Neural circuit models for evidence accumulation through choice-selective sequences

Lindsey S. Brown[1], Jounhong Ryan Cho[1], Scott S. Bolkan[1], Ilana B. Witten[1], Mark S. Goldman[2,3]

1 Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA
2 Center for Neuroscience, Department of Neurobiology, Physiology and Behavior, University of California, Davis, Davis, CA, USA
3 Department of Ophthalmology and Vision Science, University of California, Davis, Davis, CA, USA

ABSTRACT

Decision making involves accumulating evidence for or against available options. Traditional circuit models of evidence accumulation generate firing rates that ramp with evidence. However, recent decision-making experiments in rodents have observed neurons across the brain that fire sequentially as a function of position or time, rather than persistently, with the subset of neurons in this sequence depending on the animal's choice. We develop two new candidate circuit models that produce such choice-selective sequences and make competing predictions about the nature of evidence encoding. The first model consists of two chains of neurons, with the location along the chain representing the animal's position and the relative amplitude of firing in the two chains representing accumulated evidence. The second model consists of a plane of neurons, with the set of active neurons in this plane encoding evidence and position. The models make distinct experimental predictions for the shapes of neuronal tuning curves and responses to localized optogenetic stimulations. More generally, this work demonstrates how graded-value information may be precisely accumulated within and transferred between neural populations, a set of computations fundamental to many cognitive operations.

INTRODUCTION

The accumulation of evidence for different alternatives is thought to be a fundamental cognitive component of decision-making. As such, neural recordings have been performed in many tasks designed to probe how evidence is accumulated to make a decision, including tactile (Luna et al., 2005; Mountcastle et al., 1990), visual (Afraz et al., 2006; Britten et al., 1992, 1993; Britten & Van Wezel, 2002; DeAngelis & Uka, 2003; Heekeren et al., 2004; Heuer & Britten, 2004; Philiastides & Sajda, 2006; Shadlen et al., 1996), and olfactory (Uchida & Mainen, 2003) tasks, motor tasks (Carpenter, 1988), and value discrimination tasks (Sugrue et al., 2005). Experiments using these tasks have been conducted with the goal of understanding how the neural circuitry within and across different brain regions contributes to the accumulation of evidence, subsequent maintenance of this information in working memory, and ultimate commitment to a decision based on this evidence (Gold & Shadlen, 2007).

Many models have been proposed to describe the decision making process. A major class of models for two-choice forced alternative decisions is the drift diffusion model (Laming, 1968; Ratcliff, 1978; Stone, 1960), in which evidence is accumulated with a noisy drift process, and decisions are made based on when the resulting accumulation reaches some threshold. The drift-diffusion model reaches the fastest decision for a given desired accuracy, making this model optimal. Many circuit-based models have been shown to be equivalent to the

drift-diffusion model and hence also achieve this optimality (Bogacz et al., 2006). Such circuit-based models include mutual inhibition models (Usher & McClelland, 2001), feedforward inhibition models (Ditterich et al., 2003), and pooled inhibition models (Wang, 2002). These basic circuit models have been extended to explain a number of different decision-making experiments with recordings in different regions where neurons show persistent activity encoding accumulated evidence (Engel & Wang, 2011; Lam et al., 2022; Lo & Wang, 2006; Machens et al., 2005; Roach et al., 2023; Wong & Wang, 2006).

The essence of the circuit-based models is the accumulation of evidence-encoding signals along a low-dimensional attractor representing the decision variable. Thus, these models predict that activity of single neurons ramps up or down with evidence throughout a decision-making task (Goldman et al., 2009; Gold & Shadlen, 2001). Such ramping activity during decision-making tasks has been observed in a number of different brain regions in different paradigms, including in the lateral intraparietal area (LIP), medial temporal area (MT), and ventral intraparietal area (VIP) of monkeys in a random dot motion discrimination task (Cook & Maunsell, 2002; Doya, 2007; Roitman & Shadlen, 2002; Shadlen & Newsome, 2001), in the dorsal premotor cortex of monkeys during discrimination tasks (Chandrasekaran et al., 2017; Peixoto et al., 2021), and in the posterior parietal cortex (PPC) and frontal orienting fields (FOF) of rats in an auditory evidence accumulation task (Boyd-Meredith et al., 2022; Hanks et al., 2015). These classes of neural circuit models that accumulate evidence through persistent, ramping activity successfully capture much of the observed experimental data.

However, a mounting body of evidence from large scale recordings in rodents during a range of decision-making tasks suggests that in many decision-making contexts, neurons across cortical and subcortical regions are not persistently active. Instead, neurons fire transiently and sequentially in a choice-specific manner. This occurs in decision-making tasks both with navigation (Akhlaghpour et al., 2016; Gage et al., 2010; Harvey et al., 2012; Koay et al., 2022; Krumin et al., 2018; MacDonald et al., 2011; Nieh et al., 2021) and without navigation (Cox et al., 2023; Parker et al., 2022; Pastalkova et al., 2008; Siniscalchi et al., 2019; Terada et al., 2017). In a recently developed task (Koay et al., 2022; Nieh et al., 2021), correct decisions are dependent upon integrating evidence during a single trial, rather than just holding a cue in working memory or learning choice value over several trials. These data challenge existing models of evidence accumulation by contradicting a core tenet of these models: the presence of persistent, rather than sequential, neural activity. Therefore, new models are needed that accumulate evidence through sequential neural activity.

For a model to accumulate evidence through sequential neural firing, it must perform two key computations. First, it must transfer information between neurons at sequential positions. Second, it must accumulate evidence received at each position. Traditional models only perform the latter computation. Here, we develop two classes of neural circuit models that perform these computations. In each model class, the transfer of accumulated evidence between neurons encoding sequential positions is accomplished via a position-modulated gating signal. The key difference in the models is in their encoding of evidence. The first class is a chains-based model, consisting of two feedforward chains of neurons that represent accumulated evidence in the relative *amplitude of their firing rates* at a given position, so that accumulated evidence is encoded monotonically in the neuronal firing rates. The second class consists of a planar configuration of neurons that jointly encodes evidence and position in the network *location* of a

stereotypically shaped, unimodal pattern ('bump') of neuronal firing. Thus, neurons in this model class have non-monotonic tuning curves.. We demonstrate that both models produce choice-selective sequences in their neural activity and can perfectly accumulate evidence and transfer it between neighboring positions. We then leverage the models to make experimentally testable and differentiating predictions for the two model classes.

RESULTS

Traditional Integrator Models Fail to Explain Choice-Selective Sequences in the Accumulating Towers Task

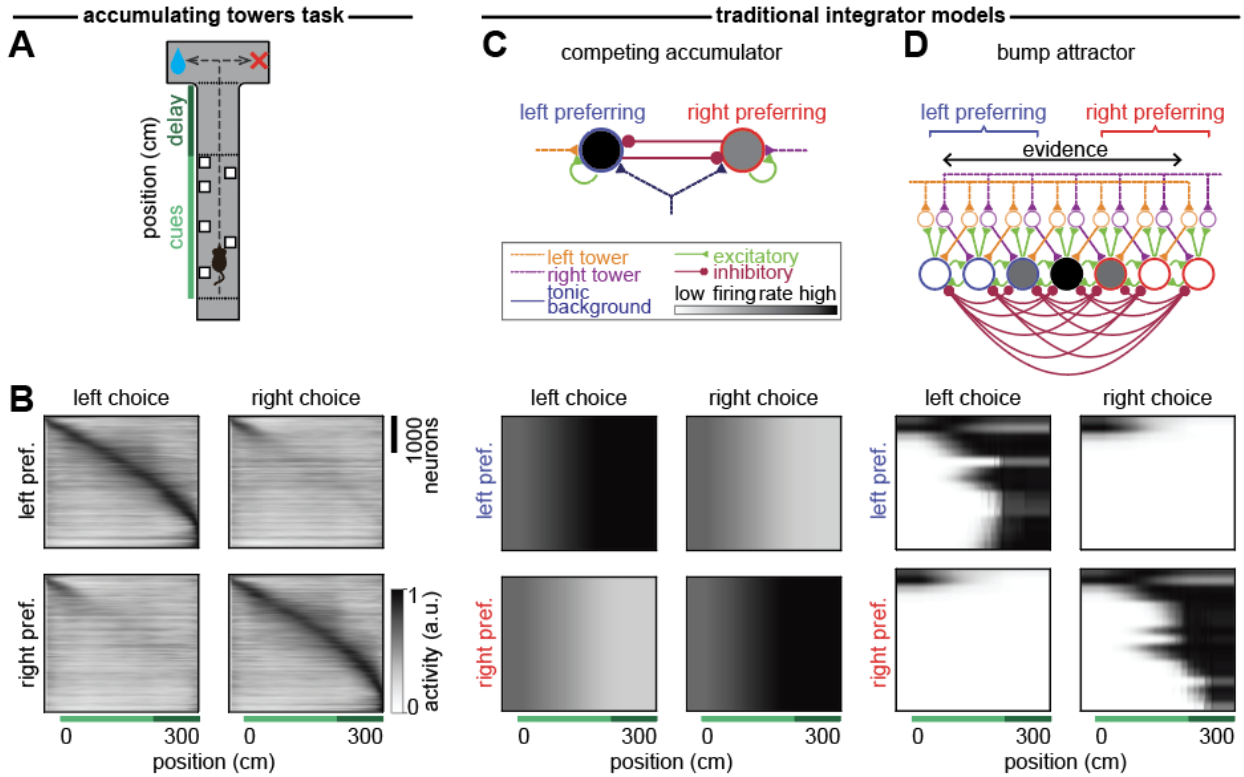

Figure 1. Choice-selective sequences are observed across brain regions but not predicted by traditional integrator models in a navigation-based, accumulation-of-evidence task. (A) Schematic of a behavioral task in which mice navigate a T-maze in virtual reality. Visual cues (towers, white rectangles) are presented to either side of the maze during the cue period (light green, 200 cm). Mice must remember which side had more towers during the delay period (dark green, 100 cm) and receive a water reward for turning to the side with more towers. **(B)** Each row shows the peak-normalized estimated firing rate at each position in the maze of a neuron with significant evidence tuning (see Methods) recorded during this task from primary visual cortex (V1, n = 1121 neurons, data from (Koay, 2021)), anterior cingulate cortex (ACC, n = 964 neurons, see Methods), retrosplenial cortex (RSC, n = 4190 neurons, data from (Koay, 2021)), hippocampus (HPC, n = 791 neurons, data from (Nieh et al., 2021)), and dorsal medial striatum (DMS, n = 322 neurons, see Methods), averaged across trials when the animal turned left (left choice, left column) or right (right choice, right column). Neurons were divided based on their evidence-selectivity (see Methods) and ordered based on the position of peak activity. **(C)** Schematic of a traditional mutually inhibitory integrator model for decision making (top). Average peak-normalized, z-scored simulated activity of neurons in the model while performing this task, sorted as in (B) (bottom). **(D)** As in (C), but for a traditional ring attractor model.

The models presented below are motivated by a large dataset of neural recordings obtained during a navigation-based, accumulation of evidence task in virtual reality ('accumulating towers task', Fig. 1A). In this task ((Pinto et al., 2018); see Methods), mice navigate a T-maze with Poisson counts of visual cues in the form of towers appearing to each side. After a delay region with no cues, the mouse reaches the arms of the maze and is rewarded for turning to the side with more cues (Fig. 1A). To make the correct decision, the mouse must both integrate the number of cues and hold this information in working memory during the delay region. Mice are able to perform this task with good accuracy that increases with the magnitude of the difference between the number of cues to each side (Pinto et al., 2018). The activity of individual neurons has been recorded using two-photon calcium activity while mice performed this task in many brain regions (Koay et al., 2022; Nieh et al., 2021). These recordings show that, across both cortical and subcortical regions, neural activity is dominated by activity that is sequential rather than persistent. Such sequences are present throughout the task, beginning in the cue region and continuing through the delay region. Moreover, many of the neurons in these sequential firing patterns are choice-selective, such that neurons preferentially fire depending upon whether the animal will ultimately choose to turn left or right (Fig. 1B). For correct trials, this is equivalent to neurons showing evidence-selectivity. Previous models of evidence accumulation, based on circuit models that neurally instantiate the drift diffusion model, do not produce sequences when integrating evidence in this task (Fig. 1C), nor do bump-based integrators (Fig. 1D), commonly used to integrate velocity in the head direction system. Further, canonical sequence models such as synfire chains do not accumulate evidence (Aviel et al., 2002; Diesmann et al., 1999; Hayon et al., 2005) or accumulate external signals but are not choice-selective (Goldman, 2009; Lim & Goldman, 2012). Thus, no traditional models capture the choice-selective sequences in the accumulating towers and other tasks, requiring new models.

Chains-Based Models

Our first class of models are chains-based models, where chains of neurons accumulate evidence, encoded in the amplitude of the firing activity. In these classes of models, we consider two competing chains, where each chain represents evidence to one side (Fig. 2A). Here, we consider a model with mutual inhibition between neurons at the same position in the chains, but other chains-based models are possible (see Supplementary Information for a model with uncoupled, rather than mutually inhibiting chains). The mutually inhibitory chains-based model can be written in the form,

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} - er_{i,R} + f\mathbf{1}_{\text{left}}(t) + P_i(t) - T]^+ \quad (1)$$

where $r_{i,L}$ is the firing rate of the ith neuron in the left chain and $r_{i,R}$ is the firing rate of the ith neuron in the right chain. Progression along each chain is controlled by a position-gating signal, where only a single pair of neurons (one from each chain) has an active position-gating signal ($P_i(t)>T$, where T is the firing threshold) at any position in the maze (cyan gating signal, Fig. 2A). At the active position, a neuron receives excitatory input from itself (green self-excitation) and the previous neuron in the same chain (green feedforward connections), inhibitory input from the neuron at the same position in the opposing chain (red inhibitory connections), and external

cues from the ipsilateral side (purple and orange inputs), with input pulses of the form $\mathbf{1}_{left}(t)$, which is 1 if there is a left cue within 0.5 cm of the current position. The firing rate of neurons not at the active position decay exponentially with time constant *1/a*. See Methods for complete details and parameterization of the model.

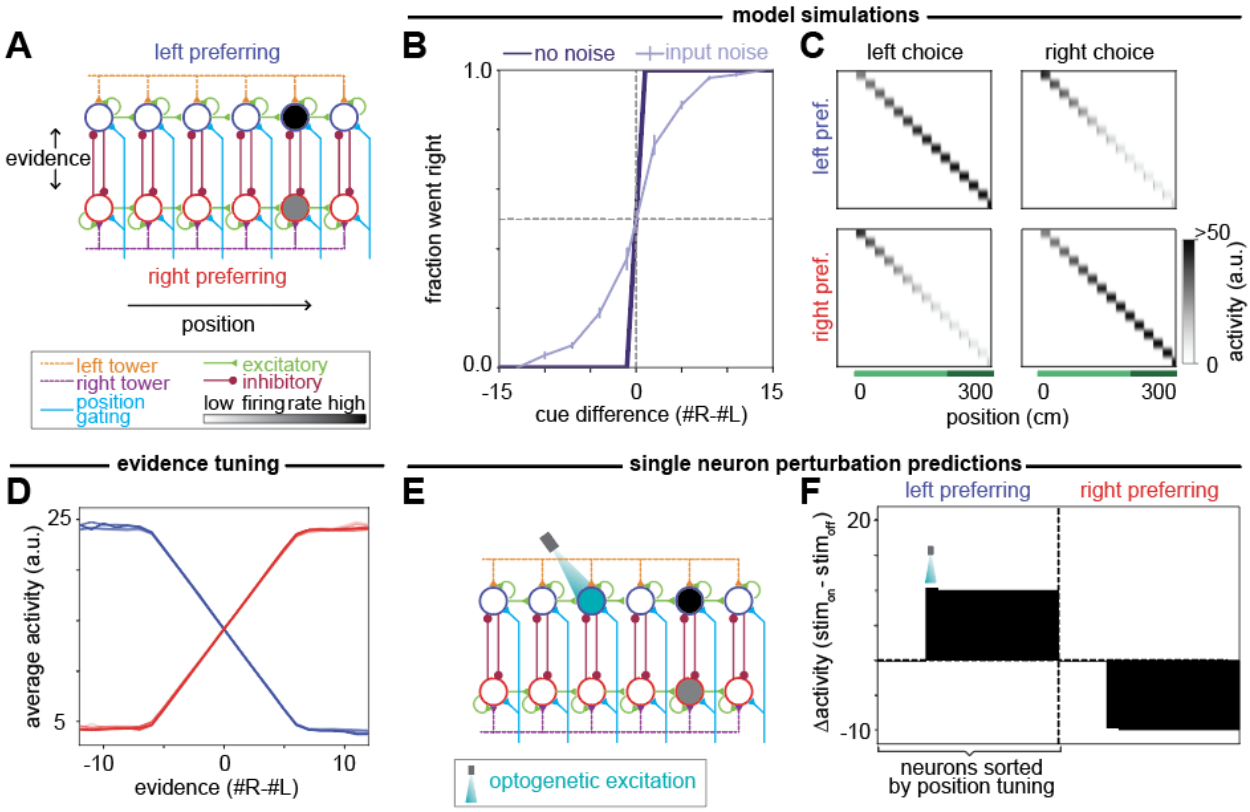

Figure 2. Mutually inhibitory competing chains model of evidence accumulation through sequences. (A) Schematic of neural circuit model for the mutually inhibitory competing chains model, showing excitatory (green) and inhibitory (red) connections between neurons (circles) as well as inputs to the circuit from the external left (orange) and right (purple) towers and a position gating signal (cyan). **(B)** Psychometric data for model simulated trials describing how often the amplitude of the final neuron in the left chain was greater than that of the final neuron in the right chain for cases where the model was simulated with (light purple) and without (dark purple) noise in the input. Error bars show s.e.m. **(C)** Each row shows the non-normalized amplitude of a model neuron at each position in the maze, averaged across simulated trials without input noise when the greater final amplitude was in the left (left choice, left column) or right (right choice, right column) chain. Neurons were divided based on their choice-selectivity (equivalent to the left and right chains in this model; see Methods) and ordered based on the position of peak activity (equivalent to position in the chain). **(D)** Tuning curves of individual neurons to evidence by finding the average activity for different evidence levels, calculated when the neurons are activated by the position-gating signal, for left-preferring (blue) and right-preferring (red) cells. **(E)** Schematic of a single neuron perturbation experiment where optogenetic excitation is applied to a single neuron in the left chain. **(F)** Simulated changes of the firing rates of all neurons in the absence of cues when a single cell (denoted with the laser as in (E)) is optogenetically stimulated.

Note that position should be interpreted broadly as a sequentially progressing state of the animal (e.g., it could represent time), so these models apply to navigation or non-navigation tasks. While not explicitly modeled here, such spatio-temporal gating could result from inputs from the basal ganglia (Monteiro et al., 2020; O'Reilly et al., 2008; Paton & Buonomano, 2018) or place cells in the hippocampus. Other models have also considered how time may be accumulated to provide a timing signal (Zemlianova et al., 2022). The models presented here could easily be expanded to include more explicit models of such spatio-temporal signals.

Together, these connections produce a neural circuit in which the location of the actively firing neurons along the chain corresponds to the position in the environment. Evidence is represented in the relative amplitude of firing of the two chains, with greater amplitude indicating greater accumulated evidence. We note that while a single chain could be used to accumulate evidence, this is not consistent with the data, in which different sets of neurons fire most actively dependent upon whether the accumulated evidence favors a left or right choice (Fig. 1B).

In this model, each chain only receives external input from one side, but the chains compete with each other through mutual inhibition between neurons at the same position within the chain (Fig. 2A). Thus, at a given position, this model has the same structure as the traditional mutual inhibition model (Usher & McClelland, 2001), and this motif is repeated to form two chains. Position gating causes the transitions between these repeated motifs, where only the neurons receiving the position gating signal actively represent and accumulate evidence. This competition between the chains results in the growth of activity in the dominant chain while activity in the other chain decreases, and the position gating results in the computation occurring in a sequence of cells.

We simulated the model for 1000 trials of the accumulating towers task (see Methods), demonstrating that this model performs the task perfectly in the absence of noise for appropriately tuned parameters (Fig. 2B). Previous analysis of the accumulating towers task and other evidence accumulation tasks suggest that noise in the animals' performance is dominated by sensory noise, so that only a fraction of input cues are integrated (Brunton et al., 2013; Pinto et al., 2018). When we correspondingly introduce input noise into our model (see Methods), the performance also decreases; as in the animals, performance increases with the magnitude of the difference between cues (Fig. 2B). The activity of the simulated neurons produces choice-selective, sequential activity, with neurons in each chain showing greater activity for the side of the evidence that it receives (Fig. 2C).

Such sequences appear qualitatively similar to those observed in the neural data, necessitating better ways to compare model predictions to neural data. One way of doing so is by analyzing the tuning curves of individual neurons. Since the position-gating hypothesized by our model makes it such that neurons are only active at some peak position, we only consider tuning to accumulated evidence levels present at these positions (see Methods). We find that neurons are monotonically tuned to left or right evidence dependent upon the chain to which the neuron belongs (Fig. 2D). Thus, evidence is encoded monotonically (up to saturation) in the relative firing rates of the two chains in this class of models.

We can also use our model to make predictions for the effects of single neuron perturbations, which could be probed experimentally by targeted optogenetic techniques. To do so, we simulate the delivery of an excitatory input to a neuron (or similarly tuned neuronal population) while the animal is at the neuron's active position (or equivalently

sub-position-threshold excitation is delivered to an individual neuron throughout the trial) for a trial with no external cues so that the accumulated evidence prior to perturbation is zero (Fig. 2E). Due to the gating of neural activity by the forward traveling position signal, the model predicts that only cells tuned to later positions will be affected by this stimulation. Subsequent neurons in the same chain will have increased activity due to this feedforward excitation while subsequent neurons in the opposing chain will have decreased activity due to the mutual inhibition between the chains (Fig. 2F).

We note that other architectures of chains-based models are possible, including a model where competition between opposing inputs occurs without inhibition between the chains (Fig. S1A). At a given position, this model has the same structure as traditional models based on a 'push-pull' feedforward arrangement of evidence inputs with cues favoring one side exciting the ipsilateral chain and inhibiting the contralateral chain (Ditterich et al., 2003), and this motif is repeated to form two chains. Both models are able to perform the task accurately (Fig. 2B, Fig. S1B), produce choice-selective sequences (Fig. 2C, Fig. S1C), and have neurons with monotonic tuning curves to evidence (Fig. 2D, Fig. S1D). We chose to focus on the mutual inhibition model here because it is the most robust to the tuning of the position gating. Specifically, this model requires only that P_i be greater than or equal to T at the active position, whereas the uncoupled chains model requires that P_i is exactly equal to T to avoid continual growth of neural activity from temporally integrating the position-gating signal (Supplementary Information; (Cannon et al., 1983)). However, such models instantiate different hypotheses that could be differentiated by the responses of different cells to optogenetic perturbations (Fig. 2E, Fig. S1E). Specifically, when the anti-correlated activity of the two chains is solely due to opposing external cue inputs, but the chains are otherwise uncoupled (Fig. S1), an excitatory perturbation in one chain only increases the firing within the same chain but has no impact on the other chain (Fig. S1F).

Taken together, these chains-based models represent a class of models in which position-gating gives rise to choice-selective sequences that encode evidence monotonically in the neuronal firing rates at each position. Such monotonic responses provide a simple prediction of such models that can be easily tested experimentally.

A Position-Gated Bump Attractor

Our second class of models is a position-gated bump attractor that represents accumulated evidence and position jointly in the location of a stereotypical, unimodal pattern ("bump") of neural activity in the population. Bump attractor models have had much success in describing the neural circuitry that computes head direction in the rodent (Ajabi et al., 2023; Clark & Taube, 2012; Redish et al., 1996; Skaggs et al., 1995; Song & Wang, 2005; Zhang, 1996) and the fly (Cope et al., 2017; Green et al., 2017; Hulse & Jayaraman, 2020; I. S. Kim & Dickinson, 2017; S. S. Kim et al., 2019; D. Turner-Evans et al., 2017; D. B. Turner-Evans et al., 2021), where the location of the bump corresponds to the animal's heading. Similarly, two-dimensional bump attractor models have been proposed to describe path integration in the hippocampal place cell system (Conklin & Eliasmith, 2005; McNaughton et al., 2006; Samsonovich & McNaughton, 1997). In these models, motion is temporally integrated to determine head direction or to perform path integration, and analogously our models temporally integrate evidence from the input visual cues. A single dimensional bump model that only

accumulates evidence does not produce sequences like those observed in the neural data as the bump would recur at multiple different positions whenever the accumulated evidence is at the same level and would remain in a fixed position throughout the delay period since the amount of evidence does not change (Fig. 1E). Thus, we modify these traditional bump attractor models to be planar-based models with separate dimensions encoding position and evidence (Fig. 3).

The activity $r_{i,j}$ of a neuron at position *i* and evidence level *j* evolves according to

$$\frac{dr_{i,j}}{dt} = -ar_{i,j} + \Theta(P_i(t) - T)\left[F\left(\sum_k W_{j,k} r_{i,k} + br_{i-1,j} + c\left(I_{j-1,L}(t) + I_{j+1,R}(t)\right)\right)\right] \quad (2)$$

where

$$I_{j,L} = r_{i,j}\mathbf{1}_{left}(t) \quad (3)$$
$$I_{j,R} = r_{i,j}\mathbf{1}_{right}(t) \quad (4).$$

Here, position gating (cyan inputs, Fig. 3A) is controlled by the Heaviside function $\Theta(P_i(t) - T)$, with activity decaying exponentially away from the active position with time constant $1/a$, and the synaptic connections are as shown in Figure 3A. In the active position, the neuron receives excitatory input from itself and its immediate neighbors (green excitatory connections between neurons at the active position), as well as the neuron at the previous position of the same evidence level (green feedforward connections), and receives inhibitory input from all other cells at that position (red inhibitory connections). The neuron receives external input mediated by neurons (Fig. 3A, small orange ($I_{j,L}$) and purple ($I_{j,R}$) circles) that are only active when coincidently receiving input from the corresponding evidence accumulation neuron and an external cue ($\mathbf{1}_{left}(t)$ for left cues and $\mathbf{1}_{right}(t)$ for right cues, which equal 1 if there is a left or right cue respectively within 0.5 cm of the current position). The multiplication in equations (3) and (4) enforces that the global external input signal only activates the neurons close to the location of the current bump, allowing the global cue signal to move the localized bump of activity (Redish et al., 1996; Song & Wang, 2005). For the complete model specification, see Methods.

As in the chains-based models, activity along the position axis is gated so that only a single layer of the model receives an active position signal at a time, and away from the active position, activity decays exponentially. As the animal progresses forward in the maze, the active position changes and the next layer of cells is activated. Feedforward excitation between the layers transfers information to neurons at the next position.

Within the active position, accumulated evidence is represented by the location of the bump along the evidence axis. As in traditional bump attractor models (Amari, 1977; Ben-Yishai et al., 1995), the bump is maintained in the absence of input by local excitation while long range inhibition prevents the bump from diffusing. Asymmetric input connections cause the bump to shift left or right in response to a left or right cue respectively (Redish et al., 1996; Song & Wang, 2005).

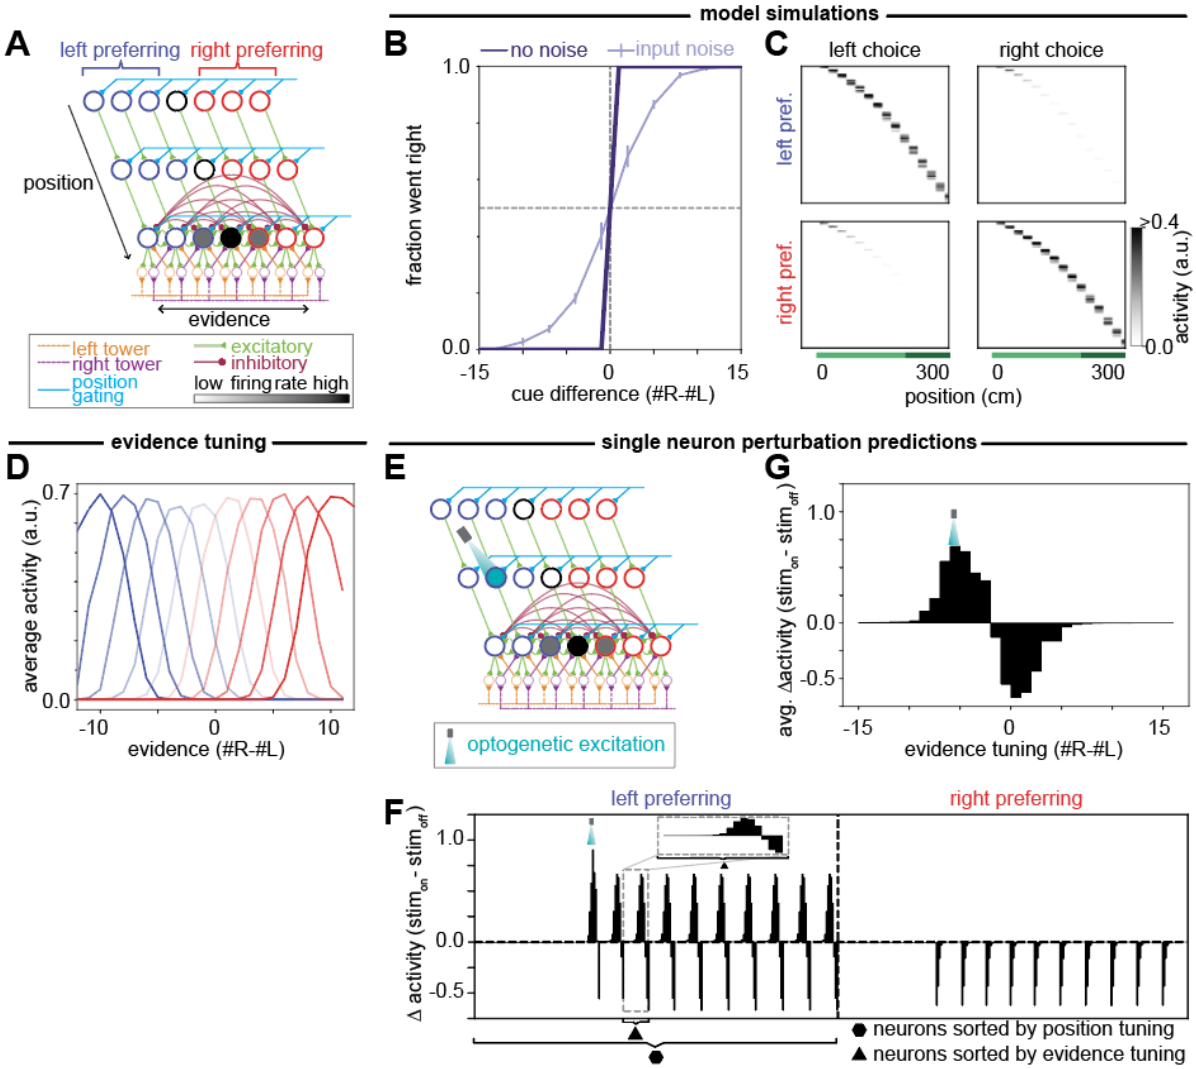

Figure 3. Position-gated bump attractor model of accumulation of evidence through sequences. **(A)** Schematic of neural circuit model for the position-gated bump attractor model, showing excitatory (green) and inhibitory (red) connections between neurons (circles) as well as inputs to the circuit from the external left (orange) and right (purple) towers and a position gating signal (cyan). Note that excitatory and inhibitory connections within a given position are only shown for the bottom row, but these same connections are present at every position. **(B)** Psychometric data for simulated trials, indicating how often the index of the maximally firing active neuron in the last set of positions was to the left of the center neuron for cases where the model was simulated with (light purple) and without (dark purple) noise in the input. Error bars show s.e.m. **(C)** Each row shows the non-normalized amplitude of a model neuron at each position in the maze, averaged across simulated trials without input noise when the greater final amplitude was in the left (left choice, left column) or right (right choice, right column) chain. Neurons were divided based on their choice-selectivity (see Methods) and ordered based on the position of peak activity. **(D)** Tuning curves of a subset of individual neurons to evidence, calculated when the neurons are activated by the position-gating signal, for left-preferring (blue) and right-preferring (red) cells. **(E)** Schematic of a single neuron perturbation experiment where optogenetic excitation is applied to a single neuron in the left preferring population. **(F)** Simulated changes of the firing rates of all neurons in the absence of cues when a single cell (denoted with the laser as in E) is optogenetically stimulated. **(G)** Average change in firing rate for neurons tuned to different evidence levels.

We simulate this model for the same set of 1000 trials as the chains-based models. As in that class of models, for appropriately tuned parameters, the position-gated bump attractor also performs the task perfectly in the absence of noise, and performance worsens when input noise is introduced (Fig. 3B). The simulated neural activity also produces choice-selective sequences (Fig. 3C). As seen by the downward curvature of the sequences in Figure 3C, the sequences are marked by the feature that more neurons respond at late positions than respond at early positions in the sequence. This is because the random walk of the bump location as evidence is accumulated leads to a greater range of evidence levels being observed at later positions in the trial; thus, on average across trials, there are more different neurons active at later positions, with each neuron corresponding to a different possible evidence level.

The most straightforward difference between the position-gated bump attractor and the chains-based models is in the shape of their tuning curves. Due to the structure of the network connectivity, the neurons in the position-gated bump attractor are non-monotonically tuned to specific evidence levels at their active position (Fig. 3D), in contrast to the monotonic tuning curves of the chains-based models. To further differentiate the position-gated bump attractor model from the chains-based models, we again make predictions for the effects of optogenetic excitation of a single cell during a trial with no cues presented (Fig. 3E). Here, we consider the regime where the stimulus is sufficiently large to cause the bump to move to the stimulated location. Due to the similar feedforward connectivity patterns between positions in both model classes, stimulation of a cell active at one position only affects the activity of cells tuned to later positions, but the other predictions are markedly different. Unlike the chains-based models, we see that rather than impacting every neuron at a later position equally, the effects are sparse and differ in magnitude across neurons (Fig. 3F). These differences are because the stimulation has caused the bump to move to a new evidence level, which is then sustained at subsequent positions. Thus, neurons with a similar evidence tuning to the cell that was stimulated increase in activity while those tuned to where the bump of activity would be located in the absence of stimulation (at 0 evidence) decrease their activity since the bump has moved to the new location (Fig. 3G). Thus, in contrast to the chains-based models for which the changes in responses following stimulation within the same choice-preferring population as the stimulated cell are unidirectional, the changes for the position-gated bump attractor are bidirectional. We note that for smaller excitatory stimuli that do not shift the bump center to the site of stimulation, and dependent upon the exact details of excitatory and inhibitory connections, smaller shifts in the bump either toward or away from the stimulated location are possible (Krishnan et al., 2018; Ma et al., 2022; Qi & Gong, 2015), but the perturbations still maintain the properties that changes are sparse and that changes within the same choice-preferring population as the stimulated neuron may be bidirectional.

Overall, the position-gated bump attractor presents a contrasting evidence encoding scheme to the competing-chains models. While both models are non-monotonically tuned to position due to the position-gating, the position-gated bump attractor encodes evidence non-monotonically while still giving rise to choice-selective sequences. These contrasting evidence encoding schemes provide an important test that may differentiate which brain regions may employ which of these computations during evidence accumulation.

DISCUSSION

The work presented here represents the first neural circuit models that accumulate evidence through sequences rather than through persistent activity, thus providing a computational framework for the mechanistic understanding of experimentally observed choice-selective sequences. Our two model classes suggest two different ways that evidence could be accumulated through choice-selective sequences, in which evidence is encoded either in the set of neurons that fire, with the amplitude of firing being stereotyped, or in the amplitude of neuronal firing rates. Both individual neuron tuning to evidence (Figs. 2D, 3D, S1D) and single neuron perturbations (Figs. 2F, 3F, S1F) could be used to distinguish between the models in neural data.

Analysis of single neuron tuning curves could suggest how different brain regions contribute differently to the decision making process because of different representations of accumulated evidence. Previous work has analyzed the relationship between tuning curves and behavioral outcomes to suggest bump attractor dynamics in prefrontal cortex during a working memory task in monkeys (Wimmer et al., 2014), suggesting the utility of analyzing tuning curves to understand underlying dynamics. Other work has speculated that the form of the tuning curve may be related to the function of the region (Salinas, 2006). The non-monotonic tuning curves of the position-gated bump attractor allow for a direct readout of the evidence based on the location of the active bump. This form of coding is more computationally intensive in the number of neurons required but provides more refined information about evidence for non-binary readouts of decision-related information, as would be needed for building a cognitive map of evidence (Kumaran & Maguire, 2005; MacDonald et al., 2011; O'Keefe & Nadel, 1978). Interestingly, previous recordings of hippocampal firing during the accumulating towers task have suggested that the tuning curves of neurons in that region are tuned to specific position and evidence values (Nieh et al., 2021), aligning with the predictions of the position-gated bump attractor model and suggesting a joint cognitive map for position and evidence. In contrast, the monotonic encoding of evidence in the chains-based models does not allow a direct readout of evidence from the identity of the active neurons, but may be ideally suited for encoding graded quantities important in decision-making, such as confidence, value, or predicted reward. Future work should test whether regions associated with these quantities encode evidence monotonically, as suggested by our chains-based models.

While the shapes of the single neuron tuning curves may be related to function and differentiate the bump attractor models from the chains-based models, they do not differentiate between the varying chains-based models. We demonstrated that targeted perturbations of specific subsets of neurons selected based on their tuning to evidence may be able to make such distinctions between different subclasses of chains-based models (Figs. 2F, S1F). While experiments have shown that optogenetic manipulations of large, undifferentiated populations of neurons does affect evidence accumulation in the accumulating towers task (Bolkan et al., 2022), more selective perturbation experiments that target neurons with specific coding properties remain to be tested experimentally. Perturbations in other neural integrator systems, where the opposing populations are anatomically separated, have successfully tested model predictions about the dynamics of the integrator (Aksay et al., 2007; Fisher et al., 2013; Gonçalves et al., 2014). In the absence of anatomical separation of neurons with different tuning curves, such perturbations will require cellular resolution stimulation, an approach that has

shown rapid development (Adesnik & Abdeladim, 2021; Pégard et al., 2017; Rickgauer et al., 2014). Targeted perturbation experiments across different brain regions could additionally provide insight into whether the observed evidence tuning in different regions results from evidence accumulation itself or a downstream readout of evidence accumulated in another region (Boyd-Meredith et al., 2022; Seung, 1996). It is also possible that evidence is encoded in a traditional, non-sequential manner at some location and then combined at a downstream readout with a sequential representation of position to give the observed neural activity. If so, then a possible experimental prediction would be that inactivating choice-selective, evidence-encoding neurons would not disrupt evidence accumulation.

In addition to looking for evidence of these models in biological circuits, the motifs and predictions of these models could also be tested for in trained recurrent neural network models, suggesting how less explicitly structured models may perform the accumulating towers task (Langdon & Engel, 2022). A similar approach has been taken in the study of the path integration circuitry of the entorhinal cortex, where core foundational motifs of connectivity of the style presented here were revealed within the complex circuitry of the trained network (Sorscher et al., 2023). Such studies also may reveal conditions favoring the emergence of evidence accumulation through sequences as opposed to traditional persistent activity, as has been analyzed for other cognitive operations (Orhan & Ma, 2019).

At the essence of accumulating evidence through sequences is the ability to both integrate evidence and transfer information between positions. While models for evidence integration have been explored extensively and form the core of our models' computations within a position, traditional models do not successfully transfer information between positions, requiring the addition of a gating mechanism. Our two models utilize two different gating mechanisms, an additive threshold in the chains-based models and a multiplicative threshold in the position-gated bump attractor, demonstrating two mechanisms for the way information flow between positions can be controlled. The ability to transfer accumulated evidence between positions is not unique to the task studied here. For example, perceptual decision making tasks may require transferring accumulated evidence between different neuronal populations when the neuronal population that needs this information to carry out an action changes (So & Shadlen, 2022). More generally, the ability to transfer and control the flow of information within and between brain regions is a fundamental cognitive operation, so the mechanisms presented here may apply across many tasks involving the transfer of analog-valued information.

In summary, our models present a new set of neural architectures for the accumulation of evidence. The models make distinguishing predictions at the single neuron level and suggest further experiments to probe the differences between the many brain regions where such choice-selective sequences are observed. Such analyses may suggest different underlying computations in different brain regions, helping to suggest the potential contribution of different brain regions to the decision making process and allowing us to better understand how decision making is coordinated across the brain.

METHODS

The Accumulating Towers Task

The accumulating towers task is an accumulation of evidence task in virtual reality, previously validated to study perceptual decision making. The complete details of the task and

experimental setup can be found in (Pinto et al., 2018), and we only briefly summarize the key features of the behavioral task here. In this task, head-fixed mice run on an 8-inch styrofoam ball with the virtual environment, generated by ViREMen (Aronov & Tank, 2014) projected onto a surrounding screen. The virtual environment has a T-maze structure. In the first 30 cm (the precue region), there are no visual cues. In the next 200 cm (cue region), visual cues appear as towers on both sides, with the total count of towers on each side determined by a Poisson distribution of different means. Towers are positioned randomly with an enforced minimal distance between them. A tower appears when the mouse is within 10 cm of its position and disappears 200 ms later. Following the cue region, there is a 100 cm delay region during which the mouse must remember which side was presented with more towers. At the end of the delay region, the mouse receives a sweetened liquid reward for turning to the side with more cues. Water restriction and shaping protocols are also described in (Pinto et al., 2018). Recording sessions were approximately 1.0-1.5 h in duration, consisting of approximately 150-250 trials per session.

In addition to the accumulation of evidence trials described above, sessions also included some warm-up maze trials (cues on one side only, no distractors, tall visual guides at the end of the maze) at the beginning of the session and interspersed in the session if the animal's performance dropped below some threshold, as described in (Pinto et al., 2018). Since these trials do not require evidence accumulation, we did not include them in our analyses.

Neural Recordings

Animal procedures were conducted in accordance with standards from the National Institutes of Health and under the approval of the Princeton University Institutional Animal Care and Use Committee.

The head-fixed setup of the accumulating towers task allows for neural imaging throughout the task. Our neural datasets consist of calcium fluorescence data previously recorded from primary visual cortex (V1) and retrosplenial cortex (RSC) (Koay et al., 2022), and hippocampus (HPC) (Nieh et al., 2021), as well as previously unpublished recordings from anterior cingulate cortex (ACC) and dorsomedial striatum (DMS) as follows: *V1:* Two-photon fluorescence imaging from mouse lines that express GCaMP6f (4 animals, 17 sessions, 2755 total neurons). *RSC:* Two-photon fluorescence imaging from mouse lines that express GCaMP6f (8 animals, 41 sessions, 8579 total neurons). For complete details of V1 and RSC recordings, see (Koay et al., 2022). *HPC:* Two-photon fluorescence imaging from mouse lines that express GCaMP6f (7 animals, 7 sessions, 3144 total neurons). For complete details, see (Nieh et al., 2021). *ACC:* Two-photon fluorescence imaging from mouse lines that express GCaMP6s (3 animals, 7 sessions, 1720 total neurons). For complete details, see below on ACC recordings. *DMS:* Two-photon fluorescence imaging from mouse lines that express GCaMP6f (8 animals, 11 sessions, 804 total neurons). For complete details, see below on DMS recordings.

Animals and Surgery for ACC Recordings

Three female mice of CaMKIIa-tTA (JAX 007004) crossed with tetO-GCaMP6s (JAX 024742) were used for two-photon imaging of ACC. These crossed lines show robust and stable expression of GCaMP6s in a subset of CaMKIIa-positive, glutamatergic neurons in the cortex. All mice were between 12-16 weeks when they underwent sterile stereotaxic surgery under

isoflurane anesthesia (4% for induction, 1-1.5% for maintenance). The skull was exposed by making a vertical cut with a scalpel. Two craniotomy holes were made, one over ACC (anteroposterior axis: 0.8 mm, mediolateral axis: -0.5 mm, relative to bregma, diameter of ~1.2 mm) and another over DMS (anteroposterior axis: 0.75 mm, mediolateral axis: -1.1 mm, relative to bregma, diameter of ~0.3 mm). Adeno-associated virus (AAV; AAVretro2 serotype) encoding mCherry under CaMKIIa promoter (diluted to 2.5 x 10^{12} genome copies/mL, packaged at Princeton Vector Core) was injected to the DMS (anteroposterior axis: 0.75 mm, mediolateral axis: -1.1 mm, dorsoventral axis: -3.1 and -2.8 mm, relative to bregma) to provide information about a projection target and to be used for better motion correction. A total of 300 nL of AAV was infused at each site along the dorsoventral axis, at a rate of 30 nL/min. After injection, the needle was held in the same place for an additional 10 minutes. The needle was slowly withdrawn over 10 minutes to prevent backflow. After injections, the dura below the craniotomy was was carefully removed within the ACC craniotomy site, and the cortical tissue above ACC was carefully aspirated up to 300 μm below the surface. Then we implanted a 1.0-mm diameter GRIN lens (3.4 mm length, Inscopix, 1050-002184) into the ACC, down to 1.5 mm in the dorsoventral axis (relative to bregma), using a 3D-printed custom lens holder. After implant, a small amount of diluted metabond (Parkell) was applied to affix the lens to the skull using a 1-mL syringe and 25-gauge needle. After 10 minutes, the lens holder grip was loosened while the lens was observed through the stereoscope to make certain that there was no movement of the lens. After applying a more diluted metabond around the lens, a titanium headplate was positioned over the skull using a custom tool and aligned parallel to the stereotaxic frame using an angle meter. The headplate was affixed to the skull using a metabond. Finally, a titanium ring to hold the immersion medium was placed on the headplate, with lens in the middle, and then affixed to the headplate using dental cement blackened with carbon. After the dental cement was fully cured, mice were unmounted from the stereotaxic frame and protective head caps were placed. Their recovery was monitored for ~2 hours. After one or two weeks of recovery, mice started water-restriction, and after 5 days behavioral training for the accumulating towers task began, as described in (Pinto et al., 2018).

Animals and Surgery for DMS Recordings

D1R-Cre (n = 3, female, EY262Gsat, MMRRC-UCD), A2a-Cre (n = 3, male, KG139Gsat, MMRRC-UCD), and D2R-Cre (n = 2, male, ER44Gsat, MMRRC-UCD) mice were used for DMS two-photon imaging. Surgical procedures were identical to those described above for the ACC save for the following details. Approximately one week prior to GRIN lens implantation, a Cre-dependent AAV expressing GCaMP6f under either the CAG promoter (AAV5-CAG-DIO-GCaMP6f, Addgene, diluted to 3.5 x 10^{12} genome copies/mL) or synapsin promoter (AAV1-Syn-DIO-GCaMP6f, diluted to 3.8 x 10^{12} genome copies/mL) was delivered (500ul per injection site; 200 ul/min) to the DMS (anteroposterior axis: 0.7 mm, mediolateral axis: 1.15 and 1.85 mm, dorsoventral axis: -3 mm, relative to bregma) and the surgical incision was closed using monofilament non-absorbable sutures (Johnson and Johnson). This allowed for expression of GCaMP6f in D1R- or D2R-expressing medium spiny neurons according to the transgenic mouse line used, and facilitated the settling of brain tissue prior to GRIN lens implantation. Following recovery (5-10 days), sutures were removed and the viral injection craniotomy was expanded and dura removed to accommodate a 1.0-mm diameter GRIN lens

with a prism (4.3 mm length, Inscopix, 1050-004601). The lens was implanted with prism window facing posterior and just anterior to viral injection sites (anteroposterior axis: 0.75 mm, mediolateral axis: 1.35 mm, dorsoventral axis: -3.25 mm, relative to bregma). Fixation of the lens to the skull, as well as the custom-printed headplate, water immersion ring, and protective covering were identical to that described above for ACC. Across all DMS targeted GCaMP6f imaging mice we obtained three distinct fields of view from one mouse, two distinct fields of view from a second mouse, and the remaining mice contributed single fields of view (8 animals, 11 sessions, 804 total neurons).

Optical Imaging and Data Acquisition for ACC and DMS Recordings

Imaging was performed using a custom-built two-photon microscope compatible with virtual reality (VR). The microscope was equipped with two high-power femtosecond lasers (Alcor), with fixed wavelengths for 920 and 1064 nm. The scanning units used a 5-mm galvanometer and a 9kHz resonant scanning mirror (Cambridge Technologies). The collected photons were split into two channels by a dichroic mirror (FF562-Di03, Semrock). The light for respective green and red channels was filtered using a bandpass filter (FF01-520/60 and FF02-607/70, Semrock), and then detected using GaAsP photomultiplier tubes (1077PA-40, Hamamatsu). The signal from the photomultiplier tubes was amplified using a high-speed current amplifier (59-179, Edmund). Black rubber tuning was attached around the objective (Zeiss, W N-Achroplan 20x/0.5 M27, 20x magnification, 0.5 NA) to shield the light from the virtual reality screen from the space from the objective to the titanium headring surrounding the implanted GRIN lens. Double-distilled water was used as the immersion medium inside the headring. The amount of laser power at the objective was ~30 mW for ACC and DMS. Control of the microscope and imaging acquisition was performed with the ScanImage software (Vidrio Technologies), on a separate computer from the VR-running one. Images were acquired at 30 Hz at a resolution of 512 x 512 pixels. Synchronization between behavioral logs and acquired images was achieved by sending behavioral information at each time the ViRMEn environment was refreshed from the VR controlling computer to the scanning computer via an I2C serial bus: these information about sessions, trials, and frames were stored in the header of TIFF image files.

Data Processing for ACC and DMS Recordings

Both ACC and DMS imaging data were processed with Suite2p. For ACC recordings in which mCherry signals are present, rigid motion correction was performed with the red channel and then applied to the functional green channel with GCaMP6s, with the following parameter settings: nimg_init = 1000, smooth_sigma = 1.15, maxregshift = 0.1. For DMS recordings, the same motion correction was performed, but with the green channel. Region-of-interests (ROI) corresponding to individual neurons were detected, and calcium signals were extracted and neuropil-corrected (neucoeff = 0.7). Subsequently, an experimenter-mediated curation was performed on these ROIs to discriminate putative neurons from noise. Neuropil-corrected fluorescence from these curated neurons are exported to MATLAB, and the $\Delta F/F$ was calculated, taking baseline fluorescence as the mode of the fluorescence within 3 minute long windows (the same baseline as applied in the V1, RSC, and HPC datasets).

Spike Inference

In order to compare neural recordings across datasets with different calcium indicators, we performed spike inference on our data based on a recently developed and validated algorithm (Fleming et al., 2021; Jewell et al., 2020). For each ΔF/F dataset, we assumed a baseline firing rate of 2 Hz and set the calcium decay rate corresponding to the calcium indicator (0.9704588 for GCaMP6f, 0.9885 for GCaMP6s). From these inferred spikes, we inferred a continuous firing rate using a gaussian window function (window length = 1 s, standard deviation = 0.25 s).

Position Binning of Neural and Evidence Data

To compare neural data along maze positions and evidence levels, we resampled the neural (continuous firing rate as calculated above) and behavioral (cumulative evidence up to that point) data in position bins. Specifically, firing rates were averaged based on the position of the animal within 5 cm bins from the start of the trial (-30 cm) to the end of the delay region (300 cm) for each trial. Cumulative evidence for each bin was defined by the number of right minus the number of left cues that had been observed before the start of the position bin.

Trial Selection Criteria

We discarded a few trials from analysis using the following criteria. As mentioned above, any trials from the warm-up maze were excluded, as these trials do not require evidence accumulation. Additionally trials with excessive travel were discarded, which were identified by finding trials where mice showed a large view angle deviation (more than 90 degrees, either left or right) before the end of the delay region.

Determining Choice Selectivity of Neurons

To identify choice-selective neurons (Fig. 1B), we first jointly fit position and evidence tuning of each neuron in our dataset as described in Data Analysis below. We then considered only cells that had significant evidence tuning to be choice-selective. To produce the sequence plots, cells were first sorted as right preferring if the mean of the neuron's evidence tuning was greater than 0 and left preferring otherwise and then sorted by their position mean. Using this sorting, we then took the average activity of each neuron at different positions for correct left and correct right trials. The resulting traces for each neuron were then peak normalized.

Parameterization of the Mutually Inhibitory Chains-Based Models

As in equation (1), the activity of the neuron at position *i* in the left chain can be described as

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} - er_{i,R} + f\mathbf{1}_{\text{left}}(t) + P_i(t) - T]^+$$

A neuron's active position is defined by a square pulse over a range of positions, where *p(t)* is the position of the animal at time *t,*

$$P_i(t) = \begin{cases} T + I_{ext} & i \leq \frac{p(t)}{P_0} < (i+1) \\ 0 & \text{otherwise} \end{cases}$$

For our simulations, we take $P_0 = 20$, which divides the total length of the maze (330 cm) into 17 different sections of active positions. Therefore, in these models we simulate 17 neurons in each chain for a total of 34 neurons.

The external input is of the form

$$\mathbf{1}_{\text{left}}(t) = \begin{cases} 1 & \min_{\ell \in L} |t - \ell| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathbf{1}_{\text{right}}(t) = \begin{cases} 1 & \min_{R \in \mathcal{R}} |t - R| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

where L is the set of times of left cues and $\mathcal{R}$ is the set of times of right cues.

In our simulations, we use parameters a =1, b = 0.2, c = 1, e =0.8, f = 4, and I_{ext} = 20.

Parameterization of the Position-Gated Bump Attractor

In Equation (2), we defined the evolution of the firing rate for a neuron at position *i* and evidence *j* as

$$\frac{dr_{i,j}}{dt} = -ar_{i,j} + \Theta(P_i(t) - T)\left[F\left(\sum_k W_{j,k} r_{i,k} + br_{i-1,j} + c\left(I_{j-1,L}(t) + I_{j+1,R}(t)\right)\right)\right].$$

The Heaviside function that acts as a multiplicative position gate is defined as

$$\Theta(x) = \begin{cases} 1 & x \geq 0 \\ 0 & x < 0 \end{cases}$$

so that the term in square brackets only affects the activity of the neuron at the active position. A neuron's active position is defined by a square pulse over a range of positions, where *p(t)* is the position of the animal at time *t,*

$$P_i(t) = \begin{cases} T & i \leq \frac{p(t)}{P_0} < (i+1) \\ 0 & \text{otherwise} \end{cases}.$$

The dynamics of the bump attractor have a nonlinearity

$$F(x) = \frac{1 + \tanh(x)}{2},$$

and synaptic connections between neurons are symmetric with excitatory connections between neurons close in evidence and inhibitory connections between neurons farther in evidence. We do this by assigning each evidence level an angle, evenly spaced between 0 and 2π, allowing us to define these sculpted synaptic connections as

$$W_{j,k} = \omega_0(\cos(\theta_j - \theta_k) + \omega_1).$$

We note that, since evidence is not periodic, a sufficiently large number of evidence levels were used such that substantial interactions do not occur between the nominally periodically connected endpoints.

When an input is present, asymmetric connections cause the active evidence neurons to shift towards the evidence. This is modeled through input nodes (Fig. 3A, purple and orange

small circles) that are activated by the coincidence of external cue signals and neural activity from the current location of the bump by

$$I_{j,L}(t) = r_{i,j}\mathbf{1}_{\text{left}}(t)$$

and

$$I_{j,R}(t) = r_{i,j}\mathbf{1}_{\text{right}}(t)$$

where

$$\mathbf{1}_{\text{left}}(t) = \begin{cases} 1 & \min_{\ell \in L} |t - \ell| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathbf{1}_{\text{right}}(t) = \begin{cases} 1 & \min_{R \in \mathcal{R}} |t - R| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

where L is the set of times of left cues and $\mathcal{R}$ is the set of times of right cues.

As in the chains-based models, the position-gating of each layer of neurons is activated over a length 20 cm (P_0 = 20), again giving us 17 different active positions. For our neurons to span the range of evidence of our experiments, we use 35 neurons in the evidence dimension, corresponding to evidence levels of -17 to 17. Thus we simulate a total of 595 neurons.

The parameters used in the simulation are: a = 1.1, b = 2.2, T = 300, ω_0 = 3, ω_1 = -0.9, and c = 5. Neurons $r_{0,-1}$, $r_{0,0}$, and $r_{0,1}$ are initialized to 0.6 activity level, and all other neurons are initialized to 0 activity.

Simulation of Differential Equations

Differential equations for each model (equations (1) and (2) and (S2)) were simulated in Python using scipy.odeint, which uses the lsoda algorithm, with a maximum stepsize of 1 ms (or equivalently 0.5 cm when assuming a constant velocity of 50 cm/s) to find the activity levels of the neurons in each model for points between -30 and 300 cm in 0.1 cm increments. In all simulations, we assume the animal travels at a constant velocity of 50 cm/s through the maze so that time and position are interchangeable. We simulated the neural activity for 1000 trials with cue positions derived from 1000 experimental trials from the behavioral and neural recording data. For the chains-based models, we say that the animal makes a left choice if the activity of the last neuron in the left chain is greater than the activity of the last neuron in the right chain. For the bump model, we say that the animal makes a left choice if the index k of the neuron with maximal firing in the last layer is less than 0.5J, where J is the total number of neurons in the layer. If the maximal firing occurs exactly at 0.5J, we determine the animal's decision by randomly choosing between the two alternatives with equal probability.

Simulations with Input Noise

For each model, we simulated 25 sessions with input noise. For each session, we selected 150 trials drawn randomly without replacement from the 1000 experimental trials that we used for simulation of the model. Let L and R be the set of left cues and right cues, respectively, in the original trial. For each cue x in L, we drew z, a random uniform number between 0 and 1, and if z<0.33, we add x to the set L'. We repeated the same process for each cue in R, yielding two reduced sets of cues when input noise is present, L' and R',

corresponding to the animal ignoring or not perceiving 0.67 of the cues. We then simulate the model with the input terms $\mathbf{1}_{left}(t)$ and $\mathbf{1}_{right}(t)$ determined according to the reduced cue sets. We then calculate the psychometric curves (as described in Psychometric Curves below) with the difference in cues based on the original cue sets, but whether the animal turns left is based on the simulation of the trial with the reduced cue sets.

Evidence Tuning Curves

In order to find the evidence tuning curves (Figs. 2D, 3D, S1D), we averaged the activity of each neuron binned by position and accumulated evidence for each of the 1000 simulated trials. Position bins were of size 0.1 cm, ranging from -30 to 300 cm. Evidence bins were in units of 1 tower, ranging from -15 to 15 towers. Bins that had less than 2 samples were not included in further analysis and not plotted. For a given neuron, we average across evidence levels to find the position with maximum firing. We then plot the average firing rate vs. evidence level at this position.

Single Cell Perturbations

Single cell perturbations were simulated in the same manner but in the absence of any cues. A subthreshold excitatory term $O_i(t)$ was added within the thresholded dynamics of each model to simulate the effects of stimulating an individual cell, where $O_i(t)$ = 1 if neuron i is the neuron being stimulated and 0 otherwise.

Psychometric Curves

For each session, we sorted trials based on the difference in the final number of cues at the end of the maze. We binned these trial differences with bin size of 3 towers, ranging from -14 to 14 towers. Within each bin, we then found the fraction of trials in each bin for which a left choice was made (Fig. 2B, Fig. 3B, Fig. S1B).

Data Analysis

Determining Position and Evidence Selectivity

In our neural recording data, we sought to identify the set of cells whose firing encoded accumulated evidence. Due to the sequential nature of the neural activity, where neurons are only active at some positions, we found that robustly identifying evidence selectivity required that we jointly identify the position range over which a neuron is active and the evidence selectivity within this set of positions. To do so, we assumed the firing rate of each neuron had the form,

$$FR(p,e) = aP(p)E(e) + b$$

where $FR(p,e)$ is the firing rate of the neuron at position p and evidence e. This assumes that the evidence tuning is independent of the position tuning and vice versa. We assume that P and E can be well described by Gaussians,

$$P(p) = \exp\left(\frac{-(p-\mu_p)^2}{2\sigma_p^2}\right)$$

and

$$E(e) = \exp\left(\frac{-(e-\mu_e)^2}{2\sigma_e^2}\right).$$

We note that this specific form of evidence tuning was not used as a detailed descriptor of neuronal firing. Rather, we found that it was effective for its purpose of identifying evidence-correlated neurons and determining if the mean evidence selectivity favors left or right evidence (as used to sort neurons by evidence-selectivity in Fig. 1B). Here, the Gaussian was allowed to extend beyond the range of observed evidence values, so that it can accommodate both monotonic and non-monotonic tuning curves (i.e., if the mean evidence occurs at the extremes of the evidence range, the evidence tuning will be monotonic throughout the range).

For each neuron, we normalized the firing rate observations (binned as in Position Binning of Neural and Evidence Data) to the range [0, 1]. We then fit the parameters of the above model using an iterative fitting procedure, alternating fitting μ_p, σ_p, a, and b with fitting μ_e, σ_e, a, and b with a nonlinear fitting method (scipy curvefit, which employs the trust region reflective algorithm, with bounds $-50 \leq \mu_p \leq 350$, $0 \leq \sigma_p \leq 200$, $-20 \leq \mu_e \leq 20$, $0 \leq \sigma_e \leq 30$, $0 \leq a \leq 10$, and $0 \leq b \leq 1$).

We then tested for significance of the evidence tuning fit using the pseudosession method (Harris, 2021). For each session of firing rate and position data, we generated new evidence for each trial in the session by resampling trials from actual sessions, and applied the iterative fitting procedure to the new pseudosession evidence and position data. We repeated this process for 50 pseudosessions and calculated the correlation coefficient between the true firing rates and the model predictions for each session. We then used a one sample t test to compare the fit of the model in the true session to the fit of the null distribution from the pseudosession, and only considered sessions for which the fit was significantly better ($p<0.05$) than the null distribution for further analysis.

Data Availability

Code for the simulation of all models and the data analysis is available upon request.

Conflicts of Interest

The authors report no conflicts of interest.

REFERENCES

Adesnik, H., & Abdeladim, L. (2021). Probing neural codes with two-photon holographic optogenetics. *Nature Neuroscience*, *24*(10), 1356–1366.

Afraz, S.-R., Kiani, R., & Esteky, H. (2006). Microstimulation of inferotemporal cortex influences face categorization. *Nature*, *442*(7103), 692–695.

Ajabi, Z., Keinath, A. T., Wei, X.-X., & Brandon, M. P. (2023). Population dynamics of head-direction neurons during drift and reorientation. *Nature*, *615*(7954), 892–899.

Akhlaghpour, H., Wiskerke, J., Choi, J. Y., Taliaferro, J. P., Au, J., & Witten, I. B. (2016). Dissociated sequential activity and stimulus encoding in the dorsomedial striatum during spatial working memory. *eLife*, *5*. https://doi.org/10.7554/eLife.19507

Aksay, E., Olasagasti, I., Mensh, B. D., Baker, R., Goldman, M. S., & Tank, D. W. (2007). Functional dissection of circuitry in a neural integrator. *Nature Neuroscience*, *10*(4), 494–504.

Amari, S. (1977). Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological Cybernetics*, *27*(2), 77–87.
Aronov, D., & Tank, D. W. (2014). Engagement of neural circuits underlying 2D spatial navigation in a rodent virtual reality system. *Neuron*, *84*(2), 442–456.
Aviel, Y., Pavlov, E., Abeles, M., & Horn, D. (2002). Synfire chain in a balanced network. *Neurocomputing*, *44-46*, 285–292.
Ben-Yishai, R., Bar-Or, R. L., & Sompolinsky, H. (1995). Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *92*(9), 3844–3848.
Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765.
Bolkan, S. S., Stone, I. R., Pinto, L., Ashwood, Z. C., Iravedra Garcia, J. M., Herman, A. L., Singh, P., Bandi, A., Cox, J., Zimmerman, C. A., Cho, J. R., Engelhard, B., Pillow, J. W., & Witten, I. B. (2022). Opponent control of behavior by dorsomedial striatal pathways depends on task demands and internal state. *Nature Neuroscience*, *25*(3), 345–357.
Boyd-Meredith, J. T., Piet, A. T., Dennis, E. J., El Hady, A., & Brody, C. D. (2022). Stable choice coding in rat frontal orienting fields across model-predicted changes of mind. *Nature Communications*, *13*(1), 3235.
Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *12*(12), 4745–4765.
Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, *10*(6), 1157–1169.
Britten, K. H., & Van Wezel, R. J. A. (2002). Area MST and heading perception in macaque monkeys. *Cerebral Cortex* , *12*(7), 692–701.
Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95–98.
Cannon, S. C., Robinson, D. A., & Shamma, S. (1983). A proposed neural network for the integrator of the oculomotor system. *Biological Cybernetics*, *49*(2), 127–136.
Carpenter, R. H. S. (1988). Movements of the eyes, 2nd rev. & enlarged ed. *Nd Rev. & Enlarged Ed*, *2*, 593.
Chandrasekaran, C., Peixoto, D., Newsome, W. T., & Shenoy, K. V. (2017). Laminar differences in decision-related neural activity in dorsal premotor cortex. *Nature Communications*, *8*(1), 614.
Clark, B. J., & Taube, J. S. (2012). Vestibular and attractor network basis of the head direction cell signal in subcortical circuits. *Frontiers in Neural Circuits*, *6*, 7.
Conklin, J., & Eliasmith, C. (2005). A controlled attractor network model of path integration in the rat. *Journal of Computational Neuroscience*, *18*(2), 183–203.
Cook, E. P., & Maunsell, J. H. R. (2002). Dynamics of neuronal responses in macaque MT and VIP during motion detection. *Nature Neuroscience*, *5*(10), 985–994.
Cope, A. J., Sabo, C., Vasilaki, E., Barron, A. B., & Marshall, J. A. R. (2017). A computational model of the integration of landmarks and motion in the insect central complex. *PloS One*, *12*(2), e0172325.
Cox, J., Minerva, A. R., Fleming, W. T., Zimmerman, C. A., Hayes, C., Zorowitz, S., Bandi, A., Ornelas, S., McMannon, B., Parker, N. F., & Witten, I. B. (2023). A neural substrate of sex-dependent modulation of motivation. *Nature Neuroscience*, *26*(2), 274–284.
DeAngelis, G. C., & Uka, T. (2003). Coding of horizontal disparity and velocity by MT neurons in the alert macaque. *Journal of Neurophysiology*, *89*(2), 1094–1111.
Diesmann, M., Gewaltig, M. O., & Aertsen, A. (1999). Stable propagation of synchronous spiking in cortical neural networks. *Nature*, *402*(6761), 529–533.

Ditterich, J., Mazurek, M. E., & Shadlen, M. N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, *6*(8), 891–898.
Doya, K. (2007). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. MIT Press.
Engel, T. A., & Wang, X.-J. (2011). Same or different? A neural circuit mechanism of similarity-based pattern match decision making. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *31*(19), 6982–6996.
Fisher, D., Olasagasti, I., Tank, D. W., Aksay, E. R. F., & Goldman, M. S. (2013). A modeling framework for deriving the structural and functional architecture of a short-term memory microcircuit. *Neuron*, *79*(5), 987–1000.
Fleming, W., Jewell, S., Engelhard, B., Witten, D. M., & Witten, I. B. (2021). Inferring spikes from calcium imaging in dopamine neurons. *PloS One*, *16*(6), e0252345.
Gage, G. J., Stoetzner, C. R., Wiltschko, A. B., & Berke, J. D. (2010). Selective activation of striatal fast-spiking interneurons during choice execution. *Neuron*, *67*(3), 466–479.
Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, *5*(1), 10–16.
Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.
Goldman, M. S. (2009). Memory without feedback in a neural network. *Neuron*, *61*(4), 621–634.
Goldman, M. S., Compte, A., & Wang, X.-J. (2009). Neural integrator models. *Encyclopedia of Neuroscience*, 165–178.
Gonçalves, P. J., Arrenberg, A. B., Hablitzel, B., Baier, H., & Machens, C. K. (2014). Optogenetic perturbations reveal the dynamics of an oculomotor integrator. *Frontiers in Neural Circuits*, *8*, 10.
Green, J., Adachi, A., Shah, K. K., Hirokawa, J. D., Magani, P. S., & Maimon, G. (2017). A neural circuit architecture for angular integration in Drosophila. *Nature*, *546*(7656), 101–106.
Hanks, T. D., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*, *520*(7546), 220–223.
Harris, K. D. (2021). Nonsense correlations in neuroscience. In *bioRxiv* (p. 2020.11.29.402719). https://doi.org/10.1101/2020.11.29.402719
Harvey, C. D., Coen, P., & Tank, D. W. (2012). Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature*, *484*(7392), 62–68.
Hayon, G., Abeles, M., & Lehmann, D. (2005). A model for representing the dynamics of a system of synfire chains. *Journal of Computational Neuroscience*, *18*(1), 41–53.
Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, *431*(7010), 859–862.
Heuer, H. W., & Britten, K. H. (2004). Optic flow signals in extrastriate area MST: comparison of perceptual and neuronal sensitivity. *Journal of Neurophysiology*, *91*(3), 1314–1326.
Hulse, B. K., & Jayaraman, V. (2020). Mechanisms Underlying the Neural Computation of Head Direction. *Annual Review of Neuroscience*, *43*, 31–54.
Jewell, S. W., Hocking, T. D., Fearnhead, P., & Witten, D. M. (2020). Fast nonconvex deconvolution of calcium imaging data. *Biostatistics* , *21*(4), 709–726.
Kim, I. S., & Dickinson, M. H. (2017). Idiothetic Path Integration in the Fruit Fly Drosophila melanogaster. *Current Biology: CB*, *27*(15), 2227–2238.e3.
Kim, S. S., Hermundstad, A. M., Romani, S., Abbott, L. F., & Jayaraman, V. (2019). Generation of stable heading representations in diverse visual scenes. *Nature*, *576*(7785), 126–131.
Koay, S. A. (2021). *Sequential and efficient neural-population coding of complex task information* [Data set]. Dryad. https://doi.org/10.5061/dryad.cvdncjt53
Koay, S. A., Charles, A. S., Thiberge, S. Y., Brody, C. D., & Tank, D. W. (2022). Sequential and

efficient neural-population coding of complex task information. *Neuron*, *110*(2), 328–349.e11.
Krishnan, N., Poll, D. B., & Kilpatrick, Z. P. (2018). Synaptic efficacy shapes resource limitations in working memory. *Journal of Computational Neuroscience*, *44*(3), 273–295.
Krumin, M., Lee, J. J., Harris, K. D., & Carandini, M. (2018). Decision and navigation in mouse parietal cortex. *eLife*, *7*. https://doi.org/10.7554/eLife.42583
Kumaran, D., & Maguire, E. A. (2005). The human hippocampus: cognitive maps or relational memory? *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(31), 7254–7259.
Laming, D. R. J. (1968). *Information Theory of Choice-reaction Times*. Academic P.
Lam, N. H., Borduqui, T., Hallak, J., Roque, A., Anticevic, A., Krystal, J. H., Wang, X.-J., & Murray, J. D. (2022). Effects of Altered Excitation-Inhibition Balance on Decision Making in a Cortical Circuit Model. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *42*(6), 1035–1053.
Langdon, C., & Engel, T. A. (2022). Latent circuit inference from heterogeneous neural responses during cognitive tasks. In *bioRxiv* (p. 2022.01.23.477431). https://doi.org/10.1101/2022.01.23.477431
Lim, S., & Goldman, M. S. (2012). Noise tolerance of attractor and feedforward memory models. *Neural Computation*, *24*(2), 332–390.
Lo, C.-C., & Wang, X.-J. (2006). Cortico–basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience*, *9*(7), 956–963.
Luna, R., Hernández, A., Brody, C. D., & Romo, R. (2005). Neural codes for perceptual discrimination in primary somatosensory cortex. *Nature Neuroscience*, *8*(9), 1210–1219.
MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal “Time Cells” Bridge the Gap in Memory for Discontiguous Events. *Neuron*, *71*(4), 737–749.
Machens, C. K., Romo, R., & Brody, C. D. (2005). Flexible control of mutual inhibition: a neural model of two-interval discrimination. *Science*, *307*(5712), 1121–1124.
Ma, H., Qi, Y., Gong, P., Lu, W., & Feng, J. (2022). Dynamics of bump attractors in neural circuits with emergent spatial correlations. In *arXiv [physics.bio-ph]*. arXiv. http://arxiv.org/abs/2212.01663
McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M.-B. (2006). Path integration and the neural basis of the “cognitive map.” *Nature Reviews. Neuroscience*, *7*(8), 663–678.
Monteiro, T., Rodrigues, F. S., Pexirra, M., Cruz, B. F., Gonçalves, A. I., Rueda-Orozco, P. E., & Paton, J. J. (2020). Manipulation of striatal population dynamics using temperature warps judgment of time. In *bioRxiv* (p. 2020.08.24.251827). https://doi.org/10.1101/2020.08.24.251827
Mountcastle, V. B., Steinmetz, M. A., & Romo, R. (1990). Frequency discrimination in the sense of flutter: psychophysical measurements correlated with postcentral events in behaving monkeys. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *10*(9), 3032–3044.
Nieh, E. H., Schottdorf, M., Freeman, N. W., Low, R. J., Lewallen, S., Koay, S. A., Pinto, L., Gauthier, J. L., Brody, C. D., & Tank, D. W. (2021). Geometry of abstract learned knowledge in the hippocampus. *Nature*, *595*(7865), 80–84.
O’Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford university press.
O’Reilly, J. X., Mesulam, M. M., & Nobre, A. C. (2008). The cerebellum predicts the timing of perceptual events. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *28*(9), 2252–2260.
Orhan, A. E., & Ma, W. J. (2019). Publisher Correction: A diverse range of factors affect the nature of neural representations underlying short-term memory. *Nature Neuroscience*, *22*(3), 505.

Parker, N. F., Baidya, A., Cox, J., Haetzel, L. M., Zhukovskaya, A., Murugan, M., Engelhard, B., Goldman, M. S., & Witten, I. B. (2022). Choice-selective sequences dominate in cortical relative to thalamic inputs to NAc to support reinforcement learning. *Cell Reports*, *39*(7), 110756.

Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsáki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, *321*(5894), 1322–1327.

Paton, J. J., & Buonomano, D. V. (2018). The Neural Basis of Timing: Distributed Mechanisms for Diverse Functions. *Neuron*, *98*(4), 687–705.

Pégard, N. C., Mardinly, A. R., Oldenburg, I. A., Sridharan, S., Waller, L., & Adesnik, H. (2017). Three-dimensional scanless holographic optogenetics with temporal focusing (3D-SHOT). *Nature Communications*, *8*(1), 1228.

Peixoto, D., Verhein, J. R., Kiani, R., Kao, J. C., Nuyujukian, P., Chandrasekaran, C., Brown, J., Fong, S., Ryu, S. I., Shenoy, K. V., & Newsome, W. T. (2021). Decoding and perturbing decision states in real time. *Nature*, *591*(7851), 604–609.

Philiastides, M. G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex* , *16*(4), 509–518.

Pinto, L., Koay, S. A., Engelhard, B., Yoon, A. M., Deverett, B., Thiberge, S. Y., Witten, I. B., Tank, D. W., & Brody, C. D. (2018). An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in Virtual Reality. *Frontiers in Behavioral Neuroscience*, *12*, 36.

Qi, Y., & Gong, P. (2015). Dynamic patterns in a two-dimensional neural field with refractoriness. *Physical Review. E, Statistical, Nonlinear, and Soft Matter Physics*, *92*(2), 022702.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Redish, A., Elga, A., & Touretzky, D. (1996). A coupled attractor model of the rodent head direction system. *Network*, *7*(4), 671–685.

Rickgauer, J. P., Deisseroth, K., & Tank, D. W. (2014). Simultaneous cellular-resolution optical perturbation and imaging of place cell firing fields. *Nature Neuroscience*, *17*(12), 1816–1824.

Roach, J. P., Churchland, A. K., & Engel, T. A. (2023). Choice selective inhibition drives stability and competition in decision circuits. *Nature Communications*, *14*(1), 147.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *22*(21), 9475–9489.

Salinas, E. (2006). How behavioral constraints may determine optimal sensory representations. *PLoS Biology*, *4*(12), e387.

Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *17*(15), 5900–5920.

Seung, H. S. (1996). How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(23), 13339–13344.

Shadlen, M. N., Britten, K. H., Newsome, W. T., & Movshon, J. A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *16*(4), 1486–1510.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, *86*(4), 1916–1936.

Siniscalchi, M. J., Wang, H., & Kwan, A. C. (2019). Enhanced Population Coding for Rewarded Choices in the Medial Frontal Cortex of the Mouse. *Cerebral Cortex* , *29*(10), 4090–4106.

Skaggs, W. E., Knierim, J. J., Kudrimoti, H. S., & McNaughton, B. L. (1995). A model of the neural basis of the rat's sense of direction. *Advances in Neural Information Processing Systems*, *7*, 173–180.

Song, P., & Wang, X.-J. (2005). Angular Path Integration by Moving "Hill of Activity": A Spiking

Neuron Model without Recurrent Excitation of the Head-Direction System. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *25*(4), 1002–1014.
So, N., & Shadlen, M. N. (2022). Decision formation in parietal cortex transcends a fixed frame of reference. *Neuron*, *110*(19), 3206–3215.e5.
Sorscher, B., Mel, G. C., Ocko, S. A., Giocomo, L. M., & Ganguli, S. (2023). A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, *111*(1), 121–137.e13.
Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, *25*(3), 251–260.
Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Reviews. Neuroscience*, *6*(5), 363–375.
Terada, S., Sakurai, Y., Nakahara, H., & Fujisawa, S. (2017). Temporal and Rate Coding for Discrete Event Sequences in the Hippocampus. *Neuron*, *94*(6), 1248–1262.e4.
Turner-Evans, D. B., Jensen, K. T., Ali, S., Paterson, T., Sheridan, A., Ray, R. P., Wolff, T., Lauritzen, J. S., Rubin, G. M., Bock, D. D., & Jayaraman, V. (2021). The neuroanatomical ultrastructure and function of a biological ring attractor. *Neuron*, *109*(9), 1582.
Turner-Evans, D., Wegener, S., Rouault, H., Franconville, R., Wolff, T., Seelig, J. D., Druckmann, S., & Jayaraman, V. (2017). Angular velocity integration in a fly heading circuit. *eLife*, *6*. https://doi.org/10.7554/eLife.23496
Uchida, N., & Mainen, Z. F. (2003). Speed and accuracy of olfactory discrimination in the rat. *Nature Neuroscience*, *6*(11), 1224–1229.
Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592.
Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*(5), 955–968.
Wimmer, K., Nykamp, D. Q., Constantinidis, C., & Compte, A. (2014). Bump attractor dynamics in prefrontal cortex explains behavioral precision in spatial working memory. *Nature Neuroscience*, *17*(3), 431–439.
Wong, K.-F., & Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *26*(4), 1314–1328.
Zemlianova, K., Bose, A., & Rinzel, J. (2022). A biophysical counting mechanism for keeping time. *Biological Cybernetics*, *116*(2), 205–218.
Zhang, K. (1996). Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *16*(6), 2112–2126.

SUPPLEMENTARY INFORMATION

SUPPLEMENTARY FIGURES

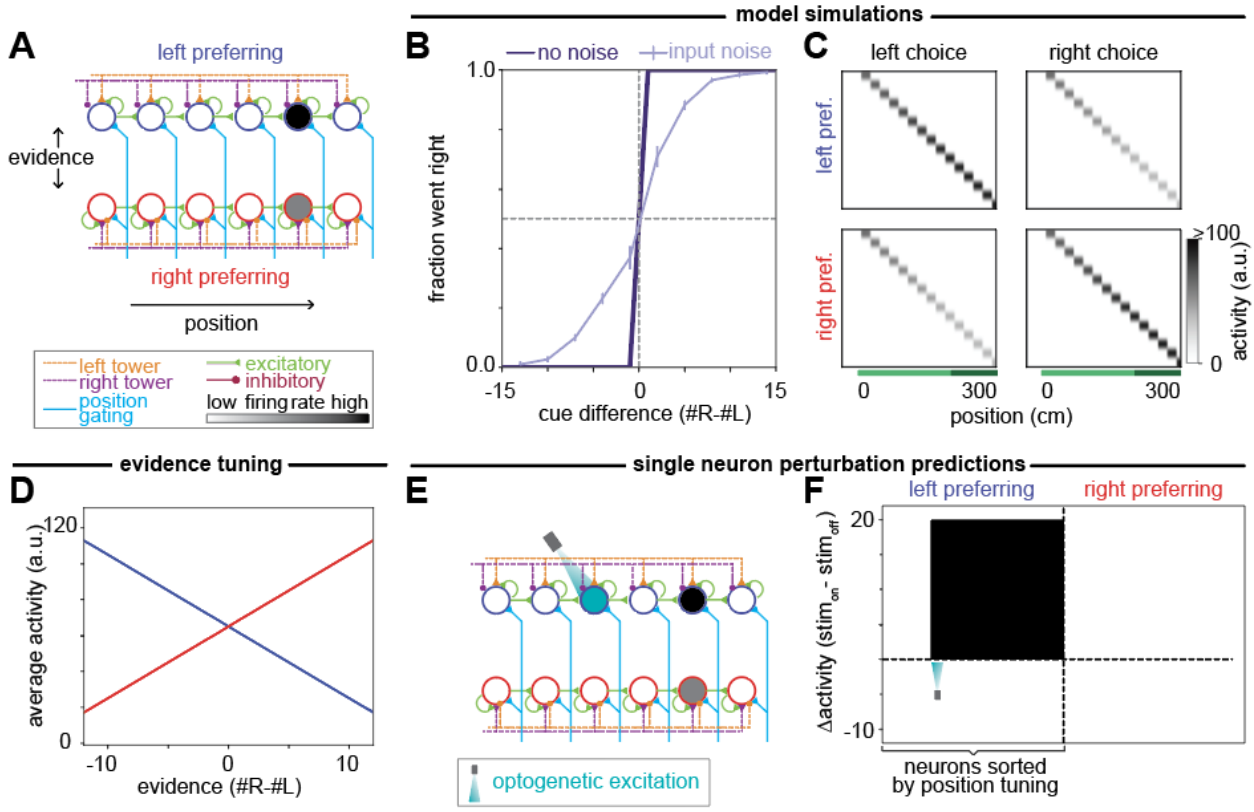

Figure S1. Uncoupled chains model of evidence accumulation through sequences. (A) Schematic of neural circuit model for the uncoupled chains model, showing excitatory (green) and inhibitory (red) connections between neurons (circles) as well as inputs to the circuit from the external left (orange) and right (purple) towers and a position gating signal (cyan). **(B)** Psychometric data for model simulated trials describing how often the amplitude of the final neuron in the left chain was greater than that of the final neuron in the right chain for cases in which the model was simulated with (light purple) and without (dark purple) noise in the input. Error bars show s.e.m. **(C)** Each row shows the non-normalized amplitude of a model neuron at each position in the maze, averaged across simulated trials without input noise when the greater final amplitude was in the left (left choice, left column) or right (right choice, right column) chain. Neurons were divided based on their choice-selectivity (see Methods) and ordered based on the position of peak activity. **(D)** Tuning curves of individual neurons to evidence by finding the average activity for different evidence levels at the neuron's peak position, for left-preferring (blue) and right-preferring (red) cells. **(E)** Schematic of a single neuron perturbation experiment in which optogenetic excitation is applied to a single neuron in the left chain. **(F)** Simulated changes of the firing rates of all neurons in the absence of cues when a single cell (denoted with the laser as in E) is optogenetically stimulated.

SUPPLEMENTARY TEXT

Below, we mathematically analyze models based on two chains, one chain representing left evidence, giving left choice-selective cells, and the other chain representing right evidence. In principle, these chains could be completely independent, with each chain integrating the cues only to its respective side, but such models would result in chains that only increase in

amplitude across the trial, inconsistent with observed data. Instead, we consider architectures with ipsilateral excitation and contralateral inhibition, where the contralateral inhibition comes either from external inputs or from the other chain. Specifically, we consider: (i) a model where the chains are uncoupled but towers appearing to the ipsilateral side are excitatory and towers appearing to the contralateral side are inhibitory (the uncoupled chains model, Fig. S1A) and (ii) a model where each chain only receives input from the ipsilateral side and competes with the other through mutual inhibition within the same position (the mutually inhibitory competing chains model, Fig. 2A).

Basic Features of Chains-Based Models

The neurons in each model can be represented by the same basic set of differential equations, where the firing rate of the ith neuron in the left chain, $r_{i,L}$, evolves according to

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} - er_{i,R} + f\mathbf{1}_{left}(t) - g\mathbf{1}_{right}(t) + P_i(t) - T]^+ \quad (S1)$$

The external cue inputs are given by $\mathbf{1}_{left}(t)$ for the left inputs and $\mathbf{1}_{right}(t)$ for the right inputs, which take the form

$$\mathbf{1}_{\text{left}}(t) = \begin{cases} 1 & \min_{\ell \in L} |t - \ell| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

and

$$\mathbf{1}_{\text{right}}(t) = \begin{cases} 1 & \min_{R \in \mathcal{R}} |t - R| < 0.5 \\ 0 & \text{otherwise} \end{cases}$$

where L is the set of times of left cues and $\mathcal{R}$ is the set of times of right cues.
The position signal is given by

$$P_i(t) = \begin{cases} T + E(t) & i \leq \frac{p(t)}{P_0} < (i+1) \\ 0 & \text{otherwise} \end{cases}.$$

where $p(t)$ is the position of the animal along the maze, so that each neuron receives an active position signal for length P_0, and only one neuron in each chain receives an active signal at any time.

For both models, we require the chains to exhibit a number of properties in order to accurately accumulate evidence, so that at the end of a trial, the chain corresponding to the side with more cues will have greater amplitude. In particular, we look to achieve the following properties: (i) Each neuron is only active transiently at a specific position, tiling space. (ii) External cue inputs are integrated linearly. (iii) In the absence of external inputs, when at a fixed position, the difference in neural activity of the two chains remains constant, without the sum growing unboundedly. (iv) In the absence of input, the difference in amplitude of the two chains is preserved across positions. We analyze the necessary conditions on the model parameters to achieve each of these properties.

Uncoupled Chains Model

For the case of the uncoupled chains model (Fig. S1), we take *e*=0, giving us

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} + f\mathbf{1}_{left}(t) - g\mathbf{1}_{right}(t) + P_i(t) - T]^+ \quad (S2)$$

We consider below the conditions needed for this model to achieve the desired properties.

Transiently Active at a Specific Position

Transient activity at a specific position is achieved through the position gating mechanism. We set the threshold *T* of the neuron sufficiently large that the term in square brackets will only affect the firing rate of the neuron in the presence of the position gating input $P_i(t)$. Otherwise, neural activity decays exponentially. This leads to the transient, sequential responses.

External Cues Integrated Linearly

In order for each cue to be integrated linearly, first, each cue should be weighted equally in its contribution to the final activity of the chain. This is achieved by giving each cue the same form of input in $\mathbf{1}_{left}(t)$ and $\mathbf{1}_{right}(t)$.

Second, this signal must be fully integrated by the chain. Since the external cue input term resides in the square brackets, only the neuron with the active position gating signal can integrate this signal. For the neurons with an active position gating signal,

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} + f\mathbf{1}_{left}(t) - g\mathbf{1}_{right}(t) + E(t)]^+$$

so we only need the term inside square brackets to be positive for the two neurons receiving the position-gating signal. We can see that this is guaranteed if $br_{i,L} + cr_{i-1,L}$ is greater than $-g$. We show below that the difference in activity of the chains is preserved across positions in the absence of external cues when *b*=*c*=*a*, so that $r_{i,L}+r_{i-1,L} = C$ for some constant *C*. Thus, by initializing the activity at the start of the trial, $r_{0,L}$, to be sufficiently large that the activity of the lower firing rate chain does not drop below *g/a*, these inputs will always be integrated.

We note that this does not require the square pulse cues that we specified here but just that each cue has the same form.

Difference in Activity Between Chains Remains Constant at a Fixed Position

When the animal maintains a fixed position, the activity of the neurons at the preceding position have decayed away, so that in the absence of external input, we have

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + \left[br_{i,L} + P_i(t) - T\right]^+.$$

By taking *E(t)*=0, we have $P_i(t) = T$ at the active position, and thus

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + \left[br_{i,L}\right]^+ = -ar_{i,L} + br_{i,L}.$$

Thus, when *a*=*b*, each chain maintains constant activity, both preserving the difference between the chains and keeping the sum constant.

Note that if *E(t)* > 0, this signal will be integrated, causing both chains to increase in amplitude. Such a signal could be designed to act as an urgency signal, but must be carefully crafted to prevent the chains from growing too rapidly.

Difference in Amplitude Preserved Across Positions

Preserving the amplitude of the chain across positions is a special case of preserving the difference across positions. To preserve amplitude across positions in the absence of cues, the next neuron in the chain must integrate the activity of the previous neuron at the same rate at which it decays. At the active position in the absence of input, and assuming *E(t)*=0, we have

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + \left[ar_{i,L} + cr_{i-1,L}\right]^{+}$$

$$\frac{dr_{i,L}}{dt} = cr_{i-1,L}$$

while we also have

$$\frac{dr_{i-1,L}}{dt} = -ar_{i-1,L}$$

To preserve the total amplitude of activity,

$$-\frac{dr_{i-1,L}}{dt} = \frac{dr_{i,L}}{dt},$$

giving us the condition *a*=*c*.

We note that neuron *i* will only fully integrate the activity of the previous position in the infinite time limit. Due to the transient nature of the position-gating signal, we have that the final amplitude of the ith neuron will be

$$r_{i,L}(P_0) = \int_0^{P_0} ar_{i-1,L}dt = \int_0^{P_0} ar_{i-1,L}(0)e^{-at}dt = r_{i-1,L}(0)\left(1 - e^{-aP_0}\right).$$

For *a* and P_0 sufficiently large, the difference in amplitude between positions is exponentially small.

Parameterization of the Uncoupled Chains Model

In our simulations, in addition to taking *E(t)* = 0, we set *a* =1, *b* =1, *c* = 1, *f* =4, *g* = 4, P_0 =20, and *T* = 300. Neurons $r_{0,L}$ and $r_{0,R}$ are initialized to 65 activity level. We note that the neurons must be initialized to a sufficiently high non-zero activity level such that an external input to the contralateral side can be negatively integrated into the activity.

Mutually Inhibitory Competing Chains Model

In this model, the chains compete through mutual inhibition, rather than receiving opposing inputs. For the mutually inhibitory competing chains model, we take *g*=0 in Equation S1, giving

$$\frac{dr_{i,L}}{dt} = -ar_{i,L} + [br_{i,L} + cr_{i-1,L} - er_{i,R} + f\mathbf{1}_{left}(t) + P_i(t) - T]^{+}.$$

We now turn to the necessary conditions for this model to achieve the desired properties.

Transiently Active at a Specific Position

Transient activity at a specific position is achieved through the same position gating mechanism as before, with *T* much larger than the other terms in square brackets so that the neuron will only be active when the position gate $P_i(t)$ is active.

External Cues Integrated Linearly

Each cue should be weighted equally in its contribution to the final activity of the chain. This is achieved by giving each cue the same form of input in $\mathbf{1}_{left}(t)$ and $\mathbf{1}_{right}(t)$. Moreover, the weights of mutual inhibition and self-excitation must be tuned such that our system of equations is a perfect integrator when the term in square brackets is above threshold, which the position gating ensures can be true for at most two neurons, both at the same position. Specifically, when the term in square brackets is above threshold, we have

$$\frac{d}{dt}\begin{pmatrix} r_{i,L} \\ r_{i,R} \end{pmatrix} = \begin{pmatrix} b-a & -e \\ -e & b-a \end{pmatrix}\begin{pmatrix} r_{i,L} \\ r_{i,R} \end{pmatrix} + \begin{pmatrix} cr_{i-1,L} + f\mathbf{1}_{left}(t) + E(t) \\ cr_{i-1,R} + f\mathbf{1}_{right}(t) + E(t) \end{pmatrix},$$

which has eigenvalues (*b-a-e*) for the common mode eigenvector, (1, 1), and (*b-a+e*) for the difference mode eigenvector, (1, -1). For a perfect integrator that accumulates evidence in the difference of firing between the two chains, we require that the eigenvalue of the difference mode eigenvector is zero, giving us e = a-b. We note that, for the assumption that both neurons are above threshold to be true, we require $E(t)+br_{i,L}+cr_{i-1,L} - er_{i,R} > 0$ and $E(t)+br_{i,R}+cr_{i-1,R} - er_{i,L} > 0$. This can be achieved by setting $E(t) = I_{ext}$ for some sufficiently large constant I_{ext}, and integration will only be perfect in the range that this is true. This common external input to both chains serves as a background about which integration occurs at each position in the chain.

Difference in Activity Between Chains Remains Constant at a Fixed Position

Holding the difference in activity between the neurons in the chain constant at a constant position is a natural consequence of having a perfect integrator for the difference mode eigenvector. Furthermore, because the common mode eigenvector is associated with a negative (decay-associated) eigenvalue, the sum of activity will not grow uncontrollably.

Difference in Amplitude Preserved Across Positions

To appropriately represent the difference in accumulated evidence between the chains, we require the difference between the chains to be preserved between positions.

$$\frac{d\Delta_i}{dt} = -\,ar_{i,L} + br_{i,L} + cr_{i-1,L} - er_{i,R} - \left(-\,ar_{i,R} + br_{i,R} + cr_{i-1,R} - er_{i,L}\right)$$

$$\frac{d\Delta_i}{dt} = (-\,a + b + e)r_{i,L} + (a - b - e)r_{i,R} + c(r_{i-1,L} - r_{i-1,R})$$

$$\frac{d\Delta_i}{dt} = (-\,a + b + e)\Delta_i + c\Delta_{i-1}$$

From our conditions on perfect integration, we have (-a+b+e)=0, giving us

$$\frac{d\Delta_i}{dt} = c\Delta_{i-1}.$$

We also have that

$$\frac{d\Delta_{i-1}}{dt} = -\,ar_{i-1,L} + ar_{i-1,R} = -\,a\Delta_{i-1}$$

For the total difference in chains to be preserved, we require

$$-\frac{d\Delta_{i-1}}{dt} = \frac{d\Delta_i}{dt},$$

giving us the condition *a*=*c*.

As in the case of the uncoupled chains model, due to the finite time of integration, we will have that

$$\Delta_i(P_0) = \Delta_{i-1}(0)\left(1 - e^{-aP_0}\right)$$

so that the difference between positions is exponentially small for *a* and P_0 sufficiently large.