HHS Public Access

Author manuscript

Nat Neurosci. Author manuscript; available in PMC 2017 April 03.

Published in final edited form as:

Nat Neurosci. 2016 December ; 19(12): 1672–1681. doi:10.1038/nn.4403.

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

Ari S. Morcos[1] and **Christopher D. Harvey**[1,2]

[1] Department of Neurobiology, Harvard Medical School, Boston, MA 02115

Abstract

We studied how the posterior parietal cortex combined new information with ongoing activity dynamics as mice accumulated 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.

Introduction

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,

[2] Correspondence should be addressed to: harvey@hms.harvard.edu.

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 statement

The authors have no competing financial interests.

including long-lasting firing rate changes in homogenous 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, alternate 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 (e.g. 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; Supplementary Fig. 1; Methods M.2). 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 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; Supplementary Fig. 2; Methods M.2.4.1).

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] (Methods M.3; Supplementary Fig. 3, 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; Supplementary Figs. 4, 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 based on a single cell's activity and a support vector regression (SVR) model to predict net evidence from a single cell's activity (Methods M. 4.4.2-3). 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 if task-relevant information was present only in the fraction of cells that had high selectivity or if 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 (see results in Methods M.2.5; Supplementary Fig. 5g-h). 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; for further results see Methods M.4.4.4 and Supplementary Fig. 5j). 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 conceptualized the population activity as 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; Methods M.4.6 and related discussion).

Specifically, the number of clusters was determined 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; Methods M.4.6.2). 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, Supplementary Fig. 6e). Transitions between clusters in adjacent epochs were marked as lines with thickness proportional to the transition probability (Methods M.4.6.3; Fig. 3b-d).

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; 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 of 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 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 neurons (Fig. 2a-b, Supplementary Figs. 4, 5a-b), the activity patterns in clusters at one epoch were largely different from the activity patterns observed in clusters at the subsequent 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 ($0.9 \pm 0.01\%$ of transitions; Supplementary Fig. 6f; Methods M.4.6.4). The activity patterns in each cluster were made up of complex combinations of activity levels in the population of individual neurons (Fig. 3j, Supplementary Figs. 7a-e). Some individual neurons thus had elevated activity in multiple clusters (Fig. 3j, 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 (e.g. correct 6-0 left 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].

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 (e.g. correct 6-0 left trial; 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 (Methods M.4.6.4). 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 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 (e.g. 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 if we could predict the future activity patterns of a trial based on 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 if 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 real data (Fig. 4c).

This example suggested that the transitions between activity patterns could be non-random 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 if, based on 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 (Methods M.4.7). This analysis therefore tests if the current activity pattern contained 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 5-6 epochs (~4-5 seconds) into the past and future (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 (see results in Methods M.2.5; Methods M.4.7.2-3; Supplementary Fig. 8a-c). Together, these results indicate that the current activity pattern contained information about past activity patterns and influenced the transition probabilities to future activity patterns, even when removing the effects of different trial events like 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. Multiple features of the data provided support for the second alternative. We found that neurons were transiently active with time-varying activity (Fig. 2a-b, Supplementary Figs. 4, 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 if 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 based on 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 prior to 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 ended. 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 if 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 (similar results were obtained with clustering). The population activity patterns at the start of a trial, following an inter-trial interval of at least two seconds, 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; 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 ten seconds 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 based on the previous trial's choice and reward (R^2: 0.02 ± 0.01, mean ± s.e.m., $p > 0.05$; Methods M.2.4.2)[34]. Also, the previous trial's choice could not be decoded from the current trial's behavioral data (e.g. running patterns; see results in Methods M.2.5; Supplementary Fig. 8d-g). 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 (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 based on 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 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 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 on a trial independent of net evidence. This prediction implies that different sequences of cues that result in the same net evidence (e.g. 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 (e.g. left) and the same net evidence (e.g. +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 KS test; Fig. 6a). Activity in an epoch could therefore be classified above chance levels based on 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; Methods M.4.4.5). While this difference was highly significant, it was relatively 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 (Methods M. 4.4.5). 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 datasets, 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 if we could predict the population activity pattern following a new evidence cue (either left or right cue) based on a) the distribution of trials across clusters alone (chance), b) the new cue type alone (cue only), c) the current activity cluster alone (cluster only), and d) both the current activity cluster and the new cue type (cue + cluster) (Methods M.4.7.1). 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 novel model for how evidence accumulation is performed in neuronal circuits. First, we have shown 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, 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 (Fig. 3h,j; Fig. 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 have shown 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 non-measured or hidden

parameters, beyond those parameters directly tested in an experiment (e.g. choice or net evidence)[19,20]. The presence of hidden signals impacts 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 non-relevant 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 and 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 (Fig. 5, 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, Supplementary Figs. 4, 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 datasets). 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 towards, 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 due 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 prior to 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 due 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 significant 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.

Data and Code Availability

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

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

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 helpful discussions, and B. Datta, D. Dombeck, J. Drugowitsch, C. Gu, and members of the Harvey lab 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.

References

1. Gold JI, Shadlen MN. The Neural Basis of Decision Making. Annu. Rev. Neurosci. 2007; 30:535–574. [PubMed: 17600525]

2. Wang X-J. Neural dynamics and circuit mechanisms of decision-making. Current Opinion in Neurobiology. 2012; 22:1039–1046. [PubMed: 23026743]
3. Shadlen MN, Newsome WT. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 2001; 86:1916–1936. [PubMed: 11600651]
4. Hanks TD, et al. Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature. 2015; 520:220–223. [PubMed: 25600270]
5. Raposo D, Kaufman MT, Churchland AK. A category-free neural population supports evolving demands during decision-making. Nat. Neurosci. 2014; 17:1784–1792. [PubMed: 25383902]
6. Licata AM, et al. Posterior parietal cortex guides visual decisions in rats. bioRxiv. 2016:066639. doi: 10.1101/066639.
7. Goard MJ, Pho GN, Woodson J, Sur M. Distinct roles of visual, parietal, and frontal motor cortices in memory-guided sensorimotor decisions. Elife. 2016; 5:471.
8. Harvey CD, Coen P, Tank DW. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. Nature. 2012; 484:62–68. [PubMed: 22419153]
9. Wong K-F, Wang X-J. A recurrent network mechanism of time integration in perceptual decisions. J. Neurosci. 2006; 26:1314–1328. [PubMed: 16436619]
10. Machens CK. Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination. Science. 2005; 307:1121–1124. [PubMed: 15718474]
11. Horwitz GD, Newsome WT. Separate signals for target selection and movement specification in the superior colliculus. Science. 1999; 284:1158–1161. [PubMed: 10325224]
12. Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. Nat. Neurosci. 2008; 11:823–833. [PubMed: 18516033]
13. Baeg EH, et al. Dynamics of Population Code for Working Memory in the Prefrontal Cortex. Neuron. 2003; 40:177–188. [PubMed: 14527442]
14. Crowe DA, Averbeck BB, Chafee MV. Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex. Journal of Neuroscience. 2010; 30:11640–11653. [PubMed: 20810885]
15. Rajan K, Harvey CD, Tank DW. Recurrent Network Models of Sequence Generation and Memory. Neuron. 2016 doi:10.1016/j.neuron.2016.02.009.
16. Vogelstein JT, et al. Fast nonnegative deconvolution for spike train inference from population calcium imaging. Journal of Neurophysiology. 2010; 104:3691–3704. [PubMed: 20554834]
17. Meister MLR, Hennig JA, Huk AC. Signal Multiplexing and Single-Neuron Computations in Lateral Intraparietal Area During Decision-Making. Journal of Neuroscience. 2013; 33:2254–2267. [PubMed: 23392657]
18. Jun JK, et al. Heterogenous Population Coding of a Short-Term Memory and Decision Task. Journal of Neuroscience. 2010; 30:916–929. [PubMed: 20089900]
19. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature. 2013; 503:78–84. [PubMed: 24201281]
20. Rigotti M, et al. The importance of mixed selectivity in complex cognitive tasks. Nature. 2013; 497:585–590. [PubMed: 23685452]
21. Maimon G, Assad JA. Beyond Poisson: increased spike-time regularity across primate parietal cortex. Neuron. 2009; 62:426–440. [PubMed: 19447097]
22. Churchland MM, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. Nat. Neurosci. 2010; 13:369–378. [PubMed: 20173745]
23. Frey BJ, Dueck D. Clustering by Passing Messages Between Data Points. 2007
24. Churchland MM, et al. Neural population dynamics during reaching. Nature. 2012 doi:10.1038/nature11129.
25. Mazor O, Laurent G. Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons. Neuron. 2005; 48:661–673. [PubMed: 16301181]
26. Briggman KL. Optical Imaging of Neuronal Populations During Decision-Making. Science. 2005; 307:896–901. [PubMed: 15705844]
27. Renart A, Machens CK. Variability in neural activity and behavior. 2014; 25:211–220.

28. Marcos E, et al. Neural Variability in Premotor Cortex Is Modulated by Trial History and Predicts Behavioral Performance. 2013; 78:249–255.
29. Churchland AK, et al. Variance as a Signature of Neural Computations during Decision Making. Neuron. 2011; 69:818–831. [PubMed: 21338889]
30. Bernacchia A, Seo H, Lee D, Wang X-J. A reservoir of time constants for memory traces in cortical neurons. Nat. Neurosci. 2011; 14:366–372. [PubMed: 21317906]
31. Donahue CH, Lee D. Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. Nat. Neurosci. 2015; 18:295–301. [PubMed: 25581364]
32. Seo H, Barraclough DJ, Lee D. Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. Cerebral Cortex. 2007; 17(Suppl 1):i110–7. [PubMed: 17548802]
33. Seo H, Lee D. Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. J. Neurosci. 2007; 27:8366–8377. [PubMed: 17670983]
34. Busse L, et al. The detection of visual contrast in the behaving mouse. J. Neurosci. 2011; 31:11351–11361. [PubMed: 21813694]
35. Safaai H, Neves R, Eschenko O, Logothetis NK, Panzeri S. Modeling the effect of locus coeruleus firing on cortical state dynamics and single-trial sensory processing. Proceedings of the National Academy of Sciences. 2015; 112:12834–12839.
36. Curto C, Sakata S, Marguet S, Itskov V, Harris KD. A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. J. Neurosci. 2009; 29:10600–10612. [PubMed: 19710313]
37. Nikolić D, Häusler S, Singer W, Maass W. Distributed Fading Memory for Stimulus Properties in the Primary Visual Cortex. PLoS Biol. 2009; 7:e1000260. [PubMed: 20027205]
38. Klampfl S, David SV, Yin P, Shamma SA, Maass W. A quantitative analysis of information about past and present stimuli encoded by spikes of A1 neurons. Journal of Neurophysiology. 2012; 108:1366–1380. [PubMed: 22696538]
39. Seo H, Barraclough DJ, Lee D. Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. J. Neurosci. 2009; 29:7278–7289. [PubMed: 19494150]
40. Sugrue LP, Corrado GS, Newsome WT. Matching behavior and the representation of value in the parietal cortex. Science. 2004; 304:1782–1787. [PubMed: 15205529]
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. 2015 doi: 10.1016/j.neuron.2015.09.008.
42. Murray JD, et al. A hierarchy of intrinsic timescales across primate cortex. Nat. Neurosci. 2014; 17:1661–1663. [PubMed: 25383900]
43. Yang Y, Zador AM. Differences in sensitivity to neural timing among cortical areas. J. Neurosci. 2012; 32:15142–15147. [PubMed: 23100435]
44. Klampfl S, Maass W. Emergence of Dynamic Memory Traces in Cortical Microcircuit Models through STDP. Journal of Neuroscience. 2013; 33:11515–11529. [PubMed: 23843522]
45. Buonomano DV, Maass W. State-dependent computations: spatiotemporal processing in cortical networks. Nat. Rev. Neurosci. 2009; 10:113–125. [PubMed: 19145235]
46. Hoerzer GM, Legenstein R, Maass W. Emergence of Complex Computational Structures From Chaotic Neural Networks Through Reward-Modulated Hebbian Learning. Cerebral Cortex. 2014; 24:677–690. [PubMed: 23146969]
47. Murakami M, Vicente MI, Costa GM, Mainen ZF. Neural antecedents of self-initiated actions in secondary motor cortex. Nat. Neurosci. 2014; 17:1574–1582. [PubMed: 25262496]
48. Jaeger H, Haas H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. Science. 2004; 304:78–80. [PubMed: 15064413]
49. Maass W, Natschläger T, Markram H. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. Neural Computation. 2002; 14:2531–2560. [PubMed: 12433288]
50. Sussillo D, Abbott LF. Generating Coherent Patterns of Activity from Chaotic Neural Networks. Neuron. 2009; 63:544–557. [PubMed: 19709635]

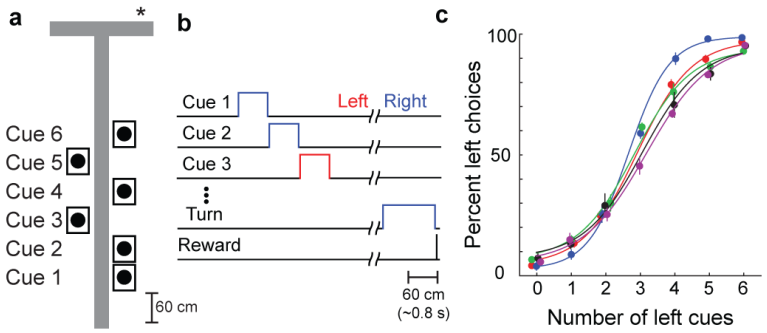

Figure 1. A navigation-based evidence accumulation task in virtual reality

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).

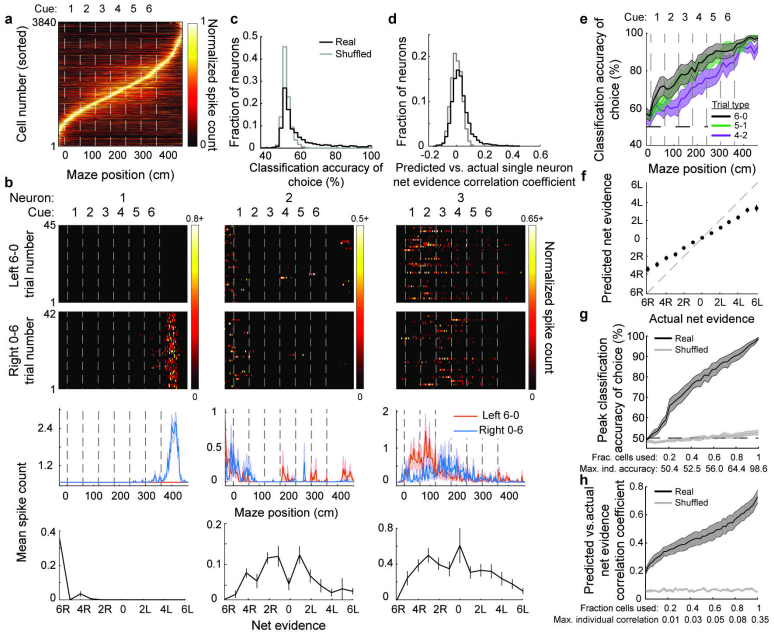

Figure 2. Distributed representation of task-relevant information across PPC neurons

a, Normalized mean activity across all trials for all neurons pooled across all datasets (n = 3840 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 (e.g. 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 datasets) 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 datasets (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 panels (**c**) and (**d**)). Real data (black) and data with shuffled trial labels (gray) are shown. Shaded error bars represent mean ± s.e.m. across datasets, and max individual neuron classification accuracies/correlations were the mean across datasets.

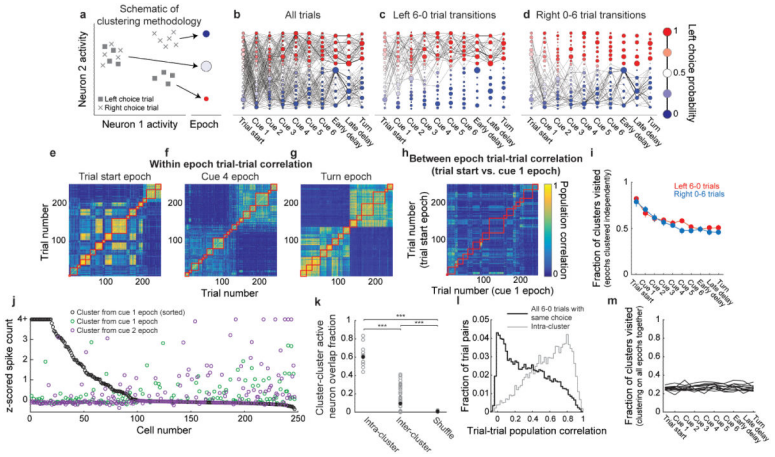

Figure 3. Clustering neuronal activity across trials reveals trial-to-trial variability

a, Schematic demonstrating clustering procedure (Methods M.4.6). 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. Nodes were colored 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 dataset. 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. **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 / number of neurons active in either cluster). Small inter-cluster overlap suggested 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 (Methods M. 4.6.4). All trial types were included. Individual lines represent datasets (n = 11).

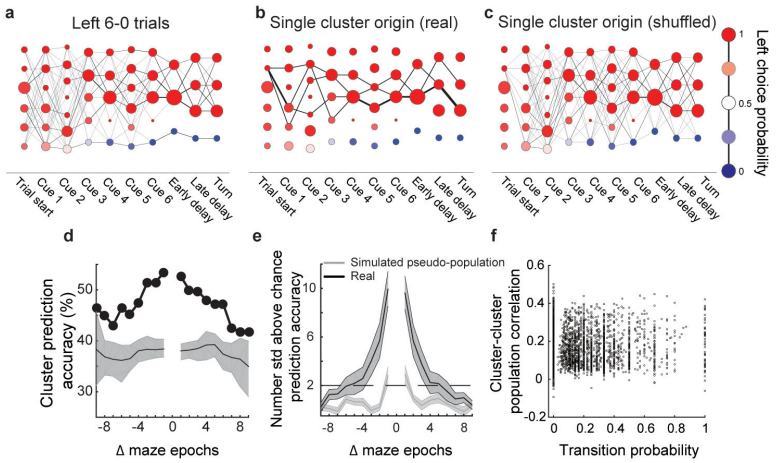

Figure 4. Long timescale temporal structure in PPC activity

a, Example transition matrix constructed only from left 6-0 (both correct and error) trials in a single dataset. 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 (Methods M.4.7.1). 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. Error bars 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 datasets (n = 11). To combine across datasets with different chance classifier performance, accuracies were converted into the number of standard deviations 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 to preserve each neuron's activity within the trial (Methods M.4.7.2). Error bars represent mean ± s.e.m. across datasets. **f,** Relationship between the population activity correlation of clusters in adjacent epochs and the transition probability between them. Transitions were not more likely between clusters with more similar population activity patterns. r = 0.02, p > 0.05, n = 3967 consecutive cluster pairs.

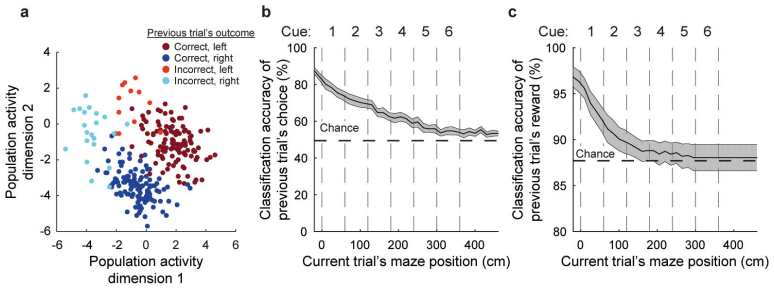

Figure 5. Neuronal population activity in the current trial reflects the previous trial's choice and outcome

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 is one trial. **b-c,** Both the previous trial's choice (**b**) and if the previous trial was rewarded (**c**) could be classified based on the population activity. Independent SVM classifiers were trained and tested at each maze position. Error bars represent mean ± s.e.m. across datasets (n = 11).

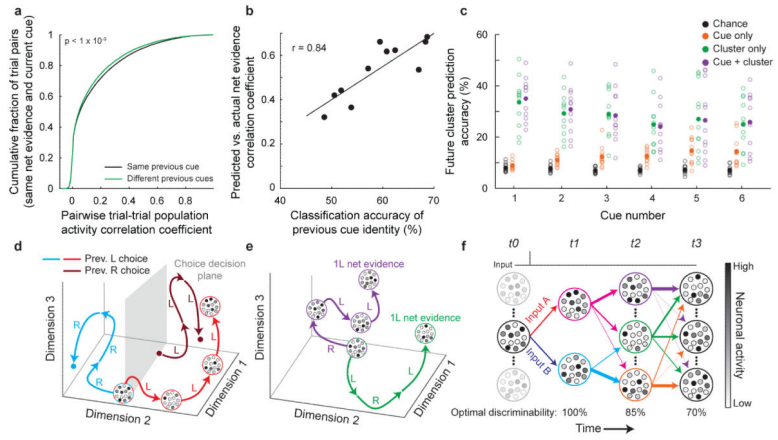

Figure 6. Analysis of neuronal activity related to evidence accumulation

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 (e.g. *LRLXXX* vs. *RLLXXX* trials at cue 3). $p < 10^{-9}$, two-sample KS test, n = 11 datasets. **b,** For each dataset, 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 datasets. **c,** For a single trial at a given epoch, the accuracy of predicting the next epoch's cluster identity for that trial based on chance (black), the evidence cue type only (orange), the current cluster identity (green), or both (purple). Individual datasets (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 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 red trajectory). Despite multiple activity patterns for choice, a decision plane (gray) can be drawn which separates activity patterns according to a given variable. Different decision planes can exist for other variables (e.g. previous trial's choice). **e,** Schematic depicting that trials with the same starting point and net evidence, but different sequences of cues, will take different paths through activity space, resulting in multiple, distinguishable activity patterns. **f,** Schematic demonstrating that 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 *t1* are most likely to transition to the orange activity pattern at *t2*, 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 *t2* as a result of these non-random transitions. The thickness of each arrow indicates the transition probability. The transitions between *t0* and *t1* 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.