Figure S1

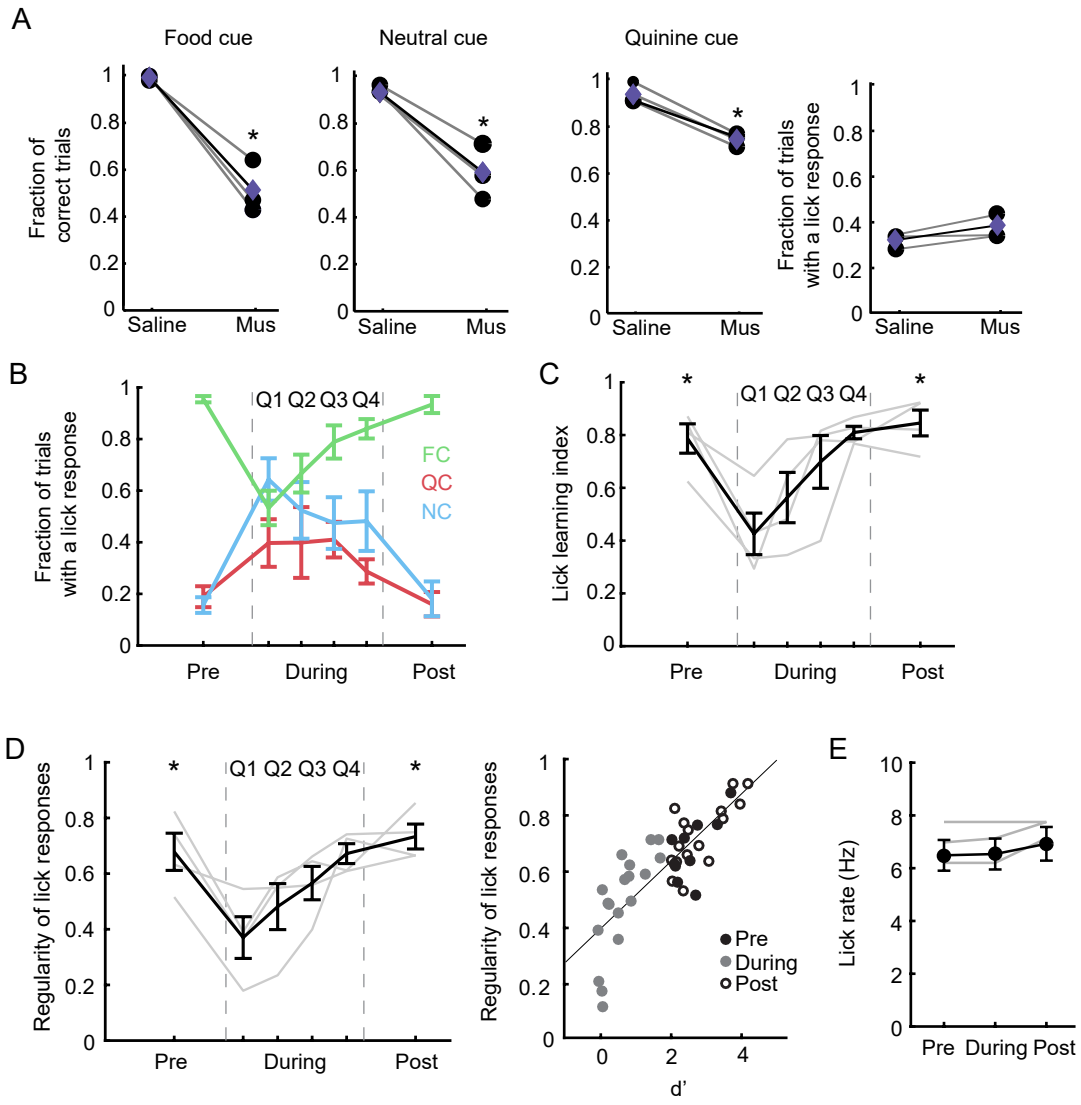
A
Food cue
Neutral cue
Quinine cue
Fraction of correct trials
Fraction of trials with a lick response
Saline
Mus
B
Fraction of trials with a lick response
Q1 Q2 Q3 Q4
FC
QC
NC
Pre
During
Post
C
Lick learning index
D
Regularity of lick responses
d'
Pre
During
Post
E
Lick rate (Hz)

Figure S1. Behavioral performance across a reassignment of cue-outcome associations, related to Figure 1.

A. Injection of the $GABA_A$ agonist muscimol ('Mus') into lateral visual association cortex decreased behavioral performance for all three cues (FC: food cue; QC: quinine cue; NC, neutral cue; * $p < 0.02$ for all comparisons, paired Student's t-test). However, the overall fraction of trials with a cue-evoked lick response did not change following injection of muscimol ($p > 0.05$, paired Student's t-test).
B. After the reassignment of cue-outcome associations, the mice initially licked roughly equally to all three cues. Over time, the licking to non-rewarded cues decreased, while the fraction of trials with licking to the rewarded cue returned to a high level. We divided up the "during-Reversal" period into quartiles ('Q1', 'Q2', 'Q3', 'Q4'; labeled above) to show the evolution of behavior with time. FC: food cue; QC: quinine cue; NC, neutral cue; Pre: pre-Reversal; During: during-Reversal; Post: post-Reversal.
C. Reassignment of cue-outcome associations also caused an increase in pre-stimulus licking (i.e. not cue-evoked). We quantified this using a lick learning index ($[Licks_{Cue} - Licks_{PreCue}]/[Licks_{Cue} + Licks_{PreCue}]$). This metric compares the number of licks in the 1-s period prior to cue presentation ($Licks_{PreCue}$) vs. the number of licks in the last second of cue presentation ($Licks_{Cue}$; pre- vs. during- vs. post-Reversal: $p < 0.01$, pre- vs. during-: $p < 0.05$, post- vs. during-: $p < 0.01$, 1-way repeated measures ANOVA, Tukey-Kramer method).
D. Left: the regularity of lick responses decreased after reassignment of cue-outcome associations but returned to pre-Reversal levels once the mice learned the new cue-outcome associations (pre- vs. during- vs. post-Reversal: $p = 0.0005$, pre- vs. during-Reversal and post- vs. during-Reversal, $p < 0.05$, 1-way repeated measures ANOVA, Tukey-Kramer method; see STAR Methods). Right: Improved behavioral performance strongly correlates with more stereotyped licking responses (Pearson correlation with d': 0.81, $p < 0.0001$).
E. The overall lick rate did not change between pre-, during-, and post-Reversal epochs ($p = 0.12$, 1-way repeated-measures ANOVA), suggesting similar levels of task engagement. Error bars denote s.e.m.

Figure S2

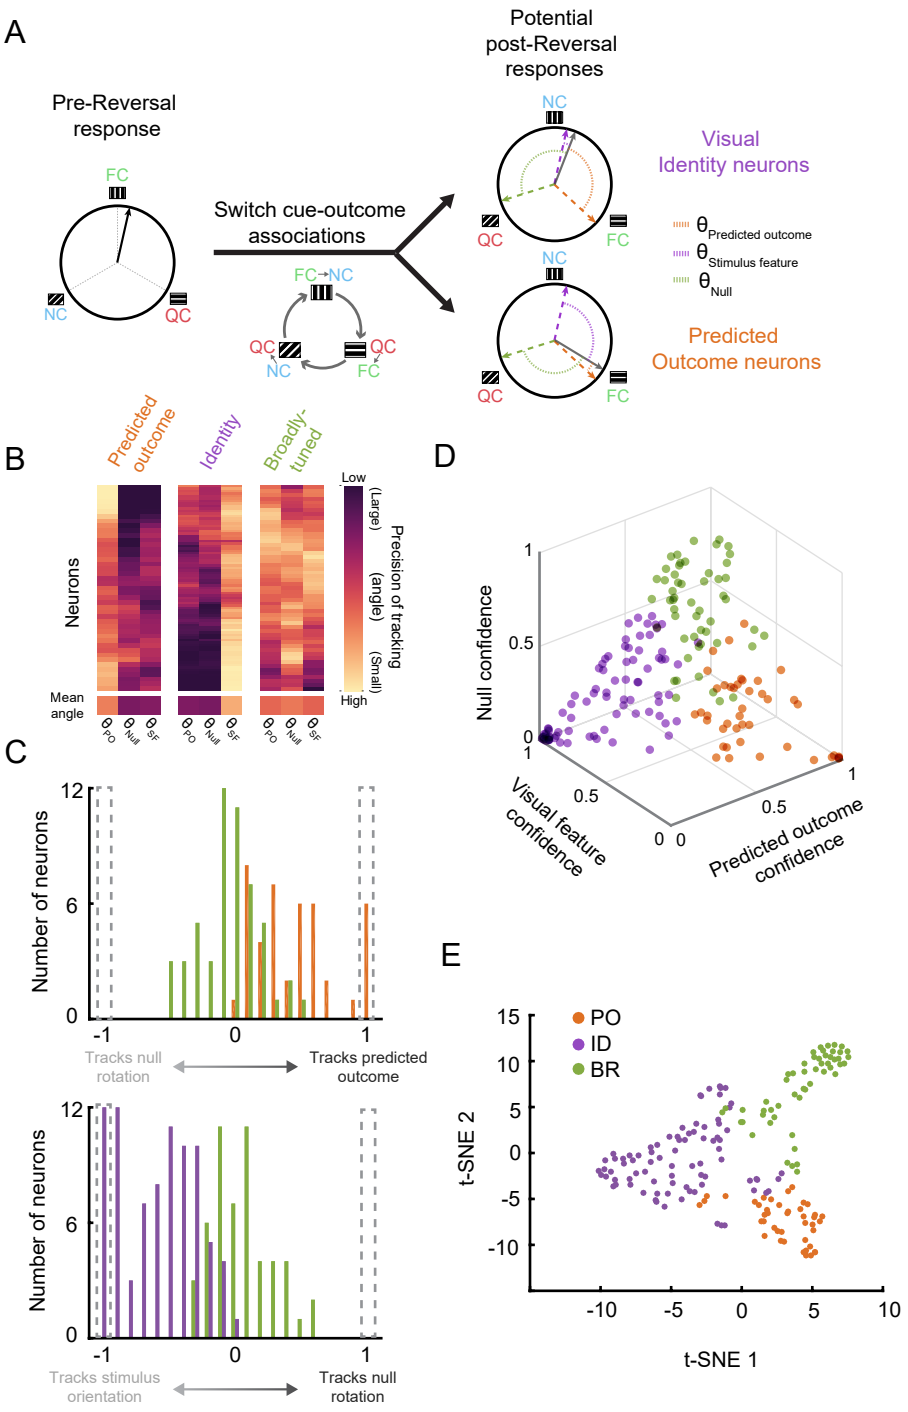
A
Pre-Reversal response
FC
NC
QC
Switch cue-outcome associations
FC→NC
QC
NC
QC
Potential post-Reversal responses
NC
QC
FC
NC
QC
FC
Visual Identity neurons
θ Predicted outcome
θ Stimulus feature
θ Null
Predicted Outcome neurons
B
Predicted outcome
Identity
Broadly-tuned
Neurons
Mean angle
θ PO
θ Null
θ SF
Low
(Large)
(angle)
(Small)
High
Precision of tracking
C
Number of neurons
12
6
0
-1
0
1
Tracks null rotation
Tracks predicted outcome
Number of neurons
Tracks stimulus orientation
Tracks null rotation
D
Null confidence
Visual feature confidence
Predicted outcome confidence
E
PO
ID
BR
t-SNE 2
t-SNE 1

Figure S2. Identity and Predicted Outcome neurons track specific stimulus orientations or predicted outcomes in lateral visual association cortex, related to Figure 2.

A. We estimated how the net response preference of a neuron pre-Reversal (black arrow) would "rotate" post-Reversal if the neuron tracked low level stimulus features (no rotation; dashed purple arrow), tracked predicted outcome (clockwise rotation of net preference; dashed orange arrow), or tracked neither stimulus identity nor predicted outcome (assessed by testing for a rotation in the counterclockwise "null" direction; dashed green arrow). We classified the neuron as a Predicted Outcome (PO), Identity (ID), or Broadly-tuned (BR) neuron based on which of these three hypothetical rotations best matched the post-Reversal net response preference (gray arrow; formally, we estimated the minimum angle among $\theta_{\text{Stimulus features}}$, $\theta_{\text{Predicted Outcome}}$, and θ_{Null}; angle values for each neuron are shown in B, below). BR neurons were defined as those neurons whose angles were similarly small for all "rotations" due to their similar responsivity to all visual stimuli. FC: food cue; QC: quinine cue; NC: neutral cue.
B. Heatmaps of angles defined in A, for PO, ID, or BR neurons ('Small' angle: 0; 'Large' angle: $\pi/2$). PO neurons with very small $\theta_{\text{Predicted Outcome}}$ tracking angles suggested selective tracking of a given predicted outcome. ID neurons with very small $\theta_{\text{Stimulus features}}$ tracking angles suggested selective tracking of stimulus identity. PO: predicted outcome; SF: stimulus features.
C. Top: incidence of neurons demonstrating different relative levels of confidence in tracking of the null rotation vs. the rotation related to the change in predicted outcomes (i.e. comparison of θ_{Null} vs. $\theta_{\text{Predicted Outcome}}$; for details, see STAR Methods). Colors indicate PO, ID, and BR neurons. X-axis values indicate the confidence of the estimate that a neuron specifically tracks a specific stimulus orientation (see STAR Methods), tracks a specific predicted outcome, or tracks the hypothetical rotation of outcomes in the opposite direction of the actual rotation (i.e. rotation in the 'null' direction). Bottom: same plot, for the comparison of relative tracking of stimulus identity vs. tracking of the null rotation. For a comparison of relative tracking of stimulus identity vs. predicted outcome, see Figure 2F. Note the substantial fraction of neurons that track either a stimulus feature or a predicted outcome with high confidence (dashed gray boxes; see STAR Methods). In contrast, *no* neurons specifically tracked the rotation in the null direction following reassignment of cue-outcome associations with high confidence.
D. Three-dimensional scatter plot of categorization confidence for each neuron that responded to at least one visual stimulus both pre- and post-Reversal. PO, ID, and BR neurons are displayed in orange, purple and green, respectively. The confidence metric was calculated as 1- $\theta_{\text{normalized SF, PO, Null}}$ (see STAR Methods). Note the clustering of purple and orange neurons (highly overlapping dots) at the visual feature confidence vertex and at the predicted outcome confidence vertex, respectively.
E. Dimensionality reduction was used to visualize the differentiability of the three categories of PO, ID, and BR neurons.

Figure S3

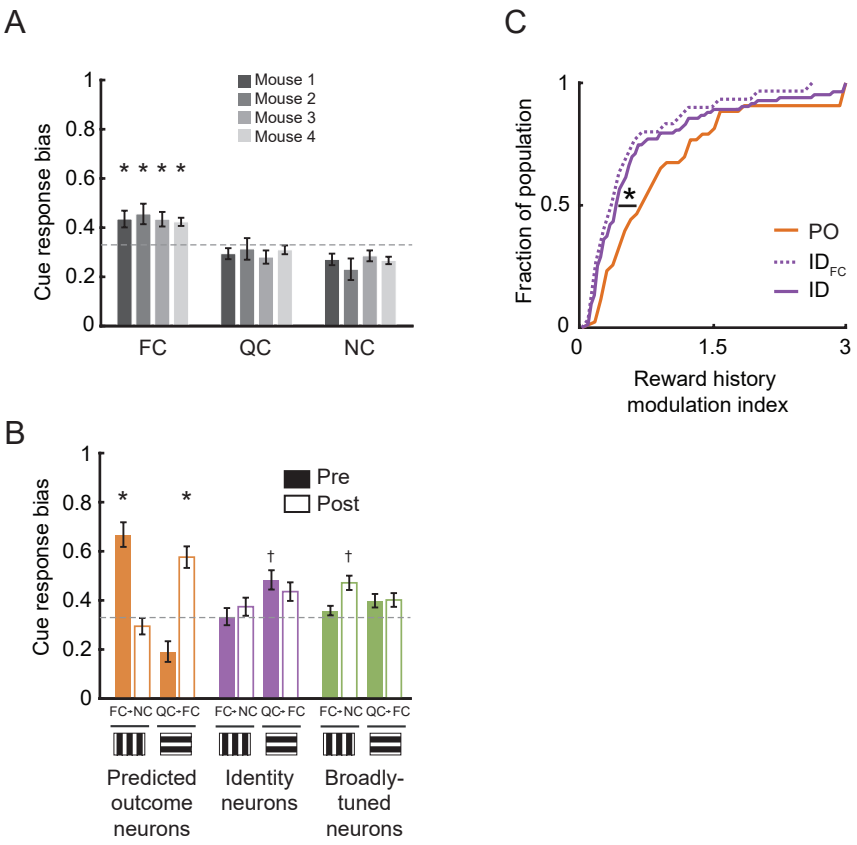
A
Mouse 1
Mouse 2
Mouse 3
Mouse 4
Cue response bias
FC
QC
NC
C
Fraction of population
PO
ID_FC
ID
Reward history modulation index
B
Pre
Post
Cue response bias
FC-NC
QC-FC
Predicted outcome neurons
Identity neurons
Broadly-tuned neurons

Figure S3. Additional analyses of response biases and modulation by reward history, related to Figure 3.

A. In each mouse recorded across the entire reversal paradigm, we observed a significant population bias towards the food cue. Of all neurons tracked pre- and post-Reversal (n = 179, which includes all PO, ID and BR neurons analyzed in Figures 2-6), 101 responded most strongly to the food cue, 42 to the quinine cue, and 36 to the neutral cue. On average, the percentages were: FC preferring: 56% $\pm$ 1%, QC preferring: 21% $\pm$ 1%, NC preferring: 23% $\pm$ 1%, mean $\pm$ SEM across 4 animals). Individual percentages per mouse were: FC preferring: 54%, 54%, 57%, and 60%, QC preferring: 22%, 23%, 19%, and 18%, and NC preferring: 24%, 23%, 23%, and 22%. * $p < 0.01$ one-tailed Wilcoxon Sign-Rank vs. chance bias (0.33), Bonferroni corrected. FC: food cue; QC: quinine cue; NC: neutral cue.
B. PO neurons, but not ID or BR neurons, showed a strong response bias to the food cue both pre- and post-Reversal. * $p < 0.0001$, † $p < 0.05$, Wilcoxon Sign-Rank vs. chance bias (0.33), Bonferroni corrected. Pre: pre-Reversal; Post: post-Reversal. FC: food cue; QC: quinine cue; NC: neutral cue.
C. Cumulative distribution of reward history modulation index values for the same data as in Figure 3F, but restricting analysis to food cue-preferring ID neurons (ID_{FC}). Data from PO and ID neurons are plotted for comparison (same data as in Figure 3F, *right*). Note that, as with the entire population of ID neurons, the subset that preferred the FC also showed significantly less sensitive to reward history than PO neurons (the vast majority of PO neurons also preferred the food cue). * $p < 0.025$, Kruskal-Wallis, Bonferroni corrected. Error bars denote s.e.m. Thus, even neurons that respond to the *same* food cue exhibit different sensitivity to reward history depending on category membership.

Figure S4

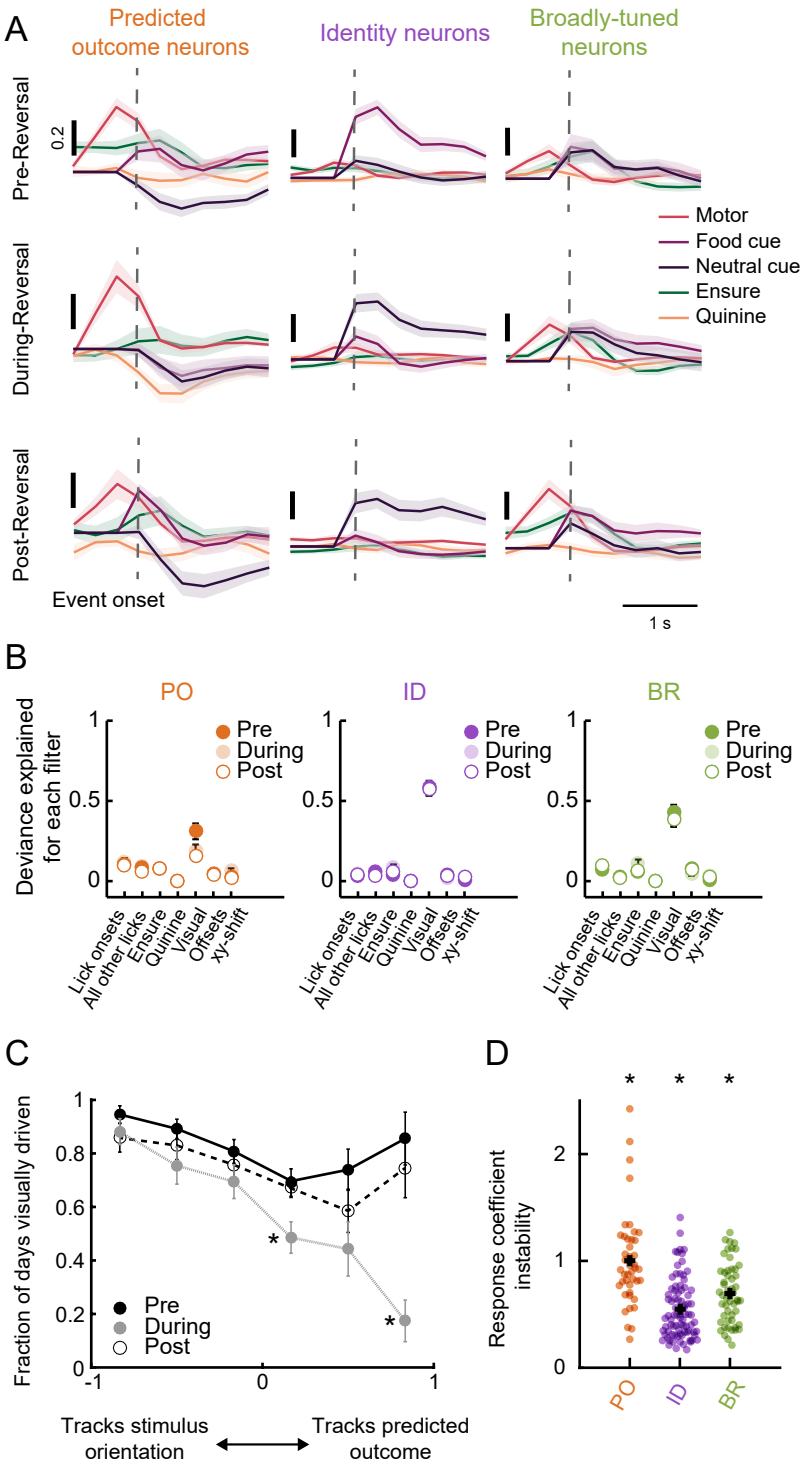
A
Predicted outcome neurons
Identity neurons
Broadly-tuned neurons
Pre-Reversal
During-Reversal
Post-Reversal
0.2
Motor
Food cue
Neutral cue
Ensure
Quinine
Event onset
1 s
B
PO
ID
BR
Deviance explained for each filter
Pre
During
Post
1
0.5
0
Lick onsets
All other licks
Ensure
Quinine
Visual
Offsets
xy-shift
C
Fraction of days visually driven
1
0.8
0.6
0.4
0.2
0
-1
0
1
Pre
During
Post
Tracks stimulus orientation
Tracks predicted outcome
D
Response coefficient instability
2
1
0
PO
ID
BR
*

Figure S4. Using a generalized linear model (GLM) to quantify behavioral modulation of neural activity, related to Figure 4.

A. The fitted beta coefficients pre-, during-, and post-Reversal for Predicted Outcome (PO), Identity (ID), and Broadly-tuned (BR) neurons that prefer the food cue pre-Reversal. The dashed gray line represents the onset for each component, including onset of motor response (i.e. time of first lick following stimulus onset), onset of food cue or neutral cue, or onset of Ensure or quinine delivery. We sub-selected food cue preferring ID neurons to illustrate the switch in cue preference pre- vs. post-Reversal (the same grating changed from being a food cue pre-Reversal to being a neutral cue during- and post-Reversal; see Figure 1D). For PO neurons, note the response component prior to motor onset and prior to Ensure delivery. For BR neurons, note the relatively larger response component prior to motor onset as compared to ID neurons, and the response component to both the food cue and the neutral cue (for clarity, the quinine cue component was not shown in these plots).
B. We quantified the fraction of the deviance explained by each behavioral event included in the GLM (normalized by the total deviance explained for each neuron). Note the use of the "xy-shift" variable (see STAR Methods) to account for brain motion. Here, we grouped the visual response component across all visual cues. Note that the highest deviance fraction belongs to the visual component. Pre: pre-Reversal; During: during-Reversal; Post: post-Reversal.
C. Fraction of days visually-driven, for the six subgroups of neurons ranging from those that mostly tracked the orientation of a stimulus (left; see Figure 2F) to those that mostly tracked the predicted outcome associated with a stimulus. This category-free analysis shows that those neurons that track predicted outcome cease to be reliably responsive during-Reversal and regain responsivity (now to the new food-predicting cue) post-Reversal.
D. Plot of the instability of beta coefficients across time (see STAR Methods). Both the visual responses as well as the responses to other behavioral event types (see x-axis in B) were more unstable in PO neurons. This response instability is quantified as the Euclidean distance between the coefficients on a single day and the mean coefficients across all daily sessions (each dot corresponds to a single neuron). * denotes significance vs. all other groups: $p < 0.015$, Kruskal-Wallis, Bonferroni corrected. Error bars denote s.e.m. This suggests higher plasticity in PO neuron response profiles vs. ID and BR neuron response profiles across sessions.

Figure S5

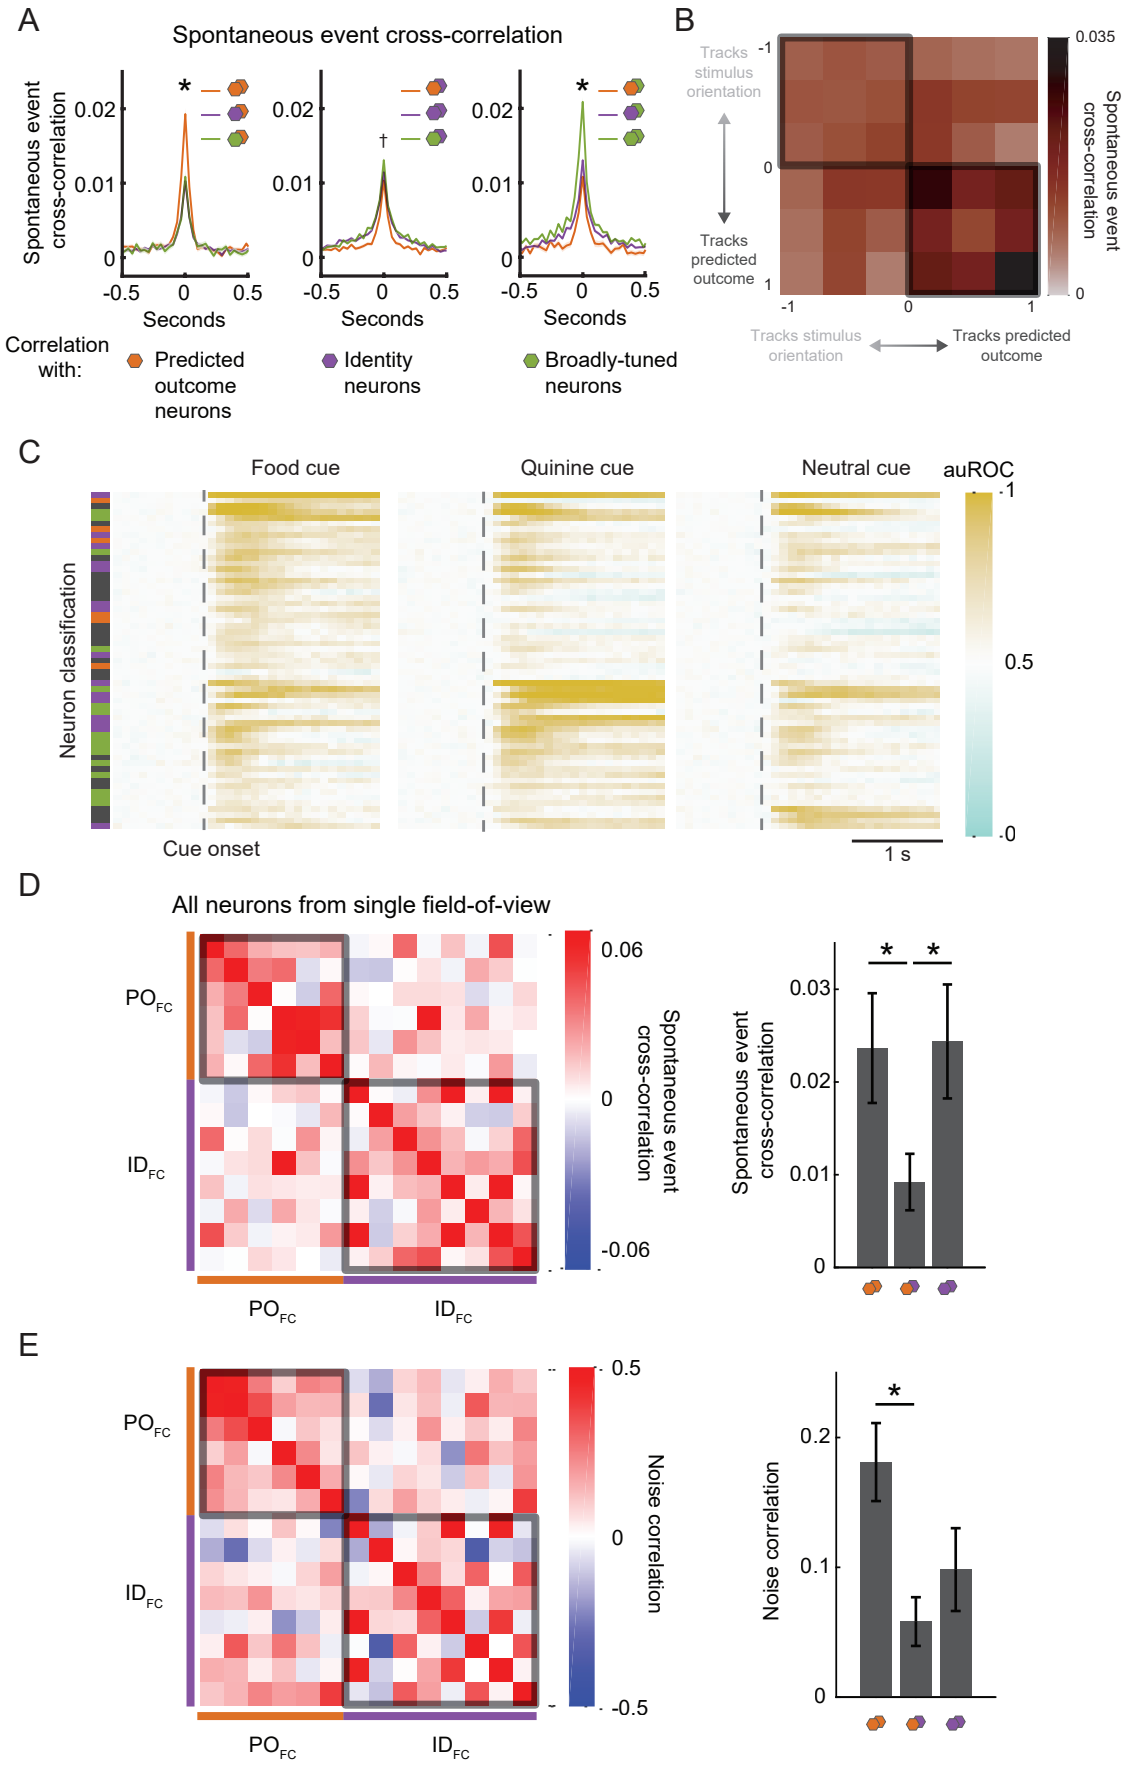
A
Spontaneous event cross-correlation
Spontaneous event cross-correlation
0.02
0.01
0
-0.5
0
0.5
Seconds
*
†
Correlation with:
Predicted outcome neurons
Identity neurons
Broadly-tuned neurons
B
Tracks stimulus orientation
Tracks predicted outcome
-1
0
1
0.035
Spontaneous event cross-correlation
C
Food cue
Quinine cue
Neutral cue
auROC
Neuron classification
Cue onset
1 s
1
0.5
0
D
All neurons from single field-of-view
PO_FC
ID_FC
0.06
0
-0.06
Spontaneous event cross-correlation
0.03
0.02
0.01
E
0.5
-0.5
Noise correlation
0.2
0.1

Figure S5. Neurons in lateral visual association cortex can be subdivided into ensembles based on correlated variability, related to Figure 5.

A. All pairwise spontaneous event cross-correlations of Predicted Outcome (PO), Identity (ID), and Broadly-tuned (BR) neurons. Color of cross-correlation corresponds to correlation of cells from a given group (left hexagon in each pair) with cells from each group (right hexagon in each pair; *left panel*: with PO neurons; *middle panel*: with ID neurons; *right panel*: with BR neurons). * denotes that correlation coefficient at zero phase lag for within-group pairs (two hexagons of same color, e.g. PO-PO, orange-orange) is significantly higher than all other across-group comparisons: $p < 0.0001$. † denotes that Identity-Identity correlation coefficient is significantly higher than the ID-PO coefficient: $p < 0.01$, Kruskal-Wallis, Tukey-Kramer multiple comparisons test.
B. A spontaneous cross-correlation heatmap for all neurons recorded across both pre- and post-Reversal, divided (as in Fig. 3B) into 6 subgroups of neurons ranging from those that mostly tracked the orientation of a stimulus (left) to those that mostly tracked the predicted outcome of the stimulus (for comparison, see Figure 5C). Each datapoint is derived from the mean cross-correlation of all relevant pairs of cells. Those neurons that tracked stimulus identity tended to have higher spontaneous cross-correlations with other neurons that tracked stimulus identity. Similarly, those neurons that tracked the predicted outcome tended to have higher spontaneous cross-correlations with other neurons that tracked predicted outcome. Note the lack of correlation between those neurons that weakly tracked stimulus orientation and those that weakly tracked predicted outcome, supporting our choice of boundary in Figure 2 for categorizing neurons as ID vs. PO neurons.
C. The response timecourse of each visually driven neuron from a single field of view during a single imaging session (n = 59 neurons). Neurons that could be tracked across the Reversal are denoted by the color associated with that category (left; gray denotes neurons that did not contain identified masks both pre- and post-Reversal, or that were not found to be visually responsive both pre- and post-Reversal). Colorbar at right denotes auROC response (see STAR Methods).
D. Left: All pairwise spontaneous event cross-correlations between PO neurons that prefer the food cue (PO_{FC}) with ID neurons that prefer the food cue (ID_{FC}) from the same session as in S5C. Within each of the two groups (gray squares), each neuron was sorted by the relative 'confidence' (as in Figure S2C) in tracking a specific predicted outcome (upper rows) vs. a specific orientation (lower rows). More specifically, within the group of PO neurons (upper left gray outlined square), the top left corner contains those PO neurons that most strongly and specifically tracked the predicted outcome, while the bottom right corner contains those PO neurons that were least distinguishable from ID neurons. Note the higher within-group correlations as compared to across-group correlations, even for those neurons that prefer the same cue. Right: the existence of higher spontaneous correlations within-group (hexagons with the same color, PO-PO and ID-ID) vs. across-group (PO-ID) was not only evident when pooled across sessions and mice (Figure 5F), but was also evident in individual imaging sessions, as shown here. * $p < 0.05$, Kruskal-Wallis, Bonferroni corrected.
E. Left: Same as D, but for noise correlations instead of spontaneous event cross-correlations. All pairwise noise correlations of FC-preferring PO neurons and ID neurons. Neurons are sorted as in D. Right: even in data from a single imaging session, we observed higher noise correlations within-group vs. across-group. * $p < 0.05$, Kruskal-Wallis, Bonferroni corrected. Error bars and discs denote s.e.m.

Figure S6

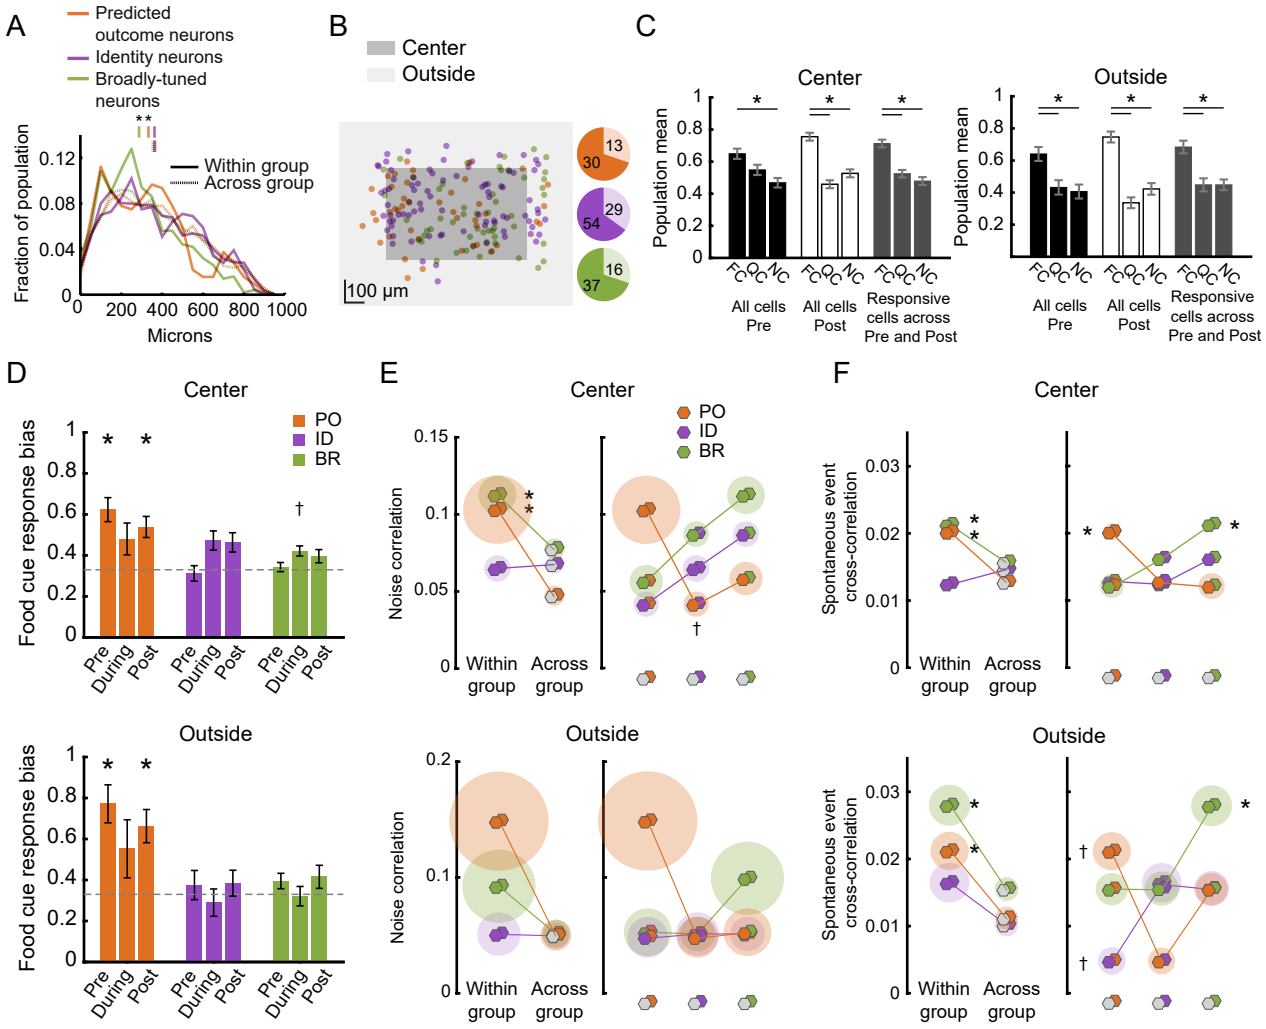
A
Predicted outcome neurons
Identity neurons
Broadly-tuned neurons
Fraction of population
Within group
Across group
Microns
B
Center
Outside
100 μm
30
13
54
29
37
16
C
Center
Outside
Population mean
All cells Pre
All cells Post
Responsive cells across Pre and Post
D
Center
Outside
Food cue response bias
PO
ID
BR
Pre
During
Post
E
Center
Outside
Noise correlation
Within group
Across group
F
Center
Outside
Spontaneous event cross-correlation

Figure S6. Spatial organization of neuronal ensembles in visual association cortex, related to Figure 5.

A. A comparison of distributions of distances between all pairs of neurons in each functional category demonstrated that, on average, there was a small but significant trend towards spatial clustering of Predicted Outcome (PO) neurons and Broadly-tuned (BR) neurons compared to Identity (ID) neurons. * denotes $p < 0.05$ Wilcoxon Rank-Sum test, Bonferroni corrected.
B. As our imaging fields were centered around postrhinal cortex (visPOR) but encompassed a broader area of visual association cortex, we assessed the spatial distribution of neurons belonging to each functional category. The proportion of neurons belonging to each category was similar in the central portion of the imaging field (dark shading) vs. in the surrounding regions (light shading). Here, data was combined across all fields of view. Right: pie chart shows number of neurons within the central region vs. those outside of the central region, across all fields of view.
C. The population response bias to the food cue across all recorded neurons (see Figure 3A) was similar to that observed both for neurons within the central portion of the imaging field (left) and for those outside of this central region (right). Pre: pre-Reversal; During: during-Reversal; Post: post-Reversal. FC: food cue; QC: quinine cue; NC: neutral cue.
D. The larger food cue response bias in PO neurons was evident both for cells located inside (top) and outside (bottom) of the central region of our imaging fields of view, and both were similar to the combined result plotted in Figure 3C.
E. The ensemble noise correlation analyses also yield similar results for cells inside (top) and outside (bottom) of the central portion of our imaging field. Both results were similar to the overall ensemble noise correlations in Figure 5B.
F. Spontaneous event cross-correlation results were also similar when separately considering cells inside (top) and outside (bottom) of the central portion of our imaging fields of view. Both results are similar to the overall ensemble spontaneous correlations in Figure 5F.

Figure S7

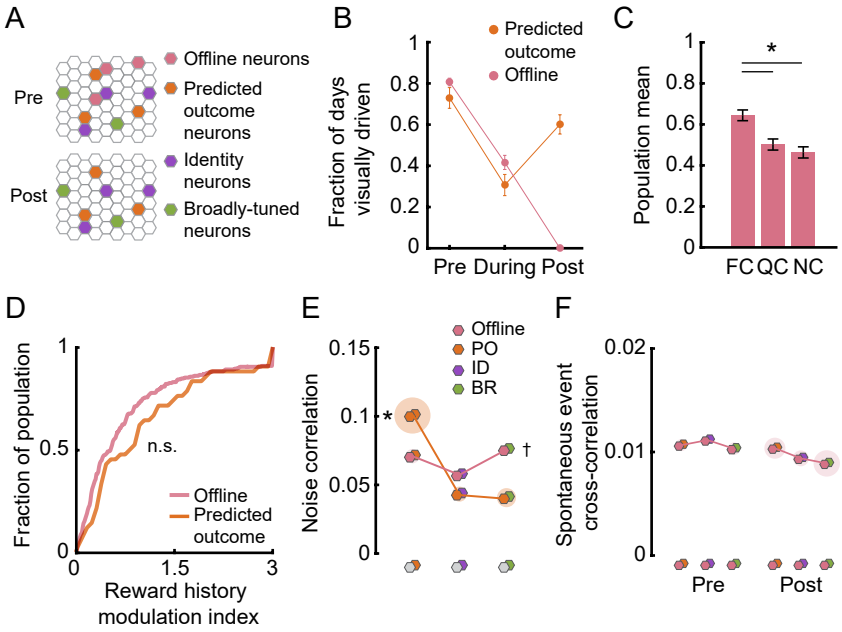
A
Pre
Post
Offline neurons
Predicted outcome neurons
Identity neurons
Broadly-tuned neurons
B
Predicted outcome
Offline
Fraction of days visually driven
1
0.8
0.6
0.4
0.2
0
Pre During Post
C
Population mean
*
1
0.8
0.6
0.4
0.2
0
FC QC NC
D
Fraction of population
1
0.5
0
n.s.
Offline
Predicted outcome
0
1.5
3
Reward history modulation index
E
Noise correlation
0.15
0.1
0.05
0
*
†
Offline
PO
ID
BR
F
Spontaneous event cross-correlation
0.02
0.01
0
Pre
Post

Figure S7. A subset of neurons ceases to respond to visual stimuli during new learning, related to Figure 7.

A. A schematic demonstrating the hypothesized activity of “Offline” neurons during our Reversal task.
B. Offline neurons had similar responsivity to visual cues as Predicted Outcome (PO) neurons before and during Reversal learning, but ceased to respond to visual cues post-Reversal. Pre: pre-Reversal; During: during-Reversal; Post: post-Reversal.
C. This group of neurons was slightly but significantly biased to the food cue. FC: food cue; QC: quinine cue; NC: neutral cue.
D. Offline neurons had similar response modulation by reward history as PO neurons pre-Reversal.
E. In contrast to Recruited neurons (Figure 7F), Offline neurons did not show stronger noise correlations with PO neurons than with ID or BR neurons. This provides a possible explanation as to why Offline neurons did not remain responsive across a change in cue-outcome associations. * denotes significant difference from all other groups: $p < 0.01$, Kruskal-Wallis, Bonferroni corrected. † denotes $p < 0.05$ with Offline-BR vs. Offline-ID neurons. ID: Identity; BR: Broadly-tuned.
F. In contrast to Recruited neurons (Figure 7G), Offline neurons also did not show increased spontaneous event cross-correlations with PO neurons compared to ID and BR neurons.

Figure S8

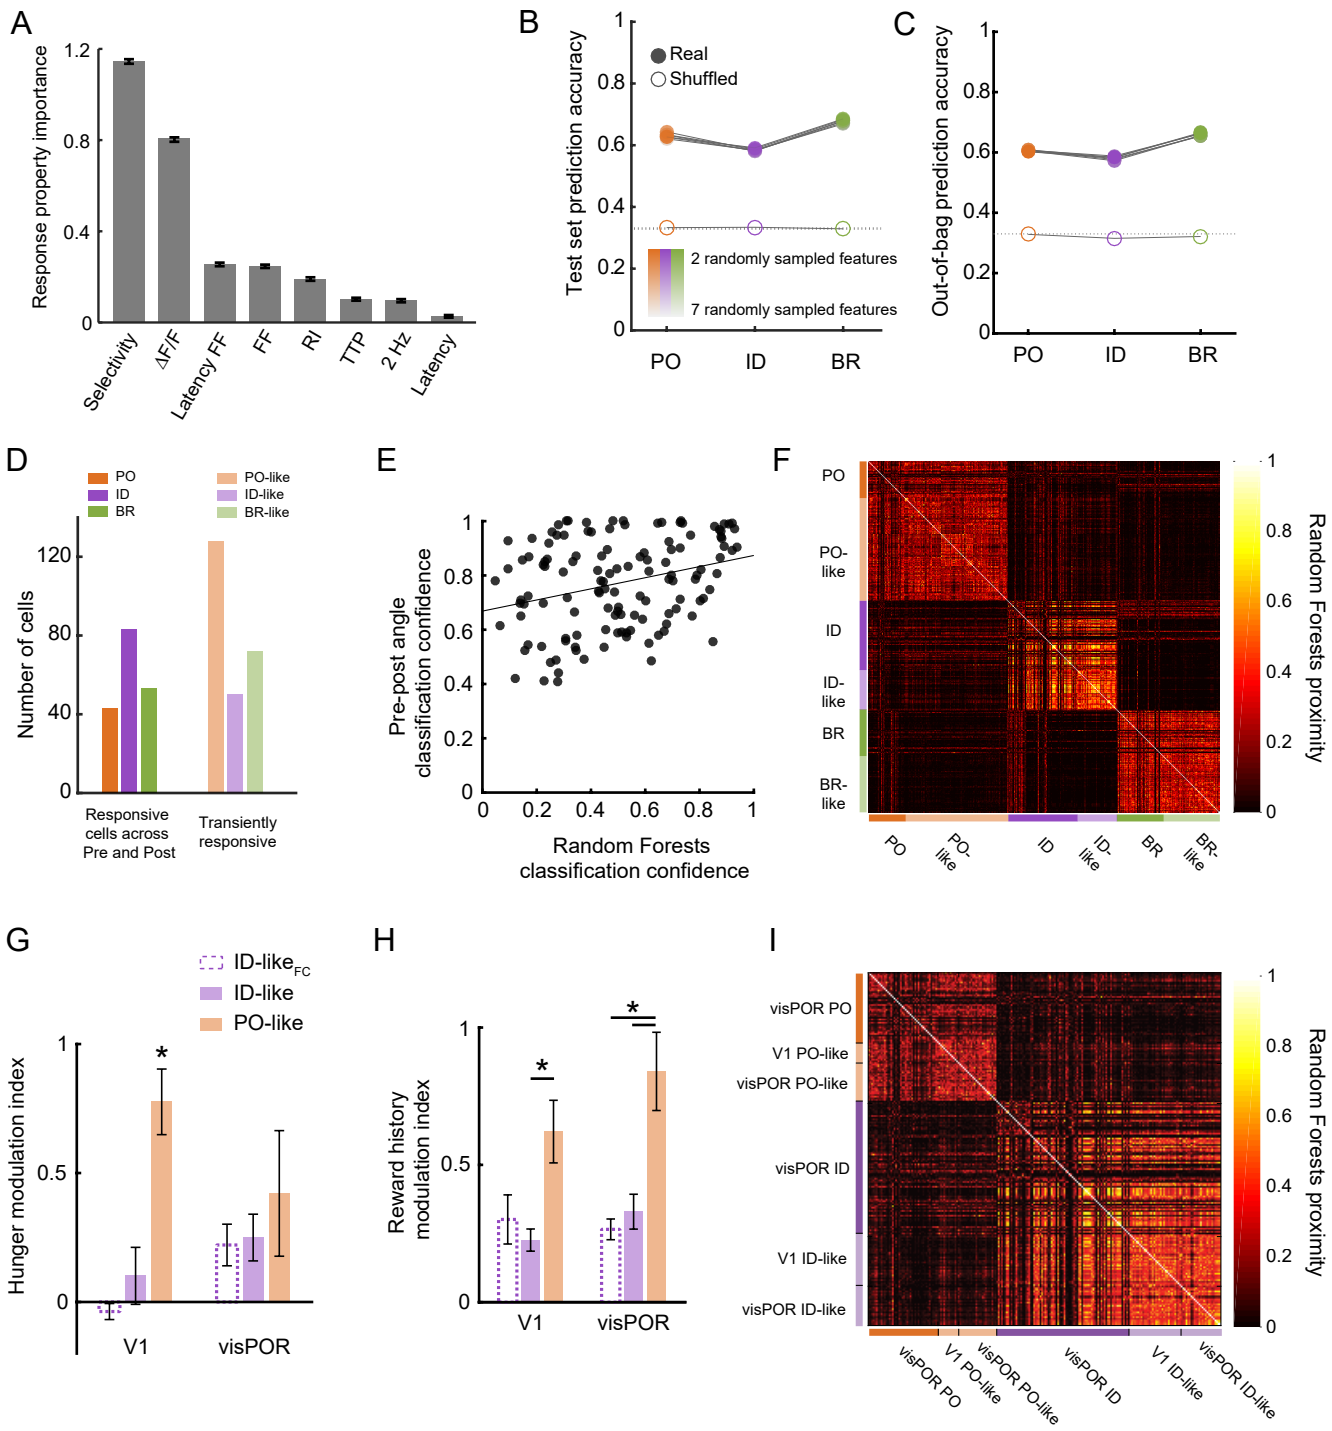
A
Response property importance
1.2
0.8
0.4
0
Selectivity
ΔF/F
Latency FF
FF
RI
TTP
2 Hz
Latency
B
Test set prediction accuracy
Real
Shuffled
2 randomly sampled features
7 randomly sampled features
PO
ID
BR
C
Out-of-bag prediction accuracy
PO
ID
BR
D
PO
ID
BR
PO-like
ID-like
BR-like
Number of cells
120
80
40
0
Responsive cells across Pre and Post
Transiently responsive
E
Pre-post angle classification confidence
Random Forests classification confidence
F
PO
PO-like
ID
ID-like
BR
BR-like
Random Forests proximity
G
ID-like$_{FC}$
ID-like
PO-like
Hunger modulation index
V1
visPOR
H
Reward history modulation index
V1
visPOR
I
visPOR PO
V1 PO-like
visPOR PO-like
visPOR ID
V1 ID-like
visPOR ID-like
Random Forests proximity

Figure S8. Validation of the Random Forests classifier and application to transiently-responsive neurons, related to Figure 8.

A. Transiently-responsive neurons with low-level visual response characteristics that were similar to either PO, ID, or BR neurons were termed 'PO-like', 'ID-like', or 'BR-like' neurons, respectively. Here, we calculated the relative importance of individual low-level visual response characteristics for the classification of neurons as Predicted Outcome (PO)-like, Identity (ID)-like, and BR-like by the Random Forests classifier (see also STAR Methods). The selectivity value was estimated as the breadth of the 3-point tuning curve. The $\Delta F/F$ value quantified the magnitude of the evoked response. The Latency Fano factor (Latency FF) quantified the trial-to-trial reliability of the latency of the cue-evoked response relative to stimulus onset. The Fano factor (FF) quantified the trial-to-trial reliability in the magnitude of the evoked response to the preferred stimulus. The ramp index (RI) quantified whether the response magnitude ramped up or down during stimulus presentation (Makino and Komiyama, 2015). The time-to-peak (TTP) quantified when the peak of the response occurred following stimulus onset. The 2 Hz temporal frequency locking (2 Hz) quantified the degree to which a neuron's activity tracked the 2 Hz temporal frequency of the visual stimulus (see Figure 8B). The latency was estimated as the first time bin after stimulus onset for which the stimulus-evoked response was significantly above baseline. See STAR Methods for further details as to how each response characteristic was derived.
B. We considered how sensitive each classification tree was to the number of features sampled at each decision node. Regardless of how many features were sampled, test set prediction accuracy was stable. We used three randomly-sampled features for all subsequent analyses (traditionally, the number of features sampled to construct each classification tree is the square root of the total number of features; see STAR Methods; Breiman, 2001). Importantly, shuffling the labels of PO, ID, and BR neurons reduced performance to chance levels (0.33).
C. As expected, the out-of-bag prediction accuracy (the prediction accuracy of those cases held out during construction of a tree) matched the test set prediction accuracy (see Figure S8B and STAR Methods).
D. Left: total numbers of experimentally-identified PO, ID, and BR neurons. These neurons were required to exhibit a significant visual response both pre- and post-Reversal. Right: numbers of PO-like, ID-like, and BR-like neurons.
E. There was a strong correlation between the Random Forests classification confidence and the confidence metric we generated in our initial determination of PO, ID, and BR neurons (Pearson correlation coefficient between confidence measures: 0.3, $p < 0.001$; as described in Figure S2 and STAR Methods).
F. Random Forests proximities (see STAR Methods) demonstrate that most neurons classified as PO-like, ID- like, or BR-like showed high similarity in low-level response features to other neurons in their training group (i.e. PO, ID, or BR neurons, respectively), but not to other training groups. Importantly, these response features were independent of those used for the initial classification (Figure S2A), which depended on whether a neuron tracked stimulus orientation or predicted outcome *across sessions* pre- vs. post-Reversal. Error bars denote s.e.m.
G. In primary visual cortex (V1), food cue responses in PO-like neurons are more modulated by hunger state than food cue responses in ID-like neurons. This was also the case when analysis was restricted to food cue-preferring ID-like neurons (ID-like_{FC}). * denotes significance vs. both other V1 groups: $p < 0.02$, Kruskal-Wallis, Bonferroni corrected (we observed a similar trend in visPOR).

H. Food cue responses in V1 and in visPOR were more modulated by recent reward history in PO-like neurons than in ID-like neurons. * denotes $p < 0.01$ for highlighted comparisons, Kruskal-Wallis, Bonferroni corrected. Food cue responses in visPOR were also more modulated by recent reward history in PO-like neurons than in ID-like$_{FC}$ neurons. If ID-like neurons were to jointly encode stimulus identity and value, we would predict that those ID-like neurons that prefer the food cue would be modulated by changes in hunger state and recent reward history. Instead, we found that ID-like$_{FC}$ neurons were not more sensitive to hunger-state or reward history than other ID-like neurons. This further suggests a *lack* of value coding by ID-like neurons and ID neurons. * denotes $p < 0.01$ for highlighted comparisons, Kruskal-Wallis, Bonferroni corrected. Error bars denote s.e.m.

I. Having applied our Random Forests classifier to previous recordings in primary visual cortex (V1) and in postrhinal cortex (visPOR), we calculated proximities based on low-level response characteristics, in a similar manner as above. This revealed clustering of neurons with similar low-level visual response characteristics, even across two entirely different sets of experiments (dark orange and purple: current data set; light orange and purple: data set from Burgess et al., 2016).

Supplemental Table 1
(related to STAR Methods section: *Identification of functional groups of neurons*)

	Pre-Reversal response			Post-Reversal response			Neuron classification
Cues	FC	QC	NC	NC	FC	QC	
Orientations	0°	270°	135°	0°	270°	135°	
Candidate neural response tuning #1	1	0.2	0.3	1	0.2	0.3	Identity (no rotation)
Candidate neural response tuning #2				0.3	1	0.2	Predicted outcome (clockwise rotation)
Candidate neural response tuning #3				0.2	0.3	1	Null (counter-clockwise rotation)

Supplemental Table 2
(related to STAR Methods section*: Joint tracking index*)

	Pre-Reversal response			Post-Reversal response			Residual calculation		
Cues	FC	QC	NC	NC	FC	QC	$[PO_{residual}]$	$[SF_{residual}]$	$[BR_{residual}]$
Orientations	0°	270°	135°	0°	270°	135°			
Normalized response (Pre-Reversal FC preferring, FC=1)	1	a	b	x	y	z	y – a	1-x	z – b
				1	a	b			
Normalized response (Pre-Reversal QC preferring, QC=1)	a	1	b	x	y	z	z - b	1-y	x - a
				a	1	b			
Normalized response (Pre-Reversal NC preferring, NC=1)	a	b	1	x	y	z	x - a	1-z	y - b
				a	b	1			