nature neuroscience

ARTICLES

https://doi.org/10.1038/s41593-022-01151-0

Behaviorally relevant decision coding in primary somatosensory cortex neurons

Christina Buetfering[1,2] ✉, Zihui Zhang[1,3], Margarita Pitsiani[1], John Smallridge[1,4], Ellen Boven[1,5], Sacha McElligott[1] and Michael Häusser[1] ✉

Primary sensory cortex is thought to process incoming sensory information, while decision variables important for driving behavior are assumed to arise downstream in the processing hierarchy. Here, we used population two-photon calcium imaging and targeted two-photon optogenetic stimulation of neurons in layer 2/3 of mouse primary somatosensory cortex (S1) during a texture discrimination task to test for the presence of decision signals and probe their behavioral relevance. Small but distinct populations of neurons carried information about the stimulus irrespective of the behavioral outcome (stimulus neurons), or about the choice irrespective of the presented stimulus (decision neurons). Decision neurons show categorical coding that develops during learning, and lack a conclusive decision signal in Miss trials. All-optical photostimulation of decision neurons during behavior improves behavioral performance, establishing a causal role in driving behavior. The fact that stimulus and decision neurons are intermingled challenges the idea of S1 as a purely sensory area, and causal perturbation suggests a direct involvement of S1 decision neurons in the decision-making process.

Understanding the neural circuit mechanisms that transform sensory information into behavior is a fundamental goal of neuroscience. Where and how decision signals, that is, internal variables that carry choice-relevant activity, are generated along the pathway to make an informed choice is key to understanding this transformation. The anatomical and conceptual separation of input-related sensory areas and output-related motor areas, has led to the prevailing idea that sensory areas extract information about the stimulus[1–3] whereas higher cortical areas use this information to generate decision variables[4–8]. This framework of hierarchical and anatomically defined processing steps has fostered ground-breaking studies in the visual system describing the transformation of sensory signals along the visual pathway[1,9]. In the somatosensory system of monkeys, seminal work by Romo and colleagues suggested that neurons in S1 represent primarily stimulus information whereas decision signals appear downstream in secondary somatosensory cortex (S2)[4,10–12]. As a consequence, S1 has only rarely been discussed as a potential locus of decision variables in perceptual decision making[5–8].

These findings were corroborated by work in rats and mice that reported stimulus-specific coding and a lack of decision coding activity in S1 (refs. [13–15]). However, a series of recent studies found that activity of neurons in L2/3 of S1 in mice performing a 'go/no go' task can be correlated with behavioral choice[16–19]. These demonstrations of choice-related activity in S1 seem to be in direct conflict with earlier work in monkeys, rats and mice. In these more recent studies, the behavioral report of detecting a stimulus is either a response (go) or the withhold of a response (no go). In trials where the animals responded, L2/3 S1 neurons exhibited higher activity. It is unclear whether this signal difference across trial types reflects behaviorally relevant decision coding or a modulation of neural activity by action-related variables such as motivation, movement preparation, feedback from the movement, sensory input that results from the movement or reward-related activity, which have been shown to be widespread across the brain[20–22]. For example, in primary visual cortex (V1), similar choice-related activity was observed during a go/no go task[23], which has been challenged by studies reporting the absence of choice-related activity in symmetric choice tasks that avoid the imbalance in behavior across trial types[22,24]. In summary, studies across mice, rats and monkeys suggest that neurons in S1 only carry stimulus information[2,13,15,25] in line with the idea of a strict processing hierarchy. However, recent studies reported putative choice signals in superficial neurons in S1 (refs. [16–18]) referring to a difference in activity in trials with a different choice action (go versus no go).

To test whether neurons in L2/3 of S1 carry behaviorally relevant decision information, it is crucial to use a task that features choice actions with symmetrical movement patterns. Moreover, we need a cued task design to separate the stimulus period from the choice action, as well as population recordings to allow dense sampling from as many neurons in L2/3 as possible. In a task with symmetric choice actions, differential neural activity before the motor action that aligns with the behavioral outcome can be distinguished from modulation of neural activity by action-related variables. Finally, to establish whether choice-related neural activity is indeed causally involved in the decision-making process, it is necessary to directly manipulate neuronal activity[26,27] during behavior.

Here, we have addressed these challenges by training mice to perform a cued two-choice texture discrimination task in which the behavioral choice was indicated by licking one of two lickports[19]. Mice had to discriminate different texture stimuli while we recorded neuronal activity simultaneously from hundreds of neurons using

[1]Wolfson Institute for Biomedical Research and Department of Neuroscience, Physiology and Pharmacology, University College London, London, UK. [2]Present address: Institute of Pathophysiology, University Medical Center of the Johannes Gutenberg University Mainz, Mainz, Germany. [3]Present address: Department of Psychiatry and Behavioral Sciences, Stanford University School of Medicine, Stanford, CA, USA. [4]Present address: Neurophenomenology of Consciousness Laboratory, Department of Psychiatry, Psychotherapy and Psychosomatics, Psychiatric Hospital, University of Zurich, Zurich, Switzerland. [5]Present address: School of Physiology, Pharmacology and Neuroscience, Faculty of Life Sciences, University of Bristol, Bristol, UK.
✉e-mail: christina.buetfering@gmail.com; m.hausser@ucl.ac.uk

two-photon calcium imaging of L2/3 neurons spanning multiple barrels. We then used an all-optical combination of two-photon optogenetic stimulation and two-photon calcium imaging in behaving mice to specifically activate ensembles of neurons that encode task-related information while reading out the activity from nearby neurons[28–32]. We show that in Correct trials, where stimulus and choice are correlated, the activity of many neurons (trial-coding neurons) in L2/3 is modulated by the trial type ('Smooth' (S) or 'Rough' (R)). When comparing the activity of trial-coding neurons in Correct and Incorrect trials across the two stimuli, we find that the activity of most trial-coding neurons in Incorrect trials was either aligned to the same stimulus as in Correct trials, or to the other stimulus. Activity in a subset of trial-coding neurons tracked stimulus identity irrespective of behavioral choice (stimulus neurons) whereas a different subset encoded behavioral choice irrespective of stimulus identity (decision neurons). On Miss trials, stimulus information was variable but a conclusive decision signal across the decision neuron population was absent. We find that stimulus and decision neurons exhibit differential activity patterns over time during the trial with the peak of activity of decision neurons following stimulus neuron activity. We also demonstrate that activity in decision neurons is not encoding a directional lick signal. Using targeted photostimulation and analysis of shared trial-by-trial variability, we show that stimulus and decision neurons exhibit different functional connectivity and local circuit integration. We further train animals to discriminate four textures and find that categorical coding of the choice develops during learning, which might reflect the formation of a stimulus-choice association. Finally, to test the behavioral relevance of decision-related activity in L2/3 S1, we use targeted photostimulation during behavior and find that this early decision variable in L2/3 S1 is causally involved in perceptual decision making. Thus, S1 is not just a sensory area extracting stimulus information but also encodes choice-related information that is involved directly in the decision-making process.

Results

Task-dependent activity in L2/3 of barrel cortex. We trained head-fixed mice to perform a two-choice texture discrimination task in which one of two textures was presented in each trial. Mice reported the identity of the stimulus by licking one of two lickports associated with each stimulus (Fig. 1a). At the start of each trial, one texture was moved into contact with the whiskers of the mouse to allow free sampling. Then, 3 s after the beginning of the trial, and approximately 1 s after the first touch between whiskers and the texture, an auditory go cue signaled the start of the response window (2 s) (Fig. 1b). Licking the lickport associated with the texture within the response window resulted in a 'Correct' trial and released a sugar water reward. Licking the wrong lickport was neither rewarded nor punished (Incorrect). Licking a lickport within the sampling period but before the go cue triggered an auditory punishment (1 s white noise) and the trial was aborted (Early response). Failure to lick any lickport during the withhold window or response window resulted in a 'Miss' trial. In all, 13 mice learned to perform the two-texture discrimination task with high accuracy (Fig. 1c and Extended Data Fig. 1). Mice performed 436 ± 114 (mean ± s.d.) trials per session. Analysis was performed on Correct and Incorrect trials in periods during which the mouse did not show a consistent lickport bias (trials without lickport bias: 355 ± 99 (mean ± s.d.); Methods), amounting to, on average, 197 ± 57 (mean ± s.d.) Correct trials and 80 ± 44 (mean ± s.d.) Incorrect trials per session.

We used two-photon calcium imaging to record population activity in barrel cortex across multiple barrels from excitatory neurons expressing GCaMP6s[17,18,33] in a 798 × 798 µm² field-of-view (FOV) (Fig. 1d). We captured the activity of 774 ± 240 (mean ± s.d.) neurons per imaging plane in each recording (total number of neurons = 61,895). We found that a subset of neurons in L2/3 barrel cortex showed trial-type-dependent responses (Fig. 1d). We then applied receiver operating characteristic (ROC) analysis to find neurons whose activity in the second before the lick differed between Correct trials with the smooth stimulus (Stim S) and Correct trials with the rough stimulus (Stim R) presented. We found that the activity of 25.2% ± 7.9% (mean ± s.d.) of neurons within a FOV encoded the trial type in Correct trials (Stim/Lickport S versus Stim/Lickport R trials) and refer to them as trial-coding neurons (Fig. 1e).

Stimulus and decision coding in L2/3 barrel cortex neurons. It is unclear whether a trial-coding neuron encodes the stimulus identity or the choice because a priori stimulus and choice are perfectly aligned in Correct trials. To test for choice coding in comparison with stimulus coding in the neural activity we calculated stimulus and choice selectivity for all neurons using a similar ROC analysis as described above but separated by trial type[34] (Fig. 2a and Extended Data Fig. 2a–c; stimulus selectivity, Stim S versus Stim R in Correct or Incorrect trials; choice selectivity, Correct versus Incorrect in Stim S or Stim R trials). We found that 36% of all neurons encoded the stimulus type and 31% of all neurons encoded the choice. To assess whether neurons with high stimulus selectivity encode the same stimulus in Correct and Incorrect trials, we looked at the activity patterns across trial types and found that a subset of these neurons encoded the same stimulus identity irrespective of the animal's choice (Fig. 2b). We refer to these neurons as 'stimulus neurons'. The activity of a different subset of neurons aligned with the behavioral choice of the mouse across different trial types. We refer to these neurons that exhibit selective firing to a specific choice irrespective of the type of stimulus that was presented as 'decision neurons' (Fig. 2b). We reliably found stimulus and decision neurons across all mice performing the two-choice texture discrimination task and stimulus and decision neurons show higher response reliability than other neurons in the FOV (Supplementary Fig. 1). Using ROC analysis (Stim S versus Stim R in Correct and Incorrect trials), we find that 10.1% ± 1.7% (mean ± s.e.m.) of all neurons in the FOV identify as stimulus neurons and 4.3% ± 0.6% (mean ± s.e.m.) as decision neurons (Fig. 2c).

To further characterize the local network integration of stimulus and decision neurons in S1 we analyzed their shared variability and spatial clustering. If stimulus and decision neurons are part of distinct functional subnetworks in the L2/3 barrel cortex, we hypothesized that shared trial-by-trial response variability between neuron pairs would be stronger within each group compared with across group[35–37]. Indeed, we found that trial-by-trial response variability is stronger within either group than across groups or across all neurons in L2/3 (Fig. 2d). Despite no obvious clustering with respect to barrel centers in either subgroup, the population of stimulus neurons exhibits a slight spatial clustering across the FOV whereas decision neurons are distributed uniformly (Extended Data Fig. 3). Together, these results suggest that stimulus and decision neurons form distinct subgroups that are distinguishable in more than just the trial-coding dimension.

To rule out that the decision signal in decision neurons is driven primarily by movements associated with the choice action, we performed a series of analyses. First, we repeated our ROC analysis of stimulus and decision neurons but restricted the analysis window to 100 ms before the lick and show that we can detect similar numbers of decision neurons (Extended Data Fig. 2d; 9.5% ± 1.6% stimulus neurons and 3.9% ± 0.5% decision neurons, mean ± s.e.m.). This means that the decision signal is present before a potential motor signal. To test the possibility that the decision signal in decision neurons is driven by whisking or running, we tracked whisker movement and running in a subset of sessions and built a generalized linear model (GLM) to dissect the impact of four regressors: stimulus type, behavioral choice, running speed and whisking (nine mice, two sessions each; Fig. 2e and Extended Data Fig. 4a,b). To identify the unique

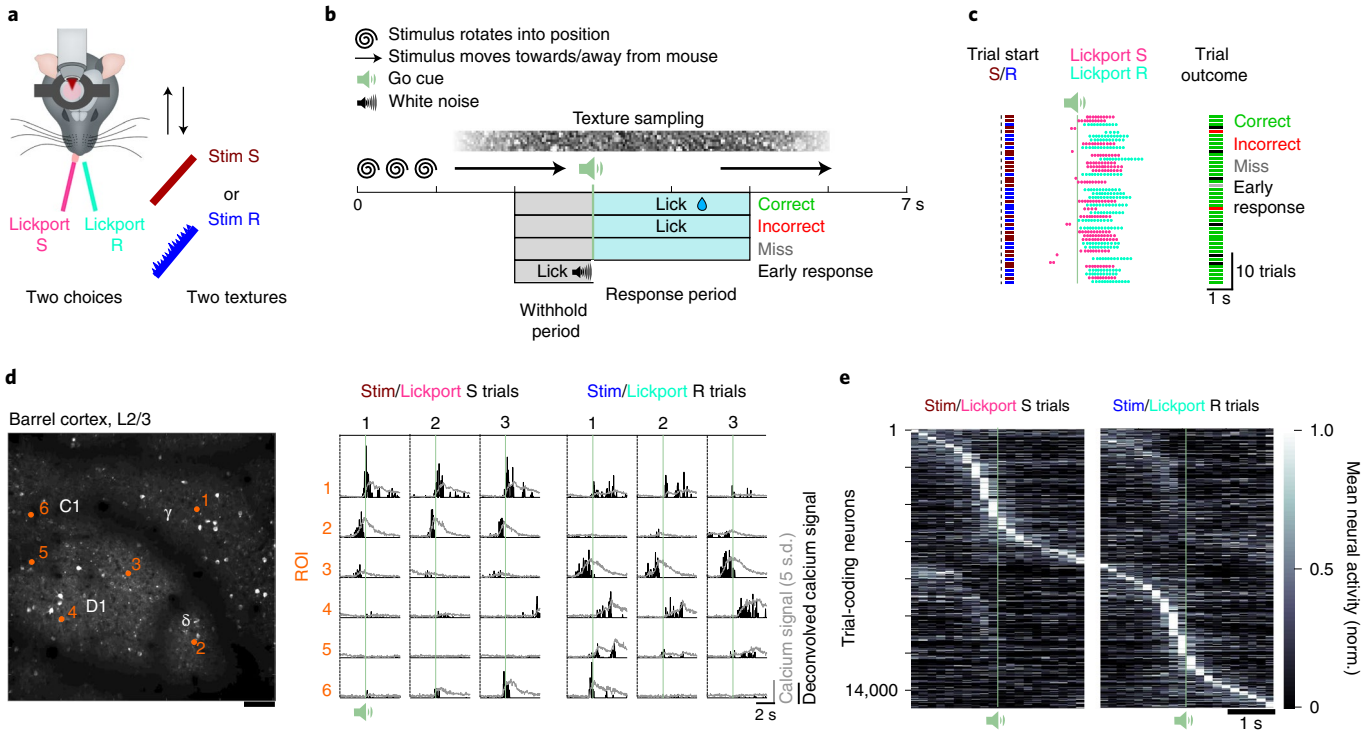

Fig. 1 | Imaging task-dependent activity in L2/3 barrel cortex during a two-choice texture discrimination task. a, Experimental set-up of the two-choice texture discrimination task with simultaneous two-photon imaging. Stim S, smooth sandpaper; Stim R, rough sandpaper. **b**, Trial structure and trial outcomes. **c**, Behavioral data from 44 consecutive trials of a single session with two textures. **d**, Left: imaging FOV (depth = 140 µm) with located barrel centers (C1, D1, γ, δ) and selected neurons (orange numbers). Repeated for all mice ($n = 13$). Right: fluorescence (gray) and deconvolved fluorescence (black) traces from selected neurons. Traces are aligned to the trial start. Fluorescence traces have been corrected for neuropil contamination, baselined and z-scored. Scale bar, 100 µm. **e**, Normalized (norm.) mean activity of all trial-coding neurons, that is neurons that distinguish between correct Stim/Lickport S and correct Stim/Lickport R trials. Neurons sorted by trial type preference and time of maximum activity. Single-plane recordings (30 Hz sampling rate) have been binned to match multiplane recordings (5 Hz sampling rate).

contribution of each regressor, we calculated how much of the average trial activity of a given neuron is explained by any of the regressors using semipartial regression with randomized regressors (Fig. 2f and Extended Data Fig. 4c,d). We found that 8.3 ± 2.5% (mean ± s.e.m.) of neurons encode running information, 9.0 ± 1.5% (mean ± s.e.m.) encode whisking information and 6.9 ± 2.0% (mean ± s.e.m.) encode stimulus information. However, in 3.6 ± 1.2% (mean ± s.e.m.) of neurons, the choice regressor was able to explain neural activity that was not explained by stimulus, running or whisking (Fig. 2g and Extended Data Fig. 4e,f). Therefore, decision signals in L2/3 are not simply a result of licking, running or whisking movements correlated with the choice of the mouse.

Distinct activity patterns in stimulus and decision neurons. We found that stimulus neurons are more active than decision neurons during the sampling period, whereas decision neurons show more activity outside of the sampling period compared with stimulus neurons (Fig. 3a). However, activity in both neuron types rises sharply after the onset of sensory input and most stimulus and decision neurons peak before the go cue and the lick (Extended Data Fig. 5a). When we aligned the activity of stimulus and decision neurons to the go cue as a proxy for stimulus presentation as well as to the lick of the mouse (Fig. 3b), the activity in both stimulus and decision neurons rose more quickly when aligned to the stimulus presentation than when aligned to the lick (Fig. 3c; stimulus neuron activity rise time (Stimulus): 400 ms; rise time (Lick): 633 ms; decision neuron activity rise time (stimulus): 600 ms; rise time (Lick): 767 ms; rise time from 10% to 90% of peak activity). Aligning stimulus and decision neuron activity to the first whisker touch in a subset of sessions in which whisker kinematics were recorded reveals similar results (Extended Data Fig. 5b–d). This suggests that activity in both neuron types is driven preferentially by stimulus presentation. The peak of decision neuron activity followed the peak of stimulus neurons in both conditions with a difference in peak time of 120 ms or 270 ms when aligned to the stimulus or the lick, respectively. For the decay of the signal, decision neuron activity fell more quickly when aligned to the lick than the stimulus and started to decay before the first lick (decision neuron activity decay time (Stimulus): 567 ms; decay time (Lick): 433 ms; decay to 50% of peak activity). Together, these analyses indicate that stimulus and decision neuron activity follows sensory input and precedes the lick, that decision neuron activity decays quickly once the lick is initiated and that stimulus neuron activity peaks before decision neuron activity.

Categorical coding in decision neurons develops with learning. If stimulus neurons encode stimulus properties, and decision neurons are associated with the categorical choice of the mouse, then a selective extension of the stimulus space should predominantly affect stimulus neurons while decision neurons encode the binary choice. To study specifically how the decision signal develops during the formation of a stimulus-choice association, it is necessary to separate learning of this association from the initial stages of task learning accompanied by changes affecting whisking and running patterns that directly influence the incoming sensory signal[33,38,39]. To this end, we trained four two-texture discrimination expert mice to associate an additional intermediate texture (S2 and R2) to each of the lickports (Fig. 4 and Extended Data Fig. 6a). On the first day

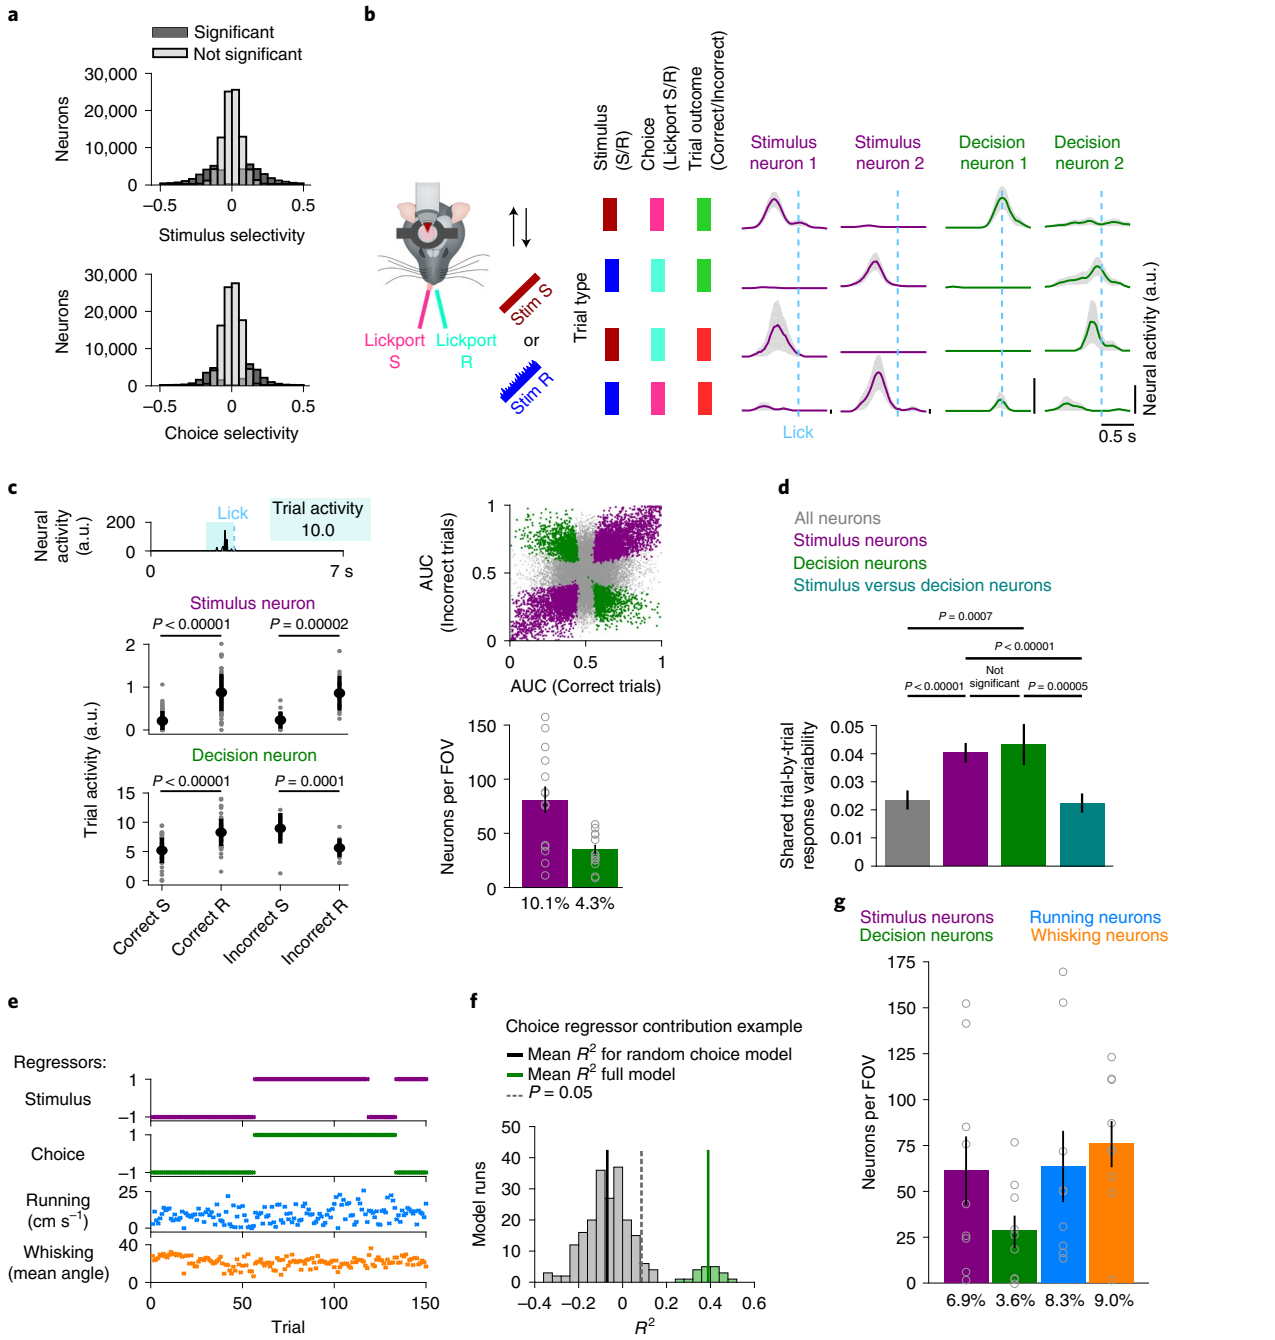

Fig. 2 | Stimulus and decision coding in L2/3 barrel cortex neurons. a, Stimulus and choice selectivity for all neurons. Two conditions for each neuron (stimulus selectivity, Correct and Incorrect trials; choice selectivity, Stim S and Stim R trials). Neurons with significant selectivity in a condition are labeled in dark gray. **b**, Trial-averaged activity of four example neurons recorded during the two-choice two-texture task. Activity split by trial type. Activity was aligned to the lick of the mouse; mean ± s.e.m. Neural activity is neuropil-corrected and deconvolved calcium signals. **c**, Top left: trial activity is calculated by averaging neural activity 1 s before the lick. Bottom left: trial activity of an example stimulus and an example decision neuron across trial types. Top right: ROC AUC values for all neurons. Neurons with significant stimulus coding in purple and neurons with significant decision coding in green. Bottom right: number and percentage of stimulus and decision coding neurons per FOV ($693 \times 693\,\mu m^2$). $n = 14$ FOVs (13 mice), mean ± s.e.m. Gray open circles denote individual FOVs. **d**, Pairwise correlation of the response to repeated presentations of identical stimuli for all neurons, stimulus neurons, decision neurons, and between stimulus and decision neuron pairs. $n = 63$ FOVs, two-sided Wilcoxon signed-rank test, mean ± s.e.m. **e**, The four regressors, sorted by stimulus and choice across trials, used to train the Gaussian GLM: stimulus, choice, running and whisking. **f**, Identification of variance uniquely explained by each of the regressors using semipartial regression. The mean variance explained for each neuron using all regressors is compared to the performance of a model with one randomized regressor. In this example the choice regressor was randomized 400× to generate a distribution of random choice regressor models. Neurons whose activity is predicted better with the regressor of interest ($P < 0.05$, one-sided t-test) are considered to encode information about the regressor of interest. **g**, Number and percentage of neurons identified to contain information on each of the four regressors per FOV. Gray open circles denote individual mice. $n = 9$ mice, mean ± s.e.m. a.u., arbitrary units.

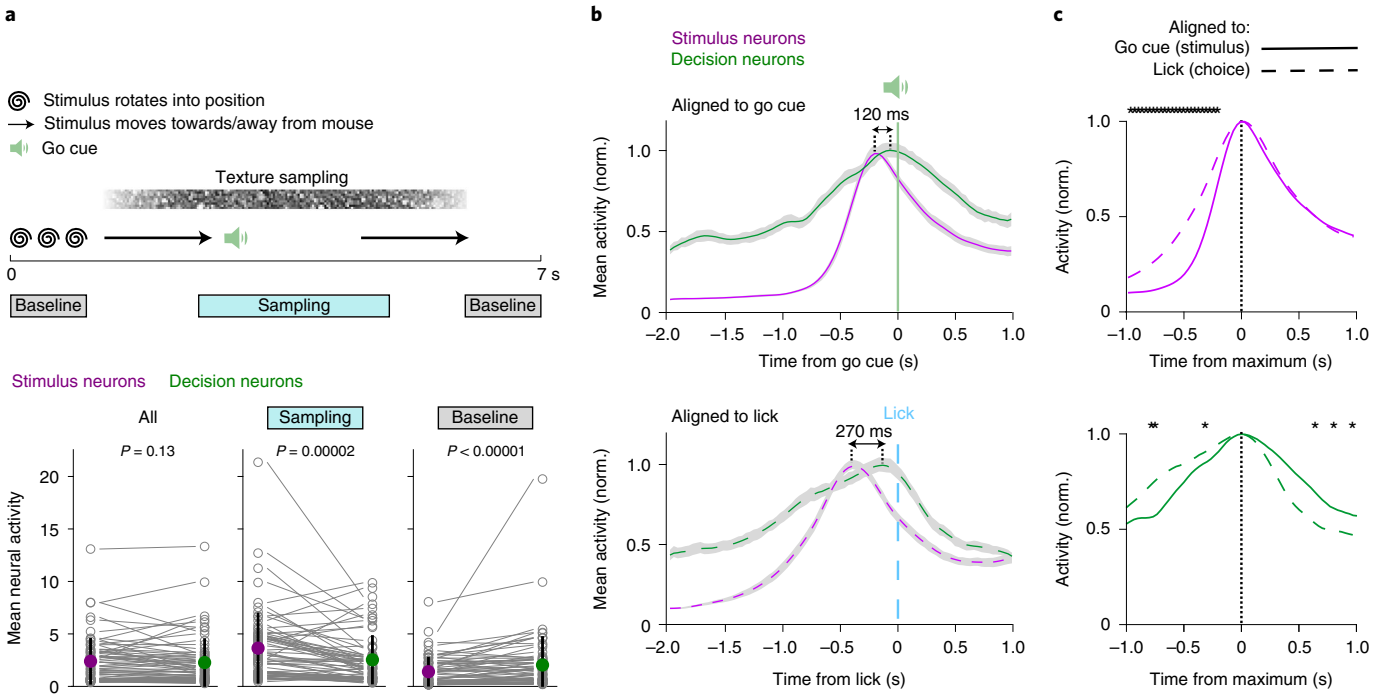

Fig. 3 | Subset-specific activity patterns and timing in stimulus and decision neurons. a, Top: trial structure to indicate 'Baseline' and 'Sampling' periods. Bottom: mean neural activity of stimulus and decision neurons throughout the trial and during the Sampling and Baseline period. Neural activity is neuropil-corrected and deconvolved calcium signals. $n = 63$ sessions, mean ± s.e.m., two-sided Wilcoxon signed-rank test. Gray open circles denote individual sessions. **b**, Activity of stimulus and decision neurons averaged across sessions and aligned to the go cue/stimulus onset (top) or the lick/choice of the mouse (bottom). The go cue serves as a proxy for the stimulus onset that is perfectly correlated with the go cue but with a small jitter due to trial-specific whisker movements. Line and shaded area represent mean ± s.e.m. **c**, The same activity traces but sorted by neuron type and aligned to the peak. Top: the sharp activity onset in stimulus neurons aligned to the stimulus differs from the rise in activity when aligned to the lick. Bottom: decision neuron activity falls off faster when activity is aligned to the lick in comparison to activity aligned to the stimulus. But activity rises more sharply when aligned to the go cue in comparison to activity aligned to the lick. Bin-wise comparison, $n = 35$ sessions, stars indicate bins with $P < 8.2 \times 10^{-4}$ (Bonferroni corrected significance level), two-sided Wilcoxon signed-rank test.

of four-texture training, we identified decision neurons based on their responses to the previously trained textures (S and R). Their responses to the newly introduced stimuli (S2 and R2) were variable and they lacked a binary response across the four textures (Fig. 4a). After training, decision neurons in four-texture mice showed categorical coding across the four stimuli, reflecting the association of two different stimuli to one lickport. During learning, the categorical coding of decision neurons, that is, the difference of responses between stimuli associated with either lickport (Stim S and S2 versus Stim R and R2) increased with an improvement of behavioral texture discrimination, whereas the discrimination between textures associated with the same lickport decreased in all four mice (Fig. 4b,c). This categorical coding was specific to decision neurons while stimulus neurons differentiated more strongly between textures associated with the same lickport (Fig. 4d,e). Furthermore, to test that the decision signal is dissociable from a lick-related signal, we took advantage of the fact that, at the beginning of the four-texture training, mice initially randomly lick the right and left lickport in the presence of a texture. We find that during this initial learning phase, before an association between the stimulus and the lickport has been made, decision neurons do not encode the lickport identity (Extended Data Fig. 6b,c).

In summary, decision neurons exhibit categorical coding, while stimulus neurons encode stimulus identity and the occurrence of categorical decision coding in decision neurons develops with learning. Therefore, the existence of a categorical decision signal in decision neurons in L2/3 barrel cortex is likely to be crucial for behavioral performance in our texture discrimination task.

Miss trials lack a conclusive decision signal. If the decision signal is relevant for the mouse to make a choice, we wondered whether an inconclusive decision signal would inhibit the mouse from making a choice on a trial-by-trial basis. To test this possibility, we analyzed stimulus and decision neuron activity in Miss trials (15% ± 10% (mean ± s.d.) of trials across all sessions, Fig. 5a). To distinguish trials in which the lack of a lick response was accompanied by a lack of stimulus information (for example, due to insufficient whisking at the end of the session), we separated Miss trials into 'Miss Stimulus+' and 'Miss Stimulus−' trials based on whether a linear classifier trained on stimulus neuron activity in Correct trials could predict the stimulus type in Miss trials (Fig. 5b). Miss Stimulus− trials are accompanied with a lower whisking amplitude, lower running speed and lower overall activity in stimulus neurons compared with Correct trials. In contrast, on Miss Stimulus+ trials, behavioral variables are unchanged and stimulus neuron activity is even higher compared with Correct trials (Fig. 5c,d and Extended Data Fig. 7). The overall activity in decision neurons is similar in Correct and Miss Stimulus+ trials. Therefore, a reduction in activity alone does not explain the lack of a lick in Miss Stimulus+ trials. But what if the decision signal across the decision neuron population is inconclusive? To address this, we used a classifier trained on decision neuron activity in Correct trials and asked if it can predict choice on Miss Stimulus+ or Miss Stimulus− trials. Performance was at chance level for both trial types (Fig. 5d). Thus, an informative sensory signal in stimulus neurons in L2/3 barrel cortex is not sufficient for the mouse to trigger a choice action. However, the lack of a decision signal is associated with the absence of a choice action in Miss trials.

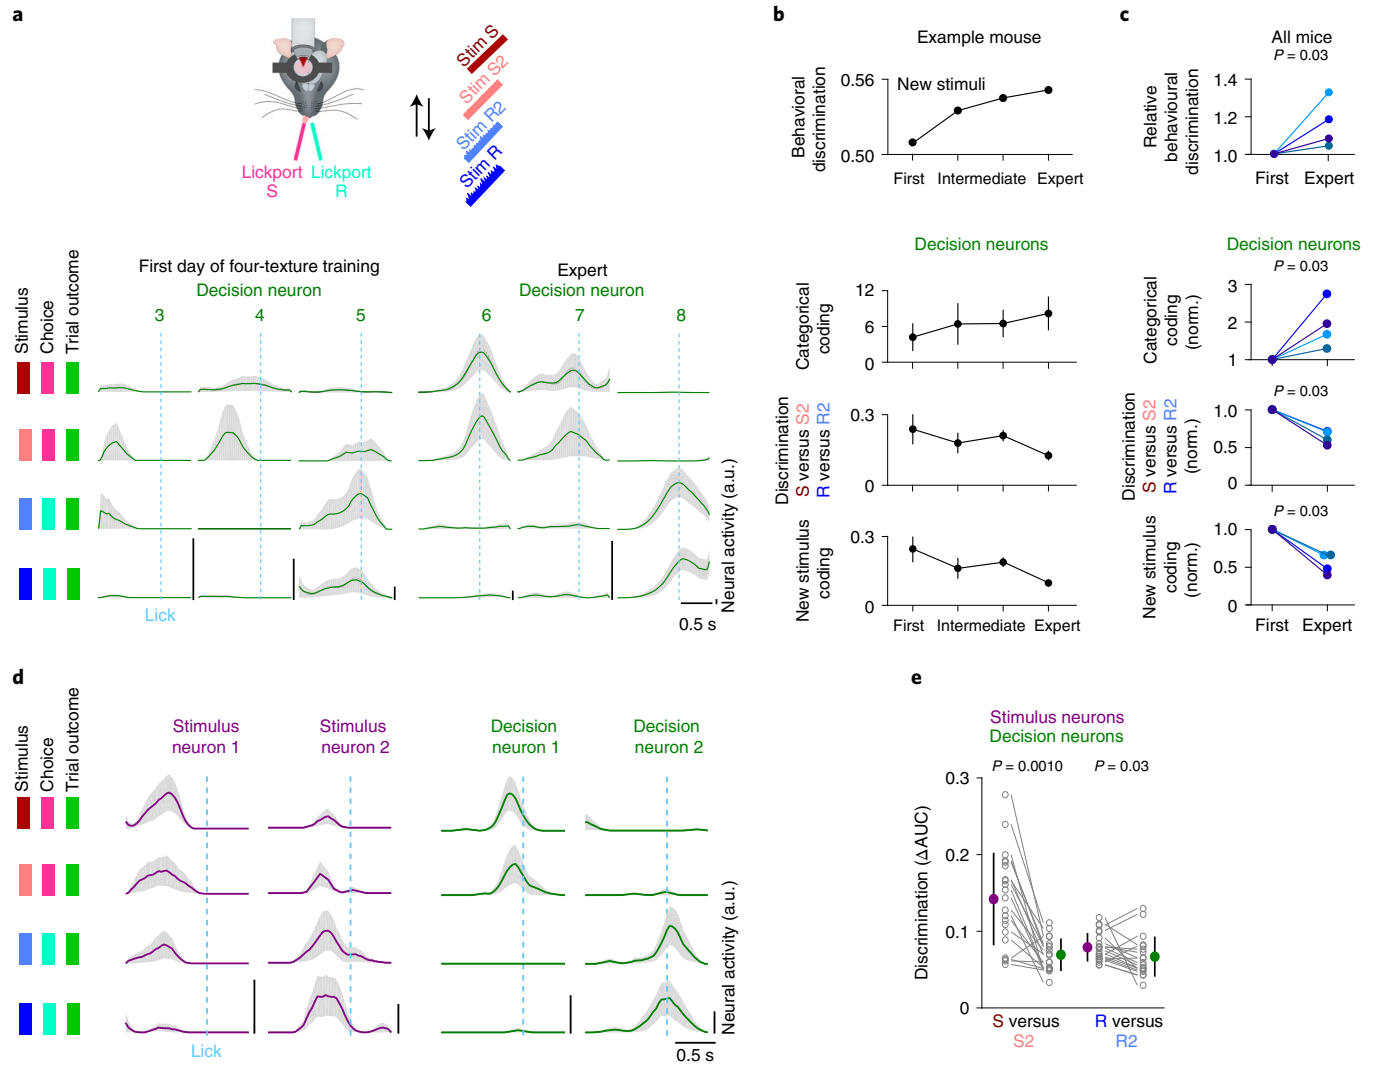

Fig. 4 | Categorical coding in decision neurons develops with learning. a, Activity of decision neurons across trial types in six example neurons: three example neurons on the first day of four-texture training and three example neurons in expert mice. Mean activity sorted by trial type. Line and shaded area represent mean ± s.e.m. **b**, Change in behavioral performance ($n = 1$ mouse) and coding of decision neurons ($n = 5/6/43/16$ decision neurons in first, two intermediate and expert session, mean ± s.e.m.) during learning in an example mouse. **c**, Change in behavioral performance and coding of decision neurons in all mice that improved in performance during learning (colors indicate different mice). The first session with four textures (First) is compared with the session with the best performance (Expert). $n = 4$ mice, two-sided Wilcoxon signed-rank test. The coding parameters are the strictly standardized mean difference $(\mu 1 - \mu 2/\sqrt{(s.d.1^2 + s.d.2^2)})$ for Stim S and Stim S2 versus Stim R and Stim R2 (Categorical coding), Stim S versus Stim S2 and Stim R versus Stim R2 (Discrimination), and Stim S and Stim R versus Stim S2 and Stim R2 (New stimulus coding). Behavioral discrimination = number of Correct trials/(number of Correct + Incorrect trials) for S2 and R2. **d**, Trial-averaged activity for four example neurons recorded during the four-texture discrimination task. Activity split by trial types. Activity was aligned to the lick of the mouse. Mean ± s.e.m. **e**, Discrimination of textures associated to the same lickport in decision neurons and stimulus neurons; 21 FOVs, mean ± s.d., two-sided Wilcoxon rank sum test.

Cell-type specific functional connectivity in L2/3 neurons. To test if the decision signal we observed is indeed causally involved in decision making and not due to feedback related to the choice action, it is important to directly manipulate the activity of selected neurons in L2/3 of barrel cortex. Therefore, we performed targeted two-photon photostimulation and simultaneous two-photon calcium imaging in behaving mice[28–32,40–42] to probe the network integration of stimulus and decision neurons, and to establish causal links between neural activity and behavior.

Stimulus and decision neurons show higher shared trial-by-trial response variability within than between subgroups, which could be due to increased recurrent connectivity or common feedback (Fig. 2d). To directly probe the local network integration, we selectively stimulated either of two different subsets of trial-coding neurons in a subset of 'catch' trials where no texture stimulus was presented to study the recruitment of stimulus and decision neurons depending on the activated target ensemble (Targets S and Targets R; Fig. 6a and Extended Data Fig. 8). We trained mice coexpressing the calcium indicator GCaMP7f and the excitatory soma-targeted opsin C1V1 in L2/3 barrel cortex to perform the two-choice texture discrimination task (Fig. 6b,c). Neurons were selected for target ensembles based on trial selectivity, and target ensembles could be activated with cellular resolution[41,43] (Fig. 6d,e). Each target ensemble across experiments contains neurons with a varying degree of stimulus and choice selectivity. We then analyzed the functional properties of 'followers', that is, neurons across the FOV that were not directly activated by the photostimulation laser but were activated or suppressed as a consequence of activating the

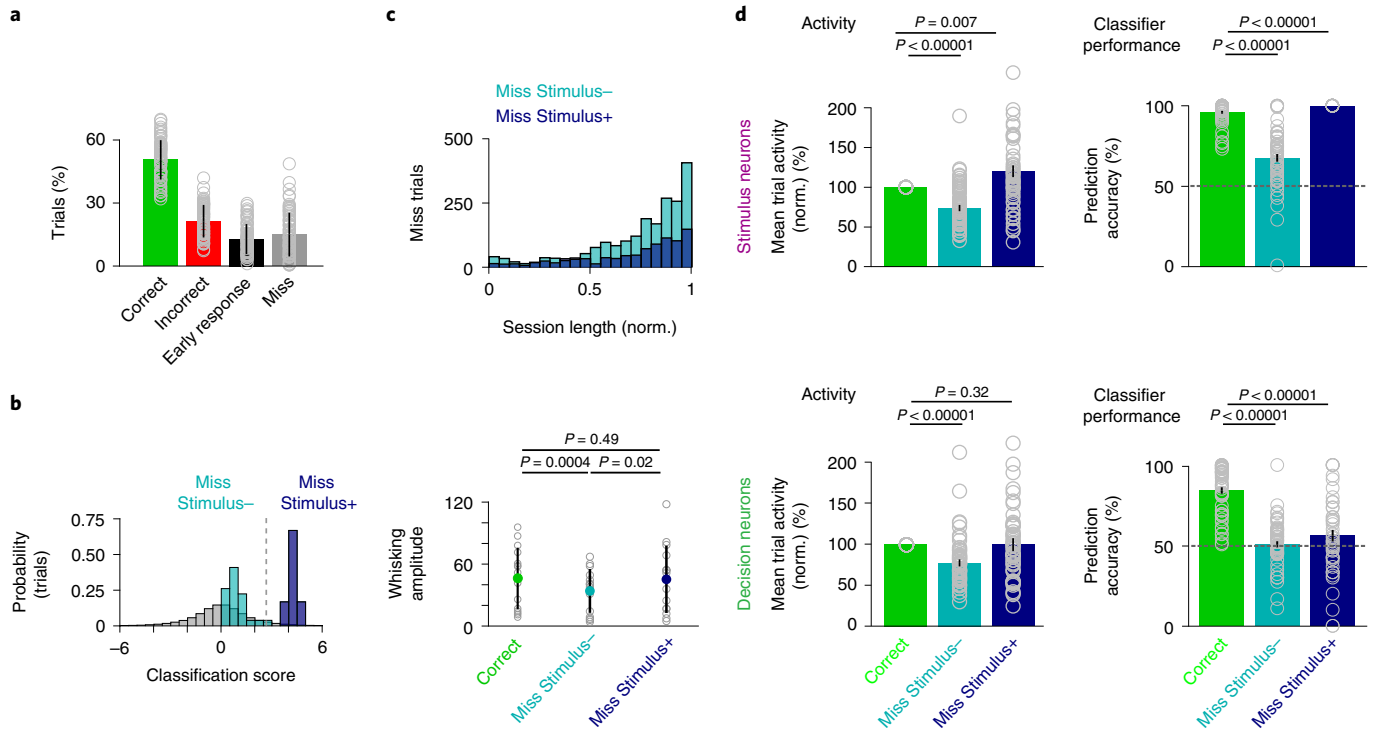

Fig. 5 | Miss trials lack a conclusive decision signal. a, Percentage of Miss trials across all sessions ($n = 66$ sessions, 13 mice, Figs. 1 and 2, mean ± s.d.). Gray open circles denote individual sessions. **b**, Example session to show classification of Miss trials into Miss Stimulus– (linear classifier trained with stimulus neuron activity predicts stimulus type at chance level) and Miss Stimulus+ trials (linear classifier trained with stimulus neuron activity predicts stimulus). Dotted line represents z-score = 1.64. Gray distribution are classification scores achieved by a linear classifier trained with shuffled trial labels (500×). **c**, Top: distribution of Miss Stimulus+ and Miss Stimulus– trials across the normalized session length ($n = 66$ sessions). Miss Stimulus– trials without stimulus information ramp up towards the end of the session. Bottom: whisking before the go cue in Correct, Miss Stimulus– and Miss Stimulus+ trials in a subset of sessions with whisker kinematics. $n = 14$ sessions, mean ± s.d., two-sided Wilcoxon signed-rank test. **d**, Left: Mean trial activity in stimulus neurons (top) and decision neurons (bottom) in Miss Stimulus– and Miss Stimulus+ trials relative to activity in Correct trials. Mean ± s.e.m., $n = 51$ sessions. Two-sided Wilcoxon signed-rank test. Right: prediction accuracy of a classifier trained to predict stimulus type with stimulus neuron activity (top) or decision neuron activity (bottom) in Correct trials. Decision neurons carry no trial information in Miss Stimulus– and Miss Stimulus+ trials and the classifiers perform at chance level in Miss trials. Mean ± s.e.m., $n = 57$ sessions. Gray open circles denote individual sessions. Testing against chance: $n = 57$ sessions, Stimulus neurons: Miss Stimulus– trials versus chance, $P = 1.1 \times 10^{-8}$; Miss Stimulus+ trials versus chance, $P = 2.3 \times 10^{-10}$. Decision neurons: Miss Stimulus– trials versus chance, $P = 0.37$; Miss Stimulus+ trials versus chance, $P = 0.08$. Two-sided Wilcoxon signed-rank test.

target ensemble, with respect to the stimulus and choice selectivity in the target ensemble (Extended Data Fig. 9a).

We found that the stimulus selectivity of followers that showed a positive (or negative) response to photostimulation is positively (or negatively) correlated with the average stimulus selectivity of the target ensemble, whereas the choice selectivity of targets and positive followers was not correlated (Fig. 6f and Extended Data Fig. 9b). This implies that the network encoding the stimulus involves more like-to-like functional connectivity, where neurons encoding the same stimulus are connected more strongly than the local circuit encoding the decision. With a difference in intragroup functional connectivity, we wondered if local activation of our target ensembles triggers different responses in stimulus and decision neurons. The global effect of photostimulation on background neurons across the FOV is suppressive[43] (Fig. 6g). When dissecting the effect of photostimulation on stimulus and decision neurons, however, we found that, whereas stimulus neurons follow the overall trend, decision neurons on average showed a positive response to photostimulation (Fig. 6g). This indicates that decision neurons might receive more net excitatory inputs when local targets are activated than stimulus neurons. This cannot be explained by a bias in the target ensembles because the targets and background cells tend to be more selective to stimulus than choice (Extended Data Fig. 9c). Taken together, the network effects following targeted photostimulation suggest a higher like-to-like functional connectivity between neurons encoding the same stimulus, while choice-selective neurons show lower recurrent connectivity but can be more effectively recruited by the local network.

Targeted optogenetic activation modulates behavior. Next, we tested how targeted photostimulation affects behavioral outcome during the task. Here, we stimulated the same ensembles as before, but in trials with texture presentation ('Texture trials'; Fig. 7a). We either stimulated the target ensemble whose trial preference aligns with the correct choice (PhotoBoost) or the target ensemble that prefers the opposite trial type (PhotoDisrupt) with a brief photostimulation during the withhold period (10.5 ± 3.9 targets per photostimulation pattern, Fig. 7a,b; Methods). On average, behavioral performance did not change consistently in photostimulation trials compared to control trials (where photostimulation was set to zero) (Extended Data Fig. 10a,b). However, target ensembles across sessions and photostimulation conditions differ with respect to the amount of stimulus and choice selectivity present in the target neurons. We therefore asked whether the stimulus and choice coding properties of the target cells can explain the seemingly heterogeneous behavioral effect. We found that a change in performance is positively correlated with the choice selectivity, but not the stimulus selectivity of the activated targets (Fig. 7c,d).

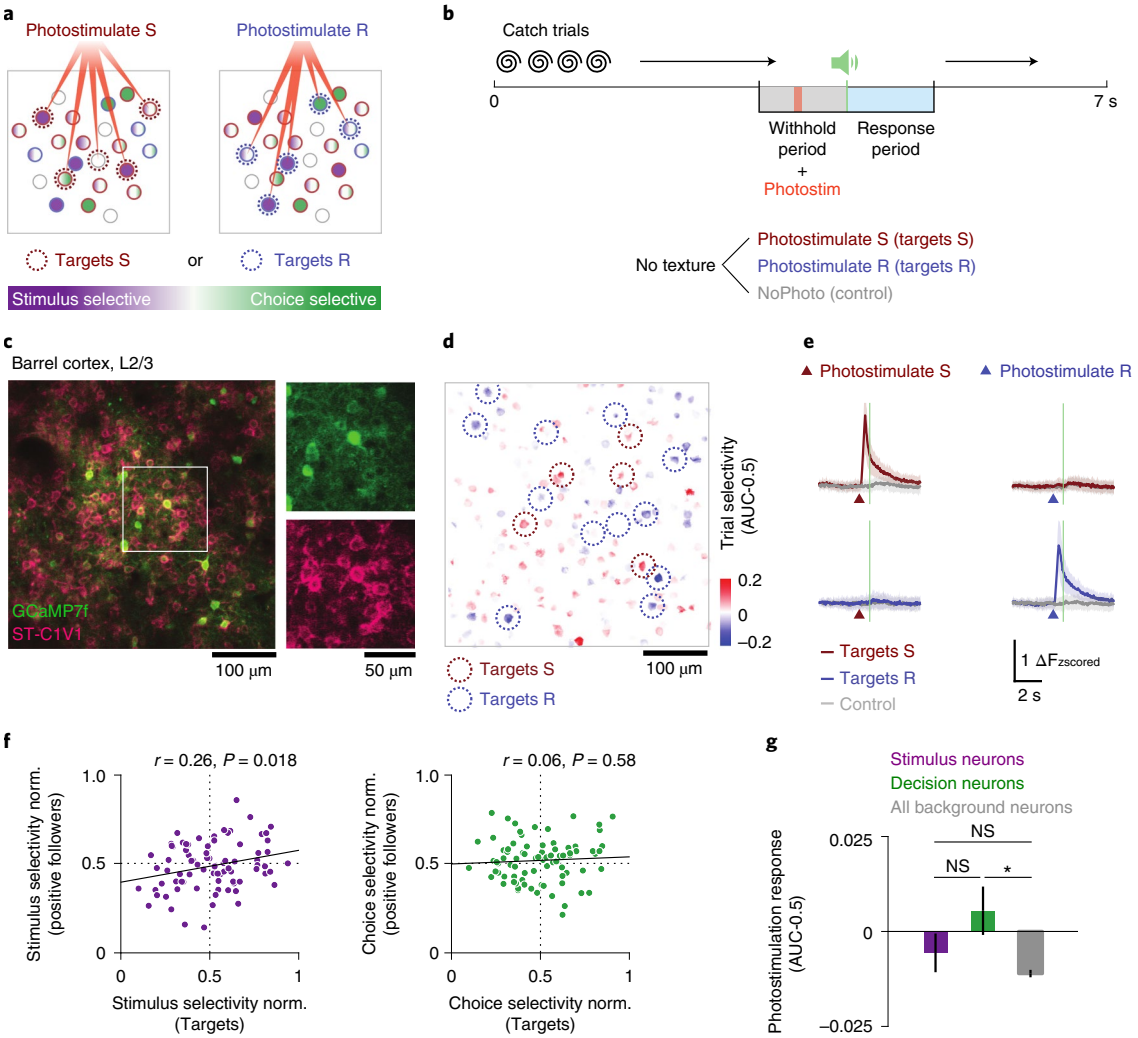

Fig. 6 | Cell-type specific functional connectivity of stimulus and decision coding neurons. a, Schematic of targeted photostimulation experiment. A subset of trial-coding neurons preferring correct Stim S/Lickport S trials or correct Stim R/Lickport R trials was selectively activated by two-photon optogenetic photostimulation. **b**, Task structure for catch trial experiments. Two-photon photostimulation (9 × 10 ms spirals) of either target ensemble was enabled in a subset of trials in the middle of the withhold period. **c**, An example FOV (depth = 200 μm) with neurons coexpressing calcium indicator (GCaMP7f) and soma-targeted opsin (ST-C1V1). Representative results of 31/60 surgeries. **d**, Trial selectivity of neurons in the FOV in **c**. Red and blue circles mark the two trial selective target ensembles. **e**, Calcium timecourses of the two target ensembles during different trial types (mean ± s.d.). $n = 80$ target ensembles, 40 sessions and seven mice. **f**, Left: Stimulus selectivity of the target ensemble is positively correlated with the stimulus selectivity of indirectly photo-activated neurons (positive followers). Right: no significant correlation between the choice selectivity of the followers and that of the targets (r and P are the Pearson correlation coefficient and the P value, respectively). Selectivity was normalized to all other background cells in the same FOV. **g**, Photostimulation response of stimulus and decision neurons (excluding targets) compared with all background cells (boxes are mean, whiskers are s.e.m.). $n = 363$ stimulus, 248 decision and 10,458 background neurons, 40 sessions and seven mice. Kruskal–Wallis one-way ANOVA test with Tukey–Kramer critical values were used for multiple comparisons; not significant (NS), stimulus versus decision neurons $P = 0.53$; NS, stimulus versus background neurons $P = 0.27$; *P stimulus versus background neurons $P = 0.022$.

The more the target population preferred the correct choice on that trial, the greater the positive effect of stimulation. Thus, targeted photostimulation of choice-selective neurons improves behavioral performance selectively, depending on the functional properties of the targeted neurons. The behavioral effect size and sign does not depend on the number or the spatial clustering of the target neurons or the response in the background population (Extended Data Fig. 10c–e). A linear regression model that takes into account the functional selectivity and photostimulation response of individual target neurons, background neuron activity and whisking state of the animal confirms our finding that the behavioral effect of photostimulation depends on the choice selectivity of directly activated targets and not stimulus selectivity or whisking (Extended Data Fig. 10f). These results provide direct evidence that the decision signal in L2/3 barrel cortex is causally involved in decision making.

Discussion

Understanding where and how behaviorally relevant decision signals are present in neural circuits is key to unraveling how sensory information is transformed into behavioral output, one of the most fundamental questions in neuroscience. Using population two-photon imaging and targeted two-photon photostimulation of neurons in S1 of mice performing a cued two-choice texture discrimination task, we have identified a decision signal in L2/3 neurons of S1 that is causally linked to behavior. Decision neurons,

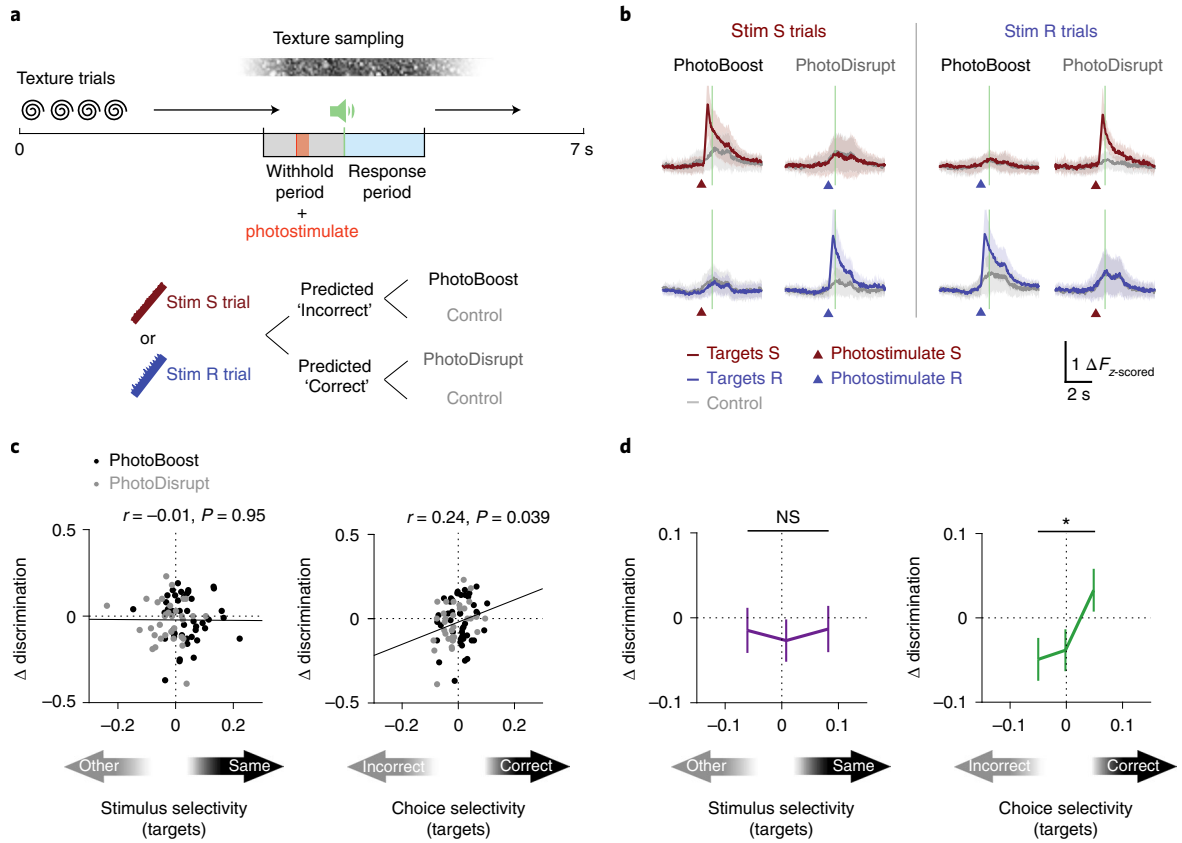

Fig. 7 | Activation of neurons with correct choice selectivity improves behavioral performance. a, Task structure for texture trial experiments. Two-photon photostimulation (9 × 10 ms spirals) of either target ensemble was enabled in a subset of trials in the middle of the withhold period. Before the photostimulation, activity of trial-coding neurons was readout online to generate a prediction on the trial outcome (Predicted Correct or Predicted Incorrect). Trials predicted Incorrect either received PhotoBoost or Control photostimulation. Trials predicted Correct received either PhotoDisrupt or Control photostimulation (Methods). **b**, Calcium timecourses of the two target ensembles during different trial types (mean ± s.d.). $n = 67$ target ensembles, 37 sessions and seven mice. **c**, Photostimulation-induced change in task performance (Δ) is positively correlated with the choice selectivity but not the stimulus selectivity of directly activated targets. $n = 77$ photostimulation conditions, 37 sessions and seven mice. Discrimination = Correct trials/(Correct + Incorrect trials); r and P are the Pearson correlation coefficient and the P value, respectively. **d**, Same data as in **c** but binned into three groups by ascending selectivity with equal number of samples per bin. Data are presented as mean ± s.e.m., Wilcoxon rank sum test, $^*P = 0.038$.

a small but distinct subgroup, carry the strongest decision signal within the population. The categorical decision signal that develops with learning is temporally aligned to sensory input, and is lacking in trials when animals fail to make a choice. Using targeted photostimulation, we demonstrate that stimulus and decision neurons are differentially integrated into the local circuit, providing independent confirmation that they represent functionally distinct groups of cells. Crucially, targeted optogenetic activation of the decision signal improves behavioral performance and therefore directly links the decision signal to behavior. These findings provide conclusive evidence that a decision signal is present at the level of primary sensory cortex and that it is not merely a copy of a decision signal generated elsewhere in the brain, but rather a signal that is readout by downstream circuits to drive behavior.

Experimental requirements to identify decision signals. There is strong evidence that the details of task design can crucially affect cortical computations as well as the involvement of brain areas necessary for task performance[12,13,23,33,44,45]. Perceptual decision-making tasks differ with respect to the stimuli presented, the number of choices, the actions associated with the choices and delay periods, as well as general environmental features. Task design and species differences might explain why a series of earlier studies in monkey and rodents reported only stimulus coding in S1 (refs. [2,10,11,13–15,25]).

We refer to choice-relevant activity as decision signals to emphasize that this activity is not a motor signal that triggers the choice action, but rather a signal that is used to make a choice. We consider decisions to be internal variables that carry task-relevant information, for example, on the identity of the lickport associated to the stimulus presented, the occurrence of the go cue or the motivational state. Mice may make a decision with respect to the lickport that needs to be licked in a given trial but choose not to lick because the go cue has not occurred, or because they are not motivated to lick.

The hallmarks of the task we used to study decision signals in L2/3 of S1 are complex texture stimuli that require temporal integration by active sampling, two choices that are reported in similar ways and a cued task design with a temporal delay between stimulus and choice action. We favored the discrimination of complex stimuli due to the observation that simple tactile tasks might not critically depend on barrel cortex[46]. Unlike the detection of a touch event by a single whisker, the discrimination of textures of different roughness requires active sampling and temporal integration of complex whisker kinetics such as acceleration and curvature[3,47]. In line with this, lesion studies in rats have shown that the ability to discriminate textures depends on barrel cortex[48]. Texture discrimination might require computations in L2/3 of S1 to convert complex and varying patterns of whisker kinetics into a simpler object-related code. Stimulus neurons, as described in our study, could report the output

of that transformation. The detection of a decision signal in S1 may be related to the fact that task-relevant signal transformation from sensory to an early and abstract decision variable is located in L2/3 S1 under these task conditions.

To identify decision signals in primary sensory cortex, we need to be able to dissociate signals triggered by the choice action from decision signals that inform the choice. In widely used go/no go tasks, only one choice is associated with an action whereas the other choice is signaled by withholding of movement. A consistent activity modulation favoring 'go' trials[16–19] could reflect choice-relevant activity or other aspects related to the choice action in 'go' trials, such as movement preparation, motivation, feedback from the movement or sensory input that results from the movement[22]. In V1, studies have demonstrated the existence of an activity difference between 'go' and 'no go' trials[23] but an absence of choice-related activity in symmetric choice tasks[22,24]. Therefore, go/no go tasks cannot definitively identify choice-related activity. Here, we recorded neural activity in mice performing a cued two-choice texture discrimination task featuring choice actions with nearly identical motor patterns, which allowed us to extract decision signals in S1. We also used large-scale recordings across multiple barrels to detect even small functional subgroups of neurons. Furthermore, we used a GLM to demix running, whisking and choice signals and find that the timing of decision neurons is temporally aligned to the stimulus rather than the lick. We consider it unlikely that decision neurons are indeed sensory neurons weakly correlated with the choice action given the numerous categorical differences we find between stimulus and decision neurons, that is, differences in spatial clustering, temporal activity patterns, trial-by-trial response variability and their functional connectivity within the local network. This distinguishes the stimulus neurons described here from sensory neurons found in, for example, monkey area MT that carry sensory information and show a weak correlation to the choice of the monkey[34,49].

The ultimate support for the claim that the decision signal described here is indeed involved in the decision-making process, and is not a signal carrying motor feedback or sensory input related to the choice action after the decision was made, is provided by our direct manipulation of the decision signal during behavior[50,51]. We find that when we increase the correct decision signal in L2/3 S1 using targeted photostimulation in behaving mice, behavioral performance is improved. This experiment establishes a causal link between neural activity and behavior and confirms a functional role of the decision signal in decision making.

The circuit organization of stimulus and decision neurons. Does the identification of stimulus and decision neurons imply two distinct neuronal subnetworks in S1? Classical literature suggests that decision signals could be generated in a purely feedforward manner along the cortical processing stream by pooling across large groups of weakly choice-selective sensory neurons[51,52]. This implies that stimulus signals and a weak decision signal are encoded in the same neurons and that the strength of the choice signal increases as we move along the processing hierarchy. Here we find that L2/3 of barrel cortex, a supposedly purely sensory processing area, contains neurons with stimulus and choice-related information—more akin to what has been found in associative cortices like the lateral intraparietal area[49]. Even though stimulus and decision neurons are intermingled within the same local circuit they show differences in their local network integration. Stimulus neurons exhibit some spatial clustering, while decision neurons in the same FOV are distributed randomly. Both subgroups show higher trial-by-trial response variability within each group than between groups, which could be due to recurrent connectivity or common feedback[35–37].

Using targeted photostimulation to directly probe the functional connectivity in the local circuit, we find that stimulus selectivity of the photostimulation targets and their followers correlates. This suggests high like-to-like recurrent connections for stimulus coding in L2/3, and is consistent with recent studies in S1, V1 and ALM[29–31,53]. The recurrent excitation in L2/3 S1 may support stimulus identification by maintaining and amplifying the 'bottom-up', feedforward inputs from L4 containing information about the stimulus features relevant to the task. In contrast, choice coding of the followers does not depend on that of the targets, which implies that the recurrent connectivity among neurons representing the same decision might be lower than those encoding the same stimulus. The shared trial-by-trial response variability between decision neurons might in this case be driven by common inputs from the local circuit, with contributions from long-range feedback connections from S1, M1, M2 and higher-order thalamus and neuromodulatory inputs[38,54]. This idea is further supported by higher activity of decision neurons outside of the sampling period in comparison to stimulus neurons, as well as the existence of more efficient excitatory inputs from local trial-coding neurons when activating trial-coding neurons. Taken together, although there is likely to be a continuum with respect to stimulus and choice coding across L2/3 S1, the ends of this putative continuum, that is, stimulus and decision neurons, show categorical differences not only in the trial-coding dimension, but also in spatial clustering, activity patterns, timing, network integration and functional connectivity.

Implications for sensory processing. The classical concept that hierarchical processing involves a step-by-step feature extraction before a decision is made in higher-order areas has been extremely influential in guiding the field to explore the computations performed during sensory processing[1,9,55]. Recent studies have suggested the possibility that choice coding might already be present in L2/3 of S1 but were either lacking appropriate task design or direct manipulation to establish causal relevance of these signals for behavior. Our results show that a decision signal can indeed be present as early as L2/3 of S1 and, crucially, that the signal in decision neurons is causally relevant for task performance. This decision signal seems to represent a transformation of the sensory signal into choice categories. Freedman and Assad suggested the existence of such a categorical and abstract early decision signal in the primate brain, in contrast to intentional decision signals that are directly linked to the motor action[56]. We suggest that the signal carried by decision neurons in L2/3 S1 might represent such an abstract categorical decision signal.

Instead of just feature extraction following a strict processing hierarchy, primary sensory cortex might be passing on information about extracted features in parallel with a categorical and task-relevant interpretation of the sensory signal. Following the concept suggested by Panzeri and colleagues[57], this would mean that decision neurons carry features with a high intersection information and are therefore involved directly in driving behavior, whereas stimulus neurons only present the stimulus identity and carry little intersection information. Alternatively, as suggested by Churchland and colleagues[58], feature extraction along the sensory pathway may not be isolated but could be modulated by top-down projections very early in the pathway. Interim decision variables could then influence downstream feature extraction as well as modulate behavior such as whisking to optimize task performance. The decision signals we describe could reflect the output of an intersection between sensory feature extraction and task-specific top-down modulation entering L2/3 S1. Studying the task-dependence of our results, the influence of top-down modulation on decision neuron activity as well as the temporal relationship between stimulus and decision neurons measured with high-density electrophysiological recordings will provide many avenues for future experiments.

In summary, our results demonstrate the existence and behavioral relevance of decision coding as early as S1, challenging the

longstanding idea that S1 provides information on sensory inputs only[2,10,11,13,14]. Instead, an early decision signal carried by neurons intermingled with stimulus neurons within the same circuit, which has a direct impact on behavior suggests a key involvement of S1 in the decision-making process through the encoding of an early categorical decision signal.

Online content

Any methods, additional references, Nature Research reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41593-022-01151-0.

Received: 2 May 2019; Accepted: 21 July 2022;
Published online: 30 August 2022

References

1. Hubel, D. H. & Wiesel, T. N. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol.* **160**, 106–154 (1962).
2. Hernandez, A., Zainos, A. & Romo, R. Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc. Natl Acad. Sci. USA* **97**, 6191–6196 (2000).
3. Arabzadeh, E., Zorzin, E. & Diamond, M. E. Neuronal encoding of texture in the whisker sensory pathway. *PLoS Biol.* **3**, e17 (2005).
4. Romo, R. & de Lafuente, V. Conversion of sensory signals into perceptual decisions. *Prog. Neurobiol.* **103**, 41–75 (2013).
5. Brody, C. D. & Hanks, T. D. Neural underpinnings of the evidence accumulator. *Curr. Opin. Neurobiol.* **37**, 149–157 (2016).
6. Hanks, T. D., Ditterich, J. & Shadlen, M. N. Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* **9**, 682–689 (2006).
7. Kim, J.-N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature* **2**, 176–185 (1999).
8. Horwitz, G. D. & Newsome, W. T. Separate signals for target selection and movement specification in the superior colliculus. *Science* **284**, 1158–1161 (1999).
9. Barlow, H. B. in *Sensory Communication* (ed. Rosenblith, W. A.) Ch. 13 (MIT Press, 2012).
10. Rossi-Pool, R. et al. Emergence of an abstract categorical code enabling the discrimination of temporally structured tactile stimuli. *Proc. Natl Acad. Sci. USA* **113**, E7966–E7975 (2016).
11. de Lafuente, V. & Romo, R. Neural correlate of subjective sensory experience gradually builds up across cortical areas. *Proc. Natl Acad. Sci. USA* **103**, 14266–14271 (2006).
12. Romo, R., Hernandez, A., Zainos, A., Lemus, L. & Brody, C. D. Neuronal correlates of decision-making in secondary somatosensory cortex. *Nat. Neurosci.* **5**, 1217–1225 (2002).
13. Fassihi, A., Akrami, A., Pulecchi, F., Schonfelder, V. & Diamond, M. E. Transformation of perception from sensory to motor cortex. *Curr. Biol.* **27**, 1585–1596 (2017).
14. Guo, Z. V. et al. Flow of cortical activity underlying a tactile decision in mice. *Neuron* **81**, 179–194 (2014).
15. McGuire, L. M. et al. Short time-scale sensory coding in S1 during discrimination of whisker vibrotactile sequences. *PLoS Biol.* **14**, e1002549 (2016).
16. Sachidhanandam, S., Sreenivasan, V., Kyriakatos, A., Kremer, Y. & Petersen, C. C. Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat. Neurosci.* **16**, 1671–1677 (2013).
17. Kwon, S. E., Yang, H., Minamisawa, G. & O'Connor, D. H. Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. *Nat. Neurosci.* **19**, 1243–1249 (2016).
18. Yang, H., Kwon, S. E., Severson, K. S. & O'Connor, D. H. Origins of choice-related activity in mouse somatosensory cortex. *Nat. Neurosci.* **19**, 127–134 (2016).
19. Chen, J. L., Carta, S., Soldado-Magraner, J., Schneider, B. L. & Helmchen, F. Behaviour-dependent recruitment of long-range projection neurons in somatosensory cortex. *Nature* **499**, 336–340 (2013).
20. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
21. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
22. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
23. Poort, J. et al. Learning enhances sensory and multiple non-sensory representations in primary visual cortex. *Neuron* **86**, 1478–1490 (2015).
24. Scott, B. B. et al. Imaging cortical dynamics in GCaMP transgenic rats with a head-mounted widefield macroscope. *Neuron* **100**, 1045–1058 (2018).
25. Brecht, M., Roth, A. & Sakmann, B. Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. *J. Physiol.* **553**, 243–265 (2003).
26. Nienborg, H. & Cumming, B. Correlations between the activity of sensory neurons and behavior: how much do they tell us about a neuron's causality. *Curr. Opin. Neurobiol.* **20**, 376–381 (2010).
27. Katz, L. N., Yates, J. L., Pillow, J. W. & Huk, A. C. Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* **535**, 285–288 (2016).
28. Rickgauer, J. P., Deisseroth, K. & Tank, D. W. Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nat. Neurosci.* **17**, 1816–1824 (2014).
29. Carrillo-Reid, L., Han, S., Yang, W., Akrouh, A. & Yuste, R. Controlling visually guided behavior by holographic recalling of cortical ensembles. *Cell* **178**, 447–457(2019).
30. Daie, K., Svoboda, K. & Druckmann, S. Targeted photostimulation uncovers circuit motifs supporting short-term memory. *Nat. Neurosci.* **24**, 259–265 (2021).
31. Marshel, J. H. et al. Cortical layer-specific critical dynamics triggering perception. *Science* **365**, eaaw5202 (2019).
32. Russell, L. E. et al. All-optical interrogation of neural circuits in behaving mice. *Nat. Protoc.* **17**, 1579–1620 (2022).
33. Peron, S. P., Freeman, J., Iyer, V., Guo, C. & Svoboda, K. A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron* **86**, 783–799 (2015).
34. Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A. A relationship between behavioral choice and the visual response of neurons in macaque MT. *Vis. Neurosci.* **13**, 87–100 (1996).
35. Okun, M. et al. Diverse coupling of neurons to populations in sensory cortex. *Nature* **521**, 511–515 (2015).
36. Ko, H. et al. Functional specificity of local synaptic connections in neocortical networks. *Nature* **473**, 87–91 (2011).
37. Shadlen, M. N. & Newsome, W. T. The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896 (1998).
38. Chen, J. L. et al. Pathway-specific reorganization of projection neurons in somatosensory cortex during learning. *Nat. Neurosci.* **18**, 1101–1108 (2015).
39. Keller, G. B. & Mrsic-Flogel, T. D. Predictive processing: a canonical cortical computation. *Neuron* **100**, 424–435 (2018).
40. Russell, L. E. et al. The influence of visual cortex on perception is modulated by behavioural state. Preprint at *bioRxiv* https://doi.org/10.1101/706010 (2019).
41. & Dalgleish, H. W. P. et al. How many neurons are sufficient for perception of cortical activity?. *eLife* **9**, e58889 (2020).
42. Robinson, N. T. M. et al. Targeted activation of hippocampal place cells drives memory-guided spatial behavior. *Cell* **183**, 1586–1599 (2020).
43. Chettih, S. N. & Harvey, C. D. Single-neuron perturbations reveal feature-specific competition in V1. *Nature* **567**, 334–340 (2019).
44. Hong, Y. K., Lacefield, C. O., Rodgers, C. C. & Bruno, R. M. Sensation, movement and learning in the absence of barrel cortex. *Nature* **561**, 542–546 (2018).
45. Kawai, R. et al. Motor cortex is required for learning but not for executing a motor skill. *Neuron* **86**, 800–812 (2015).
46. Park, J. M. et al. Deep and superficial layers of the primary somatosensory cortex are critical for whisker-based texture discrimination in mice. Preprint at *bioRxiv* https://doi.org/10.1101/2020.08.12.245381 (2022).
47. Wolfe, J. et al. Texture coding in the rat whisker system: slip-stick versus differential resonance. *PLoS Biol.* **6**, e215 (2008).
48. Guic-Robles, E., Jenkins, W. M. & Bravo, H. Vibrissal roughness discrimination is barrelcortex-dependent. *Behav. Brain Res.* **48**, 145–152 (1992).
49. Shadlen, M. N. & Newsome, W. T. Motion perception: seeing and deciding. *Proc. Natl Acad. Sci. USA* **93**, 628–633 (1996).
50. Nienborg, H., Cohen, M. R. & Cumming, B. G. Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu Rev. Neurosci.* **35**, 463–483 (2012).
51. Parker, A. J. & Newsome, W. T. Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* **21**, 227–277 (1998).
52. Shadlen, M. N., Britten, K. H., Newsome, W. T. & Movshon, J. A. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* **16**, 1486–1510 (1996).

53. Peron, S. et al. Recurrent interactions in local cortical circuits. *Nature* **579**, 256–259 (2020).
54. Condylis, C. et al. Context-dependent sensory processing across primary and secondary somatosensory cortex. *Neuron* **106**, 515–525 (2020).
55. Felleman, D. J. & Van Essen, D. C. Distributed hierarchical processing in the primate. *Cereb. Cortex* **1**, 1–47 (1991).
56. Freedman, D. J. & Assad, J. A. Neuronal mechanisms of visual categorization: an abstract view on decision making. *Annu Rev. Neurosci.* **39**, 129–147 (2016).
57. Panzeri, S., Harvey, C. D., Piasini, E., Latham, P. E. & Fellin, T. Cracking the neural code for sensory perception by combining statistics, intervention, and behavior. *Neuron* **93**, 491–507 (2017).
58. Churchland, P. S., Ramachandran, V. S. & Sejnowski, T. J. in *Large-Scale Neuronal Theories of the Brain* (eds Koch, C. & Davis, J. L.) Ch. 2 (MIT Press, 1994).

Methods

Mice. All animal procedures approved by the local Animal Welfare and Ethical Review Board at University College London and performed under license from the UK Home Office in accordance with the Animals (Scientific Procedures) Act 1986. We used transgenic mice expressing GCaMP6s[59] in excitatory neurons (EMX1-Cre; CaMKIIa-tTA; Ai94 (Jax catalog nos. 027784, 007004 and 024104) and CaMKIIa-tTA; tetO-G6s (Jax catalog nos. 007004 and 024742)) for calcium imaging experiments and C57Bl/6J (Charles River Laboratories) for targeted photostimulation experiments. All EMX1-Cre; CaMKIIa-tTA; Ai94 mice were treated with doxycycline administered orally in the drinking water to prevent expression of GCaMP6s during development[60] and were checked for aberrant activity (Wide-field imaging). A solution of 5% sucrose and 2 mg ml^{-1} doxycycline in tap water was given as drinking water from birth up until 7 weeks of age. Mouse age ranged from 8 to 36 weeks on the days of experiments and were of either sex. Mice were group-housed before surgery and single-housed after surgery. Mice were kept at a normal or reversed 12 h dark/light cycle at a temperature of 22 °C and 62% humidity.

Surgery. Mice were anesthetized using Isoflurane (0.5–2%) and injected with 0.1 mg kg^{-1} buprenorphine hydrochloride (Vetergesic) and 5 mg kg^{-1} Carprofen (Rimadyl). A metal headplate with a 5 mm circular imaging well was fixed to the skull overlying somatosensory cortex with dental acrylic (Super-Bond C&B, Sun-Medical). A craniotomy was drilled above S1 (right hemisphere, 1.6 mm posterior and 3.5 mm lateral of bregma). For targeted photostimulation experiments, 0.9–1 µl of a mixture of AAV1-syn-jGCaMP7f-WPRE[61] (dilution 1:15), AAV9-FLEXED-C1V1-Kv2.1-mRuby2 (ref. [43]) (dilution 1:30) and pENN-AAV-CaMKII-0.4Cre-SV40 (dilution 1:15) at a ratio of 2:1:2 was injected into L2/3 (~300 µm deep at 0.1 µl min^{-1}). The dura was removed after virus injection. A cranial window, composed of either a 3 mm circular glass coverslip glued to a 2 mm square glass with UV-curable optical cement (NOR-61, Norland Optical Adhesive) or a 4 mm circular glass coverslip was press-fit into the craniotomy and sealed using Vetbond before fixing it with dental acrylic. At the earliest 1 week after the surgery mice were water restricted to increase motivation for task training.

Wide-field imaging. After a minimum of 1 week after surgery wide-field imaging was performed at 15 Hz using a 470 nm LED (M470L3, Thorlabs) to illuminate the area. Imaging was performed using an ORCA-Flash 4.0 v.3 (Hamamatsu) camera and a ×4 microscope objective (×4 Nikon Plan Fluorite Imaging Objective, 0.13 numerical aperture (NA), 17.2 mm working distance). Excitation light passed through an aspheric condenser lens (ACL2520U-DG15, Thorlabs), a filter (ET470/40, Chroma) and was reflected into the light path by a 495 nm longpass dichroic (FF495-Di03-25×36, Semrock) to reach the brain. Emitted light passed through the same 495 nm longpass dichroic as well as a 749 nm shortpass dichroic (FF749SDi01-25×36×3, Semrock) and an emission filter (HQ525/50, Chroma) before reaching the camera. Spontaneous activity of awake mice running freely on a treadmill was acquired for 2–4 min to exclude the occurrence of aberrant activity in transgenic mice[60]. Calcium signal time series were obtained from the average pixel intensity.

To locate the barrels in the FOV, single whiskers were threaded into a glass capillary and deflected 10–30× for 1 s with a piezo oscillating at 10 Hz (sinusoidal) every 10 s while performing wide-field imaging as described above. To locate the barrel centers, we calculated event-triggered averages of the movies by aligning the frames to the start of the whisker deflections using custom-written Python scripts.

High-speed videography. We used a Mako U-029B (Allied Vision) camera to record whisker movements under IR illumination at 100–200 Hz. The data were streamed and recorded using the software package 2ndLook (IO Industries). We used DeepLabCut v.1.0 (ref. [62]) to label the two most prominent whiskers in every frame across a subset of sessions (nine mice, two sessions each). We calculated the average whisker angle based on the average of the angle of the two whiskers in each frame to reflect the movement of the whisker pad. To extract the first whisker touch we trained another set of models to detect the texture stimulus in the movie. First whisker touch was defined as the first imaging frame of the trial in which texture stimulus and whiskers were closer than 22 pixels (based on visual inspection).

Texture discrimination task. Behavioral training was carried out in training boxes or under the two-photon microscopes which were equipped with the same training apparatus[63]. A 2.5×2.5 cm^2 piece of sandpaper (DOM's DIY direct, Amazon) was attached to four of the six arms of a custom-built plastic hexagon array with protruding arms. Sandpapers used for training and imaging were S of P1200 or P3000, R of P60 or P100 and S2 of 400 or 600, R2 of 180. To adjust the ratio of Correct and Incorrect trials, sandpapers used in targeted photostimulation experiments varied with S of P180/320/600/1500 or 3000 and R of 60 or 120. To rotate the selected sandpaper into position the hexagon was mounted onto a rotating stepper motor (X-NMS17C-PTB2, Zaber Technologies Inc.). The number of steps of the rotating stepper motor to rotate the selected sandpaper into position was matched, such that every transition featured the same number of movements. Once it reached its final position, the sandpaper was moved towards the mouse by mounting the stepper motor onto a linear stage (X-LSM050B-KX12AG or LSM100B-T4-MC04 with X-MCB1-KX12BG Controller, Zaber Technologies Inc.). Both motors were controlled using a microcontroller and Arduino software (Arduino IDE). The mouse was allowed to lick one of the two lickports in front of it with the sound of go cue (100–200 ms, 6 kHz). Licks were detected by closing an electrical circuit. All signals were controlled and displayed using the custom-written software PyBehavior (Lloyd Russell, https://github.com/llerussell/PyBehavior). Licking the correct lickport triggered a solenoid valve (225PNC1-11, NResearch) to release water rewards consisting of 2–4 µl drops of sugar water (10% sucrose in drinking water). Licking in the withhold window 1 s before the go cue triggered 1 s of white noise as well as a 3–5 s time-out. Mice were in the dark and freely running on a circular Styrofoam treadmill during training and imaging sessions. The behavior and imaging data were synchronized using PackIO[64]. In targeted photostimulation experiments, the motor carrying the sandpaper was moved in at a slightly slower speed and the go cue was triggered at 4 s into the trial instead of 3 s.

Training was performed in stages. After habituation to the trainer and the head fixation, water drops were delivered until mice readily licked both lickports. In the next training stage, the sandpapers were presented to the whiskers followed by an auto-reward 200 ms after the go cue to passively pair the texture with the lickport. Textures were presented in blocks with decreasing repeats. Once a mouse respected the withhold period and started to lick the lickports after the go cue the auto-reward was removed. Mice that were licking the lickports after the go cue and showed signs of correctly predicting the lickport that is associated with the presented texture were moved on to the final training stage. Here, the textures were presented randomly with equal probability. Trials lasted for 7 s with a variable delay of 0–1 s in between trials. Mice were trained or imaged no more than once per day until most Miss trials indicated satiety.

Two-photon calcium imaging. Two-photon calcium imaging was performed on a custom-build microscope equipped with a resonant scanning module (LSK-GR08, Thorlabs), GaAsP photomultiplier tube (Hamamatsu) and a ×16, 0.8 NA microscope objective (Nikon) using ThorImage 3.1 (Thorlabs). We used a Ti-Sapphire laser (Mai Tai HP, Spectra Physics) to excite GCaMP6 at 920 nm. The FOV was 798×798 µm^2 (512×512 pixels) and images were acquired at 30 Hz for single-plane recordings and at 5 Hz for multiplane recordings. For multiplane recordings, the objective was moved between frames using a piezo objective scanner (PFM450E, Thorlabs). The piezo was allowed to settle for 35 ms in the new z position before the next frame was recorded.

Targeted photostimulation and simultaneous two-photon calcium imaging. Simultaneous calcium imaging and photostimulation was performed using an 'all-optical' microscope[32,65]. Two-photon imaging (512×512 pixels per frame, 30 Hz) of L2/3 barrel cortex (around 100–300 µm deep) was performed by resonant-galvanometer raster scanning a femtosecond-pulsed laser beam (wavelength 920 nm or 765 nm, Chameleon Ultra II, Coherent, output power 2–4 W) across FOVs of 416×416 µm^2. A ×16, 0.8 NA objective (Nikon) was used for all experiments. GCaMP7f was imaged with an excitation wavelength of 920 nm and mRuby (fluorophore coupled with opsin) with 765 nm (power on sample, 40–50 mW). The excitation source for the photostimulation path was a femtosecond-pulsed laser fixed at 1,030 nm (Satsuma HP2, Amplitude; average output, 20 W; pulse width, 280 fs; repetition rate, 2 MHz). A reflective multilevel phase spatial light modulator was used to display holograms (OverDrive Plus SLM, Meadowlark Optics/Boulder Nonlinear Systems; 7.68×7.68 mm^2 active area, 512×512 pixels, optimized for 1,064 nm). The weighted Gerchberg-Saxton algorithm[66] was used to calculate holograms to be displayed on the SLM and the weights were adjusted to compensate for the difference in diffraction efficiency between holographic spots. An acousto-optic modulator (AOM, MCQ80-A2-L1064-Z32) was used for modulating photostimulation power.

A custom Python-based, real-time all-optical interface (pyRTAOI, written in Python v.3.6 with PyQt5, developed in Spyder v.3.2.4) was used for online image analysis and photostimulation control. Raw data streaming from the microscope system (Prairie View, Bruker Corporation) was first organized into 512×512 frames accelerated by a GPU (NVIDIA GeForce GTX 750 Ti) and then passed to the calcium imaging analysis toolbox, CaImAn (2018 version)[67], for online motion-correction, neuronal signal extraction and denoising. Photostimulation pattern and power were controlled by a custom SLM control interface (written in C++ with Qt 5.9, developed in Microsoft Visual Studio 2013). The power control is synchronized with spiral scanning by the microscope system by voltage pulse triggers.

Calcium traces extracted online by CaImAn were used for quick analysis on the day of experiments. Quality of each region of interest (ROI) was assessed with a convolutional neural network based method[67]. Cells were first detected from the live imaging stream while the mouse started performing the task (normally around 150–250 ROIs were detected in 15–25 min). Imaging and behavioral recording started after the animal made consecutive correct transitions between the two lickports and the rate of Early lick trials decreased. Before experiments, we trimmed the whiskers, leaving three to six principal whiskers corresponding to the barrels in the FOV with coexpression of GCaMP and C1V1 to minimize task-relevant activity outside of the photostimulatable area (Fig. 6c). First, a

baseline imaging session consisting of 150 trials was used for mapping functional identity of cells in the FOV and for building choice decoders. Next, to test which trial-coding neurons can be activated optogenetically, a short photostimulation pulse (9×10 ms spirals, 15–20 µm in diameter, 7–8 mW per cell) was sent onto each texture-selective cell in a random sequence (0.5 s interstimulus interval, 10–15 repeats per cell). The stimulation power was chosen such that the two-photon activation efficacy matches that in the calibrated data[65] (20–80 mW, pulse repetition rate, 80 MHz, pulse width, 300 fs; Supplementary Figs. 2 and 3). Trial-coding neurons that show a reliable calcium response to photostimulation (*z*-scored area under the ROC curve (AUC) > 1 and response amplitude > 1) were grouped by trial preference into two functional target ensembles (Extended Data Fig. 8). In a subset of sessions, extra photostimulation conditions were added in the catch trials, where photostimuli were targeted directly at identified stimulus or choice-selective neurons to increase the range of choice and stimulus selectivity in the target ensemble. In a random set of catch trials (around 15–20% of all trials in a session), either of the two target ensembles was selectively stimulated 600 ms before go cue (9×10 ms spirals, 7–8 mW per cell; Fig. 6).

During texture trials (Fig. 7), calcium activity recorded from trial-coding neurons was projected onto a choice decoder that attempts to predict the upcoming choice that the animal is going to make. Specifically, two-choice decoders (for S and R trials, respectively) were built based on logit models fit to the first 4–5 demixed principal components of the calcium activity of trial-coding neurons in the withhold window (dPCA[68]). In each trial, the decoder prediction will be queried during a short time-window in the withhold period (five imaging frames, 600 ms to 433 ms before go cue; Extended Data Fig. 10g). In half of the trials when the choice was predicted to be incorrect during the query window, a brief photostimulation (9×10 ms spirals, 7–8 mW per cell, 567 ± 41 ms before go cue, mean ± s.d.) was delivered to the target ensemble that prefers the same texture as that being presented to the animal (PhotoBoost). In one-third of the trials when the choice was predicted to be correct (that is, the absence of an incorrect prediction during the query window), photostimulation of the same power and duration will be delivered immediately at the end of the query window to the target ensemble that encodes the other texture which is associated with the incorrect choice (PhotoDisrupt). To increase the chance of a behavioral effect despite limitations in trial numbers, we used a classifier for online prediction of the trial outcome based on trial-coding neuron activity before the photostimulation[40,69]. Using this classifier, it was slightly more likely for PhotoBoost photostimulation to occur in Incorrect trials and PhotoDisrupt photostimulation to occur in Correct trials. In five sessions, the conditions were swapped with respect to the PhotoBoost and PhotoDisrupt stimulation, meaning that the target ensemble preferring the texture associated with the incorrect (or correct) choice was activated when the predicted choice is about to be correct (or incorrect). Overall, the online predictor performance does not correlate with the targets' choice selectivity (Extended Data Fig. 10h). Again, this was done to extend the sampling range of choice and stimulus selectivity of the target ensemble. The online choice prediction accuracy is slightly better than chance (the fraction of Correct trials in the control trials for PhotoDisrupt is 4% higher than random, $P = 0.016$; the fraction of Incorrect trials in the control trials for PhotoBoost is 15% higher than random, $P = 0.009$, Wilcoxon signed-rank test). Catch trials and texture trials were randomly interleaved in a session.

To rule out the possibility that the animal might be able to hear the photostimulation laser scanners, spiral scanning started 1 s before the withhold period and continued until the end of the reward period in every trial, such that the only difference between photostimulation and corresponding control trials is the photostimulation laser power controlled by the AOM via an analog output device (NI-DAQmx. PCI-6713, National Instruments).

Data analysis. *Data processing, session selection and trial selection.* For Figs. 1–5, ROIs corresponding to individual neurons were selected and calcium signals were extracted, neuropil-corrected and deconvolved using Suite2P for MATLAB[70]. All further measures of neural activity are based on deconvolved calcium traces unless indicated otherwise. Behavioral data were analyzed with custom-written MATLAB scripts. Imaging sessions were selected based on imaging depth, behavioral performance and number of trials. All imaging sessions were recorded at a depth between 100 µm and 200 µm below the surface (Supplementary Fig. 4). Sessions were included if mice performed more than 80% discrimination (Correct/(Correct + Incorrect)) in at least one window of 50 trials with at least 12 response trials (Correct and Incorrect), which indicates that they are proficient in the task. Every mouse that reached this criterion at least once is considered an expert mouse. Every imaging session included had at least 20 Correct trials and 5 Incorrect trials for Stim S and Stim R each. The difference in timing of the first lick with respect to the trial start when comparing Stim S and Stim R trials was below 500 ms. Unless otherwise indicated, trials that fall within a period of strong lickport bias were removed from the analysis. This is done to remove variability in the population of Correct trials and reduce the number of false-positive Correct trials that are random licks of a certain lickport and not a lick that was triggered as a result of perceptual decision making. Behavioral discrimination (Correct trials/(Correct + Incorrect trials)) is higher in unbiased trials compared with biased trials but the procedure has only a very small effect on the identification of stimulus and decision neurons (Supplementary Fig. 4). To identify lickport bias we analyzed periods of ten trials in a sliding window. Whenever there were more than five trials that were either Correct or Incorrect (that is with lick responses only in the response window) but more than 90% of all first licks were directed to the same port—this period of ten trials was labeled as biased. Trials on which a reward was triggered manually for training purposes were removed from the analysis.

For Figs. 6 and 7, calcium traces from all imaging sessions recorded on the same day were concatenated, extracted and neuropil-corrected using Suite2P. Calcium traces were *z*-scored and normalized to the baseline level for each trial (1 s from the beginning of the trial) to obtain $\Delta F_{z\text{-scored}}$. Cells with spatial footprints more than 30% overlapping with a 40-pixel (32.5 µm) diameter circular mask centered at the centroids of the photostimulation spirals were taken as light-targeted cells (Extended Data Fig. 8b). Sessions with overall texture discrimination (number of Correct trials/(number of Correct trials + number of Incorrect trials)) below 0.6 were excluded from analysis. All trials in the baseline behavioral session and one-quarter of the control trials randomly selected in the stimulation behavioral session were used for evaluating functional identity (that is, trial, stimulus and choice selectivity) of each neuron. The remaining trials in the all-optical session were used for assessing photostimulation and behavioral effects. Photostimulation or control conditions with fewer than eight trials where the animal made a choice during the reward period were excluded from behavioral analysis. Signals in frames during photostimulation (90 ms) may contain light-stimulation artifacts and thus were removed and linearly interpolated with values pre- and postphotostimulation. Frames in a 300-ms window immediately before the go cue were used for measuring photostimulation response.

All light-targeted cells that showed significantly higher response in the photostimulation trials than the no-photostimulation trials (*z*-scored AUC > 1.64) were taken as directly activated targets (Extended Data Fig. 9f). Photostimulation conditions where the number of directly activated cells is less than 40% of the number of target spots were excluded from analysis. Background cells (that is, all cells outside the light-targeted zones) that showed significantly higher or lower response in photostimulation trials compared with no-photostimulation trials (*z*-scored AUC > 1.64 or <−1.64) were taken as positive or negative followers, respectively (on average, 7.6 ± 4.1 positive followers and 13.9 ± 7.7 negative followers per photostimulation condition, mean ± s.d., Extended Data Fig. 10a).

Neuron type identification. Neurons were classified using ROC analysis as well as a Gaussian GLM in a subset of sessions with whisking and running information. Classification obtained using ROC analysis was used for all figures unless otherwise indicated. For the ROC analysis, we computed an ROC curve by varying the activity (1 s before the lick, Figs. 1–5; or 1 s before go cue, Figs. 6 and 7) threshold and sorting responses into S and R trials. We then calculated the absolute difference of the AUC from 0.5. To evaluate significance, we shuffled the trial labels and calculated 200 AUCs. If the true AUC was more than 1.64 s.d. away from the mean of the shuffled distribution, then the neuron was considered to discriminate S from R trials. Neurons that can discriminate S from R trials only considering Correct trials were classified as 'trial-coding neurons'. Neurons that can discriminate S from R trials in Correct and Incorrect trials and prefer the same stimulus type were classified as 'stimulus neurons'. Preference of the same stimulus is defined as a higher mean firing rate for the same stimulus in Correct and Incorrect trials. Neurons that can discriminate S from R trials in Correct and Incorrect trials and prefer different stimulus types but the same choice were classified as 'decision neurons'. We computed the false-positive rate for each neuron type by shuffling the trial labels 100 times per neuron. The false-positive rate using this ROC analysis is 0.27% ± 0.004% for stimulus neurons and decision neurons and 4.6% ± 0.06% for trial-coding neurons ($n = 77$ FOVs, mean ± s.e.m.).

We built a Gaussian GLM (with identity link function) with four regressors ('stimulus', 'choice', 'running' and 'whisking') to predict the average trial activity 1 s before the lick. For each neuron, we trained ten models using all four regressors in Correct and Incorrect trials with twofold crossvalidation for an estimate of the mean prediction accuracy of the full model. We used ridge regression with a lambda of 0.1. We quantified the performance of the model by calculating the Pearson correlation between the real trial activity and the predicted trial activity. The mean Pearson correlation coefficient of the full model was then compared with the mean Pearson correlation of four reduced models in which one of the regressor vectors was randomized each. To test the unique contribution of each regressor to the trial activity of the neuron, we shuffled the regressor vector 400 times and trained the model as described before. We considered a neuron to code for stimulus, choice, running or whisking if the mean coefficient of the full model was more than 1.64 s.d. higher than the mean coefficient of the reduced model.

To visualize the distribution of significant correlation coefficients (that is, unique regressor associations) across all neurons, we reduced dimensionality using *t*-distributed stochastic neighbor embedding (*t*-SNE). We used the 'tsne' MATLAB function with a Mahalanobis distance metric and otherwise default settings.

In eight out of ten sessions, we identified more decision neurons using the ROC algorithm compared with the GLM analysis. In these sessions, 65% ±3% of neurons identified using GLM analysis were also identified using ROC analysis (mean ± s.e.m., $n = 8$ sessions).

Stimulus and choice selectivity. To compute stimulus selectivity we obtained the AUC values (Stim S versus Stim R) in Correct and Incorrect trials as described above and calculated the difference from 0.5. Each neuron therefore has a stimulus selectivity in Correct and a stimulus selectivity in Incorrect trials. To compute choice selectivity, we obtained the AUC values (Correct versus Incorrect) in Stim S and Stim R trials and calculated the difference from 0.5. Each neuron has a choice selectivity in Stim S and a choice selectivity in Stim R trials.

Shared trial-by-trial response variability. To compute the trial-by-trial variation[71,72] to the same stimulus we subtracted the mean activity in a 1-s window before the lick across all Correct trials with the same stimulus before calculating the pairwise correlation coefficient between all neurons, all stimulus neurons, all decision neurons or between pairs of stimulus and decision neurons. All sessions with more than five stimulus and decision neurons were included.

Spatial clustering. To quantify spatial clustering, we calculated the pairwise spatial distances between stimulus neurons or decision neurons in each FOV with at least five neurons of each type. The mean pairwise spatial distance was compared to the mean pairwise distance of a shuffled distribution (200×) containing an equivalent number of neurons from the entire population that was selected randomly. We then calculated the z-score between the pairwise spatial distance of stimulus or decision neurons and the shuffled distribution. A distribution of z-scores around 0 indicates that pairwise distances are not different from randomly selected neurons in the FOV.

Analysis of the timing of activity. To analyze the timing of activity in stimulus and decision neurons (Fig. 3), only single-plane imaging sessions at 30 Hz were used (nine mice, 35 sessions). Mean-event amplitude was averaged across trials for every neuron and then averaged for every session. Statistics were computed across sessions.

Analysis of four-texture learning sessions. For four mice that were successfully trained on the four-texture discrimination task (Fig. 4), we selected the first four-texture training session as well as the session with best task performance for analysis (expert). Neuron identification was conducted as before using only Correct S and R trials.

Analysis of Miss trials. To separate Miss trials into trials with and without stimulus information in each session (Fig. 5), we trained a linear classifier to predict the stimulus presented in Correct trials based on the activity of stimulus neurons. Correct stimulus R and stimulus S trials were balanced before training. Only sessions with at least five stimulus neurons were included. For each Miss trial, we compared the classification score achieved by the linear classifier to the distribution of classification scores achieved by a linear classifier trained with shuffled trial labels (500×). If the stimulus prediction was better than chance as represented by the shuffled distribution of classification scores, a Miss trial was labeled Miss Stimulus+, that is, a Miss trial with stimulus information in the population of stimulus neurons. Otherwise it was labeled Miss Stimulus–, that is, a Miss trial without stimulus information in the population of stimulus neurons. The number of Miss trials increases towards the end of a session with more Miss Stimulus– trials than Miss Stimulus+ trials later in the session (Fig. 5c and Extended Data Fig. 7a). We used the same classifier to calculate the prediction accuracy in Correct, Miss Stimulus– and Miss Stimulus+ trials based on stimulus neuron activity. We then trained a similar linear classifier using decision neuron activity to predict the trial type. Despite the lack of an active choice, we argue that expert mice can discriminate the textures. Therefore, most trials with a conclusive stimulus signal should lead to a correct choice. If decision neuron activity contained choice information, it should yield a prediction accuracy different from chance. To test the performance of the classifiers against chance, we computed shuffled distributions (100×) by shuffling trial labels and compared the accuracy of the classifiers with the shuffled distribution.

Analysis of behavioral performance following targeted photostimulation. To assess the impact of targeted photostimulation on behavior (Fig. 7), we stimulated the target ensembles as described above. 'PhotoBoost' trials were compared with the other putatively Incorrect trials where the identical triggering procedure was enabled but the photostimulation power was set to zero (control trials). The 'PhotoDisrupt' trials were compared with the other putatively Correct trials as controls (control trials). We used 'leave-one-out' crossvalidation to measure behavioral change and target selectivity with different sets of trials to rule out the possibility that any observed correlation between the change in texture discrimination performance and the functional identity of the activated targets may reflect a natural cofluctuation between the variables. Specifically, for each pair of photostimulation control conditions, we take one trial from the photostimulation trial set and one trial from the control trial set and compare their trial outcome (giving −1, 0 or 1). We then used the rest of the photostimulation and control trials in the two trial sets to find the activated targets and calculated their stimulus or choice selectivity using the trials previously preserved for measuring functional identity (that is, all trials in the baseline behavioral session and one-quarter of the control trials in the all-optical session, which were not used for measuring photostimulation effects). This procedure was repeated 100 times and the average change in discrimination and the average target selectivity were taken as the final result.

Statistics and reproducibility. No statistical methods were used to predetermine sample size. Our sample sizes are similar to those reported in previous publications[19,29,30]. Data analysis was conducted in MATLAB. Statistical tests were done using a Wilcoxon rank sum test (MATLAB 'ranksum'), Wilcoxon signed-rank test for paired data (MATLAB 'signrank') or Kruskal–Wallis one-way ANOVA test for comparison between more than two groups (MATLAB 'kruskalwallis'). No randomization of experimental subjects was necessary as all mice were trained and recorded under the same conditions. Data collection and analysis were not performed blind to the conditions of the experiment, but analysis relied on code that was standardized for all experimental conditions. Data distributions were not assumed to be normally distributed and all statistical comparisons between groups of continuous variables were performed using nonparametric tests. Adjustments for multiple comparisons were not made unless stated otherwise. No mice were excluded from the analysis unless they did not learn the task, data could not be collected due to poor GCaMP expression or occlusion of the chronic window. All other exclusion of trials or sessions are mentioned in the methods section.

Reporting summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

Data availability

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

Code availability

The analysis code used in this study are available from the corresponding authors upon reasonable request.

References

59. Chen, T. W. et al. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300 (2013).
60. Steinmetz, N. A. et al. Aberrant cortical activity in multiple GCaMP6-expressing transgenic mouse lines. *eNeuro* https://doi.org/10.1523/ENEURO.0207-17.2017 (2017).
61. Dana, H. et al. High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nat. Methods* **16**, 649–657 (2019).
62. Mathis, A. et al. DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).
63. Guo, Z. V. et al. Procedures for behavioral experiments in head-fixed mice. *PLoS One* **9**, e88678 (2014).
64. Watson, B. O., Yuste, R. & Packer, A. M. PackIO and EphysViewer: software tools for acquisition and analysis of neuroscience data. Preprint at *bioRxiv* https://doi.org/10.1101/054080 (2016).
65. Packer, A. M., Russell, L. E., Dalgleish, H. W. & Häusser, M. Simultaneous all-optical manipulation and recording of neural circuit activity with cellular resolution in vivo. *Nat. Methods* **12**, 140–146 (2015).
66. Gerchberg, R. W. & Saxton, W. O. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* **35**, 237–246 (1972).
67. Giovannucci, A. et al. CaImAn an open source tool for scalable calcium imaging data analysis. *eLife* **17**, e38173 (2019).
68. & Kobak, D. et al. Demixed principal component analysis of neural population data. *eLife* **5**, e10989 (2016).
69. Zhang, Z., Russell, L. E., Packer, A. M., Gauld, O. M. & Hausser, M. Closed-loop all-optical interrogation of neural circuits in vivo. *Nat. Methods* **15**, 1037–1040 (2018).
70. Pachitariu, M. et al. Suite2p: beyond 10,000 neurons with standard two-photon microscopy. Preprint at *bioRxiv* https://doi.org/10.1101/061507 (2016).
71. Cohen, M. R. & Kohn, A. Measuring and interpreting neuronal correlations. *Nat. Neurosci.* **14**, 811–819 (2011).
72. Averbeck, B. B., Latham, P. E. & Pouget, A. Neural correlations, population coding and computation. *Nat. Rev. Neurosci.* **7**, 358–366 (2006).

Acknowledgements

We are grateful to S. Chun, M. Hoddinott and O. Houghton for help with animal husbandry; A. O'Leary, E. Ordonneau, D. Vardalakis, R. Ratto and O. Houghton for help with training mice; and A. Roth, D. Kostadinov, N. Robinson, M. Fisek, B. Bicknell and M. London for helpful discussions and comments on the manuscript. We thank

L. Russell for providing PyBehavior software and P. Dzialecka for contributing to the initial development of the pyRTAOI real-time all-optical control software. C.B. was supported by an EMBO Long-Term Fellowship and a Marie Curie Fellowship from the EU. M.P. was supported by a LIDO scholarship from the BBSRC. E.B. was supported through the Erasmus+ Traineeship programme by the EU. M.H. was supported by the Wellcome Trust (PRF 201225/Z/16/Z and 224668/Z/21/Z), BBSRC (BB/N009835/1) and ERC (AdG 695709). The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript.

Author contributions

C.B. and M.H. conceived the project. C.B. and Z.Z. designed the experiments. Z.Z. developed software and performed all-optical experiments. C.B., Z.Z., M.P., J.S., E.B. and S.M. performed experiments. C.B., S.M. and Z.Z. analyzed whisker data. C.B. and Z.Z. analyzed the data. C.B., Z.Z. and M.H. wrote the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Extended data is available for this paper at https://doi.org/10.1038/s41593-022-01151-0.

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41593-022-01151-0.

Correspondence and requests for materials should be addressed to Christina Buetfering or Michael Häusser.

Peer review information *Nature Neuroscience* thanks Michael Brecht, Samuel Hires and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

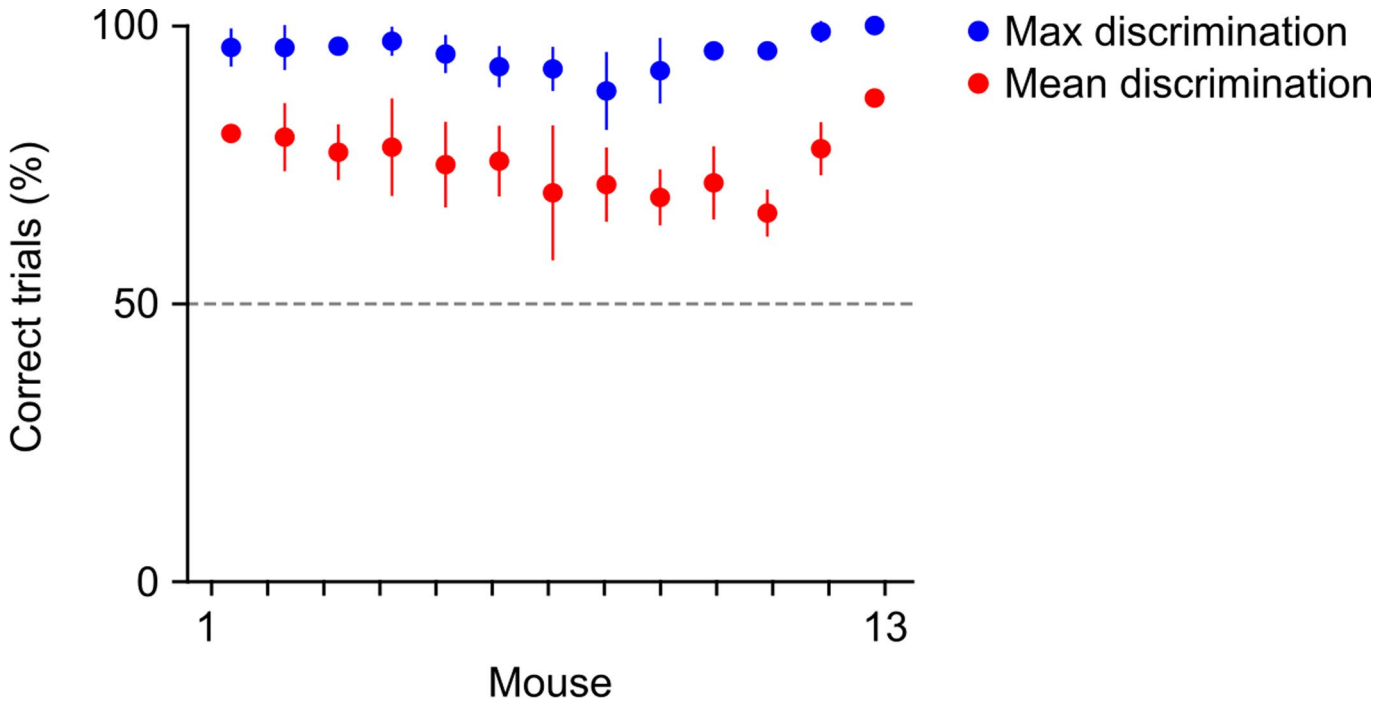

Extended Data Fig. 1 | Behavioral performance of expert mice in the two-choice two-texture discrimination task. Behavioral performance in all sessions of all mice (n = 13 mice) included in the analysis (excluding experiments using targeted photostimulation). Mean discrimination was calculated as percentage of Correct trials over all Correct and Incorrect trials. Maximum discrimination was calculated in a sliding window of 50 trials containing at least 12 trials with responses (Correct or Incorrect). Data are presented as mean values ± s.e.m.

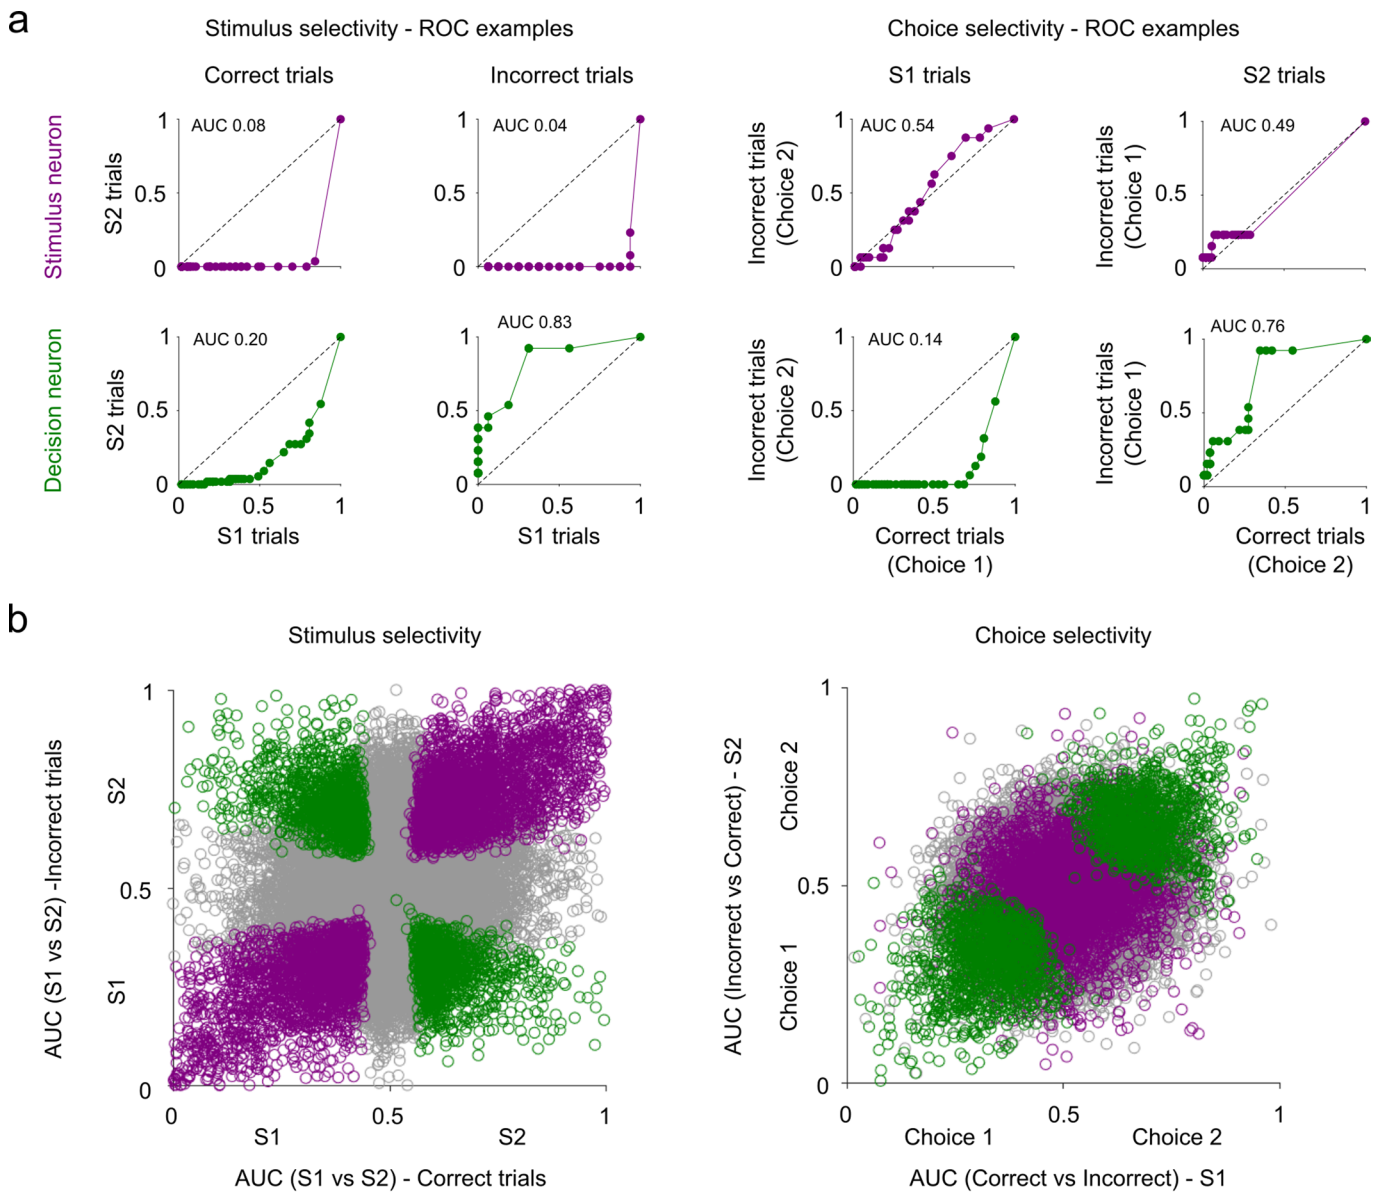

Extended Data Fig. 2 | See next page for caption.

Extended Data Fig. 2 | Stimulus and choice selectivity in L2/3 S1 neurons. (**a**) ROC curves in an example stimulus and an example decision neuron. The stimulus neuron prefers Stimulus 1 trials in Correct and Incorrect trials (top left) but does not differentiate between Correct and Incorrect trials with the same stimulus type (top right). The decision neuron has a strong preference for Stimulus 1 in Correct trials but prefers Stimulus 2 in Incorrect trials, that is follows the choice of the mouse and not the stimulus (bottom left). This also shows when comparing activity between Correct and Incorrect trials of the same Stimulus type (bottom right). (**b**) ROC AUC values for all neurons when comparing activity between Stim 1 and Stim 2 trials (left) and Correct and Incorrect trials (right). (**c**) The percentage of stimulus, decision and other neurons that show significant choice selectivity in at least one stimulus condition (left) or both stimulus conditions (right). (**d**) Number and percentage of stimulus and decision coding neurons per FOV (693 ×693 μm). N = 14 FOVs (13 mice), mean ± s.e.m.. Average trial activity 1 s before the lick (left, Fig. 2c), average trial activity from 1 s before the lick to 0.1 s before the lick (right). Gray open circles denote single FOVs.

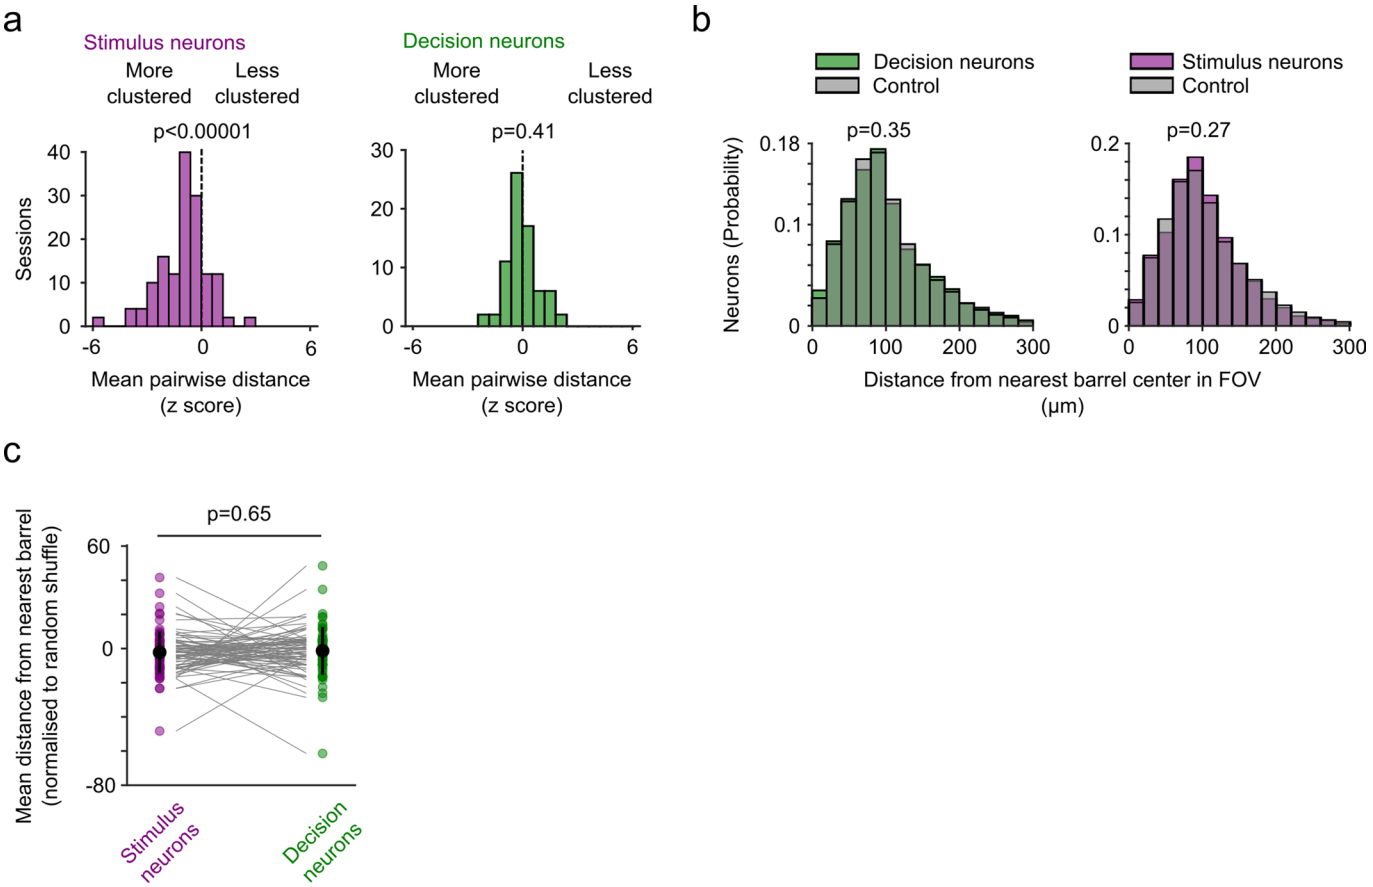

Extended Data Fig. 3 | Spatial clustering of stimulus and decision neurons. (**a**) Distribution of z-scores for mean pairwise distances. Z-score < 0 means that neurons are more clustered, z-score > 0 means that neurons are less clustered than a randomly selected group of neurons in the FOV. N = 73 FOVs, Test against 0, two-sided Wilcoxon rank sum test. (**b**) Distribution of absolute distances to the nearest barrel center for all decision neurons (left) and stimulus neurons (right). Control distribution has been calculated from the same number of randomly selected neurons across the FOV. Two-sided Wilcoxon rank sum test. (**c**) Mean distance from the nearest barrel (normalized to a shuffled distribution of randomly selected neurons across the FOV) for each session. Data are presented as mean values ± s.d., two-sided Wilcoxon signed-rank test.

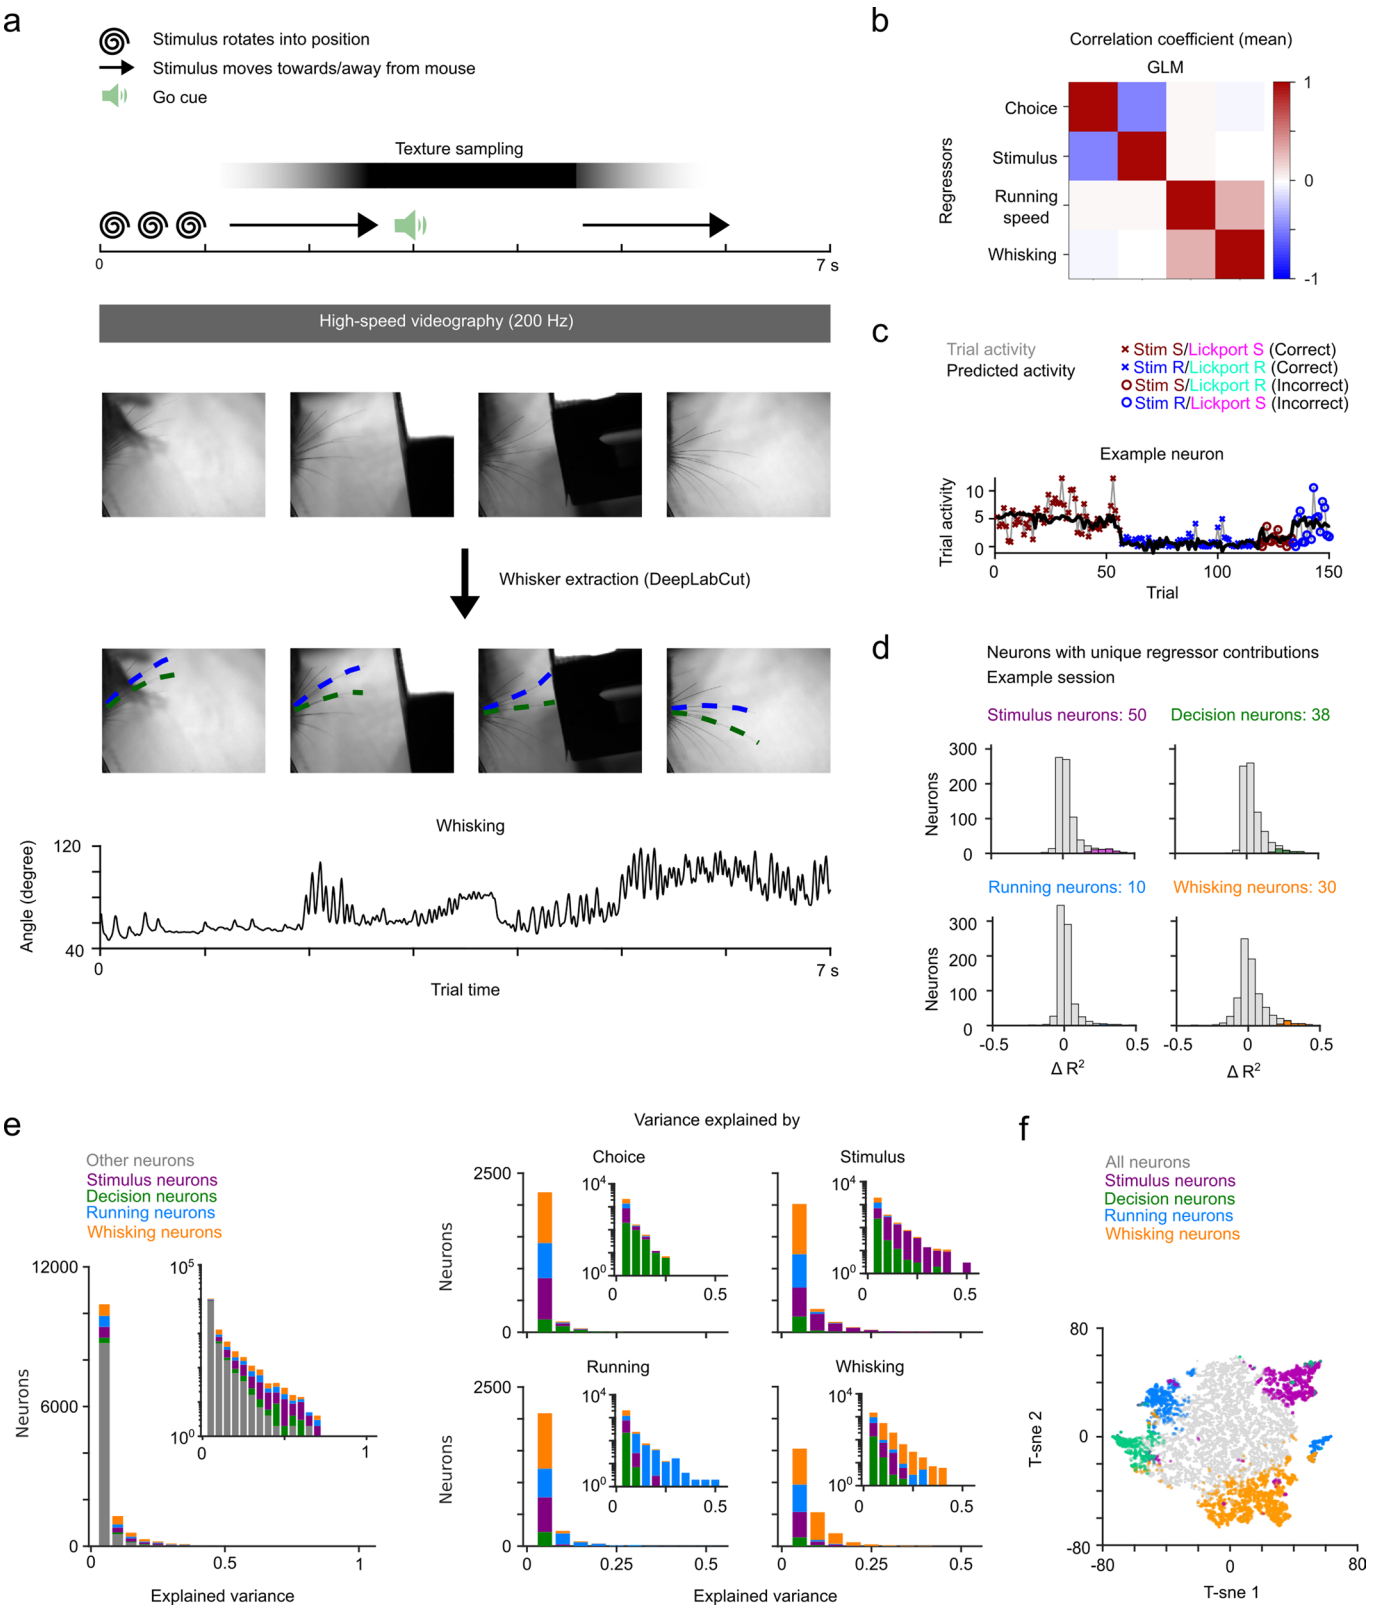

Extended Data Fig. 4 | See next page for caption.

Extended Data Fig. 4 | Gaussian GLM to identify neurons whose activity is predicted by choice. (**a**) Extraction of whisking kinematics from high-speed videography using DeepLabCut in an example trial. (**b**) Covariance matrix for GLM regressors. (**c**) Gaussian GLM predicting trial activity in an example neuron. (**d**) All neurons of an example session and their delta R2, that is the difference between R2 of the full model and the R2 of the mean of the distribution of random regressor models. (**e**) Left, explained variance by full model for all neurons. Inlay is the same data but plotted on a log scale to see that functionally identified neurons have more variance explained by the GLM analysis than other neurons. Right, Variance explained by each regressor individually for neurons identified as stimulus, choice, running or whisking neurons by the GLM analysis. Inlay is the same data but plotted on a log scale to see that the neurons with more variance explained by the respective regressor are identified as choice, stimulus, running or whisking neurons respectively. (**f**) Two-dimensional embedding of the explained variance for each regressor for all neurons. Significant contribution of a regressor to the prediction of neural activity is colored (Purple = Stimulus; Green = Choice; Blue = Running; Orange = Whisking). Neurons with multiple significant regressors show mixed colors.

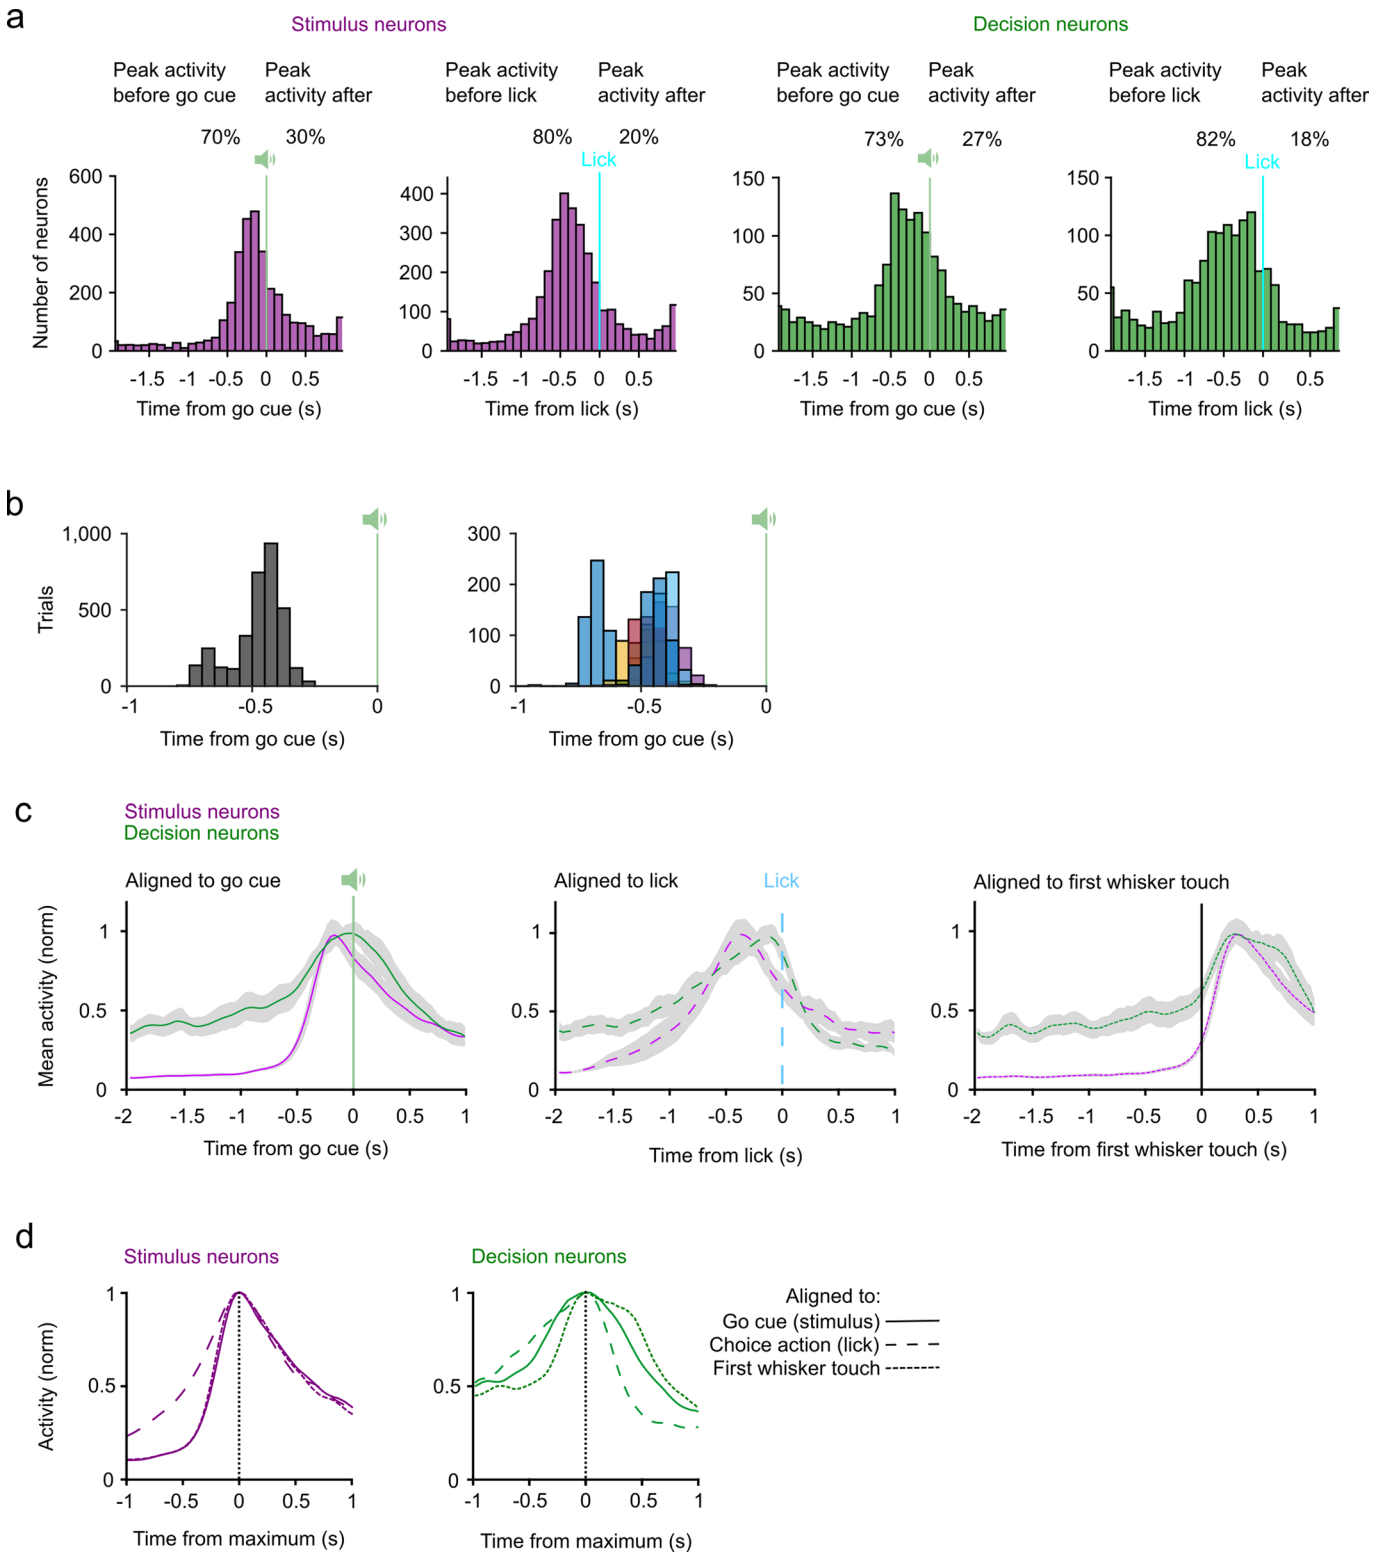

Extended Data Fig. 5 | Timing of stimulus and decision neuron activity. (**a**) Peak activity of single neurons with respect to the go cue or the first lick. The majority of neurons exhibits peak activity before the go cue and lick in the stimulus and decision neuron subgroup. (**b**) The time of first touch in all trials in a subset of sessions. Left, All first touch times across sessions. Right, First touch times color-coded by session. N = 8 sessions. (**c**) Activity of stimulus and decision neurons averaged across sessions and aligned to the stimulus onset (left), the lick of the mouse (center) or the first whisker touch (right). Line and shaded area represent mean ± s.e.m.. (**d**) The same activity traces but sorted by neuron type and aligned to the peak.

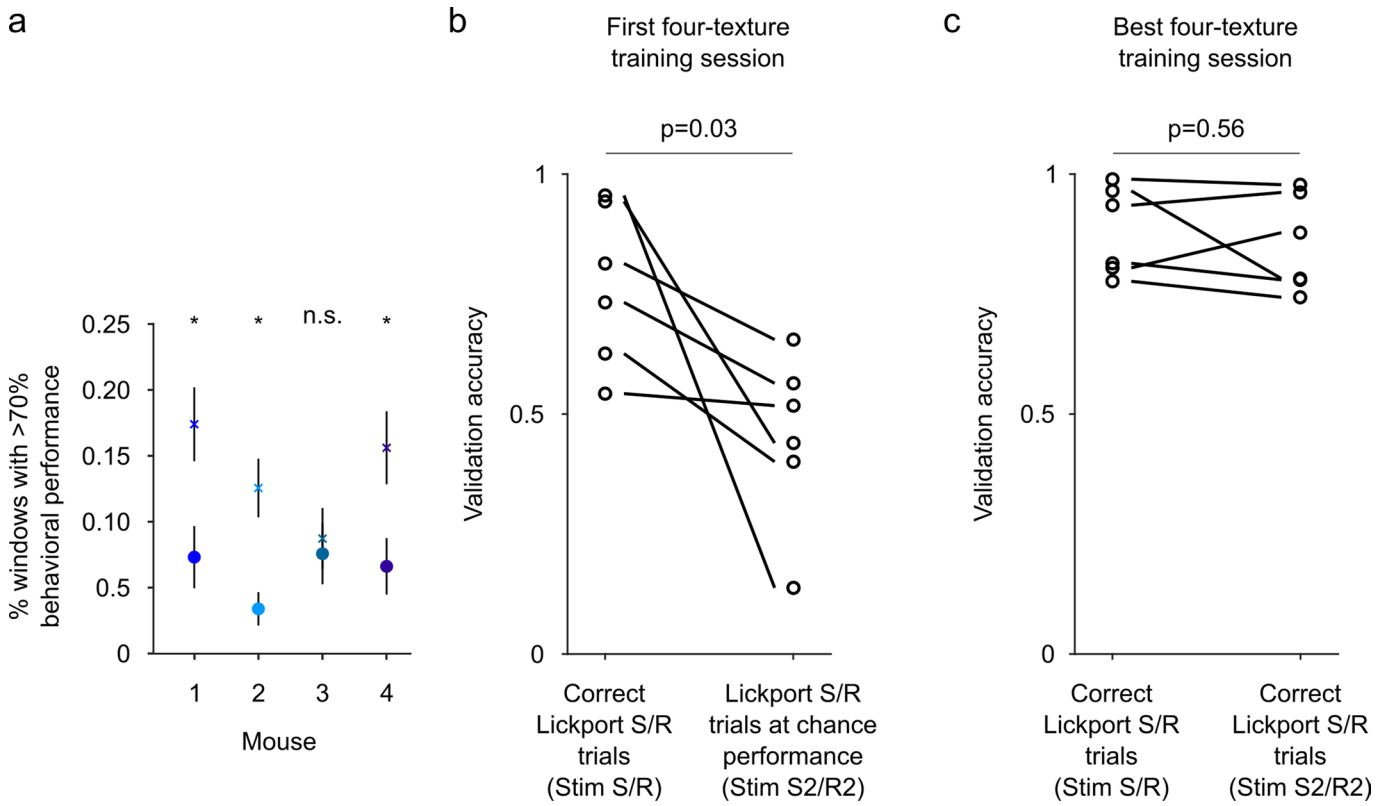

Extended Data Fig. 6 | Learning and decision coding in trials with uninformed licks. (**a**) Increase in performance between the first day of training of the four texture task and the expert session. Each session was divided into windows of trials containing 5 trials (step size 1) with the new textures and we calculated the percentage of windows that contain a percentage of Correct Trials that is above 70%. Round marker: first session; cross marker: expert session. Data are presented as mean values ± s.e.m.. P-values of a two-sided Wilcoxon rank sum test are: Mouse 1: p = 0.01, mouse 2: p = 0.0005, mouse 3: p = 0.72, mouse 4: p = 0.01. N = number of windows: Mouse 1: 123/184, mouse 2: 206/223, mouse 3: 132/149, mouse 4: 136/173. Colors indicate different mice similar to the color code in Fig. 4c. (**b**) Validation accuracy of a linear classifier trained using decision neuron activity in correct Lickport S/R trials with Stim S/R presented or random Lickport S/R trials with Stim S2/R2 presented on the first day of four-texture training (n = 6 mice). Random S2/R2 trials occurred during stretches of behavior (50 trials) with chance level performance in S2 and/or R2 trials. Two-sided Wilcoxon signed-rank test. (**c**) The same as b) but using decision neuron activity in correct trials with either Stim S/R or Stim S2/R2 presented in the best four-texture session for each mouse (n = 6 mice). Two-sided Wilcoxon signed-rank test.

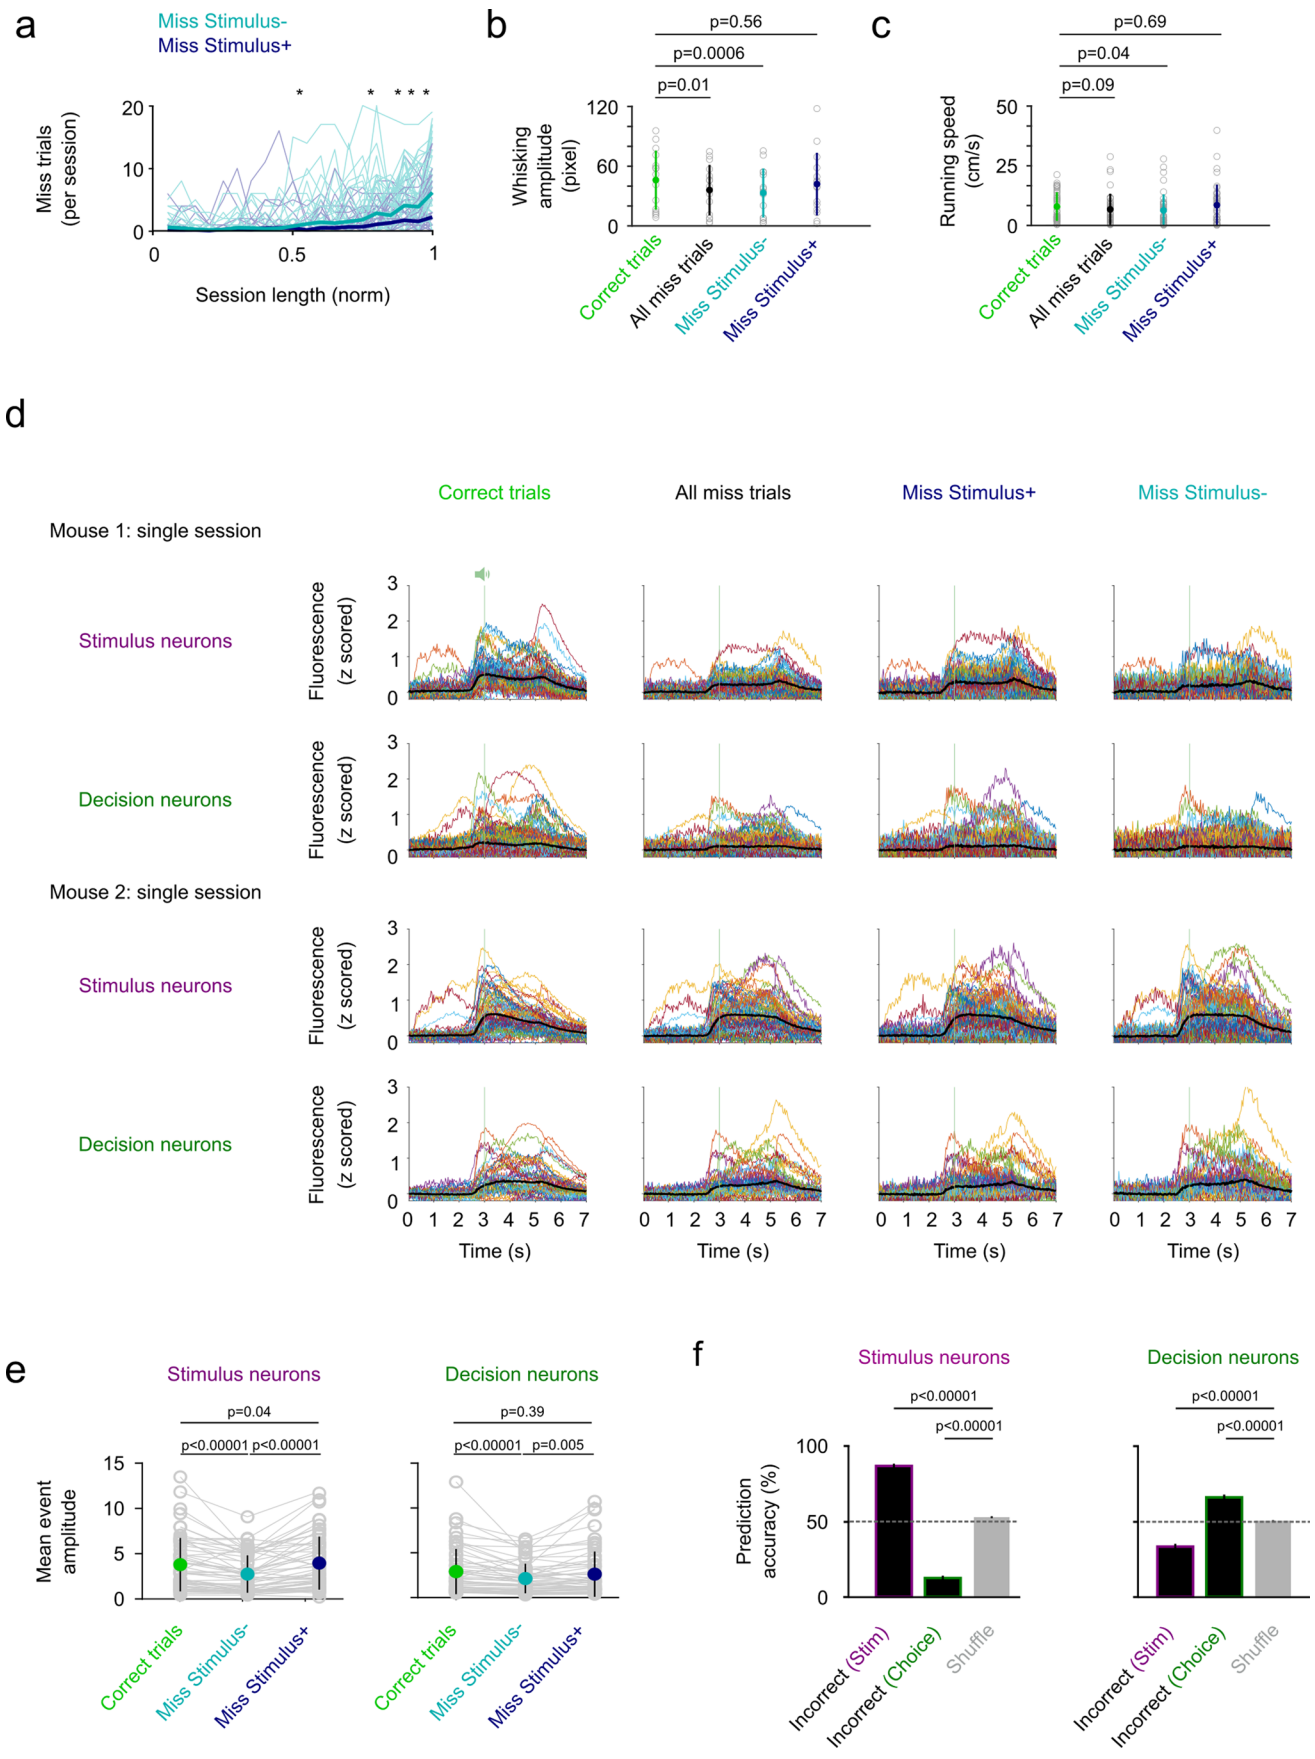

Extended Data Fig. 7 | See next page for caption.

Extended Data Fig. 7 | Distribution of Miss trials in the session and whisking, running and fluorescence during Miss trials. (**a**) Distribution of Miss Stimulus- and Miss Stimulus+ trials in each session. Session length has been normalized. N = 66 sessions, stars indicate bins with $p < 0.0025$ (Bonferroni corrected significance level), two-sided Wilcoxon signed-rank test. (**b**) Whisking amplitude before the go cue in Correct, Miss, Miss Stimulus- and Miss Stimulus+ trials in a subset of sessions with whisker kinematics. Data are presented as mean values ± s.d., n = 15 sessions, two-sided Wilcoxon signed-rank test. (**c**) Running speed before the go cue in Correct, Miss, Miss Stimulus- and Miss Stimulus+ trials in a subset of sessions with whisker kinematics. Data are presented as mean values ± s.d., n = 15 sessions, two-sided Wilcoxon signed-rank test. (**d**) Fluorescence traces (z-scored) in two example sessions split by neuron type and trial type. Thick black line indicates the mean. (**e**) Mean trial activity in stimulus and decision neurons. Data are presented as mean ± s.d., n = 57 sessions, two-sided Wilcoxon signed-rank test. (**f**) Prediction accuracy of a classifier trained on correct trials to predict stimulus type or choice with stimulus neuron activity (left) or decision neuron activity (right) in incorrect trials. Mean ± s.e.m., n = 63 sessions, two-sided Wilcoxon signed-rank test.

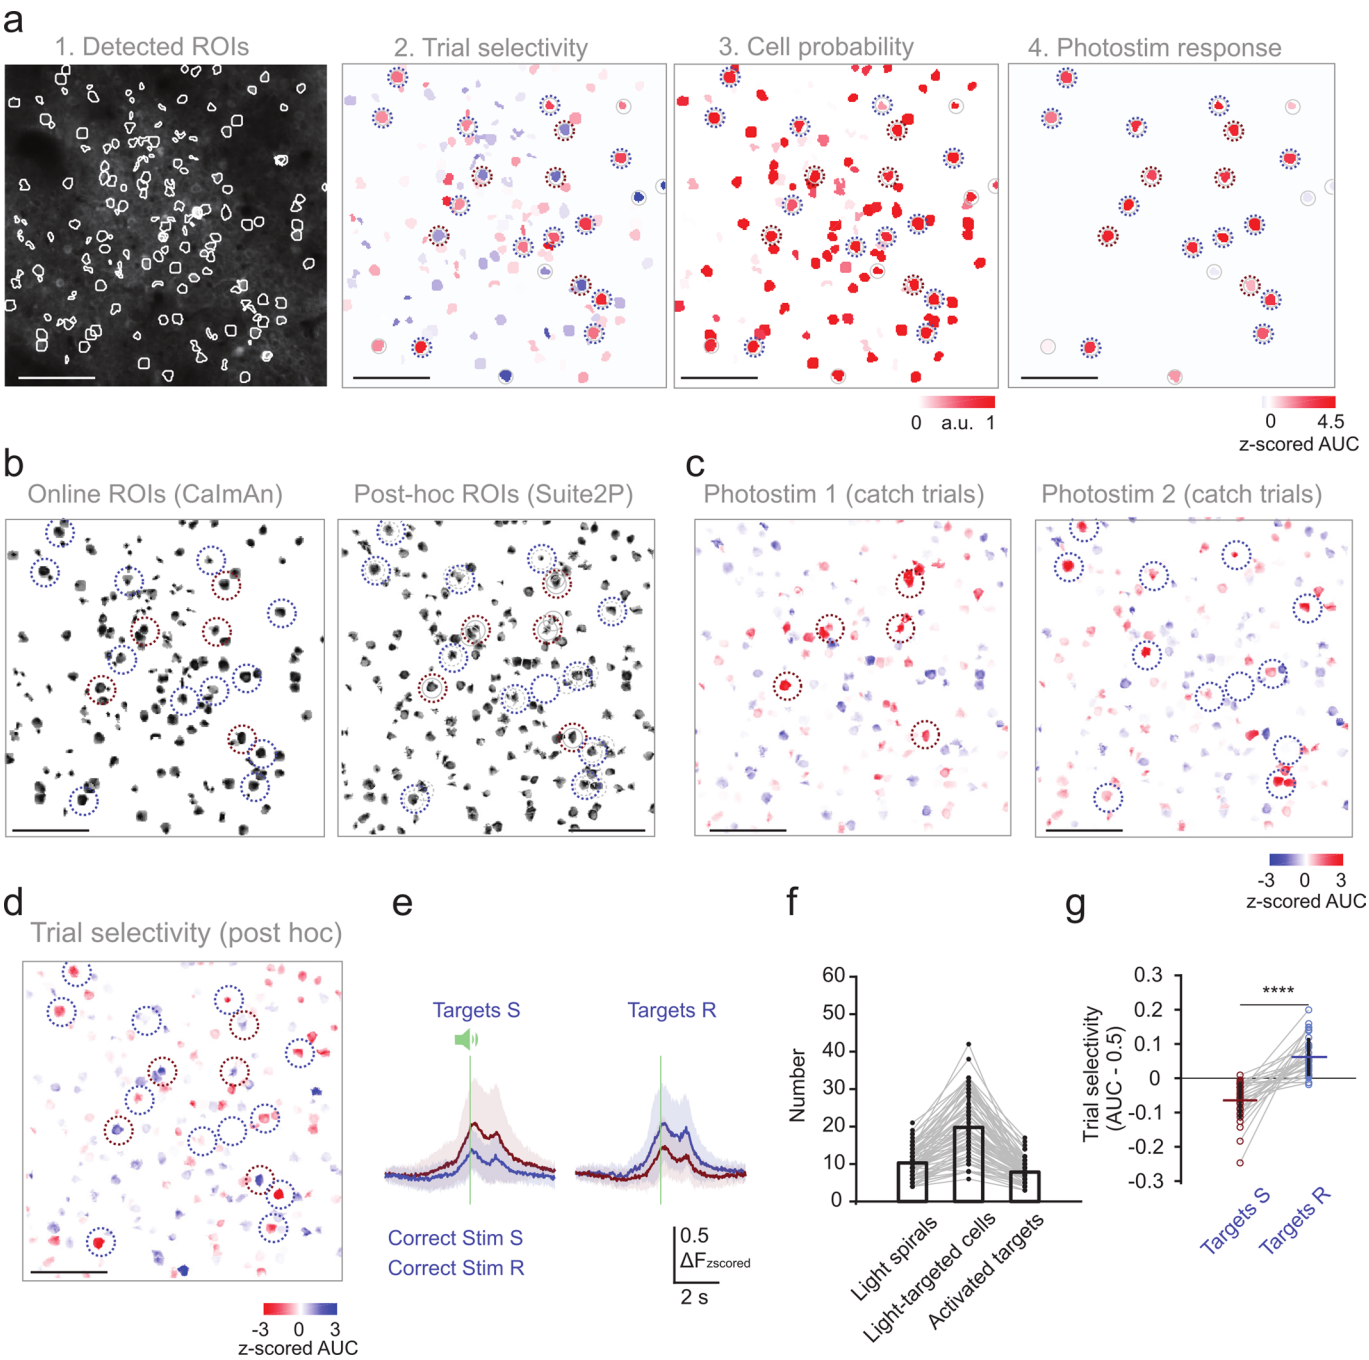

Extended Data Fig. 8 | Online target selection for photostimulation. (**a**) Online target selection procedure. Left to right: (1) ROIs were detected from live imaging stream (15 to 25 minutes). (2) Trial selectivity of each ROI was estimated from a baseline imaging session (during behavior, 150 trials). (3) Quality of each ROI was assessed with a convolutional neural network (CNN) based method. ROIs with low probability of being a neuron were excluded from target selection (threshold 0.15–0.3). (4) Photostimulation response of the trial-coding neurons that were screened for photo-excitability. Dashed and solid circles mark the target ensembles selected by this procedure. Representative of 7 mice. (**b**) Comparison of ROIs detected online and post-hoc. Thick dashed and solid circles mark the loci of light spirals used during the experiment for activating the target ensembles. Thin circles mark cells within the photostimulation spatial resolution. Pixel intensity indicates the weight of each pixel in generating the signal extracted from each ROI (same for c and d). (**c**) Photostimulation response in the field-of-view when the two target ensembles were stimulated during the all-optical behavioral session (comparing photostim and no-photostim catch trials). (**d**) Trial selectivity of ROIs detected post-hoc. Scale bars in (**a**–**d**), 100 μm. (**e**) Calcium time course of the directly activated cells in correct Stim S trials and correct Stim R trials. Thick line is mean; shaded area is s.d.. N = 80 target ensembles, 40 sessions, 7 mice. (**f**) Number of light spirals (10.3 ± 3.8), cells under spirals (light-targeted, 19.8 ± 8.2) and directly activated targets (7.9 ± 3.1) per photostimulation condition in catch trials (mean ± s.d.). N = 80 target ensembles, 40 sessions, 7 mice. (**g**) Trial selectivity of directly activated target ensembles (averaged across all activated targets in each ensemble. Data are presented as mean ± s.d., ****$p < 0.0001$, two-sided Wilcoxon signed-rank test). N = 27 out of the 40 sessions where both target ensembles meet the activation criteria (see Methods).

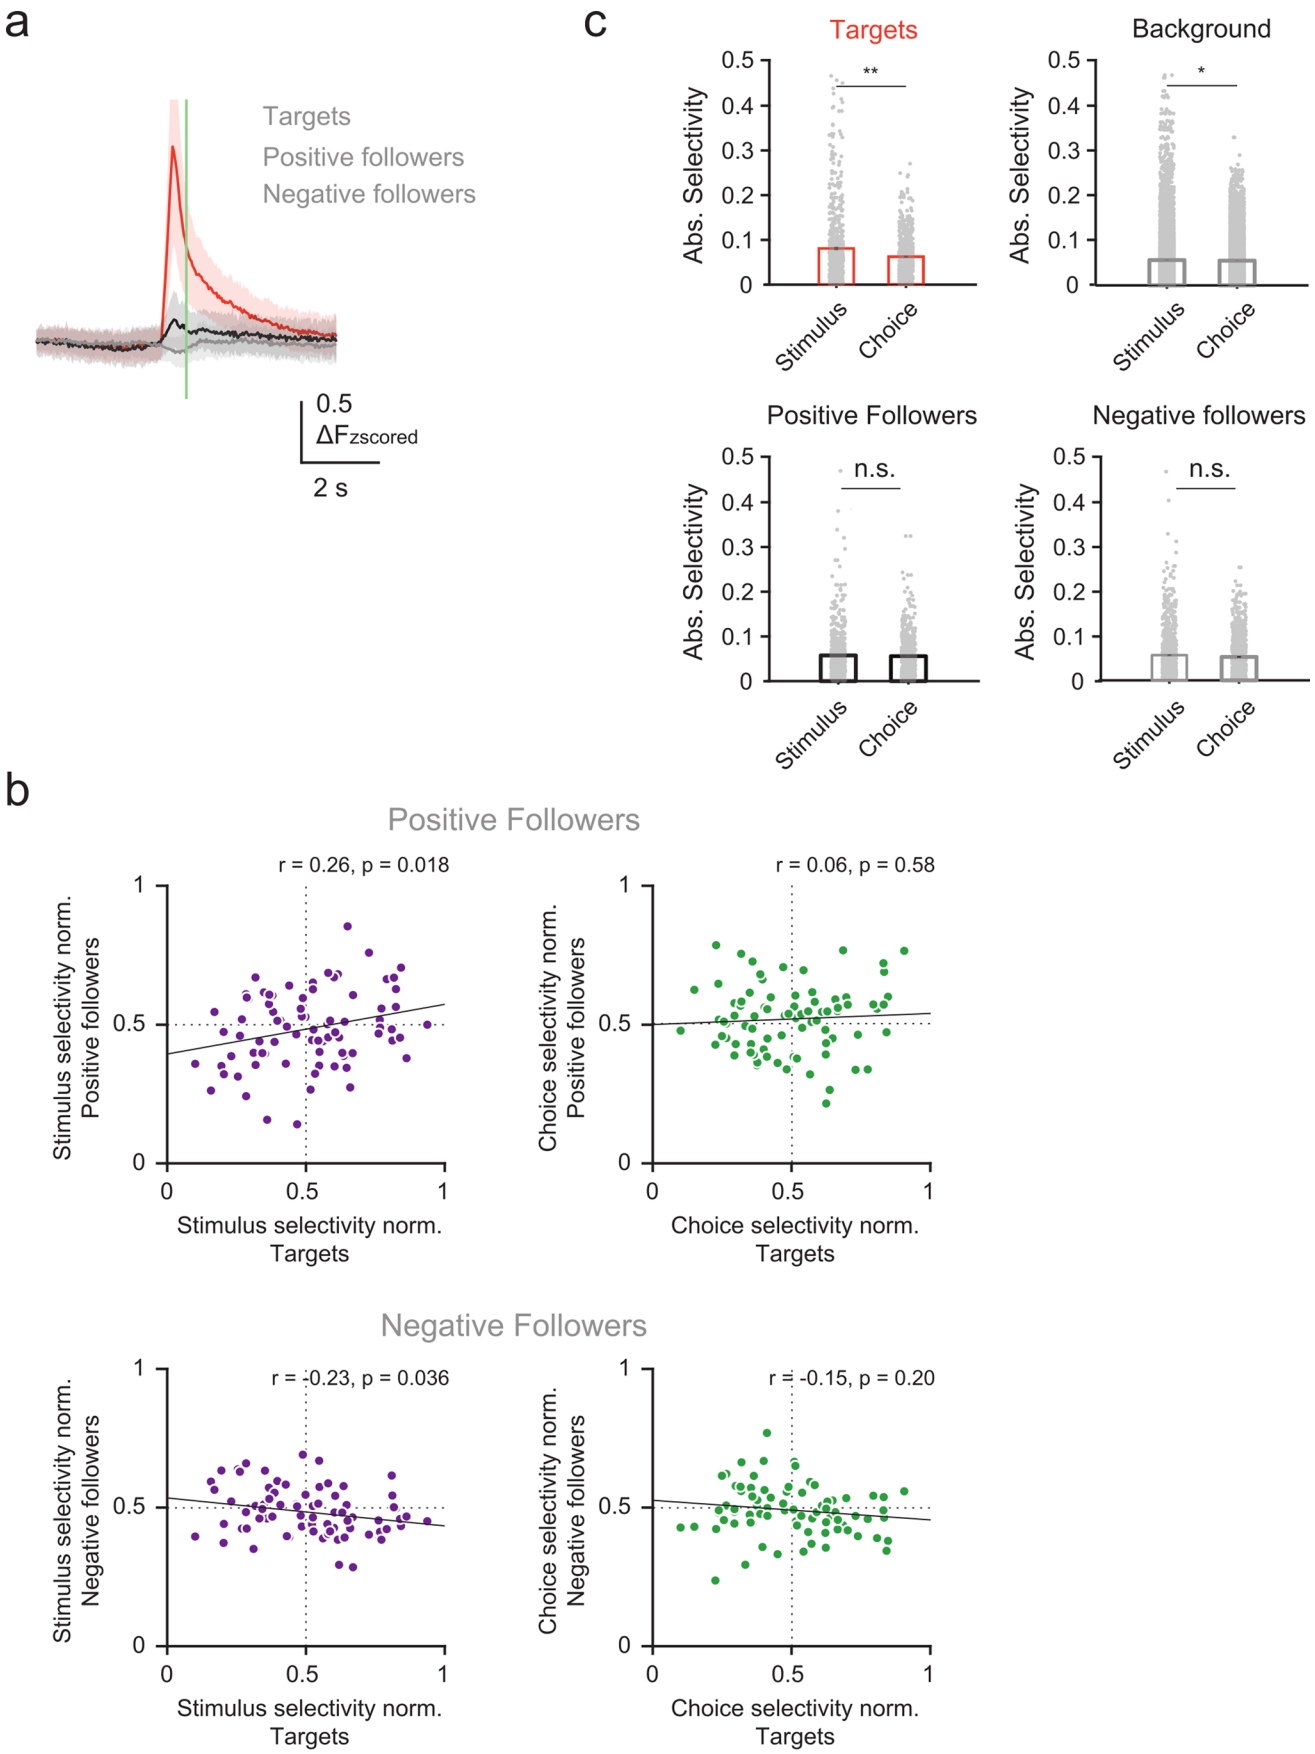

Extended Data Fig. 9 | See next page for caption.

Extended Data Fig. 9 | Stimulus and choice selectivity of targets and followers in targeted photostimulation experiments. (**a**) Calcium timecourses of the directly activated target, positive follower and negative follower ensembles aligned to go cue (green vertical line). Thick line is mean; shaded area is s.d.. N = 80 target ensembles, 40 sessions, 7 mice. (**b**) The relative stimulus selectivity of the positive or negative followers compared with other background cells (non-followers) is positively or negatively correlated with the stimulus selectivity of the target ensembles, respectively (same dataset as in **a**). There is no significant correlation between the choice selectivity of the followers and that of the targets. R and p are the Pearson correlation coefficient and the p-value, respectively. (**c**) The absolute value of stimulus selectivity is higher than that of choice selectivity in target and background cells but is comparable to choice selectivity in the followers (targets, **p = 0.0046, n = 569 neurons; background cells, *p = 0.018, 10458 neurons; positive followers, n.s., p = 0.97, n = 558 neurons; negative followers, n.s. p = 0.29). 962 neurons, 40 sessions, 7 mice. Bars are mean, two-sided Wilcoxon signed-rank test.

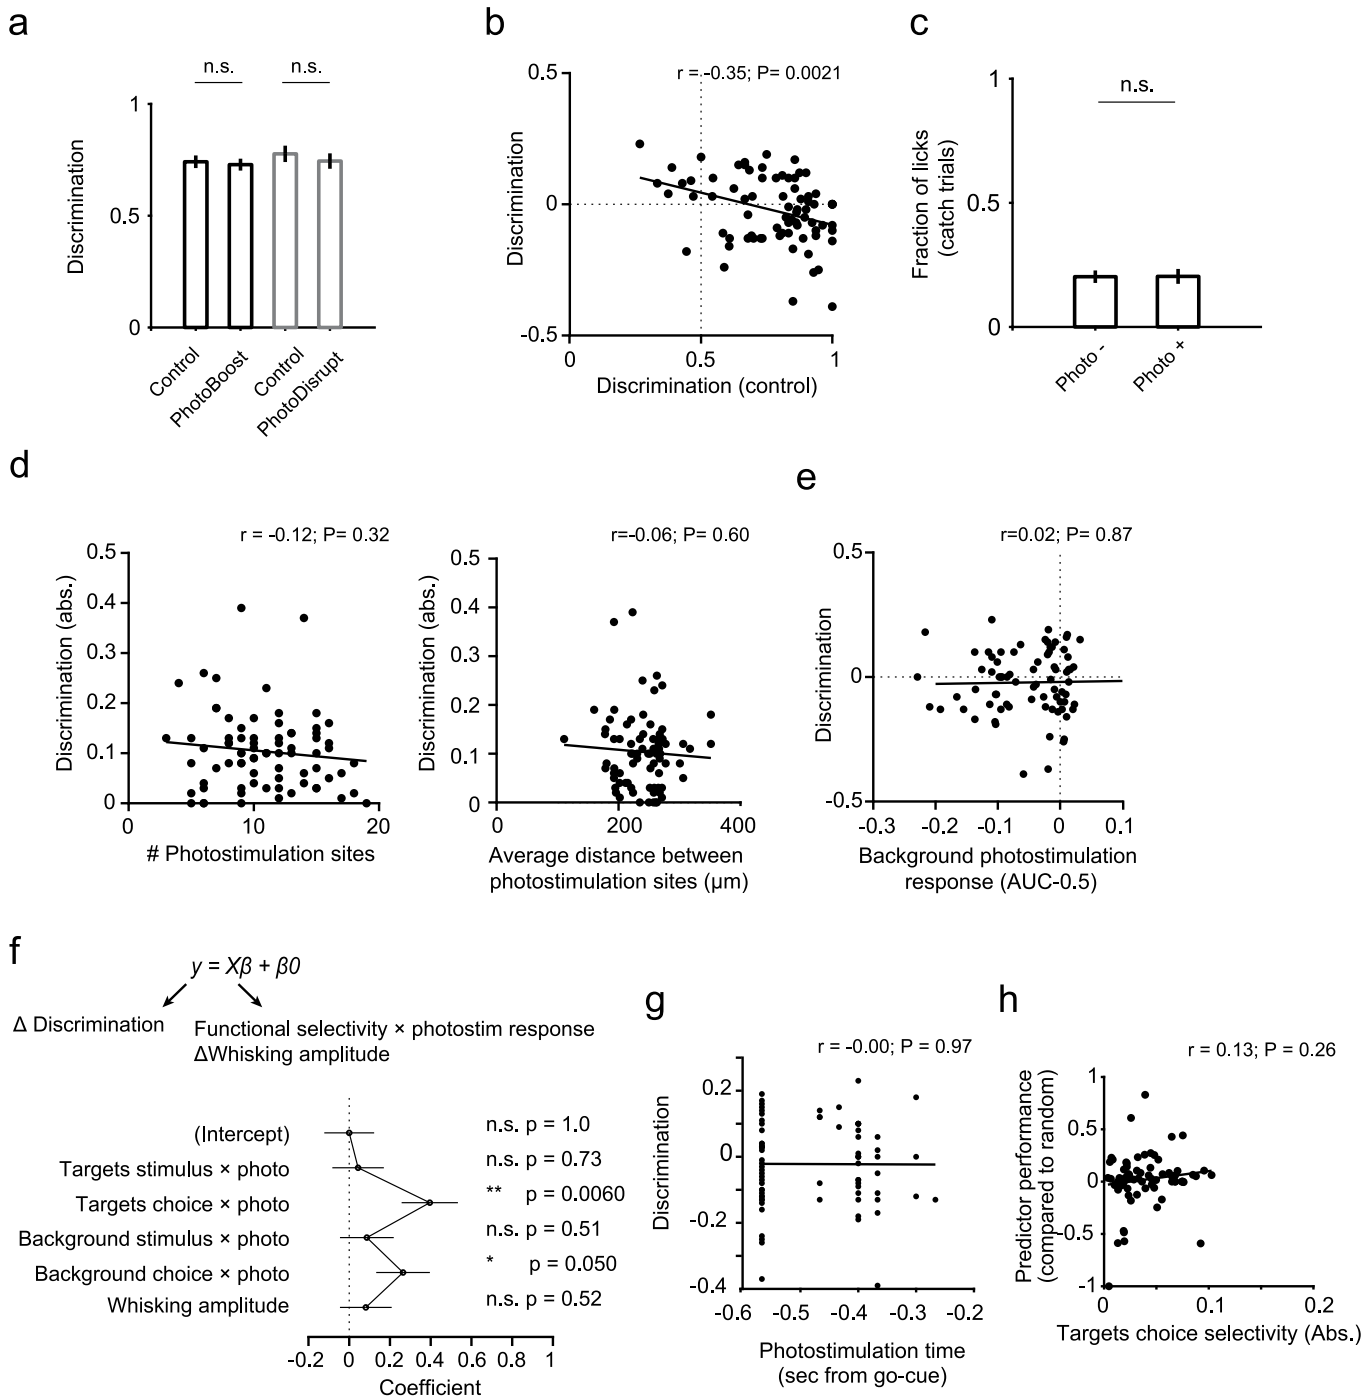

Extended Data Fig. 10 | See next page for caption.

Extended Data Fig. 10 | Analysis of the behavioral effect of targeted photostimulation. (**a**) Task performance measured as the percentage of correct choices over the total number of correct and incorrect choices the animal made in different types of trials (0.74 ± 0.028 and 0.73 ± 0.027 in PhotoBoost control trials and PhotoBoost trials, n.s., $p = 0.59$, $n = 38$ condition pairs; 0.78 ± 0.037, 0.75 ± 0.034 in PhotoDisrupt control trials and PhotoDisrupt trials, n.s., $p = 0.73$, $n = 28$ condition pairs; 32 sessions, 7 mice, two-sided Wilcoxon signed-rank test. 5 sessions where target ensembles were swapped in PhotoBoost and PhotoDisrupt were not included in this plot. Boxes are mean, whiskers are s.e.m.. (**b**) The change in texture discrimination performance in photostimulation trials is negatively correlated with the discrimination performance in the control trials, but apart from the conditions with perfect texture discrimination in control trials, photostimulation could induce both positive and negative change. (**c**) Lick rate in catch trials that received photostimulation (Photo+) is similar to the catch trials without photostimulation (Photo-). $N = 44$ photostimulation ensembles, n.s. $p = 1$, two-sided Wilcoxon signed-rank test. (**d**) The effect size of photostimulation on discrimination performance does not depend on the number or the spatial clustering (measured as the average pairwise distance between light spirals) of photostimulation sites. (**e**) The change in texture discrimination performance in photostimulation trials does not correlate with the average photostimulation response of the background cells. Same experiments as in Fig. 7b. (**f**) A linear regression model to predict discrimination performance change based on the weighted average photostimulation response by the stimulus or choice selectivity of targets and background neurons, as well as whisking amplitude. $N = 60$ photostimulation conditions, 29 sessions, 7 mice in which whisker movements were recorded by high-speed videography. Data are presented as mean $\pm$ s.e.m.. The p-values are for the t-statistic of the hypothesis test that the corresponding coefficient is equal to zero. (**g**) The change in discrimination performance does not correlate with the onset of photostimulation. (**h**) The performance of the online predictor does not correlate with the choice selectivity of the targets. g-h, Same experiments as in Fig. 7b. R and p are the Pearson correlation coefficient and the p value, respectively.

nature research

Corresponding author(s): Buetfering/Häusser

Last updated by author(s): Jul 8, 2022

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see Authors & Referees and the Editorial Policy Checklist.

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a	Confirmed	
☐	☒	The exact sample size (*n*) for each experimental group/condition, given as a discrete number and unit of measurement
☐	☒	A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
☐	☒	The statistical test(s) used AND whether they are one- or two-sided *Only common tests should be described solely by name; describe more complex techniques in the Methods section.*
☐	☒	A description of all covariates tested
☐	☒	A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
☐	☒	A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
☐	☒	For null hypothesis testing, the test statistic (e.g. *F*, *t*, *r*) with confidence intervals, effect sizes, degrees of freedom and *P* value noted *Give P values as exact values whenever suitable.*
☒	☐	For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
☒	☐	For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
☐	☒	Estimates of effect sizes (e.g. Cohen's *d*, Pearson's *r*), indicating how they were calculated

Our web collection on statistics for biologists contains articles on many of the points above.

Software and code

Policy information about availability of computer code

Data collection	Behavoiural data was acquired using PyBehaviour (www.github.com/llerussell/PyBehavior). Imaging data were acquired using ThorImage 3.1. PackIO was used to synchronize imaging data and behavioural data (reference X). Arduino IDE was used to control setup-components such as stimulus presentation. pyRTAOI, a custom-based real-time all-optical interface was used for all-optical experiments. CaImAn was used for online ROI detection in all-optical experiments (reference). Highspeed videography was recorded using the software package 2ndLook (IO Industries). A custom SLM control interface was written in C++ with Qt 5.9, developed in Microsoft Visual Studio 2013.
Data analysis	Data analysis was performed in Matlab using custom written scripts and available software. Segmentation and event detection of calcium imaging data was performed using Suite2p (www.github.com/cortex-lab/Suite2P, reference). DeepLabCut 1.0 (www.github.com/DeepLabCut/DeepLabCut, reference) was used to analyze whiskers in highspeed videography.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors/reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research guidelines for submitting code & software for further information.

Data

Policy information about availability of data

All manuscripts must include a data availability statement. This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

The data and analysis code that support the findings of this study are available from the corresponding authors upon reasonable request

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

☒ Life sciences ☐ Behavioural & social sciences ☐ Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see nature.com/documents/nr-reporting-summary-flat.pdf

Life sciences study design

All studies must disclose on these points even when the disclosure is negative.

Sample size	No statistical methods were used to pre-determine sample sizes.
Data exclusions	No mice were exluded from the analysis unless they did not learn the task, data could not be collected due to poor transgenic GCaMP or viral expression or occlusion of the chronic window. All other reasons for exclusions of trials or sessions are mentioned in the Methods sections. For example groups of trials were excluded from the analysis if the mouse developed a lickport bias as described in the Methods section "Data analysis". For some analyses, sessions were only included if they contained at least 5 stimulus and 5 decision neurons (see e.g. "Spatial clustering" and "Shared trial-by-trial response variability" in the Methods section).
Replication	Decision neurons, as described in this manuscript, were found in all mice that we were able to record and that learnt the task. Therefore, the core finding that these neurons exist was replicated multiple times. Moreover, every other finding in the manuscript was shown in multiple mice and therefore replicated across mice.
Randomization	No randomization of mice was done as all mice were trained and recorded under the same conditions. Presentation of texture stimulus type was randomized within each training session. There was also a variable delay between trials to randomize the trial start time.
Blinding	Data collection and analysis were not performed blind to the conditions of the experiment, but analysis relied on code that was standardized for all experimental conditions.

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involved in the study
☒	☐ Antibodies
☒	☐ Eukaryotic cell lines
☒	☐ Palaeontology
☐	☒ Animals and other organisms
☒	☐ Human research participants
☒	☐ Clinical data

Methods

n/a	Involved in the study
☒	☐ ChIP-seq
☒	☐ Flow cytometry
☒	☐ MRI-based neuroimaging

Animals and other organisms

Policy information about studies involving animals; ARRIVE guidelines recommended for reporting animal research

Laboratory animals	We used wildtype mice (C57/Bl6), transgenic mice expressing GCaMP6s in excitatory neurons (EMX1-Cre; Camk2a-tTA; Ai94 (Jax #027784; #007004; #024104) and Camk2a-tTA; tetO-G6s (Jax #007004; #024742)) mice. Mice were of either gender and between 8 and 36 weeks old at the start of experiments.
Wild animals	No wild animals were used in this study.
Field-collected samples	No field-collected samples were used in this study.
Ethics oversight	All animal procedures were approved by the local Animal Welfare and Ethical Review Board at University College London and performed under license from the UK Home Office in accordance with the Animals (Scientific Procedures) Act 1986

Note that full information on the approval of the study protocol must also be provided in the manuscript.