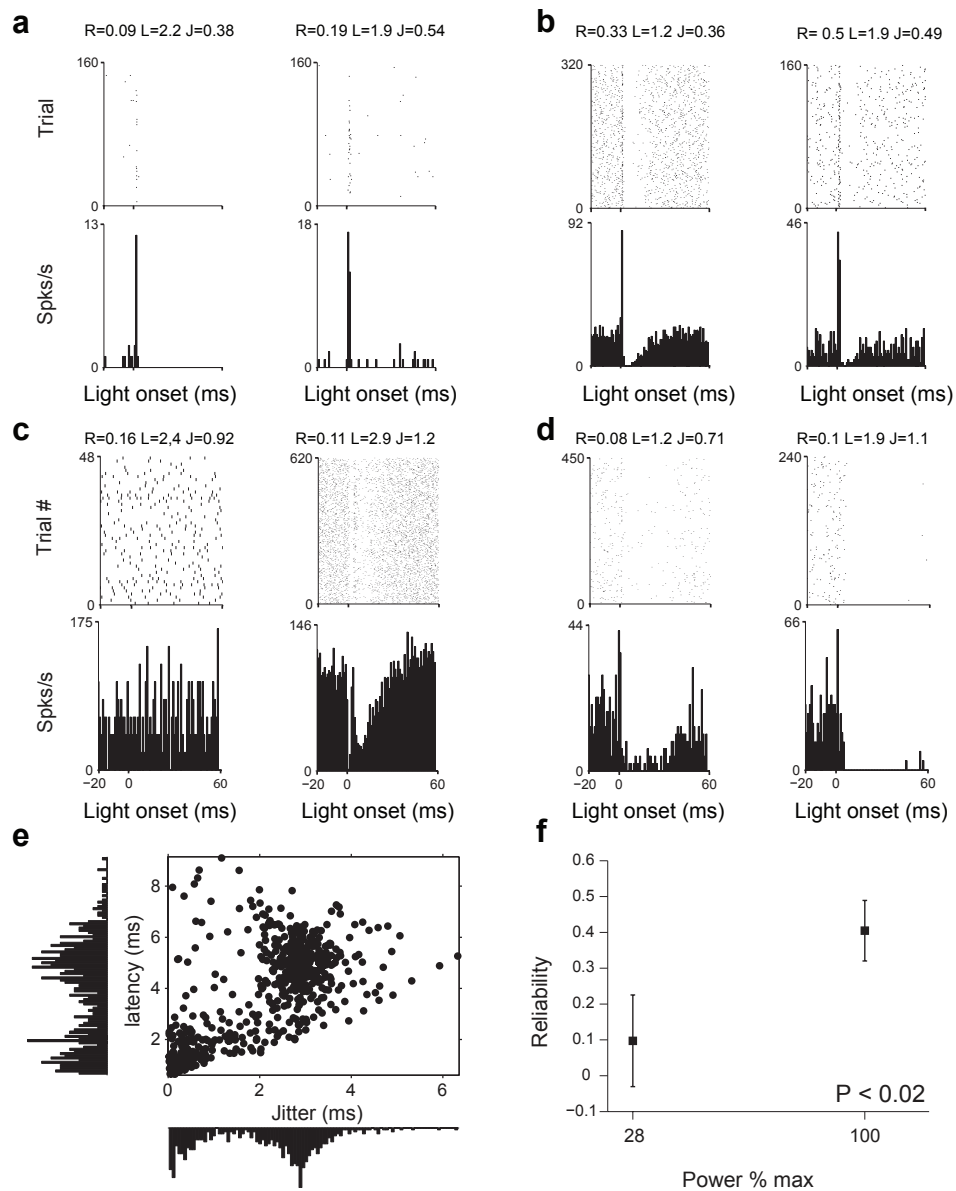


Supplementary Figure 8 | Firing rate dynamics for Pv and Som interneurons. a, PETHs of Pv (n=14, left panel), NS-Som (n=10, middle panel) and Pv interneurons recorded in the port task (Port-Pv, n=12, right panel) aligned to reward-zone-exit (left panel), reward-zone-entry (middle panel) and reward-port-exit (right panel) events, respectively. Individual neurons are shown in different colors and the averaged PETH is shown in the lower panels. **b,** Comparison of average peri-event firing rates for Pv, NS-Som, WS-Som and NT (not-tagged population) neurons at reward-zone-exit / reward-port-exit (left panel) and reward-zone-entry event (right panel). Port-Pv neurons are shown separately for reward-port-exit in the left panel. For each population of neurons the average firing rate in a baseline period (0.5 s before event) was compared with that in a peri-event period (0.5 s after event) using Wilcoxon signed-rank test, except for reward-zone-exit, baseline period was 0.5 - 1 s after event and peri-event period was 0.5 s before event. Statistically significant effects are denoted with horizontal lines * $P < 0.05$, ** $P < 0.01$ and *** $P < 0.0001$. Bar height denotes mean and error bar standard error of the mean.



Supplementary Figure 9 | Responses of optically tagged and unidentified neurons to specific behavioral events.

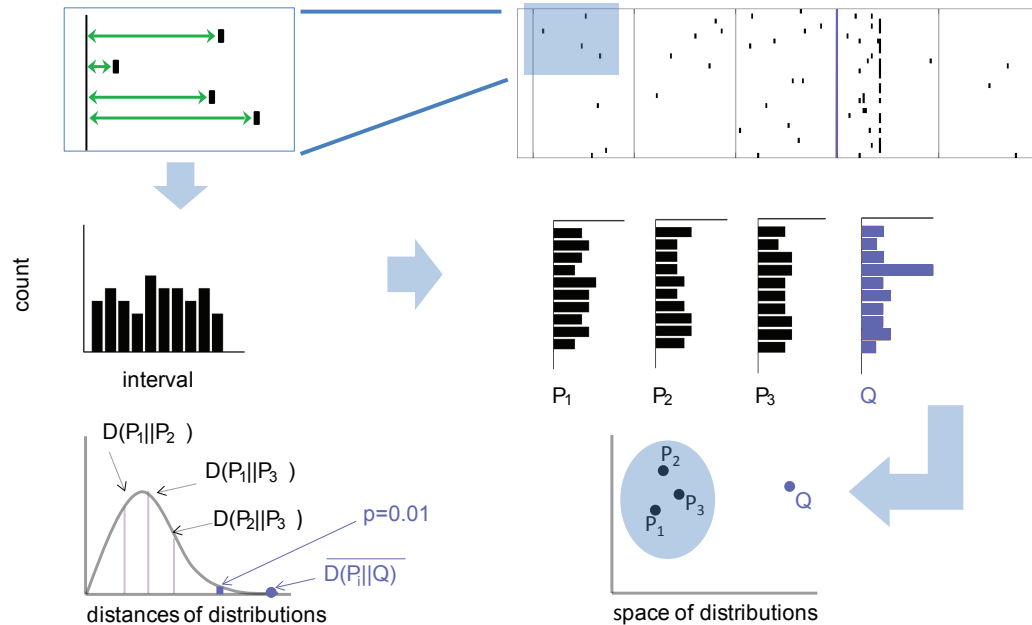
a, Responses of Pv population aligned to reward-zone entry. Each row represents a single neuron and their average PETH plotted below. **b**, Responses of wide-spiking (WS) and narrow spiking (NS) Som interneurons and average PETH (bottom panel) at reward-zone exit. **c**, Responses of Pv neurons aligned to trigger-zone exit and their average. Notice similarity of responses to this trigger-zone-exit and reward-zone-exit events, suggesting that Pv neurons were encoding more abstract decision to leave the platform (Fig.2 c,d). **d**, Z-scored PETHs for all unidentified neurons grouped by selectivity. Significant activation (green sidebar, $n=105$, $P < 0.01$, permutation test), suppression (red sidebar, $n=187$, $P < 0.01$, permutation test) and no significant modulation (gray sidebar, $n=452$, $P > 0.01$, permutation test) of firing rates around reward-zone-exit (from top to bottom). Averaged z-scored PETHs for the three groups are plotted (bottom panel).



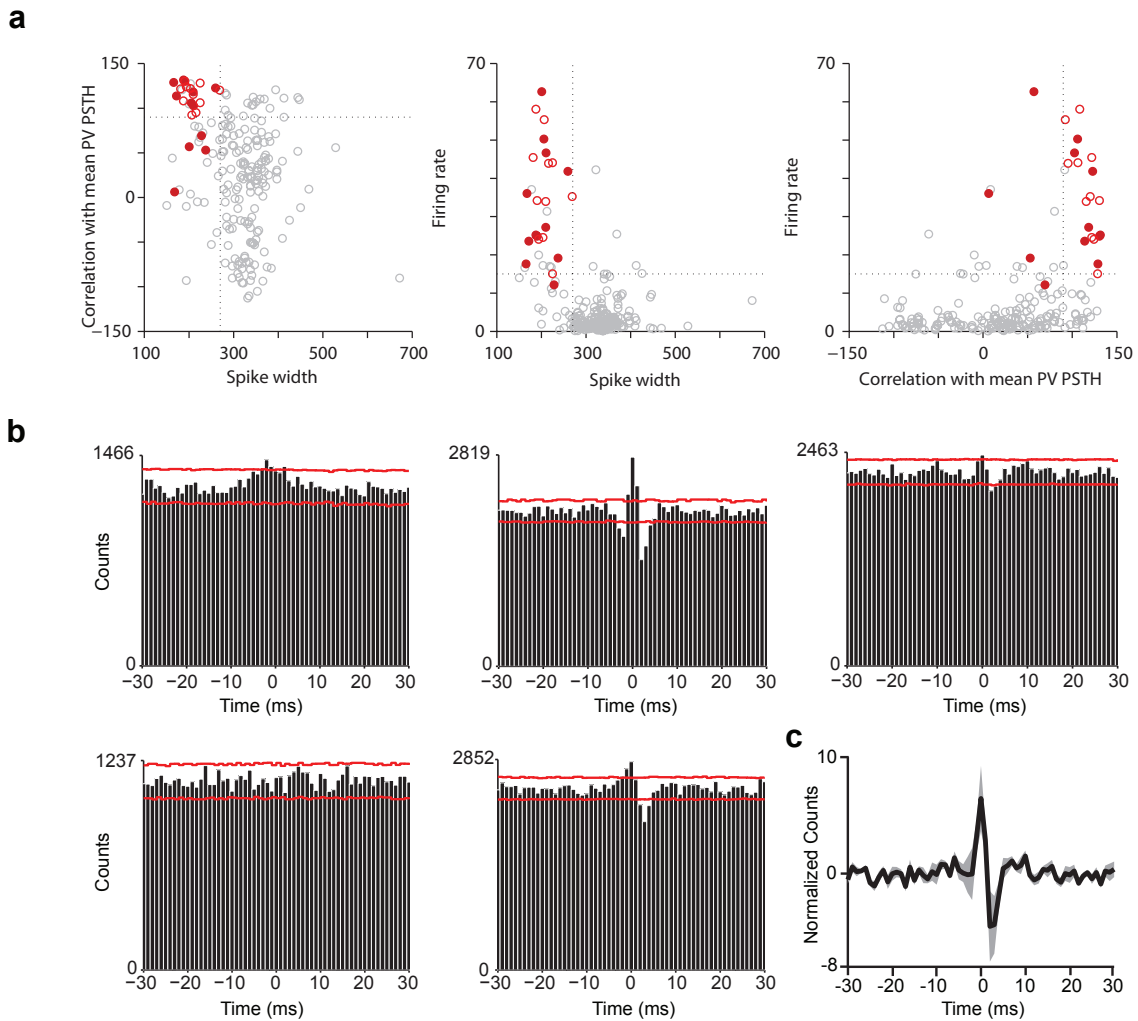
Supplementary Figure 10 | Advantages of statistical measure for optical tagging.

Spike raster and PSTH examples for directly activated and not directly activated neurons.

a, Responses of two tagged neurons with low baseline firing and low light evoked first spike reliability, latency and jitter (R , L and J respectively). **b**, Tagged cells based on SALT with high baseline firing rates and light induced excitation followed by inhibition. **c-d**, Raster/PSTH of not tagged cells that showed comparable light evoked first spike reliability, latency and jitter. **c**, High firing rate neuron (left panel) and neuron with a complex inhibition-excitation response (right). **d**, Two examples of precise light evoked inhibition. **e**, Scatter plot of first spike jitter versus latency for all neurons recorded. **f**, Light evoked spike reliability as a function of light power (max power = 2-6mW at the tip of 55 μ m fiber). Response reliability decreased with lower light intensity.



Supplementary Figure 11 | Visual description of optical tagging test (SALT). Top, stimulus-aligned raster plots were segmented to 10 ms bins and the latency of the first spike from bin start was determined for each stimulus trial. Bin starts were arbitrary in the baseline period (gray vertical lines) and aligned to light stimulus onset for test period (blue). Middle, histograms of these first spike latencies were constructed for baseline and light stimulation (blue). Light-evoked spikes appeared as an over-representation of a characteristic latency compared with the baseline. Bottom, modified Jensen-Shannon divergence (distance metric on distributions) between light-stimulation and baseline histograms was tested against divergence values derived from the baseline. If light-induced distribution of spike times was significantly 'far away' from baseline ($P < 0.01$), the cell was considered light activated.



Supplementary Figure 12 | Identification of putative Pv interneurons and their short-timescale interactions. **a**, Putative Pv (pPv) neurons were selected on the basis of 3 properties, correlation with average Pv PSTH in response to reward zone exit, average baseline firing rate and spike width. Scatter plots show distribution of these properties in the recorded neurons and co-segregation of optically tagged Pv neurons (red, filled circles) and pPv neurons (red, open circles). **b**, Cross-correlogram (CCG) of firing between pairs of concurrently recorded Pv neurons (both pPv and Pv) showing significant short timescale interactions (4/5 pairs, $P < 0.001$ bootstrap test). **c**, Normalized z-scored CCG for all recorded pPv pairs.