



DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Population dynamics in parietal cortex during evidence accumulation for decision-making

The Harvard community has made this article openly available.
[Please share](#) how this access benefits you. Your story matters.

Citation	Morcos, Ari Simon. 2016. Population dynamics in parietal cortex during evidence accumulation for decision-making. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.
Accessed	May 6, 2018 9:54:22 AM EDT
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493459
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

(Article begins on next page)

Population dynamics in parietal cortex during evidence accumulation for decision-making

A dissertation presented

by

Ari Simon Morcos

to

The Division of Medical Sciences

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Neurobiology

Harvard University

Cambridge, Massachusetts

April 2016

©2016 Ari Simon Morcos

All rights reserved.

Population dynamics in parietal cortex during evidence accumulation for decision-making

Abstract

Cortical circuits combine new inputs with ongoing activity during a variety of behaviors, including evidence accumulation during decision-making. However, the neural circuit mechanisms underlying how populations of neurons perform the computations necessary for this process and the dynamics which govern how neuronal populations change from moment-to-moment during evidence accumulation remain unclear. Here, we trained mice to perform several novel virtual-navigation decision tasks, including one which requires the accumulation of multiple, discrete evidence cues. As mice accumulated evidence, the posterior parietal cortex (PPC) transitioned between distinguishable and largely uncorrelated activity patterns, often involving mostly different sets of active neurons from moment-to-moment. These activity patterns contained task-relevant information distributed across the neuronal population. Because animals make decisions on single trials, we chose to analyze these activity patterns on a trial-by-trial basis. As single trials unfolded, each event — whether a new evidence cue or a behavioral choice — modified the dynamics of the PPC for seconds, even across trials. These events did not change the tonic activity of a specific set of neurons; rather, each event altered the probabilities that govern how one activity pattern transitions to the next, constraining the possible future patterns of activity. Thus, representations of ongoing events were influenced both by the sequence of previous evidence cues within the current trial and by the outcome of the previous trial, thereby generating multiple distinguishable activity patterns for the same level of accumulated evidence. These observations suggest that evidence accumulation

does not rely upon the explicit competition between groups of neurons (as would be predicted by winner-take-all models), but instead reflects dynamical properties of the PPC that may instantiate a form of short-term memory consistent with reservoir computing.

Table of contents

Abstract	iii
List of figures	ix
Acknowledgments	x
1 General introduction and background	1
1.1 Drift-diffusion models	3
1.2 Single neuron correlates of evidence accumulation for decision-making	4
1.3 Sensory, motor, or both?	6
1.4 Heterogeneity across neurons and variability across trials	8
1.5 Transient or persistent activity dynamics?	10
1.6 Neural trajectories and methods for analyzing high-dimensional neural data	12
1.7 Neural algorithms for memory and decision-making	14
1.7.1 Winner-take-all models	14
1.7.2 Reservoir computing	15
1.8 Methods for recording from multiple neurons simultaneously	18
1.9 A virtual reality paradigm for mouse behavior	20
2 Development of novel virtual-navigation decision-making behaviors	22
2.1 Introduction	22
2.2 A virtual reality system for mouse behavior	23
2.2.1 Detailed description	24
2.3 Fixed association evidence accumulation task	25
2.3.1 Task description	25
2.3.2 Training procedure	26

2.3.3 Behavioral characterization	29
2.4 Delayed-match-to-sample task	33
2.4.1 Task description	33
2.4.2 Training procedure	36
2.4.3 Behavioral characterization	40
2.5 Delayed-match-to-sample evidence accumulation task	41
2.5.1 Task description	41
2.5.2 Training procedure	42
2.5.3 Behavioral characterization	44
2.6 Behavioral analysis suite	46
2.7 Thoughts on task design	48
3 History-dependent variability in population dynamics during evidence accumulation in cortex	51
3.1 Introduction	51
3.2 Results	53
3.2.1 Single neuron responses during evidence accumulation	55
3.2.2 Task-relevant information is distributed across neuronal populations	56
3.2.3 Clustering neuronal activity patterns across trials	63
3.2.4 Temporally structured trial-trial variability	72
3.2.5 A memory of past events is maintained in PPC population activity over seconds and across trials	80
3.2.6 A novel model for evidence accumulation based on history-dependent dynamics	82
3.3 Acknowledgements	90
3.4 Author contributions	90
4 Discussion and future experiments	91
4.1 History-dependent population dynamics	91
4.2 Importance of single-trial population analyses	93
4.3 Inconsistencies with winner-take-all models	94

4.4	A general rule for dynamics in the PPC	96
4.5	Developing neuronal implementations of history-dependent dynamics	98
4.6	Additional experimental tests	99
4.7	Recreating history-dependent dynamics <i>in silico</i> with recurrent neural networks	102
5	Methods	105
5.1	Subjects	105
5.2	Imaging	106
5.2.1	Surgical procedure	106
5.2.2	Two-photon microscope design	107
5.2.3	Imaging data acquisition	108
5.2.4	Pre-processing of imaging data	108
5.3	Data analysis	109
5.3.1	General analysis procedures	109
5.3.2	Choice selectivity for individual neurons	113
5.4	Classifiers (without clustering)	113
5.4.1	General procedures	113
5.4.2	Two-class classifications	114
5.4.3	Multi-class classifications	115
5.4.4	Classifiers built by adding-in subsets of neurons	116
5.4.5	Classification of cue sequences	116
5.5	Analysis of activity in high-dimensional state space	117
5.5.1	Factor analysis	117
5.5.2	Relationship between pairwise trial-trial distances before and after cue presentation	117
5.5.3	Relationship between pairwise trial-trial distances and population activity vectors	118
5.6	Clustering methods	119
5.6.1	Pre-processing for clustering	119
5.6.2	Clustering via affinity propagation	120
5.6.3	Transition probabilities between clusters	121

5.6.4	Clustering based on all time points together (rather than epoch-by-epoch clustering)	122
5.7	Classifiers based on activity in cluster-space	122
5.7.1	Classification of the cluster identity at past and future epochs based on the cluster identity at the current epoch	122
5.7.2	Classification of past and future cluster identities with simulated pseudo-populations	124
5.7.3	Classification of past and future cluster identities based on behavioral data	124
5.7.4	Classification of behavioral errors based on cluster identity	125
5.8	Analysis of the overlap of active neurons across clusters	126
	References	140

Listing of figures

2.1 Schematic of virtual reality behavior system.	23
2.2 A fixed association evidence accumulation task in virtual reality.	25
2.3 Behavioral training procedure for the fixed association evidence accumulation task	27
2.4 Behavioral characterization of the fixed association evidence accumulation task.	30
2.5 A delayed-match-to-sample task in virtual reality.	34
2.6 Behavioral training procedure for the delayed-match-to-sample task	36
2.7 A delayed-match-to-sample evidence accumulation task in virtual reality. . .	41
2.8 Behavioral training procedure for the delayed-match-to-sample evidence accumulation task	43
2.9 Behavioral characterization of the delayed-match-to-sample evidence accumulation task.	45
2.10 Behavioral analysis suite.	47
3.1 Example imaging field of view and activity traces.	54
3.2 Summary of datasets analyzed.	54
3.3 Distributed representation of task-relevant information across PPC neurons. .	57
3.3 (Continued)	58
3.4 Mean population activity patterns in PPC for all cells and selective cells. . . .	59
3.5 Analyses of single neuron- and population-level representations of task-relevant features.	62
3.6 Clustering neuronal activity across trials to reveals trial-to-trial variability. . .	65
3.6 (Continued)	66
3.7 Characterization of behavioral and neuronal patterns across clusters.	69

3.7 (Continued)	70
3.8 Visualizations of neuronal activity across clusters.	73
3.9 Long timescale temporal structure in PPC activity.	75
3.10 Influence of the current population activity pattern on future population activity patterns and the change in population activity in n-dimensional activity space.	76
3.11 Contribution of behavioral variability to temporally structured trial-trial variability.	78
3.12 Neuronal population activity in the current trial reflects the previous trial's choice and outcome.	81
3.13 Contribution of behavioral variability to classification of the previous trial's outcome.	83
3.14 Visualizations of trial trajectories.	83
3.15 Analysis of neuronal activity related to evidence accumulation.	87
3.15 (Continued)	88
5.1 Main results re-analyzed using $\Delta F/F$ values.	111
5.1 (Continued)	112

Acknowledgments

This dissertation would not be possible without the support and advice of my mentors, friends, and family. First and foremost, I owe a tremendous debt of gratitude to my dissertation advisor, Chris Harvey. When I first considered choosing Chris as my advisor, I worried about joining a brand new lab with no previous students or post-docs. Over the last five years, Chris has time and time again shown me that those worries were unfounded. I have been consistently impressed with his dedication, clarity of thought, and scientific insight. It has been a privilege to learn from him, and I will miss his advice and mentorship.

I also want to thank all the members of the Harvey lab for our many discussions about science and life, and for their constant advice and support. Our lab environment has truly felt like one where time and advice are given without thought of return, where ideas are discussed freely, and where criticism is always constructive. So, to Selmaan Chettih, Laura Driscoll, Shin Kira, Matthias Minderer, Noah Pettit, Caroline Runyan, and Alice Wang, I thank you. I am a better scientist because of you.

I would also like to thank all the members of the Program in Neuroscience for creating an environment which always felt challenging, yet secure. In particular, I would like

acknowledge Karen Harmin, who in addition to ensuring that the trains run on time, has also served as an invaluable friend and mentor throughout my time here.

I thank the members of my Dissertation Advisory Committee — Mark Andermann, John Assad, and Nao Uchida — for their thoughtful advice and feedback. I also thank my Thesis Examination Committee — Jan Drugowitsch, Mehrdad Jazayeri, and Venkatesh Murthy — for taking the time to provide feedback on this dissertation.

To all of my friends who have been with me during my time here: thank you. Thank you for complaining with me, for helping me with manuscripts and talks, and teaching me about football (baseball's still better, but football's ok, I guess). In particular, I thank Alex Trott, who has been the best lunch buddy I could have asked for.

And last, but most certainly not least, I thank my girlfriend, Julia Caine, my dog, Dash, and my family for their endless love and support.

Chapter 1

General introduction and background

Each of us makes hundreds of decisions every day, and these decisions often require the combination of incoming sensory evidence with an internal representation of the world. For example, in order to decide whether to turn left or right at an upcoming intersection, one must integrate incoming sensory information about the world, such as street names or landmarks, with internally represented information, such as directions to the destination or a memory of the path taken so far.

The accumulation of evidence for decision-making is essential to survival. Many of the decisions that we and other animals make rely on this process by requiring neuronal circuits to integrate novel inputs with ongoing activity. Behaviorally, evidence accumulation for decision-making has been extensively characterized and modeled. However, the mechanisms neuronal circuits use to perform the computations underlying this process remain poorly understood. We are interested in understanding how cortical circuits perform these computations. In this thesis, we have developed a variety of new behavioral and analysis techniques to probe the population dynamics underlying evidence

accumulation for decision-making. Using these methods, we sought to understand the rules that govern how populations of neurons change from moment-to-moment as evidence is accumulated, and how these dynamics contribute to the computations underlying perceptual decision-making.

This chapter provides a broad overview of previous work to characterize evidence accumulation for decision-making on both a behavioral and neuronal level, as well as a discussion of several key methodological developments. Chapter 2 explains the training protocols and characterization of three novel navigation-based decision tasks in virtual reality: one which requires mice to accumulate discrete pieces of evidence, one with a delayed-match-to-sample design to dissociate activity related to sensory information from that related to motor planning, and one which combines the two. Chapter 3 includes an extensive description of the population dynamics in the PPC while mice perform an evidence accumulation task. This chapter also includes the description of several new analysis techniques to examine how neuronal activity changes from moment-to-moment on single trials. Finally, Chapter 4 includes a discussion of the results, including how they relate to winner-take-all and reservoir computing models of decision-making, and proposes a set of future experiments to further elucidate the population dynamics in parietal cortex during evidence accumulation.

1.1 Drift-diffusion models

The accumulation of evidence over time for decision-making has been described by so-called ‘drift-diffusion’ models (Ratcliff & Smith, 2004; Smith, 2000). In these models, a decision variable (representing the sum of evidence acquired) begins a trial with a neutral value, and drifts slowly as evidence is accumulated until a threshold or ‘decision bound’ is reached, at which point a decision is made. One prediction of such models is a tradeoff between the speed and accuracy of perceptual decisions. For example, in a system with decision bounds close to the system’s starting point, the decision variable would reach one of the decision bounds rapidly, resulting in a speedy decision, but one that is based on less evidence and therefore more susceptible to random fluctuations. Alternatively, moving the decision bound away from the system’s starting point would increase the probability of an accurate decision, but also increase the decision time by requiring more evidence to be accumulated. This speed-accuracy tradeoff has been observed in humans, nonhuman primates, and rodents (Rinberg et al., 2006; Roitman & Shadlen, 2002; Wickelgren, 1977), though it is absent in some behavioral contexts (Uchida & Mainen, 2003). Drift-diffusion models also predict that increasing the strength of the evidence increases the slope of the accumulation curve, resulting in faster decisions while decreasing evidence strength reduces the slope and results in slower decisions.

Drift-diffusion and related models can also be used to effectively model behavior (Brunton et al., 2013; Kira et al., 2015; Drugowitsch et al., 2012; Shadlen & Kiani, 2013).

For example, one recent study used a drift-diffusion framework to model the behavior of rats and humans accumulating auditory clicks (Brunton et al., 2013). Surprisingly, the authors found that behavioral variability did not originate from noise in the accumulator (e.g., the memory of previous sensory stimuli); rather variability seemed to originate purely from the presentation of sensory evidence. This result suggests that the brain acts as an optimal accumulator and demonstrates the power of such purely behavioral models to describe the dynamics of decision-making processes.

1.2 Single neuron correlates of evidence accumulation for decision-making

Neural correlates of evidence accumulation for decision-making have been found throughout the brain, but the neural computations underlying this process remain unclear (Shadlen & Newsome, 1996; Gold & Shadlen, 2000; Yang & Shadlen, 2007; Hanks et al., 2015; Britten & Shadlen, 1992; Mante et al., 2013; Horwitz & Newsome, 1999; Gold & Shadlen, 2007; Roitman & Shadlen, 2002). The posterior parietal cortex (PPC) has emerged as a prime candidate region for this process, in large part due to its extensive connections to both sensory and motor regions (Blatt et al., 1990) and a variety of studies in non-human primates (Bisley & Goldberg, 2010; Gold & Shadlen, 2007; Andersen & Cui, 2009). Much of this work has focused on the lateral intraparietal area (LIP), a region within the PPC which exhibits substantial saccade-related activity (Bisley & Goldberg, 2010; Gold & Shadlen, 2007). The majority of such studies have used variants of

two-alternative forced choice (2AFC) random-dot motion tasks, in which monkeys view a stimulus with varying amounts of net motion due to the biased motion of a set of rapidly moving dots. Following a delay of several seconds, monkeys are instructed to make a saccade toward one of two targets, one of which is within the receptive field of a recorded LIP neuron (Shadlen & Newsome, 1998, 1996). In response to their preferred direction of motion during this task, LIP neurons increase their firing rate steadily in a ramping fashion throughout stimulus presentation and maintain a high firing rate throughout the delay period before returning to a baseline firing rate after the saccade is made (Shadlen & Newsome, 1996). Similar results have also been observed in the PPC of rats performing evidence accumulation tasks (Hanks et al., 2015) as well as in other brain regions, including the superior colliculus and the prefrontal cortex (Mante et al., 2013; Horwitz & Newsome, 1999).

While causal experiments have been difficult to perform, several studies have shown that the PPC influences decision-making. For example, microstimulation of LIP during stimulus presentation has induced biases in both the direction and the reaction time of the saccade that were correlated with the total accumulated evidence, implying a causal role for LIP in evidence accumulation during decision-making (Gold & Shadlen, 2000; Hanks et al., 2006). In one study, monkeys performed an evidence-accumulation task in which they viewed several discrete pieces of evidence (different shapes), each of which was associated with a fixed reward probability on one of two targets (Yang & Shadlen, 2007). To make accurate perceptual decisions and maximize their reward probability, the monkeys had to

integrate the multiple pieces of evidence. During this task, LIP neurons exhibited changes of activity correlated not only with which target each piece of evidence preferred, but with the magnitude of the reward probability change as well. These results suggest that neurons in LIP not only integrate evidence over time, but also calculate a sum of probabilities favoring possible alternatives to best compute an accurate decision.

These studies and others have led to hypotheses that LIP accumulates evidence and contains the decision variable described in decision-making models (Mazurek, 2003). One key prediction of these models is that stronger evidence should increase the slope of the decision variable curve and result in a faster decision. To test this, a reaction time version of the random-dot motion task was developed in which monkeys can respond at any time, thereby providing a rough estimate of decision speed (Roitman & Shadlen, 2002). During this task, the slope of the LIP neurons' increase in firing rate is steeper during trials in which the monkey's reaction time is lower and in which the motion coherence is higher (stronger evidence), suggesting that the faster reaction time is due to stronger evidence causing the decision variable to reach the decision bound more rapidly. Activity of LIP neurons during a reaction time version of the discrete shapes task was also consistent with the predictions of drift-diffusion models (Kira et al., 2015).

1.3 Sensory, motor, or both?

While the work described above has effectively demonstrated that LIP contains activity correlated with a developing decision, it has been difficult to assign an exact function to

these changes in firing. In particular, the PPC has been found to contain activity correlated with both movement intention (Andersen & Cui, 2009; Buneo & Andersen, 2006; Cui & Andersen, 2007; Quian Quiroga, 2006) and spatial attention (Bisley & Goldberg, 2010; Bisley, 2003; Ipata et al., 2006; Balan, 2006; Kusunoki et al., 2000). In the movement intention view, activity in the PPC primarily predicts motor actions, rather than the visual cues. In contrast, the spatial attention view proposes that activity in the PPC is best explained by attention to the visual stimuli, independent of the resulting motor action.

This ambiguity is in large part due to the fixed relationship between the decision and the motor plan in the majority of previous studies (e.g. in the random-dot motion task, net motion to the left always requires a saccade to the left for reward) (Freedman & Assad, 2011). To test whether the PPC can represent an abstract decision, several groups have designed behavioral tasks with a delayed-match-to-sample design in which there is a variable relationship between the motor response and the decision (Freedman & Assad, 2006; Fitzgerald et al., 2011; Bennur & Gold, 2011). Importantly, these tasks feature a period of time in which the animal must make a perceptual decision without knowing how to report the decision, allowing for the isolation of decision related neuronal activity from activity related to motor preparation. These studies have demonstrated that LIP contains activity correlated with categorization of visual stimuli independent of motor actions (Freedman & Assad, 2006; Fitzgerald et al., 2011). Additionally, during a modified version of the random-dot motion task with a delayed-match-to-sample design, activity in LIP correlated with a developing, abstract perceptual decision similar to that observed in

the traditional random-dot motion task (Bennur & Gold, 2011). Taken together, these studies suggest that, while it may play a role in motor planning as well, LIP is involved in performing computations necessary for abstract decision-making.

1.4 Heterogeneity across neurons and variability across trials

Due primarily to technical limitations, the majority of the results described above were derived from pseudo-populations in which each cell's activity was averaged across all trials with similar stimuli and/or choices. In many cases, activity was further averaged across cells to construct an activity trace averaged across both neurons and trials. Additionally, such datasets are often constructed from a biased sampling cells, with only those neurons that exhibited persistent activity during a delayed saccade task included. While these techniques have revealed important features of the neuronal activity, they also obscure potentially meaningful heterogeneity across neurons and trials. In fact, recent studies that have analyzed the activity of individual neurons during decision tasks have found that the dynamics of single neuron responses are highly diverse (Meister et al., 2013; Park et al., 2014; Jun et al., 2010; Raposo et al., 2014; Mante et al., 2013; Rigotti et al., 2013). Many of these neurons exhibited complex response dynamics. For example, some neurons responded to different features of the stimulus at different times in the task, while many others had responses that were seemingly unrelated to task features. What role, if any, such heterogeneity across neuronal populations may play in the computations underlying decision tasks, however, remains unclear.

Much of the reason that previous studies of neuronal activity during decision-making have relied upon pseudo-populations has been due to the difficulty of recording from multiple neurons simultaneously. However, the emergence of new technologies, including multi-electrode arrays and optical imaging techniques, has allowed for many studies to probe the activity of neuronal populations on single trials. These studies have revealed that the activity of many neurons is weakly, but positively correlated on a trial-to-trial basis, even after removing correlations due to task differences (Zohary et al., 1994). These correlations, often termed ‘noise correlations’, have been observed across cortex in both anaesthetized and behaving animals (Cohen & Kohn, 2011). From the perspective of optimal coding of information, the role of these correlations can be difficult to parse (Averbeck et al., 2006). For example, if neurons with similar tunings to stimuli, correlated noise will make stimulus decoding more difficult (Zohary et al., 1994). In contrast, correlated noise between neurons with opposite stimulus tuning can aid stimulus decoding (Romo et al., 2003). These correlations have been found to be modulated by visual attention (Cohen & Maunsell, 2009), stimulus onset (Churchland et al., 2010), and associative learning (Jeanne et al., 2013), suggesting that correlated trial-trial variability may contribute to neuronal computation. As with neuronal heterogeneity, however, the structure of such inter-neuronal correlations as well as what role they might play in neuronal computation is still unknown.

1.5 Transient or persistent activity dynamics?

The majority of decision-making studies have found neurons to be active throughout the trial. However, a variety of recent studies have found neurons which are active only transiently, at specific times during a behavior (Harvey et al., 2012; Crowe et al., 2010; Pastalkova et al., 2008; Fujisawa et al., 2008). In one recent study, populations of neurons in the PPC were simultaneously recorded while mice performed a navigation-based decision task in virtual-reality (Harvey et al., 2012). During this task, neurons were active only briefly and at specific times in the trial (e.g., during the beginning of the delay), with a small set of neurons having clear choice selectivity. Across the population, the activity of individual neurons tiled the entire trial duration, leading to the description of these dynamics as ‘sequences’. Similar sequences of neuronal activation have been found in the hippocampus (Pastalkova et al., 2008) and prefrontal cortex of rodents (Fujisawa et al., 2008), as well as in LIP of non-human primates (Crowe et al., 2010). However, most of these studies analyzed the trial-averaged activity of neuronal subsets during relatively simple decision tasks and how the activity of neuronal populations may change from moment-to-moment on single trials during more complicated decision tasks involving evidence accumulation over time remains vague.

The discrepancy between the persistent activity traditionally reported and the transient activity patterns more recently observed could be caused by differences in task design, cortical layer, and model organism (i.e., non-human primates vs. rodents). However, each

of these possibilities seems unlikely to individually account for this difference. For example, while many of the studies that reported transient activity were performed in rodents (Harvey et al., 2012; Pastalkova et al., 2008; Fujisawa et al., 2008), transient activity has also been observed in non-human primates (Crowe et al., 2010) and in songbirds (Hahnloser et al., 2002). Moreover, neurons with ramping, persistent activity patterns have also been reported in rodents (Hanks et al., 2015; Raposo et al., 2014). Many of the rodent studies were also performed while animals performed navigation-based tasks (Harvey et al., 2012; Pastalkova et al., 2008; Fujisawa et al., 2008), yet transient activity patterns have also been observed during stationary tasks (Crowe et al., 2010). Together, these results suggest that both transient and persistent activity dynamics are likely present in the brain.

How these sequential activity patterns are generated by neuronal circuits remains unclear, though a variety of models have proven capable of generating such dynamics. In the simplest models, neurons are connected in a feedforward chain, such that neurons which fire at a given time point in the sequence preferentially activate the neurons at the next time point. Theoretical work has demonstrated that, if properly tuned, such chains can sustain a memory of a stimulus for a prolonged time period (Goldman, 2009). These networks, however, require a highly constrained network architecture. In contrast, recurrently connected networks subjected to an appropriate learning rule have also been shown to be capable of generating transient dynamics (Rajan et al., 2016; Klampfl & Maass, 2013). While these models often make clear predictions about the connectivity of

neurons with different activity patterns, these predictions have proved challenging to test experimentally.

1.6 Neural trajectories and methods for analyzing high-dimensional neural data

These studies have led to the conceptualization of neuronal population activity at a given time as a point in the n -dimensional space defined by the activity of each recorded neuron (where n is the number of neurons) (Harvey et al., 2012; Mazor & Laurent, 2005; Briggman, 2005; Raposo et al., 2014; Churchland et al., 2012). As neuronal population activity changes from moment-to-moment within a trial, it moves through this space, creating a trial-specific neuronal ‘trajectory’. The path taken by these neuronal trajectories, as well as the location of the activity pattern in the n -dimensional space, may represent task-relevant information. Importantly, analyses based on the concept of neural trajectories allow for the analysis of the moment-to-moment changes in neuronal populations on single trials. When animals perform behaviors, they must do so based on the single trial activity of neuronal populations. However, the majority of studies of decision-making and other behaviors have analyzed neuronal activity as averages across neurons and trials, potentially obscuring valuable information about cortical dynamics that may be present in the correlations between neurons. The methods underlying these analyses, while challenging and comparatively new, can also scale to accommodate large populations of neurons. This will become increasingly important as the simultaneous

recording of larger and larger numbers of neurons becomes feasible.

However, the power of these analysis comes at the cost of complexity. As the number of dimensions in the data increases (e.g., as more neurons are added), it becomes more and more difficult to understand the dynamics of the system. As an intuitive example of this issue, it is trivial to visualize a one-, two-, or three-dimensional system. A four-dimensional system can also be visualized, though it is more difficult (as a movie of how a three-dimensional system changes over time). However, the raw visualization of any system with more than four dimensions becomes extremely challenging. Additionally, the number of possible configurations of a system (e.g., the number of possible neuronal activity patterns) increases exponentially with the number of dimensions. This challenge has been termed the ‘curse of dimensionality’ (Bellman, 1961).

A variety of techniques have been developed to interpret high-dimensional data. In neuroscience, dimensionality reduction methods have been especially popular. These methods attempt to find structure in the data that can be represented with fewer dimensions than there are neurons. They include unsupervised techniques such as principal components analysis (PCA) and factor analysis (Murphy, 2012; Harvey et al., 2012; Briggman, 2005; Mazor & Laurent, 2005), as well as targeted techniques which attempt to find dimensions relevant for specific task variables (Mante et al., 2013). While dimensionality reduction is an essential tool, more powerful techniques will need to be developed and applied to neuronal data to decipher neuronal population dynamics.

1.7 Neural algorithms for memory and decision-making

1.7.1 Winner-take-all models

A variety of computational models have been developed to explain the neural activity dynamics observed during evidence accumulation and other decision tasks. However, so-called ‘winner-take-all’ models have become the most widespread (Wong & Wang, 2006; Wang, 2012; Machens, 2005; Wang, 2002). In these models, distinct, recurrently connected pools of neurons each receive inputs in favor of one of several behavioral alternatives. Each pool of neurons indirectly inhibits each of the other pools. As a result, as one pool receives input in favor of its preferred alternative, its firing rate increases, increasing both its recurrent excitatory drive and increasing the inhibition of competing pools of neurons. Inhibition of the competing neuronal pools relieves the inhibition of those pools onto the preferred pool, further increasing its firing rate. This process results in a positive feedback loop, eventually leaving only one pool of neurons active. Hence, there can only be one ‘winner’. Winner-take-all models have successfully explained a large variety of behavioral and neuronal results (Wang, 2008).

However, these models have several disadvantages. First, they require a highly constrained network connectivity, with precise connectivity patterns of each input to each pool of neurons as well as within and across pools. How such a constrained network architecture might be learned, however, especially through relatively sparse rewards in naturalistic environments, remains unknown. Second, because of connectivity constraints,

winner-take all models are specific to individual tasks. A unique winner-take-all circuit would therefore be necessary for each unique task. While a generalized winner-take-all circuit in which inputs for different tasks are routed to the same computational circuit might be possible, such networks have not been thoroughly explored. Third, these models generalize poorly to decision-making with multiple alternatives. While winner-take-all models have been adapted to decisions with three and four alternatives, they do so at the cost of substantially increased complexity (Churchland et al., 2008; Churchland & Ditterich, 2012; Niwa & Ditterich, 2008). Finally, winner-take-all models predict that neurons involved in the decision-making process will have highly homogeneous activity dynamics, both across neurons and across trials. However, a variety of studies have found heterogeneous and highly variable neurons that nevertheless represent decision-related information (Meister et al., 2013; Park et al., 2014; Jun et al., 2010; Raposo et al., 2014; Mante et al., 2013; Rigotti et al., 2013).

1.7.2 Reservoir computing

The computations necessary for short-term memory and decision-making may also be performed by networks exemplified by chaotic and complicated dynamics. In reservoir networks (also termed liquid state machines or echo-state networks), the network activity is determined as a function of both incoming sensory inputs and the ongoing activity of the network (Jaeger & Haas, 2004; Maass et al., 2002; Buonomano & Maass, 2009; Buonomano & Merzenich, 1995; Verstraeten et al., 2007). As a result, the effect of the same input on the network activity may be different depending on the network activity

present immediately prior to the input's presentation.

One intuitive way to conceptualize such networks is to imagine the surface of a pond. If a black pebble is dropped from a given height, a specific set of ripples will emerge on the surface of the pond. If the pond is allowed to return to equilibrium and the same pebble is dropped from the same location, the same set of ripples will occur. If, however, a red pebble is dropped one second before the black pebble, the ripples occurring as a result of the red pebble will interact in some complicated fashion with the ripples created by the black pebble, generating a unique pattern of ripples. Importantly, this pattern depends on a large variety of factors, including the features of the pebbles themselves, the locations from which they were dropped, and the time interval between them. All of this information might be represented simultaneously in the pattern of ripples generated. In the context of a neural circuit, the surface of the pond would represent the neural circuit itself, the pebbles inputs to the network, and the ripples the precise pattern of activity present in the network at a given time.

Reservoir networks therefore contain a 'reservoir' in which information about inputs (e.g., sensory stimuli) 'echoes' for some period of time as it gradually decays. These models do not require that the various activity patterns which represent an input look similar to one another; in contrast, these activity patterns may be highly different, with the differences caused by the representation of information related to other inputs or by differences in time. Information about a given input would therefore not be represented explicitly by the activity of individual neurons (e.g., high activity means input A, while low activity means

input B). Instead, information would be represented implicitly by the pattern of activity in the high-dimensional activity space. This information could be read out by downstream networks which project the high-dimensional activity onto a single dimension which best separates specific inputs (Hoerzer et al., 2014; Jaeger & Haas, 2004; Sussillo & Abbott, 2009; Buonomano & Maass, 2009; Maass et al., 2002; Natschläger & Maass, 2005). Reservoir networks could therefore represent multiple streams of information simultaneously. Additionally, such networks have been shown to be capable of producing transient, sequential dynamics as have been observed previously (Rajan et al., 2016; Klampfl & Maass, 2013).

Reservoir computing makes several testable experimental predictions. First, because the ongoing network activity influences how inputs are represented at later time points, trial-trial variability in network activity should be predictable based on the variability at earlier time points in a trial. Second, events in previous trials should influence the representation of events in the current trial when the time interval between trials is sufficiently short. A variety of studies, across cortical areas, have demonstrated that events from the previous trial can be decoded from neuronal activity during the current trial (Bernacchia et al., 2011; Donahue & Lee, 2015; Seo et al., 2007; Seo & Lee, 2007; Nikolić et al., 2009; Klampfl et al., 2012; Seo et al., 2009; Sugrue et al., 2004; Chaudhuri et al., 2015; Murray et al., 2014). Third, in this framework, the representation of inputs would occur generally, such that a memory of all inputs is maintained, rather than specific inputs being privileged. Network activity in the reservoir should therefore be independent

of how it is read out. In a recent study, monkeys were trained to perform a variant of the random-dot motion task in which the stimulus varied along two dimensions independently: motion and color (Mante et al., 2013). On alternating trial blocks, monkeys had to distinguish either the coherent motion (as in traditional random-dot motion tasks) or the coherent color. Recording from neuronal populations in prefrontal cortex, the authors found that both color and motion were represented simultaneously, with only subtle differences in their representation during different trial blocks, as would be predicted by reservoir computing. On different trial blocks, different readout networks may have been recruited to discriminate either color or motion. Finally, because multiple sources of information would be represented simultaneously in a general-purpose manner, information in reservoir networks is likely to be distributed implicitly across populations of neurons, rather than explicitly across sparse subsets of neurons. Sparse representation of information could still be consistent with reservoir networks; such a representation would, however, likely require a more constrained network architecture.

1.8 Methods for recording from multiple neurons simultaneously

Most studies of the PPC have been performed using single-cell recordings. However, methods have been recently developed to record from large populations of neurons simultaneously behaving animals (Miller & Wilson, 2008). To isolate single units, sets of independently movable tetrodes have been implanted in rodents, allowing for the recording of up to 20 units simultaneously from freely moving rodents (Kepecs et al.,

2008; Davidson et al., 2009). Alternatively, multi-electrode arrays have been developed which can be used to record the activity of up to 96 isolated units simultaneously (Churchland et al., 2012). However, these recordings provide little anatomical information about the recorded units, recordings cannot be limited to cells of a specific subtype, and the same neurons can only rarely be recorded from over several recording sessions. Optical imaging techniques using two-photon laser scanning microscopy (2PLSM) and genetically encoded calcium indicators (GECIs) (Chen et al., 2013; Looger & Griesbeck, 2012) can overcome many of these limitations, allowing the simultaneous recording of the activity of hundreds of neurons (Dombeck et al., 2010; Harvey et al., 2012; Huber et al., 2012; Li et al., 2015). This technique provides two key advantages over other techniques for recording large populations of neurons in awake, behaving animals. First, imaging provides anatomical information about the cells recorded from, allowing experimenters to ask questions about anatomical segregation of neurons based on their activity patterns and for identification of specific cell types either in vivo or post-mortem. Second, because this technique can be performed with chronically implanted cranial windows, the same population of neurons can be recorded from over multiple recording sessions. Third, while single-unit electrophysiology often results in a higher sampling of task-related and high firing neurons, optical imaging provides a relatively unbiased sampling of neurons. Fourth, optical imaging allows for the sampling of a large number of cells from a local region, while traditional recording methods sample cells over a much larger region. One caveat, however, is that genetically encoded calcium

indicators exhibit a non-linear relationship between spike number and fluorescence change, making it difficult to determine the precise number of spikes fired by a neuron. There is, however, a clear general relationship between fluorescence and spiking activity (Tian et al., 2009; Akerboom et al., 2012; Dombbeck et al., 2010; Chen et al., 2013). Additionally, a variety of deconvolution algorithms to extract spiking information from the GECI signal, with increasing success (Vogelstein et al., 2010; Pnevmatikakis et al., 2016; Theis et al., 2014). Together, these features make optical imaging a powerful method for recording the simultaneous activity of large populations of neurons in vivo.

1.9 A virtual reality paradigm for mouse behavior

The majority of studies focused on decision-making have used head-restrained, stationary monkeys who respond by performing a saccade, a reach, or a button press (Gold & Shadlen, 2007; Bisley & Goldberg, 2010; Freedman & Assad, 2006, 2011; Buneo & Andersen, 2006). These tasks, however, are low throughput and cannot yet be paired with optical imaging techniques which enable the recording of activity from hundreds of neurons simultaneously. A virtual reality system for mice was recently developed which can overcome these technical limitations to understand microcircuit function during decision-making (Harvey et al., 2009). In this system, a head-restrained mouse is surrounded by a large screen on which a visual virtual environment is displayed. By running on a spherical treadmill, the mouse can navigate through the virtual environment – effectively allowing the mouse to run a virtual maze. Because the mouse’s head remains

stationary, this paradigm can be paired with advanced optical imaging techniques. This virtual reality system therefore allows mice to perform complex behavioral tasks employing navigation while advanced optical imaging techniques are used to record from large populations of neurons simultaneously.

Chapter 2

Development of novel virtual-navigation decision-making behaviors

2.1 Introduction

In order to understand how cortical circuits perform the computations underlying decision-making, we must be able to record from and manipulate these circuits during decision tasks. We therefore need to develop well-characterized and controlled behaviors which enable us to accurately attribute changes in neuronal activity to specific features of the decision-making process.

Here, we have developed three novel virtual reality (VR) decision tasks based on a T-maze. In the first task, mice are required to accumulate six discrete evidence cues to determine whether to turn into the left or right arm of the T-maze for a reward. In the second task, which has a delayed-match-to-sample design, mice are presented with only one cue, but the association between the cue and the turn direction can vary from trial to trial and is unknown to the mouse during cue presentation and during a delay period of several seconds. This

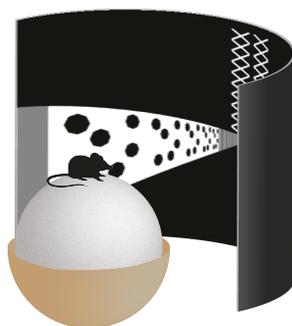


Figure 2.1 | Schematic of virtual reality behavior system.

task therefore allows for the dissociation of activity related to sensory information from that related to motor planning (Freedman & Assad, 2011). The third and most difficult task incorporates components of both previous tasks to create an evidence accumulation task with a delayed-match-to-sample design.

2.2 A virtual reality system for mouse behavior

To develop these behaviors, we have used a modified version of the previously developed VR system (Harvey et al., 2009, 2012; Dombeck et al., 2010). In this system, head-restrained mice are surrounded by a large, half-cylindrical screen on which a first-person view of a virtual environment is presented (Figure 2.1). By running on a spherical treadmill, mice can navigate through this environment, allowing them to perform behavioral tasks for a reward.

The VR system has several key advantages. First and most importantly, the use of virtual environments enables the rapid prototyping of task designs, because virtual environments can easily be changed, even within a behavioral session. Second, the comparatively small

amount of space required for VR system as well as relatively short training sessions (usually less than an hour) allow for the training of many mice in parallel. Third, because the contents environment are completely controlled by the experimenter, the only stimuli present in the virtual environment are those which are intentionally included. Finally, because mice are head-restrained during VR behaviors, this system can be paired with optical imaging (Harvey et al., 2012; Dombeck et al., 2010), enabling the recording and manipulation of large populations of neurons during behavior, as well as whole-cell electrophysiology (Harvey et al., 2009; Domnisoru et al., 2013).

2.2.1 Detailed description

First-person images of the virtual environment were back-projected onto a half-cylindrical screen with a diameter of 24 inches and a depth of 12 inches using a PicoP microprojector (MicroVision Inc.). Images were transformed to account for the shape of the screen. The spherical treadmill was a custom 8-inch ball made of open-cell Styrofoam foam. The spherical treadmill was supported by air to allow free rotation. Movement of the spherical treadmill was recorded using an optical sensor (Logitech MX518) positioned beneath the ball. Forward/backward translation in VR was controlled by changes in pitch (relative to the mouse's body), and rotation in VR was controlled by changes in roll (relative to the mouse's body). This is different from previous studies (Harvey et al., 2012, 2009; Dombeck et al., 2010), in which rotation in VR was controlled by changes in yaw rather than roll (relative to the mouse's body). This modification significantly improved the mouse's control of rotation in the virtual environment. The recorded behavioral

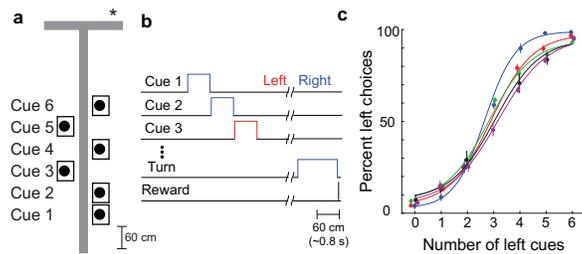


Figure 2.2 | A fixed association evidence accumulation task in virtual reality. **a**, Schematic of an example 2-4 right trial in a virtual T-maze. Asterisk marks the reward location. **b**, Sequence of trial events. **c**, Performance for five mice (mean \pm s.e.m, 7-12 sessions).

parameters were the mouse's position in the virtual environment (x/y position), the rotational velocity of the spherical treadmill (about the pitch and roll axes relative to the mouse's body), and the mouse's view angle in the environment. Virtual environments were built and run using the MATLAB-based software ViRMEn (Virtual Reality Mouse Engine) (Aronov & Tank, 2014).

2.3 Fixed association evidence accumulation task

2.3.1 Task description

We developed a navigation-based evidence accumulation task for head-restrained mice based on a T-maze in VR (Figure 2.2a-b). 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 seconds. To receive a reward, mice had to

determine if more cues were presented on the left or the right and, after a short stretch of maze without additional cues (90 cm, ~1 second), turn at the T-intersection toward the direction that had more cues. Between trials, the task difficulty was modulated by varying the difference between the number of left and right cues, or the net evidence. Within each trial, there were six total cues, with between zero and six cues presented on the left and the remainder on the right. For example, a trial with six cues on the left and no cues on the right (6-0 left, net evidence: 6 left) was easy, while a trial with two cues on the left and four cues on the right (2-4 right, net evidence: 2 right) was much more difficult. The sequence of cues was determined randomly for each trial of a given net evidence. To perform well on this task, mice must accumulate multiple, discrete pieces of evidence across long timescales (~5-6 seconds) and remember their choice during a short delay (~1 second). By accumulating multiple pieces of evidence, mice were able to perform this task at high accuracy stably for many days (Figure 2.2c).

2.3.2 Training procedure

Following at least two days of water restriction, mice began behavioral training. Behavioral sessions were performed six days per week and lasted 45-60 minutes. Mice received liquid rewards through a lick spout (4 μ l/reward, 10% sweetened condensed milk). Mice were trained to perform the evidence accumulation task over a series of eight mazes (Figure 2.3). In mazes 1-5, mice learned to steer in the virtual environment and to associate visual cues with a reward. In mazes 6-7, mice learned to remember the location of the cue during a brief delay period (~1 second). In maze 8, the initially continuous white cue on either the

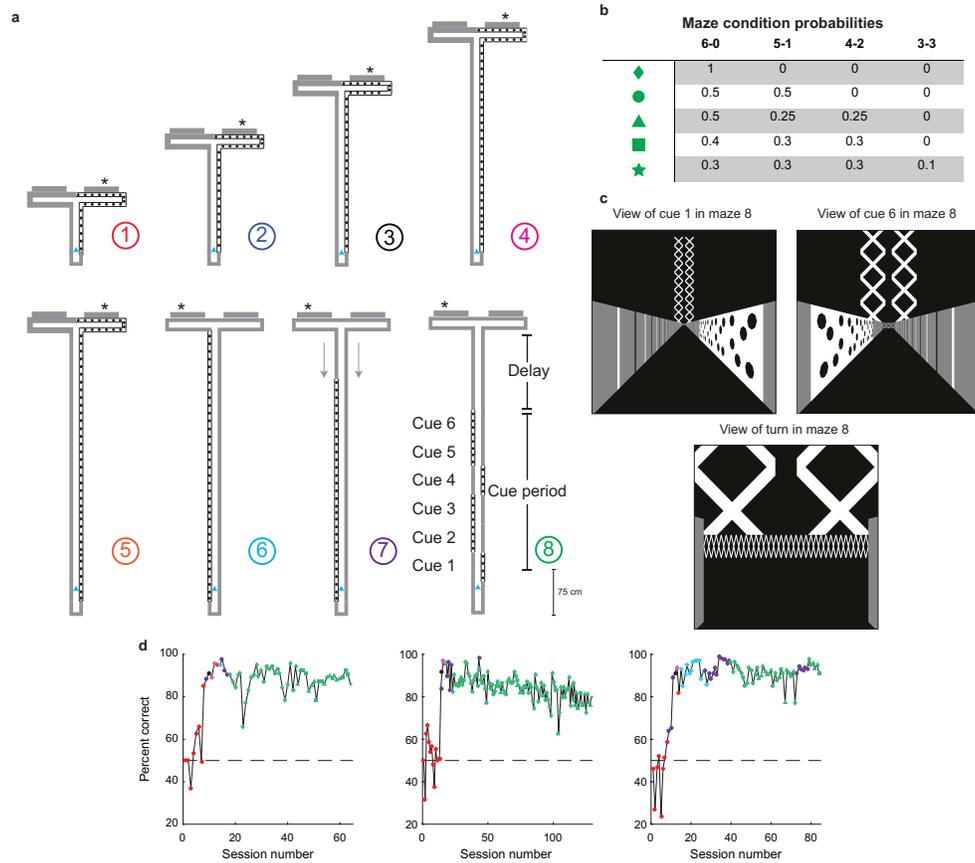


Figure 2.3 | Behavioral training procedure for the fixed association evidence accumulation task. a, Sequence of mazes used for behavioral training. Asterisks indicate reward location. Only some example mazes are shown (for example, right choice and not left choice maze in maze 1). **b,** Distribution 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**.

left or the right was divided into 6 discrete cues, all of which were on the same side (e.g. 6-0 left and 0-6 right trials). Cue visibility was locked to the mouse's position such that only one cue was visible at a time. There was no delay between the offset of one cue and the onset of the next cue. At this point, the probability of more difficult conditions (e.g. 5-1 trials) was gradually increased using the probability distributions in Figure 2.3b. Within each difficulty, the precise pattern of evidence was randomly determined before each trial. On 3-3 trials, the rewarded location was selected randomly. Trials were completed when the mouse turned into one of the arms of the T-maze. Following the completion of the trial, the screen changed to black for the duration of the inter-trial interval (ITI; 2 s for correct choice and 4 s for incorrect choice).

In some cases, mice developed biases such that they favored left or right choices. To discourage these biases, we implemented bias correction throughout training. Bias correction was not used during the imaging sessions described in Chapter 3. On each trial, we determined a probability that a trial would be a right choice trial. The probability was set to be equal to the fraction of left choices over the previous 20 trials, such that if the mouse had made many left or right choices previously then the opposite choice trial was likely selected. Once mice reached expert levels, prolonged biases rarely developed. 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 (Figure 2.3) 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. We found that the use of crutch trials was essential to establishing stable behavior both within a single session and across multiple sessions.

For mazes 1-3, the criterion for advancement to the next maze was the mouse's number of completed trials per minute, independent of performance (> 7 trials/min for advancement). Mice on mazes 4-5 were advanced to the next maze following one day of $\geq 90\%$ performance. Mice on mazes 6-7 were advanced following three days of $\geq 90\%$ performance. On maze 8, mice were advanced after three days of $\geq 85\%$ performance on easy trials (6-0 and 5-1 trials) and $\geq 65\%$ performance on hard trials (4-2 trials). At this stage of training, mice were sensitive to the rapid introduction of difficult trials, and performance on easy trials was closely monitored. If performance on easy trials fell below 80%, mice were moved to an easier distribution of maze trials. For this reason, we were only able to introduce 3-3 trials in some of the mice. Training in total required ~ 30 -60 daily training sessions (Figure 2.3d).

2.3.3 Behavioral characterization

Mice could have achieved intermediate performance on this task by only paying attention to a single cue. A mouse perfectly following this strategy would be expected to make a correct choice on 100% of 6-0 trials (because every cue matches the correct choice), 83% on 5-1 trials (because each cue has a $5/6$ chance of matching the correct choice), and 67% on 4-2 trials. The psychometric functions of each mouse would therefore be

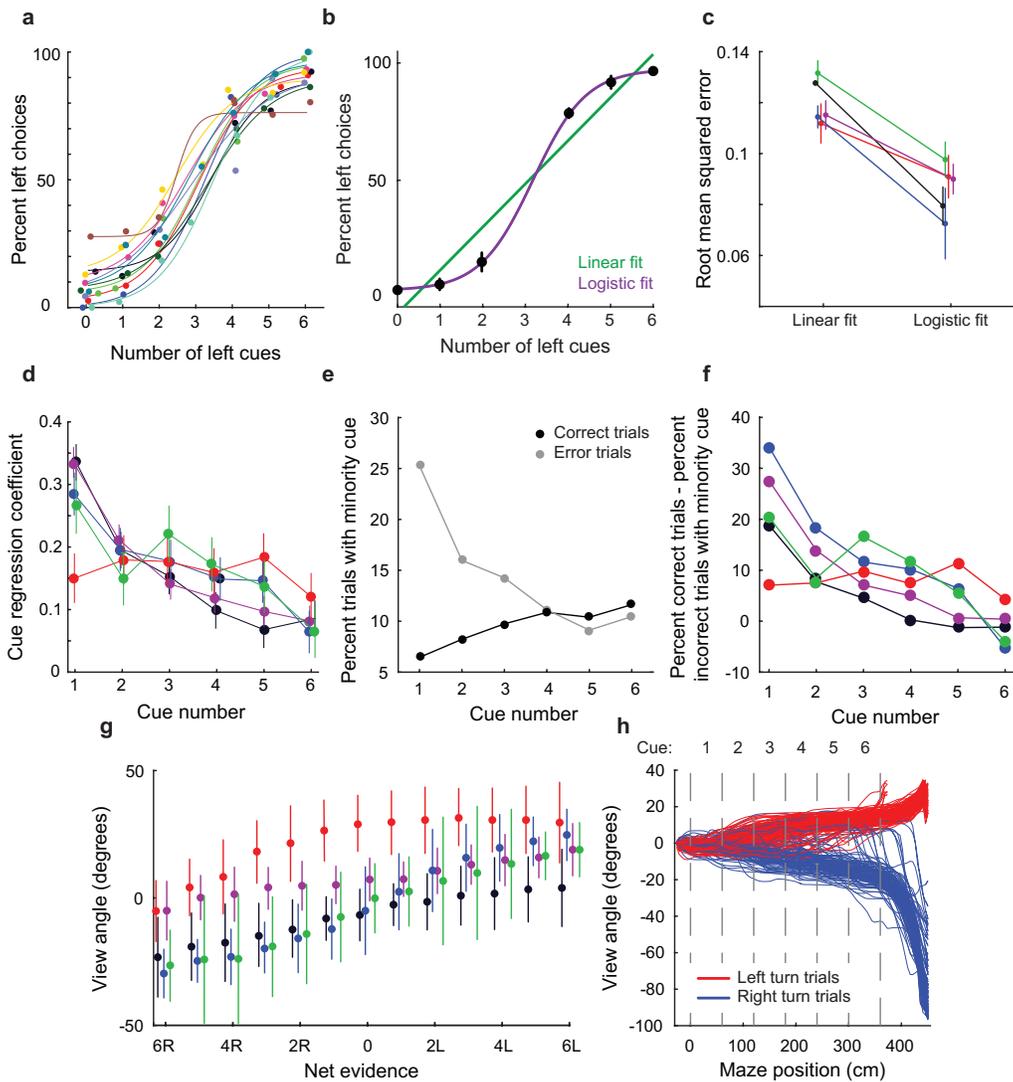


Figure 2.4 | Behavioral characterization of the fixed association evidence accumulation task. **a**, Psychometric functions describing behavioral performance on individual behavioral sessions. Each of these sessions was used to acquire imaging data in Chapter 3. Data were fit with a logistic function. **b**, Example performance from a single mouse across seven behavioral sessions fit by a linear (green) and logistic (purple) model. **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 \pm s.e.m. across datasets. Mice are colored the same as in Figure 2.2c. **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. Mice are colored the same as in Figure 2.2c. **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. Mice are colored the same as in Figure 2.2c. **g**, Relationship between net evidence and view angle for each mouse combined across all cue positions. Mice are colored the same as in Figure 2.2c. **h**, View angle as a function of maze position for individual trials within an example session. Traces are colored according to whether the mouse turned left (red) or right (blue).

approximately linear. If, however, mice used more than one cue to make their decisions, their psychometric functions should look approximately sigmoidal. We therefore 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 trials) using maximum-likelihood estimation. To compare linear and logistic model fits each behavioral day of each mouse was fit separately by both models, and the distribution of root mean squared errors (RMSE) was compared with a two-sample Student's t-test. Across mice, the logistic function fit the data significantly better than the linear function (Figure 2.4a-b), indicating that mice used more than one piece of evidence per trial. It is worth noting that this analysis does not demonstrate that mice used all six cues; it only shows that mice used more than one. Unfortunately, the precise analysis of each mouse's strategy was not possible due to limited trial numbers, especially on trials with ambiguous evidence.

Because mice did not perform the task optimally, it is possible that they weighted some evidence cues more than others. To test if cues were weighted equally, we used multivariate linear regression, with the behavioral choice as the dependent 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. We found that all cues had significant regression coefficients with a preference toward earlier segments, suggesting that mice accumulated evidence with a primacy bias (Figure 2.4d). To confirm this result, we performed a complementary analysis in which we analyzed the location of minority

cues (cues indicating the incorrect choice) on correct and error trials. We found that minority cues were uniformly distributed on correct trials, but more likely to appear in early cue positions on error trials (Figure 2.4e-f). This results suggests that mice were more likely to make an error when a minority cue was present as one of the first cues. Mice exhibit a primacy bias even though the optimal strategy for this task is to take into account all six cues (as long as at least one trial is present in which the sixth cue provides information). This suggests that mice must have placed some weight on making decisions early, before all six cues had been presented. Interestingly, this unequal weighting of cues was not present in all mice, as one mouse (colored in red in Figures 2.2 and 2.4) weighted all cues approximately equally.

We also found there to be a notable correlation between the net accumulated evidence and the mouse's view angle (Figure 2.4g). As a result, the view angles of left and right turn trials diverged well before the end of the cue period (Figure 2.4h). This correlation highlights one of the downsides of closed-loop tasks in which the association between sensory cues and choices is fixed: animals may change their behavior, and hence, the features of the sensory cues in response to prior sensory cues, introducing undesirable correlations between sensory and motor variables. These correlations can make the precise interpretation of neural responses challenging. As discussed in more detail in Chapter 3, we were therefore careful to ensure that neural results could not be wholly explained by these differences in behavior.

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 outcome (whether the previous trial was correct or incorrect) 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. Consistently, 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 ($R^2: 0.02 \pm 0.01$, mean \pm s.e.m across datasets, $p > 0.05$).

2.4 Delayed-match-to-sample task

2.4.1 Task description

This task is a navigation-based decision task with a non-fixed association between sensory evidence and choice. As with the fixed association task described above, this task is designed for head-restrained mice and is based on a T-maze in VR with two towers positioned behind the arms of the T (Figure 2.5a-c). During the cue period at the beginning of the trial, mice encounter an extended stretch of wall which is either predominantly black or white with dots of the opposite color. As mice run down the stem of the T, they enter a several second delay period in which the walls are gray, independent of the cue identity. During both the cue and delay periods, the arms and towers at the end of the T-maze are gray. At the conclusion of the delay period, the coloring of one of the arms (along with its corresponding tower) changes to white with black dots while the

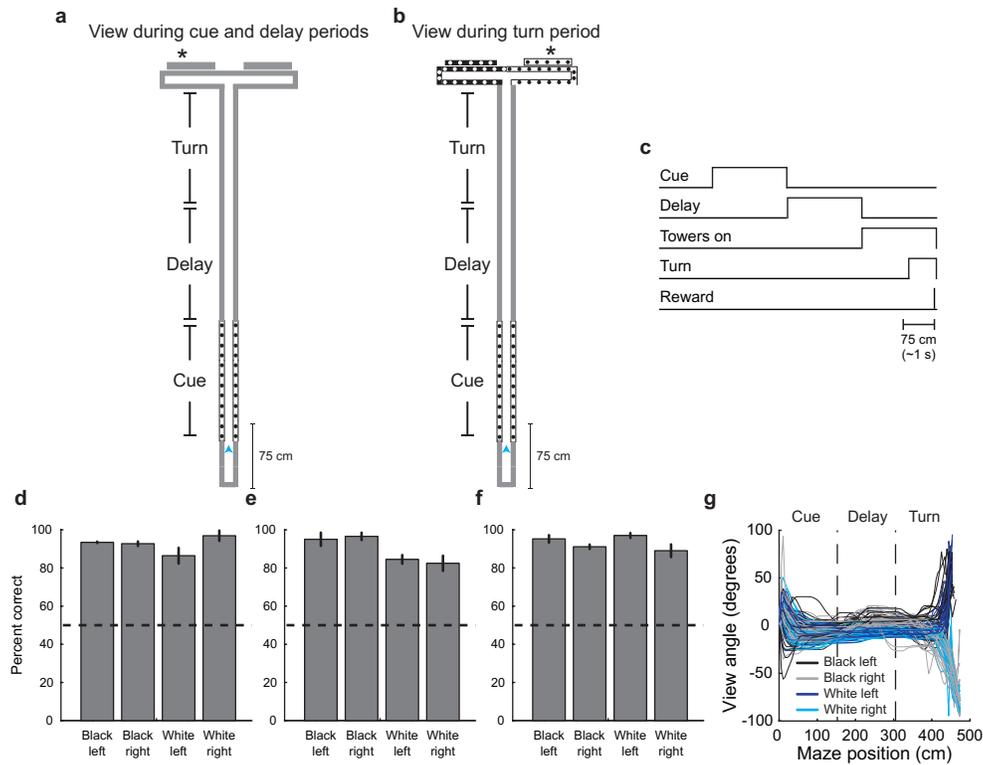


Figure 2.5 | A delayed-match-to-sample task in virtual reality. **a-b**, Schematic of an example white cue trial during the cue or delay period (**a**) and during the turn period (**b**). Asterisk marks the reward location. **c**, Sequence of trial events. **d-f**, Performance for three example mice broken up by trial configuration (mean \pm s.e.m, 4-8 sessions). **g**, View angle as a function of maze position for individual trials within an example session. Traces are colored according to trial configuration. Note that traces do not separate according to choice until the turn period and that there is a systematic difference between black and white trials during the delay period.

other changes to black with white dots. To receive a reward, the mouse must turn toward the arm which matches the cue presented during the cue period. There are therefore four trial configurations (two cues combined with two turns). Critically, the association between the cue and the rewarded turn is variable from trial to trial, such that during the cue and delay periods, the mouse does not know which direction it will need to turn for a reward. Thus, mice must remember the cue during the delay period, abstracted from a motor plan. The delayed-match-to-sample design of this task therefore allows the separation of neuronal signals related to the memory of sensory cue from those related to a motor plan (Freedman & Assad, 2011).

Mice were able to perform this task successfully across all four trial configurations (e.g., a black cue with a rewarded left turn; Figure 2.5d-f). The consistent performance across trial conditions suggests that mice were not significantly biased. For example, if mice exhibited a bias toward left choices, performance would be high on conditions in which left choices are rewarded (i.e., black left and white left) and low on conditions in which right choices are rewarded (i.e., black right and white right). Alternatively, mice may have exhibited a cue-specific bias, in which they are more likely to turn left on black cue trials and right on white cue trials. This type of bias would result in high performance on those trials, but low performance on opposite turn trials. While weak versions of these biases appeared (for example, the mouse displayed in Figure 2.5f had a weak left choice bias), these biases were not significant. Interestingly, mice often performed well on three out of the four conditions, suggesting that mice did not learn an abstract rule to match the cue presented during the cue

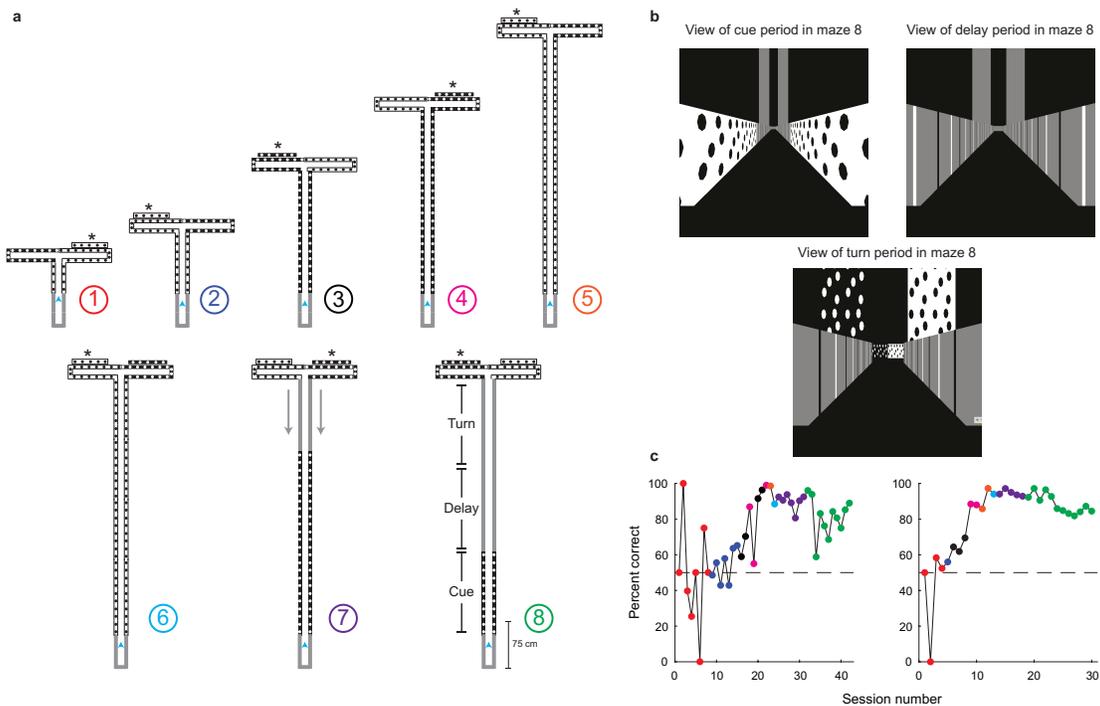


Figure 2.6 | Behavioral training procedure for the delayed-match-to-sample task. **a**, Sequence of mazes used for behavioral training. Asterisks indicate reward location. Only some example mazes are shown (for example, white right maze in maze 1). **b**, Screen captures of the virtual environment during the cue, delay and turn periods in maze 8. **c**, Behavioral performance across sessions for two example mice. Colors correspond to the maze colors indicated in **a**.

period to the arms, but instead memorized configurations (e.g., when the cue is white and white is on the left, go left).

2.4.2 Training procedure

Following at least two days of water restriction, mice began behavioral training. Behavioral sessions were performed six days per week and lasted 45-60 minutes. Mice received liquid rewards through a lick spout (4 μ l/reward, 10% sweetened condensed milk). Mice were trained to perform the delayed-match-to-sample task over a series of eight mazes (Figure 2.6). In mazes 1-5, mice progressed through a series of mazes of

increasing length which taught them to steer in the virtual environment and to associate visual cues with a reward. This progression is similar to that used in the training for the fixed association evidence accumulation task, but the mazes used to train the delayed-match-to-sample task were structured slightly differently. In each of these mazes, both walls of the stem contained either the white or black cue and the arms contained both cues, with one on the left and one on the right. The correct arm, which matched the cue present in the stem of the maze, was indicated by a tower colored with the same cue which was located behind the rewarded arm. The unrewarded arm did not have a tower behind it. Mice could therefore achieve perfect performance on these mazes by turning toward the arm with a tower without paying attention to the cues on the walls of the T-maze. The trial configuration was randomly selected from the four possibilities prior to each trial. Importantly, all four configurations were present during these early mazes to prevent mice from learning a spurious association between a cue and a choice (e.g., turn left when a black cue is present).

In maze 6, an additional tower was added behind the unrewarded arm as well, forcing mice perform a non-delayed-match-to-sample in which they must turn toward the arm whose coloring matches the cue presented in the stem of the maze. When we first attempted this transition, mice were unable to learn it, rapidly dropping to chance performance and picking up strong biases (e.g., always turn left). To ease mice into this transition, consecutive trials were ‘paired’, such that a trial with a single tower (as in maze 5) always preceded a trial with two towers (as in maze 6). Critically, these paired trials had the same configuration and mice

received three times as many rewards for making a correct choice on the two tower trial. Because consecutive trials featured the same reward location in this paired organization, this structure may have taught mice to repeat a rewarded choice. However, we found that mice rarely learned this strategy, and that the paired organization helped mice to rapidly achieve high performance on maze 6, often within a single day (see cyan points in Figure 2.6c).

Maze 7 begins identically to maze 6, except visibility of the arm configuration is obscured until mice pass the $2/3$ point of the stem. However, because there is no delay period yet, this maze still only requires a non-delayed-match-to-sample. As mice performed trials correctly, a delay (in the form of gray walls separating the cue and arms) would slowly appear, increasing in length as mice performed more and more trials correctly. Once the length of the delay reached $1/3$ of the stem, mice performed an instantaneous-match-to-sample, in which the cue and arm configuration were no longer visible simultaneously. As the delay length increased beyond this point, mice performed a delayed-match-to-sample task with an increasingly longer delay period. Maze 8 appeared identically to the final version of maze 7, in which the cue, delay, and turn periods each extended for $1/3$ of the maze stem.

As with the fixed association task, bias correction was introduced beginning with maze 4. Bias correction was applied to both turn direction and arm configuration independently. Prior to each trial, the probability that a left choice would be rewarded was set to the fraction of right choices over the preceding 20 trials. The probability that a turn towards a white arm

would be rewarded was set to the fraction of turns toward black arms over the preceding 20 trials. For example, if a mouse turned left 75% of the time and toward black arms 75% of the time, then the probability of the next trial being black left would be 0.0625 (0.25×0.25), being black right 0.1875 (0.25×0.75), being white left 0.1875 (0.75×0.25) and being white right 0.5625 (0.75×0.75). Crutch trials were introduced as well, beginning with maze 7. Crutch trials were identical to trials from maze 6 (Figure 2.6a). As with the fixed association task, the probability of a crutch trial was set to the fraction of error trials over the preceding 20 trials.

For mazes 1-3, the criterion for advancement to the next maze was the mouse's number of completed trials per minute, independent of performance (> 7 trials/min for advancement). Mice on mazes 4-5 were advanced to the next maze following one day of $\geq 90\%$ performance. Mice on maze 6 were advanced following one day of $\geq 80\%$ performance on trials with two towers. In maze 7, delay advancement was controlled according to the following rule. After the first 30 trials of a session, the delay advanced by $1/60$ of the stem whenever mice performed three consecutive correct trials. Advancement was halted for at least 30 trials every $1/6$ of the maze, creating a block structure of advancement (e.g., once the delay had reached $1/6$ of the stem, $2/6$ of the stem, etc.). To resume delay advancement, mice must have performed 75% of the preceding 30 trials correctly. On consecutive training sessions, mice were reset to the last block on which they had high accuracy. For example, if a mouse ended a session with a delay of $27/60$, the next session would begin with a delay of $2/6$. This process continued until mice had accuracy $\geq 75\%$

on sessions beginning with a delay of 2/3, at which point they were advanced to maze 8. Training in total required ~20-40 daily training sessions (Figure 2.3d). This task was difficult to train; only about one-quarter of the mice who begin training eventually reached high accuracy on the final task.

2.4.3 Behavioral characterization

In contrast to the fixed association evidence accumulation task, mice did not demonstrate a strong association between view angle and choice during the cue and delay periods (Figure 2.5g). This serves as a confirmation of the delayed-match-to-sample design of this task, as the mouse cannot know its choice until the turn period when the arm configuration is revealed (presuming the mouse is performing the task at high accuracy). Interestingly, there was a weak, but noticeable, relationship between the cue and the mouse's view angle during the delay period. For example, the example mouse in Figure 2.5g systematically turned toward the left during the delay period when it had seen a black cue and toward the right when it had seen a white cue. This effect suggests that mice may have used their view angle to remember the cue throughout the delay period, potentially complicating the interpretation of neuronal activity during the delay period, and removing some of the advantages of a delayed-match-to-sample task design.

To test whether there was a relationship between the mouse's choice on a trial and the outcome of the previous trial, we used the same multivariate logistic regression as we did for the fixed association evidence accumulation task. Consistent with that result, the model was

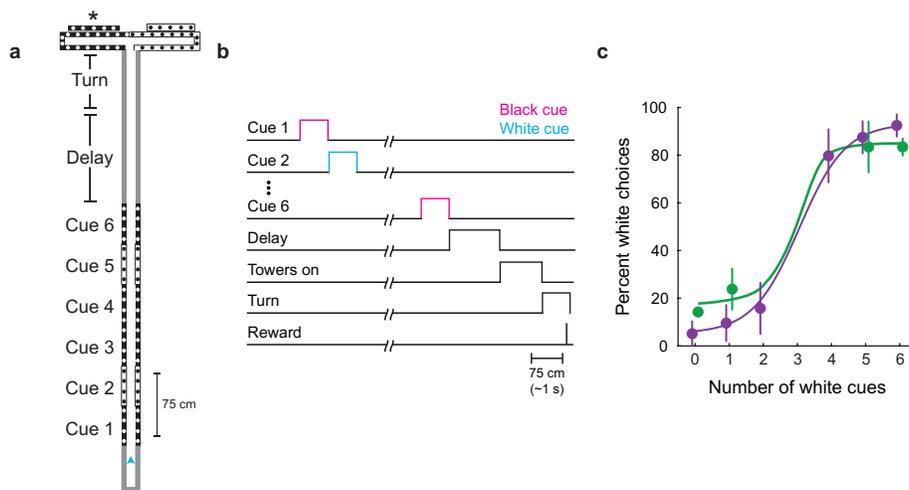


Figure 2.7 | A delayed-match-to-sample evidence accumulation task in virtual reality. a-b, Schematic of an example 4-2 trial. Asterisk marks the reward location. **b**, Sequence of trial events. **c**, Performance for two mice (mean \pm s.e.m, 3 and 12 sessions for the green and purple mice, respectively).

unable to predict the mouse's choice, again suggesting that there was no clear relationship between the mouse's choice on trial n and its choice on trial $n+1$ (R^2 : 0.01 ± 0.01 , mean \pm s.e.m across datasets, $p > 0.05$).

2.5 Delayed-match-to-sample evidence accumulation task

2.5.1 Task description

This task adds an evidence accumulation component to the delayed-match-to-sample task (Figure 2.7a-b). As in the fixed association evidence accumulation task, as mice run down the stem of a T-maze, they are presented with six, discrete visual cues. In contrast to the fixed association task, these cues appear on both walls of the T-maze and can either be black or white (with dots of the opposite color). As in the delayed-match-to-sample task, the arm configuration is obscured (colored gray) until the conclusion of the delay period, at which

point it becomes visible. To perform this task well, mice must decide whether there are more black or white cues, remember the accumulated cue identity throughout the delay, and then use the memory of the accumulated cue to make a choice toward the corresponding arm during the turn period. As in the delayed-match-to-sample task, mice cannot know which choice will be rewarded until the end of the delay period. Mice must therefore maintain a memory of the accumulated cue that is independent of the eventual motor choice throughout the cue and delay periods.

This task is extremely difficult and inefficient to train; only four out of close to one hundred mice ever reached high accuracies on the final task, and of those, the performance of all but one was unstable from day to day. However, several mice were able to perform this task with high accuracy (Figure 2.7c).

2.5.2 Training procedure

The first portion of training for this task was almost identical to the training for the delayed-match-to-sample task described in section 2.4.2. The only difference is that the cue period was lengthened to allow for the eventual discretization into six cues. Once mice performed the delayed-match-to-sample task stably across days with high accuracy, the cue period was divided into six discrete cues (maze 8). As with the fixed association evidence accumulation task, cue visibility was locked to the mouse's position such that only one cue was visible at a time and there was no delay between the offset of one cue and the onset of the next cue. Initially, only 6-0 trials (trials in which all six cues were black or white) were included.

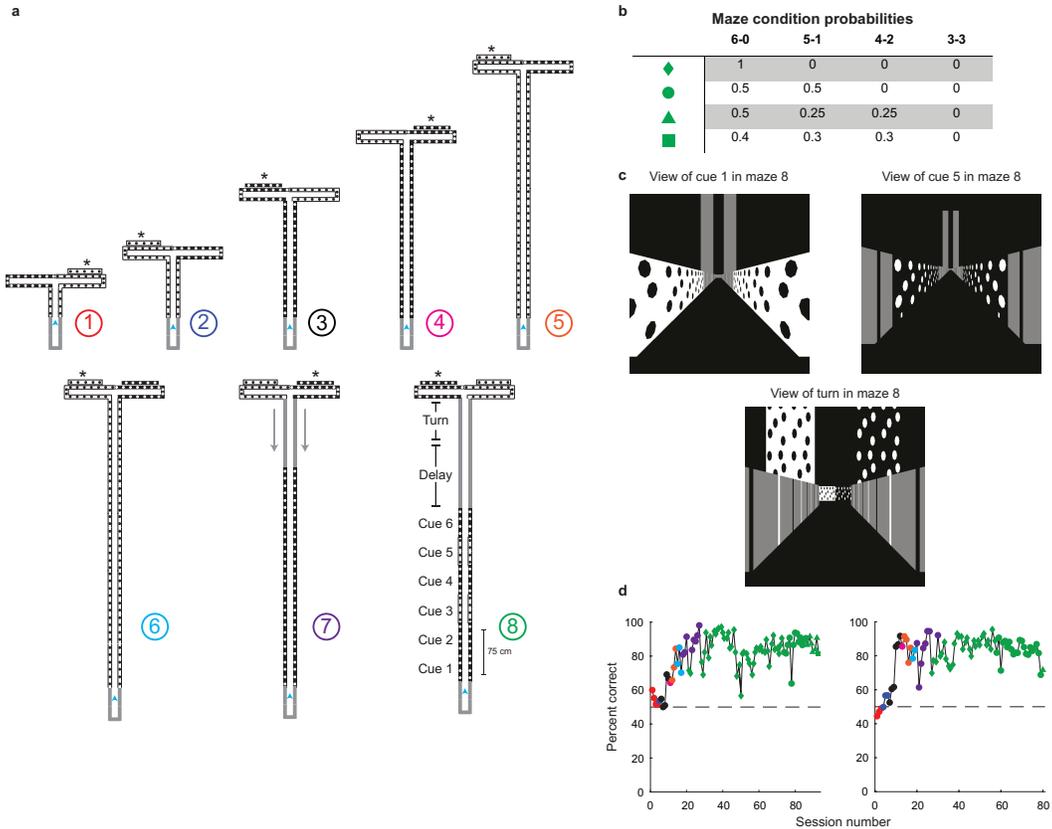


Figure 2.8 | Behavioral training procedure for the delayed-match-to-sample evidence accumulation task. **a**, Sequence of mazes used for behavioral training. Asterisks indicate reward location. Only some example mazes are shown (for example, white right maze in maze 1). **b**, Distribution of net evidence 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 5, and the turn in maze 8. **d**, Behavioral performance across sessions for two example mice. Colors correspond to the maze colors indicated in **a**. Shapes correspond to the net evidence probabilities in **b**.

More difficult conditions (5-1 and 4-2 trials) were gradually introduced according to the probability distributions in Figure 2.8b. Only one mouse was ever able to successfully perform 4-2 trials and no mice were ever able to perform 3-3 trials without performance dropping significantly on easier trial conditions. Bias correction and crutch trials were included as in the training for the delayed-match-to-sample task. Training in total required ~60-90 sessions.

2.5.3 Behavioral characterization

As in the fixed association evidence accumulation task, we tested whether mice weighted some evidence cues more than others (Section 2.3.3). In contrast to the primacy bias observed in the fixed association task, a multivariate linear regression demonstrated that mice weighted cues approximately equally (Figure 2.9a). The analysis of minority cue location, however, suggested that mice had a noticeable recency bias, as minority cues were most likely to induce an error when they were located at the fifth and sixth cue positions (Figure 2.9b-c).

We also found there to be no systematic correlation between the net accumulated evidence and the mouse's view angle (Figure 2.9d). As in the delayed-match-to-sample task, the mouse's view angle was also unable to predict the mouse's choice until the turn period (Figure 2.9e). Interestingly, the difference in view angle due to cue during the delay period that we observed during the delayed-match-to-sample task (see Section 2.4.3) was absent during this task.

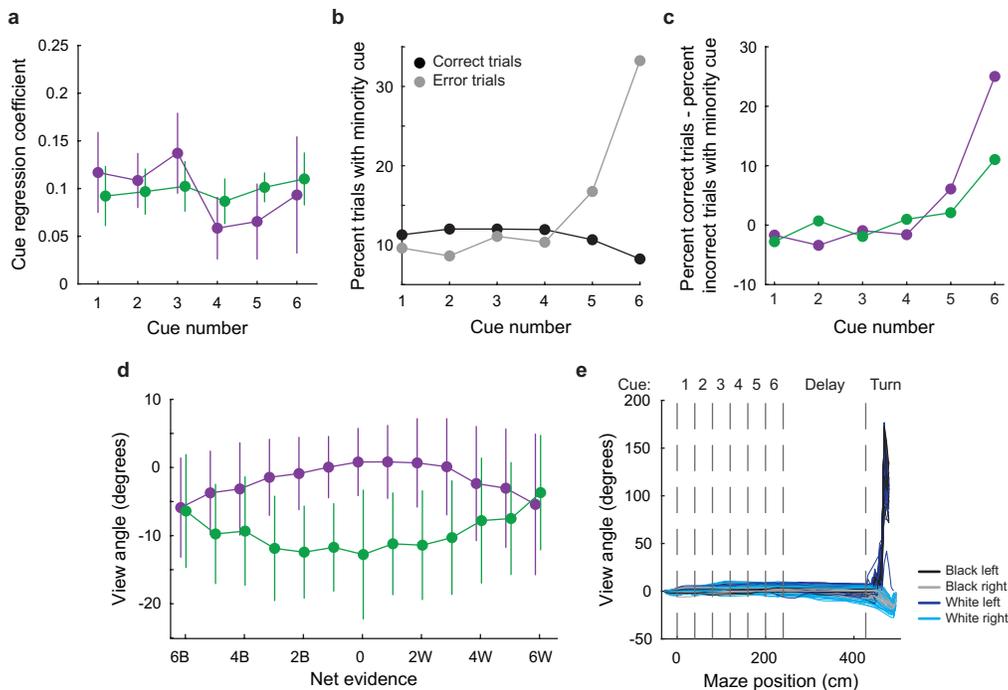


Figure 2.9 | Behavioral characterization of the delayed-match-to-sample evidence accumulation task. a, Multivariate linear regression in which the mouse's choice was the response variable and the six cue identities were separate explanatory variables. Regression coefficients for two mice (3 and 12 sessions for the green and purple mice, respectively) are shown. Error bars indicate confidence intervals. Mice are colored the same as in Figure 2.7c. **b-c,** 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 (**b**; mouse colored purple in **a**, **c**, **d**) and as the difference of the error and correct points (**c**) for both mice. Mice are colored the same as in Figure 2.7c. **d,** Relationship between net evidence and view angle for each mouse combined across all cue positions. Mice are colored the same as in Figure 2.7c. **e,** View angle as a function of maze position for individual trials within an example session. Traces are colored according to the trial configuration.

Finally, to test whether there was a relationship between the mouse's choice on a trial and the outcome of the previous trial, we used the same multivariate logistic regression as we did for both previous tasks (Sections 2.3.3 and 2.4.3). Consistent with those results, the model was unable to predict the mouse's choice, again suggesting that there was no clear relationship between the mouse's choice on trial n and its choice on trial $n+1$ (R^2 : 0.01 ± 0.01 , mean \pm s.e.m across sessions for purple mouse, $p > 0.05$; 0.02 ± 0.01 , mean \pm s.e.m across sessions for green mouse, $p > 0.05$).

2.6 Behavioral analysis suite

I have also developed the *Behavioral Analysis Suite* (BAS), a MATLAB-based software package for the analysis of behavioral data. The BAS is a modular software package that can be used both graphically and programmatically to analyze any behavioral task with a discrete trial structure¹. It has several key features. First, it enables the easy filtering of trials according to any combination of task parameters. For example, in the delayed-match-to-sample task described above, one could easily extract all black left trials. One can also subset trials according to substantially more complicated filters, such as all white left trials preceded by a white right trial, lasting longer than seven seconds, and in which the mouse made an error. The ability to rapidly create arbitrary trial subsets was essential during the development of the behaviors described in this chapter. Second, the BAS can be used online to analyze and monitor ongoing behavior as well as offline to analyze past behavior. In

¹The design of data structure underlying the BAS was inspired by discussions with Alex Trott and Alex Wiltschko.

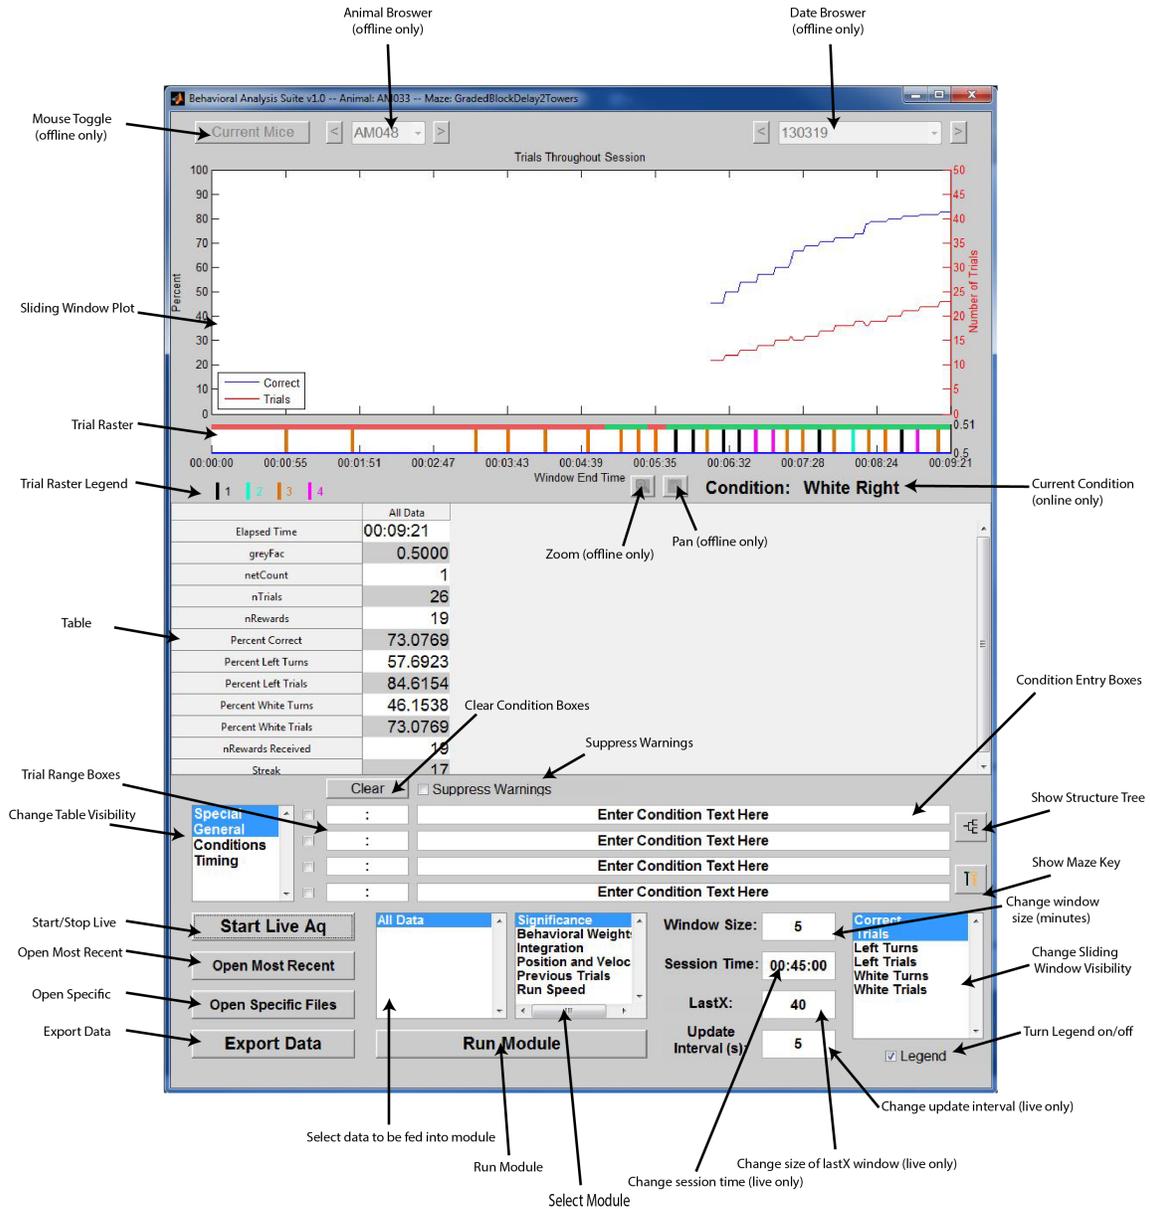


Figure 2.10 | Behavioral analysis suite. Graphical interface of the behavioral analysis suite with arrows indicating each component.

offline mode, users can easily combine multiple behavioral sessions, even from different animals. Third, filtered data and generated plots can be exported for further analysis and processing. Finally, custom, user-written modules can be added to the graphical interface to perform arbitrary computations on subsets of the data generated by the BAS. This tool allows users to add modules of their own which perform analyses specific to their use cases. An example of the graphical interface is provided in Figure 2.10.

The behavioral analysis suite is freely available on GitHub².

2.7 Thoughts on task design

The tasks described in this chapter demonstrate that mice can perform complex decision tasks requiring the extensive use of working memory. However, it also highlights the difficulty of designing, training, and interpreting such tasks. For example, during the delayed-match-to-sample task, which is designed to identify neural signals of abstract decision-making that are independent of motor actions, some mice systematically shifted their view angle during the delay period in response to the cue (Figure 2.5g). It is therefore difficult to determine whether the mouse is actually remembering the cue, or merely using its view angle as a form of ‘memory’. This example highlights the difficulty of interpreting the strategy taken by animals to perform tasks. Care must be taken to ensure that the task design does not encourage alternative strategies which are close to as effective as the desired task strategy.

²<https://github.com/arimorcos/BehaviorAnalysisSuite>

One way to reduce the potential for alternate behavioral strategies is to increase behavioral control. This will also reduce external sources of behavioral variance. For example, new tasks should endeavor to reduce variance in view angle across trial conditions. This could be accomplished in two ways. The view angle could be restricted during portions or the entirety of a trial, enabling perfect control. Alternatively, mice could be encouraged to run with repeatable patterns either by including incentives for running straight, such as extra rewards for trials completed quickly, or by punishment for deviations, such as the addition of friction to the environment, substantially slowing the pace of mice with extreme view angles. However, even with more tightly controlled tasks, we can never eliminate all sources of behavioral variance, so measured caution must always be taken in the interpretation of neural signals.

We should also seek to engineer systems allowing for the rapid prototyping of task designs. This is a critical advantage of the VR system described above. Many task designs that may at first seem unlikely to be successful can prove to be highly effective. For example, in training the delayed-match-to-sample task, the 'paired' strategy was essential to induce mice to pay attention to the cue configuration (see Section 2.4.2). This design seemed likely to induce mice to simply repeat choices which led to a reward. However, only a small fraction of mice ever learned this strategy. This example illustrates both the difficulty of predicting how animals will react to certain task designs and highlights the utility of rapid prototyping. If the cost of attempting such a strategy had been high, we likely never would have attempted it.

Because tasks can vary dramatically in the amount of effort required for their development and training, we should also seek to use the simplest task possible for a given scientific question. For example, the delayed-match-to-sample evidence accumulation task proved incredibly difficult to train reliably, while the fixed association evidence accumulation task, in comparison, was trivial to train and develop. However, almost all of the scientific questions answerable by the former task could also be answered by the latter task, with the notable exception of questions seeking to dissociate sensory responses from motor responses. The benefits of a simpler task for the majority of scientific questions may therefore outweigh the cost of reduced interpretability.

Chapter 3

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

Ari S. Morcos and Christopher D. Harvey

This and the following chapters are a modified version of a submitted manuscript.

3.1 Introduction

The activity patterns in a cortical microcircuit are determined both by the characteristics of the external inputs it receives and the dynamic, ongoing changes in its internal activity state. The computation to combine ongoing activity with new inputs is essential for many complex behaviors, including evidence accumulation during decision-making. Much work has focused on identifying the neuronal algorithms and mechanisms by which evidence accumulation occurs, with considerable emphasis on the posterior parietal cortex (PPC) (Shadlen & Newsome, 1996; Gold & Shadlen, 2000; Yang & Shadlen, 2007; Hanks et al., 2015; Britten & Shadlen, 1992). Previous work has modeled evidence accumulation as a competition between distinct pools of neurons with recurrent excitation within a pool

and mutual inhibition across pools (Wong & Wang, 2006; Wang, 2002). The activity of individual neurons in these models is long-lasting within a trial and homogeneous across neurons. These models propose a ‘winner-take-all’ competition between these neuronal groups in which activity eventually converges to one of several attractor states (Wong & Wang, 2006; Machens, 2005; Wang, 2002). Although such models are consistent with some experimental results (Shadlen & Newsome, 1996; Gold & Shadlen, 2000; Yang & Shadlen, 2007; Hanks et al., 2015; Britten & Shadlen, 1992; Horwitz & Newsome, 1999), our recent work in the mouse PPC during a navigation-based decision task found that neurons had transient, time-varying activity that was heterogeneous across neurons (Harvey et al., 2012). These results led us to conceptualize the activity in the PPC on single trials as a trajectory of time-varying neuronal population activity patterns. The apparent inconsistencies between our previous results and traditional models revealed to us the lack of competing algorithmic models for evidence accumulation.

Here, we used the conceptual framework of time-varying neuronal population activity trajectories to study neuronal population dynamics on single trials during evidence accumulation. Previous work has emphasized independent recordings from selected subsets of individual neurons, summarized as averages across trials and cells, in part because of technical challenges in measuring and interpreting neuronal population activity on single trials. We therefore developed new experimental and computational methods based on unbiased sampling of activity from large populations of neurons to reveal structure in the moment-to-moment changes in population activity.

We found that the PPC had long timescale dynamics in the form of structured transitions between transient and largely uncorrelated patterns of population activity. New inputs to the network, including evidence cues and behavioral choices, constrained the possible population activity patterns for seconds into the future. This effect occurred by changing the transition probabilities between activity patterns consisting of largely different combinations of active neurons. The population-level representation of new inputs thus depended both on the identity of the input and the near-past activity patterns in the population, such that PPC activity never reset but rather functioned as a continuous record of recent events. In addition, multiple task-relevant features were represented simultaneously, distributed across populations of heterogeneous and variable neurons, such that single task features (e.g., choice) did not converge to single activity patterns but rather were represented across trials by many different activity patterns. These results are inconsistent with evidence accumulation models that require a direct competition between two or more groups of neurons (Wong & Wang, 2006; Machens, 2005; Wang, 2002), and instead reveal novel features of PPC dynamics that motivate a new algorithmic model based on general purpose, history-dependent dynamics.

3.2 Results

We used the fixed association evidence accumulation task described in Section 2.3. Briefly, as mice ran down a virtual T-maze, they were presented with six visual cues that could each appear on the left or the right wall at fixed locations (Figures 2.2a-b, 2.3). To receive a

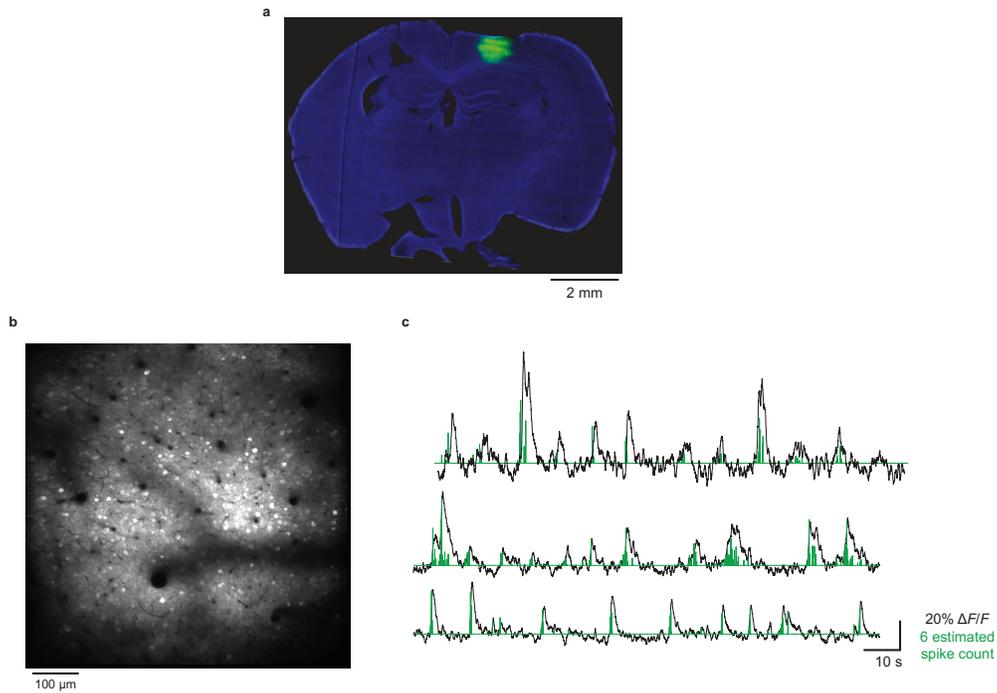


Figure 3.1 | Example imaging field of view and activity traces. **a**, Example histology image of GCaMP6m-expressing neurons in the PPC. **b**, Example two-photon image of GCaMP6m-expressing neurons in layer 2/3 of the PPC. **c**, Example $\Delta F/F$ traces (black) and deconvolved estimated spike counts (green) (Methods 5.2.4).

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

Figure 3.2 | Summary of datasets analyzed.

reward, the mouse had to turn at the T-intersection toward the direction that had more cues. Mice performed the task with high accuracy by accumulating multiple pieces of evidence, with a bias towards early cues (Figures 2.2c, 2.4).

3.2.1 Single neuron responses during evidence accumulation

We began our analysis of neuronal population activity during this task by examining 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 extracted estimated spike counts using deconvolution of the fluorescence traces (Vogelstein et al., 2010) (Methods 5.2; Figures 3.1, 3.2). 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 (Harvey et al., 2012) (Figures 3.3a, 3.4, 3.5a-b). To test for differences in activity between trials with different choices and net evidence, we calculated a selectivity index for choice and used a support vector regression (SVR) model to predict net evidence from a single cell's activity (Methods 5.3.2, 5.4.3). A fraction of neurons had a statistically significant choice selectivity index, and some neurons had a significant relationship between the actual net evidence and the net evidence predicted from their activity (choice 20.2%, net evidence 22.7%; 5% expected by chance; Figure 3.3c-d). When we plotted the mean activity patterns for the neurons with choice selectivity, we identified choice-specific sequences of activity, consistent with PPC activity patterns during less complex decision tasks (Harvey et al., 2012) (Figure 3.4c-d). However, the

majority of neurons did not have significantly different activity between trials of different choices and net evidence, consistent with our previous results (Harvey et al., 2012). The distribution of choice selectivity indices and net evidence prediction accuracies for single neurons largely overlapped with the distribution from shuffled data (Figure 3.3c-d). The low selectivity values in these neurons resulted from unreliable responses with varying activity times between trials of the same type and similar activity between trials of different types.

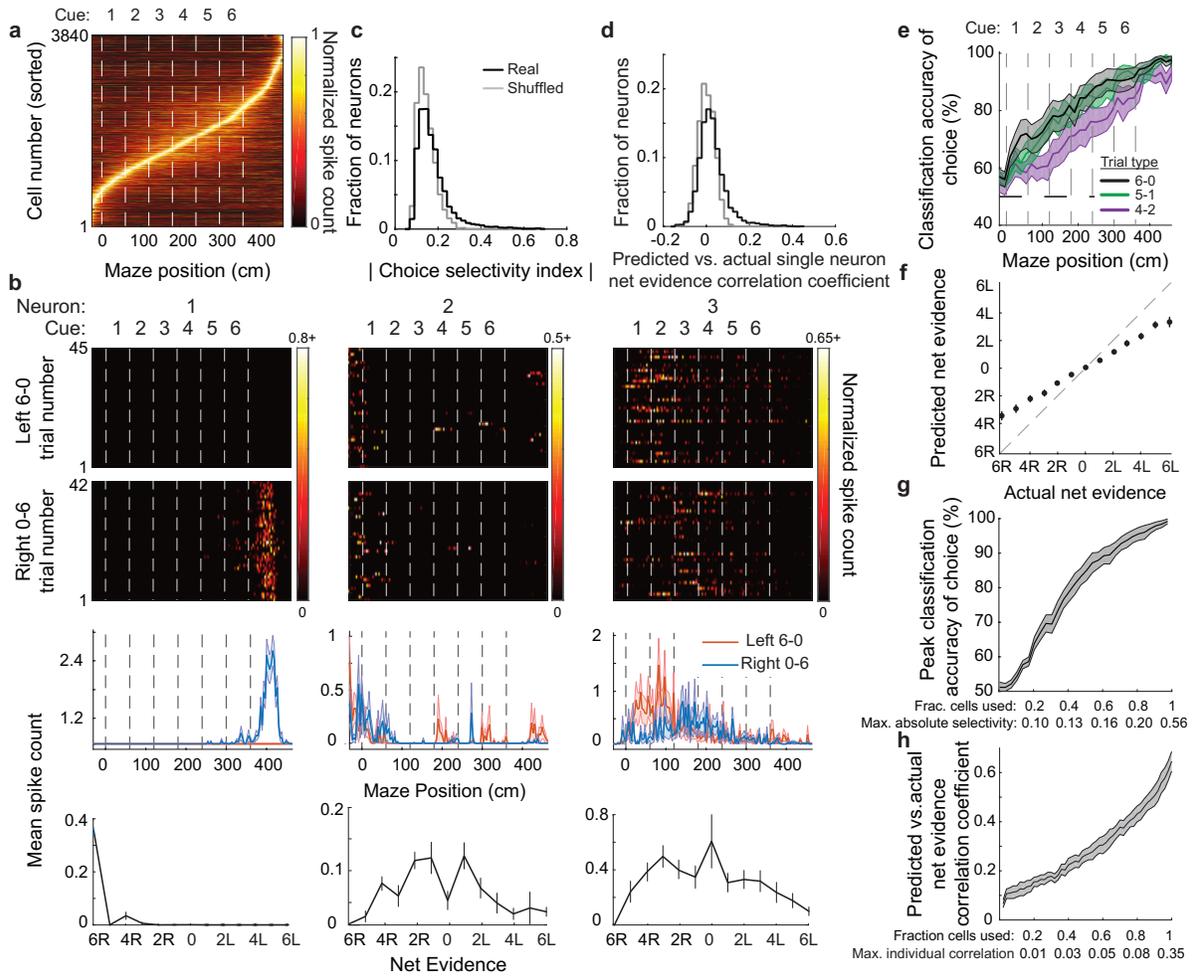
3.2.2 Task-relevant information is distributed across neuronal populations

Although the majority of individual neurons lacked strong selectivity, the population activity (concatenated activity of all individual neurons) contained information about the choice and net evidence on single trials with high accuracy. We quantified this selectivity using a support vector machine (SVM) classifier for choice and a SVR model to compare actual net evidence with the net evidence predicted by the population activity (Figure 3.3e-f).

Figure 3.3 (following page) | 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 on left 6-0 and right 0-6 trials for three example neurons. Top panels: each row is an individual trial. Bottom panels: mean \pm 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 the choice selectivity index for individual neurons (black) and with shuffled trial labels (gray). Choice selectivity was calculated separately for each spatial bin, and the maximum magnitude across bins was taken for each neuron. Choice selectivity for each neuron was calculated based on activity in left 6-0 and right 0-6 trials (Methods 5.3.2). **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 \pm s.e.m., $n = 11$ datasets) for choice using an SVM based on population activity during 6-0 (black), 5-1 (green) and 4-2 (purple) trials. 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 \pm 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**)). Classifier performance increased as neurons individually containing little task-relevant information were included, suggesting that information was distributed across neurons. Shaded error bars represent mean \pm s.e.m. across datasets, and max individual neuron classification accuracies/correlations were the mean across datasets.

Figure 3.3 | (Continued)



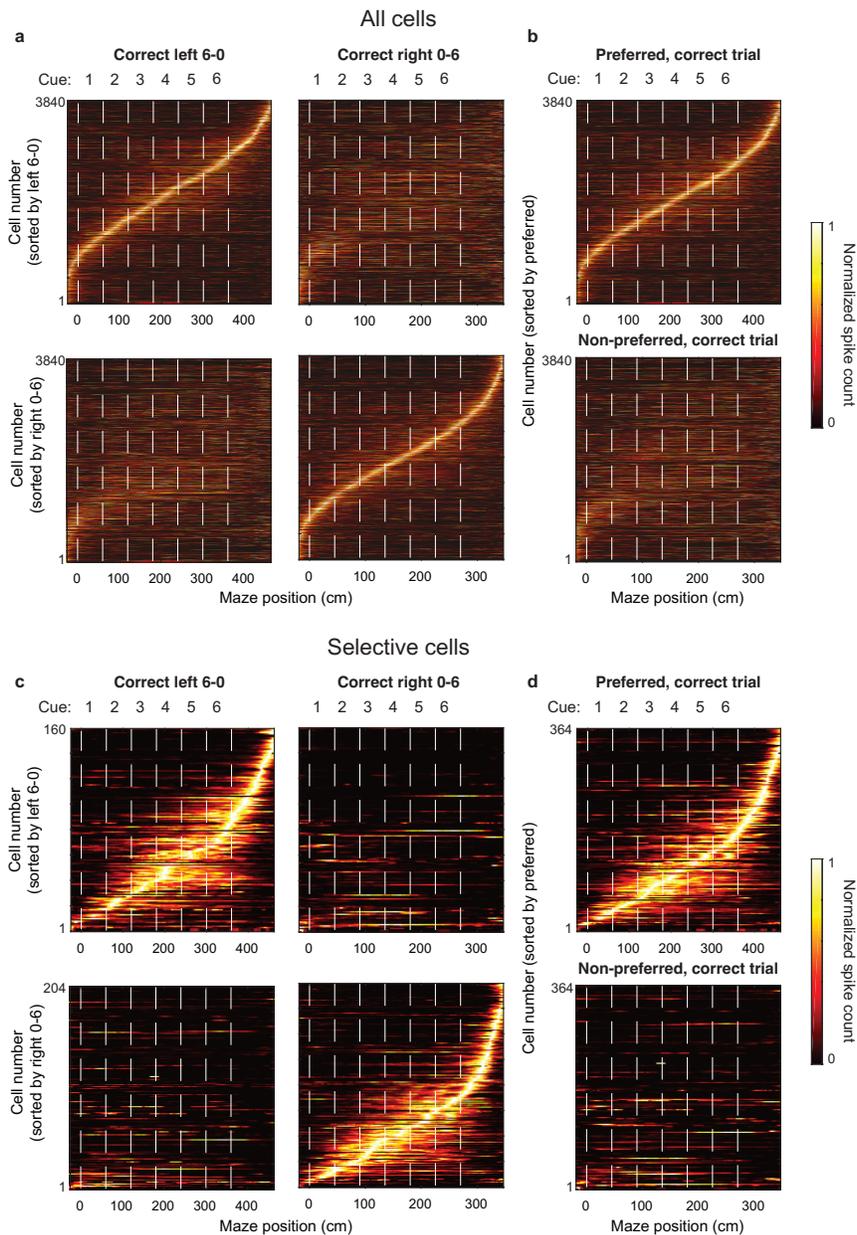


Figure 3.4 | Mean population activity patterns in PPC for all cells and selective cells. **a**, Normalized mean activity across correct left 6-0 (left) and correct right 0-6 (right) trials for all neurons pooled across all datasets ($n = 3840$ cells from 5 mice). Traces were normalized to the peak of 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 the sign of its choice selectivity index. **c-d**, Same as **a-b**, but only for selective cells. Selective cells were defined as all cells whose peak choice selectivity index magnitude exceeded 0.25.

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 evidence accumulation. In all cases, we found that behavioral variability could not entirely explain the neuronal activity patterns we observed. Our analyses included the mouse's position and view angle in the maze, which together defined the visual scene, and also included the spherical treadmill's rotational velocity, which was determined by the mouse's running pattern. Together, these parameters defined a large part of the mouse's visual and motor experience, including the parameters that were most likely to be correlated with specific task events.

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 (e.g., 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). Our SVR analysis to predict net evidence (Figure 3.3d, f, h) included all trials, for which differences in view angle across net evidences were present (Figure 2.4g). However, when we limited our analysis to only trials with similar view angles ($\pm 2.5^\circ$), the SVR analysis based on population activity predicted above chance levels the actual net evidence (Figure 3.5e-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; Figure 3.5g-h). These results suggest that a representation of net evidence was present independent of heading direction, running patterns, and direct visual input.

The choice and net evidence information might have come solely from a small fraction of neurons with high selectivity; these are the neurons that we and others have emphasized previously. Alternatively, information could have been distributed across a large group of neurons extending beyond the classically selective neurons. To distinguish these possibilities, we applied 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 (Figure 3.3g-h). Using the 40% least selective neurons, we were able to predict the mouse's choice with 75% accuracy, even though none of these neurons had a statistically significant choice selectivity index (Figure 3.3g). The improvement with the addition of more neurons could have resulted from the additive effect of weak information in each neuron or the revelation of correlated activity patterns with the addition of larger numbers of neurons or both; we plan to investigate these possibilities in future work. Together, these results suggest a

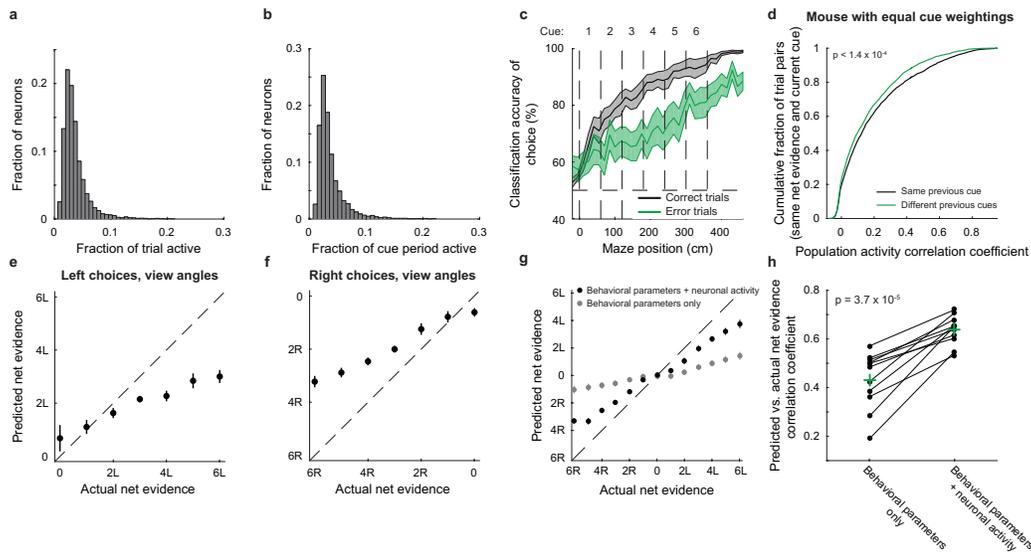


Figure 3.5 | Analyses of single neuron- and population-level representations of task-relevant features. **a-b**, Histogram of the fraction of the entire trial (**a**) and cue period (cues 1-6) (**b**) neurons were active ($n = 3840$ neurons from 5 mice). **c**, SVM classification accuracy (mean \pm s.e.m., $n = 11$ 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 Figure 3.15a 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 Figures 2.2, 2.4b, 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 ($\pm 2.5^\circ$) view angles on left choice (**e**) and right choice (**f**) trials. 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. **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 5.4.3). Error bars represent mean \pm s.e.m. across datasets ($n = 11$). **h**, Data from (**g**) shown for individual datasets. Green crosses represent means across datasets ($n = 11$).

population representation in which information is distributed across heterogeneous and variable neurons, as has been suggested previously in studies of coding in individual neurons (Meister et al., 2013; Park et al., 2014; Jun et al., 2010; Raposo et al., 2014; Mante et al., 2013; Rigotti et al., 2013; Maimon & Assad, 2009; Churchland et al., 2010). For all subsequent analyses, we therefore considered all the neurons we imaged, in contrast to many previous studies that have selected, either at the measurement or analysis stage, only those neurons that were individually selective for a specific task feature, such as choice.

3.2.3 Clustering neuronal activity patterns across trials

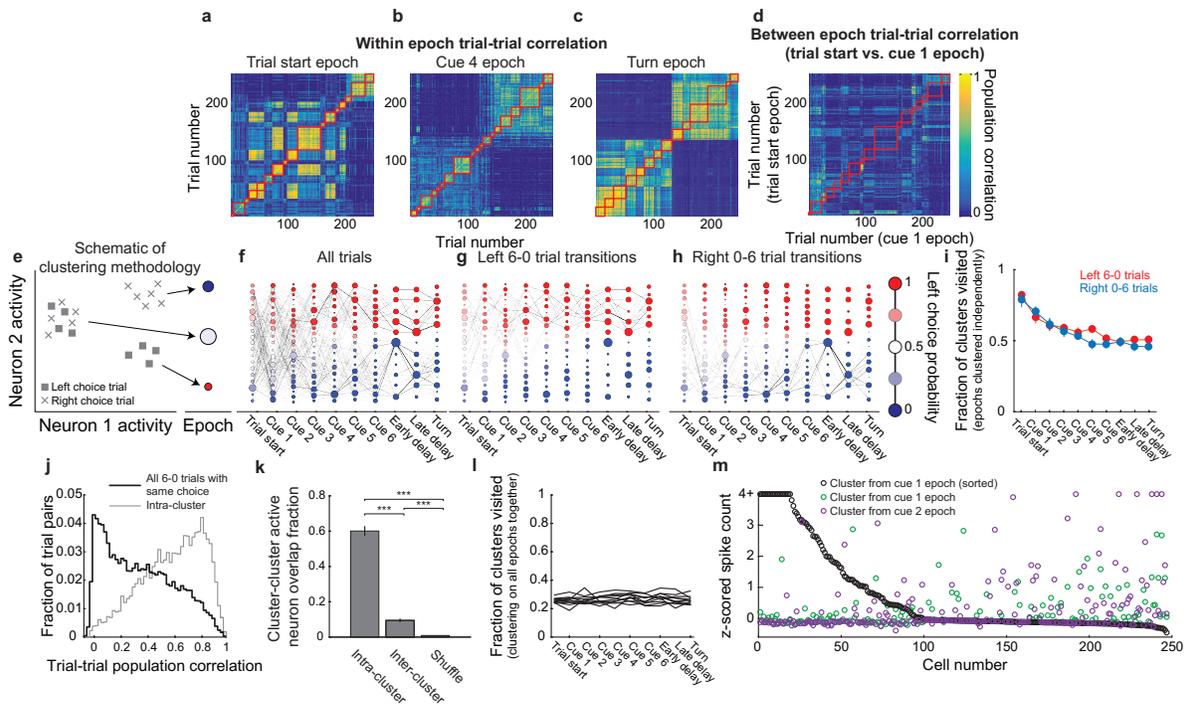
Given that neuronal activity was in large part heterogeneous across neurons and variable between trials and that task-relevant information was distributed across neurons, we reasoned that further analyses of individual neuron activities might not achieve our goal of understanding features of the population activity on single trials. We therefore focused our subsequent analyses exclusively on the population level by developing methods to measure how the population activity pattern changed from moment to moment. We defined the population activity pattern at a given time period as the vector of each neuron's estimated spike count in that period. 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 over time, we reduced the dimensionality of the population activity using a clustering algorithm. Activity patterns were clustered based on similarity in the n -dimensional activity space defined by the n simultaneously imaged neurons (Methods 5.6). We considered each cluster to represent an

activity pattern in the PPC population. The number of clusters was determined using the affinity propagation clustering algorithm (Frey & Dueck, 2007). Our results were consistent across a wide range of cluster numbers and affinity propagation settings (Figure 3.7k; Methods 5.6.2). Clustering was performed independently for ten epochs, each of which corresponded to a different time period of 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. 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 (11 ± 8.4 trials/cluster on average; Figure 3.7e). Each trial was part of a single cluster at each epoch. For visualization, each cluster was represented as a circular node with area proportional to the number of trials in the cluster (Figure 3.6e). Transitions between clusters in consecutive epochs were marked as lines with thickness proportional to the transition probability (Methods 5.6.3). Single trials could therefore be described as an activity trajectory defined by the sequence of clusters visited from epoch to epoch. 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 (Harvey et al., 2012; Churchland et al., 2012; Mante et al., 2013; Mazor & Laurent, 2005; Briggman, 2005; Raposo et al., 2014).

Figure 3.6 (following page) | Clustering neuronal activity across trials to reveals trial-to-trial variability.

a-c, Example trial-trial population activity correlation matrices at the trial start epoch (**a**), cue 4 epoch (**b**), and the turn epoch (**c**) sorted by cluster identity. Red squares indicated cluster membership such that pairs of trials within the same red square were in the same cluster. **d**, 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. **e**, Schematic demonstrating clustering procedure (Methods 5.6). At each of ten spatially-defined maze epochs, clustering was used to group together 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. Transition matrices were constructed by calculating the empirical transition probability between adjacent clusters. **f**, An example transition matrix constructed from all trials in a single dataset. Edge widths between nodes represent the forward transition probability. Nodes were colored and sorted as described in (**e**). **g-h**, Transition probabilities for left 6-0 (**g**) and right 0-6 (**h**) trials using the same clusters derived from all trials as in **f**. Despite identical cues and choices, the paths through the transition matrix of each group were highly variable. **i**, Fraction of clusters visited by left 6-0 (red) and right 0-6 (blue) trials at each epoch decreased to only 0.5, suggesting that much variability remained among these trials, even at the turn. **j**, 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). **k**, Mean overlap fraction of active neurons within the same cluster (intra-cluster) and across clusters (inter-cluster). The overlap fraction was 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. Cells with activity exceeding a z-score threshold of 1.5 were considered 'active'. *** $P < 0.001$, two-sample Student's t-test. Error bars represent mean \pm s.e.m. across datasets. **l**, Fraction of clusters visited at a given epoch when clustering was performed on all epochs together (Methods 5.6.4). All trial types were included. Individual lines represent datasets ($n = 11$). The fraction of clusters visited was similar at all epochs, even at the turn epoch. **m**, 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, such that the cell numbers were the same for each displayed cluster.

Figure 3.6 | (Continued)



Before exploring population dynamics in the cluster space, we first 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 (Figure 3.6a-c). 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 (Figures 3.3a, 3.4, 3.5a-b), the activity patterns in clusters at one epoch were largely different from the activity patterns observed in clusters at the subsequent epoch (Figures 3.6d, 3.9f). 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; Figure 3.7f; Methods 5.6.4). The activity patterns in each cluster were made up of complex combinations of activity levels in the population of individual neurons. Some individual neurons thus had elevated activity in multiple clusters (Figures 3.6m, 3.7i-j, 3.8). A cluster should therefore be considered as a pattern of activity across neurons in the population, such that the patterns between clusters are discriminable from one another. Future work will aim to identify the structure of these activity patterns within clusters, but, importantly, for the focus of this work, clusters were primarily considered an abstracted

grouping of similar activity patterns, and the precise activity patterns that defined each cluster were not important for the subsequent analyses and results.

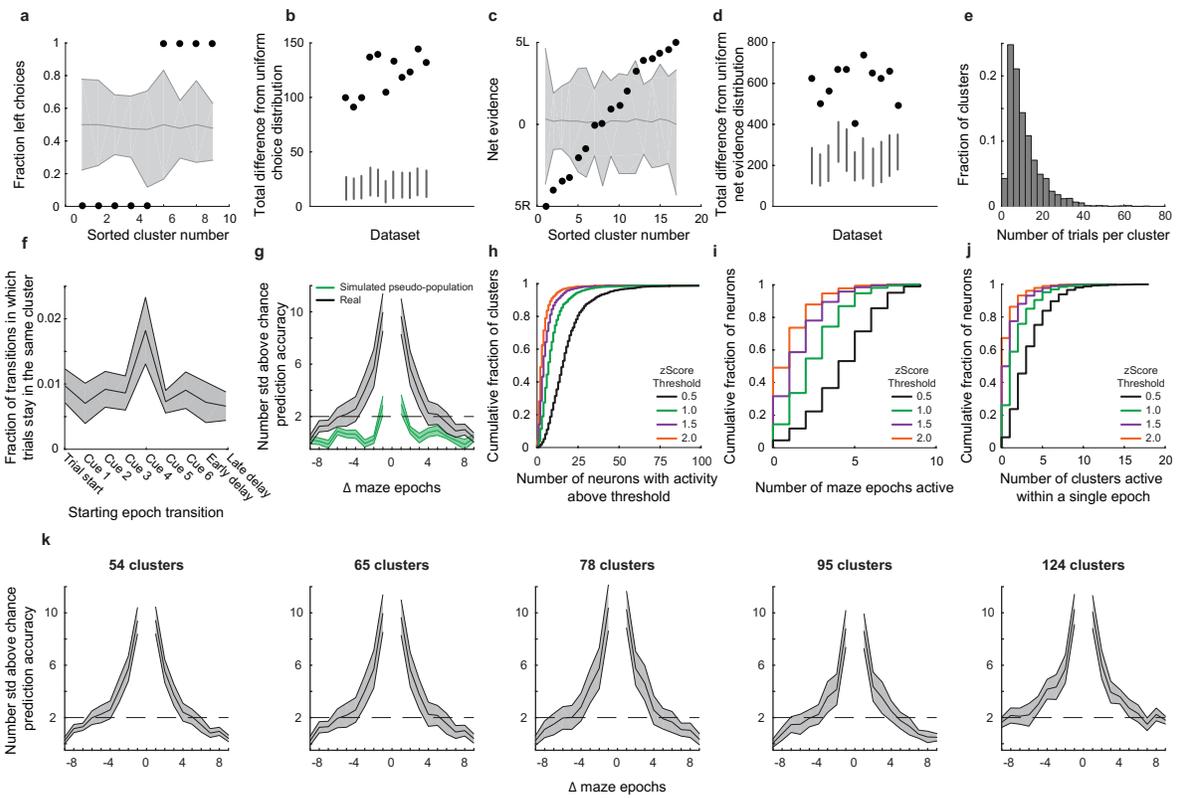
Activity patterns reflecting important task-relevant features, including choice and net evidence, were apparent in the cluster space, even though clustering was performed on neuronal activity alone without any information about behavioral parameters (Figure 3.7a-d). For choice, for example, different paths through clusters emerged for left- and right-choice trials, which is a visualization of the choice-specific activity trajectories we observed previously (Harvey et al., 2012) (Figure 3.6f-h).

We used the cluster space to visualize the population activity trajectories on single trials. This visualization revealed a high amount of trial-trial variability. Trials with the same choice, even during the turn epoch, occupied a diverse set of clusters and traversed different paths through cluster space, suggesting they had distinguishable activity patterns and trajectories. This variability could have been caused by differences in the sequences of cues presented on different trials. We therefore focused on trials with identical evidence cues and choices (e.g., correct 6-0 left trials), which were randomly interleaved with other evidence accumulation trials during imaging (5-1, 4-2, and 3-3 trials). Trials of the same type might be expected to have highly correlated activity patterns and thus converge to similar paths through a small set of clusters. In contrast, the population of trials with identical evidence cues and choices occupied more than half of all possible clusters at each epoch, even at the turn after a choice was made (Figure 3.6g-i).

Figure 3.7 (following page) | 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 \pm s.e.m across datasets. **g**, Based on the cluster identity for a trial at a given epoch, the accuracy of predicting the clusters that trial occupied in the past and future. Real data are shown in black and a simulated pseudo-population is shown in green. To create the pseudo-population, trial identities were shuffled independently for each neuron to break neuron-neuron correlation structure but to preserve each neuron's activity within the trial (Methods 5.7.2). Error bars represent mean \pm s.e.m. across datasets. **h**, Cumulative distribution of the number of neurons active in each cluster for different z-score activity thresholds. **i**, 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. **j**, Cumulative distribution of the number of clusters in which a neuron was active within a single epoch for different z-score activity thresholds. **k**, 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 \pm s.e.m. across datasets.

Figure 3.7 | (Continued)



This initial visualization of single trials in cluster space revealed that different trajectories emerged for different trial types (e.g., left vs. right choice) and that within each trial type extensive variability was present, consistent with previous studies of variability in the activity of cortical neurons (Britten & Shadlen, 1992; Renart & Machens, 2014; Marcos et al., 2013; Churchland et al., 2010, 2011; Maimon & Assad, 2009).

Although the variability in the clusters occupied revealed that trials had distinguishable activity patterns at a given epoch, we wanted to gain more insight into the neuronal basis of this variability. For example, trial-trial variability could have resulted from modulations of the tonic firing of all neurons or from major changes in which sets of neurons were active in each trial. As a first examination of the variability, for each pair of clusters in a given epoch, we calculated the fraction of neurons that were active in both clusters (active neurons were defined by a threshold in z-scored estimated spike counts: threshold = 1.5). Surprisingly, only 10% of neurons on average were active in both clusters, even when limiting our analysis to trials with identical choices and evidence cues. Many trials of the same type therefore had largely non-overlapping populations of active neurons (Figures 3.6k, 3.8). Consistently, the correlation coefficient between the population activity patterns for pairs of trials of the same type at the same epoch (e.g., 6-0 left trials at the turn period) had a wide distribution. Some trial pairs had highly correlated activity, and others had correlation coefficients near zero (Figure 3.6a-c, j). 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 5.6.4). 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 (Figure 3.61). 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. Together, these results suggest that individual trials with the same cues and choices varied greatly to the extent of having largely non-overlapping sets of active neurons.

3.2.4 Temporally structured trial-trial variability

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 therefore focused only on trials of a single type in order to remove the variability due to differences in trial events (e.g., different evidence cues and choices). Using only trials of a single type (e.g., left 6-0 or right 0-6 trials), we generated a new cluster space to examine the within trial type variability (Figure 3.9a). 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. To examine these possibilities, we tested if we could predict the future

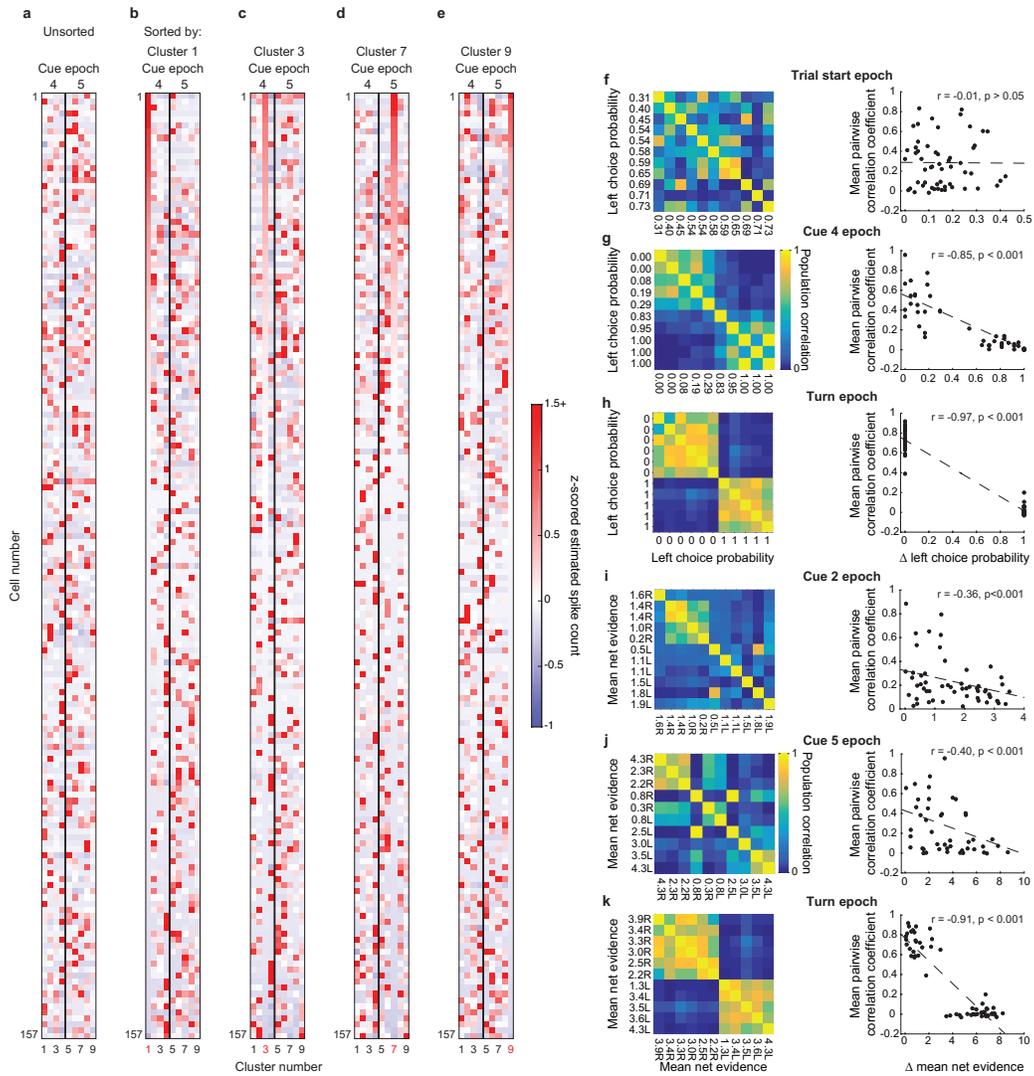


Figure 3.8 | Visualizations of neuronal activity across clusters. **a-e**, Mean z-scored spike count for individual neurons across clusters comprised only of correct left 6-0 trials at two adjacent epochs (Cues 4 and 5) from a single dataset. These plots demonstrate that the activity across clusters 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.

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 at many epochs later in the trial (Figure 3.9b). 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, which is consistent with structure in the activity pattern transitions across time. 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 ('noise') case, the trials starting in a single cluster visited all subsequent clusters, in contrast to what we observed in the real data (Figure 3.9c).

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. 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, and in some cases across the entire trial (Figure 3.9d-e). Similar predictability was

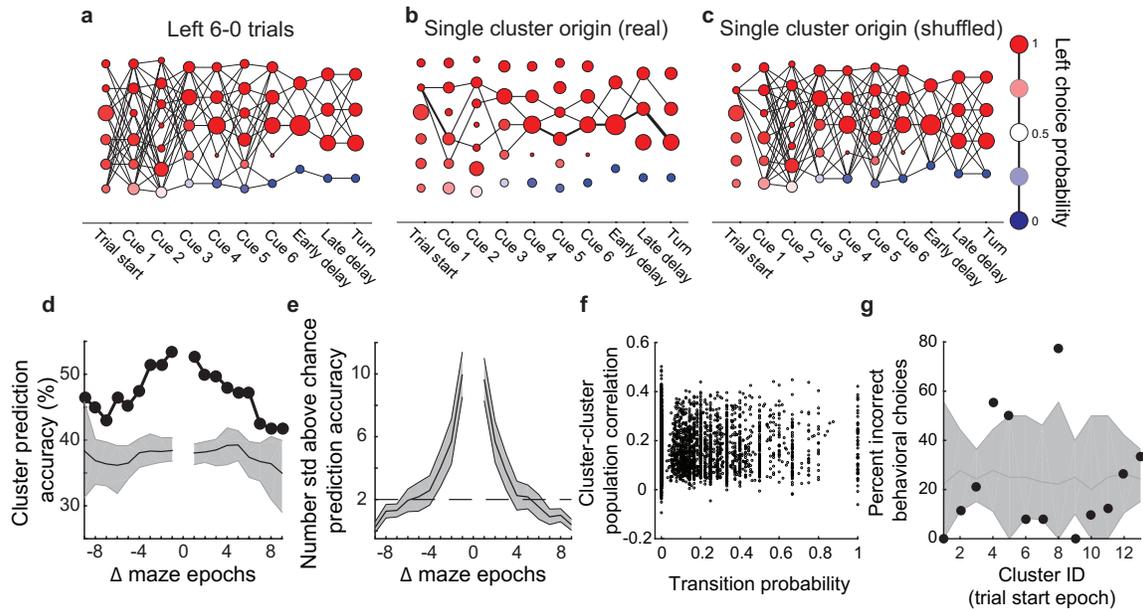


Figure 3.9 | Long timescale temporal structure in PPC activity. **a**, Example transition matrix constructed only from left 6-0 (both correct and error) trials in a single dataset. The nodes at the trial start had high left choice probabilities, even before evidence cues were presented, because the only 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 5.7.1). 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. Error bars represent mean \pm 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 ($r = 0.02$, $p > 0.05$). **g**, Distribution of behavioral error probabilities (i.e. error trials in a cluster divided by total trials in a cluster) across clusters at the trial start for a single dataset. Error bars represent 99% confidence intervals of data in which the assignment of trials to clusters was shuffled. Because many of the clusters contain substantially more or fewer error trials than the shuffled distribution, this distribution is significantly different than chance ($p < 0.01$). For example, trials beginning in cluster 8 result in an error trial close to 80% of the time. Similar results were observed in other datasets.

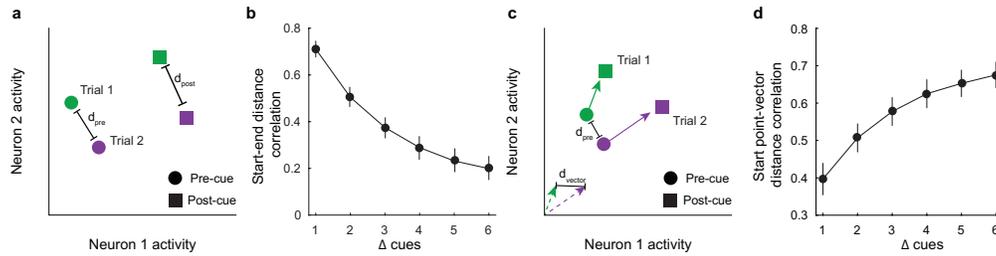


Figure 3.10 | Influence of the current population activity pattern on future population activity patterns and the change in population activity in n -dimensional activity space. **a**, Schematic illustrating the distances between two trials before (d_{pre}) and after (d_{post}) a cue was presented. **b**, Pairwise correlation coefficients between the starting and ending distances across trials with identical cues presented and for different numbers of cues ($\Delta cues$) between the start and end points. At all $\Delta cues$, correlation coefficients were significantly different from a shuffled distribution (Methods 5.5.2; $p < 0.001$), showing that the distance between trials before an identical sequence of cues is presented is predictive of their distance after the cues are presented. Error bars represent mean \pm s.e.m. across datasets. **c**, Schematic illustrating the starting (d_{pre}) distance between two trials and the distance between each trial's vector resulting from cue presentation (d_{vector}). **d**, Pairwise correlation coefficients between the starting distance and vector distance across trials with identical cues presented. Correlation increased with $\Delta cues$. At all $\Delta cues$, correlation coefficients were significantly different from a shuffled distribution (Methods 5.5.3; $p < 0.001$), suggesting that the vector defining the change in population activity in response to a cue depends on both the cue and the starting population activity pattern. Error bars represent mean \pm s.e.m. across datasets.

observed in the n -dimensional activity space (Figure 3.10a-b; Methods 5.5.2). In addition, the vector describing the change in population activity in response to a new cue was also dependent on the current activity pattern, which suggests that dynamics in the population activity affected how inputs influenced ongoing network activity (Figure 3.10c-d; Methods 5.5.3). We performed extensive tests to ensure that the temporal structure was not imposed by the clustering process (Methods 5.7.2).

It is also possible 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 if 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 (Methods 5.7.3). We found that the current behavioral parameters (view angle, maze position, treadmill rotational velocity, which together define the visual scene and running patterns) alone explained above chance, but poorly, past and future activity pattern clusters for only 1 epoch into the past or future (Figure 3.11). 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 significantly better, including across a longer timescale of $\sim 5-6$ epochs (determined by adjusted R^2 to compare models with different numbers of parameters; Figure 3.11). 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 be caused by multiple factors. First, this structure 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 structured transitions from one short-lived population activity pattern to another. Multiple features of the data provided strong support for the second alternative. We found that neurons were transiently active with time-varying activity (Figures 3.3a, 3.4, 3.5a-b). Also, clusters from different epochs had mostly distinct activity patterns (Figure 3.6d, 3.9f). Furthermore, transitions were just as likely between clusters with

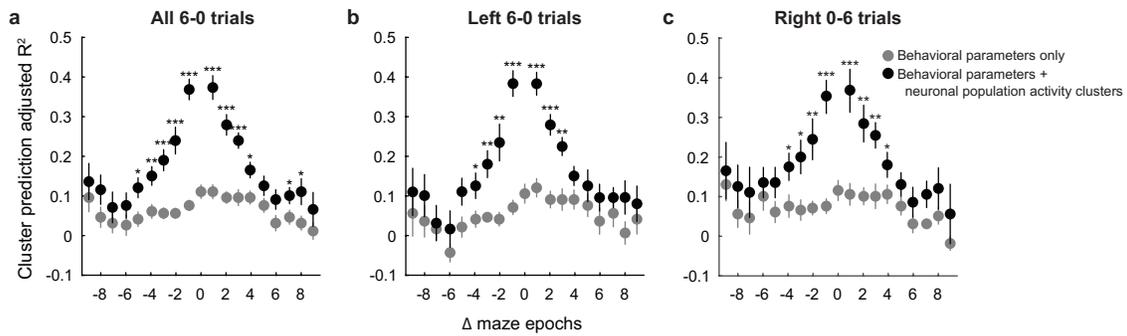


Figure 3.11 | Contribution of behavioral variability to temporally structured trial-trial variability. Our ability to predict the past and future population activity pattern based on the current population activity pattern could not be explained by behavioral variability. We performed a multivariate 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 5.7.3). Consistently, the model based on both behavioral parameters and the neuronal population 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 the results were pooled together to generate the plot for all 6-0 trials (a). The behavioral parameters used were x/y position, x/y treadmill velocity, and view angle as described in Section 2.2. Separate models were fit for each combination of previous and future cluster identities and combined based on the number of maze epochs between them (Δ maze 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.

similar activity patterns as they were between clusters with dissimilar activity patterns (Figure 3.9f). As an additional test of whether the long timescale structure could have emerged from long-lasting activity in individual neurons, we shuffled the trial identities separately for each neuron among trials of the same type. This shuffle broke neuron-neuron correlation structure but preserved 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 (Figure 3.7g). Together these results indicate that the temporal structure in single trials did not arise from long-lasting activity in individual cells; rather, structured moment-to-moment transitions occurred between transient patterns of neuronal activity with largely different sets of active neurons.

Given that long timescale structure existed in PPC activity, we considered the possibility that activity patterns may be predictive of behavioral outcomes long in advance. To test this possibility, we asked if the population activity pattern, even before the first cue was presented, was related to if the mouse made a correct or incorrect choice (Methods 5.7.4). We found that even before the mouse had seen any evidence cues, a small set of clusters contained significantly more error trials than would be expected by chance, suggesting that the activity patterns associated with those clusters were predictive of incorrect choices later in the trial ($p < 0.001$; Figure 3.9g). Also, using an SVM classifier, we could predict the mouse's choice on error trials, but not on correct trials, weakly above chance before any cues were presented ($p < 0.01$ at the first bin a trial; Figure 3.5c).

3.2.5 A memory of past events is maintained in PPC population activity over seconds and across trials

Thus far, our results indicate that long timescale structure exists in the PPC over seconds, not as sustained activity in individual neurons but rather as orderly transitions from one activity pattern to another, often with large variations in which neurons were active. We have shown, based on our ability to predict past and future activity patterns from the current activity, that the activity pattern at a given moment contained information about past activity patterns and also influenced the transition probabilities to future activity patterns over seconds. 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. This new activity pattern would depend on both the features of the event and the activity pattern transition probabilities immediately prior to the event. Because of the long temporal structure that we identified in PPC dynamics, the activity pattern following the event is then expected to influence the transition probabilities to future activity patterns. Because the event helped to generate this new activity pattern, and because this activity pattern influences the transition probabilities to future activity patterns, the event also has an effect on the transition probabilities to future activity patterns. Therefore, by helping to create the activity pattern in the population, a transient event is expected to have a long-lasting effect by constraining the possible future activity patterns. We therefore hypothesized that transient events have signatures of their occurrence long after they ended. In this case, the

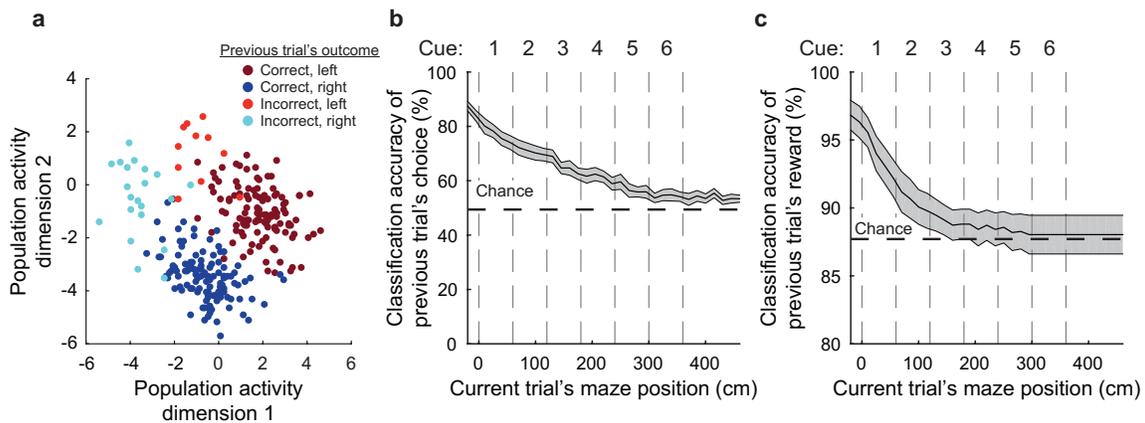


Figure 3.12 | Neuronal population activity in the current trial reflects the previous trial's choice and outcome. **a**, Population activity patterns on different trials at the trial start epoch colored by the choice and outcome (reward or no reward) of the previous trial. Dimensionality was reduced 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 \pm s.e.m. across datasets ($n = 11$).

variability we observed 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). Here, and in following analyses, because we were not directly analyzing transitions between activity patterns, we performed our analyses on the population activity without clustering for simplicity (note that 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 (Bernacchia et al., 2011; Donahue & Lee, 2015; Seo et al., 2007; Seo & Lee, 2007; Marcos et al., 2013). We visualized this result with dimensionality reduction by factor

analysis (Figures 3.12, 3.14) and quantified the result using an SVM classifier based on population activity (Figure 3.12b-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 (Figure 3.12b). 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 \pm s.e.m., $p > 0.05$; Section 2.3.3) (Busse et al., 2011). Mice may have behaved differently (e.g., had different running patterns) in the current trial depending on the choice and reward outcome of the previous trial. In this case, our ability to classify the outcome of the previous trial based on neuronal population activity may simply reflect the representation of current motor variables. Consistently however, 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 (Figure 3.13). 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 (Figure 3.15d).

3.2.6 A novel model for evidence accumulation based on history-dependent dynamics

Our analyses to this point have focused in large part on comparisons between trials of a single type, but the features identified in these analyses have direct implications for evidence accumulation. We have shown that activity patterns in the PPC partially define

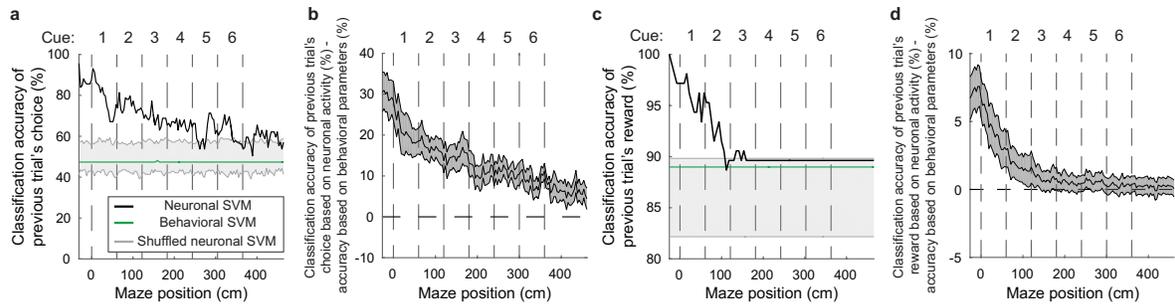


Figure 3.13 | Contribution of behavioral variability to classification of the previous trial's outcome. **a**, Comparison of a neuronal activity-based SVM (black), behavioral parameter-based SVM (green), and the 99% confidence interval of a neuronal activity-based SVM with shuffled labels (gray) for 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. **b**, 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 \pm s.e.m. across datasets. **c-d**, Same as in **(a-b)**, but with classifiers for whether or not the previous trial was rewarded.

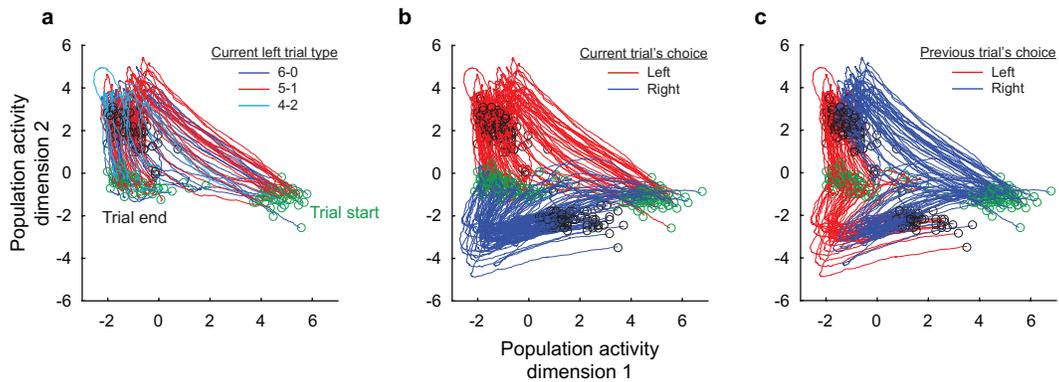


Figure 3.14 | Visualizations of trial trajectories. Trajectories of correct trials colored by the current trial type **(a)**, the current trial's choice **(b)**, and the previous trial's choice **(c)**. Trials with the same choice but different trial types were highly overlapping **(a)**, while trials 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.

the set of possible future activity patterns over seconds (Figure 3.9). 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 3.12). In this framework, we can consider how evidence accumulation might occur. In response to the first evidence cue, the network activity pattern would change based on the type of the evidence cue (left or right cue) and the set of activity pattern transition probabilities at the time of the cue. In response to the second evidence cue, the activity pattern would once again change based on the type of the evidence cue and the set of transition probabilities associated with the current activity pattern. Because the first evidence cue influenced the activity pattern at the time of the second cue, and thus the transition probabilities, the activity pattern resulting after the second cue would be in part a result of both the first and second cues. This same process would be repeated for each subsequent cue. The resulting activity pattern after all cues would therefore be influenced by each previous cue and thus contain information about each of the previous cues. Because this process cascades, the order of the evidence cues would be important. Each unique cue sequence would therefore result in a unique activity pattern, even for the same net evidence. The accumulated evidence cues would be represented generically as a sequence of inputs. In this case, 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. Our results and their extension therefore make predictions about population activity during

evidence accumulation tasks.

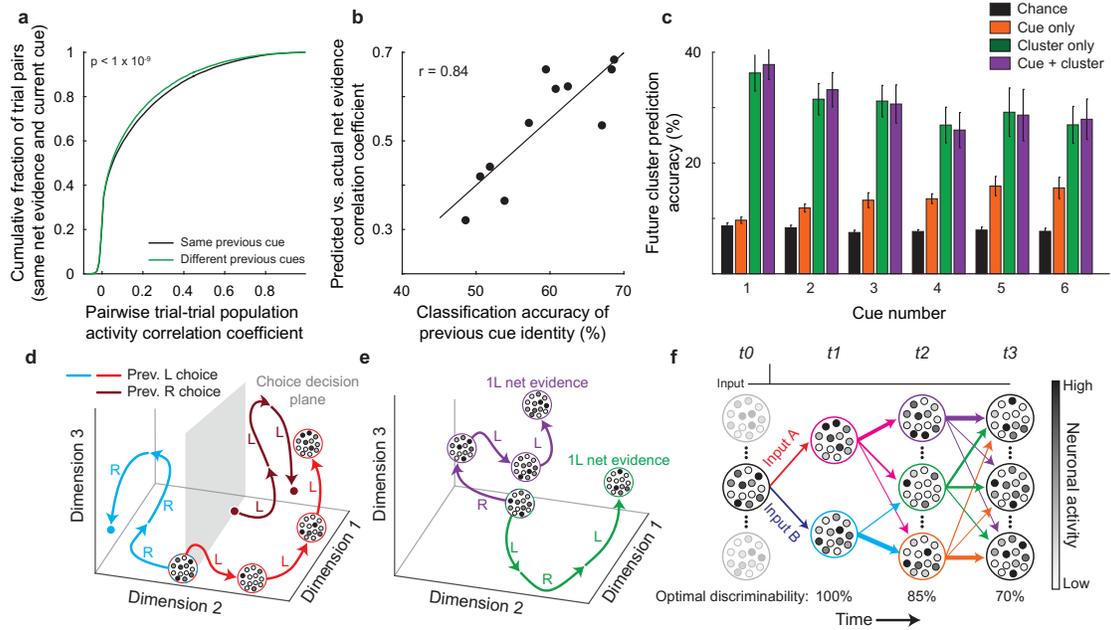
A first prediction is that the population activity pattern should reflect not only the net evidence but also the sequence of evidence cues on a trial independent of net evidence. This prediction implies that different sequences of cues that result in the same net evidence (e.g., left-right-left vs. right-left-left) should generate distinguishable activity patterns (Figure 3.15e). 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; Figure 3.15a). 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: $59.0 \pm 2.2\%$, mean \pm s.e.m.; $p < 0.001$, permutation test with shuffled trial labels). 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 activity difference for distinct evidence sequences could reflect different internal accumulated evidence values due 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, and when we considered data from a mouse that weighted all cues equally (Figure 2.4, see mouse marked in red; for both cases: $p < 0.05$, comparison of pairwise

activity correlations for trials with the same or different previous cue, two-sample KS test; Figure 3.5d). The population activity pattern therefore contained information about the sequence of past evidence cues, independent of net evidence, which is consistent with our findings of long-lasting modifications of dynamics by inputs to the PPC.

Another prediction is that the signal for the sequence of past evidence cues (independent of a signal for 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 imaging 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 ($r = 0.84$, $p < 0.001$; Figure 3.15b). This result indicates that the cue-driven modifications to activity pattern transition probabilities leading to a cue history signal might also serve as the algorithmic mechanism underlying evidence accumulation.

Figure 3.15 (following page) | 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). This analysis tested if neuronal activity at a given epoch contained information about the previous epoch's cue identity, independent of maze epoch, current cue identity, and net evidence. **b,** For each dataset, the ability to classify net evidence (correlation coefficient for predicted vs. actual net evidence using SVRs, e.g., as in Figure 3.3f) 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**). Classification of previous cue and net evidence were highly correlated ($r = 0.84$, $p < 0.001$, $n = 11$ datasets), suggesting they might have related mechanisms. **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), knowledge of the evidence cue only (orange), knowledge of the current cluster identity (green), or both (purple). Error bars represent s.e.m. across 11 datasets. **d,** Schematic illustrating that because the population activity depends on both the inputs and the near-past population activity, trials with the the same sequence of cues, but different starting points due to different past events, will take different paths through activity space and ultimately result in distinguishable activity patterns. For example, the red and dark red traces both experience the same same sequence of cues (*LLL*), but start in different locations due to different trial histories, and therefore have different trajectories through activity space. Each large circle with small circles inside of it represents the activity pattern of the network at a given time point, with each small circle indicating the schematized activity of a neuron. Note that activity patterns are transient and change over the course of a trial (see red trajectory). Despite the existence of multiple activity patterns for the same variable (e.g., 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 network activity by helping to create a new activity pattern with different transition probabilities to future activity patterns. These events therefore influence the set of activity patterns explored over seconds into the future. 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. Each large circle with small circles inside of it represents a possible activity pattern of the network, with each small circle indicating the schematized activity of a neuron. The thickness of each arrow indicates the probability of a transition between two activity patterns. The transitions between $t0$ and $t1$ indicate the change in activity due to one of two inputs. Note that the activity patterns both within each time point and across time points are highly different. Because the transition probabilities are probabilistic, memory of the inputs gradually decays as activity patterns diverge, leading to a decrease in the optimal discriminability of inputs A and B over time.

Figure 3.15 | (Continued)



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 5.7.1). We performed this analysis in the cluster space to facilitate the analysis of transition probabilities between activity patterns. Based on knowing 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). 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; Figure 3.15c). 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.

3.3 Acknowledgements

We thank Selmaan Chettih and Matthias Minderer for developing the cell selection software, Mark Andermann, John Assad, Wolfgang Maass, Ofer Mazor, Stefano Panzeri, and Alexander Trott for helpful discussions, and Bob Datta, Daniel Dombeck, Jan Drugowitsch, Chenghua 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 (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.

3.4 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.

Chapter 4

Discussion and future experiments

4.1 History-dependent population dynamics

We have used new experimental and analysis methods to explore the activity in the PPC neuronal population on single trials of an evidence accumulation task. These methods extend approaches that have been applied previously during evidence accumulation tasks in other animals (Yang & Shadlen, 2007; Gold & Shadlen, 2007; Raposo et al., 2014, 2012; Brunton et al., 2013; Hanks et al., 2015; Scott et al., 2015). Our work has identified two features of PPC activity that together motivate a novel algorithmic 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 a behavioral choice, modified the dynamics of the PPC over a timescale of seconds. Surprisingly, the effect of each event, including evidence cues, was not a change in 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. 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 (Figure 3.15d-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 had an ongoing record of recent past events thus forming a continuous, gap-free short-term memory. Our findings support and extend previous work that described the PPC as an accumulator (Shadlen & Newsome, 1996; Gold & Shadlen, 2000; Yang & Shadlen, 2007; Hanks et al., 2015; Britten & Shadlen, 1992; Horwitz & Newsome, 1999) by proposing that accumulation might occur generally by means of ‘reverberation’ of all 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 of the same type (e.g., 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 (Figure 3.15d-e). 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), with the remaining variance likely

due to biological and measurement noise. The presence of hidden signals impacts our analysis and 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.

4.2 Importance of single-trial population analyses

Our findings therefore underscore the importance of analyzing population activity on single trials. For example, averaging together highly variable trials might obscure important information in the population dynamics. When we simulated a pseudo-population of non-simultaneously recorded neurons (Methods 5.7.2), we lost the ability to detect long timescale dynamics in the PPC. We could not for a given trial predict future and past activity patterns based on the trial’s current population activity pattern (Figure 3.7g), suggesting that simultaneously measured neuron-neuron correlations are critical. Our new experimental and analysis methods thus combine and put into a new context features identified in previous studies, including heterogeneous activity patterns across neurons (Meister et al., 2013; Park et al., 2014; Jun et al., 2010; Raposo et al., 2014; Rigotti et al., 2013), distributed representations of task stimuli including for

non-relevant inputs (Mante et al., 2013; Safaai et al., 2013; Petersen & Diamond, 2000; Rigotti et al., 2013), activity-dependent processing of stimuli (Harris & Thiele, 2011; Safaai et al., 2015; Curto et al., 2009), and the encoding of previous stimuli that indicates stimulus reverberation over time (Bernacchia et al., 2011; Donahue & Lee, 2015; Seo et al., 2007; Seo & Lee, 2007; Nikolić et al., 2009; Klampfl et al., 2012; Seo et al., 2009; Sugrue et al., 2004; Chaudhuri et al., 2015; Murray et al., 2014; Yang & Zador, 2012).

4.3 Inconsistencies with winner-take-all models

Our findings are inconsistent with traditional neuronal implementations of evidence accumulation (Wong & Wang, 2006; Machens, 2005; Wang, 2002). These models propose a winner-take-all competition that evolves over time between distinct pools of neurons, reaching an attractor state once a choice has been made. These models rest on several key predictions that were not apparent in our data. First, traditional implementations of 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). 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. Second, published model implementations propose that the neurons in a given pool have homogeneous and long-lasting activity patterns. Instead, consistent with our previous results (Harvey et al., 2012), we found that neurons in the PPC were highly heterogeneous, with transient and

time-varying activity (Figures 3.3a, 3.4, 3.5a-b, 3.6d). Third, traditional winner-take-all implementations predict that activity eventually reaches the same attractor state for all identical choices, erasing the history of previous events. In contrast, we observed that, throughout a trial, history signals were present for many events, including the sequence of previous cues and outcomes from previous trials, suggesting that the population activity must not have converged to a single attractor for each choice (Figures 3.12, 3.15a). 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 \pm s.e.m. across datasets). Although our results are inconsistent with current neuronal implementations of winner-take-all dynamics, they could be consistent with emerging, but not fully explored, models in which a winner-take-all circuit is embedded within a network with history-dependent dynamics (Klampfl & Maass, 2013) or in which activity in a winner-take-all circuit is drawn towards, but never converges to, dynamically changing attractors.

Our experiments were not designed to test how accumulation might be used for decision-making, such as in relation to drift-diffusion models. Often, accumulation and drift-diffusion have been tightly linked to winner-take-all mechanisms, but we propose that alternative algorithms and implementations, such as the one suggested by our results, could underlie these processes. In this case, our results suggest that the accumulator

would be leaky over the timescale of seconds because the timescale of structured transitions we measured and of the memories of past events decayed over seconds. Also, our results provide the important constraint that a single value for a decision variable would likely be represented by many different activity patterns. Similarly, a decision bound would likely be read out from a distributed representation of many signals such that many different activity patterns might represent the same bound (Figure 3.15d-e).

4.4 A general rule for dynamics in the PPC

We propose a general rule for PPC dynamics in which any input that triggers a change in activity will have a long-lasting effect on future activity patterns due the long timescale dynamics of changes in transition probabilities. 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 (Figure 3.15a). 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 (Buonomano & Maass, 2009; Briggman, 2005; Hoerzer et al., 2014) (Figure 3.15d). The low-dimensional projection in the readout network could be consistent with previous recordings of ramping activity during evidence accumulation tasks (Shadlen & Newsome, 1996; Yang & Shadlen, 2007; Hanks et al., 2015; Britten & Shadlen, 1992; Horwitz & Newsome, 1999; Murakami et al., 2014). This model argues that the PPC has the general role of a ‘reverberator’ of its inputs and that evidence accumulation occurs as a specific example of this general feature. This new model is consistent with the theoretical framework developed in reservoir computing (Jaeger & Haas, 2004; Maass et al., 2002; Buonomano & Maass, 2009; Buonomano & Merzenich, 1995; Verstraeten et al., 2007).

Importantly, 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, such as a point attractor network, that is fine tuned for a specific type of input, such as visual cues during virtual navigation in our case (Wong & Wang, 2006; Machens, 2005; Wang, 2002). In contrast, our proposed model would allow for 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 (Legenstein et al., 2008; Gütig & Sompolinsky, 2006; Jaeger & Haas, 2004; Sussillo & Abbott, 2009; Barak et al., 2013; Sussillo, 2014; Rosenblatt, 1958; Hoerzer et al., 2014; Mante et al., 2013). We consider this advantage important for the PPC, which, as we have shown, contains many signals in the same population of neurons and thus likely contributes to many learned behaviors in parallel.

4.5 Developing neuronal implementations of history-dependent dynamics

The current work primarily advances our understanding of the algorithm performed by the PPC while providing some constraints on the neuronal implementation, such as a distributed representation of task features, time-varying activity, and predictable transitions between activity states defined by complex patterns of activity in the neuronal population. However, to further improve our understanding of this model's plausibility and how information is processed in the PPC, future work will be required to generate a variety of potential neuronal implementations of this algorithm. These neuronal implementations could be based on a variety of potential architectures. However, they should each make experimentally testable predictions which will allow us to distinguish between them.

4.6 Additional experimental tests

Our work suggests that layer 2/3 of the PPC may serve as a general purpose ‘reverberator’ whose goal is merely to maintain an implicit, distributed memory of inputs for as long as possible. This architecture would be independent of specific tasks and could therefore be generated early in life via mechanisms such as unsupervised learning. As an animal learns a task, downstream readout networks would gradually converge to the appropriate readout weights to optimally recover the identity of task-relevant inputs (Legenstein et al., 2008; Gütig & Sompolinsky, 2006; Jaeger & Haas, 2004; Sussillo & Abbott, 2009; Barak et al., 2013; Sussillo, 2014; Hoerzer et al., 2014). One potential advantage of this model is that, because the reward signals used in mechanisms such as reinforcement learning are sparse (Sutton & Barto, 1998), it may be easier to learn a limited set of readout weights than a large number of recurrent weights.

If activity in layer 2/3 of the PPC serves as a general purpose reverberator of inputs in a task-independent manner, its activity dynamics should not change much over the course of learning. One would therefore expect that, even in animals early in training, task-relevant events would remain discriminable over seconds. To test this result experimentally, one could record the activity of neuronal populations before, during, and after learning of a new task. If layer 2/3 of the PPC does in fact play the role of a general purpose reverberator of inputs, the activity dynamics should not change. However, if the internal PPC dynamics remain unchanged over learning, but the impact of task-relevant events on

network activity increases over learning, task-relevant events may not be discriminable early in training simply because they do not exert a pronounced impact on PPC activity. The results of this experiment may therefore be difficult to interpret.

As a related analysis, one could train individual mice to perform multiple tasks within the same behavioral session. If layer 2/3 of the PPC serves as a general purpose reverberator of inputs, one would again expect that the activity patterns in the PPC would be remain unchanged and therefore be similar in both tasks. Recently, Mante and colleagues observed that during a variant of the random dot motion task in which monkeys had to discriminate either color or motion on alternating blocks of trials, population activity in the PFC was similar regardless of whether motion or color was behaviorally relevant (Mante et al., 2013). This result is consistent with the predictions of our model, but further tests will be required to determine whether this result generalizes to the PPC and other tasks.

Because our model postulates that the primary change over the course of learning is the appropriate modification of readout weights, the prevalence of readout networks should gradually increase over the course of learning. If these readout networks include the neurons recorded previously (Gold & Shadlen, 2007; Hanks et al., 2015; Shadlen & Newsome, 1996), we would expect these neural populations to grow in size, increase their selectivity, or both during task learning. To test this experimentally, one could record from populations of neurons known to exhibit task selectivity before, during, and after learning. However, because these predictions are relatively unconstrained, the results of these experiments may be difficult to interpret. For example, it is possible that individual

neurons with high task selectivity exist merely as a byproduct of the history-dependent dynamics described in Chapter 3. This would suggest that an explicit readout network would not be necessary. However, an implicit readout network, in which a high-dimensional, distributed representation is transformed into a different high-dimensional, distributed representation might still be present. This result would still be consistent with both our data and our model. We must therefore work to further constrain these models to generate clear, testable predictions.

A variety of studies have recently provided support for the notion of a ‘hierarchy of timescales’ across cortex, in which the time constant of neuronal activity is shorter in areas closer to the sensory periphery, such as visual and auditory cortex, and longer in association cortex, including PPC and PFC (Bernacchia et al., 2011; Murray et al., 2014; Yang & Zador, 2012). This work has primarily focused on the time constant of activity for individual neurons. However, one could use analyses similar to those described in Chapter 3 to instead measure the time constant of *predictability* for population activity in different cortical areas. For example, if such a hierarchy exists, one might find that trial-trial variability is predictive of future variability for longer in the PPC than in V1. While this result would be consistent with the interpretation of previous work, it would suggest a substantively different mechanism based on history-dependent population dynamics.

Because our model proposes that the PPC acts in a general-purpose manner, we would expect that certain inputs to layer 2/3 of the PPC are not privileged over others. In other words, inputs which are behaviorally relevant should not have a greater impact on PPC

activity than inputs which are behaviorally irrelevant. To test this, one could design a task which contains signals which are behaviorally irrelevant by design. For example, one could use a variant of the fixed association evidence accumulation task in which cues can either be black or white, but still appear on either the left or the right side of the T-maze. The mouse's goal would still be to determine whether more cues were on the left or the right, regardless of the color of each cue. The location of each cue would therefore be behaviorally relevant, while its color would be behaviorally irrelevant. One could then ask whether cue color could be decoded from PPC population activity independent of cue location. Because cue location and color may have different visual salience, one would have to perform the experiment in the opposite contingency as well (e.g., in which cue color is behaviorally relevant, and cue location is behaviorally irrelevant).

4.7 Recreating history-dependent dynamics *in silico* with recurrent neural networks

As a proof of principle experiment, it would be useful to demonstrate that such a model could successfully perform tasks *in silico*. The first step of this experiment would be to train an artificial neural network to recapitulate the dynamics observed *in vivo*, independent of a specific task. As a first test, one could use a recurrent neural network (RNN) without spiking (i.e., a rate network). One advantage of such networks is that effective techniques for learning appropriate recurrent weights developed in computer science, such as backpropagation through time, can be used (LeCun et al., 2015). However, in contrast to

traditional supervised learning approaches in which RNNs are trained to produce a desired output, the RNN in this experiment would be trained primarily to maintain discriminability of inputs for as long as possible. Importantly, because neurons in layer 2/3 of the PPC exhibited transient dynamics (Figures 3.3a, 3.4, 3.5a-b), an additional constraint imposing transiency would be added. This constraint would significantly increase the difficulty of this task by forcing the RNN to maintain discriminability of inputs through a variety of highly different activity patterns, as we observed. These objectives can be formalized as a differentiable loss function which penalizes the prolonged activity of individual neurons and rewards discriminability of inputs over long timescales (for example, by encouraging the maximization of the Euclidean distance between trial pairs with different inputs long after input offset). To test the similarity of the RNN's dynamics with those observed *in vivo*, analyses similar to those presented in Chapter 3 could be used to ensure consistency with experimental results.

Once an RNN has been trained to produce history-dependent dynamics, the weights of the RNN can be fixed. Readout networks would then be trained to perform a variety of tasks as inputs are delivered to the RNN, including evidence accumulation and delayed-match-to-sample tasks like those described in Chapter 2. A variety of learning rules could be learned, including those which are biologically implausible, such as backpropagation, as well as biologically plausible learning rules, such as those based on local synaptic computations (Legenstein & Maass, 2014). If the same RNN with the same weights can be used to perform a wide variety of tasks effectively (with different readout

networks for each task), it would serve as an effective proof of principle that models based on reservoir computing are plausible.

RNNs trained to recapitulate history-dependent dynamics could further be used to generate novel experimental predictions. For this purpose, artificial neural networks (ANNs) possess two key advantages over actual neural networks. First, ANNs provide the experimenter with a perfect, noiseless measurement of the activity of every neuron in the network at every time point. In contrast, our methods for monitoring activity in real neural networks are limited both in the number of neurons we can simultaneously record as well as the noise injected by our measurement techniques. Second, ANNs can be easily manipulated at a fine scale. For example, in section 3.2.2, we found that task-relevant information was distributed across large populations of neurons, including those neurons that were individually non-selective for task features (Figure 3.3). However, we cannot directly test *in vivo* whether neurons which have high choice selectivity are necessary for the neuronal computation. In an ANN, however, we could easily set the activation of all highly selective neurons to zero. The deficit in network performance in response to this targeted inactivation would directly test the necessity of these neurons for these computations. Of course, ANNs also have a critical disadvantage: they are artificial and may exhibit activity dynamics which are radically different from those present in real neural networks. The advantages provided by ANNs should therefore be used primarily for rapid *in silico* experiments, with the ultimate goal of generating predictions which can be tested in real neural networks *in vivo*.

Chapter 5

Methods

5.1 Subjects

All experimental procedures were approved by the Harvard Medical School Institutional Animal Care and Use Committee and were performed in compliance with the Guide for the Care and Use of Laboratory Animals. Data were acquired from five male C57BL/6J mice (Jackson Labs), which were 8-10 weeks old at the start of behavioral training, and 14-22 weeks old during imaging. Prior to training, 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 μ l of water each day (total, including rewards). Each mouse's weight was measured daily to ensure that it was 80% of the mouse's pre-water-restriction weight.

5.2 Imaging

5.2.1 Surgical procedure

When mice performed well on maze 7 (Figure 2.3), they underwent a surgery (isoflurane anesthesia) to implant a cranial window. For three days prior to surgery, mice were given 5 mL of water per day. The behavioral training headplate was removed, and a circular craniotomy with a diameter of 3.1 mm was made over PPC on the left hemisphere (stereotaxic coordinates: 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 μ m spacing between injections). Viruses were obtained from the University of Pennsylvania Vector Core Facility. Injections were made near the center of the craniotomy, ~275 μ m below the dura, using a beveled glass pipette (~15 μ m 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 (Dombeck et al., 2010). Following at least one day of recovery, mice resumed training. Imaging began at least 4 weeks post-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; Figure 3.2).

5.2.2 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 large (20 mm) back aperture objectives. Collection optics were housed in an aluminum box to prevent light interference from the VR display. 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 (Chameleon Vision II, Coherent) operated at 920 nm. The microscope was controlled by ScanImage (version 5; Vidrio Technologies) (Pologruto et al., 2003).

5.2.3 Imaging data acquisition

Imaging data were acquired at 30 Hz at a resolution of 512 x 512 pixels ($\sim 700 \mu\text{m} \times \sim 700 \mu\text{m}$ field-of-view) using a Nikon 16x 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. Imaging data were acquired in sets of 25,000 frames with brief breaks between acquisitions to ensure alignment of the scans. Up to 100,000 frames were acquired from each imaging session over the course of ~ 1 hour. Imaging data were acquired from single planes at depths between 100 and 200 μm below the dura. Multiple fields-of-view were acquired from the same mouse across different days. Data were analyzed from 11 fields-of-view from 5 mice.

5.2.4 Pre-processing 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 MATLAB software (manuscript in preparation). In brief, following motion correction (Greenberg & Kerr, 2009), the correlation of fluorescence timeseries 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 (Shi & Malik, 2000) 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 timeseries against low-activity regions of the cell timeseries, and the scaled background timeseries was then subtracted from the cell timeseries. 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) (Vogelstein et al., 2010), which minimized the impact of the indicator's decay kinetics on our analyses. Similar results were obtained from the non-deconvolved $\Delta F/F$ traces (Figure 5.1).

5.3 Data analysis

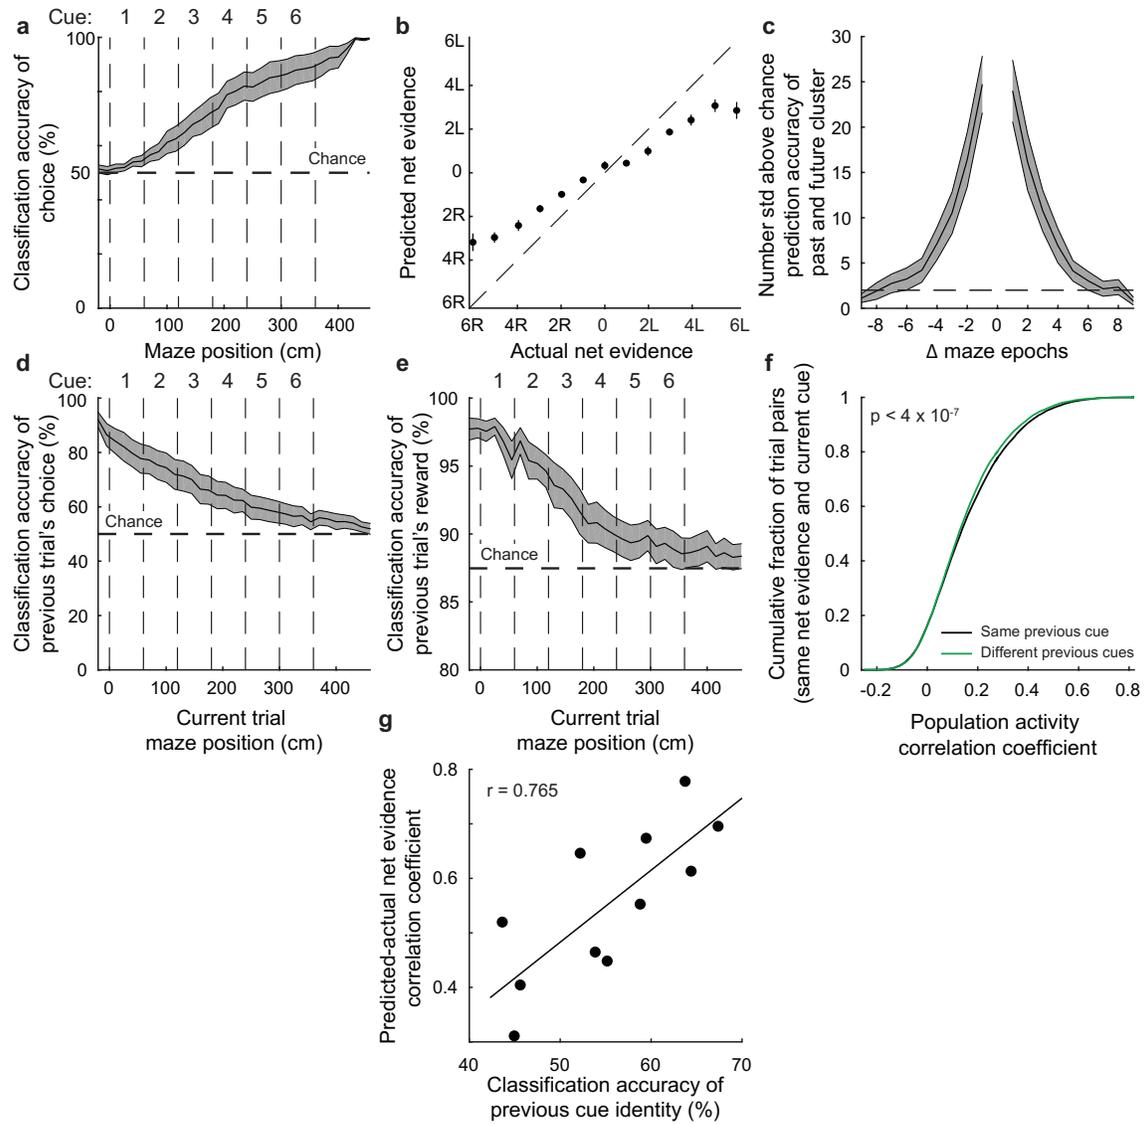
5.3.1 General analysis procedures

Data were grouped into spatial bins (3.75 cm/bin) corresponding to locations in the virtual maze. To bin data from positions in the arms of the T-maze, the T-maze was linearized prior to binning by folding the arms such that they were a continuation of the stem. Neuronal activity and behavioral parameters were averaged in each bin. On average, each bin contained 2-3 imaging frames per trial. Unless otherwise noted, all analyses were performed on both correct and error trials together. All correlation coefficients were from Pearson's correlations. Portions of this research were conducted on the Orchestra High

Performance Compute Cluster at Harvard Medical School. This NIH supported shared facility consists of thousands of processing cores and terabytes of associated storage and is partially provided through grant NCRR 1S10RR028832-01.

Figure 5.1 (following page) | Main results re-analyzed using $\Delta F/F$ 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 \pm s.e.m. across datasets. Compare to Figure 3.3e. **b**, Actual net evidence vs. net evidence predicted by a SVR classifier. Error bars represent mean \pm s.e.m. across datasets. Compare to Figure 3.3f. **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 \pm s.e.m. across datasets. Compare to Figure 3.9e. **d-e**, Classification accuracy as in (a), but for previous trial's choice and for whether the previous trial was rewarded (**e**). Compare to Figure 3.12b-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 Figure 3.15a. **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 Figure 3.15b.

Figure 5.1 | (Continued)



5.3.2 Choice selectivity for individual neurons (Figure 3.3a-c)

Normalized activity was calculated for each neuron individually by dividing its estimated spike count within each bin on each trial by its maximum across all bins of all trials (Figure 3.3a-b). To calculate the left-right choice selectivity index (Figure 3.3c), for each neuron individually, we calculated its average activity for each spatial bin separately across left 6-0 trials and right 0-6 trials. The selectivity index was defined as:

$$\text{Selectivity index} = \frac{\mu_{\text{left}6-0} - \mu_{\text{right}0-6}}{\mu_{\text{left}6-0} + \mu_{\text{right}0-6}}$$

A selectivity index of +1 corresponds to a neuron which is active exclusively on left 6-0 trials, while a selectivity index of -1 corresponds to a neuron which is active exclusively on right 0-6 trials. A selectivity index of 0 corresponds to a neuron with equivalent mean activity in both conditions. The selectivity index for each neuron was calculated separately for each spatial bin, and the peak selectivity index was calculated as the maximum of the absolute value of the selectivity index across all bins. For the shuffled distribution, we shuffled the trial labels (i.e. which trials were left 6-0 or right 0-6), and recalculated the selectivity index as above.

5.4 Classifiers (without clustering)

5.4.1 General procedures (Figures 3.3e-h, 3.5c, e-h, 3.12b-d, 3.13)

Unless otherwise noted, all population classifiers were support vector machines (SVMs) with a radial basis function (Gaussian) kernel (Murphy, 2012; Smola & Vapnik, 1997).

All population classification was performed on the concatenated activity of all individual neurons. All SVMs were implemented using the MATLAB interface of the open-source libsvm library (version 3.20) (Chang & Lin, 2011). 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 10-fold cross-validation on the training/validation set of only a single dataset, and the same hyperparameters were used for all datasets. For classifiers applied to the activity of individual neurons (Figure 3.3d), classifiers were trained with the activity of the given neuron as the only input feature. For population classifiers (Figures 3.3e-h, 3.5c, e-h, 3.12b-d, 3.13), the activity of all neurons in the population were used as the input to the SVM.

5.4.2 Two-class classifications (Figures 3.3e, g, 3.5c, 3.12, 3.13 5.1a, d-e)

For two-class classification problems, such as the classification of the choice (Figures 3.3e, g, 3.5c), previous trial's choice (Figure 3.12b), and previous trial's reward outcome (Figure 3.12c), a C-Support Vector Classification (C-SVC) approach was used. Independent SVMs were trained on each spatial bin. In cases where classification accuracy based on neuronal data was compared with accuracy based on behavioral data (Figure 3.13), hyperparameters were optimized separately for each.

5.4.3 Regression for net evidence (Figures 3.3d, f, h, 3.5e-h, 5.1b)

For the regression of net evidence (Figures 3.3d, f, h, 3.5e-h), an ϵ -Support Vector Regression (ϵ -SVR) approach was used to obtain a continuous approximation of net evidence (Murphy, 2012; Smola & Vapnik, 1997; Chang & Lin, 2006). 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. Training and testing sets were divided based on whole trials to prevent similar activity on different cues within the same trial from corrupting the results. To determine prediction accuracy, the predicted class label was compared to the actual class label via a correlation coefficient. To calculate statistical significance, the results were compared to the distribution resulting from 1000 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 ($p < 0.001$; Figure 3.3f). To exclude the impact of view angle alone, which changes the visual scene, on net evidence classification, we recalculated the ϵ -SVR separately for left and right trials with nearly identical view angles ($\pm 2.5^\circ$; Figure 3.5e-f). To exclude the impact of view angle in conjunction with other behavioral parameters, such as the x/y position in the maze and the rotational velocity of the spherical treadmill, we trained SVR classifiers to predict net evidence based on either

behavioral parameters alone or behavioral parameters in addition to neuronal population activity (Figure 3.5g-h).

5.4.4 Classifiers built by adding-in subsets of neurons (Figure 3.3g-h)

Classifiers were built similar to the population classifiers described in Methods 5.4.1, except with the input being a subset of neurons. The single neuron choice selectivity and SVR net evidence correlation were used to determine single neuron selectivity for choice and net evidence, respectively. Neurons were sorted in ascending order based on their selectivity. A separate classifier was trained for increasingly larger populations of neurons with neurons added in from least to most selective.

5.4.5 Classification of cue sequences (Figures 3.5d, 3.15a-b, 5.1g-h)

To determine if the identity of previous cues could be read out at a current epoch, given identical current cue and net evidence (Figures 3.5d, 3.15a-b), we examined pairwise trial-trial population activity correlations. At each of cues 3-6, we separated trial into those which contained the patterns *LRL* (left-right-left) and *RLL* (right-left-left) or *LRR* and *RLL* 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 eight groups (2 choices x 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 based on 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; whichever previous cue had a higher correlation with the test trial was considered the prediction of the classifier. To determine significance, accuracy was compared to the distribution of accuracies resulting from 1000 shuffles of the labels assigning previous cues to trials.

5.5 Analysis of activity in high-dimensional state space

5.5.1 Factor analysis (Figures 3.12a, 3.14)

For visualizations using factor analysis (Figures 3.12a, 3.14), dimensionality was reduced to 5-factors using built-in MATLAB toolboxes. Two factors were selected from these five for visualization.

5.5.2 Relationship between pairwise trial-trial distances before and after cue presentation (Figure 3.10a-b)

To calculate the correlation between the distance between trials before and after a cue(s) (Figure 3.10a-b), trials were divided into groups based on the number of cues presented between measurements (e.g. $\Delta 2$ cues is equivalent to measuring before the first cue and after the second cue or measuring before the third cue and after the fourth cue). Within each group, only trials with the same marginal cues and starting net evidence were

compared. The Euclidean distance in n-dimensional activity space between pairs of trials was calculated before (pairwise starting distance) and after (pairwise ending distance) the marginal cues were presented. Across all pairs of trials within each group, the Pearson correlation coefficient of the population activity on each trial was calculated. To determine significance, we generated new ending states for each trial by shuffling the distribution of vectors which define the change in activity in response to cues and adding them to the original starting states. Vectors were calculated by subtracting the n-dimensional vector corresponding to the starting state from the vector corresponding to the ending state. This step was necessary to rule out the possibility that the network activity changed little in comparison to the original distance between trials. For example, if two trials were separated by a distance of 10 estimated spikes, but the activity only changed by 0.1 estimated spikes, the pairwise ending distance would be expected to remain strongly correlated with the pairwise starting distance, regardless of the direction of the vectors. A p-value was then calculated by comparing the correlation coefficient to the distribution obtained from 1000 such shuffles.

5.5.3 Relationship between pairwise trial-trial distances and population activity vectors (Figure 3.10c-d)

To calculate the correlation between the distance between trials before the presentation of a cue and the vector which results from the response to the cue (Figure 3.10c-d), trials were divided into groups and paired in the same fashion as described above. Vectors were calculated as described above. For each pair of trials, the similarity between vectors was

calculated by determining the Euclidean distance between the two vectors. Across all pairs of trials within each group, the Pearson correlation coefficient between the pairwise starting state distance and the pairwise vector distance was then calculated. Significance was determined by comparing the correlation coefficient to the distribution obtained from 1000 shuffles of the vector trial identities.

5.6 Clustering methods (Figures 3.6, 3.7, 3.8, 3.9, 3.11 3.15c, 5.1c, f)

We used clustering to reduce the dimensionality of the population activity in as unbiased a fashion as possible, without inclusion of information about behavioral parameters and without encouraging projection of the data onto dimensions which 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 for the calculation of transition probabilities between discrete activity states. However, clusters were not considered an indication of discreteness of the underlying activity patterns. Clustering revealed groups of trials with similar population activity patterns (Figure 3.6a-c).

5.6.1 Pre-processing 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 4 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 (4 spatial bins). For the early and late delay, respectively, activity was averaged across the 4 spatial bins beginning 15 and 37.5 cm after offset of the final cue. For the turn, activity was averaged across the final 4 spatial bins in the maze. Each epoch corresponded to approximately 200 ms.

5.6.2 Clustering via affinity propagation

Within each epoch, trials were clustered into groups based on their neuronal activity using affinity propagation (Frey & Dueck, 2007). 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 is 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 datasets 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 (Figure 3.7g). Clustering was only used for analyses in which dimensionality reduction was necessary. Affinity propagation was

performed using code provided by the laboratory of Brendan Frey.

5.6.3 Transition probabilities between clusters (Figures 3.6, 3.7, 3.9)

To calculate transition probabilities between clusters at different maze epochs (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 (Figures 3.6e-h, 3.9a-c), for each cluster, we calculated the fraction of trials that had a given feature. For example, to superimpose left choice probability, we calculated the fraction of trials in each cluster that resulted in a left choice. 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 (Figure 3.7a-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 1000 shuffles (Figure 3.7b, d). To determine whether there was a relationship between the similarity of clusters at adjacent epochs and the transition probability between them, for each pair of adjacent clusters, we calculated the mean pairwise trial-trial population activity correlation and compared it to the transition probability (as defined above). Transitions were not significantly more likely between more similar clusters ($r = 0.02$, $p > 0.05$).

5.6.4 Clustering based on all time points together (rather than epoch-by-epoch clustering) (Figures 3.6l, 3.7f)

To determine if 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 (Figure 3.7f). 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 (Figure 3.6l). To remove outliers, only clusters containing three or more trials at a given epoch were counted as visited.

5.7 Classifiers based on activity in cluster-space

5.7.1 Classification of the cluster identity at past and future epochs based on the cluster identity at the current epoch (Figures 3.7k, 3.9d-e, 3.15c, 5.1c)

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 (Figure 3.9d-e). In a leave-one-out fashion, we limited our analysis to only the trials which 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 3.9d-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 calculated across both of these conditions. To establish statistical significance, accuracy was compared to 1000 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 in Figure 3.15c, except in Figure 3.15c, it was applied to all trials together, independent of evidence or choice. The other classifiers used in Figure 3.15c 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 with both the same current cue as the test trial and which were in the same cluster at epoch i were used to determine the most likely cluster at epoch j .

5.7.2 Classification of past and future cluster identities with simulated pseudo-populations (Figure 3.7g)

The temporal structure we observed could be caused by the prolonged activity of individual neurons, either due to persistent activity in the underlying neuronal activation or to the prolonged decay kinetics of the calcium indicator. To test if 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 (e.g. correct left 6-0 trials; Figure 3.7g). 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 (Figure 3.7g), suggesting that the predictability we observed in the real data could not be explained by prolonged activity or slow indicator kinetics.

5.7.3 Classification of past and future cluster identities based on behavioral data (Figure 3.11)

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 (Figure 3.11). To allow for 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 based on 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 based on the number of epochs separating the predictor and response (e.g. 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.

5.7.4 Classification of behavioral errors based on cluster identity (Figures 3.9g)

To determine whether behavioral errors were uniformly distributed across clusters (Figure 3.9g), we clustered trials based on their population activity patterns at the trial start, before any cues had been presented. Because differences in activity were likely to be subtle, we

used a conservative clustering approach, which resulted in more clusters. Thus, we used a preference parameter equivalent to the thirtieth percentile of the distance matrix. The distribution of the probability of a behavioral error across these clusters was then calculated. To determine statistical significance, we calculated the expected number of errors in each cluster based on the overall error rate and the number of trials in each cluster and calculated the total absolute difference from the uniform expected error rate in each cluster. This metric was compared to a shuffle of the assignments of trials to clusters.

5.8 Analysis of the overlap of active neurons across clusters (Figure 3.3k)

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 (Figure 3.3k), 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:

$$\text{Overlap fraction} = \frac{\text{number of neurons with activity above threshold in both clusters}}{\text{number of neurons with activity above threshold in either cluster}}$$

The mean value across all pairs of clusters at all epochs was calculated. 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, 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. The mean across all clusters was then calculated. 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 1000 times, and the inter-cluster overlap fraction was recalculated.

References

- Akerboom, J., Chen, T. W., Wardill, T. J., Tian, L., Marvin, J. S., Mutlu, S., Calderon, N. C., Esposti, F., Borghuis, B. G., Sun, X. R., Gordus, A., Orger, M. B., Portugues, R., Engert, F., Macklin, J. J., Filosa, A., Aggarwal, A., Kerr, R. A., Takagi, R., Kracun, S., Shigetomi, E., Khakh, B. S., Baier, H., Lagnado, L., Wang, S. S. H., Bargmann, C. I., Kimmel, B. E., Jayaraman, V., Svoboda, K., Kim, D. S., Schreier, E. R., & Looger, L. L. (2012). Optimization of a GCaMP Calcium Indicator for Neural Activity Imaging. *Journal of Neuroscience*, 32(40), 13819–13840.
- Andersen, R. A. & Cui, H. (2009). Intention, action planning, and decision making in parietal-frontal circuits.
- Aronov, D. & Tank, D. W. (2014). Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron*, 84(2), 442–456.
- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.
- Balan, P. F. (2006). Integration of Exogenous Input into a Dynamic Saliency Map Revealed by Perturbing Attention. *Journal of Neuroscience*, 26(36), 9239–9249.
- Barak, O., Sussillo, D., Romo, R., Tsodyks, M., & Abbott, L. F. (2013). From fixed points to chaos: Three models of delayed discrimination. *Progress in Neurobiology*, 103, 214–222.
- Bellman, R. E. (1961). *Adaptive Control Processes*. Rand Corporation.
- Bennur, S. & Gold, J. I. (2011). Distinct Representations of a Perceptual Decision and the Associated Oculomotor Plan in the Monkey Lateral Intraparietal Area. *Journal of Neuroscience*, 31(3), 913–921.
- Bernacchia, A., Seo, H., Lee, D., & Wang, X.-J. (2011). A reservoir of time constants for

- memory traces in cortical neurons. *Nature Neuroscience*, 14(3), 366–372.
- Bisley, J. W. (2003). Neuronal Activity in the Lateral Intraparietal Area and Spatial Attention. *Science*, 299(5603), 81–86.
- Bisley, J. W. & Goldberg, M. E. (2010). Attention, Intention, and Priority in the Parietal Lobe. *Annual Review of Neuroscience*, 33(1), 1–21.
- Blatt, G. J., Andersen, R. a., & Stoner, G. R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. 299(4), 421–445.
- Briggman, K. L. (2005). Optical Imaging of Neuronal Populations During Decision-Making. *Science*, 307(5711), 896–901.
- Britten, K. H. & Shadlen, M. N. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. 12(December), 4745–4765.
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science*, 340(6128), 95–98.
- Buneo, C. a. & Andersen, R. a. (2006). The posterior parietal cortex: Sensorimotor interface for the planning and online control of visually guided movements. *Neuropsychologia*, 44(13), 2594–2606.
- Buonomano, D. & Merzenich, M. (1995). Temporal information transformed into a spatial code by a neural network with realistic properties. *Science*, 267(5200), 1028–1030.
- Buonomano, D. V. & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, 10(2), 113–125.
- Busse, L., Ayaz, A., Dhruv, N. T., Katzner, S., Saleem, A. B., Schölvinck, M. L., Zaharia, A. D., & Carandini, M. (2011). The detection of visual contrast in the behaving mouse. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(31), 11351–11361.

- Chang, C.-C. & Lin, C.-J. (2006). Training ν -Support Vector Regression: Theory and Algorithms. *dx.doi.org.ezp-prod1.hul.harvard.edu*, 14(8), 1959–1977.
- Chang, C. C. & Lin, C. J. (2011). LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and ...*
- Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H., & Wang, X.-J. (2015). A Large-Scale Circuit Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron*.
- Chen, T.-W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., Schreiter, E. R., Kerr, R. a., Orger, M. B., Jayaraman, V., Looger, L. L., Svoboda, K., & Kim, D. S. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458), 295–300.
- Churchland, A. K. & Ditterich, J. (2012). New advances in understanding decisions among multiple alternatives. *Current Opinion in Neurobiology*, 22(6), 920–926.
- Churchland, A. K., Kiani, R., Chaudhuri, R., Wang, X.-J., Pouget, A., & Shadlen, M. N. (2011). Variance as a Signature of Neural Computations during Decision Making. *Neuron*, 69(4), 818–831.
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693–702.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., & Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., Priebe, N. J., Finn, I. M., Ferster, D., Ryu, S. I., Santhanam, G., Sahani, M., & Shenoy, K. V. (2010). Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369–378.
- Cohen, M. R. & Kohn, A. (2011). Measuring and interpreting neuronal correlations. *Nature*

Neuroscience, 14(7), 811–819.

Cohen, M. R. & Maunsell, J. H. R. (2009). Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*, 12(12), 1594–1600.

Crowe, D. A., Averbeck, B. B., & Chafee, M. V. (2010). Rapid Sequences of Population Activity Patterns Dynamically Encode Task-Critical Spatial Information in Parietal Cortex. *Journal of Neuroscience*, 30(35), 11640–11653.

Cui, H. & Andersen, R. a. (2007). Posterior Parietal Cortex Encodes Autonomously Selected Motor Plans. *Neuron*, 56(3), 552–559.

Curto, C., Sakata, S., Marguet, S., Itskov, V., & Harris, K. D. (2009). A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(34), 10600–10612.

Davidson, T. J., Kloosterman, F., & Wilson, M. a. (2009). Hippocampal Replay of Extended Experience. *Neuron*, 63(4), 497–507.

Dombeck, D. a., Harvey, C. D., Tian, L., Looger, L. L., & Tank, D. W. (2010). Functional imaging of hippocampal place cells at cellular resolution during virtual navigation. *Nature Neuroscience*, 13(11), 1433–1440.

Domnisoru, C., Kinkhabwala, A. a., & Tank, D. W. (2013). Membrane potential dynamics of grid cells. *Nature*, 495(7440), 199–204.

Donahue, C. H. & Lee, D. (2015). Dynamic routing of task-relevant signals for decision making in dorsolateral prefrontal cortex. *Nature Neuroscience*, 18(2), 295–301.

Drugowitsch, J., Moreno-Bote, R., Churchland, A. K., Shadlen, M. N., & Pouget, A. (2012). The Cost of Accumulating Evidence in Perceptual Decision Making. *Journal of Neuroscience*, 32(11), 3612–3628.

Fitzgerald, J. K., Freedman, D. J., & Assad, J. a. (2011). Generalized associative

- representations in parietal cortex. *Nature Neuroscience*, 14(8), 1075–1079.
- Freedman, D. J. & Assad, J. a. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature*, 443(7107), 85–88.
- Freedman, D. J. & Assad, J. a. (2011). A proposed common neural mechanism for categorization and perceptual decisions. *Nature Neuroscience*, 14(2), 143–146.
- Frey, B. J. & Dueck, D. (2007). Clustering by Passing Messages Between Data Points.
- Fujisawa, S., Amarasingham, A., Harrison, M. T., & Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, 11(7), 823–833.
- Gold, J. I. & Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404(6776), 390–394.
- Gold, J. I. & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30(1), 535–574.
- Goldman, M. S. (2009). Memory without Feedback in a Neural Network. *Neuron*, 61(4), 621–634.
- Greenberg, D. S. & Kerr, J. N. D. (2009). Automated correction of fast motion artifacts for two-photon imaging of awake animals. *Journal of Neuroscience Methods*, 176(1), 1–15.
- Gütig, R. & Sompolinsky, H. (2006). The tempotron: a neuron that learns spike timing-based decisions. *Nature Neuroscience*, 9(3), 420–428.
- Hahnloser, R. H. R., Kozhevnikov, A. A., & Fee, M. S. (2002). An ultra-sparse code underlies the generation of neural sequences in a songbird. *Nature*, 419(6902), 65–70.
- Hanks, T. D., Brunton, B. W., Kopec, C. D., Duan, C. a., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, 520(7546), 220–223.

- Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, 9(5), 682–689.
- Harris, K. D. & Thiele, A. (2011). Cortical state and attention. *Nature Reviews Neuroscience*, 12(9), 509–523.
- Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, 484(7392), 62–68.
- Harvey, C. D., Collman, F., Dombeck, D. a., & Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, 461(7266), 941–946.
- Hoerzer, G. M., Legenstein, R., & Maass, W. (2014). Emergence of Complex Computational Structures From Chaotic Neural Networks Through Reward-Modulated Hebbian Learning. *Cerebral Cortex*, 24(3), 677–690.
- Horwitz, G. D. & Newsome, W. T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science*, 284(5417), 1158–1161.
- Huber, D., Gutnisky, D. a., Peron, S., O'Connor, D. H., Wiegert, J. S., Tian, L., Oertner, T. G., Looger, L. L., & Svoboda, K. (2012). Multiple dynamic representations in the motor cortex during sensorimotor learning. *Nature*, 484(7395), 473–478.
- Ipata, A. E., Gee, A. L., Gottlieb, J., Bisley, J. W., & Goldberg, M. E. (2006). LIP responses to a popout stimulus are reduced if it is overtly ignored. *Nature Neuroscience*, 9(8), 1071–1076.
- Jaeger, H. & Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science*, 304(5667), 78–80.
- Jeanne, J. M., Sharpee, T. O., & Gentner, T. Q. (2013). Associative Learning Enhances Population Coding by Inverting Interneuronal Correlation Patterns. *Neuron*, 78(2), 352–363.

- Jun, J. K., Miller, P., Hernandez, A., Zainos, A., Lemus, L., Brody, C. D., & Romo, R. (2010). Heterogenous Population Coding of a Short-Term Memory and Decision Task. *Journal of Neuroscience*, 30(3), 916–929.
- Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, 455(7210), 227–231.
- Kira, S., Yang, T., & Shadlen, M. N. (2015). A neural implementation of Wald’s sequential probability ratio test. *Neuron*, 85(4), 861–873.
- Klampfl, S., David, S. V., Yin, P., Shamma, S. A., & Maass, W. (2012). A quantitative analysis of information about past and present stimuli encoded by spikes of A1 neurons. *Journal of Neurophysiology*, 108(5), 1366–1380.
- Klampfl, S. & Maass, W. (2013). Emergence of Dynamic Memory Traces in Cortical Microcircuit Models through STDP. *Journal of Neuroscience*, 33(28), 11515–11529.
- Kusunoki, M., Gottlieb, J., & Goldberg, M. E. (2000). The lateral intraparietal area as a salience map: the representation of abrupt onset, stimulus motion, and task relevance. *Vision Research*, 40(10-12), 1459–1468.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Legenstein, R. & Maass, W. (2014). Ensembles of Spiking Neurons with Noise Support Optimal Probabilistic Inference in a Dynamically Changing Environment. *PLoS Computational Biology*, 10(10), e1003859.
- Legenstein, R., Pecevski, D., & Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Computational Biology*, 4(10), e1000180.
- Li, N., Chen, T.-W., Guo, Z. V., Gerfen, C. R., & Svoboda, K. (2015). A motor cortex circuit for motor planning and movement. *Nature*, 519(7541), 51–56.

- Looger, L. L. & Griesbeck, O. (2012). Genetically encoded neural activity indicators. *Current Opinion in Neurobiology*, 22(1), 18–23.
- Maass, W., Natschläger, T., & Markram, H. (2002). Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation*, 14(11), 2531–2560.
- Machens, C. K. (2005). Flexible Control of Mutual Inhibition: A Neural Model of Two-Interval Discrimination. *Science*, 307(5712), 1121–1124.
- Maimon, G. & Assad, J. a. (2009). Beyond Poisson: increased spike-time regularity across primate parietal cortex. *Neuron*, 62(3), 426–440.
- Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474), 78–84.
- Marcos, E., Pani, P., Brunamonti, E., Deco, G., Ferraina, S., & Verschure, P. (2013). Neural Variability in Premotor Cortex Is Modulated by Trial History and Predicts Behavioral Performance. *78(2)*, 249–255.
- Mazor, O. & Laurent, G. (2005). Transient Dynamics versus Fixed Points in Odor Representations by Locust Antennal Lobe Projection Neurons. *Neuron*, 48(4), 661–673.
- Mazurek, M. E. (2003). A Role for Neural Integrators in Perceptual Decision Making. *Cerebral cortex (New York, N.Y. : 1991)*, 13(11), 1257–1269.
- Meister, M. L. R., Hennig, J. A., & Huk, A. C. (2013). Signal Multiplexing and Single-Neuron Computations in Lateral Intraparietal Area During Decision-Making. *Journal of Neuroscience*, 33(6), 2254–2267.
- Miller, E. K. & Wilson, M. a. (2008). All My Circuits: Using Multiple Electrodes to Understand Functioning Neural Networks. *Neuron*, 60(3), 483–488.
- Murakami, M., Vicente, M. I., Costa, G. M., & Mainen, Z. F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. *Nature Neuroscience*, 17(11), 1574–

1582.

- Murphy, K. P. (2012). Machine learning: a probabilistic perspective.
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, 17(12), 1661–1663.
- Natschläger, T. & Maass, W. (2005). Dynamics of information and emergent computation in generic neural microcircuit models. *Neural Networks*, 18(10), 1301–1308.
- Nikolić, D., Häusler, S., Singer, W., & Maass, W. (2009). Distributed Fading Memory for Stimulus Properties in the Primary Visual Cortex. *PLoS Biology*, 7(12), e1000260.
- Niwa, M. & Ditterich, J. (2008). Perceptual decisions between multiple directions of visual motion. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 28(17), 4435–4445.
- Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience*, 17(10), 1395–1403.
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science*, 321(5894), 1322–1327.
- Petersen, R. S. & Diamond, M. E. (2000). Spatial-temporal distribution of whisker-evoked activity in rat somatosensory cortex and the coding of stimulus location. *Journal of Neuroscience*, 20(16), 6135–6143.
- Pnevmatikakis, E. A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., Ahrens, M., Bruno, R., Jessell, T. M., Peterka, D. S., Yuste, R., & Paninski, L. (2016). Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data. *Neuron*, 89(2), 285–299.
- Pologruto, T. a., Sabatini, B. L., & Svoboda, K. (2003). ScanImage: flexible software for

- operating laser scanning microscopes. *Biomedical engineering online*, 2(1), 13.
- Quian Quiroga, R. (2006). Movement Intention Is Better Predicted than Attention in the Posterior Parietal Cortex. *Journal of Neuroscience*, 26(13), 3615–3620.
- Rajan, K., Harvey, C. D., & Tank, D. W. (2016). Recurrent Network Models of Sequence Generation and Memory. *Neuron*.
- Raposo, D., Kaufman, M. T., & Churchland, A. K. (2014). A category-free neural population supports evolving demands during decision-making. *Nature Neuroscience*, 17(12), 1784–1792.
- Raposo, D., Sheppard, J. P., Schrater, P. R., & Churchland, A. K. (2012). Multisensory Decision-Making in Rats and Humans. *Journal of Neuroscience*, 32(11), 3726–3735.
- Ratcliff, R. & Smith, P. L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review*, 111(2), 333–367.
- Renart, A. & Machens, C. K. (2014). Variability in neural activity and behavior. 25, 211–220.
- Rigotti, M., Barak, O., Warden, M. R., Wang, X.-J., Daw, N. D., Miller, E. K., & Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. *Nature*, 497(7451), 585–590.
- Rinberg, D., Koulakov, A., & Gelperin, A. (2006). Speed-Accuracy Tradeoff in Olfaction. *Neuron*, 51(3), 351–358.
- Roitman, J. D. & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 22(21), 9475–9489.
- Romo, R., Hernández, A., Zainos, A., & Salinas, E. (2003). Correlated Neuronal Discharges that Increase Coding Efficiency during Perceptual Discrimination. *Neuron*, 38(4), 649–657.

- Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386–408.
- Safaai, H., Neves, R., Eschenko, O., Logothetis, N. K., & Panzeri, S. (2015). Modeling the effect of locus coeruleus firing on cortical state dynamics and single-trial sensory processing. *Proceedings of the National Academy of Sciences*, 112(41), 12834–12839.
- Safaai, H., von Heimendahl, M., Sorando, J. M., Diamond, M. E., & Maravall, M. (2013). Coordinated Population Activity Underlying Texture Discrimination in Rat Barrel Cortex. *Journal of Neuroscience*, 33(13), 5843–5855.
- Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W., & Brody, C. D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife*, 4, 494.
- Seo, H., Barraclough, D. J., & Lee, D. (2007). Dynamic signals related to choices and outcomes in the dorsolateral prefrontal cortex. *Cerebral Cortex*, 17 Suppl 1(suppl 1), i110–7.
- Seo, H., Barraclough, D. J., & Lee, D. (2009). Lateral intraparietal cortex and reinforcement learning during a mixed-strategy game. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 29(22), 7278–7289.
- Seo, H. & Lee, D. (2007). Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 27(31), 8366–8377.
- Shadlen, M. N. & Kiani, R. (2013). Decision Making as a Window on Cognition. *Neuron*, 80(3), 791–806.
- Shadlen, M. N. & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences*, 93(2), 628–633.
- Shadlen, M. N. & Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18(10), 3870–3896.

- Shi, J. & Malik, J. (2000). Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8), 888–905.
- Smith, P. L. (2000). Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer. *Journal of Mathematical Psychology*, 44(3), 408–463.
- Smola, A. & Vapnik, V. (1997). Support vector regression machines. *Advances in neural information processing*
- Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science*, 304(5678), 1782–1787.
- Sussillo, D. (2014). Neural circuits as computational dynamical systems. *Current Opinion in Neurobiology*, 25, 156–163.
- Sussillo, D. & Abbott, L. F. (2009). Generating Coherent Patterns of Activity from Chaotic Neural Networks. *Neuron*, 63(4), 544–557.
- Sutton, R. S. & Barto, A. G. (1998). *Reinforcement learning: An introduction*.
- Theis, L., Berens, P., Froudarakis, E., Reimer, J., Roman-Roson, M., Baden, T., Euler, T., Tolias, A. S., & Bethge, M. (2014). *Supervised learning sets benchmark for robust spike rate inference from calcium imaging signals*. Technical report.
- Tian, L., Hires, S. A., Mao, T., Huber, D., Chiappe, M. E., Chalasani, S. H., Petreanu, L., Akerboom, J., Mckinney, S. A., Schreiter, E. R., Bargmann, C. I., Jayaraman, V., Svoboda, K., & Looger, L. L. (2009). Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nature Methods*, 6(12), 875–881.
- Uchida, N. & Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nature Neuroscience*, 6(11), 1224–1229.
- Verstraeten, D., Schrauwen, B., D’Haene, M., & Stroobandt, D. (2007). An experimental unification of reservoir computing methods. *Neural Networks*, 20(3), 391–403.

- Vogelstein, J. T., Packer, A. M., Machado, T. A., Sippy, T., Babadi, B., Yuste, R., & Paninski, L. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *Journal of Neurophysiology*, 104(6), 3691–3704.
- Wang, X.-J. (2002). Probabilistic Decision Making by Slow Reverberation in Cortical Circuits. 36(5), 955–968.
- Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. *Neuron*, 60(2), 215–234.
- Wang, X.-J. (2012). Neural dynamics and circuit mechanisms of decision-making. *Current Opinion in Neurobiology*, 22(6), 1039–1046.
- Wickelgren, W. A. (1977). Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica*, 41(1), 67–85.
- Wong, K.-F. & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 26(4), 1314–1328.
- Yang, T. & Shadlen, M. N. (2007). Probabilistic reasoning by neurons. *Nature*, 447(7148), 1075–1080.
- Yang, Y. & Zador, A. M. (2012). Differences in sensitivity to neural timing among cortical areas. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 32(43), 15142–15147.
- Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370(6485), 140–143.