History-dependent variability in population dynamics during evidence accumulation in cortex

Ari S Morcos & Christopher D Harvey

We studied how the posterior parietal cortex combines new information with ongoing activity dynamics as mice accumulate evidence during a virtual navigation task. Using new methods to analyze population activity on single trials, we found that activity transitioned rapidly between different sets of active neurons. Each event in a trial, whether an evidence cue or a behavioral choice, caused seconds-long modifications to the probabilities that govern how one activity pattern transitions to the next, forming a short-term memory. A sequence of evidence cues triggered a chain of these modifications resulting in a signal for accumulated evidence. Multiple distinguishable activity patterns were possible for the same accumulated evidence because representations of ongoing events were influenced by previous within- and across-trial events. Therefore, evidence accumulation need not require the explicit competition between groups of neurons, as in winner-take-all models, but could instead emerge implicitly from general dynamical properties that instantiate short-term memory.

In cortical microcircuits, ongoing activity patterns are combined with new inputs to perform many complex neural computations, including evidence accumulation during decision-making[1,2]. To understand how ongoing activity is combined with external inputs, considerable focus has been placed on the posterior parietal cortex (PPC)[3,4], which is thought to be necessary for visual decision-making tasks in rodents[5–8]. Previous work has emphasized models in which evidence accumulation occurs as a winner-take-all competition between neuronal activity patterns associated with different decisions[2]. This view predicts that, as evidence is accumulated, activity converges to one of several attractor states, each associated with a different decision. Winner-take-all dynamics have commonly been implemented as a highly structured competition between distinct recurrently connected pools of neurons with mutual inhibition across pools[9,10]. Predictions of these models, including long-lasting firing rate changes in homogeneous pools of individual neurons, have been supported by some experimental data[3,4,11]. However, recent work showing the prevalence of time-varying activity patterns in neuronal populations[8,12–15] provides initial suggestions of potential alternatives. For example, alternative implementations of winner-take-all competitions could also be possible, such as competitions between sequences of population activity. Or, entirely different algorithms for evidence accumulation might be present that do not require winner-take-all mechanisms.

Here we expanded the study of evidence accumulation in two ways. First, previous work has often emphasized independent recordings from selected subsets of individual neurons, typically summarized as averages across trials and cells. However, because animals make decisions on single trials using populations of neurons, we developed new experimental and computational methods to reveal structure in the moment-to-moment changes in population activity. Second, because models proposing mechanisms other than winner-take-all competitions have not emerged, we not only compared our data with winner-take-all dynamics but also took an exploratory approach aimed at uncovering results that might motivate new conceptual models for evidence accumulation. The starting point for our conceptual framework was our previous work in the mouse PPC in which neuronal activity was described as a trajectory through time-varying population activity patterns[8].

We found that the PPC had long-timescale dynamics in the form of orderly transitions between transient and largely different patterns of population activity. As a result, the representation of new inputs depended both on the identity of the input and the near-past activity patterns in the population. PPC activity never reset but rather functioned as a continuous record of recent events. In addition, multiple task-relevant features were represented simultaneously such that individual task features (for example, choice) did not converge to single activity patterns but instead were represented across trials by many different activity patterns. Our results motivate a new model in which a winner-take-all competition between distinct pools of neurons would not be necessary. Rather, evidence accumulation may emerge from general, long-timescale dynamical properties, which would naturally form a history of the sequence of past events and thus create a short-term memory from which information, such as accumulated evidence, could be read out.

RESULTS

We developed a navigation-based evidence accumulation task in which a head-restrained mouse ran down a virtual-reality T-maze. The mouse was presented with six visual cues that could each appear on the left or the right wall at fixed locations (**Fig. 1a**,**b** and **Supplementary Fig. 1**; see Online Methods). To receive a reward, the mouse had to turn toward the direction that had more cues. Task difficulty was modulated by varying the net evidence, defined as the

Department of Neurobiology, Harvard Medical School, Boston, Massachusetts, USA. Correspondence should be addressed to C.D.H. (harvey@hms.harvard.edu).

Received 13 July; accepted 6 September; published online 3 October 2016; doi:10.1038/nn.4403

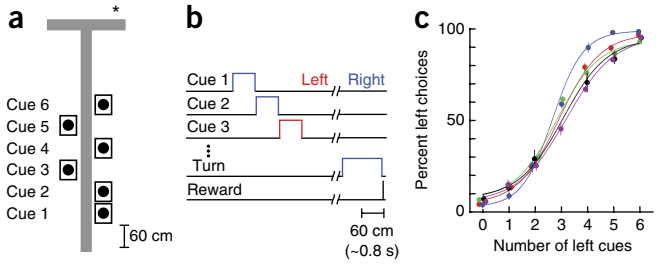

Figure 1 A navigation-based evidence accumulation task in virtual reality. (**a**) Schematic of an example 2–4 right trial in a virtual T-maze. Asterisk marks the reward location. (**b**) Sequence of trial events. (**c**) Performance for the five mice imaged (mean ± s.e.m., 7–12 sessions per mouse).

difference between the number of left and right cues (six total cues per trial). Mice performed the task with high accuracy by accumulating multiple pieces of evidence per trial, with a bias toward earlier segments (**Fig. 1c** and **Supplementary Fig. 2**; see Online Methods).

Distributed population representations of choice- and net evidence-related information

We first examined the distribution of activity patterns in individual neurons. We used calcium imaging to measure the activity of ~350 neurons simultaneously in layer 2/3 of the PPC and estimated spike counts using deconvolution of the fluorescence traces[16] (Online Methods; **Supplementary Fig. 3** and **Supplementary Table 1**). Consistent with our previous work, most neurons were transiently active for less than 10% of the trial on average, and different neurons were active at different points in the trial, such that across the population, activity tiled the full trial duration[8] (**Fig. 2a**,**b** and **Supplementary Figs. 4** and **5a**,**b**). To test for differences in activity between trials with different choices and net evidence, we used a support vector machine (SVM) to predict choice on the basis of a single cell's activity and a support vector regression (SVR) model to predict net evidence from a single cell's activity (Online Methods). Some neurons had a statistically significant choice classification accuracy, and some neurons had a significant relationship between the actual net evidence and the net evidence predicted from their activity (choice 29.4%, net evidence 22.7%; 5% expected by chance; **Fig. 2c**,**d**). When we plotted the mean activity patterns for the significantly choice-selective neurons, we identified choice-specific sequences of activity[8] (**Supplementary Fig. 4c**,**d**).

We next considered the entire population of neurons to determine whether task-relevant information was present only in the fraction of cells that had high selectivity or whether neurons that did not have statistically significant selectivity might contribute small amounts of information to a population code. The population activity (concatenated activity of all individual neurons) contained information about the choice and net evidence on single trials, as revealed using a SVM classifier for choice and a SVR model to predict net evidence (**Fig. 2e**,**f**). Information about choice and net evidence could not be explained only by behavioral differences between trials of different choices and net evidence, such as differences in running patterns in the maze (**Supplementary Fig. 5e–h** and Online Methods, "Contribution of behavioral variability to neuronal activity results"). We examined the distribution of information within the population by applying the population activity classifiers for choice and net evidence to increasingly larger subsets of neurons, beginning with neurons with the lowest individual classification accuracy. The accuracy of both classifiers increased with the incorporation of neurons that individually represented choice and net evidence poorly (**Fig. 2g**,**h**). Using the 40% least selective neurons, we were able to predict the mouse's choice with ~75% accuracy (**Fig. 2g**, **Supplementary Fig. 5j** and Online Methods). These results suggest a population representation in which information is distributed across heterogeneous and variable neurons[5,17–22].

Clustering-based methods for analyzing population activity dynamics on single trials

Given that neuronal activity was in large part heterogeneous across neurons and variable between trials and that task-relevant information was distributed across neurons, we focused on how the population activity pattern changed from moment to moment. Because methods to analyze moment-to-moment transitions between transient population activity patterns have not been commonly used previously, we developed a new analysis framework. We defined the population activity pattern within a given time period as a vector of each neuron's estimated spike count. We considered the population activity to be a trajectory involving transitions from one activity pattern to another. To facilitate the analysis and visualization of transitions between patterns, we reduced the dimensionality of the population activity using a clustering algorithm to group similar population activity patterns (**Fig. 3a** and Online Methods).

Specifically, we determined the number of clusters using the affinity propagation clustering algorithm[23]. Our results were consistent across a wide range of cluster numbers and affinity propagation settings (**Supplementary Fig. 6j** and Online Methods). Clustering was performed independently for ten epochs in the trial. For each epoch, the estimated spike count on each of m trials for each of n simultaneously imaged neurons was calculated, resulting in m points in an n-dimensional space (**Fig. 3a**). We clustered these m points such that each cluster corresponded to a different set of trials with similar population activity patterns at a given epoch. For visualization, each cluster was represented as a circular node with area proportional to the number of trials in the cluster (**Fig. 3a–d** and **Supplementary Fig. 6e**). Transitions between clusters in adjacent epochs were marked as lines with thickness proportional to the transition probability (**Fig. 3b–d** and Online Methods).

Single trials could therefore be described as an activity trajectory defined by the sequence of clusters visited from epoch to epoch (**Fig. 3b–d**). These cluster-space trajectories are conceptually identical to trajectories that have previously been described using principal component analysis and other methods; the only difference is in the dimensionality reduction algorithm used[5,8,19,24–26]. Activity patterns reflecting important task-relevant features, including choice and net evidence, were apparent in the cluster space, even though clustering was performed on neuronal activity alone without any information about behavioral parameters (**Fig. 3b–d** and **Supplementary Fig. 6a–d**). For choice, for example, different paths through clusters emerged for left- and right-choice trials, which is a visualization of choice-specific activity trajectories[8] (**Fig. 3b–d**).

Before exploring population dynamics in the cluster space, we sought to gain an intuition on how neuronal activity patterns related to the clusters. We visualized the relationship between neuronal activity and clusters by calculating for each pair of trials the correlation between their population activity patterns at a given epoch. We then sorted the matrix of trial–trial correlation coefficients by the trials that were clustered together (**Fig. 3e–g**). This visualization revealed that clustering identified structure in the trial–trial activity pattern correlations and showed that clusters varied over a wide distribution in how similar they were to one another. As expected by the transient activity we observed in individual

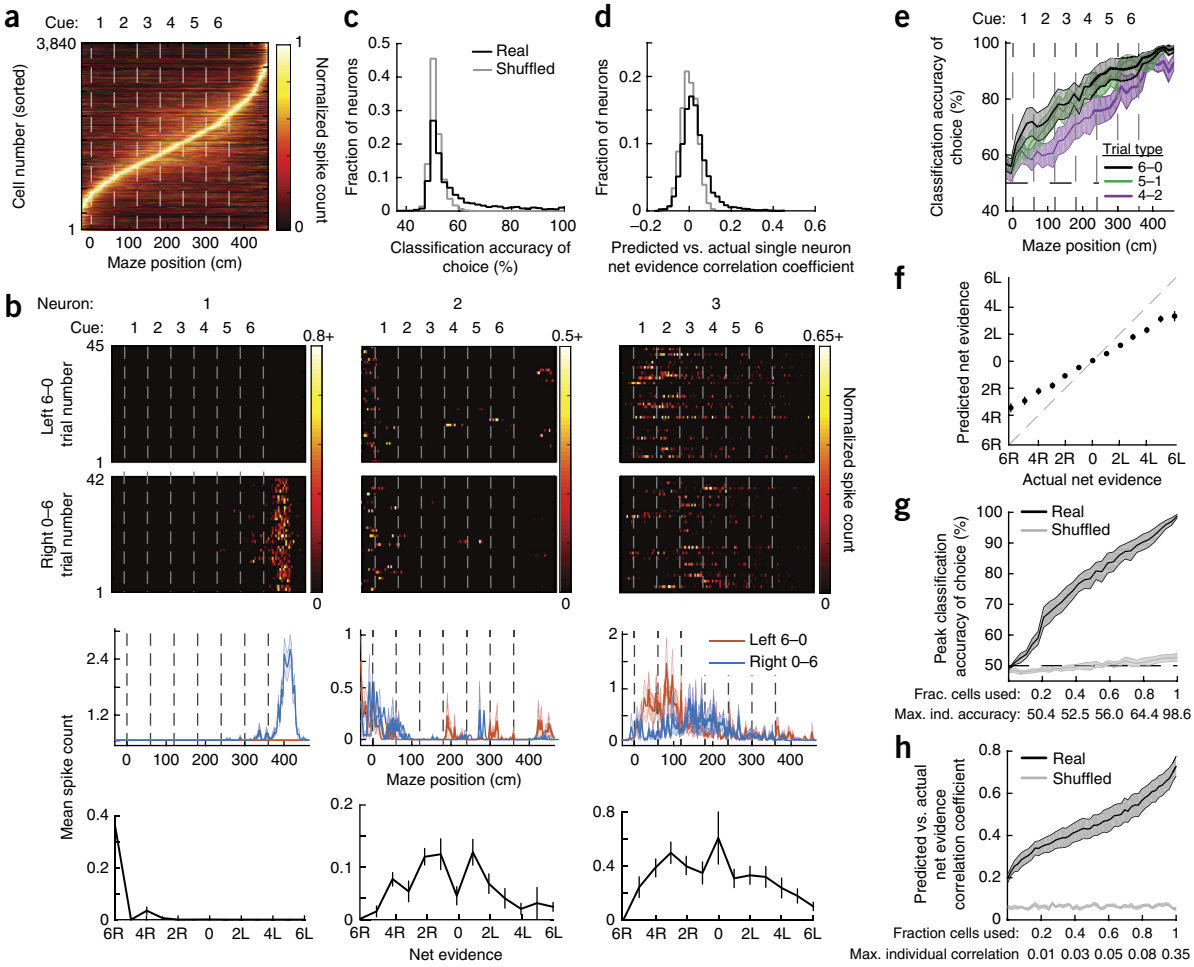

Figure 2 Distributed representation of task-relevant information across PPC neurons. (**a**) Normalized mean activity across all trials for all neurons pooled across all data sets (n = 3,840 cells from 5 mice). Traces were normalized to the peak of each cell's activity, averaged, and sorted by the peak's maze position. (**b**) Single trial activity for three example neurons. Top panels: each row is an individual trial. Bottom panels: mean ± s.e.m. For each net evidence condition (for example, 2L), the mean spike count was calculated by combining the activity at all cue epochs matching the given net evidence. (**c**) Histogram of SVM classification accuracy for choice for individual neurons (black) and with shuffled trial labels (gray), based on the average activity within each left 6–0 and right 0–6 trial. (**d**) Histogram of SVR model performance using all trial types, quantified as the correlation between the actual net evidence and the net evidence predicted by the SVR model, for individual neurons (black) and with shuffled net evidence labels (gray). (**e**) Classification accuracy (mean ± s.e.m., n = 11 data sets) for choice using an SVM based on population activity. Independent classifiers were trained and tested at each maze position. (**f**) Actual net evidence vs. net evidence predicted by a SVR classifier trained on population activity across all cue epochs and trial types. Error bars represent mean ± s.e.m. across data sets (n = 11). (**g**,**h**) Peak classifier accuracy for choice (**g**) and the predicted vs. actual net evidence correlation coefficient (**h**) for classifiers constructed with increasing numbers of neurons, added from least to most selective (based on histograms from **c** and **d**). Real data (black) and data with shuffled trial labels (gray) are shown. Lines and shaded areas represent mean ± s.e.m. across data sets, and maximum individual neuron classification accuracies and correlations were the mean across data sets.

neurons (**Fig. 2a**,**b** and **Supplementary Figs. 4** and **5a**,**b**), the activity patterns in clusters at one epoch were largely different from the activity patterns observed in clusters at the next epoch (**Fig. 3h**). Consistently, when we clustered activity patterns from all epochs together, rather than for single epochs individually, such that the clusters were the same from epoch to epoch, we found that the likelihood of a trial staying in the same cluster across consecutive epochs was rare (0.9 ± 0.01% of transitions; **Supplementary Fig. 6f** and Online Methods). The activity patterns in each cluster were made up of complex combinations of activity levels in the population of individual neurons (**Fig. 3j** and **Supplementary Fig. 7a–e**). Some individual neurons thus had elevated activity in multiple clusters (**Fig. 3j** and **Supplementary Figs. 6g–i**,**7a–e**). A cluster should therefore be considered as a pattern of activity across neurons, such that the patterns between clusters are discriminable from one another. The precise activity patterns that defined each cluster were not important for the focus of this work.

Highly variable population activity patterns on trials with identical cues and choices

We used the cluster space to visualize the population activity trajectories on single trials. This visualization revealed a high amount of trial–trial variability, as trials with the same choice and evidence cues (for example, correct trials with all cues on the left, or left 6–0 trials) occupied more than half of all possible clusters at each epoch, even at the turn after a choice was made (**Fig. 3c**,**d**,**i**). Trials of the same type therefore had distinguishable trajectories of activity patterns and did not converge to similar paths through a small set of clusters, which is consistent with previous studies of variability in the activity of cortical neurons[21,22,27–29].

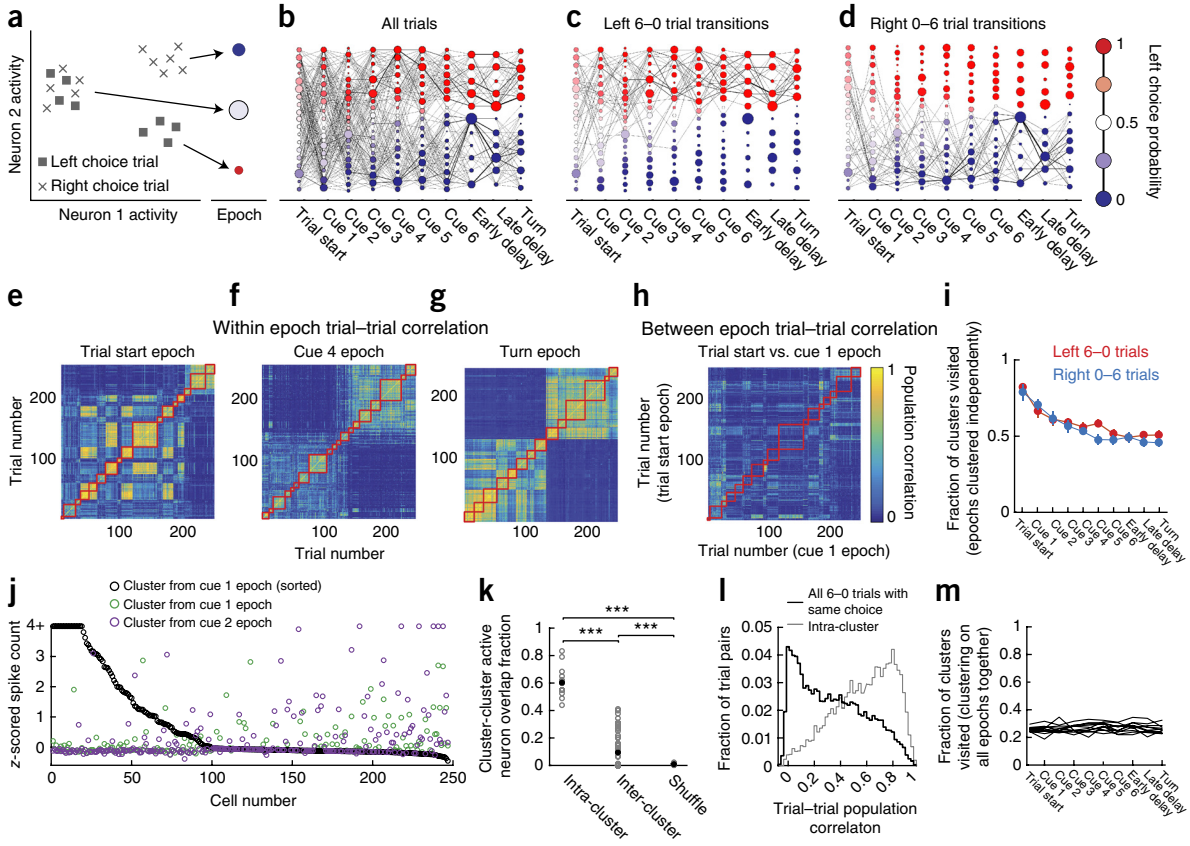

Figure 3 Clustering neuronal activity across trials reveals trial-to-trial variability. (**a**) Clustering procedure (Online Methods). At each of ten spatially defined maze epochs, clustering grouped individual trials with similar population activity patterns. Clusters at each maze epoch were represented as a column of nodes with area proportional to the number of trials in each cluster. Node colors were based on the fraction of trials within each cluster resulting in a left choice. Nodes were sorted vertically from largest to smallest left choice probability. (**b**) Example transition matrix constructed from all trials in a single data set. Edge widths between nodes represent the forward transition probability. (**c**,**d**) Transition probabilities for left 6–0 (**c**) and right 0–6 (**d**) trials using the clusters derived from all trials as in **b**. (**e**–**g**) Example trial–trial population activity correlation matrices at the trial start epoch (**e**), cue 4 epoch (**f**), and the turn epoch (**g**) sorted by cluster identity. Red squares indicate pairs of trials that were in the same cluster. (**h**) Example trial–trial population activity correlation matrix for two consecutive epochs (trial start epoch compared to cue 1 epoch). Trials were sorted according to the cluster identity during the trial start epoch. Because trials were sorted identically in both epochs, trial pairs along the diagonal would be expected to have high correlations if trial activity was similar in consecutive epochs. In contrast, the low correlations along the diagonal suggest that trials had highly different population activity in consecutive epochs. (**i**) Fraction of clusters visited at each epoch decreased to only 0.5, suggesting that much trial–trial variability remained even at the turn. Error bars represent mean ± s.e.m. across data sets ($n = 11$). (**j**) z-scored activity of cells in two clusters during correct left 6–0 trials at the cue 1 epoch (black and green) and one cluster at the cue 2 epoch (purple). Cells were sorted according to their activity in the black cluster. (**k**) Overlap fraction of active neurons (z-score > 1.5) within the same cluster (intra-cluster) and across clusters (inter-cluster), defined as number of neurons active in both clusters divided by number of neurons active in either cluster. Small inter-cluster overlap suggests that a largely distinct group of neurons was active in each cluster. Overlap fractions were calculated separately for correct left 6–0 and right 0–6 trials. Shuffled overlap index was calculated by shuffling the assignment of trials to clusters. $^{***}P < 0.001$, two-sample Student's t-test. Gray points are individual cluster pairs and black points are means within each group. (**l**) Distribution of pairwise trial–trial population activity pattern correlations for pairs of trials with identical cues and choices at the turn epoch for all 6–0 trials (black) and only trial pairs in the same cluster (gray). (**m**) Fraction of clusters visited at a given epoch when clustering was performed on all epochs together (Online Methods). All trial types were included. Individual lines represent data sets ($n = 11$).

The trial–trial variability could have resulted from modulations of the tonic firing of a specific set of neurons or from major changes in which sets of neurons were active in each trial. We found evidence for the second possibility. We calculated the similarity between the groups of neurons that were active in pairs of clusters explored on a single trial type (for example, correct 6–0 left trials; **Fig. 3j**). Specifically, for each pair of clusters in a given epoch, we quantified the fraction of neurons that were active in both clusters using a threshold in z-scored estimated spike counts (threshold = 1.5). Surprisingly, only ~10% of neurons on average were active in both clusters in a pair, even when limiting our analysis to trials with identical choices and evidence cues (**Fig. 3j**,**k**). Many trials of the same type therefore had largely non-overlapping populations of active neurons. Consistently, the correlation coefficient between the population activity patterns for pairs of trials of the same type at the same epoch had a wide distribution, with some trial pairs being highly correlated and others having correlation coefficients near zero (**Fig. 3l**). In addition, we quantified the variability as a function of time in the trial using the cluster space defined by clustering activity patterns from all epochs together, rather than clustering independently within each epoch (Online Methods). The variability was estimated at a given epoch as the fraction of clusters explored by a population of trials. Surprisingly, when considering all trial types together, the fraction of clusters visited did not decrease over the course of the trial (**Fig. 3m**). The activity therefore

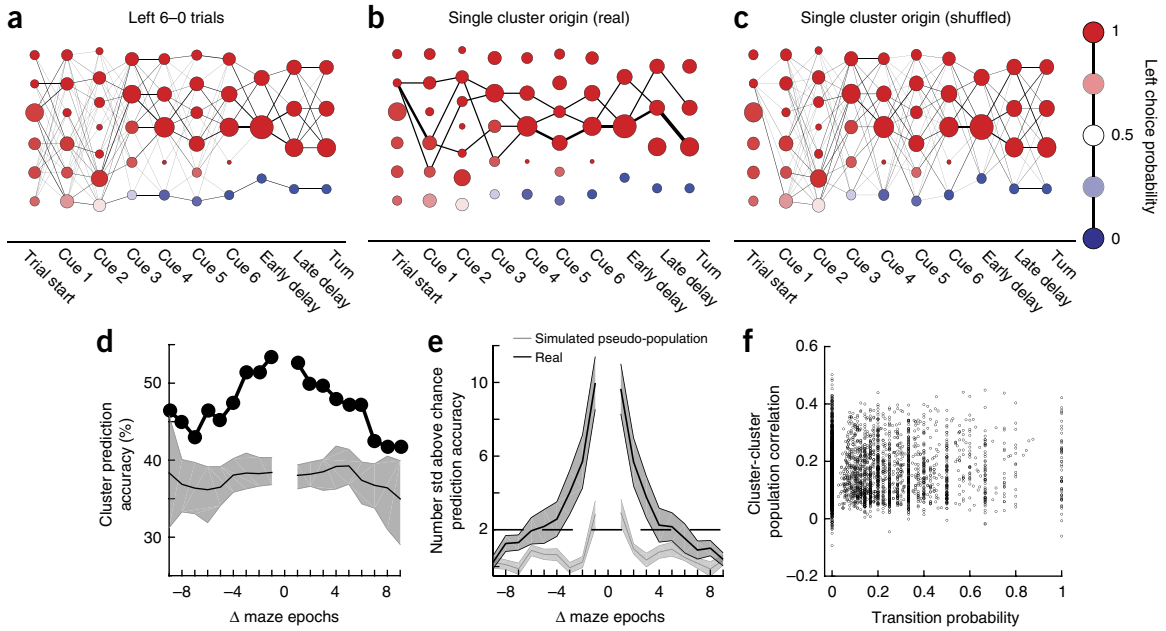

Figure 4 Long-timescale temporal structure in PPC activity. (**a**) Example transition matrix constructed only from left 6–0 (both correct and error) trials in a single data set. The nodes at the trial start had high left choice probabilities because the only trials included in this analysis were left 6–0 trials, which almost always resulted in a left choice. (**b**,**c**) Transition probabilities of all trials starting from a single cluster for real (**b**) and shuffled (**c**) data. In shuffled data, the assignment of trials to clusters was randomized, maintaining the distribution of trials across clusters. (**d**) Based on the cluster identity for a trial at a given epoch, the accuracy of correctly predicting that trial's past and future cluster occupancies (Online Methods). Predictability across many epochs suggests long-timescale temporal structure in single-trial activity trajectories. Accuracies were pooled across left 6–0 and right 0–6 trials that were clustered and considered separately. Line and shading represent the median and 99% confidence intervals from data in which the assignment of trials to clusters was shuffled. (**e**) Same as in **d**, but averaged across all data sets ($n = 11$). To combine across data sets with different chance classifier performance, accuracies were converted into the number of s.d. above the shuffled distribution. To create the pseudo-population, trial identities were shuffled (within a trial-type category) independently for each neuron to break neuron–neuron correlation structure but preserve each neuron's activity within the trial (Online Methods). Lines and shading represent mean ± s.e.m. across data sets. (**f**) Relationship between the population activity correlation of clusters in adjacent epochs and the transition probability between them. Transitions were no more likely between clusters with more similar population activity patterns. $r = 0.02$, $P = 0.21$, $n = 3{,}967$ consecutive cluster pairs.

maintained a high number of distinguishable activity patterns throughout the trial and did not collapse to a low-variability representation even at the turn epoch, after a choice had been made.

Population activity trajectories as orderly, seconds-long sequences of transitions between transient activity patterns

Given that a stereotyped sequence of activity patterns was not present for trials with identical cues and choices, we sought to understand the cause of the trial–trial variations. We generated a new cluster space using only trials of a single type (for example, left 6–0 trials) to remove the variability due to different evidence cues and choices (**Fig. 4a**). The variability in activity trajectories in this case could be due predominantly to biological or measurement noise. If so, the transitions from one activity pattern to the next are expected to be unpredictable, such that each single trial wanders through a random sequence of activity patterns. Alternatively, the variability between trials of the same type could carry information. In this case, each trial is expected to traverse an orderly set of activity patterns, such that the transition from one activity pattern to the next is predictable. We tested whether we could predict the future activity patterns of a trial on the basis of the trial's current activity pattern. As a first test, we visualized the paths of trials starting from a single cluster and found that only a subset of subsequent clusters was visited by those trials, even many epochs later (**Fig. 4b**). This example suggests that by knowing the trial's starting point, we could predict, to some extent, the clusters visited by that trial in the future. To visualize whether this structure could occur by chance, we simulated a 'noise' case by shuffling the assignment of trials to clusters at each epoch (maintaining the distribution of trials across clusters), thus creating transitions between clusters that mimic noise-driven transitions. In the shuffled case, the trials starting in a single cluster visited all subsequent clusters, in contrast to what we observed in the unshuffled data (**Fig. 4c**).

This example suggested that the transitions between activity patterns could be nonrandom and that temporal structure might exist in the variable paths traversed by single trials of the same type. We quantified this structure by developing a classifier in cluster space that asked whether, on the basis of the identity of the cluster occupied by a given trial at the current epoch, we could predict the identities of the clusters occupied by that same trial in past and future epochs (Online Methods). This analysis therefore tests whether the current activity pattern contains information about past and future activity patterns within a single trial. For trials with identical choices and evidence cues, the classifier predicted significantly above chance which cluster a trial occupied 5 or 6 epochs (~4–5 s) into the past and future ($P < 0.001$; **Fig. 4d**,**e**). Extensive analyses revealed that the temporal structure could not be explained by trial–trial differences in behavioral patterns, such as running patterns in the maze, and was not imposed by the clustering process (**Supplementary Fig. 8a–c** and Online Methods, "Contribution of behavioral variability to neuronal activity results"). Together, these results indicate that the current activity pattern contained information about past activity patterns and influenced the transition probabilities to future activity

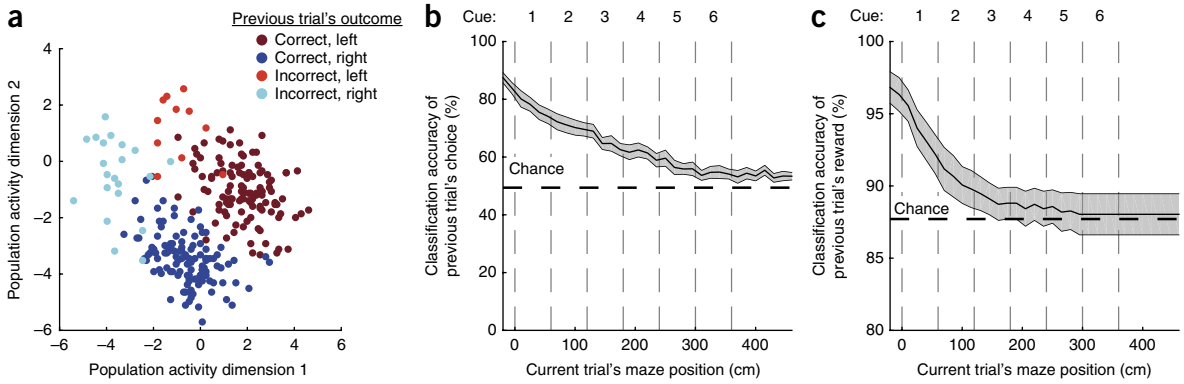

Figure 5 Neuronal population activity in the current trial reflects the previous trial's choice and outcome. (**a**) Population activity patterns on different trials at the trial start epoch colored by the choice and outcome (reward or no reward) of the previous trial. Dimensionality was reduced using factor analysis for visualization purposes. Each circle represents one trial. (**b**,**c**) Both the previous trial's choice (**b**) and whether the previous trial was rewarded (**c**) could be classified on the basis of the population activity. Independent SVM classifiers were trained and tested at each maze position. Lines and shading represent mean ± s.e.m. across data sets ($n = 11$).

patterns, even when removing the effects of different trial events such as evidence cues and choice.

The long-timescale temporal structure we observed could arise from persistent activity patterns, in which single neurons have long-lasting activity across epochs. Alternatively, there may exist predictable progressions between time-varying activity patterns, such that the PPC has long-timescale dynamics via orderly transitions from one short-lived population activity pattern to another. Several features of the data provided support for the second alternative. We found that neurons were transiently active with time-varying activity (**Fig. 2a**,**b** and **Supplementary Figs. 4** and **5a**,**b**). Also, clusters from different epochs had mostly distinct activity patterns (**Fig. 3h**,**j**). Furthermore, transitions were just as likely between clusters with similar activity patterns as they were between clusters with dissimilar activity patterns (**Fig. 4f**). To further test whether the long-timescale structure emerged from long-lasting activity in individual neurons, we shuffled the trial identities separately for each neuron among trials of the same type to disrupt neuron–neuron correlation structure while preserving activity patterns in individual neurons (simulating a pseudo-population). The removal of neuron–neuron correlations eliminated our ability to predict the past and future clusters visited by a single trial on the basis of the current cluster occupied by that trial (**Fig. 4e**). Together these results indicate that the temporal structure in single trials did not arise from long-lasting activity in individual cells; rather, orderly transitions occurred between transient patterns of neuronal activity with largely different sets of active neurons.

Long-lasting changes in population dynamics due to previous task events

Thus far, our results indicate that long-timescale structure exists in the PPC over seconds: the activity pattern at a given moment contained information about past activity patterns and also influenced the transition probabilities to future activity patterns. These results make important predictions about the timescale over which information about transient events is maintained in the PPC. An event during a trial is expected to result in a new population activity pattern that depends on both the features of the event and the activity pattern transition probabilities immediately before the event. The activity pattern that results is then expected to influence the transition probabilities to future activity patterns. Therefore, by helping to create a population activity pattern, a transient event is expected to have a long-lasting effect by constraining the possible future activity patterns, which in effect forms a short-term memory. We therefore hypothesized that transient events should have signatures of their occurrence long after they end. In this case, variability between trials of the same type could have emerged as a consequence of differences in recent past events.

To test this hypothesis, we asked whether the variability in activity patterns at the beginning of a trial could be explained by two prominent past events: the previous trial's choice and the previous trial's reward outcome (correct or incorrect). Because we were not directly analyzing transitions between activity patterns, we performed our analyses on the population activity without clustering for simplicity (we obtained similar results with clustering). The population activity patterns at the start of a trial, following an inter-trial interval of at least 2 s, were highly different for trials that had different choices and reward outcomes in the previous trial[28,30–33]. We visualized this result with dimensionality reduction by factor analysis (**Fig. 5a** and **Supplementary Fig. 9c**) and quantified the result using an SVM classifier based on population activity (**Fig. 5b**,**c**). The previous trial's choice could be decoded above chance for as long as 10 s after the conclusion of the previous trial, including well into the current trial (**Fig. 5b**). This signal did not have an easily detectable behavioral effect because a linear model with interactions could not predict the mouse's choice on the current trial on the basis of the previous trial's choice and reward ($R^2 = 0.02 \pm 0.01$, mean ± s.e.m., $P = 0.37$; Online Methods)[34]. Also, the previous trial's choice could not be decoded from the current trial's behavioral data (for example, running patterns; **Supplementary Fig. 8d–g** and Online Methods). PPC activity therefore contained information about events from previous trials many seconds after they had ended. As a result, trials with identical cues and choices had highly variable activity patterns due to differences in past events.

Predictions and tests for population activity dynamics during evidence accumulation

Although our analyses have focused in large part on comparisons between trials of a single type, the features identified have direct implications for evidence accumulation. We have shown that activity patterns in the PPC partially define the possible future activity patterns over seconds (**Fig. 4**). Events that help to establish a new activity pattern will therefore influence the transition probabilities to future activity patterns, creating a short-term memory of the event, as we have shown for choices and reward outcomes across trials

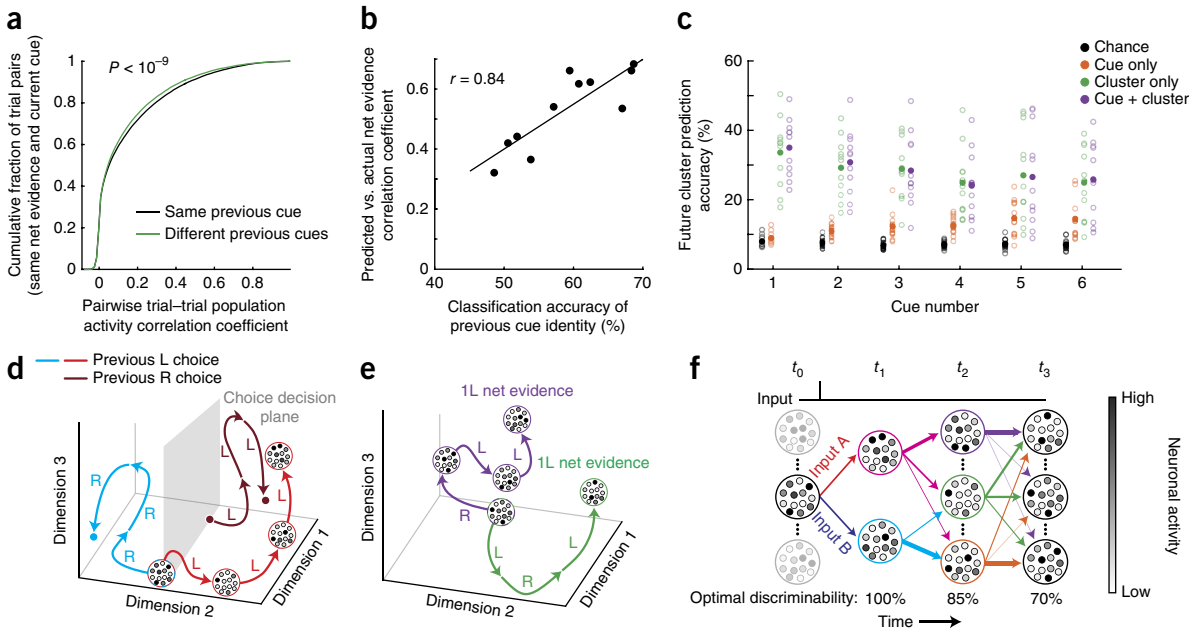

Figure 6 Analysis of neuronal activity related to evidence accumulation. (**a**) Cumulative distribution of the pairwise trial–trial population activity correlation coefficient for epochs with the same (black) or different (green) previous cues, keeping net evidence, current cue, and epoch constant (for example, LRLXXX vs. RLLXXX trials at cue 3). $P < 10^{-9}$, two-sample Kolmogorov–Smirnov test, $n = 11$ data sets. (**b**) For each data set, the ability to classify net evidence (correlation coefficient for predicted vs. actual net evidence using SVRs, as in **Fig. 2f**) was compared with the ability to classify the previous cue's identity (independent of maze epoch, current cue identity, and net evidence, as in **a**). Pearson's correlation $r = 0.84$, $P < 0.001$, $n = 11$ data sets. (**c**) For a single trial at a given epoch, the accuracy of predicting the next epoch's cluster identity for that trial on the basis of chance (black), the evidence cue type only (orange), the current cluster identity (green), or both (purple). Individual data sets (open circles) and means (filled circles) are shown. (**d**) Schematic illustrating that because the population activity depends on both the inputs and the near-past population activity, trials with the same sequence of cues but different starting points due to different past events will take different paths through activity space and result in distinguishable activity patterns (compare bright red and dark red trials). Each large circle surrounding small circles represents the population activity pattern at a given time point, with each small circle schematizing the activity of a neuron. Activity patterns are transient and change over the course of a trial (see bright red trajectory). Despite multiple activity patterns for choice, a decision plane (gray) can be drawn that separates activity patterns according to a given variable. Different decision planes can exist for other variables (for example, previous trial's choice). (**e**) Trials with the same starting point and net evidence but different sequences of cues will take different paths through activity space, resulting in distinguishable activity patterns. (**f**) Transient events have a long-lasting impact on population activity by creating a new activity pattern with different transition probabilities to future activity patterns. For example, if the network receives input B, the network transitions to the cyan activity pattern. Trials in the cyan activity pattern at t_1 are most likely to transition to the orange activity pattern at t_2, less likely to transition to the green activity pattern, and never transition to the purple activity pattern. The identity of the input can therefore be decoded at t_2 as a result of these non-uniform transition probabilities. The thickness of each arrow indicates the transition probability. The transitions between t_0 and t_1 indicate the change in activity due to one of two inputs. The circles follow the schematization from **d**. Because the transition probabilities are probabilistic, memory of the inputs decays as activity patterns diverge, leading to a decrease in the discriminability of inputs A and B.

(**Fig. 5**). In this framework, we can consider how evidence accumulation might occur. In response to an evidence cue, the network activity pattern would change on the basis of the cue type (left or right) and the set of activity pattern transition probabilities at the time of the cue. The response to a second cue would follow the same process and thus depend both on the second cue's type and on the transition probabilities for the activity pattern resulting from the first cue. Because the transition probabilities at the time of the second cue were set in part by the first cue, the activity pattern after the second cue would reflect both the first and second cues. Because this process cascades, each unique cue sequence would result in a unique activity pattern, even for the same net evidence. The activity pattern after all six cues would therefore be influenced by the sequence of previous cues. A single abstract variable for net evidence, in which the same final net evidence converges to the same activity pattern, regardless of the cue sequence, is therefore not expected to be present. Rather, the accumulated evidence cues would be represented generically as a sequence of inputs emerging from the long-timescale dynamics. Our results lead to predictions about population activity during evidence accumulation tasks.

A first prediction is that the population activity pattern should reflect not only the net evidence but also the sequence of evidence cues within a trial independent of net evidence. This prediction implies that different sequences of cues that result in the same net evidence (for example, left-right-left vs. right-left-left) should generate distinguishable activity patterns. To test this prediction, we selected trial epochs with the same current cue (for example, left) and the same net evidence (for example, +1 left) but with different cue types in the previous epoch, thus isolating effects due to the cue history. Trial epochs that had the same cue type in the previous epoch had significantly higher trial–trial population activity correlations than epochs with different cue types in the previous epoch ($P < 10^{-9}$, two-sample Kolmogorov–Smirnov test; **Fig. 6a**). Activity in an epoch could therefore be classified above chance levels on the basis of the type of cue in the previous epoch despite identical current cues and net evidences (classification accuracy: 59.0 ± 2.2%, mean ± s.e.m.; $P < 0.001$,

permutation test with shuffled trial labels; Online Methods). While this difference was highly significant, it was modest in amplitude, suggesting that it only accounted for a small fraction of the total trial–trial variability. The differences for distinct evidence sequences did not appear to reflect different internal accumulated evidence values due to unequal weighting of early and late cues (Online Methods, "Classifiers (without clustering): classification of cue sequences"). The population activity pattern therefore contained information about the sequence of past evidence cues, independent of net evidence.

Another prediction is that the signal for the sequence of past evidence cues (independent of net evidence) could underlie evidence accumulation in the population activity. Accumulated evidence would therefore be represented implicitly as a sequence of cues rather than explicitly as a single, abstract value such as net evidence. This prediction suggests that population activity with strong signals for cue histories should also have strong signals for evidence accumulation. Taking advantage of the variability across data sets, we found that our ability to decode the sequence of past cues, given the same net evidence, was strongly correlated with the decoding of net evidence ($r = 0.84$, $P < 0.001$; **Fig. 6b**). This result indicates that the cue-driven modifications to activity pattern transition probabilities leading to a cue history signal might also serve as a mechanism underlying evidence accumulation.

A final prediction is that if the current activity pattern influences the transition probabilities to future activity patterns, then both the current activity pattern and the type of evidence cue should influence the activity pattern following a new evidence cue. We compared trials with identical net evidence at the same epoch and asked whether we could predict the population activity pattern following a new evidence cue (either left or right cue) on the basis of (i) the distribution of trials across clusters alone (chance), (ii) the new cue type alone (cue only), (iii) the current activity cluster alone (cluster only), and (iv) both the current activity cluster and the new cue type (cue + cluster) (Online Methods). We performed this analysis in the cluster space to facilitate the analysis of transition probabilities between activity patterns. Based on the new cue's type, there was an increase in the ability to predict the identity of the next epoch's activity cluster, indicating that evidence cues triggered changes in population activity ($P < 0.001$ for cues 2–6; two-sample Student's t-test; **Fig. 6c**). However, the identity of the current activity cluster was more predictive of the next epoch's activity cluster than was the new cue's type ($P < 0.001$ for all cues; two-sample Student's t-test; **Fig. 6c**). Therefore, although new inputs influenced the future population activity pattern, the past population activity pattern had a larger effect, consistent with a role for the current activity pattern in defining the set of possible future activity patterns.

DISCUSSION

Our work identified two features of PPC activity that together motivate a new model for how evidence accumulation is performed in neuronal circuits. First, we found that each event during a trial, such as a new evidence cue or behavioral choice, modified the dynamics of the PPC over a timescale of seconds (**Figs. 4**,**5** and **6a–c**). Surprisingly, these events did not change the tonic activity of a specific set of neurons; rather, each event altered the set of activity patterns that the population could occupy in the future and thus the transition probabilities between complex population activity patterns, often involving transitions between different sets of active neurons (**Figs. 3h**,**j** and **4e**,**f**). This finding leads to a potentially generalizable rule in which transient inputs and activity patterns in the PPC 'reverberate' as long-lasting changes in the set of possible activity pattern transitions and trajectories, resulting in a short-term memory of each past input and activity pattern (**Fig. 6d–f**). This process was seemingly continuous in that the PPC activity pattern never appeared to reset, even after a trial was finished; rather, the PPC activity maintained an ongoing record of recent past events, thus forming a continuous, gap-free short-term memory. Our findings support and extend previous work that showed evidence accumulation signals in the PPC[3,4] by proposing that accumulation might occur generally by means of reverberation of network activity changes and by demonstrating that this accumulation could occur as long-timescale dynamics mediated by orderly transitions between transient and highly different activity patterns.

Second, we found that trials with identical evidence cues and choices were highly variable, such that these trials did not converge to a single, low-variance activity pattern, but were instead represented by widely varying patterns of population activity (**Fig. 3**). The diversity of activity patterns emerged because the PPC had information about many signals, including past events such as previous choices, reward outcomes, and evidence cues. Variability can therefore be considered, in part, as signals for unmeasured or hidden parameters, beyond those parameters directly tested in an experiment (for example, choice or net evidence)[19,20]. The presence of hidden signals affects our interpretation of neuronal activity in that it may be inaccurate to consider activity in layer 2/3 of PPC as specific for a set of measured task parameters and to think of the representation of those parameters as a small set of noisy network activity patterns. For example, the neuron–neuron activity correlation structure remaining after the subtraction of activity resulting from a selected subset of task variables, typically referred to as 'noise correlations', may reflect, in some cases, 'residual correlations' due to additional signals in the PPC. Together, our results therefore combine and put into a new context features identified in previous studies, including heterogeneous activity patterns across neurons[5,17,18,20], distributed representations of task stimuli including for irrelevant inputs[19,20], activity-dependent processing of stimuli[35,36], and the encoding of previous stimuli that indicates stimulus reverberation[30–33,37–43].

Our findings are inconsistent with key features of winner-take-all models[2,9,10]. First, traditional winner-take-all models predict that on different trials with the same choice the population activity converges to the same, low-variance pattern (attractor, which could potentially take multiple possible forms, such as a point in activity state space or a trajectory), which is predicted to erase the history of previous events. In contrast, we found that the same trial types (and choices) did not converge to a single pattern but instead consisted of highly different activity patterns (**Fig. 3**). In addition, throughout a trial, history signals were present for many events, including the sequence of previous cues and outcomes from previous trials (**Figs. 5** and **6a**). Second, published models propose that neurons have homogeneous and long-lasting activity patterns. Instead, consistent with our previous results[8], we found that neurons in the PPC were highly heterogeneous, with transient and time-varying activity (**Fig. 2a** and **Supplementary Figs. 4** and **5a**,**b**). Finally, most implementations of winner-take-all competitions involve mutual inhibition between competing pools of neurons that should result in negative population activity correlations between trials with different choices. Instead, we observed a correlation coefficient close to zero for such trial pairs ($r = -0.01 \pm 0.003$, mean ± s.e.m. across data sets). Although our results are inconsistent with current implementations of winner-take-all dynamics, they could be consistent with emerging models in which a winner-take-all circuit is embedded within a network with history-dependent dynamics[44] or in which activity in a

winner-take-all circuit is drawn toward, but never converges to, dynamically changing attractors.

We propose a potentially generalizable rule for PPC dynamics in which inputs that trigger a change in activity have a long-lasting effect on future activity patterns owing to the long-timescale dynamics of changes in transition probabilities (**Fig. 6d–f**). In the case of evidence accumulation, the evidence cues would not be privileged over other inputs; rather, evidence cues, like all other inputs, would help generate new activity patterns and thus new transition probabilities to future activity patterns. With multiple evidence cues offset in time, the changes in the transition probabilities would cascade such that the activity pattern following a sequence of cues would in part be defined by, and thus contain information about, the precise order of cues. Different sequences of cues would therefore result in unique activity patterns, as we have shown (**Fig. 6a**). As a result, the same net evidence, choice, and likely decision variable would not converge to the same activity pattern from trial to trial, but rather would form a diverse set of activity patterns. We predict that before learning of a task these activity patterns would not be associated with one another. Rather, through learning, the weights of connections onto a downstream readout network could be modified to establish a decision plane for choice or a manifold for net evidence. The readout network would therefore be able to associate the initially arbitrary sets of activity patterns with a task-specific meaning and behavioral output, as has been demonstrated in computational models[26,45,46] (**Fig. 6d**). The low-dimensional projection in the readout network could be consistent with previous recordings of ramping activity during evidence accumulation tasks[3,4,11,47]. Our model argues that the PPC has the general role of a reverberator of its inputs owing to intrinsic long-timescale dynamics and that evidence accumulation can be considered as a sequence of cues that establish a cue-sequence-dependent activity pattern. Evidence accumulation would thus occur as a specific example of a general dynamics feature. This new model is consistent with the theoretical framework developed in reservoir computing[45,48,49].

Our proposed algorithm offers advantages over a winner-take-all competition. In a winner-take-all competition, evidence accumulation would occur through an explicit, abstract signal for accumulated evidence. Such a signal is typically implemented in a highly specialized network architecture that is fine-tuned for a specific type of input, such as visual cues during virtual navigation in our case[9,10]. In contrast, our proposed model would allow the same network to flexibly scale for decision-making with multiple alternatives and to perform computations relevant to many diverse and novel tasks. This flexibility could be achieved through plasticity in readout weights, rather than through the construction of a new circuit architecture for each task[19,46,48,50]. We consider this advantage important for the PPC, which contains many signals in the same population of neurons and thus likely contributes to many learned behaviors in parallel.

METHODS

Methods, including statements of data availability and any associated accession codes and references, are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank S. Chettih and M. Minderer for developing the cell selection software; M. Andermann, J. Assad, W. Maass, O. Mazor, S. Panzeri, and A. Trott for discussions; and B. Datta, D. Dombeck, J. Drugowitsch, C. Gu, and members of the Harvey laboratory for comments on the manuscript. We also thank the Research Instrumentation Core at Harvard Medical School. This work was supported by a Burroughs-Wellcome Fund Career Award at the Scientific Interface, the Searle Scholars Program, the New York Stem Cell Foundation, the Alfred P. Sloan Research Foundation, a NARSAD Brain and Behavior Research Young Investigator Award, NIH grants from the NIMH BRAINS program (R01MH107620) and NINDS (R01NS089521), and a Stuart H.Q. & Victoria Quan Fellowship (A.S.M.). C.D.H. is a New York Stem Cell Foundation Robertson Neuroscience Investigator. Portions of this research were conducted on the Orchestra High Performance Compute Cluster at Harvard Medical School (supported by grant NCRR 1S10RR028832-01).

AUTHOR CONTRIBUTIONS

A.S.M. and C.D.H. conceived of the project, designed the experiments and analyses, and wrote the paper. A.S.M. collected and analyzed the data.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Gold, J.I. & Shadlen, M.N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
2. Wang, X.-J. Neural dynamics and circuit mechanisms of decision-making. *Curr. Opin. Neurobiol.* **22**, 1039–1046 (2012).
3. Shadlen, M.N. & Newsome, W.T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).
4. Hanks, T.D. *et al.* Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**, 220–223 (2015).
5. Raposo, D., Kaufman, M.T. & Churchland, A.K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).
6. Licata, A.M. *et al.* Posterior parietal cortex guides visual decisions in rats. Preprint at *bioRxiv* http://dx.doi.org/10.1101/066639 (2016).
7. Goard, M.J., Pho, G.N., Woodson, J. & Sur, M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. *Elife* **5**, 471 (2016).
8. Harvey, C.D., Coen, P. & Tank, D.W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484**, 62–68 (2012).
9. Wong, K.-F. & Wang, X.-J. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
10. Machens, C.K., Romo, R. & Brody, C.D. Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science* **307**, 1121–1124 (2005).
11. Horwitz, G.D. & Newsome, W.T. Separate signals for target selection and movement specification in the superior colliculus. *Science* **284**, 1158–1161 (1999).
12. Fujisawa, S., Amarasingham, A., Harrison, M.T. & Buzsáki, G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11**, 823–833 (2008).
13. Baeg, E.H. *et al.* Dynamics of population code for working memory in the prefrontal cortex. *Neuron* **40**, 177–188 (2003).
14. Crowe, D.A., Averbeck, B.B. & Chafee, M.V. Rapid sequences of population activity patterns dynamically encode task-critical spatial information in parietal cortex. *J. Neurosci.* **30**, 11640–11653 (2010).
15. Rajan, K., Harvey, C.D. & Tank, D.W. Recurrent network models of sequence generation and memory. *Neuron* **90**, 128–142 (2016).
16. Vogelstein, J.T. *et al.* Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.* **104**, 3691–3704 (2010).
17. Meister, M.L.R., Hennig, J.A. & Huk, A.C. Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J. Neurosci.* **33**, 2254–2267 (2013).
18. Jun, J.K. *et al.* Heterogenous population coding of a short-term memory and decision task. *J. Neurosci.* **30**, 916–929 (2010).
19. Mante, V., Sussillo, D., Shenoy, K.V. & Newsome, W.T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78–84 (2013).
20. Rigotti, M. *et al.* The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
21. Maimon, G. & Assad, J.A. Beyond Poisson: increased spike-time regularity across primate parietal cortex. *Neuron* **62**, 426–440 (2009).
22. Churchland, M.M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* **13**, 369–378 (2010).
23. Frey, B.J. & Dueck, D. Clustering by passing messages between data points. *Science* **315**, 972–976 (2007).
24. Churchland, M.M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
25. Mazor, O. & Laurent, G. Transient dynamics versus fixed points in odor representations by locust antennal lobe projection neurons. *Neuron* **48**, 661–673 (2005).
26. Briggman, K.L., Abarbanel, H.D. & Kristan, W.B. Jr. Optical imaging of neuronal populations during decision-making. *Science* **307**, 896–901 (2005).
27. Renart, A. & Machens, C.K. Variability in neural activity and behavior. *Curr. Opin. Neurobiol.* **25**, 211–220 (2014).

28. Marcos, E. *et al.* Neural variability in premotor cortex is modulated by trial history and predicts behavioral performance. *Neuron* **78**, 249–255 (2013).
29. Churchland, A.K. *et al.* Variance as a signature of neural computations during decision making. *Neuron* **69**, 818–831 (2011).
30. Bernacchia, A., Seo, H., Lee, D. & Wang, X.-J. A reservoir of time constants for memory traces in cortical neurons. *Nat. Neurosci.* **14**, 366–372 (2011).
31. Donahue, C.H. & Lee, D. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nat. Neurosci.* **18**, 295–301 (2015).
32. Seo, H., Barraclough, D.J. & Lee, D. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cereb. Cortex* **17** (Suppl. 1), i110–i117 (2007).
33. Seo, H. & Lee, D. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *J. Neurosci.* **27**, 8366–8377 (2007).
34. Busse, L. *et al.* The detection of visual contrast in the behaving mouse. *J. Neurosci.* **31**, 11351–11361 (2011).
35. Safaai, H., Neves, R., Eschenko, O., Logothetis, N.K. & Panzeri, S. Modeling the effect of locus coeruleus firing on cortical state dynamics and single-trial sensory processing. *Proc. Natl. Acad. Sci. USA* **112**, 12834–12839 (2015).
36. Curto, C., Sakata, S., Marguet, S., Itskov, V. & Harris, K.D. A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *J. Neurosci.* **29**, 10600–10612 (2009).
37. Nikolić, D., Husler, S., Singer, W. & Maass, W. Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biol.* **7**, e1000260 (2009).
38. Klampfl, S., David, S.V., Yin, P., Shamma, S.A. & Maass, W. A quantitative analysis of information about past and present stimuli encoded by spikes of A1 neurons. *J. Neurophysiol.* **108**, 1366–1380 (2012).
39. Seo, H., Barraclough, D.J. & Lee, D. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *J. Neurosci.* **29**, 7278–7289 (2009).
40. Sugrue, L.P., Corrado, G.S. & Newsome, W.T. Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787 (2004).
41. Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H. & Wang, X.-J. A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex. *Neuron* **88**, 419–431 (2015).
42. Murray, J.D. *et al.* A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).
43. Yang, Y. & Zador, A.M. Differences in sensitivity to neural timing among cortical areas. *J. Neurosci.* **32**, 15142–15147 (2012).
44. Klampfl, S. & Maass, W. Emergence of dynamic memory traces in cortical microcircuit models through STDP. *J. Neurosci.* **33**, 11515–11529 (2013).
45. Buonomano, D.V. & Maass, W. State-dependent computations: spatiotemporal processing in cortical networks. *Nat. Rev. Neurosci.* **10**, 113–125 (2009).
46. Hoerzer, G.M., Legenstein, R. & Maass, W. Emergence of complex computational structures from chaotic neural networks through reward-modulated Hebbian learning. *Cereb. Cortex* **24**, 677–690 (2014).
47. Murakami, M., Vicente, M.I., Costa, G.M. & Mainen, Z.F. Neural antecedents of self-initiated actions in secondary motor cortex. *Nat. Neurosci.* **17**, 1574–1582 (2014).
48. Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
49. Maass, W., Natschläger, T. & Markram, H. Real-time computing without stable states: a new framework for neural computation based on perturbations. *Neural Comput.* **14**, 2531–2560 (2002).
50. Sussillo, D. & Abbott, L.F. Generating coherent patterns of activity from chaotic neural networks. *Neuron* **63**, 544–557 (2009).

ONLINE METHODS

Subjects. Experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee. Data were acquired from five male C57BL/6J mice (Jackson Labs), age 8–10 weeks at the start of behavioral training and 14–22 weeks during imaging. Mice were housed as pairs in cages in a room with a reverse light/dark cycle. Mice had no previous history of any other experiments. A titanium headplate was affixed to the mouse's skull using dental cement (Metabond, Parkell). Mice were placed on a water schedule in which they received 800 µl of water each day. Each mouse's weight was measured daily to ensure that it was ≥80% of the mouse's pre-water-restriction weight.

Behavior. *Virtual reality system.* We used a modified version of the virtual reality (VR) system described previously[51]. Images were back-projected onto a half-cylindrical screen (24-inch diameter) using a PicoP microprojector (MicroVision). Spherical treadmill movement was recorded using an optical sensor positioned beneath the ball. Forward/backward translation in VR was controlled by changes in pitch, and rotation in VR was controlled by changes in roll (both relative to the mouse's body). The recorded behavioral parameters were the mouse's position and view angle in the virtual environment along with the rotational velocity of the spherical treadmill. Virtual environments were made using ViRMEn[52].

Task description (**Fig. 1** *and* **Supplementary Fig. 1**). While running down the stem of the T with predominantly gray walls, mice encountered six visual cues (white wall segments with black dots) at fixed locations. Each cue could appear on either the left or right wall, and only one cue was visible at a time. Cue visibility was determined by the mouse's position such that the duration of each cue was determined by the mouse's running speed. On average, each cue was visible for ~0.8 s. To receive a reward, mice had to determine whether more cues were presented on the left or the right and, after a short stretch of maze without additional cues (90 cm, ~1 s), turn at the T-intersection toward the direction that had more cues. Task difficulty was modulated by varying the difference between the number of left and right cues (net evidence; 6 total cues per trial, ranging from 0 to 6 presented on the left and the remainder on the right). The sequence of cues was determined randomly for each trial of a given net evidence. On 3–3 trials, the rewarded location was selected randomly. Following the completion of the trial, the screen changed to black for the duration of the inter-trial interval (2 s for correct choice and 4 s for incorrect choice).

Behavioral training procedure (**Supplementary Fig. 1**). Behavioral sessions lasted 45–60 min. Mice received liquid rewards through a lick spout (4 µl/reward, 10% sweetened condensed milk). Mice were trained to perform the evidence accumulation task over a series of eight mazes (**Supplementary Fig. 1**). We implemented bias correction throughout training. Bias correction was not used during imaging sessions. On each trial, we determined a probability that a trial would be a right choice trial as the fraction of left choices over the previous 20 trials. To maintain a high level of performance throughout the session, we introduced a small fraction of easy trials ('crutch trials') interleaved with the evidence accumulation trials. Crutch trials were identical to trials from maze 5 (**Supplementary Fig. 1**) in which no evidence accumulation or delay were present. The probability of a crutch trial on a given trial was equal to the fraction of error trials over the previous 20 trials. Crutch trials were used during imaging sessions and were excluded from all analyses.

Behavioral analyses. Behavioral performance was calculated as the fraction of trials in which the mouse performed the correct choice, excluding crutch trials.

Behavioral analysis of evidence accumulation (**Supplementary Fig. 2**). To test whether mice used more than one piece of evidence per trial, we fit the behavioral performance as a function of number of left cues with a logistic function (assuming more than one piece of evidence used per trial) and a linear function (assuming a single piece of evidence used per trial). To compare model fits, for each mouse, each behavioral day was fit separately by each model, and the distribution of root mean squared errors was compared with a two-sample Student's *t*-test. Across mice, the logistic function fit the data significantly better than the linear function (**Supplementary Fig. 2b**,**c**), indicating that mice used more than one piece of evidence per trial.

To test which cues mice used across trials, we used multivariate linear regression, with the behavioral choice as the response variable and the cue identities as the explanatory variables. To include large numbers of trials, multiple consecutive sessions (mean: 9, range: 7–12) were combined. All cues had significant regression coefficients with a preference toward earlier segments, suggesting that mice accumulated evidence with a primacy bias (**Supplementary Fig. 2d–f**). This unequal weighting of cues was not present in all mice.

Analysis of across-trial behavioral effects. To test whether there was a relationship between the mouse's choice on a trial and the outcome of the previous trial, we used a multivariate logistic regression with interactions with the previous trial's choice and reward as binary explanatory variables and the mouse's choice on the test trial as the response variable. We combined multiple consecutive sessions (mean: 9, range: 7–12) to include large numbers of trials. This model was unable to predict the mouse's choice, suggesting that there was no easily detectable behavioral relationship between the mouse's choice and the outcome of the previous trial (R^2: 0.02 ± 0.01, mean ± s.e.m. across data sets, $P = 0.37$).

Contribution of behavioral variability to neuronal activity results. We performed multiple analyses to test the possibility that our results were due to contributions from behavioral parameters such as changes in the visual scene or running patterns, rather than features such as evidence accumulation or variability in internally driven neuronal population activity. In all cases, we found that behavioral variability could not entirely explain the neuronal activity patterns we observed.

One possibility is that the mouse began to turn left or right as it saw evidence cues such that accumulation of evidence was performed through the mouse's viewing angle in the maze (for example, left of center viewing for more accumulated left cues) rather than through an internal representation of net evidence. In such a case, net evidence could be correlated with different heading directions (view angle in the maze), motor signals (turning on the treadmill), and direct visual input (combination of view angle and position in the maze). However, when we limited our analysis to only trials with similar view angles (±2.5 degrees), the SVR analysis based on population activity predicted above chance levels the actual net evidence (**Supplementary Fig. 5e**,**f**). Additionally, when we trained SVR models on behavioral parameters alone (view angle, maze position, two axes of treadmill rotational velocity in a single model) or on behavioral parameters in addition to neuronal population activity, models trained on both behavioral parameters and neuronal population activity consistently predicted the net evidence better than those trained on behavioral parameters alone, despite modest predictability from behavioral parameters alone (comparison of models with different numbers of parameters was made possible by the use of non-overlapping training and testing sets; **Supplementary Fig. 5g**,**h**). These results suggest that a representation of net evidence was present independent of heading direction, running patterns, and direct visual input.

Another possibility is that trial–trial differences in behavioral parameters could have generated the structured trial–trial variability in neuronal activity and the presence of history signals across long timescales. We ruled out these possibilities using a series of tests to see whether neuronal activity explained additional variability beyond what could be explained by the behavioral variability, by building both neuronal activity and behavioral features into a single logistic regression model. We found that the current behavioral parameters alone explained above chance, but poorly, past and future activity pattern clusters for only ~1 epoch into the past or future (**Supplementary Fig. 8a–c**). In contrast, models using the current behavioral parameters and the current activity pattern cluster (or the current activity pattern alone) predicted the past and future epochs substantially better, including across a longer timescale of 5 or 6 epochs (determined by adjusted R^2 to compare models with different numbers of parameters; **Supplementary Fig. 8a–c**). Consistently, using behavioral parameters for visual scene and running patterns, we were unable to classify above chance levels history signals from the previous trial in the subsequent trial (**Supplementary Fig. 8d–g**). Together these results indicate that the trial–trial variability and history signals we observed included neuronal signals that could not be explained by major behavioral variability.

Imaging. *Surgical procedure.* When mice performed well on maze 7 (**Supplementary Fig. 1a**), they underwent a surgery to implant a cranial window. For 3 d before surgery, mice were given 5 mL of water per day. A circular craniotomy (3.1 mm diameter) was made over left PPC (2 mm posterior, 1.75 mm lateral of bregma). A virus mixture containing a 4:1 volumetric ratio of tdTomato (AAV2/1-CAG-tdTomato) to GCaMP6 (AAV2/1-synapsin-1-GCaMP6f or AAV2/1-synapsin-1-GCaMP6m) was delivered by three injections of ~20 nL (~5 min/injection, ~150 µm spacing between injections). Viruses were obtained from the University of Pennsylvania Vector Core Facility. Injections were made near the center of the craniotomy, ~275 µm below the dura, using a beveled

glass pipette (~15 µm tip diameter) and a custom air pressure injection system. The pipette was advanced using a micromanipulator (Sutter MP285) at a 30-degree angle to minimize compression of the brain. A window with glass plug (5 mm diameter coverslip plus two 3 mm diameter coverslips; #1 thickness; CS-3R and CS-5R, Warner Instruments) was made using UV-curable, optically transparent adhesive (Norland Optics). The window was affixed to the brain using a drop of Kwik-Sil (World Precision Instruments) and affixed to the skull using Metabond mixed with ~5% vol/vol India ink to prevent light leakage. A headplate was affixed to the skull using Metabond mixed with India ink. A titanium ring was mounted on top of the headplate to interface with a cylinder of black rubber to surround the microscope's objective lens, thus preventing light leak from the VR display into the microscope[53]. Imaging began at least 4 weeks after injection and was continued for up to 12 weeks. Fields of view containing cells with GCaMP6 in the nucleus were excluded. In a given session, we imaged ~350 neurons simultaneously during ~300 trials (range, 188–648 neurons; range, 231–414 trials; $n = 5$ mice; **Supplementary Table 1**).

Two-photon microscope design. Imaging was performed using a custom-built two-photon microscope. The microscope scan head included a resonant scanning mirror and a galvanometric mirror separated by a scan lens–based relay telescope. Fluorescence light collection optics were based on a custom design to collect wide dispersion angles from objectives with a large (~20 mm) back aperture. The microscope was stationary, and the mouse was mounted on an *xyz* translation stage (Dover Motion). Green and red emission light were separated by a dichroic mirror (580 nm long-pass, Semrock) and bandpass filters (525/50 and 641/75 nm, Semrock) and collected by GaAsP photomultiplier tubes (Hamamatsu). Excitation light was delivered from a Ti:sapphire laser (Coherent) operated at 920 nm. The microscope was controlled by ScanImage (version 5; Vidrio Technologies)[54].

Imaging data acquisition. Imaging data were acquired at ~30 Hz at a resolution of 512 × 512 pixels (~700 µm × ~700 µm field of view) using a Nikon 16×, 0.8 NA objective lens. Imaging and behavioral data were synchronized using custom-written MATLAB software by simultaneously recording the frame clock from ScanImage and an iteration counter from ViRMEn (Online Methods). Up to 100,000 frames were acquired from each imaging session over the course of ~1 h. Imaging data were acquired at depths between 100 and 200 µm below the dura. Data were analyzed from 11 fields of view from 5 mice.

Preprocessing of imaging data. Motion correction, the definition of putative cell bodies, and extraction of fluorescence traces ($\Delta F/F$) were performed in a semi-automatic fashion using custom-written software. In brief, following motion correction[55], the correlation of fluorescence time series was calculated for each pair of pixels within ~60 µm of one another. Fluorescence sources (putative cells) were then identified by applying a continuous-valued, eigenvector-based approximation of the normalized cuts objective[56] to the correlation matrix, followed by discrete segmentation by *k*-means clustering, yielding binary masks for all identifiable fluorescence sources. For each putative cell, the local neuropil fluorescence was estimated by averaging across nearby pixels devoid of fluorescence sources. The scale of neuropil contamination of the cell fluorescence was estimated by regressing the background time series against low-activity regions of the cell time series, and the scaled background time series was then subtracted from the cell time series. Cell selection and neuropil subtraction were performed using a tool that allowed manual examination of clustering results and parameters, in combination with anatomical information and fluorescence traces corresponding to each cluster. All neuropil contamination fits were also examined by eye and adjusted when necessary. All fluorescence traces were deconvolved to estimate the probability of a spike in each frame (estimated spike count)[16]. Similar results were obtained from the non-deconvolved $\Delta F/F$ traces (**Supplementary Fig. 10**).

Data analysis. *General analysis procedures.* Data were grouped into spatial bins (3.75 cm/bin). The T-maze was linearized before binning by folding the arms such that they were a continuation of the stem. Neuronal activity and behavioral parameters were averaged in each bin (2 or 3 imaging frames per bin per trial). Unless otherwise noted, all analyses were performed on both correct and error trials together. All correlation coefficients were from Pearson's correlations.

Classifiers (without clustering): general procedures. Unless otherwise noted, all population classifiers were support vector machines (SVMs) with a radial basis function (Gaussian) kernel[57,58] implemented using the libsvm library (version 3.20)[59]. All population classification was performed on the concatenated activity of all individual neurons. Data were divided into non-overlapping training/validation and test sets (50% of trials each). To prevent overfitting, models were trained exclusively on the training/validation set, with the test set left untouched until final testing. Hyperparameter (C and γ; regularization weight and radial basis function width, respectively) selection was performed using a random search method with tenfold cross-validation on the training/validation set of only a single data set, and the same hyperparameters were used for all data sets.

Classifiers (without clustering): two-class classifications (**Fig. 2c**,**e**,**g** *and* **Supplementary Figs. 5c**,**j**,**8d–g** and **10a**,**d**,**e**). For two-class classification problems, such as the classification of the choice (**Fig. 2c**,**e**,**g** and **Supplementary Fig. 5c**), previous trial's choice (**Fig. 5b**), and previous trial's reward outcome (**Fig. 5c**), a *C*-support vector classification (*C*-SVC) approach was used. For the single neuron choice classifier, neuronal activity was averaged across each left 6–0 and right 0–6 trial (**Fig. 2c**). For population classifiers, independent SVMs were trained on each spatial bin (**Fig. 2e**,**g** and **Supplementary Figs. 5c**,**9d–g** and **10a**,**d**,**e**). In cases where classification accuracy based on neuronal data was compared with accuracy based on behavioral data (**Supplementary Fig. 8d–g**), hyperparameters were optimized separately for each. In cases where the weights placed on individual neurons were analyzed (**Supplementary Fig. 5j**), a linear kernel was used to allow for interpretability of the weights.

Classifiers (without clustering): net evidence support vector regression (**Fig. 2d**,**f**,**h** *and* **Supplementary Figs. 5e–h** and **10b**). For the prediction of net evidence on the basis of neuronal population activity (**Fig. 2d**,**f**,**h** and **Supplementary Fig. 5e–h**), an ε-support vector regression (ε-SVR) approach was used[57,58,60]. For the net evidence models, the average activity during the third quarter of each cue's presentation (~200 ms) was calculated for each neuron. For training and testing, each cue was treated as a separate trial with class labels corresponding to the net evidence including that cue. Trials were divided into training and testing sets as whole trials to prevent similar activity on different cues within the same trial from corrupting the results. To determine prediction accuracy, the predicted net evidence was compared to the actual net evidence via a correlation coefficient. To calculate statistical significance, the results were compared to the distribution resulting from 1,000 shuffles of class labels. To rule out categorization, which would result in identical guesses within left and right net evidence conditions but a positive slope across all net evidence conditions, we calculated significance separately within left and right net evidence conditions; both were statistically significant across mice for the population classifiers ($P < 0.001$; **Fig. 2f**).

Classifiers (without clustering): classifiers built by adding-in subsets of neurons (**Fig. 2g**,**h** *and* **Supplementary Fig. 5i**). Classifiers were built similarly to the population classifiers described above, except with the input being a subset of neurons. The single neuron choice SVM and SVR net evidence correlation were used to determine single neuron selectivity for choice and net evidence, respectively. Neurons were sorted in ascending order on the basis of their selectivity. A separate classifier was trained for increasingly larger populations of neurons with neurons added in from least to most selective.

The classifier's accuracy increased as neurons with low individual classification accuracy were incorporated (**Fig. 2g**,**h**). Using the least selective 40% of neurons, the classifier for choice reached an overall accuracy of ~75%, even though the most selective neuron included only reached an individual accuracy of 52.5%. While this result may appear counterintuitive at first, there are several ways in which it can occur. For example, consider a neuron that is active on only 10 out of 200 trials. However, all 10 of the trials during which the neuron is active result in a left choice. The neuron's individual classification accuracy will necessarily be low: the best it can achieve is 52.5% (chance performance on the 190 trials during which it is silent and perfect performance on the 10 trials during which it is active). However, a population classifier combining many such neurons, each active on a different subset of 10 trials, may still achieve a high accuracy overall.

As another example, consider a neuron that is active on all trials with 10% higher activity on left trials than on right trials. In a noiseless environment, such a neuron would result in perfect classification accuracy. However, as noise increases, the neuron's classification accuracy will quickly fall, and if the noise is much larger than the difference in average activity between left and right trials, the neuron will display poor individual accuracy. Nevertheless, if a population classifier averages the responses of many such neurons (assuming independent noise), its accuracy will increase steadily as more neurons are included (and as more noise is averaged out), eventually reaching perfect classification accuracy.

doi:10.1038/nn.4403

We tested whether our classifier's performance was due to weak information in single cells or correlations between neurons by shuffling the trial identities within a trial-type category separately for each neuron, thus preserving each neuron's activity but breaking inter-neuronal correlations (i.e., simulating a pseudo-population). In the shuffled case, the classifier performed with high accuracy, indicating that many cells beyond the highly selective ones contained small amounts of choice- and net evidence-related information (**Supplementary Fig. 5i**). Importantly, these analyses indicate that our classifier was not dependent on correlations but they do not rule out the possibility that population-level correlations may play a key role in information encoding.

*Classifiers (without clustering): classification of cue sequences (***Fig. 6a**,**b** *and* **Supplementary Figs. 5d** and **10g**,**h***).* To determine whether the identity of previous cues could be read out at a current epoch, given identical current cue and net evidence (**Fig. 6a**,**b** and **Supplementary Fig. 5d**), we examined pairwise trial–trial population activity correlations. At each of cues 3–6, we separated trials into those that contained the patterns LRL (left-right-left) and RLL (right-left-left) or LRR and RLR with the last cue in the pattern matching the currently analyzed cue. For example, at the third cue, the pattern LRLRRR would be a match, while at fifth cue, the pattern RRLRLR would be a match. To rule out predictability due to differences in net evidence, a subset of trials in which the distribution of net evidence was equivalent across the two groups was used. This procedure resulted in 8 groups (2 choices × 4 patterns). At each cue epoch, pairwise trial–trial correlations of the mean neuronal population activity vector for n simultaneously imaged neurons were calculated for trials with the same previous cue or different previous cues. To predict the previous cue on the basis of these correlations, in a leave-one-out fashion, the mean population activity correlations between the test trial and all other trials with a left or right previous cue were calculated. Accuracy was compared to the accuracies from 1,000 shuffles of the labels assigning previous cues to trials.

The activity difference due to different previous cues could reflect different internal accumulated evidence values owing to unequal weighting of early and late cues. However, similar results were obtained when we restricted our analysis to the fifth and sixth cues, which were weighted similarly behaviorally ($P = 0.03$), and when we considered data from a mouse that weighted all cues equally (**Supplementary Fig. 2d**, see mouse marked in red; for both cases, $P < 1.4 \times 10^{-4}$, comparison of pairwise activity correlations for trials with the same or different previous cue, two-sample Kolmogorov–Smirnov test; **Supplementary Fig. 5d**).

Visualization of activity in high-dimensional state space via factor analysis (**Fig. 5a** *and* **Supplementary Fig. 9***).* For visualizations using factor analysis (**Fig. 5a** and **Supplementary Fig. 8**), dimensionality was reduced to 5-factors and two were selected for visualization.

*Clustering methods (***Figs. 3**,**4** and **6c** *and* **Supplementary Figs. 6**,**7**,**9a**–**c** and **10c**,**f***).* We used clustering to reduce the dimensionality of the population activity without inclusion of information about behavioral parameters and without encouraging projection of the data onto dimensions that maximize variance due to specific task features, such as choice. We used clustering because it did not assume linearity in the data structure and facilitated analyses by discretizing activity patterns, allowing the calculation of transition probabilities between discrete activity states. However, clusters were not considered an indication of discreteness of the underlying activity patterns. Clustering provided the advantage of allowing analysis of the rules that govern the transitions between activity patterns from moment to moment within a trial. Although other analysis approaches may have been possible, standard methods have not to our knowledge been developed previously to study the rules governing transitions between transient and largely different patterns of population activity, as we observed here.

Preprocessing for clustering. Prior to clustering, each spatially binned trial was divided into ten non-overlapping epochs, corresponding to the start of the trial, cues 1–6, the early and late delay, and the turn. For the trial-start epoch, neuronal activity was averaged over the four spatial bins (15 cm) immediately preceding onset of the first cue. For cues 1–6, activity was averaged over the third quarter of each cue's presentation (four spatial bins). For the early and late delay, respectively, activity was averaged across the four spatial bins beginning 15 and 37.5 cm after offset of the final cue. For the turn, activity was averaged across the final four spatial bins in the maze. Each epoch corresponded to approximately 200 ms.

Clustering via affinity propagation. Within each epoch, trials were clustered into groups on the basis of their neuronal activity using affinity propagation (code from B. Frey, University of Toronto)[23]. Affinity propagation has two inputs: a distance matrix and a 'preference' for each data point. We calculated the distance as the negative sum of pairwise Euclidean distance and one minus the pairwise cosine similarity between every trial in the n-dimensional activity space (each dimension being the activity of a single neuron). In contrast to other clustering methodologies, such as k-means clustering, the preference parameter does not specify the number of clusters, but rather a general range. For example, in our experience, clustering on different data sets using the same preference parameter can result in anywhere from 1 to 30 clusters. To determine the preference parameter, we calculated the number of clusters generated across a range of preference parameter values. The tenth percentile of the difference matrix was within a stable range, such that small modifications in the preference parameter did not greatly influence the number of clusters identified. The choice of this parameter and the resulting number of clusters had little impact on our results (**Supplementary Fig. 6j**).

*Transition probabilities between clusters (***Figs. 3** and **4***).* To calculate transition probabilities between clusters (cluster a at epoch 1, cluster b at epoch 2), we calculated the fraction of trials in cluster a which were also in cluster b. To superimpose behavioral variables on clusters (**Figs. 3a**–**d** and **4a**–**c**), for each cluster, we calculated the fraction of trials that had a given feature. To validate that the clustering found meaningful groups, we analyzed whether the distribution of behavioral variables across clusters at a given epoch was significantly different than chance. To calculate chance, we shuffled the cluster labels such that the number of trials in each cluster was maintained, but with the trials assigned to a given cluster determined randomly (**Supplementary Fig. 6a**–**d**). Significance was established by summing the absolute difference in behavioral variables from the expected uniform distribution for the real data and comparing this total difference to the distribution of the same metric obtained from 1,000 shuffles (**Supplementary Fig. 6b**,**d**).

Clustering based on all time points together, rather than epoch-by-epoch clustering (**Fig. 3m** *and* **Supplementary Fig. 6f***).* To determine whether clustering separately at each epoch resulted in multiple clusters with similar population activity patterns in different epochs, we also performed clustering on all epochs together. While the total number of clusters for epoch-by-epoch-based clustering was larger than those for all-epoch-based clustering (ratio: 1.6 ± 0.06), the ratio was relatively low, suggesting that activity patterns were distinguishable across epochs. Self-transitions were identified as transitions across epochs in the all-epoch clustering in which a trial was in the same cluster at two consecutive epochs (**Supplementary Fig. 6f**). To determine whether the variability of trials in cluster space changed over the course of a trial, we also calculated the fraction of the total clusters explored at each maze epoch (**Fig. 3m**). To remove outliers, only clusters containing three or more trials at a given epoch were counted as visited.

*Classifiers based on activity in cluster space: classification of the cluster identity at past and future epochs on the basis of the cluster identity at the current epoch (***Figs. 4d**,**e** and **6c** *and* **Supplementary Figs. 6j** and **10c***).* To predict the past or future cluster identity during a trial at a certain epoch (epoch j), given the cluster identity at another epoch (epoch i), we implemented a classifier built in cluster space (**Fig. 4d**,**e**). In a leave-one-out fashion, we limited our analysis to only the trials that were in the same cluster as the test trial at epoch i. Of those trials, we then asked which cluster was most common at epoch j. The classifier then predicted that the test trial would also be in that cluster at epoch j. Prediction accuracy was calculated as the fraction of predicted clusters that matched the actual cluster. In **Figure 4d**,**e**, clustering was performed separately on correct left 6–0 and right 0–6 trials to rule out structured variability due to different cues or behavioral choices. Accuracy was compared to 1,000 shuffles of the assignment of trials to clusters. Cluster assignments were shuffled independently at each epoch. This shuffle maintains the distribution of trials across clusters. This classifier was identical to the "cluster only" classifier used for **Figure 6c**, except that in **Figure 6c** it was applied to all trials together, independently of evidence or choice. The other classifiers used in **Figure 6c** followed a similar logic. For the "chance" classifier, the same procedure was followed, except all trials except for the test trial were used to determine the most likely cluster at epoch j, not just those in the same cluster at epoch i. For the "cue only" classifier, only those trials with the same current cue (left or right) were used to determine the most likely cluster at epoch j. For the "cue + cluster" classifier, only those trials that both had the same current cue as the test trial and were in the same cluster at epoch i were used to determine the most likely cluster at epoch j.

Classifiers based on activity in cluster space: classification of past and future cluster identities with simulated pseudo-populations (**Fig. 4e**). The temporal structure we observed could be caused by the prolonged activity of individual neurons, either owing to persistent activity in the underlying neuronal activation or to the prolonged decay kinetics of the calcium indicator. To test whether prolonged activity patterns could account for temporally structured and predictable trial–trial variability, we shuffled the trial labels independently for each neuron across trials with the same choice and sequence of evidence cues (for example, correct left 6–0 trials; **Fig. 4e**). This shuffle therefore breaks neuron–neuron correlation structure but maintains the temporal structure of each neuron's individual activity, simulating a pseudo-population. Following this shuffle, clustering was performed (separately for correct 6–0 left and 0–6 right trials), and past and future activity patterns were classified as described above using the "population activity only" classifier. In contrast to the unshuffled case, we were unable to predict past and future activity patterns over more than one epoch (**Fig. 4e**), suggesting that the predictability we observed in the real data could not be explained by prolonged activity or slow indicator kinetics.

Classifiers based on activity in cluster space: classification of past and future cluster identities on the basis of behavioral data (**Supplementary Fig. 8a–c**). To determine the fraction of past and future predictability accounted for by variability in behavioral parameters across clusters, we used logistic regression to compare predictability based only on behavioral parameters to that based on both behavioral parameters and neuron activity-defined clusters (**Supplementary Fig. 8a–c**). To allow binary classification, only trials whose cluster identity contained either the most or second most trials during the prediction epoch were included. We used all recorded behavioral parameters (*x*/*y* position, spherical treadmill rotational velocities, view angle) for this analysis, which together account for the mouse's general running pattern and visual scene. At each epoch, we trained a logistic regression model to predict the activity pattern at the current epoch or at another epoch in the past or the future on the basis of either the behavioral variables alone (behavior only) or the behavioral variables in addition to the current cluster identity (behavior + neuronal activity clusters). Note that for the same epoch, we did not include neuronal clusters as an explanatory variable as they were identical to the response variable in that case. To compare model performance across the two cases, which have different numbers of predictors, we used adjusted R^2. Independent models were created for each combination of epochs and combined on the basis of the number of epochs separating the predictor and response (for example, cue 1 and cue 2 are separated by 1 epoch, while trial start and turn epochs were separated by 9 epochs). Separate models were calculated for left 6–0 and right 0–6 trials to rule out behavioral variability induced by the choice. Results were qualitatively similar when models included linear interaction terms and quadratic terms.

Analysis of the overlap of active neurons across clusters (**Fig. 3k**). To calculate the overlap fraction for active neurons between clusters, each neuron's activity across all trials was *z*-scored. Within each cluster, each neuron's mean *z*-scored activity was calculated and compared to a *z*-score activity threshold of 1.5 (**Fig. 3k**), though similar results were obtained using different thresholds. Neurons whose mean *z*-scored activity was above this threshold were determined to be active. Using correct left 6–0 and right 0–6 trials separately to rule out differences due to evidence cues, the pairwise overlap fraction between all clusters at each epoch was calculated as

$$\text{overlap index} = \frac{\text{number of neurons active in both clusters}}{\text{number of neurons active in either cluster}}$$

For the intra-cluster measure, trials within a cluster were randomly divided into two groups, the mean activity within each group was calculated for each neuron, this mean activity was compared to the *z*-score threshold, and the overlap fraction between the two groups was calculated. This process was repeated 100 times and the results were averaged. To reduce variability due to low trial numbers, only clusters (intra-cluster measure) or cluster pairs (inter-cluster measure) with greater than 20 combined trials were included. To determine the shuffled overlap fraction, the cell labels within each cluster were randomly assigned 1,000 times, and the inter-cluster overlap fraction was recalculated.

Data and code availability. The data and code that support the findings of this study are available from the corresponding author upon reasonable request.

A **Supplementary Methods Checklist** is available.

51. Harvey, C.D., Collman, F., Dombeck, D.A. & Tank, D.W. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* **461**, 941–946 (2009).
52. Aronov, D. & Tank, D.W. Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron* **84**, 442–456 (2014).
53. Dombeck, D.A., Harvey, C.D., Tian, L., Looger, L.L. & Tank, D.W. Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nat. Neurosci.* **13**, 1433–1440 (2010).
54. Pologruto, T.A., Sabatini, B.L. & Svoboda, K. ScanImage: flexible software for operating laser scanning microscopes. *Biomed. Eng. Online* **2**, 13 (2003).
55. Greenberg, D.S. & Kerr, J.N.D. Automated correction of fast motion artifacts for two-photon imaging of awake animals. *J. Neurosci. Methods* **176**, 1–15 (2009).
56. Shi, J. & Malik, J. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**, 888–905 (2000).
57. Murphy, K.P. *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
58. Smola, A. & Vapnik, V. Support vector regression machines. *Adv. Neural Inf. Process. Syst.* 281–287 (1997).
59. Chang, C.C. & Lin, C.J. LIBSVM: A library for support vector machines. *ACM Trans. Intell. Syst. Technol.* **2.3**, 27 (2011).
60. Chang, C.-C. & Lin, C.-J. Training v-support vector regression: theory and algorithms. *Neural Comput.* **14.8**, 1959–1977 (2002).

doi:10.1038/nn.4403

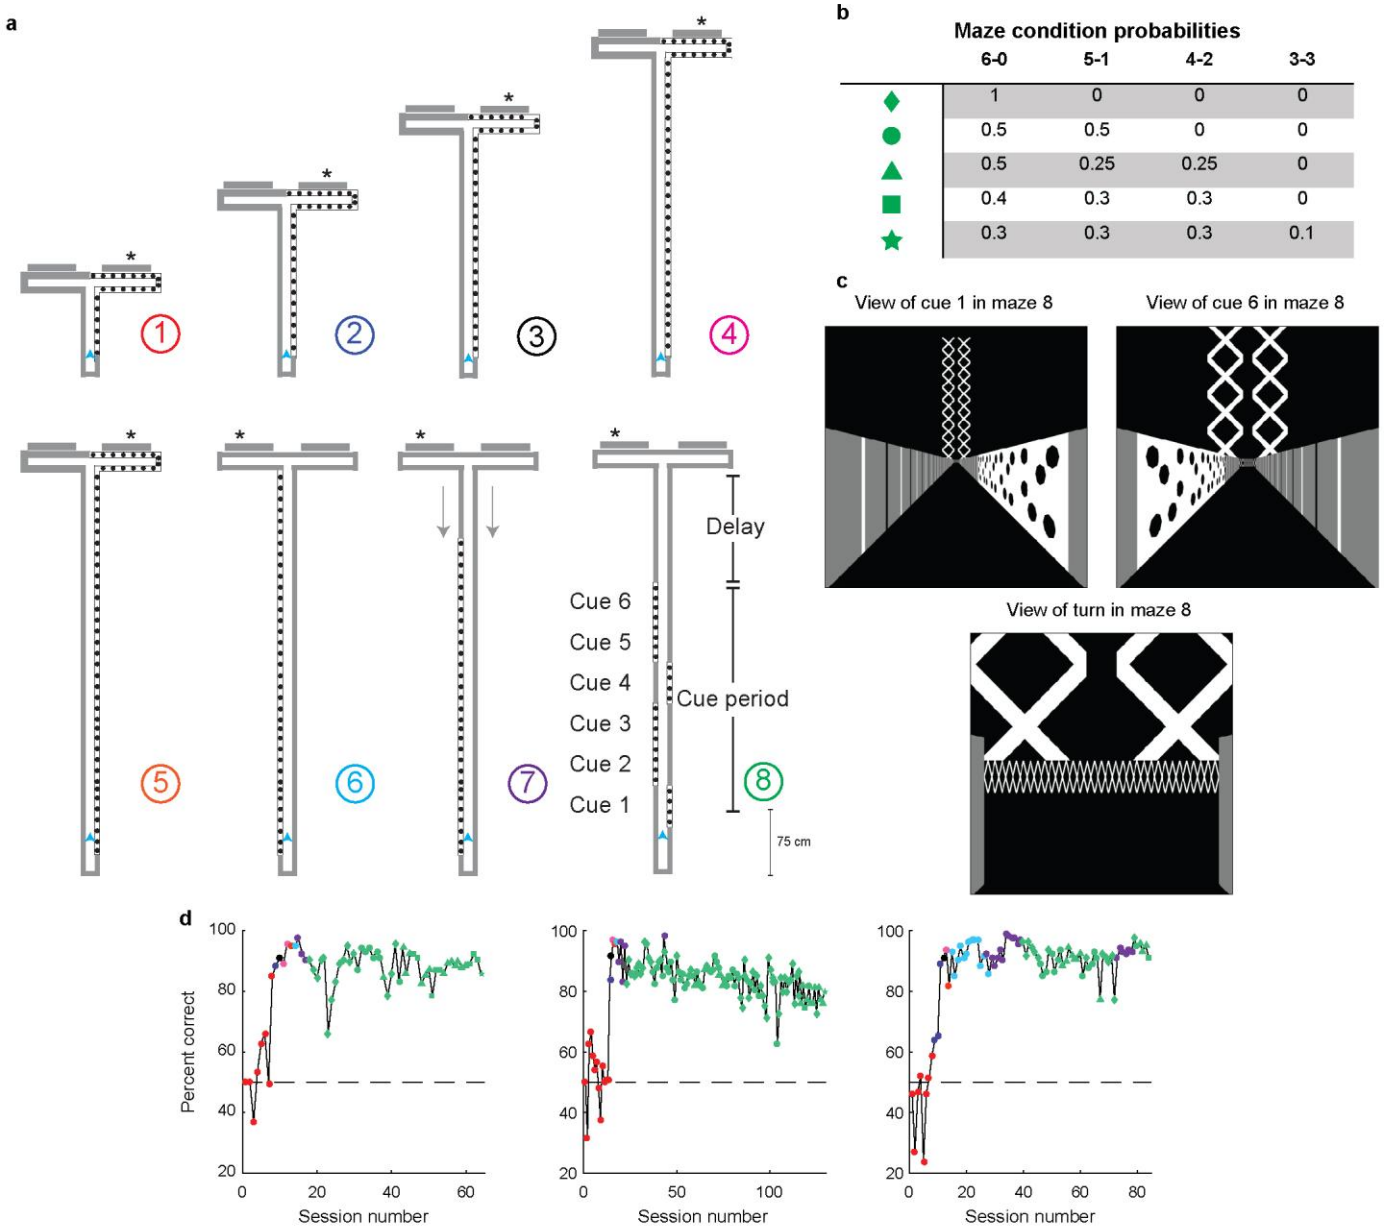

Supplementary Figure 1

Behavioral training.

a, Mazes used for behavioral training. Asterisks indicate reward location. Only some example mazes are shown (for example, right choice and not left choice maze in maze 1). **b,** Distribution of net evidence corresponding to different difficulties used in training the final task (maze 8; see **d**). **c,** Screen captures of the virtual environment at cue 1, cue 6, and the turn in maze 8. **d,** Behavioral performance across sessions for three example mice. Colors correspond to the maze colors indicated in **a**. Shapes correspond to the net evidence probabilities in **b**.

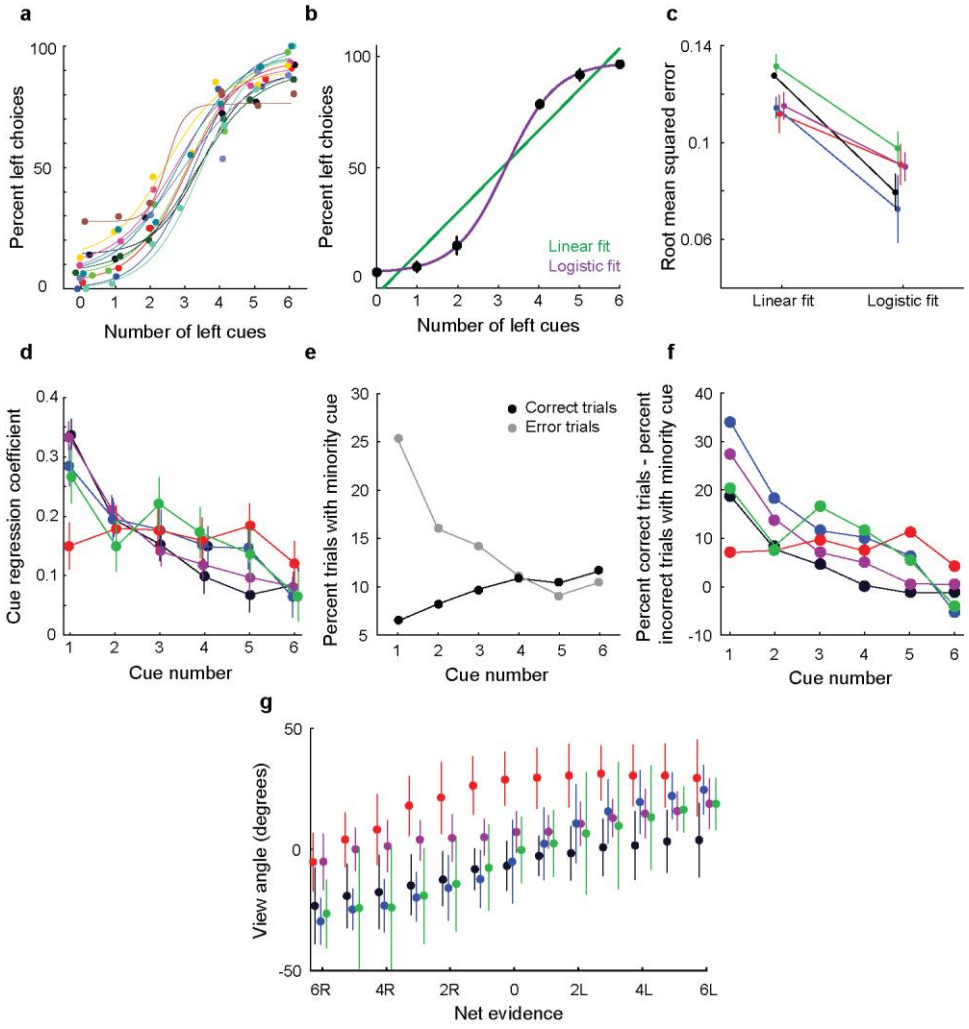

Supplementary Figure 2

Behavioral analysis of evidence accumulation.

a, Behavioral performance on each of the 11 imaging sessions, fit with a logistic function. **b,** Example performance from a single mouse across seven behavioral sessions fit by a linear (green) and logistic (purple) model (Methods). **c,** Across mice, the logistic model fit the data significantly better than the linear model ($p < 0.05$ for all mice, two-sample Student's t-test), suggesting that mice used more than one piece of evidence per trial to make a choice. Error bars represent mean ± s.e.m. across datasets. Mice are colored the same as in Fig. 1c. **d,** Multivariate linear regression in which the mouse's choice was the response variable and the six cue identities were separate explanatory variables. Regression coefficients for five mice (7-12 sessions each) are shown. Four of the five mice weighted early cues more than late cues. Error bars indicate confidence intervals. **e-f,** Fraction of correct (black) and error (gray) trials containing a minority cue (a cue indicating the incorrect choice) at each cue position, for a single mouse (**e**) and as the difference of the error and correct points (**f**) for five mice. **g,** Relationship between net evidence and view angle for each mouse combined across all cue positions.

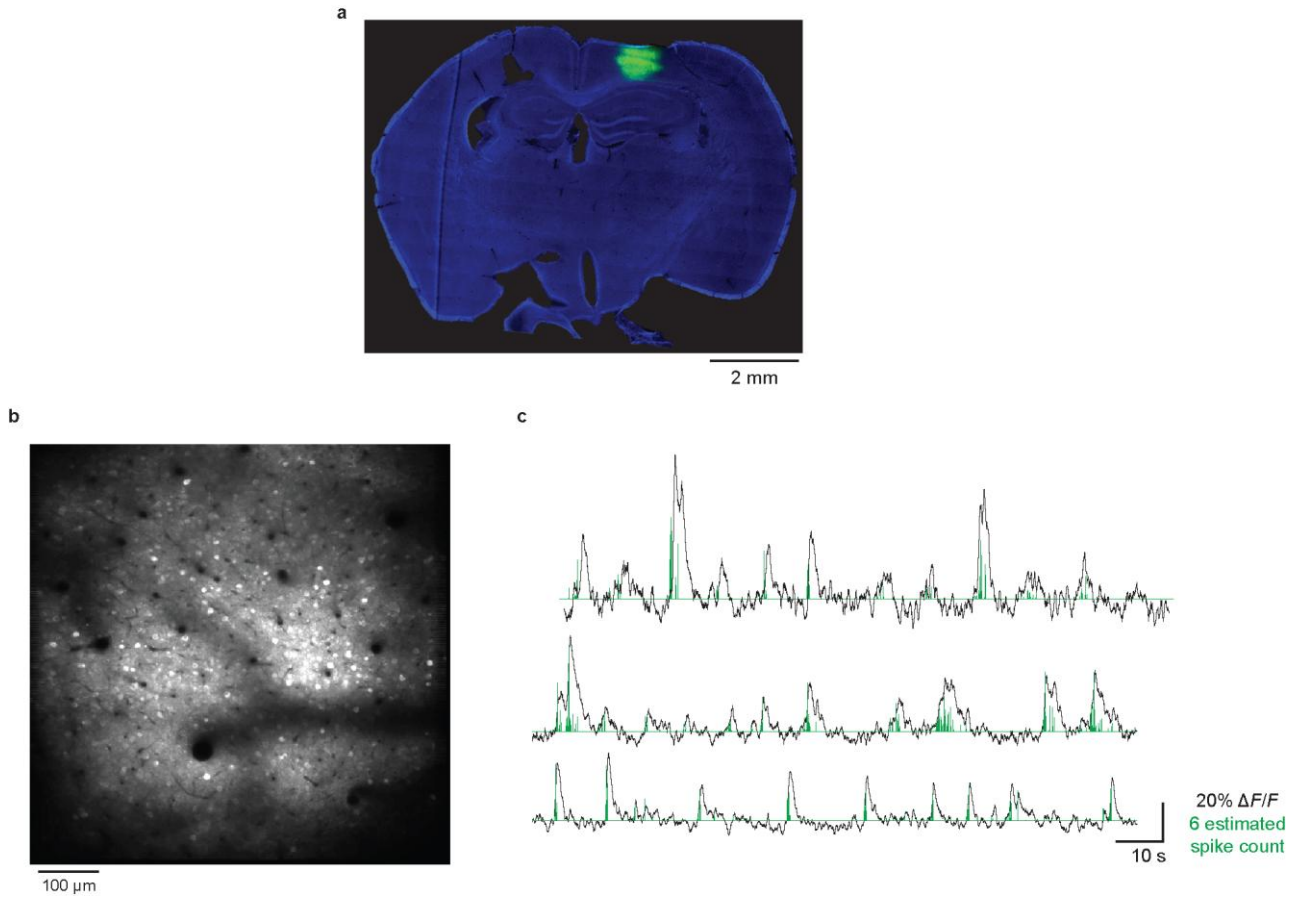

Supplementary Figure 3

Example imaging field of view and activity traces.

a, Example histology image of GCaMP6m-expressing neurons in the PPC. **b,** Example two-photon image of GCaMP6m-expressing neurons in layer 2/3 of the PPC. **c,** Example ΔF/F traces (black) and deconvolved estimated spike counts (green) (Methods).

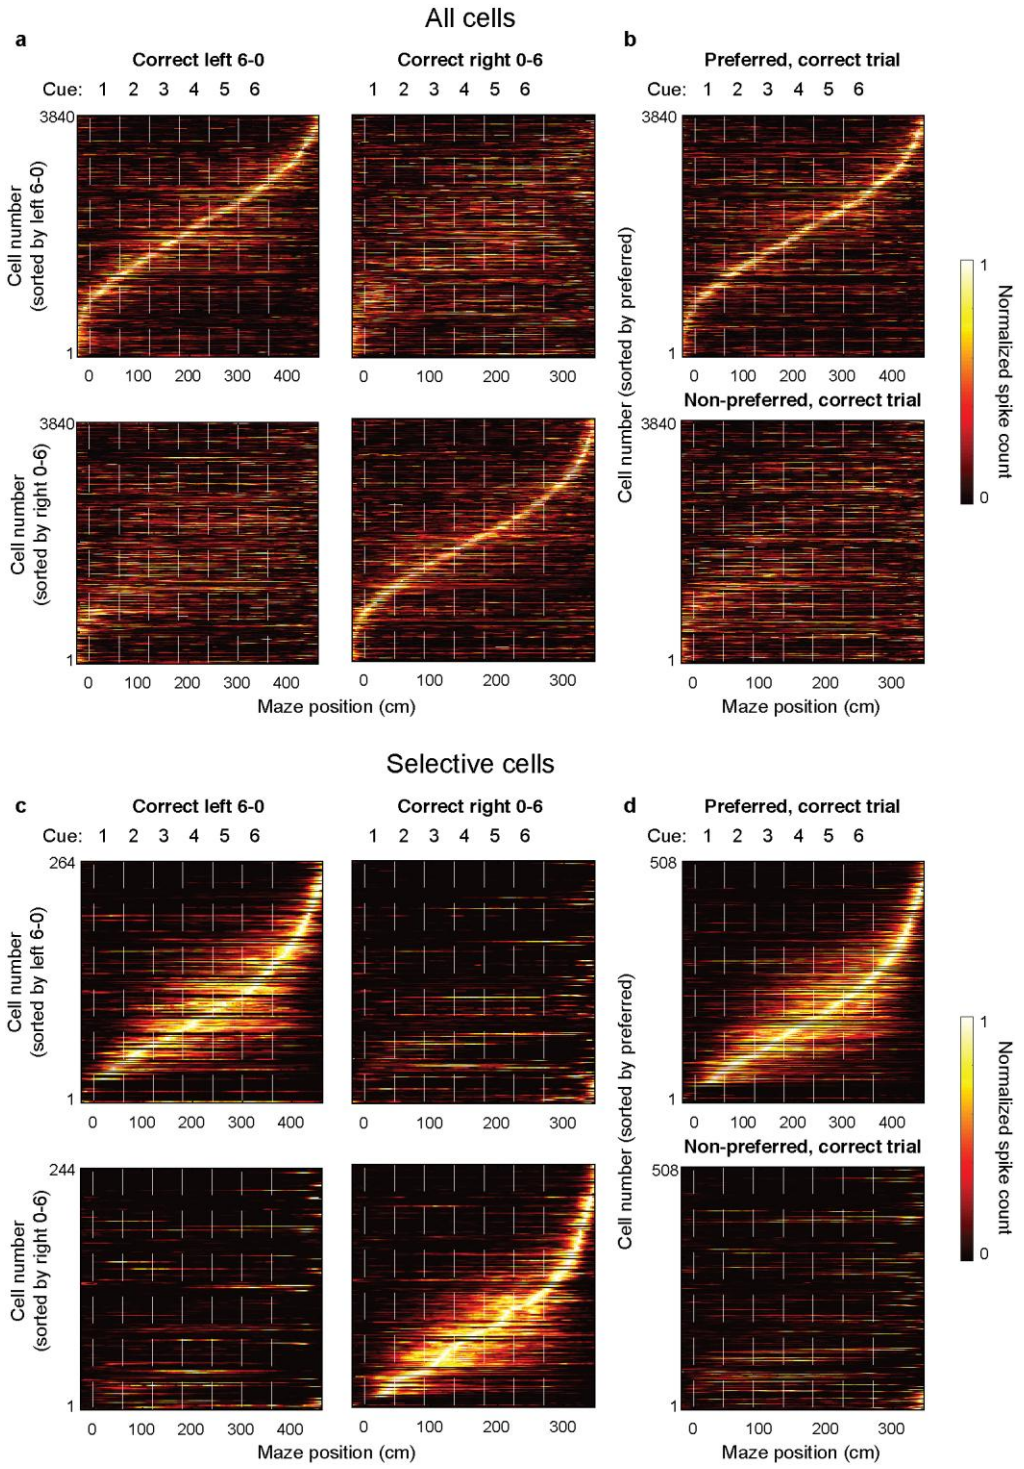

Supplementary Figure 4

Mean population activity patterns in PPC for all cells and selective cells.

a, Normalized mean activity across correct left 6-0 (left) and correct right 0-6 (right) trials for all neurons pooled across all datasets (n = 3840 cells from 11 datasets, 5 mice). Traces were normalized to the peak of each cell's activity on either correct left 6-0 (top) or correct right 0-6 (bottom) trials, averaged, and sorted by the peak's maze position. **b,** Same as in **a**, except for on preferred (top) or non-preferred (bottom) correct 6-0 trials. Cells were sorted according to each cell's activity in its preferred condition. Preferred trial type was

determined for each cell individually based on which trial type had higher mean activity. **c-d,** Same as **a-b,** but only for selective cells. Selective cells were defined as all cells with choice classification accuracy above 70%.

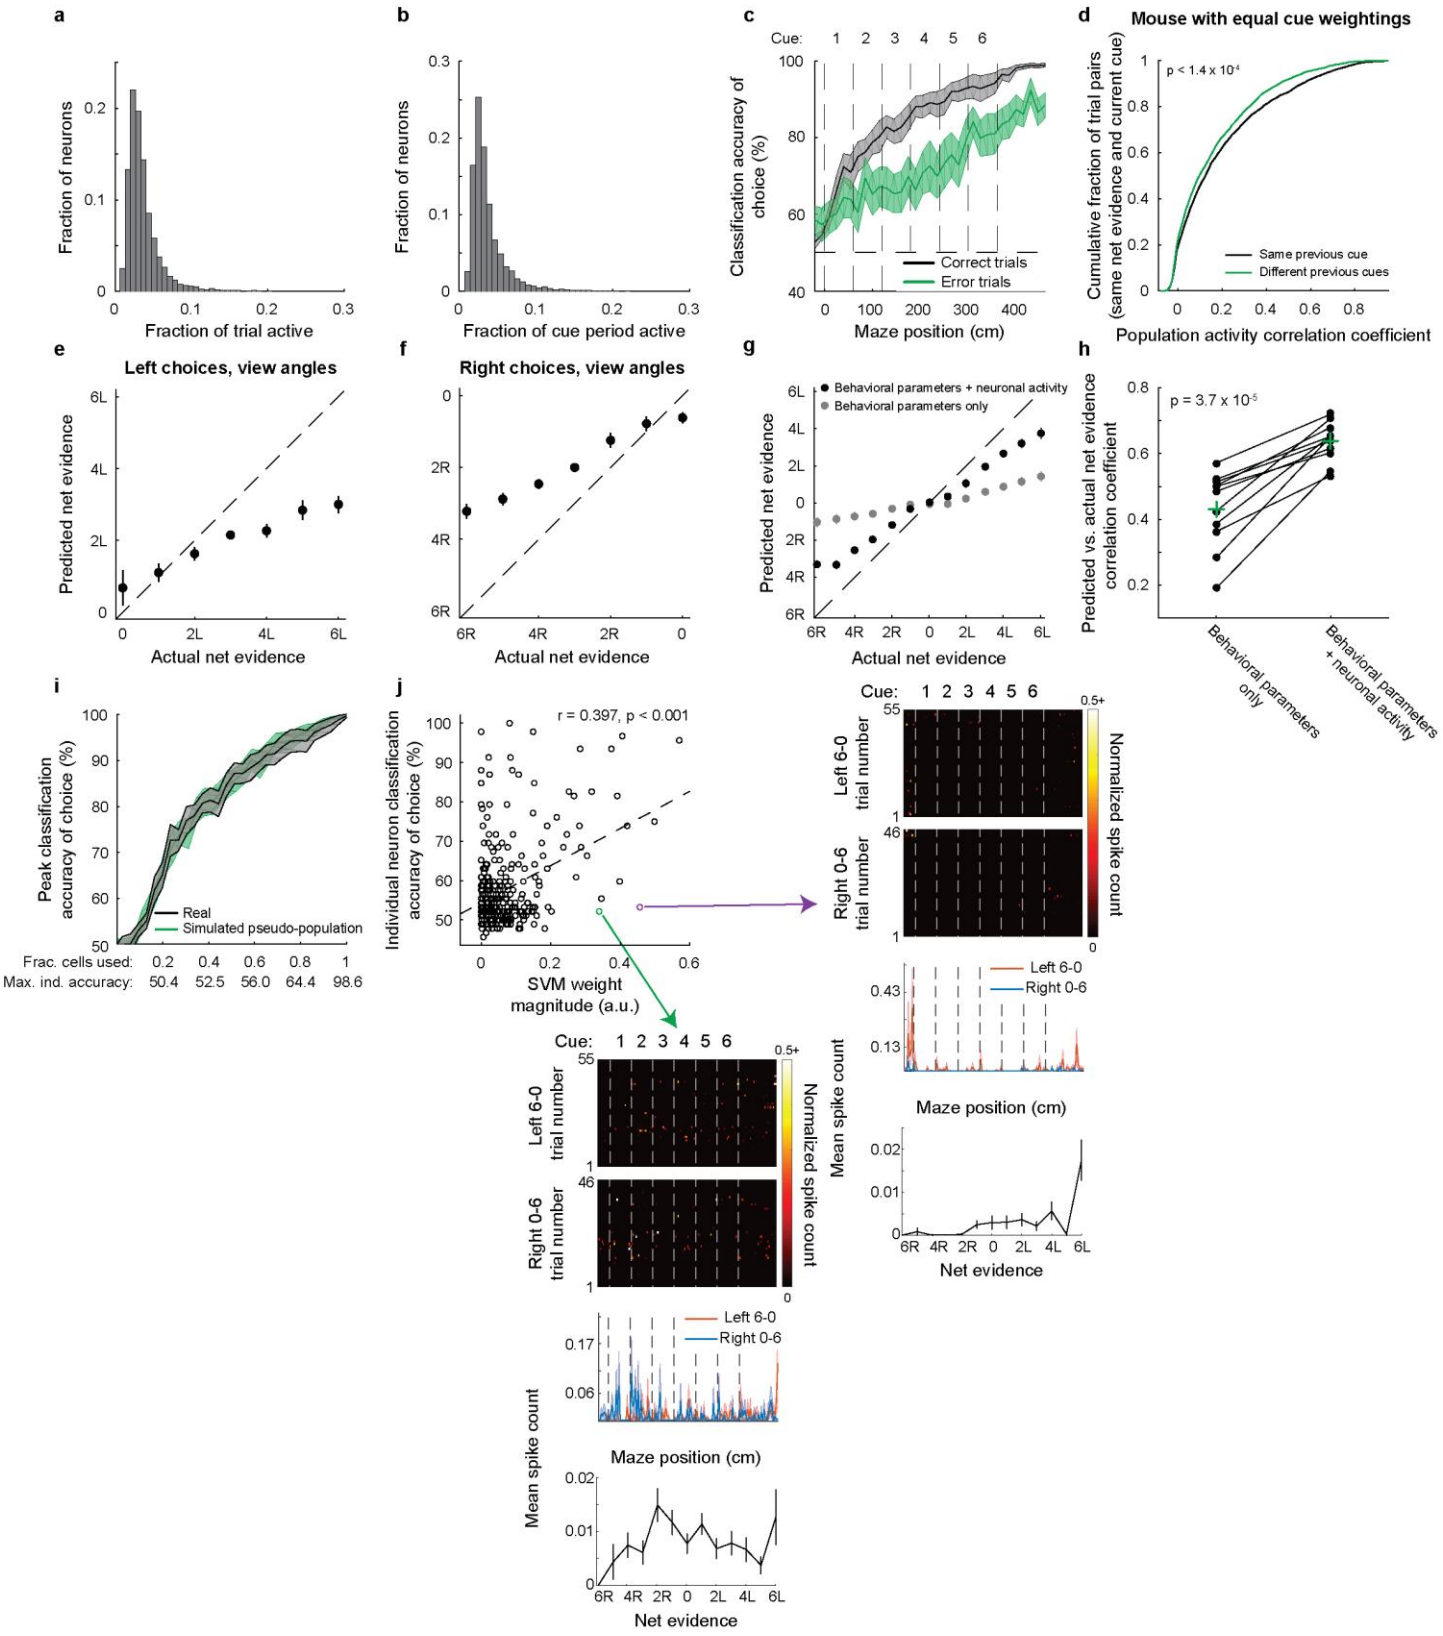

Supplementary Figure 5

Analyses of single-neuron- and population-level representations of task-relevant features.

a-b, Histogram of the fraction of the entire trial (**a**) and cue period (cues 1-6) (**b**) neurons were active (n = 3840 neurons from 5 mice). **c,** SVM classification accuracy (mean ± s.e.m., n = 11 datasets) for choice based on population activity on correct and error trials. Independent classifiers were trained and tested at each maze position. **d,** Same as Fig. 6a except for a mouse with equal cue weightings. Cumulative distribution of the pairwise trial-trial population activity correlation coefficients for epochs with the same (black) or different (green) previous cues, keeping net evidence and epoch constant (e.g. *LRLXXX* vs. *RLLXXX* trials at cue 3) ($p < 1.4 \times 10^{-4}$, two-sample KS test, n = 2 datasets; mouse colored as red in Fig. 1c, Supplementary Fig. 2b, d). This analysis tested if neuronal activity at a given epoch contained information about the previous epoch's cue, independent of maze epoch and net evidence. **e-f,** SVR classifiers for net evidence performed on trials with nearly identical (±2.5°) view angles on left choice (**e**) and right choice (**f**) trials. **g**, Actual net evidence vs. net evidence predicted by an SVR classifier trained on behavioral parameters only (gray) or both behavioral parameters and neuronal population activity (black) (Methods). Error bars represent mean ± s.e.m. across datasets (n = 11). Across mice the predicted vs. actual net evidence correlation coefficient was significantly higher for the model with behavioral parameters and neuronal activity than for the model with behavioral parameters only ($p < 0.001$ relative to shuffled net evidence labels). Net evidence therefore appeared decodable beyond information provided by view angle. **h,** Data from (**g**) shown for individual datasets. Green crosses represent means across datasets (n = 11; $p = 3.7 \times 10^{-5}$, two-sample Student's t-test). **i,** Peak classifier accuracy for choice for classifiers constructed with increasing numbers of neurons, added from least to most selective (based on histograms from Fig. 2c). Real data are shown in black and a simulated pseudo-population is shown in green. To create the pseudo-population, trial identities were shuffled (within a trial-type category) independently for each neuron to break neuron-neuron correlation structure but to preserve each neuron's activity within the trial (Methods). Shaded error bars represent mean ± s.e.m. across datasets, and max individual neuron classification accuracies/correlations were the mean across datasets. **j,** Individual neurons' choice classification accuracy as a function of the magnitude of the weight placed on each neuron by a linear SVM choice classifier trained on all neurons. The population classifier reached a peak accuracy of 100%. While neurons with higher individual classification accuracy were weighted more strongly, the SVM still weighted some neurons with low individual accuracy. Single trial activity on left 6-0 and right 0-6 trials for two example neurons with relatively high weight are also shown. These two neurons illustrate two ways that neurons with low individual selectivity can contribute to a population code. The left neuron (green) is active on both trial types with high variability, but slightly more so on right trials. The right neuron (purple) is primarily active on left trials, but is only active on a small subset of trials (see Methods). Top panels: each row is an individual trial. Bottom panels: mean ± s.e.m. For each net evidence condition (e.g. 2L), the mean spike count was calculated by combining the activity at all cue epochs matching the given net evidence.

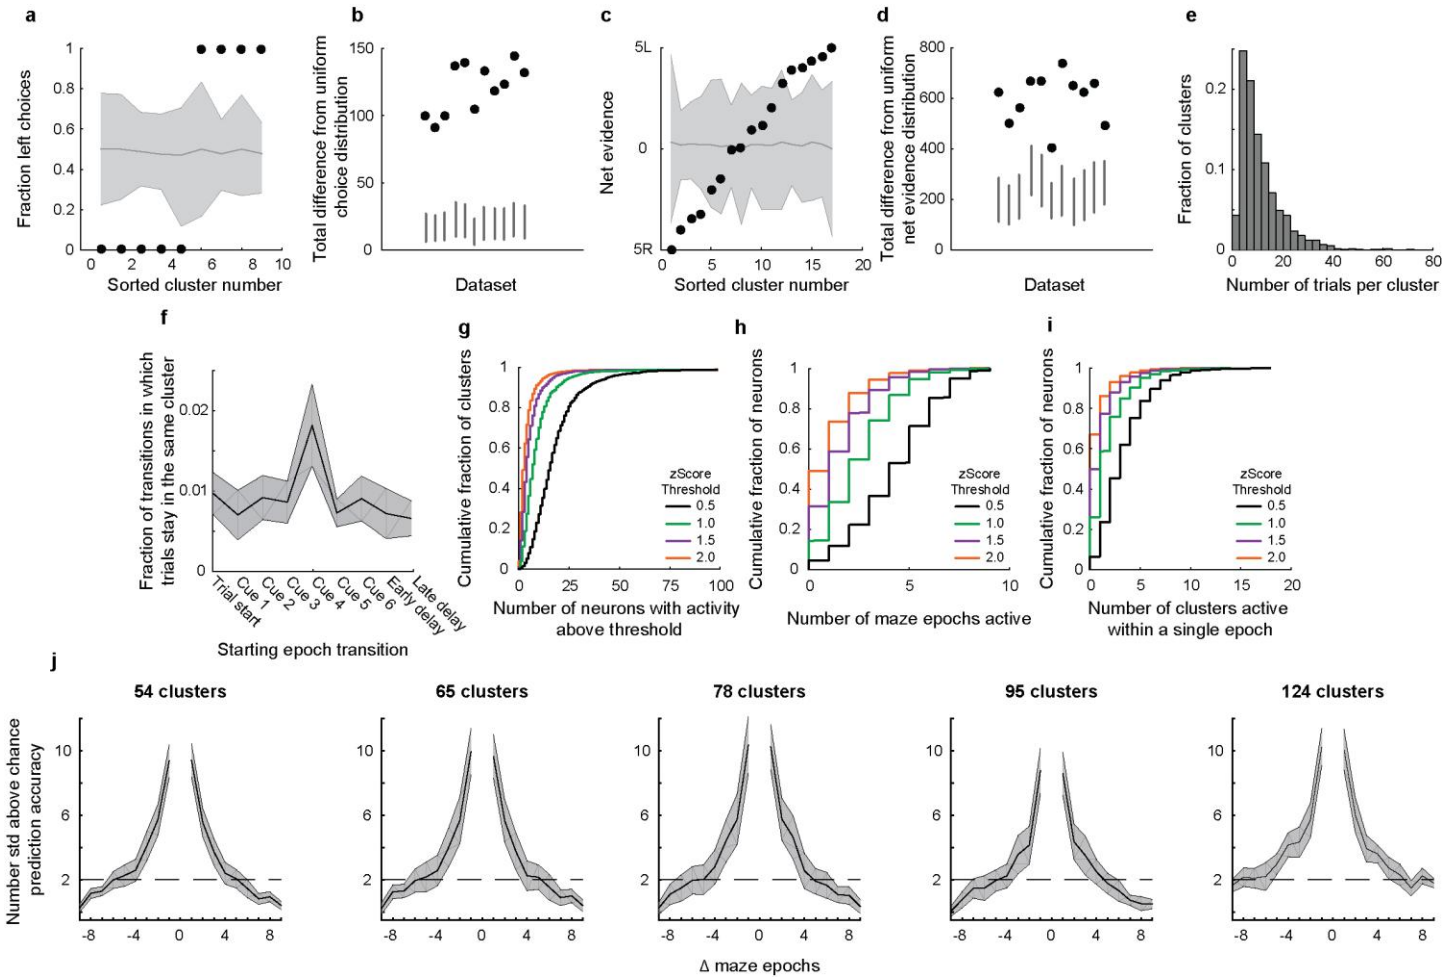

Supplementary Figure 6

Characterization of behavioral and neuronal patterns across clusters.

a, Fraction of trials in each cluster in the turn epoch that were left choice trials for an example dataset. Clustering revealed neuronal activity patterns related to behavioral choices. Gray area indicates the median and 99% confidence intervals of the shuffled distribution of trial assignments to clusters. **b,** Comparison of the total difference from a uniform distribution for the real data (circles) to the 99% confidence intervals of the corresponding shuffle for each dataset (lines). The total difference was calculated as the summed absolute difference from the shuffle median across clusters. **c-d,** Same as in **a-b,** but for net evidence during the fifth cue. **e,** Distribution of trials per cluster across all epochs and datasets (n = 2457 clusters). **f,** Cluster self-transition probabilities for clustering performed using all epochs together. Transition probabilities were considered from one epoch to the next epoch. Low self-transition probabilities suggested that activity patterns changed over the time of consecutive epochs. Error bars represent mean ± s.e.m across datasets. **g,** Cumulative distribution of the number of neurons active in each cluster for different z-score activity thresholds. **h,** Cumulative distribution of the number of maze epochs in which a neuron was active in at least one cluster for different z-score activity thresholds. **i,** Cumulative distribution of the number of clusters in which a neuron was active within a single epoch for different z-score activity thresholds. **j,** For a given trial based on the current cluster identity, the accuracy of predicting the clusters occupied by that trial in the past and future epochs did not depend greatly on the clustering preference parameters (percentile of the distance matrix used for clustering; 1st, 10th, 30th, 50th, 70th from left to right) and, hence, numbers of clusters. Cluster numbers are the mean number of clusters for each preference parameter across datasets. Error bars represent mean ± s.e.m. across datasets.

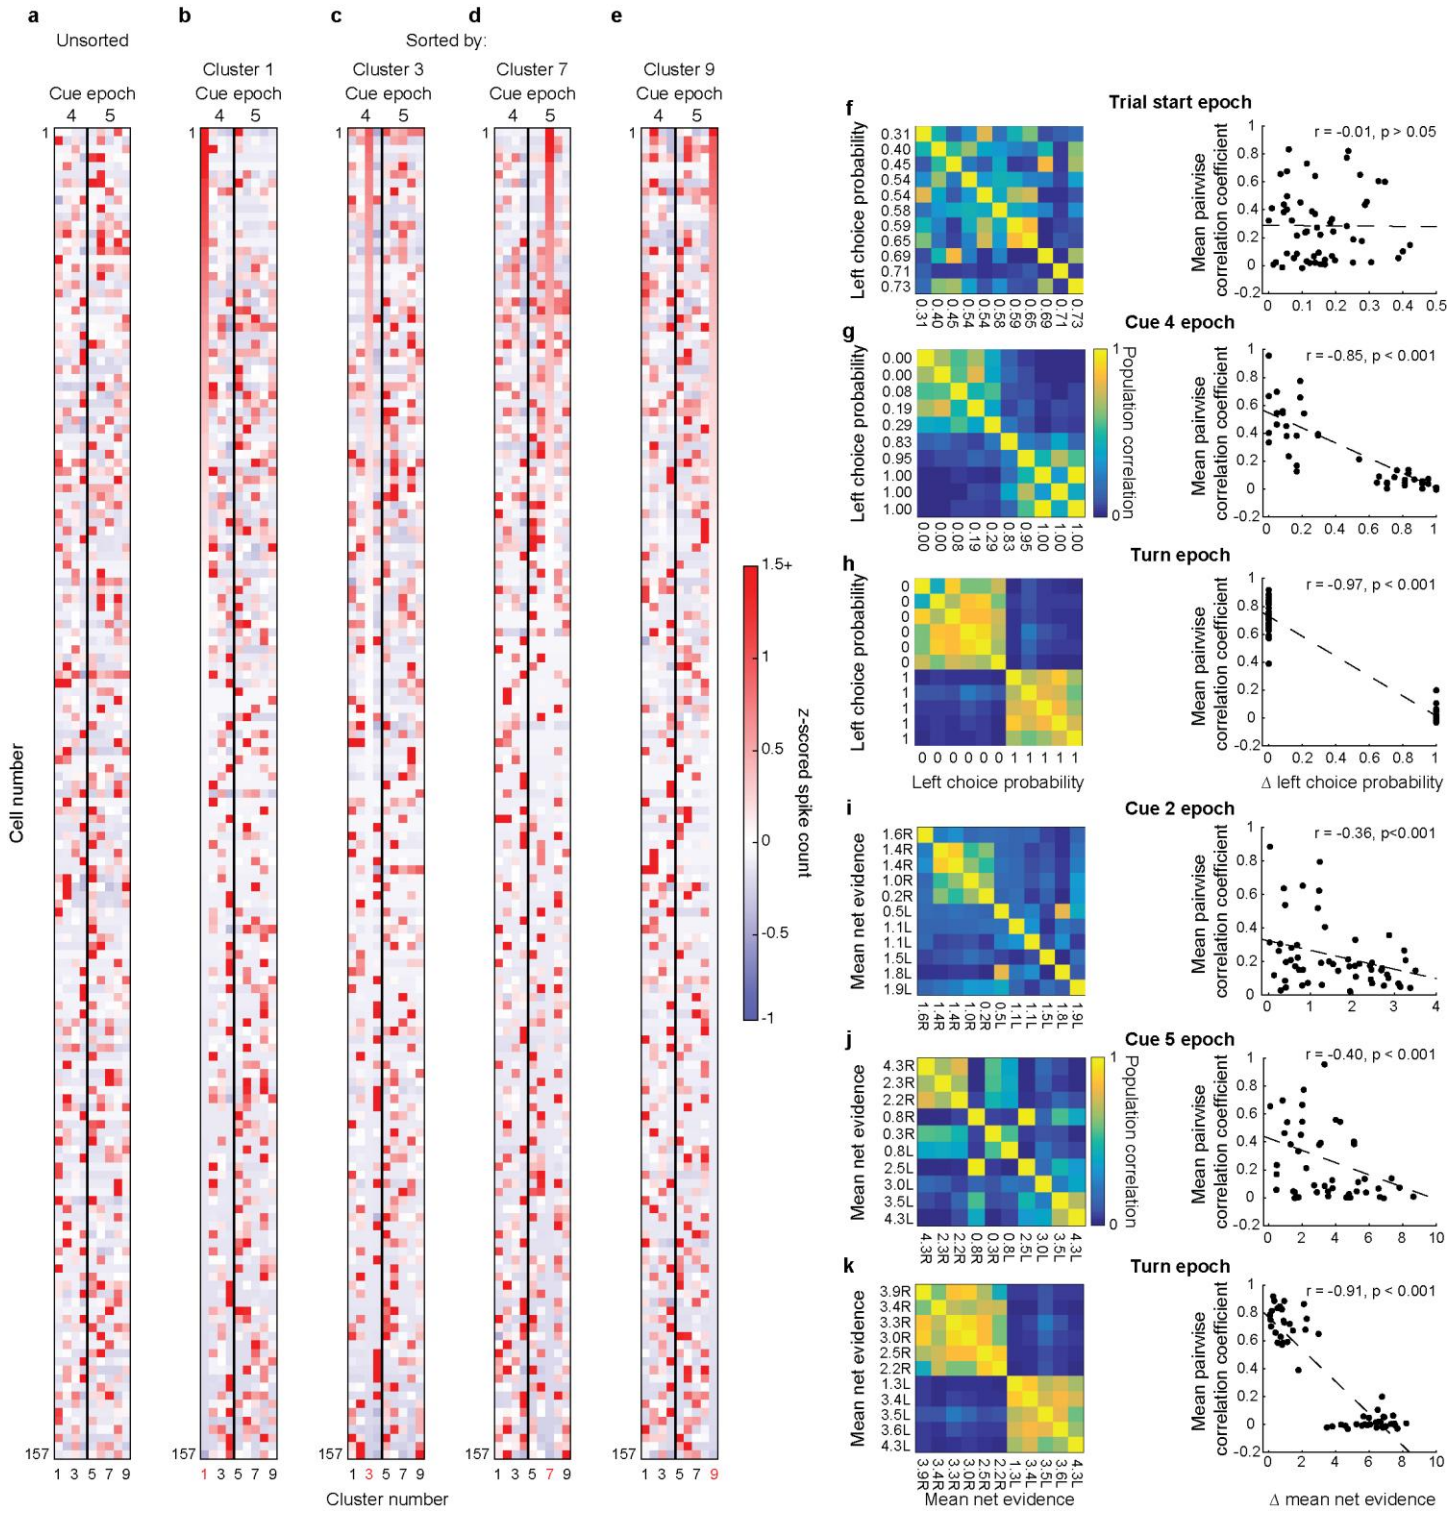

Supplementary Figure 7

Visualizations of neuronal activity across clusters.

a-e, Mean z-scored spike count for individual neurons across clusters comprised only of correct left 6-0 trials at two adjacent epochs (Cues 4 and 5) from a single dataset. These plots demonstrate that the activity across clusters and epochs featured largely different patterns of active neurons. Neurons were either unsorted (**a**) or sorted according to their activity in clusters 1, 3, 7, or 9 (**b-e**). Neurons

whose mean z-scored activity was less than 0.001 in all of the displayed clusters were excluded for display purposes (these neurons were active during a different trial epoch). Clusters were generated from correct left 6-0 trials. **f-h,** Left panels: Matrix of population activity correlations between each pair of cluster centers sorted according to the cluster's left choice probability at three different maze epochs. For each cluster, the population activity was calculated as the mean activity vector across trials for each cluster. Right panels: Population activity correlation between each pair of clusters as a function of their difference in left choice probability. **i-k,** Same as in **f-h**, but for net evidence.

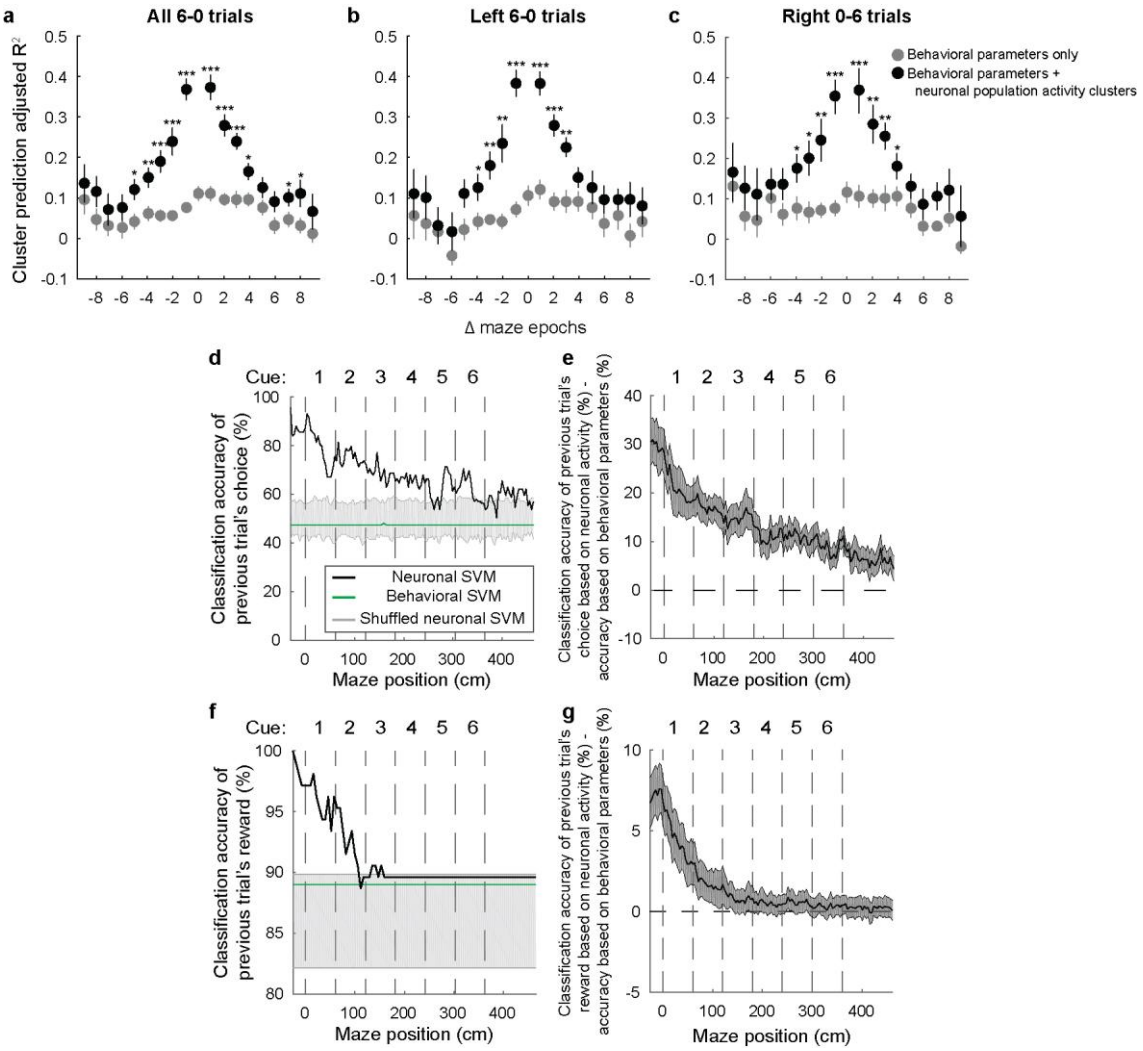

Supplementary Figure 8

Contribution of behavioral variability to trial–trial variability and classification of the previous trial's outcome.

a-c, Our ability to predict the past and future population activity pattern based on the current population activity pattern could not be explained by behavioral variability. We performed a multivariate logistic regression to predict a trial's cluster identity at a given epoch based on only the behavioral parameters at another epoch (gray) or both the behavioral parameters and the cluster identity at another epoch (black). To allow for a binary classifier, we only included those trials whose cluster identity contained either the most or second most trials during the prediction epoch (Methods). Consistently, the model based on both behavioral parameters and the previous cluster identity outperformed the model based on only behavioral parameters. This analysis was performed on left 6-0 trials (b) and right 0-6 trials (c) separately, and pooled together for all 6-0 trials (a). The behavioral parameters used were x/y position, x/y treadmill velocity, and view angle. Separate models were fit for each combination of previous and future cluster identities and combined based on the number of maze epochs between them (Δmaze epochs). Adjusted R^2 values were used to compare the predictive power of models with different numbers of explanatory variables. *P < 0.05, **P < 0.01, ***P < 0.001, two-sample Student's t-test. **d,** Comparison of a neuronal activity-based SVM (black), behavioral parameter-based SVM (green), and the 99% confidence interval of a neuronal activity-based SVM with shuffled labels (gray) for the previous trial's choice for a single dataset. The behavioral parameter-based SVM could not discriminate the previous trial's choice. Classifiers were trained to distinguish the mouse's choice on the previous trial independently at each bin in the current trial. **e,** Difference between the classification accuracy of the neuronal activity-based SVM and the behavioral parameter-based SVM for the previous trial's choice. Error bars represent mean ± s.e.m. across datasets. **f-g,** Same as in (**d-e**), but with classifiers for whether or not the previous trial was rewarded.

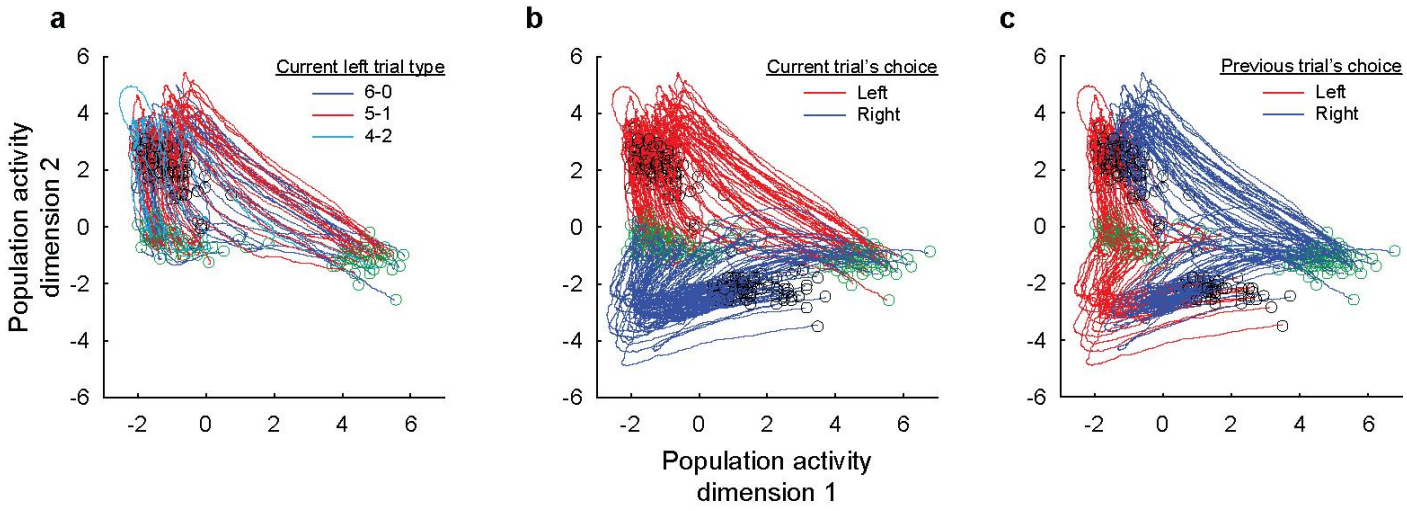

Supplementary Figure 9

Visualizations of trial trajectories.

a-c, Trajectories of correct trials colored by the current trial type (**a**), the current trial's choice (**b**), and the previous trial's choice (**c**). Trials with the same choice but different trial types were highly overlapping (**a**), while trials with different choices were highly different (**b**). Much of the variance within a choice could be explained by the outcome of the previous trial (**c**). Green and black circles mark the trial start and trial end, respectively. For visualization purposes, the dimensionality of the data was reduced using factor analysis.

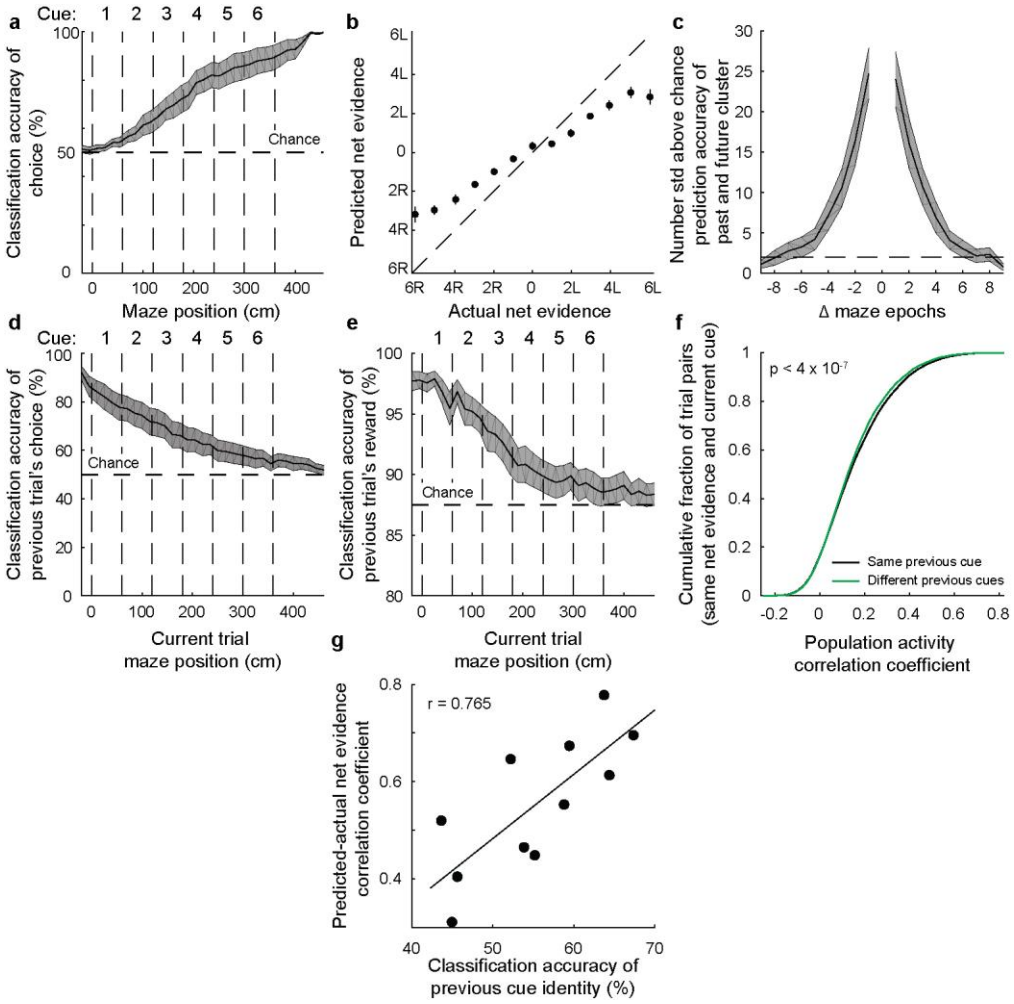

Supplementary Figure 10

Main results reanalyzed using Δ*FIF* values.

a, Classification accuracy for choice as a function of maze position (SVM, radial basis function kernel). Independent classifiers were trained and tested at each maze position. Error bars represent mean ± s.e.m. across datasets. Compare to Fig. 2e. **b,** Actual net evidence vs. net evidence predicted by a SVR classifier. Error bars represent mean ± s.e.m. across datasets. Compare to Fig. 2f. **c,** For a given trial based on the current epoch's cluster identity, the accuracy of predicting the clusters occupied by that trial in the past and future epochs, compared to shuffled assignments of trials to clusters. Error bars represent mean ± s.e.m. across datasets. Compare to Fig. 4e. **d-e,** Classification accuracy as in (**a**), but for previous trial's choice and for whether the previous trial was rewarded (**e**). Compare to Fig. 5b-c. **f,** Cumulative distribution of the pairwise trial-trial population activity correlation coefficients for trials with the same (black) or different (green) previous cues given the same maze epoch and same net evidence ($p < 4 \times 10^{-7}$, two-sample KS test). Compare to Fig. 6a. **g,** Relationship between classification accuracy of the previous cue and the classification accuracy of net evidence across datasets ($r = 0.76$, $p < 0.001$). Compare to Fig. 6b.

Session ID	# cells	6-0 accuracy / # trials	5-1 accuracy / # trials	4-2 accuracy / # trials	3-3 accuracy / # trials	GCaMP Variant
131_140911	194	100% / 107	91.1% / 90	82.4% / 34	N/A / 0	f
131_140916	188	94.7% / 94	85.7% / 105	81.8% / 33	N/A / 0	f
136_140820	270	93.1% / 87	89.8% / 79	66.0% / 47	N/A / 0	f
142_141212	434	90.0% / 181	85.0% / 181	76.9% / 52	N/A / 0	m
142_141218	381	89.6% / 135	81.7% / 120	67.5% / 40	N/A / 0	m
142_150103	215	90.7% / 86	83% / 94	74.3% / 105	60.5% / 38	m
144_141203	648	87.8% / 106	79.5% / 146	73.8% / 42	N/A / 0	m
144_141204	585	88.6% / 140	82.6% / 155	66.7% / 72	N/A / 0	m
144_141228	330	80.9% / 173	77.9% / 163	77.6% / 49	N/A / 0	m
150_141128	323	100% / 114	88.3% / 102	73.1% / 108	41.7% / 12	f
150_141207	313	98.1% / 107	82.2% / 107	78.4% / 116	57.1% / 49	f

Supplementary Table 1

Summary of datasets analyzed.