bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

A network model of path control

Mohammed Abdal Monium Osman[1, 2], Wenyi Zhang[1],
Cengiz Pehlevan[2, 3, 4, *], and Rachel I. Wilson[1, 4, *]

[1]Department of Neurobiology, Harvard Medical School
[2]Kempner Institute for the Study of Natural and Artificial Intelligence, Harvard University
[3]John A. Paulson School of Engineering and Applied Sciences, Harvard University
[4]Center for Brain Science, Harvard University
[*]Correspondence to: cpehlevan@seas.harvard.edu, rachel_wilson@hms.harvard.edu

Abstract

We present a network model that enables a navigating agent to shape the spatial and temporal structure of its paths through its environment. The model is inspired by the insect fan-shaped body. In it, a recurrent neural network receives an input that encodes an agent's velocity vector. Meanwhile, network activity determines the setpoint of a control loop that influences this velocity vector through motor outputs. The network thus guides future movement using memory of past movement. We illustrate how the network's dynamics determine the trajectory features it extracts to produce the setpoint. We then relate the joint dynamics of network activity and the agent's velocity vector to the structure of the resulting paths. This network can produce paths suitable for different forms of exploration. It can also stabilize movements initiated by lower-level systems. Our analyses illustrate how recurrent neural dynamics and memory can shape behavior when embedded in closed loops.

1 Introduction

The fan-shaped body (FB) of the insect central complex is a brain region downstream of the head direction system. It contains a set of neural arrays in which sinusoidal activity patterns encode vectors relevant to navigation [1–5]. Each array constitutes a "columnar cell type" of the FB (Fig. 1A). Activity across these arrays is coupled by highly structured, array-specific patterns of recurrent connectivity (Fig. 1B,C; [3]). Whereas feedforward networks of such arrays can implement arithmetic operations like vector addition and rotation (Fig. 1D,E; [1, 2]), it is not entirely clear what computations are performed by recurrent ensembles of sinusoidal arrays.

To understand this network, we need to consider the inputs and outputs that couple network activity to behavior (Fig. 1F,G). One of the chief inputs to the FB is a sinusoidal array whose activity encodes the animal's velocity vector, i.e. its current direction and speed of travel in the ground plane (PFNd/PFNv cells) [4, 5]. It has been proposed that the FB might use this input for path integration, the estimation of positional displacement from a time series of velocity measurements [6–10]. Path integration is commonly thought to underlie homing behavior in insects like bees and ants [11–13]. It has been suggested that a sinusoidal array could path-integrate a velocity input to produce a sinusoid that encodes position [14].

While path integration remains a plausible role for the FB, recent work in *Drosophila* suggests a more complicated picture. Based on recent connectomes [3], the FB actually contains many parallel pathways downstream of the velocity input, making the network more intricate than we might expect from minimal models of path integration (e.g. [14]). In addition, recent calcium imaging experiments have found no columnar cell types that encode a classical path integral; instead, many of the relevant cell types seem to encode lagged and smoothed versions of the velocity input [15–18], consistent with relatively transient, leaky integration dynamics. These results suggest that the FB may temporally filter its velocity input in

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

more varied ways than classical path integration, but it remains unclear how these more transient forms of integration might facilitate navigation.

On the output side, FB activity contributes to locomotor control. Specifically, PFL2/PFL3 cells continuously compare the animal's heading to a setpoint specified by activity in the FB; whenever the animal's heading deviates from this setpoint, PFL2/PFL3 cells slow the animal's forward speed and drive corrective turning toward the setpoint [15, 19]. This setpoint is a two-dimensional vector encoded by the summed activity patterns of multiple arrays (columnar cell types) in the FB. These arrays are themselves downstream from the velocity vector input signal. Thus, the FB is part of a closed loop in which behavior (the animal's velocity vector) influences neural activity, and vice versa.

Finally, it is also important to recognize that the FB does not completely specify locomotor behavior. Instead, behavior emerges from the combined influence of the FB and other circuits acting in parallel, such as lower-level sensorimotor systems that steer the animal toward or away from salient stimuli. It remains unclear how these parallel pathways are coordinated with the effects of the FB.

To explore the FB's role in navigation, we developed a network model based on its architecture. In this model, a network of recurrently interacting sinusoidal arrays receives an external input that encodes an agent's velocity vector. We first illustrate how the network's recurrent connectivity determines the temporal filter that each array applies to the velocity input. This filtering allows the network to extract information about the agent's trajectory over a variety of time scales and frequency bands, and it generalizes classical path integration to more varied forms of signal processing.

We then couple the agent's locomotor behavior to network activity through a control policy based on that of the FB [6, 15, 19, 20], thus defining a closed neural-behavioral loop. Within our model, agent behavior is also influenced by a forcing function meant to capture lower-level sensorimotor pathways that are independent of the FB. Analyzing this neural-behavioral system, we relate the structure of the network and forcing function to the shapes of the agent's paths through its environment. The resulting paths can combine meandering, oscillatory, and persistent movement, potentially supporting a variety of exploratory behaviors. The network can also stabilize movements initiated by the lower-level pathways, thereby enabling the agent to persistently move toward an intermittently sensed cue.

Based on this modeling, we propose that the FB may use its velocity input to extract and store the features of past movement needed to generate a path with a desired spatial structure. These results demonstrate how memory of past movement can be used to organize structured behavioral sequences. More generally, they illustrate how recurrent neural circuits embedded in closed loops can shape behavior.

2 Results

2.1 A network model inspired by the FB

The FB is divided into horizontal rows and vertical columns. Different rows (or "layers") serve different functions [3]. In this study, we will focus on the intermediate layers of the FB, which contain the circuitry most directly implicated in spatial navigation [4–6, 15, 19]. Each layer is horizontally tiled by "columnar cells". Each columnar cell type forms a complete set of tiles, and cells with adjacent horizontal positions are tuned to adjacent compass directions (Fig. 1A; [3]). We will model the FB as N populations of neurons, each representing a columnar cell type. Each neuron in this network model will be indexed by a cell type label $i \in \{1, \ldots, N\}$ indicating the population it belongs to and an angular position $\theta \in [0, 2\pi)$ indicating its column.

Within the FB, columnar neurons are coupled by recurrent connections whose strength depends on their respective cell types and relative positions (Fig. 1B; [3]). These connections come in two varieties: phase-aligned connections between neurons that share the same column, and phase-inverted connections between neurons that occupy opposing columns (columns separated by exactly half the horizontal extent of the FB) (Fig. 1B; [3]). These phase-inverting connections are enabled by a class of columnar cell types called hΔ (Fig. 1C; [3, 21, 22]).

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

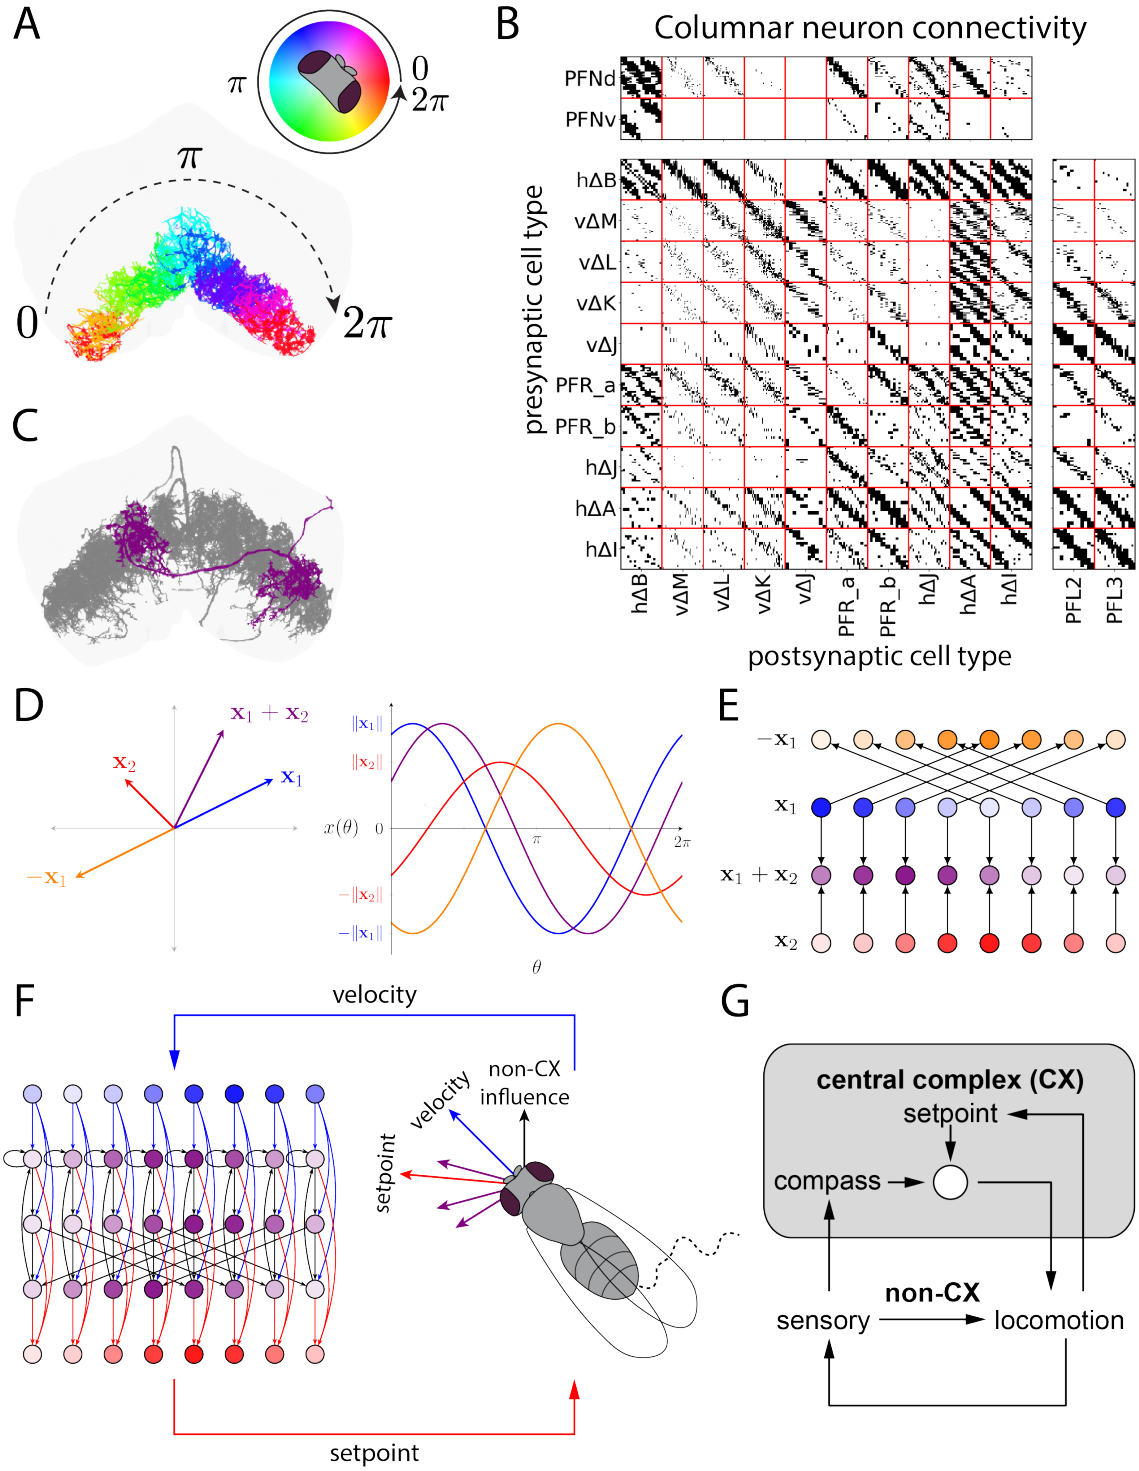


Figure 1: Columnar organization of the FB. (A) Neurons of a given columnar cell type tile the FB, subdividing it into columns that correspond to the 2π angular extent of compass directions. (B) Binarized connectivity of a subset of the columnar cells downstream of the network's velocity input (PFNd/PFNv) and upstream of the network's motor output (PFL2/PFL3). Note that within each block, neurons of a given pre- and post-synaptic type, sorted by their column, topographically connect with some characteristic strength. Each block consists of phase-preserving (diagonal) and/or phase-inverting (off-diagonal) connections. (C) Phase-inverting connections are permitted by the morphology of hΔ neurons, which contain arbors in pairs of opposing columns. (D) A two-dimensional vector can be associated with a sinusoidal activity pattern whose phase and amplitude match the vector's direction and length, an encoding scheme called a phasor representation [2]. (E) Phasor representations can be instantiated by spatial activity patterns in arrays of columnar neurons, with simple feedforward motifs implementing operations like vector addition and inversion [1]. (F) The animal's behavior and neural activity define a closed loop. In it, behavior influences neural activity through an input that encodes the animal's velocity vector (blue), this input is processed by a recurrent network of sinusoidal arrays (purple), and network activity influences behavior through a setpoint signal (red) produced by summing the vectors encoded by these activity patterns. Finally, non-CX circuits independent of this loop influence the velocity vector in parallel (represented by the black arrow). (G) Block diagram schematizing the system's control architecture.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Within our network model, the strengths of these connections are represented by two matrices: $\mathbf{W}^+, \mathbf{W}^- \in \mathbb{R}_+^{N\times N}$. W_{ij}^+ parameterizes the strength of the phase-aligned connections from neurons of cell type j to those of cell type i, and W_{ij}^- is defined analogously for phase-inverted connections. In our model, all these connections are non-negative, a choice motivated by experimental evidence that FB columnar neurons are excitatory [23, 24].

Next, we consider the input that informs this network about the agent's behavior. In *Drosophila*, neurons in the intermediate layers of the FB receive input from cells (PFNd/PFNv cells; [4, 5]) that encode the animal's velocity vector via a sinusoidal spatial pattern, an encoding scheme called a phasor representation (Fig. 1D,E; [1, 2]). We model this input as a spatial activity profile that encodes the current velocity $\mathbf{v}(t) \in \mathbb{R}^2$ of an agent moving through 2D space (Fig. 2A-C):

$$v(\theta, t) = k_V \|\mathbf{v}(t)\| \left(1 + \cos\left(\theta - \arg(\mathbf{v}(t))\right)\right). \tag{1}$$

Here, $v(\theta, t)$ denotes the firing rate of the input neuron in column θ at time t. The phase of this input profile encodes the agent's direction of movement $\arg(\mathbf{v}) \in [0, 2\pi)$, and its amplitude encodes the agent's translational speed $\|\mathbf{v}(t)\|$, scaled by a constant k_V that translates units of speed into firing rate. Because firing rates cannot be negative, the sinusoid includes a non-zero mean that ensures its trough is at zero, rather than a negative value. This velocity input is delivered in parallel to multiple columnar cell types, each of which weights it by a cell-type-specific connection strength b_i'.

We then model the neural dynamics that result from this connectivity using a linear recurrent circuit model. In it, each columnar neuron acts as a leaky integrator of its phase-aligned, phase-inverting, and velocity inputs (Fig. 2A,B):

$$\tau \frac{\partial}{\partial t} x_i(\theta, t) = -x_i(\theta, t) + \sum_{j=1}^{N} \left[W_{ij}^+ x_j(\theta, t) + W_{ij}^- x_j(\theta - \pi, t) \right] + b_i' v(\theta, t), \tag{2}$$

where τ is a neuronal time constant and $x_i(\theta, t)$ denotes the firing rate at time t of the neuron of cell type i in column θ.

Due to the sinusoidal shape of the external input and the structure of the network's dynamics, each cell type's spatial activity profile will always be a sinusoid of period 2π, consistent with the kinds of spatial activity patterns observed experimentally. Thus, each activity profile is fully described by its first harmonic $\hat{\mathbf{x}}_i^{(1)}(t) \in \mathbb{R}^2$ (which describes the sinusoid's phase and amplitude) and its mean $\mu_i(t) \in \mathbb{R}_+$ (Fig. 2C; Appendix A):

$$x_i(\theta, t) = \mu_i(t) + \left\| \hat{\mathbf{x}}_i^{(1)}(t) \right\| \cos\left(\theta - \arg\left(\hat{\mathbf{x}}_i^{(1)}(t)\right)\right). \tag{3}$$

In the following section, we will illustrate how this network temporally processes the velocity input. Next, we will specify how network activity influences the agent's velocity, thereby closing the loop.

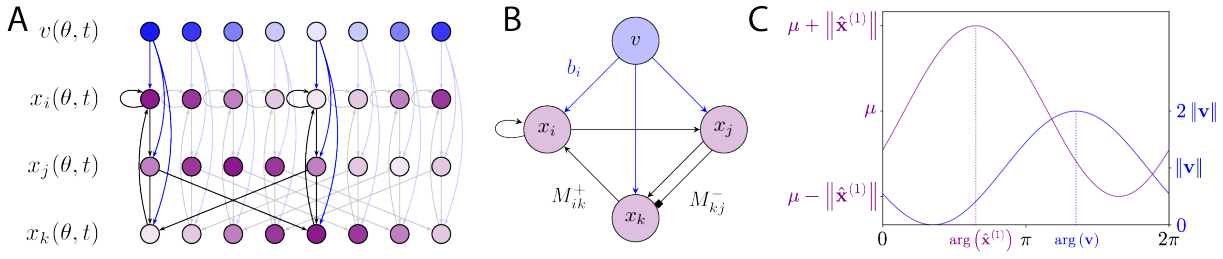


Figure 2: Overview of network model. (A) The network model consists of a set of neural arrays, each meant to represent a columnar cell type. Each array receives an external input that encodes the agent's velocity. Neurons are recurrently coupled by phase-aligned and phase-inverted connections whose strength depends on their respective cell types. Connectivity is shown for one example instantiation of the network, with highlighted connections for a pair of opposing columns (at positions 0 and π). (B) Cell-type-level summary of connections in (A). (C) At each time point, the activity profiles of the arrays in the network are sinusoids of period 2π.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

2.2 Temporally filtering the velocity input

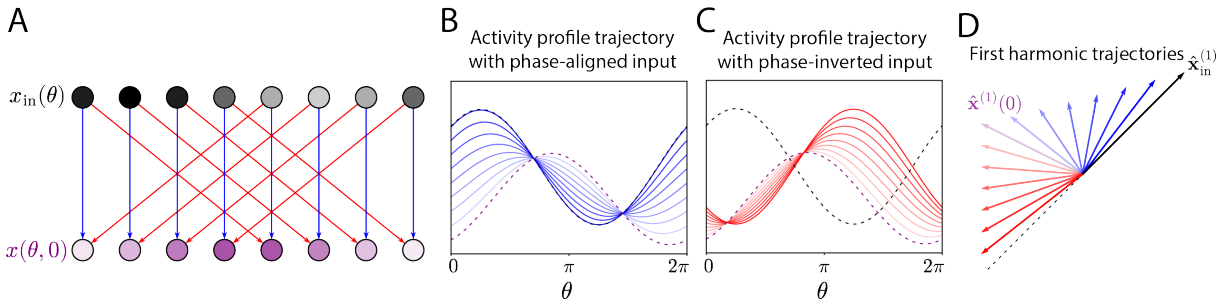


Figure 3: Dynamics of the first harmonic. (A) In this example, a neural array (purple) receives either phase-aligned (blue) or phase-inverted (red) connections from an input array (black) whose activity profile is constant over time. (B) The array's activity profile is a time-varying sinusoidal pattern. When the input is phase-aligned, the activity profile flows from its initial state (purple dotted line) to that of the input array (black dotted line). (C) When the input is phase-inverted, the phase of the activity profile at steady state is opposite that of the input. (D) The nature of the phase shift determines whether the first harmonic of the activity profile is attracted towards or repelled by the first harmonic of the input.

In this section, we will explore how the agent's velocity shapes network activity. Recall that this network is composed of N arrays, and the activity of each array can be decomposed into two components: the first harmonic and the mean. These two components of network activity do not influence one another (Appendix A). Here, we will restrict our focus to the first harmonic (though see Appendix B for an overview of the mean component).

Let $\hat{\mathbf{X}}^{(1)}(t) \in \mathbb{R}^{N\times 2}$ denote the state matrix whose rows contain the first harmonics of the N activity profiles. The dynamics of $\hat{\mathbf{X}}^{(1)}$ follow

$$\tau \frac{d}{dt}\hat{\mathbf{X}}^{(1)}(t) = -\hat{\mathbf{X}}^{(1)}(t) + (\mathbf{W}^{+} - \mathbf{W}^{-})\hat{\mathbf{X}}^{(1)}(t) + \mathbf{b}\mathbf{v}(t)^{\top}, \tag{4}$$

where $\mathbf{b} \in \mathbb{R}^N$ is a vector of input weights, with $b_i = k_V b_i'$. Note that the recurrent couplings between the first harmonics depend on the differences between the phase-aligned and phase-inverted connection strengths through the coupling matrix $\hat{\mathbf{W}}^{(1)} = \mathbf{W}^{+} - \mathbf{W}^{-}$. This is because phase-inverting connections shift the influence of the presynaptic activity profile by π radians, which is equivalent to multiplying its first harmonic by -1 (Fig. 1E; [1]). Thus, even though all the single-neuron connections are excitatory, phase-inverting connections act as negative couplings between the first harmonics. Negative couplings promote repulsion between first harmonics, whereas positive couplings promote attraction (Fig. 3).

To understand the computations being performed by this network, we next consider its responses to a time-varying input. Because the dynamics are linear and time-invariant, each array will temporally filter the network's input using a kernel that depends on the network's connectivity [25, 26]. The first harmonic of this input encodes the agent's velocity vector. Consequently, the first harmonic of each activity profile will temporally filter the velocity history to compute a displacement-like quantity (Fig. 4A,B):

$$\hat{\mathbf{x}}_i^{(1)}(t) = \left(\mathcal{H}\left[\hat{\mathbf{W}}^{(1)} - \mathbf{I}, \mathbf{b}, \mathbf{e}_i, \tau\right] * \mathbf{v}\right)(t) = \int_0^t dt'\, \mathcal{H}\left[\hat{\mathbf{W}}^{(1)} - \mathbf{I}, \mathbf{b}, \mathbf{e}_i, \tau\right](t-t')\mathbf{v}(t'). \tag{5}$$

Above, $\mathcal{H}[\dots]$ is a functional that takes the parameters of a linear dynamical system (its internal couplings, input weights, readout weights, and time constant) and returns the temporal filter implemented by that system (Appendix C). The internal couplings here correspond to the network's recurrent connectivity and leak, and they determine the filter's characteristic time scales and oscillation frequencies. Meanwhile, the input and readout weights determine how strongly each of these components shapes a given feature of the system's response to the input. In this case, the feature of interest is the first harmonic of the ith array, rather than a linear combination of multiple activity profiles, so the readout weights simply select that array using a one-hot vector $\mathbf{e}_i \in \mathbb{R}^N$.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

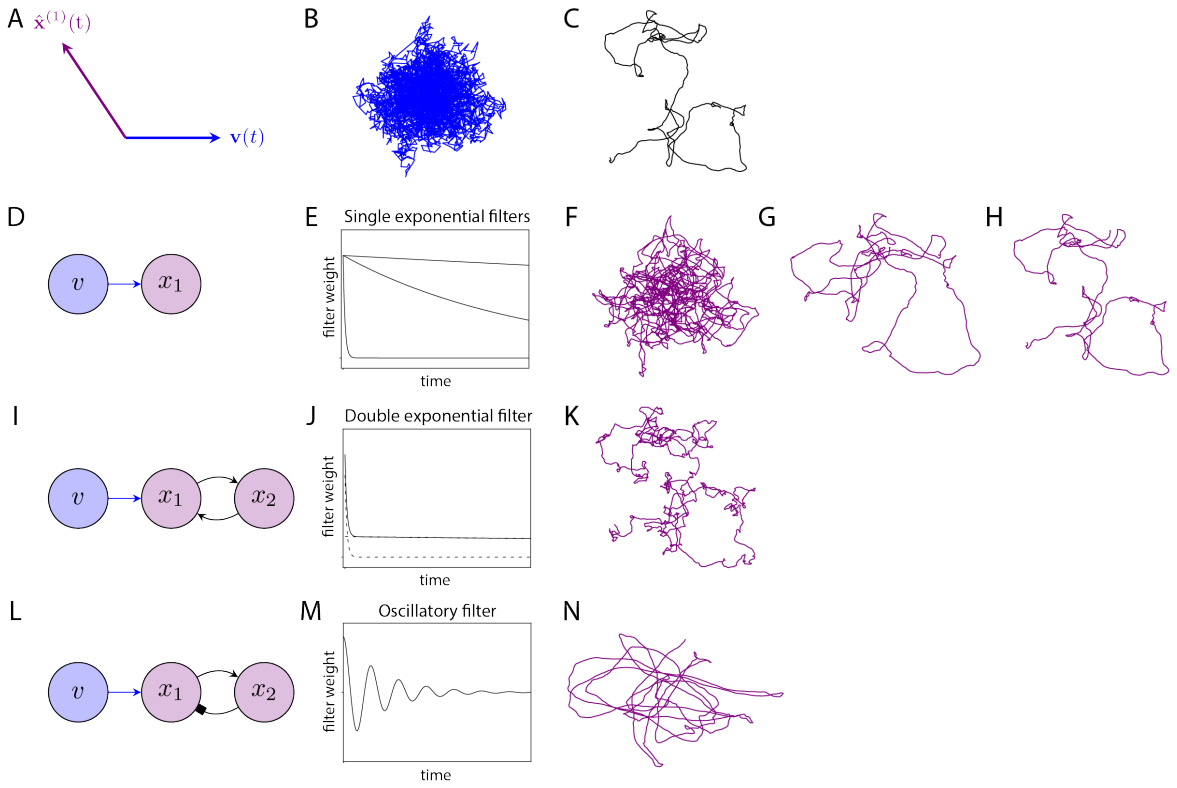


Figure 4: Temporally filtering the velocity input. (A) The first harmonic of each activity profile is a time-varying vector that depends on the sequence of velocity inputs and the network's connectivity. (B) Example sequence of velocity vectors that acts as input in the examples that follow. (C) The associated positional trajectory. (D) A single self-coupled neural array integrating the velocity vector input. (E) Resulting temporal filters for different coupling strengths. (F-H) Resulting trajectories of the first harmonic of the neural array, ordered by increasing filter persistence. The most transient filter more closely tracks velocity, while the most persistent filter more closely tracks position. (I) A pair of rings coupled by phase-aligned connections. (J) The resulting filter is a double exponential. (K) This produces a first harmonic trajectory with structure on two distinct time scales. (L) A pair of asymmetrically coupled rings, in which one connection is phase-preserving and the other is phase-inverting. (M) The resulting filter is band-pass. (N) This generates a first harmonic trajectory that is oscillatory.

When this filter is constant in time (e.g. $\mathcal{H}[\dots](t) = 1$ for all t), the first harmonic tracks the agent's positional displacement (Fig. 4C); this is classical path integration ([14]; see also [27–29]). However, the network can also realize more varied temporal filters, the shapes of which will depend on its connectivity (Appendix C). For example, a single-array network will act as a leaky integrator of the input, producing a smoothed estimate of recent displacement over the filter's effective time window (Fig. 4D-H). In a multi-array network, each array can instead filter the input with a sum of exponentials, thereby mixing information about displacement over different time scales (Fig. 4I-K). Finally, a network with phase-inverted connections and oscillatory dynamics can also apply sinusoidal band-pass filters to the input (Fig. 4L-N). These networks generalize path integration to the wider variety of signal processing operations that can be realized by linear dynamical systems, enabling the network to extract information about the agent's trajectory over various time scales and frequency bands.

2.3 Coupling behavior to neural activity through motor output

To close the loop, we need to specify how this network influences behavior. Recent experimental work on walking *Drosophila* has shown that the FB influences the animal's heading and speed through two specialized populations of columnar output neurons: PFL2 and PFL3 [6, 15, 19]. Specifically, these output neurons continuously compare the fly's current heading to a vector setpoint signal. When the heading is not aligned with the setpoint (e.g., when the fly turns away from the setpoint direction to inspect an object it encounters along its path, or to avoid an obstacle), these output neurons drive corrective turning

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

and reduce forward speed, until the fly's heading is once again aligned with the setpoint [15, 19]. The setpoint vector is encoded by a sinusoidal activity pattern that reflects the summed contributions of the FB columnar cell types that synapse onto PFL2 and PFL3 cells [6, 15, 18, 19].

Motivated by this experimental work, we model the setpoint signal as a weighted sum of activity profiles:

$$u(\theta, t) = \sum_{i=1}^{N} c_i' x_i(\theta, t), \tag{6}$$

where $c_i' \in \mathbb{R}_+$ is the strength of the connections from neurons of cell type i to output neurons (PFL2 and PFL3 neurons) in the same column. Since this setpoint signal is a sum of sinusoids, it too is a sinusoid. Within our model, the first harmonic $\hat{\mathbf{u}}^{(1)}(t)$ of this sinusoid determines the instantaneous setpoint of the agent's velocity vector:

$$\tau \frac{d}{dt}\mathbf{v}(t) = -\gamma \mathbf{v}(t) + k_U \hat{\mathbf{u}}^{(1)}(t) + \mathbf{f}(t), \tag{7}$$

where γ is a drag-like coefficient that represents the agent's intrinsic tendency to slow down, and $k_U \in \mathbb{R}_+$ is a constant that specifies the magnitude of the locomotor drive that results from the setpoint signal. Equation 7 describes a control policy that continuously compares the velocity and setpoint vector and drives corrections when the two are misaligned. When rewritten in polar coordinates, this recapitulates the basic features of the heading and speed control performed by PFL2 and PFL3 cells (Appendix D). In addition, the velocity dynamics also depend on a forcing function $\mathbf{f}(t) \in \mathbb{R}^2$ that is independent of the FB-like network. This forcing function captures the time-varying influence of sensory stimuli and motor drives that influence the agent's velocity vector through pathways that bypass the central complex. Finally, note that these velocity dynamics are linear. In Appendix G, we explore how nonlinearity in the transfer function that relates neural output to behavior affects the system's dynamics.

2.4 Three descriptions of the closed-loop system

While neural activity influences the agent's behavior through motor output, agent behavior also influences neural activity through the network's velocity vector input (Fig. 5A). In this section, we consider three ways of reasoning about these reciprocal interactions.

First, we can describe the agent's behavior as an autoregressive process that depends on the full history of past behavior. Here, we can eliminate neural variables altogether by rewriting the setpoint vector $\hat{\mathbf{u}}^{(1)}(t)$ as a filtered version of the velocity vector:

$$\tau \frac{d}{dt}\mathbf{v}(t) = -\gamma \mathbf{v}(t) + \left(\mathcal{H}\left[\hat{\mathbf{W}}^{(1)} - \mathbf{I}, \mathbf{b}, \mathbf{c}, \tau\right] * \mathbf{v}\right)(t) + \mathbf{f}(t), \tag{8}$$

where $\mathbf{c} \in \mathbb{R}^N$ is a vector of output weights, with $c_i = k_U c_i'$. The filter value $\mathcal{H}[\ldots](t - t')$ determines how strongly the current change in velocity is driven by the velocity at time t'. Thus, the shape of the agent's path depends on the shape of the filter, with more persistent filtering leading to more persistently correlated (or anti-correlated) behavior.

Second, we can describe behavior and neural activity as different dimensions of the same joint system:

$$\tau \frac{d}{dt}\begin{bmatrix}\mathbf{v}(t)^\top \\ \hat{\mathbf{X}}^{(1)}(t)\end{bmatrix} = \underbrace{\begin{bmatrix}-\gamma & \mathbf{c}^\top \\ \mathbf{b} & \hat{\mathbf{W}}^{(1)} - \mathbf{I}\end{bmatrix}}_{\mathbf{A}} \begin{bmatrix}\mathbf{v}(t)^\top \\ \hat{\mathbf{X}}^{(1)}(t)\end{bmatrix} + \begin{bmatrix}\mathbf{f}(t)^\top \\ \mathbf{0}\end{bmatrix}. \tag{9}$$

In this equation, the history dependence that was explicitly captured by the filter in equation 8 is instead summarized by the current activity pattern of each array. The dynamics of this system are set by the closed-loop coupling matrix $\mathbf{A} \in \mathbb{R}^{(N+1)\times(N+1)}$. $\mathbf{A}$ jointly describes the network's recurrent dynamics, the intrinsic (velocity-drag) dynamics of behavior, and the couplings between neural activity and behavior. From this perspective, the network's connectivity influences behavior by influencing the properties of $\mathbf{A}$, an idea known as pole placement in control theory [26].

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Finally, we can describe the agent's velocity as a filtered version of the motor drives captured by the forcing function:

$$\mathbf{v}(t) = \left(\mathcal{H}\left[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau\right] * \mathbf{f}\right)(t). \tag{10}$$

In this view, the dynamics of the joint system determine how the forcing function is transformed into a trajectory through the environment. As a result, the agent is not merely reactive, in that the forcing function alone does not dictate its instantaneous behavior. Rather, the network "observes" and "remembers" the agent's past behavior, and this memory influences moment-to-moment responses to the forcing function.

In short, we can describe this system at several levels of abstraction. Across these levels, the FB network augments behavior through memory, determining how past behavior influences future behavior. In the sections that follow, we will illustrate how this can be useful.

2.5 Path control at multiple temporal and spatial scales

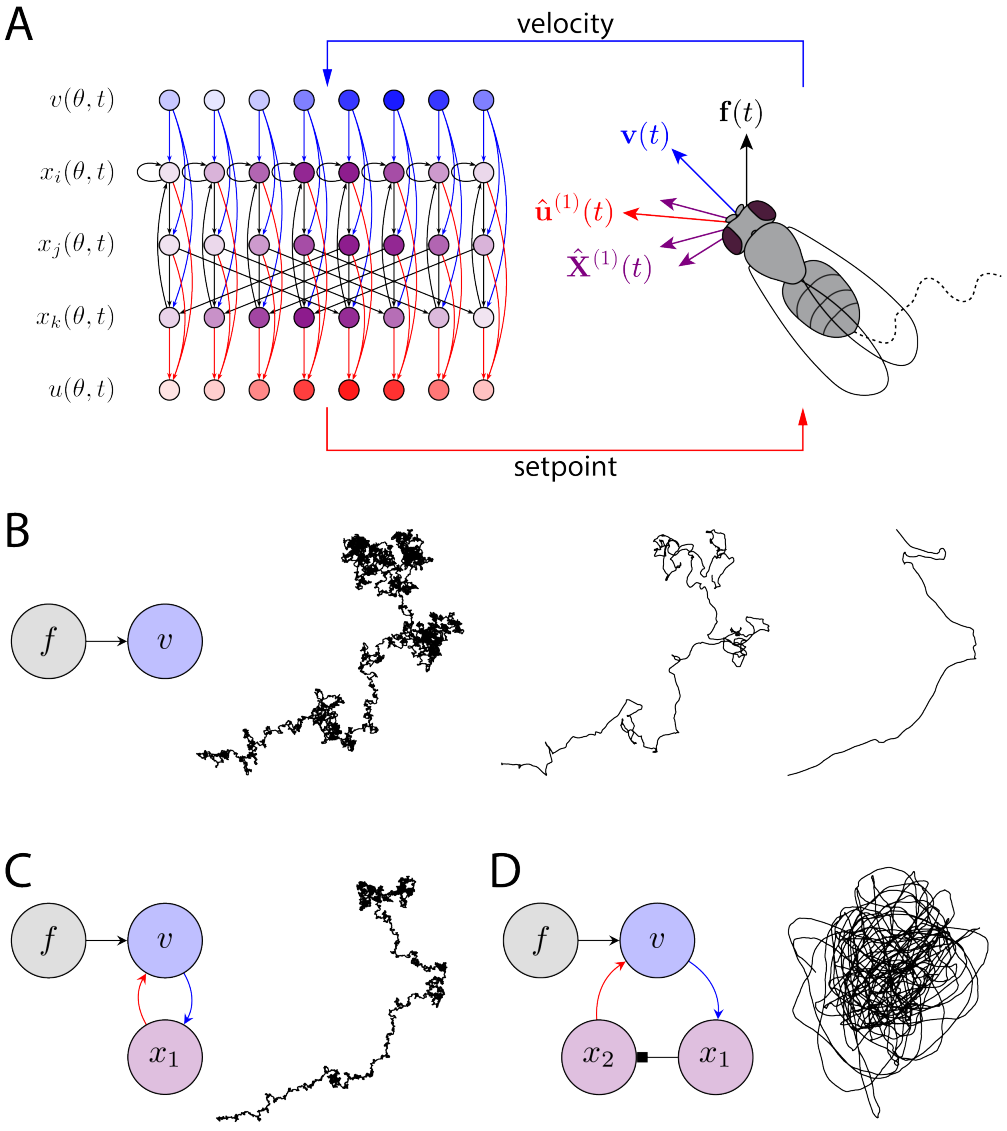


Figure 5: Shaping path structure. (A) The agent's velocity vector and neural activity define a closed-loop system externally driven by the forcing function. (B) An agent with no arrays that is subject to random forcing produces a correlated random walk with a single correlation time scale that depends on γ. Decreasing γ increases path straightness. (C) The addition of one array produces a trajectory with structure on two time scales, resulting in small, diffusive excursions around a mostly straight path. (D) In response to the same forcing function, an agent with a phase-inverted connection produces an oscillatory trajectory. By inverting the phase of the velocity input in this manner, this network's setpoint encourages periodic reversals.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

One potential use for this system is the control of the agent's path structure, an important aspect of search behaviors like foraging. To illustrate this capacity, we will model the forcing function as an unstructured and random source of movement variability (white noise). This stochasticity might represent the random appearance of stimuli along the agent's path, or random noise in the nervous system. Despite this unstructured forcing, the agent's trajectory through the environment will possess a characteristic spatial and temporal structure that depends on the parameters of the closed-loop system. Thus, by adjusting the system's couplings, the agent can tailor the structure of its exploratory behavior to its present needs and environment.

For example, an agent with no neural arrays will produce a trajectory with a single correlation time scale that depends on the drag coefficient γ. As in classic work on random walks [30, 31], varying this correlation time scale allows the agent to interpolate between highly diffusive movement that densely explores the area traversed and ballistic, straight-line movement that more efficiently moves between points (Fig. 5B). When neural arrays are introduced, the agent can generate richer paths that have structure on multiple spatial and temporal scales. For example, an agent with a single array can generate a trajectory that is relatively straight on long time scales but diffusive on short time scales, precisely because memory of recent travel direction allows it to resume course after an excursion (Fig. 5C). Finally, phase shifts enable an agent to generate oscillatory trajectories characterized by regular backtracking, which might be useful for more systematic local searches (Fig. 5D). These differences in path structure and their relation to the closed-loop system's parameters are captured by summary statistics like the velocity autocovariance and mean squared displacement (Appendix E; equations 30, 31).

Note that in addition to altering the shape (e.g. straightness) of the trajectory, varying a given parameter in isolation can also alter the agent's average speed. Independent control of path shape and speed can be achieved by a compensatory rescaling of the forcing function, which can be used to set the length scale of the trajectory while preserving its shape.

2.6 Stabilizing movement towards an intermittent cue

Another potential use for this system is the stabilization of movement guided by an intermittent cue, a previously proposed function of the FB [9, 17]. To illustrate this capacity, we augment the forcing function with a temporally sparse signal that biases the agent towards a fixed cue direction when that signal is present (Appendix F; Fig. 6A,B). For example, this signal might represent intermittent glimpses of a distant visual cue, or intermittent gusts of wind carrying an attractive odor. An agent that cannot stabilize movement in the cue direction will orient towards the signal when it is present but meander during the gaps between signals (Fig. 6C).

In contrast, an agent whose behavior integrates the forcing function over a longer time scale will more persistently move in the cue direction during gaps in the signal (Fig. 6D-I). The specific time scale realized by the agent's behavior will depend on the dynamics of the closed-loop system, and it can be manipulated in two ways. First, strengthening the network's recurrent connections increases the time scale over which the setpoint signal integrates the agent's velocity and, by extension, the forcing function. This enables the agent to accumulate evidence of the cue's salience and amplify biases in the forcing function over longer time scales (Fig. 6D-F). Second, we can increase the gain of the control loop (determined by the strength b of the network's velocity input and the strength c of the network's motor output), which enables the agent to more strongly correct moment-to-moment differences between its velocity and setpoint. Thus, even if the setpoint signal itself only integrates the velocity over a small window, the FB network can increase persistence in the cue direction by more aggressively combating drift in its heading direction (Fig. 6G-I). In either case, by tuning the dynamics of the closed-loop, the agent can balance random exploration and exploitation of the intermittent cue.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

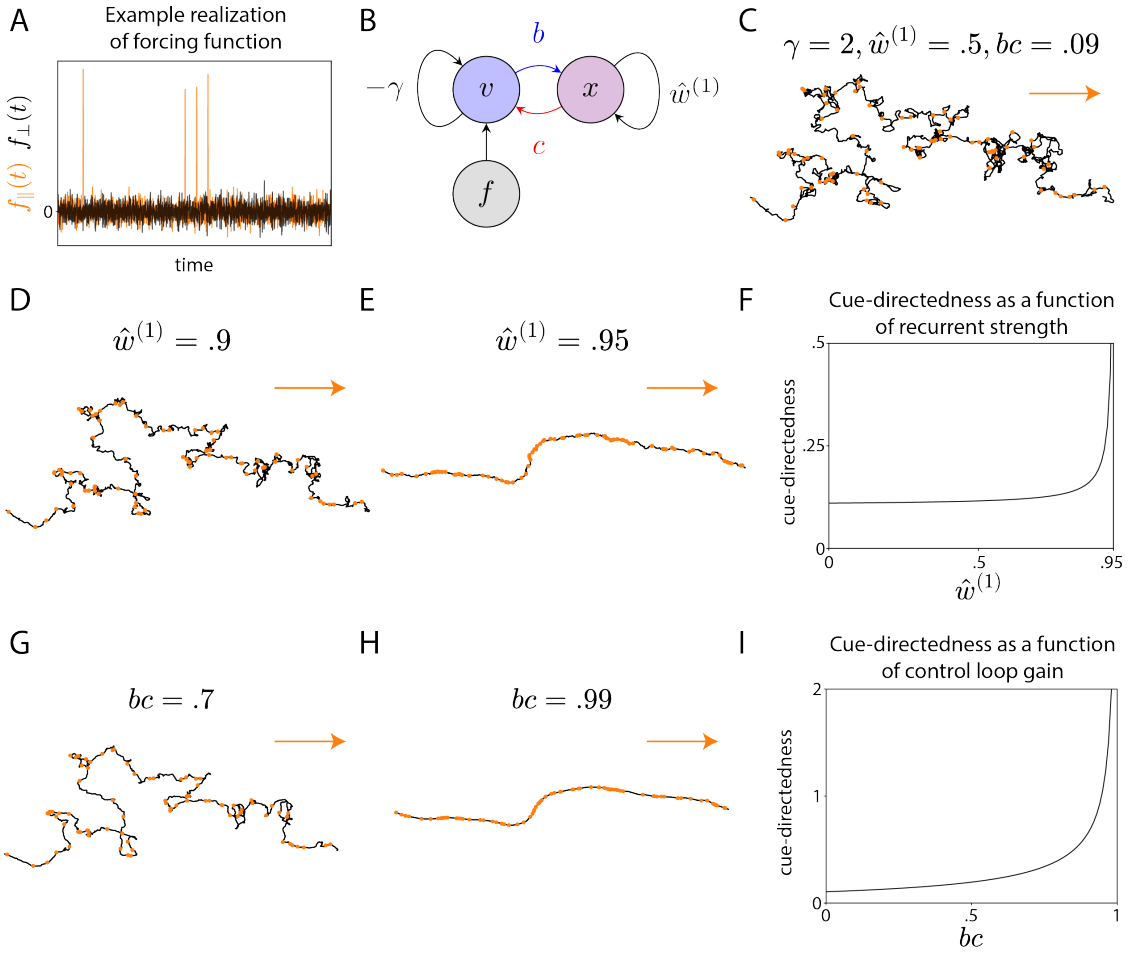


Figure 6: Stabilizing movement towards an intermittent cue. (A) The forcing function is a time-varying, two-dimensional signal that can be decomposed into components aligned (orange) and orthogonal (black) to the cue direction. It consists of a directionally-unbiased component that continuously injects variability into the agent's behavior, and a cue-directed component that is intermittently present. (B) The properties of the closed-loop system determine the trajectory that results from this forcing. In this example, we consider a network with a single array. (C) Example positional trajectory. Orange dots correspond to moments when the intermittent cue is present. This cue transiently biases the agent in the cue direction, which points to the right. Here, the couplings are relatively weak, so the agent's trajectory is meandering. (D,E) Example trajectories with the same forcing function but stronger recurrent connectivity, corresponding to a longer neural integration time scale. (F) The resulting path becomes more cue-directed (ratio of movement aligned and orthogonal to the cue direction; Appendix F) as the neural recurrence increases. (G,H) Example trajectories with higher control loop gain, which represents the strength of the network's velocity input and/or steering output (Appendix F). (I) The path becomes more cue-directed as the loop gain increases.

3 Discussion

In this study, we have shown how an FB-like network can control the spatial and temporal structure of an agent's paths through its environment. This recurrent network temporally filters a velocity vector input, and it uses this filtered vector as the setpoint of a controller that continuously adjusts the agent's velocity. Importantly, the agent's velocity is also influenced by lower-level sensorimotor pathways. These lower-level pathways allow the agent to react to local environmental features and inject variability into its behavior, while the higher-level recurrent network maintains a memory of past movement and governs the overall path structure.

This path control perspective may clarify several otherwise puzzling features of the FB, such as the multitude of pathways downstream of the velocity input, the prevalence of leaky integrator dynamics, and the circular relationship between the setpoint signal and the animal's velocity vector. Our model suggests that the multitude of columnar cell types downstream of the velocity input may enable the FB to extract a variety of temporal or spectral features from the velocity trajectory. Meanwhile, using a combination of

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

these features as a time-varying velocity setpoint may allow the FB to influence path structure at multiple spatial and temporal scales. This work thus builds on previous models of path integration in the FB [6, 8, 9, 14, 32] by exploring circuit implementations of more varied temporal filters.

This work also leverages the idea that sinusoidal activity patterns can represent two-dimensional vectors [1, 14]. In previous work, these activity patterns have encoded variables like an agent's velocity [4, 5], positional displacement [6, 8, 9, 14], and the direction and speed of the wind [33]. In contrast, individual cell types in our model do not explicitly encode a specific environmental variable. Instead, each array simply applies a cell-type-specific linear filter to the velocity input, and this filtering results in activity patterns that help generate behavior. This distinction between representation and generation loosely echoes discussions of neural dynamics in areas like motor cortex [34, 35].

In addition to its first harmonic (sinusoidal component), each activity profile also has a time-varying mean activity level. Within the FB literature, the mean has received less attention, in part because it lacks a clear interpretation that is analogous to the first harmonic's interpretation as a vector. In Appendix B, we analyze the dynamics of mean activity in our model. There, we show that much like the first harmonics temporally filter the agent's velocity vector, the means temporally filter the agent's translational speed. Thus, much like the first harmonic acts as a displacement-like vector variable, the mean acts as a path-length-like scalar variable. We also note that the first harmonic and mean differentially interact with phase-inverting connections, which negatively couple the first harmonics of the activity profiles but positively couple their means. Finally, while the mean plays no role in our model's simple control loop, this simplification is unlikely to hold in practice: optogenetically increasing the mean activity of FB output neurons can influence locomotion [19], and several cell types downstream of the FB are well-poised to read out the mean activity of particular FB columnar cell types [3].

Another idealized feature of the behavioral dynamics (equation 7) we analyzed in the main text was their linearity. An unrealistic consequence of this linearity is that the agent's speed is unbounded. To address this, in Appendix G, we augment the behavioral dynamics with a saturating nonlinearity that prevents the agent's speed from exceeding some maximum value. There, we note that the introduction of this simple nonlinearity can result in an interesting emergent phenomenon: self-sustaining neural activity and behavior. In this regime, the neural circuit's persistent activity is sustained by the sensory input that results from its own motor output. This mirrors, and to some degree contrasts with, mechanisms of persistent activity that instead rely on strong recurrent interactions between neurons [36, 37]. It also functions as a simple example of structure jointly emerging from the dynamics of "the body" (the agent's behavior) and the nervous system, a principle commonly stressed in embodied cognition [38–40].

Regardless of the precise mapping from neural activity to behavior, the general function of this model is to control path structure. The resulting paths can combine persistent, meandering, and oscillatory movement, and their exact structure depends on the model's coupling parameters. Interestingly, these paths are examples of correlated random walks. Correlated random walks are commonly used descriptive models in the fields of animal movement and computational ethology [30, 41, 42]. Here, we illustrate one way that a correlated random walk with a given statistical structure can directly arise from an underlying network model.

Correlated random walks with different structures can be useful in different contexts of foraging and search. According to foraging theory, organisms need to employ a contextually appropriate mix of persistent and meandering movement to search an environment [43, 44]. Meandering allows an organism to densely explore a rich environment, while straight movement allows it to rapidly leave a barren one. The right strategy depends on the organism's environment and internal state. This contextual information may reach the FB via tangential cells [3]. The arbors of individual tangential cells innervate all the columns of a given FB layer. This morphology makes these cells well-poised to modulate columnar activity dynamics and thereby tailor the organism's trajectories to its present circumstances [17, 45].

Of course, organisms that lack a nervous system can still forage. For example, bacteria employ a run-and-tumble foraging strategy, in which straight runs are punctuated by random turns whose frequency depends on environmental cues [46]. Whereas idealized models of run-and-tumble foraging are memoryless,

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

our network model temporally filters the velocity input to maintain a memory of past motion. As a result, this network can influence path structure on a variety of time scales and spatial scales. For example, the path can be straight on long time scales but meandering on short ones, precisely because memory allows it to resume course after an excursion.

Another function of this network is to stabilize useful actions, based on accumulated evidence in favor of those actions [9, 17, 47]. This capacity is useful because natural sensory signals can be discontinuous in time. For example, an attractive visual object can be intermittently visible, or a breeze carrying a food odor can be intermittently sensed. Within this model, orientation towards such cues is initiated by lower-level sensorimotor systems that are independent of the FB-like network. These lower-level systems are represented by the forcing function. When the forcing function exhibits a directional bias due to the intermittent signal, the network can stabilize movement in this direction regardless of the signal's modality or valence. Thus, the system can sustain actions initiated by more "distractable" low-level pathways, a capacity loosely reminiscent of executive function. Interestingly, *Drosophila* will navigate toward remembered cue directions, but genetic disruptions of the compass system disrupt this ability [48]. This dependence on the central complex is consistent with the idea that persistent navigation towards a remembered cue emerges from FB network dynamics.

In many ways, the FB exhibits the rudiments of cognition: it contains allocentric, almost symbolic representations of certain variables; can manipulate these representations using phasor-based arithmetic operations; exhibits working memory; accumulates evidence over time; and enables persistent action that is robust to distractors. Interestingly, analogous high-level functions are sometimes attributed to brain regions like the prefrontal cortex [49]. It is often natural to view such regions as the apex of a control hierarchy, issuing abstract commands that low-level systems translate into motor output. In this closed-loop model, however, the low-level pathways represented by the forcing function are the real engine of behavior, while the high-level network mainly influences the time scales over which the behaviors they initiate persist. Recent findings suggest that the prefrontal cortex may function similarly during self-directed exploratory behavior [50].

Within this model, there is also no direct neural communication between these low-level pathways and the network. Instead, the network simply observes the influence of these pathways on behavior itself, and it bases its setpoint for future behavior on current behavior. The idea that neural controllers operating in parallel might interact through their mutually observed effects on behavior—rather than internally agreeing on some coherent course of action prior to carrying out—has some precedent in cognitive science and robotics [51, 52].

Unlike other recent models of agents controlled by recurrent neural networks [53, 54], we did not train this network to perform a task. Instead, we aimed to understand the neural dynamics and behavior that could emerge from observed patterns of neural connectivity in the insect central complex. Closed-loop models of network-driven behavior also have precedents in the motor control literature. Here, neuromechanical models have explored how biomechanics and neural circuitry jointly produce locomotion [55–58]. Our work extends closed-loop modeling from the sensorimotor periphery (e.g. the spinal cord and brainstem) to a higher-order structure, the central complex. Although the central complex is deep in the brain, the input and output pathways that link its activity to behavior are known explicitly and thus well-constrained.

In our model, the network's structure resulted in a low-dimensional space of neural activity patterns, defined by the first harmonic of each activity profile. The network's sensory input and motor output could both be defined with respect to this space (rather than single neuron activities), and the simplicity of the resulting interactions enabled us to directly characterize the joint system consisting of both neural and behavioral variables. This approach therefore connects models of neural computation based on the population dynamics of recurrent circuits [59, 60] to models of feedback control based on reciprocal interactions between the brain, body, and external world [38–40, 61].

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Acknowledgments

Many thanks to colleagues in the Wilson and Pehlevan Labs for their feedback on this work, and to Caleb Weinreb and Jacob Zavatone-Veth for their comments on an earlier version of this manuscript. During the preparation of this manuscript, M.A.M.O. used GPT-5 to generate code and proofread the text, via a ChatGPT Edu license provided by the Harvard Kenneth C. Griffin Graduate School of Arts and Sciences. C.P. is supported by an NSF CAREER Award (IIS-2239780) and the Simons Collaboration on the Physics of Learning and Neural Computation. This work has been made possible in part by a gift from the Chan Zuckerberg Initiative Foundation to establish the Kempner Institute for the Study of Natural and Artificial Intelligence. This work was also supported by NIH grant R01NS129647 to R.I.W. R.I.W. is an HHMI Investigator.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

References

[1] David S Touretzky, A David Redish, and Hank S Wan. Neural representation of space using sinusoidal arrays. *Neural Computation*, 5(6):869–884, 1993.

[2] John O'Keefe. An allocentric spatial model for the hippocampal cognitive map. *Hippocampus*, 1(3):230–235, 1991.

[3] Brad K Hulse, Hannah Haberkern, Romain Franconville, Daniel Turner-Evans, Shin-ya Takemura, Tanya Wolff, Marcella Noorman, Marisa Dreher, Chuntao Dan, Ruchi Parekh, Ann M Hermundstad, Gerald M Rubin, and Vivek Jayaraman. A connectome of the *Drosophila* central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife*, 10:e66039, oct 2021.

[4] Cheng Lyu, LF Abbott, and Gaby Maimon. Building an allocentric travelling direction signal via vector computation. *Nature*, 601(7891):92–97, 2022.

[5] Jenny Lu, Amir H Behbahani, Lydia Hamburg, Elena A Westeinde, Paul M Dawson, Cheng Lyu, Gaby Maimon, Michael H Dickinson, Shaul Druckmann, and Rachel I Wilson. Transforming representations of movement from body-to world-centric space. *Nature*, 601(7891):98–104, 2022.

[6] Thomas Stone, Barbara Webb, Andrea Adden, Nicolai Ben Weddig, Anna Honkanen, Rachel Templin, William Wcislo, Luca Scimeca, Eric Warrant, and Stanley Heinze. An anatomically constrained model for path integration in the bee brain. *Current Biology*, 27(20):3069–3085, 2017.

[7] Stanley Heinze, Ajay Narendra, and Allen Cheung. Principles of insect path integration. *Current Biology*, 28(17):R1043–R1058, 2018.

[8] Florent Le Moël, Thomas Stone, Mathieu Lihoreau, Antoine Wystrach, and Barbara Webb. The central complex as a potential substrate for vector based navigation. *Frontiers in psychology*, 10:690, 2019.

[9] Roman Goulard, Stanley Heinze, and Barbara Webb. Emergent spatial goals in an integrative model of the insect central complex. *PLoS computational biology*, 19(12):e1011480, 2023.

[10] Ilaria D'Atri and Shamik DasGupta. Neural basis of odometry in drosophila. *bioRxiv*, pages 2025–03, 2025.

[11] Martin Müller and Rüdiger Wehner. Path integration in desert ants, cataglyphis fortis. *Proceedings of the National Academy of Sciences*, 85(14):5287–5290, 1988.

[12] Matthew Collett and Thomas S Collett. How do insects use path integration for their navigation? *Biological cybernetics*, 83(3):245–259, 2000.

[13] Rüdiger Wehner, Mandyam V Srinivasan, et al. Path integration in insects. *The neurobiology of spatial behaviour*, pages 9–30, 2003.

[14] Thomas Wittmann and Helmut Schwegler. Path integration—a network model. *Biological Cybernetics*, 73(6):569–575, 1995.

[15] Peter Mussells Pires, Lingwei Zhang, Victoria Parache, LF Abbott, and Gaby Maimon. Converting an allocentric goal into an egocentric steering signal. *Nature*, 626(8000):808–818, 2024.

[16] Andrew F Siliciano, Sun Minni, Chad Morton, Charles K Dowell, Noelle B Eghbali, Silas E Busch, Juliana Y Rhee, LF Abbott, and Vanessa Ruta. A vector-based strategy for olfactory navigation in drosophila. *Nature*, pages 1–12, 2026.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

[17] Nicholas D Kathman, Aaron J Lanz, Jacob D Freed, and Katherine I Nagel. Neural dynamics for working memory and evidence integration during olfactory navigation in drosophila. *Nature Communications*, 2026.

[18] Wenyi Zhang. *Mechanisms of Setpoint Control in Drosophila Navigation System*. Harvard University, 2026.

[19] Elena A Westeinde, Emily Kellogg, Paul M Dawson, Jenny Lu, Lydia Hamburg, Benjamin Midler, Shaul Druckmann, and Rachel I Wilson. Transforming a head direction signal into a goal-oriented steering command. *Nature*, 626(8000):819–826, 2024.

[20] Chuntao Dan, Brad K Hulse, Ramya Kappagantula, Vivek Jayaraman, and Ann M Hermundstad. A neural circuit architecture for rapid learning in goal-directed navigation. *Neuron*, 112(15):2581–2599, 2024.

[21] Uwe Homberg. Interneurones of the central complex in the bee brain (apis mellifera, l.). *Journal of insect physiology*, 31(3):251–264, 1985.

[22] Ulrike Hanesch, K-F Fischbach, and Martin Heisenberg. Neuronal architecture of the central complex in drosophila melanogaster. *Cell and Tissue Research*, 257(2):343–366, 1989.

[23] Nils Eckstein, Alexander Shakeel Bates, Andrew Champion, Michelle Du, Yijie Yin, Philipp Schlegel, Alicia Kun-Yang Lu, Thomson Rymer, Samantha Finley-May, Tyler Paterson, et al. Neurotransmitter classification from electron microscopy images at synaptic sites in drosophila melanogaster. *Cell*, 187(10):2574–2594, 2024.

[24] Tanya Wolff, Mark Eddison, Nan Chen, Aljoscha Nern, Preeti Sundaramurthi, Divya Sitaraman, and Gerald M Rubin. Cell type-specific driver lines targeting the drosophila central complex and their use to investigate neuropeptide expression and sleep regulation. *Elife*, 14:RP104764, 2025.

[25] Joao P Hespanha et al. *Linear systems theory*, volume 41. Princeton university press Princeton, NJ, 2009.

[26] Karl Johan Åström and Richard Murray. *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2021.

[27] H Sebastian Seung. How the brain keeps the eyes still. *Proceedings of the National Academy of Sciences*, 93(23):13339–13344, 1996.

[28] Mark S Goldman. Memory without feedback in a neural network. *Neuron*, 61(4):621–634, 2009.

[29] Jiali Yu, Zhang Yi, and Jinsong Leng. Line attractors of coupled ring neural networks with block circulant weight matrix. *International Journal of Computer Mathematics*, 97(3):687–699, 2020.

[30] Edward A Codling, Michael J Plank, and Simon Benhamou. Random walk models in biology. *Journal of the Royal society interface*, 5(25):813–834, 2008.

[31] Crispin Gardiner. *Stochastic methods*, volume 4. Springer Berlin Heidelberg, 2009.

[32] Robert John Vickerstaff and EA Di Paolo. Evolving neural models of path integration. *Journal of Experimental Biology*, 208(17):3349–3366, 2005.

[33] Itzel G Ishida, Sachin Sethi, Thomas L Mohren, Mia K Haraguchi, LF Abbott, and Gaby Maimon. Neuronal calcium spikes enable vector inversion in the drosophila brain. *Cell*, 189(3):748–764, 2026.

[34] Krishna V Shenoy, Maneesh Sahani, and Mark M Churchland. Cortical control of arm movements: a dynamical systems perspective. *Annual review of neuroscience*, 36(1):337–359, 2013.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

[35] David G Clark, Blake Bordelon, Jacob A Zavatone-Veth, and Cengiz Pehlevan. Structure, disorder, and dynamics in task-trained recurrent neural circuits. *bioRxiv*, pages 2026–03, 2026.

[36] Rani Ben-Yishai, R Lev Bar-Or, and Haim Sompolinsky. Theory of orientation tuning in visual cortex. *Proceedings of the National Academy of Sciences*, 92(9):3844–3848, 1995.

[37] Haim Sompolinsky, Andrea Crisanti, and Hans-Jurgen Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259, 1988.

[38] Francisco J Varela, Eleanor Rosch, and Evan Thompson. The embodied mind. *The embodied mind: Cognitive science and human experience.*, 1991.

[39] Hillel J Chiel and Randall D Beer. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in neurosciences*, 20(12):553–557, 1997.

[40] Paul Cisek. Beyond the computer metaphor: Behaviour as interaction. *Journal of Consciousness Studies*, 6(11-12):125–142, 1999.

[41] Eliezer Gurarie, Christen H Fleming, William F Fagan, Kristin L Laidre, Jesús Hernández-Pliego, and Otso Ovaskainen. Correlated velocity models as a fundamental unit of animal movement: synthesis and applications. *Movement ecology*, 5(1):13, 2017.

[42] William Bialek. On the dimensionality of behavior. *Proceedings of the National Academy of Sciences*, 119(18):e2021860119, 2022.

[43] David W Stephens and John R Krebs. *Foraging theory*, volume 6. Princeton university press, 1986.

[44] William J Bell. *Searching behaviour: the behavioural ecology of finding resources*. Springer Science & Business Media, 1990.

[45] Aaron J Lanz, Nick D Kathman, Emily Hao, Bard Ermentrout, and Katherine I Nagel. Disinhibition of a recurrent attractor gates a persistent goal signal for navigation. *bioRxiv*, pages 2025–10, 2025.

[46] Howard C Berg. *Random walks in biology*. Princeton University Press, Princeton, N.J, 1983.

[47] Joshua I Gold and Michael N Shadlen. The neural basis of decision making. *Annu. Rev. Neurosci.*, 30(1):535–574, 2007.

[48] Kirsa Neuser, Tilman Triphan, Markus Mronz, Burkhard Poeck, and Roland Strauss. Analysis of a spatial orientation memory in drosophila. *Nature*, 453(7199):1244–1247, 2008.

[49] Earl K Miller and Jonathan D Cohen. An integrative theory of prefrontal cortex function. *Annual review of neuroscience*, 24(1):167–202, 2001.

[50] Caleb Weinreb, Lakshanyaa Thamarai Kannan, Alia Newman-Boulle, Tim Sainburg, Winthrop F Gillis, Alex Plotnikoff, Sofia Makowska, Jonah E Pearl, Mohammed Abdal Monium Osman, Scott W Linderman, et al. Spontaneous behavior is a succession of self-directed tasks. *Neuron*, 2026.

[51] Daniel C. Dennett. *Consciousness Explained*. Penguin Books, 1991.

[52] Rodney A Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.

[53] Josh Merel, Diego Aldarondo, Jesse Marshall, Yuval Tassa, Greg Wayne, et al. Deep neuroethology of a virtual rodent. *arXiv preprint arXiv:1911.09451*, 2019.

[54] Satpreet H Singh, Floris Van Breugel, Rajesh PN Rao, and Bingni W Brunton. Emergent behaviour and neural dynamics in artificial agents tracking odour plumes. *Nature machine intelligence*, 5(1):58–70, 2023.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

[55] Örjan Ekeberg. A combined neuronal and mechanical model of fish swimming. *Biological cybernetics*, 69(5):363–374, 1993.

[56] Gentaro Taga. A model of the neuro-musculo-skeletal system for human locomotion: I. emergence of basic gait. *Biological cybernetics*, 73(2):97–111, 1995.

[57] R Kukillaya, J Proctor, and P Holmes. Neuromechanical models for insect locomotion: Stability, maneuverability, and proprioceptive feedback. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 19(2), 2009.

[58] Cengiz Pehlevan, Paolo Paoletti, and L Mahadevan. Integrative neuromechanics of crawling in d. melanogaster larvae. *Elife*, 5:e11031, 2016.

[59] David Sussillo. Neural circuits as computational dynamical systems. *Current opinion in neurobiology*, 25:156–163, 2014.

[60] Saurabh Vyas, Matthew D Golub, David Sussillo, and Krishna V Shenoy. Computation through neural population dynamics. *Annual review of neuroscience*, 43(1):249–275, 2020.

[61] Norbert Wiener. *Cybernetics or Control and Communication in the Animal and the Machine*. The MIT Press, 1948.

[62] Hugh R Wilson and Jack D Cowan. A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2):55–80, 1973.

[63] Shun-ichi Amari. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biological cybernetics*, 27(2):77–87, 1977.

[64] Bard Ermentrout. Neural networks as spatio-temporal pattern-forming systems. *Reports on progress in physics*, 61(4):353–430, 1998.

[65] Paul C Bressloff and Jack D Cowan. An amplitude equation approach to contextual effects in visual cortex. *Neural computation*, 14(3):493–525, 2002.

[66] Haozhe Shan and Ashok Litwin-Kumar. Graph embeddings for identifying symmetries in connectomes. *bioRxiv*, 2025.

[67] Surya Ganguli, Dongsung Huh, and Haim Sompolinsky. Memory traces in dynamical systems. *Proceedings of the national academy of sciences*, 105(48):18970–18975, 2008.

[68] Yoshiki Kuramoto. Self-entrainment of a population of coupled non-linear oscillators. In *International symposium on mathematical problems in theoretical physics*, pages 420–422, 1975.

[69] Stephen O Rice. Mathematical analysis of random noise. *The Bell System Technical Journal*, 23(3):282–332, 1944.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Appendix

A Dynamics of the Fourier coefficients

In this section, we derive the governing equations for the Fourier coefficients (first harmonic and mean) of the activity profiles. The model in the main text is an example of a spatially organized network with multiple cell types and translationally invariant connections [62–64]. This connectivity is organized around an angular variable, more specifically making the model a ring network [29, 36, 65, 66]. In the linear case, the dynamics of such systems can be expressed as follows:

$$\tau \frac{\partial}{\partial t} x_j(\theta, t) = -x_j(\theta, t) + \sum_{k=1}^{N} \int_0^{2\pi} d\theta' w_{jk}(\theta - \theta') x_k(\theta', t) + b_j v(\theta, t). \tag{11}$$

Here, $v(\theta, t)$ is some external input and $w_{jk}(\cdot)$ is a connectivity kernel (periodic over 2π) that relates the strength of a connection to the cell types and relative positions of the two neurons. For simplicity, we model the network as a continuum, though translating to the discrete case is straightforward (see [29]).

As is standard in the study of such networks, we can use the Fourier transform to express the shape of each population's activity profile in the frequency domain. Let $\hat{z}_j^{(\ell)} \in \mathbb{C}$ denote the ℓth (with $\ell \in \mathbb{N}$) unnormalized Fourier coefficient of the jth activity profile:

$$\hat{z}_j^{(\ell)}(t) = \int_0^{2\pi} d\theta \, e^{i\ell\theta} x_j(\theta, t), \tag{12}$$

where $i = \sqrt{-1}$ denotes the imaginary unit (used elsewhere in this manuscript as a discrete indexing variable). Similarly, let $\hat{w}_{jk}^{(\ell)}$ and $\hat{v}^{(\ell)}(t)$ denote the ℓth Fourier coefficients of the connection kernels and input pattern. The dynamics of $\hat{z}_j^{(\ell)}$ follow

$$\tau \frac{d}{dt} \hat{z}_j^{(\ell)}(t) = -\hat{z}_j^{(\ell)}(t) + \sum_{k=1}^{N} \hat{w}_{jk}^{(\ell)} \hat{z}_k^{(\ell)}(t) + b_j \hat{v}^{(\ell)}(t). \tag{13}$$

Note that $\hat{z}_j^{(\ell)}(t)$ interacts with the ℓth Fourier coefficients of the other activity profiles and the input but is completely uncoupled from Fourier coefficients corresponding to other frequencies ($m \neq \ell$). Each Fourier mode thus defines an independent dynamical system driven by the corresponding Fourier coefficient of the input and coupled through the corresponding Fourier coefficients of the connectivity kernels:

$$\tau \frac{d}{dt} \hat{\mathbf{z}}^{(\ell)}(t) = -\hat{\mathbf{z}}^{(\ell)}(t) + \hat{\mathbf{W}}^{(\ell)} \hat{\mathbf{z}}^{(\ell)}(t) + \hat{v}^{(\ell)}(t)\mathbf{b}, \tag{14}$$

where $\hat{\mathbf{z}}^{(\ell)}(t) \in \mathbb{C}^N$ is the population vector with entries $\hat{z}_j^{(\ell)}$ and $\hat{\mathbf{W}}^{(\ell)} \in \mathbb{C}^{N \times N}$ is the coupling matrix with entries $\hat{w}_{jk}^{(\ell)}$ (see [29] and [64] for analogous modal reductions).

Because the external input in the main text is a simple sinusoid lacking any high frequency content ($\hat{v}^{(\ell)}(t) = 0$ for all $\ell > 1$), its activity profiles will not exhibit high frequency content unless it is present at initialization. In addition, assuming a given mode's dynamics are stable ($\operatorname{Re} \lambda_j \left(\hat{\mathbf{W}}^{(\ell)} \right) < 1$ for all j), any high frequency content present at initialization will dissipate. In any case, we assume an initial state of zero ($x_j(\theta, 0) = 0$) throughout, thus guaranteeing that the activity profiles are always simple sinusoids and reducing the dynamics to modes 0 and 1. In the main text, for clarity, we expressed these modes using real (rather than complex) numbers and rescaled them so that their magnitudes directly equal the associated features (mean and amplitude) of the activity profile:

$$\mu_j(t) = \frac{1}{2\pi} \hat{z}_j^{(0)}(t) \in \mathbb{R}, \tag{15}$$

$$\hat{\mathbf{x}}_j^{(1)}(t) = \frac{1}{\pi} \begin{bmatrix} \operatorname{Re} \hat{z}_j^{(1)}(t) \\ \operatorname{Im} \hat{z}_j^{(1)}(t) \end{bmatrix} \in \mathbb{R}^2. \tag{16}$$

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Finally, to obtain the governing equations for these variables, we consider the specific connectivity kernel used in the main text. This connectivity kernel is a weighted sum of two delta functions, respectively corresponding to the phase-preserving and phase-inverting connections:

$$w_{jk}(\theta) = W^{+}_{jk}\delta_{2\pi}(\theta) + W^{-}_{jk}\delta_{2\pi}(\theta - \pi). \tag{17}$$

The Fourier coefficients of this kernel depend on the parity of the mode:

$$\hat{w}^{(\ell)}_{jk} = W^{+}_{jk} + (-1)^{\ell} W^{-}_{jk}.$$

This parity dependence reflects the fact that a phase shift of π leaves modes of even wavenumber unchanged while multiplying modes of odd wavenumber by -1. Consequently, phase-inverting connections act as negative, "effectively inhibitory" couplings among the first (and other odd-numbered) harmonics of the activity profiles, even though all the units in the network are excitatory.

B Dynamics of the mean

The activity of each neural array in our model can be decomposed into two components: the first harmonic and the mean. In this section, we consider the dynamics of the latter. Let $\boldsymbol{\mu} \in \mathbb{R}^N$ denote the mean activity in each array. The dynamics of $\boldsymbol{\mu}$ follow

$$\tau \frac{d}{dt}\boldsymbol{\mu}(t) = -\boldsymbol{\mu}(t) + \left(\mathbf{W}^{+} + \mathbf{W}^{-}\right)\boldsymbol{\mu}(t) + \mathbf{b}\|\mathbf{v}(t)\|. \tag{18}$$

Whereas the first harmonic of the network's input encodes the agent's velocity vector, the mean of the input encodes the agent's translational speed (due to the non-zero offset in equation 1). As a result, whereas the first harmonic computes a displacement-like quantity by temporally filtering its velocity vector input, the mean computes a path-length-like quantity by temporally filtering its speed input. In addition, note that phase-inverting connections positively couple the means but negatively couple the first harmonics. Due to this coupling difference, the two components will generally apply different filters to their respective inputs when the network contains phase-inverting connections.

Finally, since the network's connections and input are non-negative, single-neuron activity must also be non-negative. As a result, the mean of each array's activity profile must always equal or exceed the amplitude of its first harmonic. Since the mean bounds the magnitude of the first harmonic, the stability of the network solely depends on the stability of the mean dynamics, with stability condition $\lambda_{\max}(\mathbf{W}^{+} + \mathbf{W}^{-}) < 1$.

C Linear filtering

Many of the results in the main text rely on the fact that time-invariant linear dynamical systems act as temporal filters of their inputs. In this section, we review the relevant mathematics, which constitute standard linear systems theory (see e.g. [25, 26]). Consider a system of the form

$$\tau \frac{d}{dt}\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) + u(t)\mathbf{b}, \tag{19}$$

$$y(t) = \mathbf{c}^{\top}\mathbf{x}(t). \tag{20}$$

Here $\mathbf{x}(t) \in \mathbb{R}^N$ is the system state, $u(t) \in \mathbb{R}$ is some time-varying input to the system, $y(t) \in \mathbb{R}$ is some observed variable, $\mathbf{A} \in \mathbb{R}^{N\times N}$ is a matrix of recurrent couplings, $\mathbf{b} \in \mathbb{R}^N$ is a vector of input weights, and $\mathbf{c} \in \mathbb{R}^N$ is a vector of the readout weights. Again, we will assume an initial condition of $\mathbf{x}(0) = \mathbf{0}$.

We can express $y(t)$ as some convolution of $u(t)$ with a kernel that depends on the system's coupling parameters:

$$y(t) = (\mathcal{H}[\mathbf{A}, \mathbf{b}, \mathbf{c}, \tau] * u)(t) = \int_0^t dt'\, \mathcal{H}[\mathbf{A}, \mathbf{b}, \mathbf{c}, \tau](t - t')u(t').$$

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

Here, $\mathcal{H}$ is a functional that takes in the system parameters and returns the associated temporal kernel:

$$\mathcal{H}[\mathbf{A}, \mathbf{b}, \mathbf{c}, \tau](t) = \mathbb{1}\{t > 0\} \cdot \frac{1}{\tau} \mathbf{c}^\top \exp\left(\frac{t}{\tau}\mathbf{A}\right) \mathbf{b}, \tag{21}$$

where the coefficient $\mathbb{1}\{t > 0\}$ ensures that the kernel is causal (equals 0 for $t < 0$). In general, such a filter can be expressed as a sum of up to N terms, each consisting of a polynomial times an exponential. Thus, for $t > 0$:

$$\mathcal{H}[\mathbf{A}, \mathbf{b}, \mathbf{c}, \tau](t) = \frac{1}{\tau} \sum_i P_i\left(\frac{t}{\tau}\right) \exp\left(\lambda_i(\mathbf{A})\frac{t}{\tau}\right), \tag{22}$$

where $P_i\left(\frac{t}{\tau}\right)$ is a polynomial function of time and $\lambda_i(\mathbf{A}) \in \mathbb{C}$ denotes the ith eigenvalue of $\mathbf{A}$.

If the matrix $\mathbf{A}$ is diagonalizable, then each $P_i\left(\frac{t}{\tau}\right) = p_i$ will be a constant and the filter will simply be a weighted sum of exponentials. Non-diagonalizable systems with nontrivial polynomial terms can exhibit interesting filtering properties that have been explored elsewhere [28, 67].

The behavior of exponential terms will in turn depend on the eigenvalues of $\mathbf{A}$. Real-valued eigenvalues will contribute exponential decay terms that (leakily) integrate the input with an effective time constant of $\tau_i = -\frac{\tau}{\mathrm{Re}\,\lambda_i(\mathbf{A})}$. Complex-valued modes will come in conjugate pairs, resulting in exponentially damped (or driven) oscillatory terms of the form

$$\exp\left(-\frac{t}{\tau_i}\right) \cos\left(2\pi\omega_i \frac{t}{\tau} + \phi_i\right), \tag{23}$$

where $\omega_i \in \mathbb{R}$ determines the oscillation's frequency and $\phi_i \in [0, 2\pi)$ determines its phase. These oscillatory terms can realize more complex (e.g. band-pass) filtering properties.

With this background in mind, we can summarize many of the variables in the main text as temporally filtered responses to external inputs:

Variable	Filter	Input
$\hat{\mathbf{x}}_i^{(1)}(t)$	$\mathcal{H}[\hat{\mathbf{W}}^{(1)} - \mathbf{I}, \mathbf{b}, \mathbf{e}_i, \tau]$	$\mathbf{v}(t)$
$\mu_i(t)$	$\mathcal{H}[\mathbf{W}^+ + \mathbf{W}^- - \mathbf{I}, \mathbf{b}, \mathbf{e}_i, \tau]$	$\lVert\mathbf{v}(t)\rVert$
$k_U \hat{\mathbf{u}}^{(1)}(t)$	$\mathcal{H}[\hat{\mathbf{W}}^{(1)} - \mathbf{I}, \mathbf{b}, \mathbf{c}, \tau]$	$\mathbf{v}(t)$
$\mathbf{v}(t)$	$\mathcal{H}[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau]$	$\mathbf{f}(t)$

Above, $\mathbf{e}_i$ denotes the ith standard basis vector (the vector with entries δ_{ij}).

D Polar coordinate description

In this section, we reexpress the dynamics of the closed-loop system in polar coordinates. This change of coordinates yields governing equations for the agent's heading and forward speed, thereby making connections to the literature on central complex output circuits more apparent. It also yields governing equations for the activity profile phases and amplitudes, thereby providing a different perspective on the neural dynamics.

To start, we will define polar representations for each of the variables. Let $\hat{\mathbf{x}}_i^{(1)} \leftrightarrow (X_i, \theta_i)$ where X_i denotes the amplitude of the ith activity profile and θ_i denotes its phase; $\hat{\mathbf{u}}^{(1)} \leftrightarrow (U, \psi)$, where U denotes the amplitude of the setpoint signal and ψ denotes its phase; and $\mathbf{v} \leftrightarrow (V, \phi)$, where V denotes the agent's forward speed and ϕ denotes its heading. Note that "heading" simply refers to the direction of the velocity vector in this model. Since the agent is modeled as a particle without an intrinsic axis of orientation, the model does not distinguish between heading and direction of travel.

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

When equation 7 is rewritten in polar coordinates, the resulting heading dynamics recapitulate the basic features of the steering control loop implemented by PFL3 cells [6, 15, 19]:

$$\begin{aligned}\tau V(t)\frac{d}{dt}\phi(t) &= -k_U U(t)\sin\left(\phi(t)-\psi(t)\right)+f_\phi(t)\\ &= -\sum_{i=1}^{N} c_i X_i(t)\sin\left(\phi(t)-\theta_i(t)\right)+f_\phi(t).\end{aligned} \tag{24}$$

This equation continuously compares the agent's current heading $\phi(t)$ to the setpoint phase $\psi(t)$, yielding an "alignment error" $\sin(\phi(t)-\psi(t))$. The resulting heading update is then proportional to the negative error, which causes the agent to turn towards the setpoint phase. In our model, the gain of this turning drive scales with $U(t)$, the amplitude of the setpoint signal. Additionally, as the second line of equation 24 shows, we can also decompose the full turning output into a weighted sum of analogous contributions from the N arrays. From this latter perspective, each array steers the agent's heading towards its activity phase $\theta_j(t)$. The strength of this pull depends on the amplitude $X_j(t)$ and the output weight c_j. In our model, turning gain is also inversely proportional to forward speed $V(t)$. Finally, $f_\phi(t)\in\mathbb{R}$ denotes the effect of the forcing function orthogonal to the current heading.

The radial component of the polar dynamics relates the agent's forward speed to the degree of alignment between the heading and setpoint, as well as the amplitude of the setpoint signal. There are populations of output neurons—namely, PFL2 cells [19]—well-poised to implement such a speed controller. More formally,

$$\begin{aligned}\tau\frac{d}{dt}V(t) &= -\gamma V(t)+k_U U(t)\cos(\phi(t)-\psi(t))+f_V(t)\\ &= -\gamma V(t)+\sum_{i=1}^{N} c_i X_i(t)\cos(\phi(t)-\theta_i(t))+f_V(t).\end{aligned} \tag{25}$$

This equation consists of a drag-like term that drives forward speed to 0 in the absence of input. The neural drive that combats this tendency is proportional to $\cos(\psi(t)-\phi(t))\in[-1,1]$, which quantifies the degree of alignment between the setpoint signal and heading. Ignoring the forcing function and treating the other state variables as fixed, this equation will drive forward speed to a maximal value $\frac{k_U U(t)}{\gamma}$ when heading and setpoint are perfectly aligned, a speed of 0 when the two are orthogonal, and a speed of $-\frac{k_U U(t)}{\gamma}$ when the two are oppositely oriented. We can view moments of transition from positive to negative V as instantaneous inversions of the heading rather than backwards movement, since the points $(-V,\phi)$ and $(+V,\phi+\pi)$ are mathematically equivalent in polar coordinates. Like the turning dynamics, the forward speed dynamics can also be decomposed into a linear combination of analogous per-array contributions (second line of equation 25). In addition, $f_V(t)\in\mathbb{R}$ denotes the effect of the forcing function aligned with the current heading.

Finally, to understand the joint dynamics of neural activity and behavior, we rewrite the neural dynamics along the first harmonic in polar coordinates. The phase dynamics are given by

$$\tau X_i(t)\frac{d}{dt}\theta_i(t) = -\sum_{j=1}^{N}(W_{ij}^{+}-W_{ij}^{-})X_j(t)\sin(\theta_i(t)-\theta_j(t))-b_i V(t)\sin(\theta_i(t)-\phi(t)), \tag{26}$$

which is analogous to the governing equation that relates neural activity to changes in heading. Similarly, the amplitude dynamics are given by

$$\tau\frac{d}{dt}X_i(t) = -X_i(t)+\sum_{j=1}^{N}(W_{ij}^{+}-W_{ij}^{-})X_j(t)\cos(\theta_i(t)-\theta_j(t))+b_i V(t)\cos(\theta_i(t)-\phi(t)), \tag{27}$$

which is analogous to the equation relating neural activity to changes in forward speed. In summary, the phases of the neural activity bumps and the agent's heading define a set of angular variables coupled by alignment (and amplitude-dependent) interactions, somewhat akin to the synchronizing forces in a Kuramoto model [68].

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

E Statistical structure of random walks

To illustrate how the network's couplings can shape the agent's trajectories through its environment, we model the forcing function as a minimally structured source of random variability:

$$\mathbf{f}(t) = \sigma \boldsymbol{\xi}(t), \tag{28}$$

where $\boldsymbol{\xi}(t) \in \mathbb{R}^2$ is a white noise process with $\langle \boldsymbol{\xi}(t) \rangle = 0$ and $\langle \boldsymbol{\xi}(t) \boldsymbol{\xi}(t')^\top \rangle = \delta(t - t')\mathbf{I}$, where $\langle \cdot \rangle$ denotes an expectation over noise realizations. In addition, $\sigma \in \mathbb{R}_+$ is a constant that sets the noise strength. With this stochastic forcing, the agent will perform a correlated random walk through its environment [30, 31].

The overall trajectory will reflect a weighted sum of the $N + 1$ modes present in the joint system. For concreteness, we will assume the closed-loop coupling matrix $\mathbf{A}$ is diagonalizable and its eigenvalues are real, so that the temporal filter realized by the agent's velocity is a sum of exponentials of the form

$$\mathcal{H}\left[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau\right](t) = \frac{1}{\tau} \sum_{i=1}^{N+1} a_i \exp\left(-\frac{t}{\tau_i}\right), \tag{29}$$

where $a_i \in \mathbb{R}$ and $\tau_i = -\frac{\tau}{\lambda_i(\mathbf{A})} \in \mathbb{R}$ are the coefficient and effective time constant associated with the ith mode of the filter. We further assume that the joint system is stable ($\lambda_i(\mathbf{A}) < 0$ for all i).

With such a filter, the random walk will exhibit a multiplicity of time scales that are evident in its autocovariance. The autocovariance relates the expected similarity (inner product) of the velocity vector at two time points to the duration $\Delta \geq 0$ separating them. In this case, the autocovariance of the agent's random walk will contain up to $N + 1$ modes that are inherited from the filter:

$$C(\Delta) = \lim_{t \to \infty} \langle \mathbf{v}(t)^\top \mathbf{v}(t + \Delta) \rangle = \frac{2\sigma^2}{\tau^2} \sum_{i=1}^{N+1} \left(\sum_{j=1}^{N+1} \frac{a_j \tau_j}{\tau_i + \tau_j} \right) a_i \tau_i \exp\left(-\frac{\Delta}{\tau_i}\right). \tag{30}$$

Each term in this weighted sum contributes a correlation time scale of τ_i. Terms with positive coefficients promote persistent movement over lags $\Delta < \tau_i$, while terms with negative coefficients promote reversals over the same time scale. The overall path structure will reflect the combined contributions of each mode.

The filter's time scales are also present in the agent's mean squared displacement (MSD), which characterizes how the agent's distance from the starting point of an interval scales with the interval's duration:

$$\mathrm{MSD}(\Delta) = \lim_{t \to \infty} \left\langle \left\| \int_t^{t+\Delta} dt' \, \mathbf{v}(t') \right\|^2 \right\rangle = \frac{4\sigma^2}{\tau^2} \sum_{i=1}^{N+1} \left(\sum_{j=1}^{N+1} \frac{a_j \tau_j}{\tau_i + \tau_j} \right) a_i \tau_i^2 \left(\Delta - \tau_i \left(1 - \exp\left(-\frac{\Delta}{\tau_i}\right) \right) \right). \tag{31}$$

Finally, if the joint system has complex modes, the filter will contain exponentially damped oscillatory terms. Such modes will encourage periodic backtracking in the agent's trajectory and contribute oscillatory terms with the same frequency and decay rate to the autocovariance.

F Navigation towards an intermittent signal

To model navigation towards an intermittent sensory cue, we augment the forcing function with a temporally intermittent signal—mathematically equivalent to "shot noise" [31, 69]—that biases the agent towards some cue direction $\mathbf{z} \in \mathbb{R}^2$ (with $\|\mathbf{z}\| = 1$):

$$\mathbf{f}(t) = \sigma \boldsymbol{\xi}(t) + J \sum_i \delta(t - t_i)\mathbf{z}. \tag{32}$$

In this equation, $\{t_i\}$ defines a set of observation times sampled from a Poisson process with rate parameter ν, and $J \in \mathbb{R}_+$ parameterizes the strength of each signal's immediate effect on behavior. To characterize the agent's propensity to move in the cue direction, we decompose its velocity into components aligned

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

with and orthogonal to $\mathbf{z}$, denoted $v_{\parallel}(t)$ and $v_{\perp}(t)$. Let $\bar{v}_{\parallel} = \lim_{t\to\infty}\langle v_{\parallel}(t)\rangle$ denote the mean cue-aligned velocity and $\sigma_{\perp} = \lim_{t\to\infty}\sqrt{\langle v_{\perp}(t)^2\rangle}$ denote the standard deviation of the cue-orthogonal velocity. The ratio $\frac{\bar{v}_{\parallel}}{\sigma_{\perp}}$ is then a natural measure of cue-directedness, somewhat akin to a signal-to-noise ratio.

For a system with closed-loop coupling matrix $\mathbf{A}$, the cue-directedness is given by

$$\frac{\bar{v}_{\parallel}}{\sigma_{\perp}} = \frac{J\nu}{\sigma}\frac{\int_0^{\infty} dt\, \mathcal{H}\left[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau\right](t)}{\sqrt{\int_0^{\infty} dt\, \mathcal{H}\left[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau\right](t)^2}}, \tag{33}$$

where $\mathcal{H}\left[\mathbf{A}, \mathbf{e}_1, \mathbf{e}_1, \tau\right]$ is the kernel with which the agent's velocity filters the forcing function (equation 10). This equation consists of two factors. The first factor captures the contribution of the forcing function itself. It is directly proportional to the cue frequency and strength and inversely proportional to the magnitude of the isotropic noise. The second factor captures the closed-loop interactions between the network and behavior. Its numerator corresponds to the DC gain with which the system amplifies signal in the forcing function, while its denominator captures the variance generated by filtering the white noise.

In Figure 6, to explore how the network's dynamics can influence the agent's behavior, we varied the parameters in a network consisting of a single array ($N = 1$) and measured the cue-directedness of the resulting trajectories. More formally,

$$\mathbf{A} = \begin{bmatrix} -\gamma & c \\ b & \hat{w}^{(1)} - 1 \end{bmatrix}, \tag{34}$$

where $\gamma, b, c > 0$, $\hat{w}^{(1)} < 1$, and $\gamma(1 - \hat{w}^{(1)}) > bc$ (which ensures stability). We then showed that given the same forcing, the cue-directedness of the resulting path increases as $\hat{w}^{(1)}$ (which sets the neural integration time scale) and bc (which sets the gain of the control loop) increase. Thus, varying both the integration time scale of the network and the gain of the closed loop can enable the agent to exploit this intermittent source of information and stabilize coherent movement in the cue direction.

G Dynamics with velocity nonlinearities

One unrealistic feature of the linear velocity dynamics defined by equation 7 is that the agent's speed is unbounded. This can be addressed by augmenting the velocity dynamics with a saturating nonlinearity that ensures the agent stays below some maximum speed:

$$\tau\frac{d}{dt}\mathbf{v}(t) = -\gamma\mathbf{v}(t) + \mathbf{g}(k_U\hat{\mathbf{u}}^{(1)}(t) + \mathbf{f}(t)). \tag{35}$$

Here $\mathbf{g} : \mathbb{R}^2 \to \mathbb{R}^2$ is a vector-valued function that rescales its input vector by applying a saturating (tanh) nonlinearity to its magnitude while preserving its direction:

$$\mathbf{g}(\mathbf{x}) = \begin{cases} \dfrac{\tanh(\|\mathbf{x}\|)}{\|\mathbf{x}\|}\mathbf{x} & \mathbf{x} \neq \mathbf{0} \\ \mathbf{0} & \mathbf{x} = \mathbf{0}. \end{cases} \tag{36}$$

Note that this nonlinearity solely alters the dynamics of the velocity vector. The neural dynamics remain linear, following equations 2 and 4. When this speed nonlinearity is introduced, the basic phenomena noted in the preceding sections still hold: increasing the strength of the neural recurrence and the gain of the control loop both increase the straightness of the simulated trajectories.

However, new phenomena also emerge. For example, in the absence of forcing ($\mathbf{f}(t) = \mathbf{0}$), a linear agent (with stable closed-loop couplings) will remain stationary, since its dynamics are dissipative and its behavior is input-driven. The same is true of a nonlinear agent with weak couplings. However, as the nonlinear agent's coupling strengths increase, it can undergo a transition from this quiescent state to an active state in which neural activity and behavior are self-sustaining. This transition can be observed in the single array setup defined by equation 34. To analyze it, we can linearize the dynamics of the

bioRxiv preprint doi: https://doi.org/10.64898/2026.10.02.756268; this version posted October 9, 2026. The copyright holder for this preprint (which was not certified by peer review) is the author/funder, who has granted bioRxiv a license to display the preprint in perpetuity. It is made available under a CC-BY 4.0 International license.

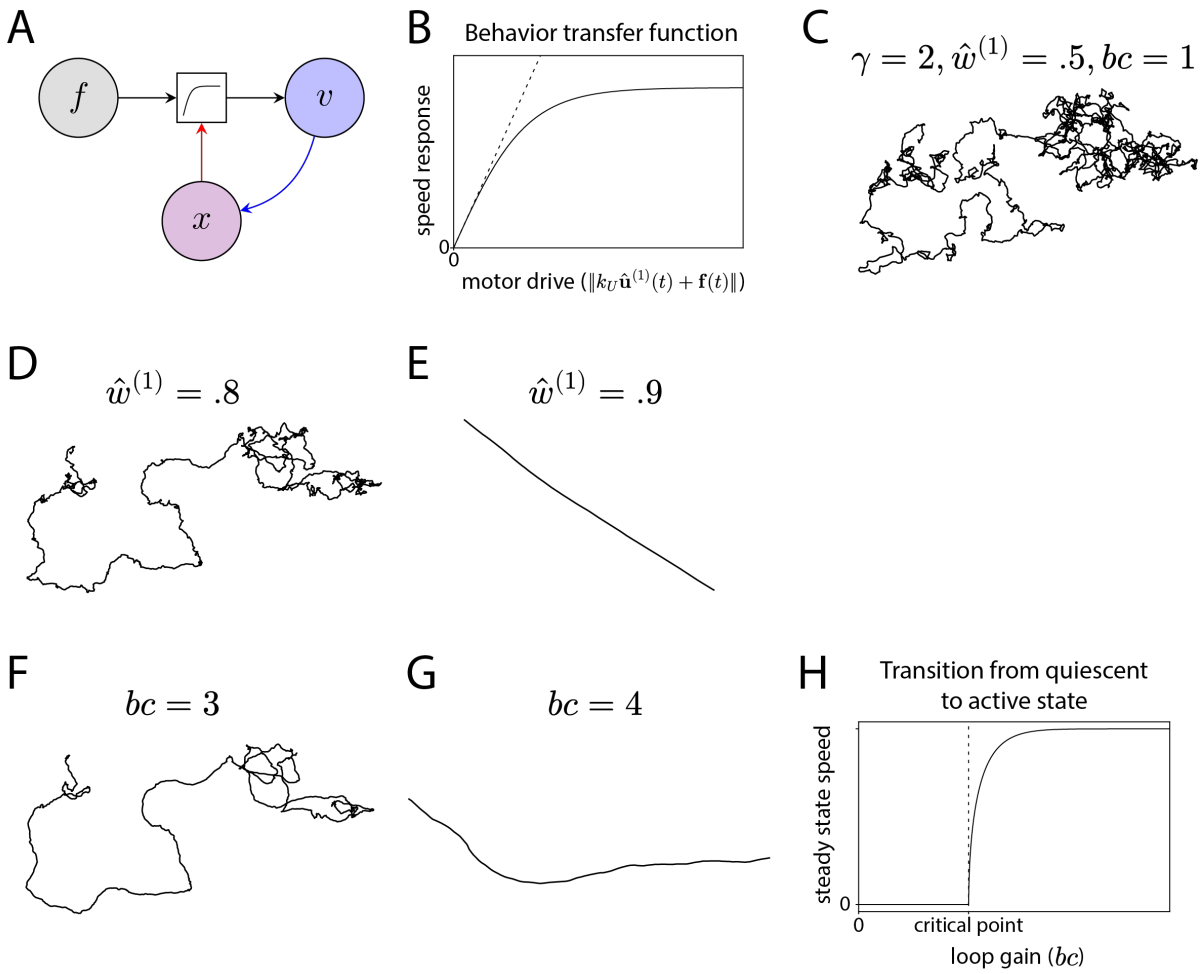


Figure S1: Agent behavior with a speed nonlinearity (A) The motor drive that the agent inherits from the network and forcing function is now mediated by a nonlinear transfer function. (B) This transfer function is a saturating nonlinearity (hyperbolic tangent) that ensures that the agent's speed remains bounded as the magnitude of the motor drive grows. (C) With this nonlinearity, as before, an agent with relatively weak coupling parameters produces a relatively meandering path. (D,E) In response to the same forcing, a network with stronger recurrence can generate a straighter path. (F,G) Increasing the gain of the control loop can also increase path straightness. (H) In the absence of forcing, the network undergoes a transition from a quiescent state lacking movement and neural activity to an active state in which activity is self-sustaining as the loop gain crosses a critical point.

agent's translational speed $V(t)$ and neural amplitude $X_1(t)$ (derived from equations 25 and 27) around the stationary fixed point $V(t) = X_1(t) = 0$:

$$\tau \frac{d}{dt} \begin{bmatrix} V(t) \\ X_1(t) \end{bmatrix} = \begin{bmatrix} -\gamma & c \\ b & \hat{w}^{(1)} - 1 \end{bmatrix} \begin{bmatrix} V(t) \\ X_1(t) \end{bmatrix}. \tag{37}$$

The stability of this fixed point then depends on the leading eigenvalue $\lambda_1(\mathbf{A})$ of the coupling matrix, which is given by

$$\lambda_1(\mathbf{A}) = \frac{\hat{w}^{(1)} - 1 - \gamma}{2} + \sqrt{\left(\frac{\hat{w}^{(1)} - 1 + \gamma}{2}\right)^2 + bc}. \tag{38}$$

When the couplings are weak, this eigenvalue is negative and the fixed point is stable. However, as the eigenvalue crosses $\lambda_1(\mathbf{A}) = 0$, the quiescent state becomes unstable and self-sustaining activity emerges. Notably, this transition can be brought about by increasing the loop gain bc. Thus, even if neural activity and behavior are both dissipative in isolation, they can be sustained by strong sensorimotor feedback between the two. In other words, in this active state with strong feedback, neural activity is persistently maintained by the sensory input that continuously results from the circuit's own motor output.