

Neuron

Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales

Highlights

- Voluntary restraint allows two-photon Ca^{2+} imaging during decision-making in rats
- Neurons with diverse dynamics collectively represent accumulated sensory evidence
- Cortical dynamics support the existence of multiple weakly coupled accumulators
- Neurons across fronto-parietal cortex encode the memory of past behavioral choices

Authors

Benjamin B. Scott,
Christine M. Constantinople,
Athena Akrami, Timothy D. Hanks,
Carlos D. Brody, David W. Tank

Correspondence

brody@princeton.edu (C.D.B.),
dwtank@princeton.edu (D.W.T.)

In Brief

Leading models of decision-making postulate that individual fronto-parietal neurons encode accumulated sensory evidence with stable changes in firing rate. Using cellular resolution calcium imaging during a pulse-based accumulation task, Scott et al. reveal that stable representations of accumulated evidence in rat fronto-parietal cortex instead arise from neuronal populations with temporally diverse responses.



Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales

Benjamin B. Scott,^{1,2,7} Christine M. Constantinople,^{1,2,7} Athena Akrami,^{1,2,3} Timothy D. Hanks,^{5,6} Carlos D. Brody,^{1,2,3,*} and David W. Tank^{1,2,4,8,*}

¹Princeton Neuroscience Institute

²Department of Molecular Biology

³Howard Hughes Medical Institute

⁴Bezos Center for Neural Circuit Dynamics

Princeton University, Princeton, NJ 08544, USA

⁵Center for Neuroscience, University of California, Davis, Davis, CA 95618, USA

⁶Department of Neurology, University of California, Davis, Sacramento, CA 95817, USA

⁷These authors contributed equally

⁸Lead Contact

*Correspondence: brody@princeton.edu (C.D.B.), dwtank@princeton.edu (D.W.T.)

<http://dx.doi.org/10.1016/j.neuron.2017.06.013>

SUMMARY

Decision-making in dynamic environments often involves accumulation of evidence, in which new information is used to update beliefs and select future actions. Using *in vivo* cellular resolution imaging in voluntarily head-restrained rats, we examined the responses of neurons in frontal and parietal cortices during a pulse-based accumulation of evidence task. Neurons exhibited activity that predicted the animal's upcoming choice, previous choice, and graded responses that reflected the strength of the accumulated evidence. The pulsatile nature of the stimuli enabled characterization of the responses of neurons to a single quantum (pulse) of evidence. Across the population, individual neurons displayed extensive heterogeneity in the dynamics of responses to pulses. The diversity of responses was sufficiently rich to form a temporal basis for accumulated evidence estimated from a latent variable model. These results suggest that heterogeneous, often transient sensory responses distributed across the fronto-parietal cortex may support working memory on behavioral timescales.

INTRODUCTION

Perceptual decision-making often requires integration of sensory information over time. Behaviorally, this is frequently modeled as an accumulation process, in which a subject's perceptual judgment is described as a latent variable evolving under the influence of sensory events, comparable to the movement of particles in drift diffusion processes (Forstmann et al., 2016). Over the last several decades, this accumulation process

has been investigated using electrophysiological recordings from single neurons in primate and rodent cortex. Continuous presentation of noisy sensory evidence produces ramping activity in neurons in posterior parietal cortex (PPC; Roitman and Shadlen, 2002) and frontal eye field (FEF; Kim and Shadlen, 1999), while brief pulses of sensory evidence produce sustained firing rate increases in PPC (Hanks et al., 2015; Huk and Shadlen, 2005), leading to the hypothesis that the brain represents the accumulated sensory evidence in forebrain neuronal activity.

Recently, the frontal orienting field (FOF) and PPC have been implicated in accumulation of evidence and memory-guided decision-making in rats. FOF and rat PPC share similar thalamocortical and corticothalamic projections with primate FEF and PPC, respectively (Erlich et al., 2011; Whitlock et al., 2012), and inactivation of FOF causes disruptions in accumulation of evidence and memory-guided orienting tasks, but not in cued sensory-motor responses (Erlich et al., 2011, 2015; Hanks et al., 2015; Kopec et al., 2015). Tetrode recordings in these regions revealed neurons tuned to the strength of evidence in a pulse-based auditory accumulation task (Hanks et al., 2015).

Most previous electrophysiological studies of fronto-parietal cortex during accumulation of evidence have focused on how firing rates of individual neurons correlate with the magnitude of sensory evidence in favor of a particular choice as predicted by an accumulator model. Typically, the responses of different neurons are analyzed as amplitude-scaled versions of a prototype temporal waveform representing the latent variable of the accumulator, implying relatively homogeneous dynamics across the neural population (Figures 1A and 1B). However, theoretical work has suggested that neural networks could also encode the memory of sensory events using a population code with heterogeneous dynamics (Figure 1B). Whether these responses could also underlie an accumulation process has not been explored. For example, Goldman (2009) demonstrated that networks of simulated neurons with complex, temporally heterogeneous responses of increasing response width can encode memories

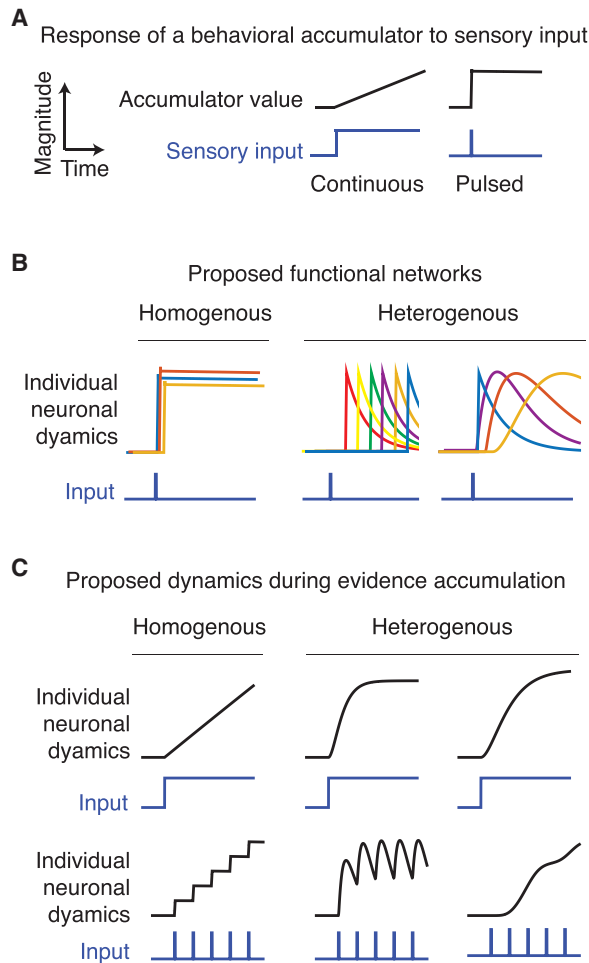


Figure 1. Predicted Neural Dynamics for Accumulation of Sensory Evidence during Perceptual Decision-Making

(A) Change in magnitude of an idealized accumulator (black line) in response to sensory input (blue line). When input is presented continuously (left), an accumulator responds with a linear ramp. When the input is pulsed (right), an accumulator responds with a step.

(B) Predicted neural dynamics from two different types of networks that have been proposed to encode the value of the accumulator. The upper row illustrates the predicted dynamics of individual neurons (colored lines) from each network in response to a single pulse of sensory input. In the homogeneous network (left column), each neuron responds with a similar step-like increase in activity. In the heterogeneous network (right column), neurons respond to input with different delays and different time courses. Two examples of heterogeneous dynamics are shown, sequential dynamics (left) and complex time varying (right). In the homogeneous network, the activity of a single neuron is sufficient to encode the memory of the input, while in the sequential and heterogeneous networks, the memory of the input is encoded in the population ensemble.

(C) Predicted responses of a single neuron from each of the two proposed networks during gradual accumulation of evidence with either continuous (upper row) or pulsed (lower row) sensory input. While continuous input produces a similar ramping response in neurons from each network, the pulsed input produces qualitatively different responses, a staircase (homogeneous, left), an oscillation, or a ramp (heterogeneous, right).

of transient events in their combined activity. Based on two-photon imaging of calcium dynamics in mouse PPC during a memory-guided virtual navigation task, Harvey et al. (2012) suggested that sequential activation of neurons following transient sensory events could encode the memory of reward location. More recently, Morcos and Harvey (2016) have shown that more complex state transitions in PPC population activity can encode both the upcoming and past choices of the animal. However, the specific nature of the encoding in each state and the responses to individual pulses of evidence in the accumulation process were not evaluated. Each of the above models makes a distinct prediction about individual neural response functions following sensory events during accumulation (Figures 1C). Motivated by these conceptual ideas, we examined neural responses to single pulses of evidence.

We used cellular resolution imaging in behaving rats to record population dynamics in FOF and PPC during a visual pulse-based accumulation of evidence task (Scott et al., 2015). In this task, rats viewed two streams of light pulses, presented from a left and a right LED, and were rewarded for orienting to the side with more pulses. Voluntary head restraint (Scott et al., 2013) provided the stability necessary for two-photon imaging of GCaMP6f-labeled neurons in PPC and FOF. Analysis of calcium-dependent fluorescence transients in layer 2/3 neurons revealed activity that predicted the animal's subsequent choice, the animal's choice on the previous trial, and graded responses reflecting the strength of accumulated evidence. Neurons responded, on average, to each evidence pulse with a sustained increase in activity. This increase mirrored the timescale of behavioral memory of each pulse as predicted by drift diffusion behavioral models (Figure S1; Brunton et al., 2013; Scott et al., 2015). However, analysis of the responses of individual neurons to pulses of evidence revealed extensive temporal heterogeneity across the population. Re-analysis of electrophysiological recordings from an auditory version of the task (Hanks et al., 2015) also revealed heterogeneous, often transient spiking responses of neurons to individual pulses of evidence, confirming the generality of this finding. These results are consistent with a model of evidence accumulation in which the combined activity of neurons with diverse time-varying responses represents the memory of sensory events.

RESULTS

Cellular Resolution Imaging during Accumulation of Evidence in Voluntarily Head-Restrained Rats

We trained rats to perform an accumulation of evidence task during voluntary head restraint. Seven adult male rats were surgically implanted with custom titanium headplates. After recovery, rats were trained to perform a two-alternative forced choice task for water reward. The details of this task have been previously described (Scott et al., 2015). Rats initiated a behavioral trial by sliding their headplate into a custom headport at one end of a behavioral chamber (Scott et al., 2013) (Figure 2A). Once inserted, the headplate activated a kinematic clamp that registered and stabilized the position of the plate and head. Rats were required to maintain head fixation for 2–3 s, until an auditory cue instructed them to remove their head from the headport and orient to a side reward port. During the

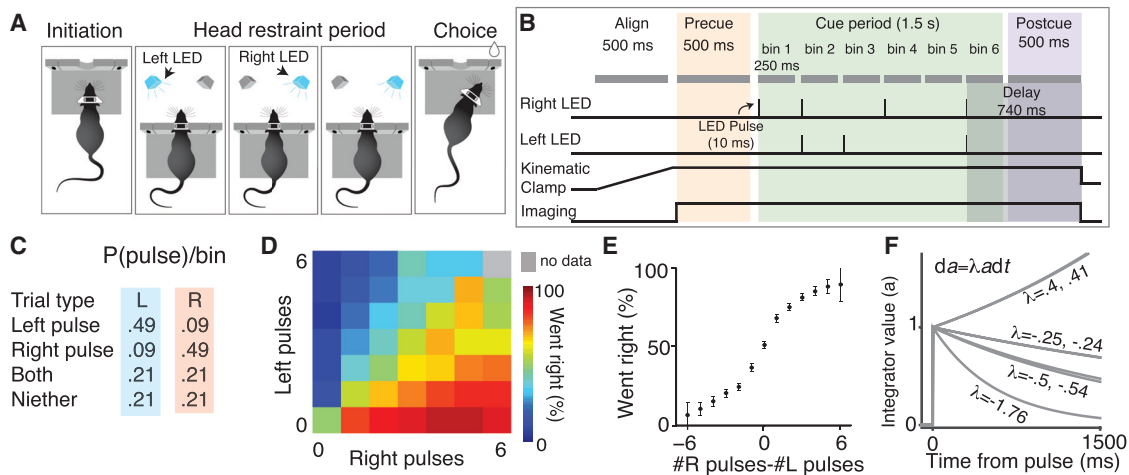


Figure 2. Behavioral Performance of Voluntary Head-Restrained Rats on a Pulse-Based Accumulation of Evidence Task

(A) Schematic of the behavioral task during an example trial (see Scott et al., 2015). Initiation: a rat initiates a behavioral trial by inserting its head and headplate into a custom headport along one wall of an operant training box. Head-restraint period: the rat is presented with a series of pseudo-randomly timed flashes from two LEDs positioned within the animal's left and right visual hemifields. Choice: after presentation of the cues, the subject is released from restraint and is free to make a behavioral report by inserting its nose to one of the side ports. Water reward (25 μ L) was baited on the side that had the greater number of flashes.

(B) The head-restraint period is divided into a series of segments. Align: a 500 ms alignment period in which the kinematic clamp registers the position of the head and headplate. Precue: a 500 ms imaging period in which the baseline fluorescence of the cells in the field of view is recorded. Cue: the cue presentation period consists of 2–6 bins, each starting with a 10 ms pulse period, followed by a 240 ms refractory period. Postcue: a 500 ms before the subject is released from head restraint. Delay between the final LED pulse and release can last 740–1,990 ms.

(C) Probability of left and right pulses in each cue bin. Probability of each event (left pulse only, right pulse only, both left and right pulses, or no pulses) depends on whether the trial was a left generative trial (blue) or right generative trial (pink).

(D) Behavioral performance on all stimulus types (no. right and no. left) pooled across seven subjects. Color of each small square indicates the percentage of trials in which the animals went right for that combination of right and left flashes.

(E) Mean psychophysical performance across all subjects for trials with different evidence strengths (no. right pulses – no. left pulses).

(F) Predicted response of the accumulator to a pulse of evidence, based on these λ values. A previously described behavioral model (Brunton et al., 2013) was used to find the time constant of the accumulator ($1/\lambda$, in units of s⁻¹) for each rat (see STAR Methods).

head-restraint period, rats were presented with a series of pseudo-randomly timed light pulses from a left and right LED (Figures 2B and 2C). Pulses on the same side were separated by a minimum of 240 ms to prevent the perceptual fusion of consecutive flashes, leading to a maximum of six flashes on each side. Pulses on opposite sides could occur simultaneously, leading to a possible maximum of 12 total flashes. Following release from restraint, rats received a water reward (25–75 μ L) for orienting to the side with more flashes.

Semi-automated training facilitated the generation of a behavioral dataset comprising 13,438 trials from 7 animals, which in turn enabled statistical characterization of the decision-making process (Figures 2D–2F and S1). We previously determined that subjects were sensitive to a difference of a single pulse and accumulated all pulses (Scott et al., 2015). Behavioral modeling suggested that the memory of each flash exhibited long time constants, on the order of hundreds of milliseconds to seconds (Figures 2F and S1; Scott et al., 2015).

Once trained to perform accumulation during head restraint (see Scott et al., 2015 for shaping procedures), rats underwent a second surgery in which an optical window was implanted over either FOF (+2 AP [anteroposterior], 1.5 ML [mediolateral]) or PPC (–3.8 AP, 2.5 ML) (Figure 3A). We also performed imaging in medial V2 (mV2; –6.0 AP, 2.0 ML) in two rats. mV2 is known to exhibit transient elevated firing rates in response to

visual stimuli (Espinoza and Thomas, 1983) and is strongly inter-connected with PPC (Kolb and Walkey, 1987; Reep et al., 1994; Wilber et al., 2015). Neuronal expression of the genetically encoded calcium sensor GCaMP6f was accomplished using an adeno-associated viral vector (see STAR Methods).

Following recovery from the optical window surgery and expression of the GCaMP6f indicator, we performed cellular resolution imaging of calcium concentration dynamics via two-photon laser-scanning microscopy (Figures 3B and 3C). The laser shutter was synchronized to the animal's behavior, opening at trial initiation and closing during the inter-trial interval, limiting laser illumination to the brain region under investigation. Activation of the kinematic clamp, synchronized with the head-fixation period, provided the stability necessary for cellular resolution imaging and registered the position of the headplate and head to within a few microns. This registration enabled imaging of the same field of view on subsequent head insertions, so that the same neurons could be imaged across trials in an imaging session (typically 100–200 trials).

Using this approach, we recorded from 500 (81 in mV2, 151 in PPC, and 268 in FOF) morphologically identified neurons 150–200 μ m below the cortical surface while rats ($n = 7$) performed the visual accumulation of evidence task (Figures 3D–3F). On average, 10.5 (range = 2–28) GCaMP6f-positive neurons could be visually detected per field of view. In total, 21% of

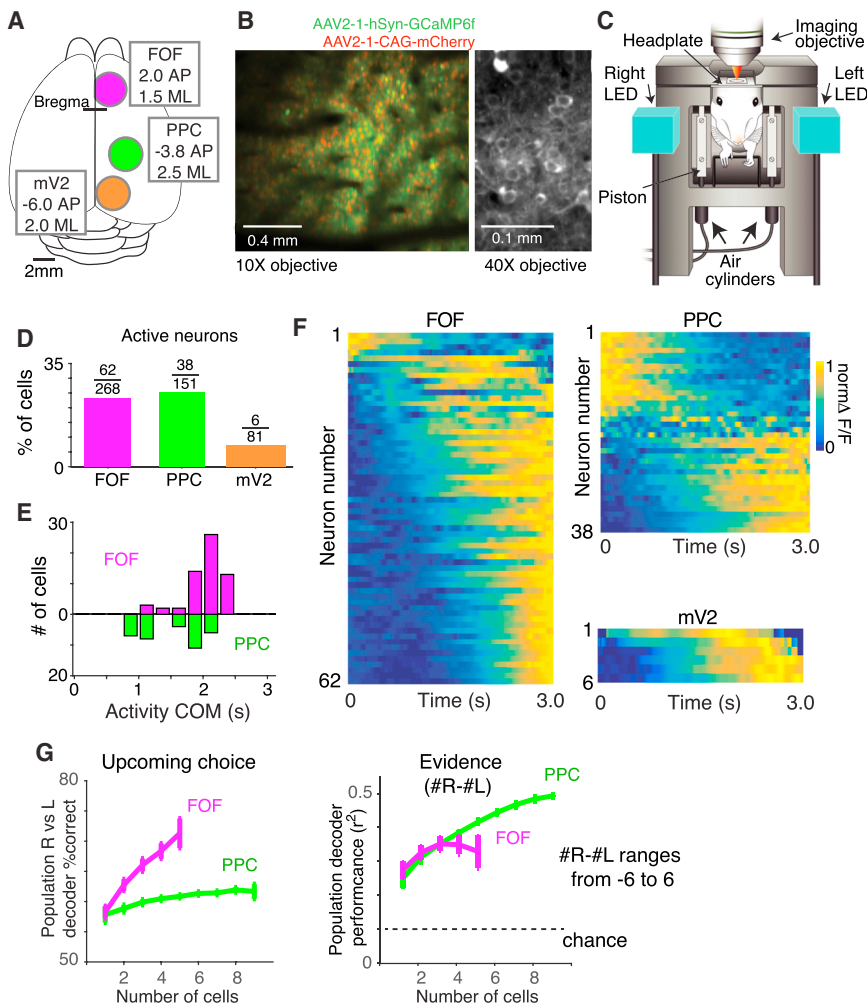


Figure 3. Neuronal Calcium Dynamics in Rat mV2, PPC, and FOF during Accumulation of Evidence

(A) Positions of the optical windows for mV2 (orange circle), PPC (green circle), and FOF (magenta circle) relative to bregma.

(B) Images obtained by two-photon microscopy of GCaMP6f and mCherry expression in rat cortical layer 2/3 neurons. Left panel shows the labeling of GCaMP6f (green) and mCherry (red) in an example field of view recorded during voluntary head restraint using a 10× air objective. Right panel shows the labeling of GCaMP6f (green) recorded in an anesthetized rat using a 40× water immersion objective.

(C) Illustration indicating the positions of the animal's head, headplate, right and left LEDs, imaging objective, and kinematic clamp (pistons and air cylinders) during imaging.

(D) Percentage of active cells recorded in each area: mV2 (orange), PPC (green), and FOF (magenta).

(E) Histogram of the response times across all active neurons in PPC (green) and FOF (magenta) evaluated by computing the center of mass (COM) of the average calcium response within the trial period. Activity of FOF neurons is centered near the end of the trial. Neurons in PPC tend to be active early in the trial or late in the trial.

(F) Calcium signals from all active neurons in mV2, PPC, and FOF averaged across all completed trials. Each row represents a neuron and the color represents the normalized $\Delta F/F$ at each point in the trial.

(G) Performance of a linear population classifier in decoding task events based on the activity of simultaneously recorded neurons in FOF and PPC. The upcoming choice of the animal and the accumulated evidence ($\#R-\#L$; ranging from -6 to 6) on each trial were decoded based on the average activity during the cue and delay period. Left panel: a linear population classifier, trained

using SVM, is significantly better at decoding the upcoming choice of the animal based on the activity of neurons in FOF than in PPC. Right panel: performance of a linear classifier, trained using ridge regression, to decode accumulated evidence ($\#R-\#L$; ranging from -6 to 6) based on the activity of neurons in FOF and PPC.

morphologically identified cells exhibited significantly modulated activity during the task (see STAR Methods). The percentage of active cells was not significantly different between PPC (25%, 38/151) and FOF (23%, 62/268), but both were significantly greater than mV2 (7%, 6/81, binomial test, $p < 0.05$; Figure 3D). Across the population, neurons exhibited complex, time-varying responses that peaked at different times within the imaging period (Figures 3E and 3F). The number of active cells we observed was low compared with previous experiments in which neurons in PPC were imaged in head-fixed mice (for example, 47% in Harvey et al., 2012). This difference may reflect the fact that imaging was restricted to behavioral trials (during voluntary head restraint). Neurons were not recorded during inter-trial intervals, limiting the window in which activity could be observed.

We evaluated how neural populations in FOF and PPC encode choice and other task-relevant parameters. We used a support vector machine (SVM)-based decoding approach to train a clas-

sifier to decode the upcoming choice of the animal based on simultaneously recorded cells in PPC and FOF (Figure 3G). Decoders based on FOF neurons strongly predicted the upcoming choice of the animal, while population decoders based on PPC neurons did not perform well. We also examined how well each area encoded accumulated evidence (no. right flashes – no. left flashes) using a linear regression-based approach (see STAR Methods). Decoders for accumulated evidence based on the activity of FOF and PPC performed significantly better than chance and improved with increasing numbers of neurons.

Population Dynamics in FOF Predict Upcoming Choice

Previous tetrode recordings identified neurons in PPC and FOF, termed “side-selective neurons,” whose firing rates were significantly different for left versus right choice trials (Figures 4A and 4B; Erlich et al., 2011; Hanks et al., 2015). Imaging from the superficial layers of FOF and PPC also revealed side selectivity (Figures 4 and S2). Consistent with previous reports (Erlich

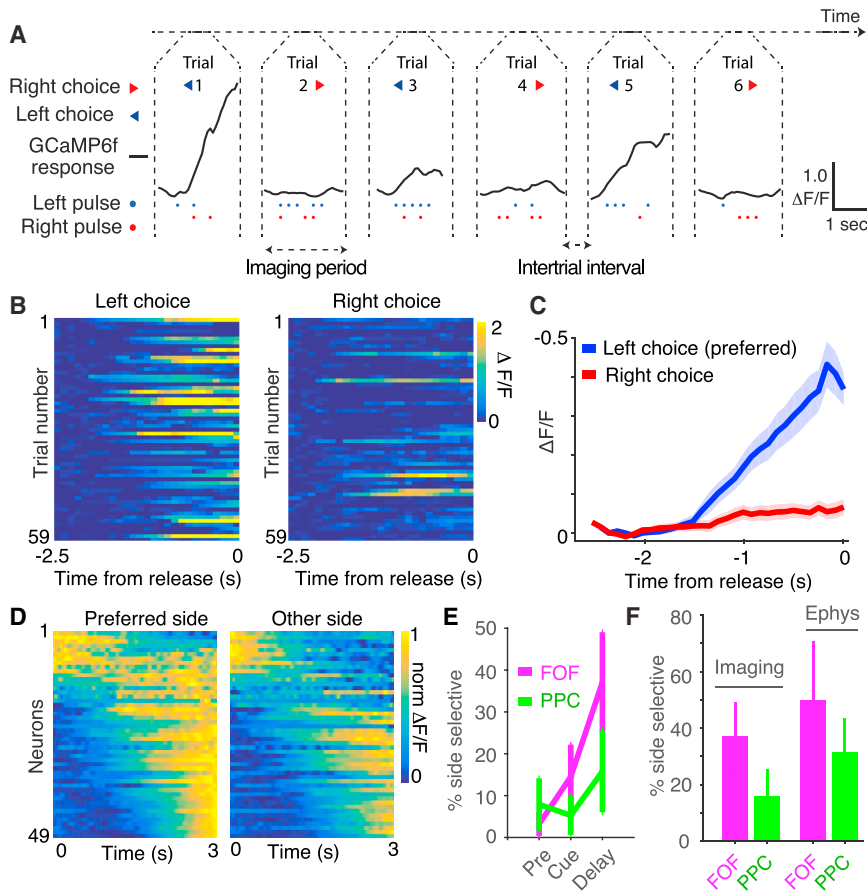


Figure 4. Side-Selective Neurons in FOF and PPC

(A) Representative calcium signals ($\Delta F/F$, black line) from a single neuron in FOF on six consecutive completed trials. The animal's choice (arrow heads, blue for left choice and red for right choice) and the timing of the left pulses (blue dots) and right pulses (red dots) are also indicated. This neuron exhibited transients on trials in which the animal subsequently chose the left reward port.

(B) Responses of the neuron shown in (A) on 118 correct trials from a single imaging session. Color indicates the strength of the response ($\Delta F/F$). Trials are separated based on whether the animal chose left (left panel) or right (right panel).

(C) Average activity of the neuron shown in (A) and (B) across all correct trials. Blue line indicates the average response ($\Delta F/F$) on left choice trials; red line indicates response on right choice trials. Shaded area indicates SEM.

(D) Dynamics of all side-selective neurons separated by the choice of the rat and sorted by the time of the center of mass of activity on the preferred side.

(E) Comparison of the percentage of side-selective neurons detected using the activity of different trial epochs. Magenta line indicates the percentage of side-selective cells detected in FOF during the pre-cue, cue, and delay periods of the task. Green line indicates the percentage of side-selective cells detected in PPC during the same periods. Error bars indicate the 95% binomial confidence interval.

(F) Comparison of the percentage of side-selective neurons in FOF and PPC recorded using imaging of calcium dynamics or using tetrode recording of action potentials. Error bars indicate the 95% binomial confidence interval.

et al., 2011; Hanks et al., 2015), neurons with preferences for ipsilateral and contralateral movements were intermixed in a single hemisphere. Most side-selective neurons exhibited ramp-like increases in calcium responses throughout the cue and delay periods (Figures 4C and 4D). In these neurons, the slope of the ramp was different on left and right choice trials and the ramps began to separate during the cue period and continued to diverge throughout the trial until release from voluntary restraint. The percentage of side-selective neurons detected during the delay period was 16% (6/38) in PPC and 37% (23/62) in FOF, which is similar to those observed in superficial recordings with tetrodes during an auditory accumulation of evidence task (Figure 4F). A smaller population of side-selective neurons that ramped down during the cue period was also observed (Figure 4D) in the calcium imaging experiments.

To determine whether neurons that were not identified as side-selective collectively contain information about the upcoming choice, we trained a linear decoder to predict choice based on the activity of groups of neurons that were not side selective (Figures S2E and S2F). Decoders based on individual non-side-selective neurons in FOF and PPC failed to predict the animal's choice above chance ($p > 0.01$). However, larger groups of simultaneously recorded neurons in both PPC and FOF that were not side selective could be used to decode the upcoming choice of the animal ($p < 0.01$).

Sensory Evidence Modulates Activity in PPC and FOF

We next asked how the dynamics of individual PPC and FOF neurons were modulated by each pulse of sensory evidence (Figure 5). By averaging the activity of neurons aligned to the timing of the flashes (pulse-triggered average), we found that neurons exhibited transient responses to individual pulses of evidence, with varied dynamics (Figures 5A–5C). Consistent with exhibiting sensory responses, individual neurons responded earlier on trials with earlier pulses and later on trials when pulses occurred later (Figures 5D and 5E).

Next we examined the effect of sensory evidence on the magnitude of neural responses in PPC and FOF. We performed linear regression between evidence at the end of the trial and the average response of the neuron during the delay period, which was at least 740 ms but often longer depending on when the last flash occurred (Figure 5F). Evidence was defined as either the difference between the number of right and left flashes ($\#R - \#L$), the total number of right flashes ($\#R$), or the total number of left flashes ($\#L$). A cell was defined as “evidence tuned” if the slope of the line that represented the best fit linear relationship between evidence and delay period activity was significantly different from zero and the correlation coefficient (r^2) was greater than 0.5 (Figures 5F and S3). In total, 39% (15/38) of active PPC neurons and 24% (15/62) of active FOF neurons had significant evidence

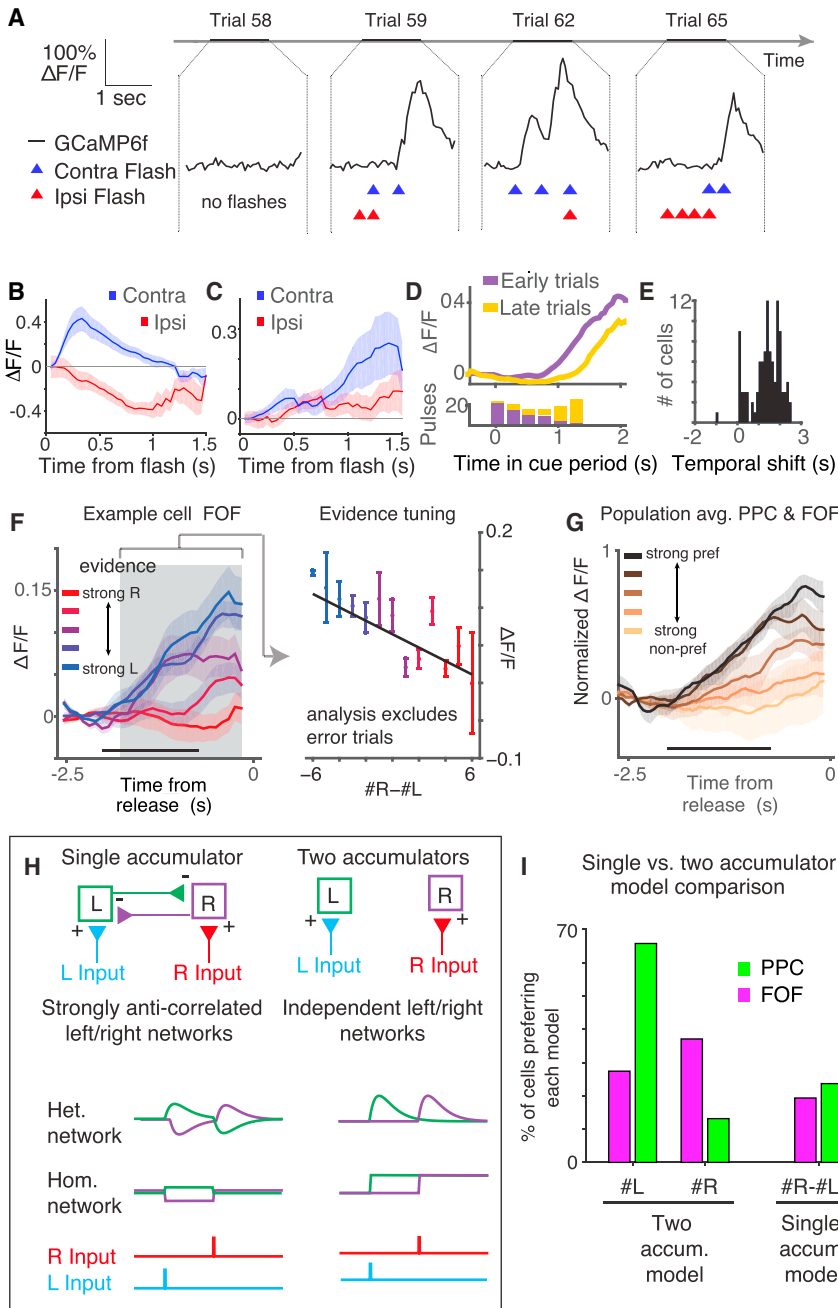


Figure 5. Sensory Evidence Modulates Activity in mV2, PPC, and FOF

(A) Calcium dynamics recorded on four example trials of an mV2 cell that exhibited robust responses to pulses on the contralateral (left) side. Black line indicates the response of the neuron ($\Delta F/F$); blue and red triangles represent the timing of contra and ipsilateral pulses, respectively.

(B) Average response of the neuron shown in (A) aligned to the time of each pulse and averaged across all pulses on the contralateral side (blue line) and ipsilateral side (red line). Shaded areas indicate SEM.

(C) Average responses to pulses on the contralateral (blue) and ipsilateral (red) side for an example neuron in FOF. This neuron exhibits a peak in the pulse-triggered average on the preferred side (contra) over 1 s after pulse presentation.

(D) Upper panel: response of an example FOF neuron averaged across trials in which the left flashes were presented early in the trial (purple line) and trials in which the pulses were presented late in the trial (yellow line). Lower panel: histogram of the timing of the flashes that were presented on early (purple bars) and late (yellow bars) trials.

(E) Temporal shift between early evidence trials and late evidence trials across all neurons. For each neuron, temporal shift was defined as the difference between the mean of the times to 0.5 of max and 0.75 of max on trials where the preferred pulses came late and those trials when the preferred pulses came early.

(F) Averaged response for a side-selective neuron on correct trials with different levels of evidence. Color of the line indicates the strength of the evidence, ranging from strong rightward (red, #R-#L > 4) to strong leftward (blue, #R-#L < -4). Shaded areas indicate SEM. Horizontal black bar indicates the maximal duration of the cue period. Gray shaded area indicates the maximal extent of the delay period. Activity during the delay period was used to calculate evidence tuning index as shown in the right panel. Right panel: response during the delay period averaged across trials with different amounts of evidence (#R-#L). Vertical bars indicate SEM. Black line indicates the best fit using least-squares regression.

(G) Population average of evidence-tuned cells ($n = 15$) separated across trials with different levels of evidence from strong evidence on the preferred side (black line, #pref side-#non-pref side > 4) to strong evidence on the non-preferred side (copper line, #pref side-#non-pref side < -4).

Black horizontal bar indicates the maximum duration of the cue period. Responses are aligned to the end of the trial.

(H) Schematic of two models that describe how multiple streams of sensory input could be accumulated during perceptual decision-making. In the single accumulator model (left), networks of neurons that process left and right sensory input (blue and red lines) strongly inhibit each other. This model predicts responses that encode the difference in sensory input on the left and right side. In the two accumulator model (right), the networks that encode left and right input are independent. The exact dynamics of the neurons depend the type of accumulator network (i.e., homogeneous or heterogeneous).

(I) Comparison of the single accumulator (#R-#L) and two accumulator (#R or #L) models schematized in (H). Bar height indicates the percentage of cells that were best fit by each model. In both PPC and FOF, the percentage of cells preferring either #R or #L was greater than the percentage preferring the difference (#R-#L).

tuning. These results demonstrate that layer 2/3 neurons in rat fronto-parietal cortex reflect the upcoming choice of the animal with dynamics that are modulated by the magnitude and timing of sensory evidence.

Two types of networks, single and two accumulator networks, have been proposed to implement neural accumulation in decision-making tasks (Figure 5H). In single accumulator-based models, subjects would maintain a memory of the difference in

pulses, while in two accumulator-based models, subjects would accumulate left and right pulses independently. To discriminate between these models, we repeated our regression analysis (shown in Figure 5F) using a model in which neurons encoded only the #R or #L pulses and compared the goodness of fit of both models to the data from each neuron (Figure 5I). The delay period activity of most neurons was better predicted by the number of flashes on a single side, rather than by the difference in the number of flashes (PPC, 63% [25/38] for #L, 13% [5/38] for #R, 24% [9/38] for #R-#L; FOF, 27% [17/62] for #L, 37% [23/62] for #R, 34% [21/62] for #R-#L; 1/62 were not best fit by a single model; Figure 5I). These data demonstrate the existence of a population of neurons that accumulate left and right evidence independently; however, the behavioral relevance of these neurons remains to be demonstrated.

Heterogeneous Dynamics across Individual FOF and PPC Neurons in Response to Pulses of Evidence

Next we exploited the pulsatile nature of the stimuli to determine the temporal profile of individual neural responses to single pulses of sensory evidence. In our previous analysis (e.g., Figures 5B and 5C), pulse-triggered responses were evaluated by averaging the $\Delta F/F$ aligned to the time of the pulse. However, since multiple left and right pulses could be presented in sequence, the response of a neuron on each trial likely reflects a combination of responses to multiple flashes. Moreover, our previous analysis revealed that many neurons exhibit responses that depend on the upcoming choice. Therefore, since the response of a neuron at each point in time is a complex combination of multiple task-related variables that are themselves correlated in time, a simple pulse-triggered average would be contaminated by responses to other events.

We implemented multiple linear regression to estimate each neuron's average pulse-triggered response, independent of subsequent pulses and upcoming choice (Figures 6A and S4). We assumed that the response of an individual neuron, at each time point, could be described by the linear function

$$\hat{y} = \mathbf{X}\beta + \epsilon,$$

where \hat{y} is the response of the neuron at each time point, \mathbf{X} is a matrix of regressors for task events, β is the weights of each regressor, and ϵ is a Gaussian noise term. The task events we parameterized were the timing of the light pulses (left and right) and the time within the trial for either left or right choice trials. In order to capture potentially complex temporal dynamics, each of the four task events was parameterized by a set of regressors that formed a temporal sequence, each at a different lag from the time of the event. Weights were calculated using ridge regression. Once fit to the data, the sets of weights corresponding to each task event represent the average response of each neuron to that particular event. The regression model was evaluated by comparing data and model predictions for cross-validated test sets (see STAR Methods). We subsequently focused on the regression weights corresponding to the different task events (e.g., pulses, choice), which we refer to as kernels (Figures S4 and S5).

Examination of the pulse kernels revealed the temporal dynamics of neuronal responses to a single pulse (Figures 6C, 6D, and S5). On average, evidence-tuned neurons responded to a pulse on the preferred side with a large amplitude increase in fluorescence that persisted during the measurable imaging period (~ 2 s; Figure 6B). However, examination of the dynamics of individual cells revealed extensive heterogeneity in the time courses of pulse kernels across the population (Figures 6C and 6D).

To quantify this diversity, we fit the time course of each pulse kernel with a model that consisted of a lag followed by the difference of two exponentials (see STAR Methods). This analysis revealed that neurons from PPC and FOF exhibited responses to flashes on the preferred side with a diversity of time courses and largely overlapping distributions of lag and rise times (Figure 6E). These dynamics were not correlated with the brightness of the neurons, indicating that the temporal heterogeneity was not an artifact of levels of GCaMP expression ($p > 0.05$; see STAR Methods). Viewed in this framework, the neurons we previously defined as “evidence tuned” may represent a subset of a population of cells whose temporal dynamics vary in a continuous way.

Are heterogeneous temporal dynamics unique to visual stimuli, or might they also play a role in representing accumulated evidence in tasks with different sensory modalities? To address this question, we used multiple linear regression to evaluate the click kernel from the firing rates of neurons previously obtained by tetrode recordings in PPC and FOF during an auditory (“Poisson clicks”) accumulation of evidence task (Hanks et al., 2015). Across the population, FOF and PPC exhibited sustained responses to auditory pulses on the time-scale predicted by behavioral models. Averaged click kernels were similar to previously reported “click-triggered averages,” which were estimated based on the analysis of residual responses after taking into account the average response of a neuron on each trial (Figures S6A and S6B; Hanks et al., 2015). As we observed in our imaging data, individual neurons exhibited a wide range of temporal dynamics in response to clicks (Figures 6F and S6I). These results demonstrate that heterogeneous temporal responses to sensory events, present in both layer 2/3 and putative deeper layers, may encode accumulated evidence across different modalities and behavioral tasks.

Diversity of Flash Responses Could Provide a Temporal Basis for Encoding Accumulated Evidence

Previous behavioral analysis has suggested that the animals' memory of sensory evidence is longer than the behavioral trial (> 2 s; Scott et al., 2015). However, the pulse responses of individual cells in PPC and FOF were transient, but started at different times. How might the brain use transient but temporally lagged responses to stably represent the memory of a flash over behaviorally relevant timescales? One possibility is that sufficiently diverse neuronal responses can be used as basis functions to construct a stable memory trace (Medina and Mauk, 2000; Goldman, 2009), or complex time-varying signals that can be used to predict sensory events (Kennedy et al., 2014). To determine whether the flash impulse response

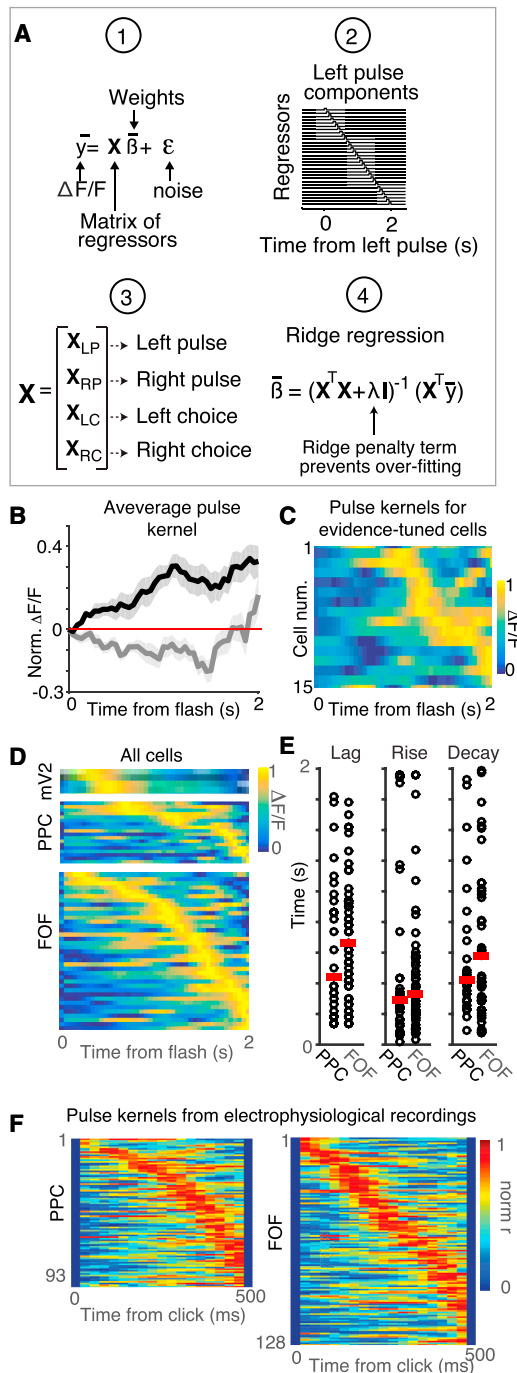


Figure 6. Neuronal Responses to Pulses in Layer 2/3 Exhibit Diverse Temporal Dynamics

(A) Neuronal calcium dynamics described using a linear regression-based approach. (1) The model describes the time-varying neuronal calcium level (y) as a linear combination of a matrix of regressors (X) that depend on task events with weights (β) and Gaussian noise (ϵ). (2) To construct the design matrix (X), each task event (left pulses, right pulses, left choice, and right choice) was parameterized as a set of regressors that take the value 1 at a specific lag relative to time of the event and the value 0 everywhere else. Shown here is an example design matrix for a single left pulse (X_{LP}). Rows correspond to the different regressors. Columns correspond to time points. Regressors

functions we observed were sufficiently diverse to represent possible behavioral accumulation time constants, we constructed a neural network model in which the output of the model was a weighted sum of each neuron's pulse kernel (Figures 7A and 7B). We trained this network to generate an output to match the time course of the sensory memory as predicted by behavioral analysis (Scott et al., 2015). The sensory memory was modeled by a step followed by the exponential function $da = dF + \lambda_1 a dt$, where λ_1 is a leak term fit on a trial-by-trial basis to the behavior of each rat (Brunton et al., 2013). We previously found that λ_1 values ranged from -1.7 (leaky) to 0.5 (unstable) across all rats with a mean of 0.006 (near perfect; Figures 2F and S1; Scott et al., 2015). Using a weighted sum of PPC and FOF impulse response functions, the network was able to match response functions of leaky, impulsive, and nearly perfect accumulators over the range of λ_1 values observed in trained rats (Figures 7C and 7D). This result suggests that the impulse response functions of PPC and FOF were sufficiently diverse to provide a temporal basis for a latent variable behavioral model of evidence accumulation.

We next wondered what circuit mechanisms could generate the diversity of timescales that we observed in PPC and FOF. Recent theoretical work has suggested that effectively feedforward connected networks could produce similar responses to the heterogeneous dynamics we observe (Goldman, 2009; Rajan et al., 2016). Two popular feedforward-based circuit mechanisms include iterative convolution through sequential neural activity, which relies on the synaptic (Bullock et al., 1994) or membrane time constants (Tank and Hopfield, 1987), and delay lines, which rely on axon conductance times (Carr and Konishi, 1988) or delayed transmitter release (Atluri and Regehr, 1998) to

corresponding to pulses represent time points following a flash, whereas regressors corresponding to choices represent time points preceding a choice. Each pulse produces a similar sequence in the complete design matrix (X). (3) The design matrix (X) was constructed from a set of sub-matrices corresponding to the each of the task events (X_{LP} , X_{RP} , X_{LC} , and X_{RC}). (4) The weights for all regressors were fit using ridge regression. The ridge regularizer (λ) was used to prevent over-fitting of the model.

(B) Average response kernels for a pulse on the preferred side (black) and non-preferred side (gray) averaged across all evidence-tuned neurons in PPC and FOF ($n = 15$). Shaded areas indicate SEM.

(C) Response kernels for a pulse on the preferred side for each evidence-tuned neuron. Cells were sorted by the time of the peak in their response. Color indicates the normalized response ($\Delta F/F$).

(D) Pulse kernels on the preferred side for each of the active neurons in mV2, PPC, and FOF. Rows correspond to individual cells sorted by the time of the peak in their response. Color indicates the normalized response ($\Delta F/F$).

(E) Comparison of lag (left panel), rise (middle panel), and decay (right panel) for each cell across PPC and FOF. Circles indicate the best-fit parameters for each neuron. Red line indicates the mean across each group. Note that PPC and FOF exhibited a wide range of best-fit parameters and that parameters were largely overlapping between these areas.

(F) Left panel: response kernels for a pulse on the preferred side for each of the 93 neurons recorded in PPC with tetrodes in the auditory "Poisson clicks" task. Rows correspond to individual cells. Cells were sorted by the time of the peak in their response. Color indicates the normalized response. Right panel: response kernel for a pulse on the preferred side for each of the 128 neurons recorded in FOF. Rows correspond to individual cells. Cells were sorted by the time of peak response. Color indicates the normalized response.

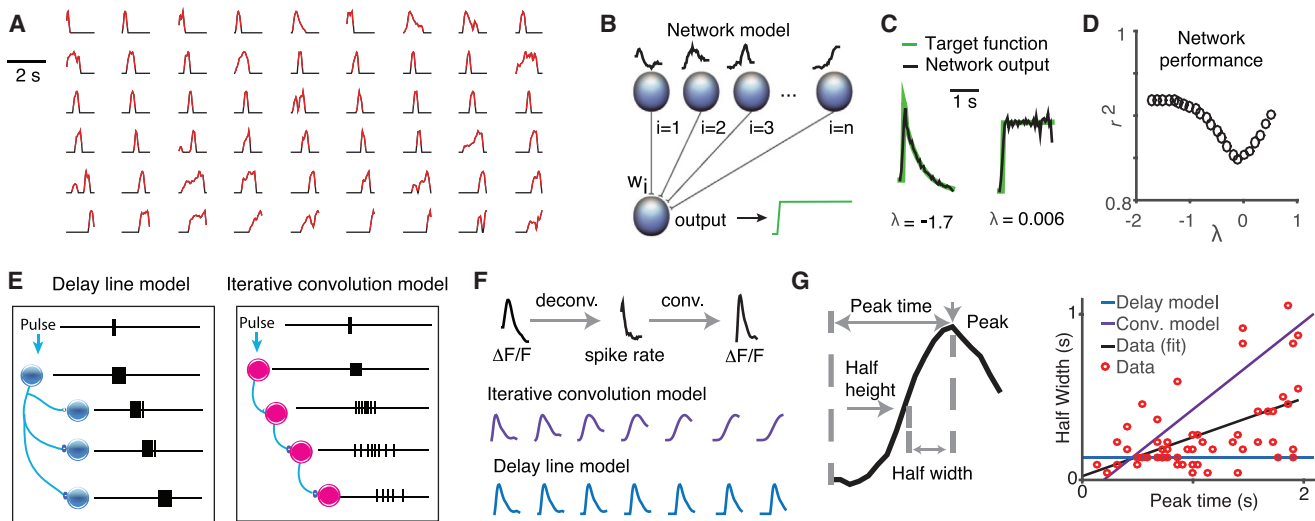


Figure 7. Diverse Dynamics in Fronto-parietal Cortex Could Act as a Temporal Basis Set for Evidence Accumulation

(A) Pulse kernels for neurons in FOF, PPC, and mV2 with positive-going responses to evidence. Periods when the response was statistically significant are colored in red. Waveforms were normalized to the magnitude of the peak response of the cell and sorted by the time of the peak. Periods when the response was not significant were set to zero (colored in black).

(B) Schematic of the architecture of the neural network used to evaluate whether impulse response functions could form a basis set for evidence accumulation. Network output is a weighted sum of the deconvolved impulse response functions of each neuron. The weights for each neuron (W_i) were adjusted using gradient descent to minimize the distance between the network output and the target function. The target functions were constructed using predictions of the time course of the accumulator value, a , derived from a drift diffusion-based behavioral model (Brunton et al., 2013). The time course of the target function was defined by the function $da = Fdt + \lambda a dt$, in which F is the timing of the flashes on the preferred side and λ is a leak parameter that was fit to the behavioral data.

(C) Examples of network output and target functions for $\lambda = -1.7$ (left panel) and $\lambda = 0.006$ (right panel).

(D) Performance of a neural network trained against a range of target function λ values ranged from -1.7 to 0.5 , which is the range observed in rats in our study. Network performance was quantified by the variance explained (r^2) between the trained network output and the target function for $\lambda = -1.7$ through 0.5 .

(E) Schematic indicating two proposed models for the generation of diverse temporal responses in fronto-parietal cortex: delay lines (left panel) and iterative convolution (right panel).

(F) Examples of the waveforms of seven successive simulated neurons in a chain shown for each model. For the iterative convolution model (purple; middle panel), the waveform of the i th neuron was produced by convolving the waveform of $i-1$ th neuron with the waveform of the first neuron in the chain. For the delay line model (blue; lower panel), the waveform of each neuron was produced by adding a 100 ms delay to the response of the previous neuron's waveform. Waveforms were then convolved with the GCaMP6f kernel to allow direct comparison with recorded GCaMP6f signal.

(G) Left panel: schematic of the approach used to quantify the timescale of the impulse response function for each neuron. Right panel: correlation between impulse response function half-width and peak time across all neurons with significant visual components. Each red dot represents the time course of the visual filter on the preferred side for each neuron. Black line indicates the best fit (least-squares) to the data. Purple line represents the prediction of the iterative convolution model, and blue line represents the prediction of the delay line model.

extend responses in time (Figure 7E). Each of these biophysically inspired models makes a different prediction about the relationship between the peak firing time and the half-width of the impulse response function (Figures 7F and 7G, blue and purple lines). We compared these model predictions to the time courses of neurons' impulse response functions from PPC and FOF (Figure 7G). We observed significant correlation between peak time and half-width, consistent with iterative convolution (Pearson's correlation, $p < 0.05$). This correlation was also significant ($p < 0.05$) when we compared the lag and rise times obtained from the parametric fit to the impulse response functions (as in Figure 6E). However, with increasing lags, there was also increased variability in the time course of the flash responses (i.e., the spread of the red dots in Figure 7G). The diversity of responses across the population is not well fit by either the delay line or the iterative convolution model individually and may reflect a combination of these mechanisms or a recurrent-based mechanism.

Heterogeneous Dynamics Encode the Memory of Responses on Previous Trials

Rats make decisions based not only on events in the current trial, but also on their memory of previous choices and reward on previous trials (Scott et al., 2015). PPC has been recently shown to encode the subjects' choices on previous trials in mice (Morcos and Harvey, 2016), and inactivations of rat PPC reduce the influence of trial history on decision-making in a working memory task (A. Akrami et al., 2016, Soc. Neurosci., abstract).

Consistent with studies in mice (Morcos and Harvey, 2016), we observed that populations of neurons in PPC encoded the choice of the animal on the previous trial (Figures 8A and S7). In addition, we found that FOF neurons also encoded the rat's previous choice. Notably, we could decode the current choice more accurately than the previous choice in FOF, while we found the opposite trend in PPC. Next we examined the responses of individual neurons, which often exhibited different responses depending on the choice of the animal on the previous trial

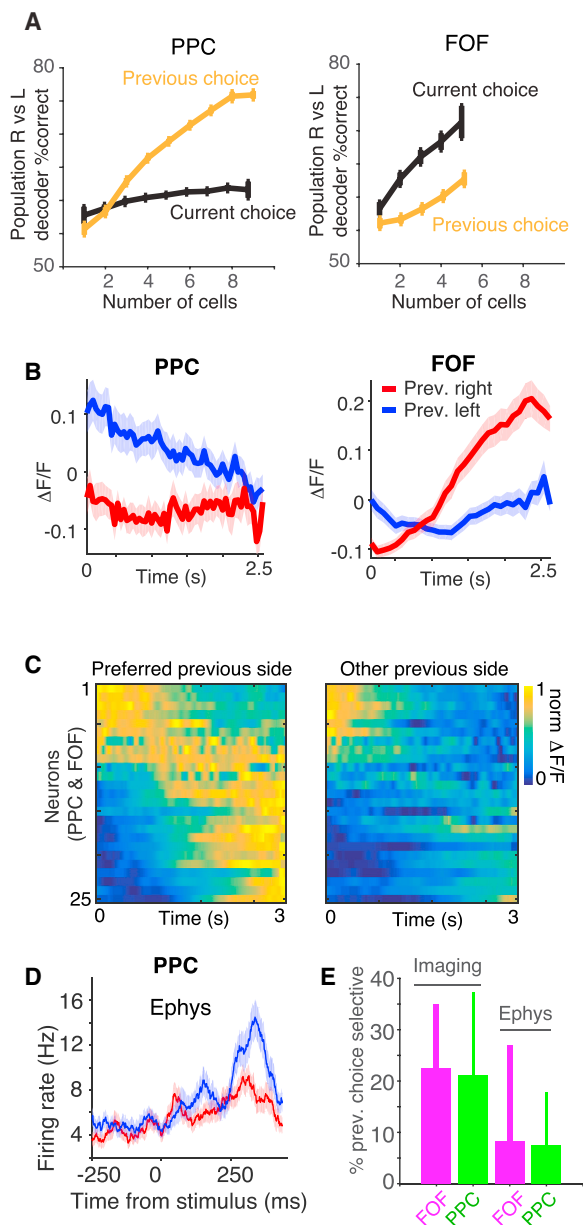


Figure 8. Previous Choice Tuning in FOF and PPC

(A) Performance of a linear classifier in decoding task events based on the activity of simultaneously recorded neurons in FOF and PPC. The choice of the animal on the previous trial was decoded based on the average activity during the entire trial. Left panel: in PPC, the classifier is significantly better at decoding the previous choice of the animal (orange line) than the upcoming choice (black line). Right panel: in FOF, the classifier performs significantly better at decoding the upcoming choice (black line) than the previous choice of the animal (orange line).

(B) Example GCaMP6f responses of two neurons recorded in PPC (left panel) and FOF (right panel) that exhibited different dynamics depending on the previous choice of the animal. Red line indicates the $\Delta F/F$ on trials in which the animal chose right on the previous trial; blue line indicates the average $\Delta F/F$ on trials in which the animal chose left on the previous trial. Shaded region represents SEM.

(C) Dynamics of all previous choice-selective neurons ($n = 25$) separated by the previous choice of the subject and sorted by the time of the center of mass of

(Figures 8B and S7D–S7F). A neuron was defined as “previous choice selective” if the distributions of fluorescence values on trials in which the previous choices were left versus right were significantly different from each other (two-sided t test, $p < 0.05$). Using this metric, 23% of neurons in FOF and 21% of neurons in PPC were previous choice selective. Across the population, previous choice selective neurons displayed heterogeneity in their temporal responses to trial history (Figure 8C). Similar analysis (see STAR Methods) revealed the presence of previous choice selective neurons in our electrophysiological dataset in the auditory version of the task (Figures 8D and 8E): 13.9% (55/397) of neurons in FOF and 17.8% (70/394) of neurons in PPC were previous choice selective. Based on these findings, we modified our linear regression analysis to include a set of regressors that represented the influence of previous choice on the response of the neuron (Figure S8). The addition of regressors that encoded the previous trial did not significantly affect the pulse kernels (Figure S8; also see the Non-parametric permutation test section in STAR Methods). This is not surprising as the response of the animal on the previous trial is uncorrelated with timing of flashes on the current trial.

DISCUSSION

Behavioral Analysis and Neural Recordings Support the Existence of Two Weakly Coupled Accumulators

Two general classes of accumulator-based behavioral models, relative and absolute accumulator models, have been proposed to describe the decision-making process during accumulation of evidence tasks. In relative accumulator models (e.g., the drift diffusion model or DDM), the decision is represented by a single value that reflects the difference in the evidence for and against a particular decision. In absolute accumulator models (e.g., race models), the decision is represented by multiple accumulators that encode the strength of evidence for each option independently. If two separate accumulators are negatively correlated—for example, by receiving the opposite of each other’s input—then the two-dimensional absolute accumulator model reduces to a one-dimensional relative accumulator model (Bogacz et al., 2006), and for simplicity, many neurophysiological studies of evidence accumulation have assumed a relative accumulator framework (Brody and Hanks, 2016; Gold and Shadlen, 2007). However, we previously demonstrated that errors in both visual and auditory pulse-based accumulation of evidence tasks are correlated with both the total number of flashes and the difference in the number of flashes (Scott et al., 2015). In a limited

activity on the preferred side. Each row represents an individual neuron. Color indicates the normalized $\Delta F/F$.

(D) Similar effects are seen in PPC electrophysiological recordings during auditory evidence accumulation. Responses of an example neuron recorded in PPC during the auditory accumulation of evidence task. Red line indicates the average firing rate on trials in which the animal chose right on the previous trial; blue line indicates the average firing rate on trials in which the animal chose left on the previous trial. Shaded region represents SEM.

(E) Comparison of the percentage of previous choice-selective neurons recorded using calcium imaging versus using tetrode recordings. Error bars indicate the 95% binomial confidence interval.

regime of simple models, such a dependency was better explained by a behavioral model in which rats independently accumulated left and right flashes (in which case they would have access to both the total and difference in flash number) rather than the difference between flashes. In the current study, we evaluated whether the activity of neurons is better predicted by the number of either left or right pulses, or by their difference. The activity of most evidence-tuned neurons in PPC and FOF was better predicted by the number of pulses on one side rather than the difference in pulses, consistent with an absolute evidence accumulator model. However, examination of the pulse kernels revealed a small amplitude negative response to pulses on the non-preferred side (Figure 6B). This suggests that neurons tuned to evidence on one side may be weakly inhibited by evidence on the other, and consequently that the accumulators for left and right evidence are neither completely independent nor strongly coupled, but instead weakly coupled (Usher and McClelland, 2001).

Temporally Diverse Sensory Responses and Evidence Encoding in PPC and FOF

Two broad classes of dynamics, homogeneous and heterogeneous, have been proposed for how networks of neurons might represent sensory evidence during decision-making (Figure 1). (1) The homogeneous dynamics model predicts that neurons would respond similarly to pulses of evidence with a sustained increase in activity. While average population data have been used to support this idea (Huk and Shadlen, 2005), the temporal heterogeneity of responses of individual neurons in primate LIP during a continuous accumulation of evidence has been taken as evidence against this model (Meister et al., 2013). (2) The heterogeneous dynamics model predicts that either individual neurons would respond transiently at different times relative to pulses (Harvey et al., 2012), or that individual neurons would respond with diverse, time-varying responses at different times relative to pulses (Brody et al., 2003; Goldman, 2009). Sequential and other forms of heterogeneous dynamics have previously been proposed for memory-based tasks, but whether these responses could also underlie an accumulation process has not been explored.

The diversity in the regression kernels of estimated pulse responses in FOF and PPC was consistent with heterogeneous dynamics, and responses were sufficiently diverse to potentially provide a basis for representing accumulated evidence over several seconds. In this framework, the observation that the responses of individual neurons can be linearly combined to reproduce the dynamics of the decision variable (Figure 7) indicates that FOF and PPC populations represent sensory memory on behaviorally relevant timescales. Whether or not the representations reported here are specific to sensory evidence, or might also apply to accumulation of more abstract forms of evidence (Yang and Shadlen, 2007), is an interesting question that remains to be addressed in future experiments.

Recently, Morcos and Harvey (2016) recorded dynamics in the mouse PPC during an accumulation of evidence task requiring virtual navigation. They observed that in a behavioral trial, PPC transitioned through a sequence of states defined by population activity, and the trajectory through this state space was

correlated with the mouse's upcoming choice. Further work is required to determine the relationship between the sequences observed in mouse PPC and the diversity of timescales in response to individual pulses that we observed in FOF and rat PPC. We note that the data presented in the current study are consistent with sequential or other, more complex heterogeneous dynamics (see Figures 1 and 7G).

Network Models Involving Heterogeneous Dynamics

Neural network models that incorporate time-varying dynamics during working memory appear to confer a number of computational advantages, including increased robustness to perturbations and reduced sensitivity to tuning parameters (Goldman, 2009; Sussillo, 2014; Sussillo and Abbott, 2009). Networks that encode sensory memory in feedforward-like dynamics, in which the network progresses through a sequence of states with distinct activity patterns, have recently been shown to be particularly robust to perturbations and noise (Ganguli et al., 2008; Goldman, 2009; Lim and Goldman, 2012). In these models, sensory stimuli trigger activity in a population of cells that fire at different times relative to the event. Downstream neurons can decode or recover the history of past events by monitoring activity across the population (Klampfl and Maass, 2013; Laje and Buonomano, 2013; Sussillo and Abbott, 2009; Tank and Hopfield, 1987). The neuronal dynamics that we observed bear striking similarities to predictions of these feedforward networks, such as time-varying responses to sensory stimuli and diversity in the timescales of these responses.

Previous studies have also demonstrated that neurons in other frontal cortical areas in rats, such as anterior cingulate cortex and prefrontal cortex, exhibit diverse, time-varying task-relevant responses, suggesting that temporally heterogeneous dynamics may underlie core computations in frontal cortex during behavior (Horst and Laubach, 2012; Jung et al., 1998; Lajtha et al., 2008; Ma et al., 2016; Powell and Redish, 2014; Young and Shapiro, 2011). Moreover, Shankar and Howard (2012) have argued that a set of many neurons responding with different time courses can represent complete information about stimulus history, whereas homogeneous dynamics (Figure 1A) cannot (Shankar and Howard, 2012; Tiganj et al., 2015).

The long timescales we observed in the responses of PPC and FOF neurons to sensory evidence place constraints on mechanistic models of accumulation and working memory in these regions. Previous studies have proposed a variety of biophysically plausible mechanisms that generate temporal diversity in feedforward neural networks. For example, axonal conduction velocities (Braitenberg, 1961; Jeffress, 1948) and synaptic delays in chains of neurons (Licklider, 1951) can produce synaptic events at different times relative to a sensory event. Additionally, the time constant of the cell membrane (~10 ms) or NMDA receptor activation (~100 ms) could act as a convolutional filter to lengthen the timescale of synaptic events in feedforward networks (Tank and Hopfield, 1987; Wang, 2002). Theoretical work has demonstrated that sequential and other feedforward-like dynamics can also be produced by artificial neural networks with recurrent architecture. Given the strong recurrent connectivity frequently observed in rat cortical layer 2/3 (e.g., Barbour and Callaway, 2008; Mason et al., 1991), such artificial networks

suggest a biologically plausible mechanism for the generation of heterogeneous dynamics.

Sequential activation of populations of neurons has been observed in both the rodent neocortex and hippocampus during navigation (Harvey et al., 2012; Wilson and McNaughton, 1993). However, similar dynamics have also been observed in tasks that require working memory but that do not involve navigation. For example, in response to discrete presentations of sensory stimuli, neurons in the rat hippocampus (MacDonald et al., 2013) and mouse somatosensory cortex (Peron et al., 2015) exhibit sequential activity that reflects the identity or location of the remembered stimulus during a delay period even when animals were not navigating. Similarly, in the cerebellum, temporally diverse responses of granule cells have been proposed to provide a basis set for sensory motor learning (Medina and Mauk, 2000) and the cancellation of self-generated sensory signals (Kennedy et al., 2014). Our work suggests that similar principles could allow populations of PPC and FOF neurons with diverse timescales to encode not just single sensory events, but also complex sequences of sensory events, as demonstrated with simulated networks (Klampfl and Maass, 2013; Laje and Buonomano, 2013; Sussillo and Abbott, 2009; Tank and Hopfield, 1987).

The Roles of the Rat PPC and FOF in Decision Commitment

It has been suggested that PPC and FOF may play distinct roles in the decision-making process (Brody and Hanks, 2016). Hanks et al. (2015) reported that neurons in FOF exhibited a more categorical representation of accumulated evidence than neurons in PPC, which the authors suggested reflected FOF's putative role in action selection (Erlich et al., 2015; Hanks et al., 2015; Kopeck et al., 2015). While we observed a greater percentage of side-selective neurons in FOF than in PPC, we did not observe a statistical difference between PPC and FOF in terms of whether representations were categorical or graded (i.e., whether the relationship between activity and evidence was sigmoidal or linear, respectively; see STAR Methods). One important distinction is that in this study, rats were presented with fewer evidence pulses than in Hanks et al. (2015). FOF firing rates at the accumulator values in the current study (up to ± 6 flashes) appear linear in the Hanks et al. (2015) dataset (see their Figure 3C). The larger range of accumulator values examined in that study may have been critical to reveal the categorical nature of FOF responses. In other words, even if representations in FOF were more categorical, the stimuli in our task may have been in a regime in which they appeared linear. Another possibility is that the increased temporal resolution of the electrophysiological measurements of Hanks et al., combined with their use of a model-based estimate of the internal accumulator variable, enabled distinguishing linear versus sigmoidal encoding.

One striking difference between FOF and PPC that we did observe was related to the memory of the previous trial. We confirmed, as reported in a previous mouse study (Morcos and Harvey, 2016), encoding of previous trial choice in PPC. We also found previous trial encoding in FOF. Furthermore, we found a difference between the two areas. While linear decoders based on FOF neurons were slightly better at predicting the upcoming

choice than the previous choice, we found the opposite to be true in PPC: decoding accuracy based on PPC was higher for the previous choice than for the upcoming choice. Distinctions between neuronal responses in parietal and frontal cortices can be elusive (Chafee and Goldman-Rakic, 1998; Hanks et al., 2015). Our results suggest a novel axis on which to explore functional differences between these two regions.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Surgery
 - Behavioral training
 - Image acquisition
 - Offline motion correction
 - ROI identification, neuropil subtraction and fluorescence signal extraction
 - Identification of positive-going and negative-going transients
 - Exclusion of frames due to Z-motion and identification of active cells
 - Electrophysiological recordings
 - $\Delta F/F$ estimation, side selective neurons and evidence tuned neurons
 - Linear decoder analysis
 - Evidence tuning model comparison
 - Multiple linear regression of neural activity
 - Non-parametric permutation test
 - Time course of the pulse kernels
 - Comparing fluorescence and time course of pulse kernels
 - Categorical tuning of FOF and PPC neurons to accumulated evidence
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information includes eight figures and can be found with this article online at <http://dx.doi.org/10.1016/j.neuron.2017.06.013>.

A video abstract is available at <http://dx.doi.org/10.1016/j.neuron.2017.06.013#mmc3>.

AUTHOR CONTRIBUTIONS

B.B.S. and C.M.C. contributed equally to this work. B.B.S., C.D.B., and D.W.T. designed the study. B.B.S. and C.M.C. collected imaging data, performed analysis, and wrote the initial draft of the manuscript. A.A. inspired the previous trial analysis. T.D.H. contributed electrophysiological recordings. All authors were involved in interpreting the data and revising the article.

ACKNOWLEDGMENTS

We thank Mikio Aoi and members of the Tank and Brody labs for useful conversations, and Jovanna Teran and Klaus Osorio for animal husbandry. Jonathan Pillow provided code for implementing evidence optimization, and

helpful guidance for implementing and evaluating the multiple linear regression analysis. Marino Pagan provided the helpful comments and code for the SVM-based decoder. Julia Kuhl provided the illustrations of the rat in Figures 2A and 3C. B.B.S. was supported by a post-doctoral NRSA fellowship, NIH award number 5F32NS78913; C.M.C. was supported by a Helen Hay Whitney fellowship; and T.D.H. was supported by a post-doctoral NRSA fellowship, NIH award number F32MH098572. The work was supported by NIH grant numbers R21NS082956 and U01NS090541 to C.D.B. and D.W.T.

Received: May 21, 2016

Revised: April 4, 2017

Accepted: June 6, 2017

Published: June 29, 2017

REFERENCES

- Atluri, P.P., and Regehr, W.G. (1998). Delayed release of neurotransmitter from cerebellar granule cells. *J. Neurosci.* *18*, 8214–8227.
- Barbour, D.L., and Callaway, E.M. (2008). Excitatory local connections of superficial neurons in rat auditory cortex. *J. Neurosci.* *28*, 11174–11185.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* *113*, 700–765.
- Braitenberg, V. (1961). Functional interpretation of cerebellar histology. *Nature* *190*, 539–540.
- Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. *Curr. Opin. Neurobiol.* *37*, 149–157.
- Brody, C.D., Hernández, A., Zainos, A., and Romo, R. (2003). Timing and neural encoding of somatosensory parametric working memory in macaque prefrontal cortex. *Cereb. Cortex* *13*, 1196–1207.
- Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science* *340*, 95–98.
- Bullock, D., Fiala, J.C., and Grossberg, S. (1994). A neural model of timed response learning in the cerebellum. *Neural Netw.* *7*, 1101–1114.
- Carr, C.E., and Konishi, M. (1988). Axonal delay lines for time measurement in the owl's brainstem. *Proc. Natl. Acad. Sci. USA* *85*, 8311–8315.
- Chafee, M.V., and Goldman-Rakic, P.S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memory task. *J. Neurophysiol.* *79*, 2919–2940.
- Chang, C.C., and Lin, C.J. (2011). LIBSVM: a library for support vector machines. *ACM Trans. Intell. Syst. Technol.* *2*, 27.
- Dombeck, D.A., Khabbaz, A.N., Collman, F., Adelman, T.L., and Tank, D.W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron* *56*, 43–57.
- Erich, J.C., Bialek, M., and Brody, C.D. (2011). A cortical substrate for memory-guided orienting in the rat. *Neuron* *72*, 330–343.
- Erich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* *4*, e05457.
- Espinoza, S.G., and Thomas, H.C. (1983). Retinotopic organization of striate and extrastriate visual cortex in the hooded rat. *Brain Res.* *272*, 137–144.
- Forstmann, B.U., Ratcliff, R., and Wagenmakers, E.J. (2016). Sequential sampling models in cognitive neuroscience: advantages, applications, and extensions. *Annu. Rev. Psychol.* *67*, 641–666.
- Ganguli, S., Huh, D., and Sompolinsky, H. (2008). Memory traces in dynamical systems. *Proc. Natl. Acad. Sci. USA* *105*, 18970–18975.
- Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* *30*, 535–574.
- Goldman, M.S. (2009). Memory without feedback in a neural network. *Neuron* *61*, 621–634.
- Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* *520*, 220–223.
- Harvey, C.D., Coen, P., and Tank, D.W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* *484*, 62–68.
- Horst, N.K., and Laubach, M. (2012). Working with memory: evidence for a role for the medial prefrontal cortex in performance monitoring during spatial delayed alternation. *J. Neurophysiol.* *108*, 3276–3288.
- Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* *25*, 10420–10436.
- Jeffress, L.A. (1948). A place theory of sound localization. *J. Comp. Physiol. Psychol.* *41*, 35–39.
- Jung, M.W., Qin, Y., McNaughton, B.L., and Barnes, C.A. (1998). Firing characteristics of deep layer neurons in prefrontal cortex in rats performing spatial working memory tasks. *Cereb. Cortex* *8*, 437–450.
- Kennedy, A., Wayne, G., Kaifosh, P., Alviña, K., Abbott, L.F., and Sawtell, N.B. (2014). A temporal basis for predicting the sensory consequences of motor commands in an electric fish. *Nat. Neurosci.* *17*, 416–422.
- Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* *2*, 176–185.
- Klampfl, S., and Maass, W. (2013). Emergence of dynamic memory traces in cortical microcircuit models through STDP. *J. Neurosci.* *33*, 11515–11529.
- Kolb, B., and Walkley, J. (1987). Behavioural and anatomical studies of the posterior parietal cortex in the rat. *Behav. Brain Res.* *23*, 127–145.
- Kopec, C.D., Erich, J.C., Brunton, B.W., Deisseroth, K., and Brody, C.D. (2015). Cortical and subcortical contributions to short-term memory for orienting movements. *Neuron* *88*, 367–377.
- Laje, R., and Buonomano, D.V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* *16*, 925–933.
- Lapish, C.C., Durstewitz, D., Chandler, L.J., and Seamans, J.K. (2008). Successful choice behavior is associated with distinct and coherent network states in anterior cingulate cortex. *Proc. Natl. Acad. Sci. USA* *105*, 11963–11968.
- Licklider, J.C. (1951). A duplex theory of pitch perception. *Experientia* *7*, 128–134.
- Lim, S., and Goldman, M.S. (2012). Noise tolerance of attractor and feedforward memory models. *Neural Comput.* *24*, 332–390.
- Ma, L., Hyman, J.M., Durstewitz, D., Phillips, A.G., and Seamans, J.K. (2016). A quantitative analysis of context-dependent remapping of medial frontal cortex neurons and ensembles. *J. Neurosci.* *36*, 8258–8272.
- MacDonald, C.J., Carrow, S., Place, R., and Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *J. Neurosci.* *33*, 14607–14616.
- Mason, A., Nicoll, A., and Stratford, K. (1991). Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. *J. Neurosci.* *11*, 72–84.
- Medina, J.F., and Mauk, M.D. (2000). Computer simulation of cerebellar information processing. *Nat. Neurosci.* *3* (Suppl), 1205–1211.
- Meister, M.L., Hennig, J.A., and Huk, A.C. (2013). Signal multiplexing and single-neuron computations in lateral intraparietal area during decision-making. *J. Neurosci.* *33*, 2254–2267.
- Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. *Nat. Neurosci.* *19*, 1672–1681.
- Park, M., and Pillow, J.W. (2011). Receptive field inference with localized priors. *PLoS Comput. Biol.* *7*, e1002219.
- Park, I.M., Meister, M.L., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat. Neurosci.* *17*, 1395–1403.
- Peron, S.P., Freeman, J., Iyer, V., Guo, C., and Svoboda, K. (2015). A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron* *86*, 783–799.
- Pinto, L., and Dan, Y. (2015). Cell-type-specific activity in prefrontal cortex during goal-directed behavior. *Neuron* *87*, 437–450.

- Powell, N.J., and Redish, A.D. (2014). Complex neural codes in rat prelimbic cortex are stable across days on a spatial decision task. *Front. Behav. Neurosci.* *8*, 120.
- Rajan, K., Harvey, C.D., and Tank, D.W. (2016). Recurrent network models of sequence generation and memory. *Neuron* *90*, 128–142.
- Reep, R.L., Chandler, H.C., King, V., and Corwin, J.V. (1994). Rat posterior parietal cortex: topography of corticocortical and thalamic connections. *Exp. Brain Res.* *100*, 67–84.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* *22*, 9475–9489.
- Scott, B.B., Brody, C.D., and Tank, D.W. (2013). Cellular resolution functional imaging in behaving rats using voluntary head restraint. *Neuron* *80*, 371–384.
- Scott, B.B., Constantinople, C.M., Erlich, J.C., Tank, D.W., and Brody, C.D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *eLife* *4*, e11308.
- Shankar, K.H., and Howard, M.W. (2012). A scale-invariant internal representation of time. *Neural Comput.* *24*, 134–193.
- Sussillo, D. (2014). Neural circuits as computational dynamical systems. *Curr. Opin. Neurobiol.* *25*, 156–163.
- Sussillo, D., and Abbott, L.F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron* *63*, 544–557.
- Tank, D.W., and Hopfield, J.J. (1987). Neural computation by concentrating information in time. *Proc. Natl. Acad. Sci. USA* *84*, 1896–1900.
- Tiganj, Z., Hasselmo, M.E., and Howard, M.W. (2015). A simple biophysically plausible model for long time constants in single neurons. *Hippocampus* *25*, 27–37.
- Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* *108*, 550–592.
- Vintch, B., Zaharia, A.D., Movshon, J.A., and Simoncelli, E.P. (2012). Efficient and direct estimation of a neural subunit model for sensory coding. *Adv. Neural Inf. Process. Syst.* *25*, 3113–3121.
- Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* *36*, 955–968.
- Whitlock, J.R., Pfuhl, G., Dagslott, N., Moser, M.B., and Moser, E.I. (2012). Functional split between parietal and entorhinal cortices in the rat. *Neuron* *73*, 789–802.
- Wilber, A.A., Clark, B.J., Demecha, A.J., Mesina, L., Vos, J.M., and McNaughton, B.L. (2015). Cortical connectivity maps reveal anatomically distinct areas in the parietal cortex of the rat. *Front. Neural Circuits* *8*, 146.
- Wilson, M.A., and McNaughton, B.L. (1993). Dynamics of the hippocampal ensemble code for space. *Science* *261*, 1055–1058.
- Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. *Nature* *447*, 1075–1080.
- Young, J.J., and Shapiro, M.L. (2011). Dynamic coding of goal-directed paths by orbital prefrontal cortex. *J. Neurosci.* *31*, 5989–6000.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and Virus Strains		
AAV1.Syn.GCaMP6s.WPRE.SV40	Penn Vector Core	Cat.# AV-1-PV2824
AAV2/1.CB7.Cl.mCherry.RBG	Penn Vector Core	Cat # AV-1-PV1969
AAV1.CAG.GCaMP6f.WPRE.SV40	Penn Vector Core	Cat # AV-1-PV3081
Experimental Models: Organisms/Strains		
Rat: Long Evans	Taconic	RRID: RGD_1566430
Rat: Sprague Dawley	Taconic	Model #: SD-M
Software and Algorithms		
MATLAB	MathWorks	RRID: SCR_001622
ScanImage	Vidrio Technologies	RRID: SCR_014307
Bcontrol	Erlich et al., 2011	http://brodywiki.princeton.edu/bcontrol/index.php/Main_Page
ImageJ	https://imagej.nih.gov/ij/index.html	RRID: SCR_003070
LIBSVM	Chang and Lin, 2011	http://www.csie.ntu.edu.tw/~cjlin/libsvm/
Other		
Blue LEDs	Super Bright LEDs	Part #: RL5-B3245
Tetrode wire	Neuralynx	Pt1r 18 um
Polyimide Tubing, 34 AWG	https://www.amazon.com/	Part #: TWPT-0063-30
Threaded Rod, #0-80	https://www.amazon.com/	Part #: B000FP0JA6

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for reagents may be directed to and will be fulfilled by the Lead Contact, Dr. David Tank (dwtank@princeton.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Animal use procedures were approved by the Princeton University Institutional Animal Care and Use Committee (IACUC; Protocols #1837 and #1853) and carried out in accordance with National Institutes of Health standards. Seven adult male rats, four Long-Evans and three Sprague Dawley (obtained from Taconic Biosciences), were used in this study. Rats were typically housed either in pairs or singly in a reverse light cycle room. Access to water was scheduled to within-box training, 2 hr per day, and between 0 and 1 hr ad lib following training.

METHOD DETAILS

Surgery

Rats underwent two different types of surgical procedures during the course of the experiment described here: headplate implantation and optical window implantation. Headplate implantation typically occurred before behavioral training, although in one rat it was necessary to move the headplate to a new location after it had become well-trained. In this surgery rats were surgically implanted with a custom-made, 5-g, titanium headplate that was bonded to the skull using a dental cement adhesive (Metabond). The headplate contained two important design features (see [Scott et al., 2013](#) for details). First, it contained a rectangular window, approximately 8.8 mm by 15mm, centered in the headplate, that allowed access to the skull and second it contained a conical depression and V-groove that interacted with the kinematic clamp to provide mechanical registration and stability during head restraint. The headplate was implanted so that the center window lay over the stereotaxic coordinates for one of the three brain regions in this study.

After training animals underwent a second surgery in which the optical window was implanted. In this surgery, a 3.5mm diameter circulate trephination was made centered on the stereotaxic coordinates for PPC and FOF. The dura was dissected away and

25–50 nL of viral vector was injected. In some animals the gene for the red fluorescent protein marker mCherry was also expressed in neurons in order to determine the magnitude and timing of fluorescence changes due to brain motion and other imaging artifacts. In 4 rats the viral vector was a cocktail containing 1 part AAV-hsyn-2.1-GCaMP6f and 1 part AAV2.1-mCherry diluted in 1 part saline. In another 4 rats 1 part AAV cag 2.1 GCaMP6f was diluted in 1 part saline. All viruses were obtained from the Penn Vector Core (Gene Therapy Program, Perelman School of Medicine, University of Pennsylvania). After injection of the viral vector an optical implant, consisting of a 3.5 mm glass coverslip bonded to a 3.5 mm diameter, 0.8mm high, stainless steel ring using optical adhesive, was placed on the surface of the cortex with the glass side down. The implant was positioned so that it depressed the surface of the cortex by about 0.2mm in order to reduce brain motion and bonded to the skull using a combination of medical grade cyanoacrylate adhesive (Vetbond) and dental cement adhesive (Metabond). Extreme care was taken to minimize trauma to the dura and cortical surface during the surgical procedure and also to maintain sterile conditions. GCaMP6f and mCherry labeling could be observed beginning 1–2 weeks after surgery (Figure 2B) and calcium-dependent transients could be recorded in GCaMP6f-labeled neurons using two-photon imaging during voluntary head restraint, typically for 2–4 weeks and in one case, up to 6 months. Finally, we note that although circular 3.5mm optical windows were implanted over the stereotaxic coordinates for PPC and FOF, clipping of the optical beam by the edges of the cannula made it very difficult, if not impossible, to image cell bodies more than a millimeter away from the center of the window (i.e., our coordinates).

Behavioral training

Following recovery rats trained in a high-throughput, computer controlled, behavioral-training chamber. Over the course of two months, rats acclimated to human handling and progressed through 3 initial stages of training. In the first stage rats learned to initiate a behavioral trial by inserting their nose into a center nose port. After trial initiation a water reward became available at either the left or right side reward port, indicated by a light illuminated at the port. Initially the nose port was flush with the front wall of the behavioral chamber, but after each trial the nose port, mounted on a linear translation stage, moved further away from the interior of the chamber so that the rat had to insert its head and headplate further into the head port to initiate a trial. Once the rat inserted far enough to trigger a pair of contact sensors at the end of the head port slot they transitioned to the second stage of training. In the second stage rats were trained to wait in the head port for increasingly long durations before a go cue instructed them to break fixation. In the third stage of training rats learned to associate LED flashes with reward. For a more detailed description of behavioral training and analysis see [Scott et al. \(2015\)](#).

Image acquisition

Two-photon imaging was performed using a custom movable objective microscope that has been previously described ([Scott et al., 2013](#)). Frame rate was 11 or 22 Hz. Typical FOV size was 150 microns by 300 microns. ScanImage software controlled the movement of the scan mirrors as well as digitization and recording of four analog input channels. Two input channels were used to record signals from the two PMTs (one for green light, one for red light) and the other two channels were used to record the timing of task events (LED flashes and the choices of the animal) from the behavioral control software. To prevent light leakage from the LEDs to the PMTs we used blue shifted LEDs (470nm mean wavelength, Super Bright LEDs) combined with a 450nm lowpass filter.

Offline motion correction

Following acquisition, image stacks were motion corrected using custom scripts written in MATLAB. First a template frame was identified manually. On sessions where expression was low, the template frame was typically composed of the average of up to 10 manually selected individual frames. Next, 2D cross correlation was performed between the template frame and each individual frame. Individual frames were aligned to the template based on the X and Y shifts that produced the highest frame-to-template cross correlation value. After correction we inspected each image stack for residual motion. Frames with evident Z motion were identified and their peak correlation value was recorded. We then excluded all frames with a peak correlation value less than the peak correlation value of the identified frame. This process was repeated until the resulting image stack contained no visible Z-motion.

ROI identification, neuropil subtraction and fluorescence signal extraction

After motion correction, region of interests (ROIs) corresponding to the neuronal somata were identified as previously described ([Scott et al., 2013](#)). We used visual inspection to identify the position of cell bodies of the image stack, demarcating pixels corresponding to the edges of cells using the ImageJ ROI manager feature. Fluorescence traces were obtained by averaging the intensity values of the pixels corresponding to each ROI for each frame. Linear interpolation was used to estimate the fluorescence values of frames that had been excluded by the motion correction algorithm. For neuropil analysis we selected a rectangular 10x10 pixel region centered on each neuron that excluded all pixels that had been identified as belonging to a neuronal ROI ([Scott et al., 2013](#)).

Identification of positive-going and negative-going transients

Once fluorescence traces were extracted, we calculated the baseline fluorescence level for each ROI. The baseline fluorescence (F) was defined as the mode fluorescence value across all time points. The standard deviation (σ_F) in the baseline was defined as

$$\sqrt{\frac{2}{N} * \sum (X - F)^2},$$

where X are all values in the fluorescence trace less than F . A positive-going significant transient was defined as a portion of the fluorescence trace that exceeded the baseline fluorescence value by more than 3 standard deviations ($> F+3\sigma_F$) for more than 0.75 s. A negative going transient was defined as a portion of the fluorescence trace that fell below the baseline by more than 3 standard deviations ($< F-3\sigma_F$) for more than 0.75 s.

Exclusion of frames due to Z-motion and identification of active cells

Brain motion, particularly Z-motion, can produce fluorescence changes as cells move into and out of the field of view. Previous studies (Dombeck et al., 2007; Scott et al., 2013) have used negative-going transients to estimate the rate of such artifacts. We compared the rate of positive going and negative going transients for each ROI. If greater than 5% of all transients for an ROI were negative-going, we raised the correlation threshold value and excluded frames whose peak correlation fell below that threshold. Linear interpolation was used to estimate the fluorescence values of the excluded frames. This process was repeated until negative going transients constituted less than 5% of all transients. If a threshold correlation could be identified at which negative going transients constituted less than 5% of all transients and positive going transients occurred at a rate of more than one per minute that ROI was defined as an active neuron.

Electrophysiological recordings

Electrophysiological recordings characterized here were presented previously (Hanks et al., 2015). Once trained, rats were implanted with custom microdrives containing tetrodes. Tetrode positions for FOF were centered at +2 anterior–posterior (AP), ± 1.3 medio–lateral (ML) (mm from Bregma); and those for PPC were centered at -3.8 to -4.1 AP and ± 2.2 ML. Recordings were made with insulated platinum iridium wire (16.6 μm diameter; Neuralynx) twisted into tetrodes. Each tetrode was threaded into a polyimide tube (34 AWG triple wall) which was part of a movable bundle of eight tubes. Within each bundle, tetrodes were spaced ~ 250 μm from each other. The tetrodes could be advanced by turning a nut against a spring on a 0-80 threaded rod so that a 1/8 turn drove the tetrodes down about 40 μm . The tetrodes were advanced at the end of sessions so that the brain tissue had time to stabilize before recording the next day. Depth relative to the surface of the brain was estimated based on the number of turns of the screw required to reach that recording site, and recording sites were verified with post-mortem histology.

To fit the linear regression model to the electrophysiology data, data were pre-processed as described in Hanks et al. (2015). Briefly, trial firing rate functions were generated by smoothing the spike trains with a causal half-Gaussian filter, with a width of 25ms. The model was fit to this smoothed firing rate. Task-related model components (left and right clicks, left and right choice) were parameterized by temporal bases defined by delta functions at discrete lags (in steps of 25ms) relative to the event. Data collected before the stimulus start time was discarded.

$\Delta F/F$ estimation, side selective neurons and evidence tuned neurons

$\Delta F/F$ for each ROI was defined as follows:

$$\frac{\Delta F}{F} = \frac{F(t) - F_0}{F_0},$$

where F_0 is the mean fluorescence value of all frames before the cue period, and $F(t)$ is the fluorescence value at each time point. To identify side-selective neurons and evidence-tuned neurons, we examined the average $\Delta F/F$ during the delay period on all correct trials. If the distribution of $\Delta F/F$ values were significantly different between left and right trials, according to a two-tailed two-sided t test with $\alpha = 0.05$, that neuron was termed side-selective. Next, for each neuron we performed robust linear regression (MATLAB Statistics Toolbox `robustfit.m`) to determine the relationship between the average delay period response and the numerical differences between flashes ($\#R-\#L$). Evidence index (EI) was defined as the slope of the best fit a linear relationship between $\Delta F/F$ and $\#R-\#L$. A cell was defined as evidence tuned if its evidence index was significantly different from 0 and the correlation coefficient (r^2) was greater than 0.5.

Linear decoder analysis

A support vector machine based linear decoder (Chang and Lin, 2011) was implemented in MATLAB using an existing library available online: <http://www.csie.ntu.edu.tw/~cjlin/libsvm/>.

The decoder was trained to predict either the choice of the animal on the current trial (choice), the choice of the animal on the previous trial (previous choice) or the side with the greater number of flashes (correct side) on each trial using the mean activity during the delay period for multiple neurons. A randomly selected subset of half of the trials was used to train the decoder (training set) and the other half was used to evaluate the decoder (test set). The performance of the decoder for each group of cells reported in the text reflects the mean performance across 100 different training and test set combinations. Groups of cells were constructed from all possible combinations of simultaneously reported neurons for example if cells A, B and C were recorded simultaneously then performance would be evaluated for the following combinations: A, B, C, AB, BC and ABC.

Evidence tuning model comparison

To evaluate whether the data was better fit by a single and two accumulator model we performed the following statistical test: For each neuron we created 1000 surrogate datasets using a bootstrapping method where each surrogate dataset consisted of a randomly selected subset of 75% of the trials. For each surrogate dataset we computed the r-squared value for the best fit linear relationship between the average activity of the neuron during the delay period and the accumulated evidence on each trial. Accumulated evidence was defined as number of right flashes (#R), number of left flashes (#L) or the difference in the number right and number left (#R-#L). Next, we performed a pairwise permutation test (see below and [Scott et al., 2015](#)) to compare the distribution of R-squared values for the fits of the data to each of the three models. The model with the highest R-squared value was considered the best fit model.

Multiple linear regression of neural activity

Each active neuron's response ($\Delta F/F$ for calcium imaging; instantaneous firing rate, r , for tetrode recordings) was modeled as a linear combination of task-related events (assuming Gaussian noise). Importantly, randomness and variability in the timings of task-related events enabled dissociating their effects on the neural response. Task-related events (here called model components: left and right flashes, left and right choice) were parameterized by temporal bases defined by delta functions at discrete lags (in steps of frame rate, usually 23 Hz [50 ms], but sometimes 12 Hz [100 ms]) relative to each event, following each left and right flash, and preceding the end of the trial (or the choice epoch). Each model component covered a specified range of time: for flashes, the temporal basis extended to 2 s after flash onset, and for choice, bases covered the maximum trial duration in each imaging session, aligned to the end of the trial. The total number of parameters for each neuron depended on the frame rate and maximum trial duration in the imaging session. The median number of parameters was 212, and ranged from 94 to 222.

The coefficients were fit using ridge regression. The ridge regularizer (λ) was selected using a procedure called evidence optimization, or type II maximum likelihood estimation, which is a maximum-likelihood procedure for estimating the prior distribution of model parameters (governed by λ) from the data (see [Park et al., 2014](#); [Park and Pillow, 2011](#)). Evidence optimization maximizes the *marginal likelihood* of the data given λ . The standard deviation estimates of the regressors were obtained by taking the square root of the diagonal of the posterior covariance matrix (which, in our linear regression model was $(X^T X + \lambda I)^{-1}$).

Model predictions were generated by convolving the task-related model components with delta functions at the times of the task events. Goodness-of-fit (cross correlation, variance explained) was computed using the predicted and actual response ($\Delta F/F$ or r) at each time point, from the cross-validated test sets. The cross validation procedure was as follows: on each of ten iterations, 75% of the data was partitioned into a training set, 25% partitioned into a test set, and the model prediction was generated for the test set. The average model prediction over all iterations was computed, and used to evaluate model performance on the average test set over all iterations. To evaluate whether regression could recapitulate the evidence tuning of individual neurons, we performed 20-fold cross validation to ensure that sufficiently diverse trial types were represented in the test set.

Model predictions were compared against the responses of each of the neurons averaged across different trial types. The regression model reproduced the temporal profile of the average responses on left and right choice trials. To quantify this, we calculated the correlation between the average model prediction and corresponding trial-averaged data on left and right choice trials. For neurons in PPC the model and data exhibited a correlation coefficient of 0.71 (± 0.23), and for neurons in FOF, the model and data exhibited a correlation coefficient of 0.88 (± 0.16), compared to 0.0027 ($\pm .01$) for shuffled controls. We also found that the model faithfully predicted the graded response to accumulated evidence for each cell. While both the choice components and visual components were required to reproduce the average trial dynamics of each neuron, the visual components alone were sufficient to predict the evidence tuning.

We evaluated the visual components of the model in the mV2 neurons we recorded. Typical mV2 cells exhibited strong, reliable responses to visual pulses that were evident on individual trials. We fit the regression model to these cells, and compared the visual components to model-free estimates of the neurons' flash-triggered averages. Model fits predicted visual components with similar time courses to those seen in the flash-triggered averages in mV2, and also in PPC and FOF.

Next we evaluated how well the model predicted the activity of each neuron at each time point. For PPC neurons the data and model at each time point exhibited an average correlation coefficient of 0.21 (± 0.17), and for FOF neurons, a correlation coefficient of 0.41 (± 0.16). Although these correlation values are comparable to what has been observed in similar encoding models ([Pinto and Dan, 2015](#)), we sought to evaluate the influence of trial-to-trial variability on model performance. We reasoned that the model could not perform better than the average response on identical trials ([Vintch et al., 2012](#)). Therefore, for identical trials, in which the exact same stimuli and behavioral choice were present, we compared the correlation between the data on each trial and the trial average, as well as the data and the model on those same trials. The average response on repeated trials exhibited a correlation coefficient of 0.78 with the data, compared to a correlation between the model and data of 0.39 on those same trials. The ratio of these values, 0.50, reflects the performance of the model normalized by the predictive power of the data on repeated trials.

Non-parametric permutation test

To evaluate whether trial history affected the estimated pulse responses, we performed the following permutation test for each neuron. On each of 100 iterations, we computed the R-squared values between the estimated pulse responses fit using randomly subsampled trials from all of the data and from subsampled trials following left choices (i.e., in which the choice history was fixed).

We compared this distribution of R-squared values to R-squared values obtained from randomly subsampling all trials. The null hypothesis is that the difference between the estimated pulse responses following left choices and the estimated pulse responses fit to all trials is not significantly different from the pulse responses obtained from two different random subsamples of the data. Therefore, the area of the distribution of R-squared values from pulse responses fit to randomly subsampled trials at the average R-squared value for post-left-choice-pulse response and randomly subsampled pulse response corresponded to the p value. For all neurons, the p value was greater than 0.05, indicating that there was no significant difference between the estimated pulse responses obtained from all trials and from those with fixed trial history (i.e., following left choices).

Time course of the pulse kernels

To quantify the diversity in the time course of sensory responses we fit the following five-parameter $Y(t | L, O, D, R, M)$ model to each neuron's pulse triggered response:

for $t \leq L$,

$$Y(t) = O;$$

for $t > L$,

$$Y(t) = O + MX \left(e^{-\frac{t-L}{D}} - e^{-\frac{t-L}{R}} \right).$$

This difference of exponentials model includes a parameter for lag (L), decay time (D), rise time (R), magnitude (M) and offset (O). Model parameters were fit using an interior-point algorithm that attempted to minimize the mean squared error between the pulse-triggered response and the model.

Comparing fluorescence and time course of pulse kernels

To evaluate whether neurons with longer flash responses had different calcium buffer concentrations (because of GCaMP6f-expression), we computed the average raw fluorescence of the soma during the pre-cue (i.e., baseline) period, and the average fluorescence of the surrounding neuropil during the pre-cue period. The ratio of these values was not significantly correlated with the lag, rise time, or decay of the flash impulse response function ($p > 0.05$, Pearson's correlation).

Categorical tuning of FOF and PPC neurons to accumulated evidence

For each evidence-tuned neuron we computed the average $\Delta F/F$ during the delay period pooled across trials with the same evidence, i.e., $\#R-\#L$. We then found the line and sigmoid function that best characterized the relationship between $\Delta F/F$ and evidence. The line was described by

$$Y = \beta_1 + E\beta_2$$

and the sigmoid was described by

$$Y = \beta_1 + \frac{\beta_2}{1 + e^{-\frac{E - \beta_3}{\beta_4}}}$$

$\beta_1 \dots \beta_4$ represent the best-fit parameters, using the MATLAB functions `robustfit.m` for the linear model and `nlinfit.m` for the sigmoidal model. E is the difference in the number of flashes ($\#R-\#L$ if the neuron preferred right evidence, $\#L-\#R$ if the neuron preferred left evidence). Y is the model prediction of the neuronal response ($\Delta F/F$).

Next, we compared the goodness of fit for these two models by computing the chi-square probability, Q . Note that Q depends both the chi-square statistic and the degrees of freedom of the model, which enables a direct comparison of the two models even though the number of parameters is different for each model (2 for linear, 4 for sigmoid). The model that had the greater Q value was considered to be the better fit. Using this criterion, 100% (4/4) of evidence tuned neurons in PPC neurons were better fit by the linear model. 94% (16/17) of evidence tuned neurons in FOF neurons were better fit by the linear model. We also computed the slope predicted by the sigmoid function, β_4 in the equation above. The mean slope for PPC neurons was 8.3 and the mean for FOF was 2.6, however this difference was not statistically significant ($p > 0.05$).

QUANTIFICATION AND STATISTICAL ANALYSIS

The details of all statistical tests performed are described in the [Method Details](#) section above. The number of cells, animals and behavioral trials used for each analysis are reported in the text and figure captions. Error bars and shaded regions on plots reflect the standard error of the mean unless otherwise noted in figure captions.