Theoretical principles explain the structure of the insect head direction circuit

Pau Vilimelis Aceituno [*1], Dominic Dall'Osto [*1], and Ioannis Pisokas [2]

[*]These authors contributed equally to this work
[1]Institute of Neuroinformatics, University of Zürich and ETH Zürich
[2]School of Informatics, University of Edinburgh, United Kingdom

Abstract

To navigate their environment, insects need to keep track of their orientation. Previous work has shown that insects encode their head direction as a sinusoidal activity pattern around a ring of neurons arranged in an eight-column structure. However, it is unclear whether this sinusoidal encoding of head direction is just an evolutionary coincidence or if it offers a particular functional advantage. To address this question, we establish the basic mathematical requirements for direction encoding and show that it can be performed by many circuits, all with different activity patterns. Among these activity patterns, we prove that the sinusoidal one is the most noise-resilient, but only when coupled with a sinusoidal connectivity pattern between the encoding neurons. We compare this predicted optimal connectivity pattern with anatomical data from the head direction circuits of the locust and the fruit fly, finding that our theory agrees with experimental evidence. Furthermore, we demonstrate that our predicted circuit can emerge using Hebbian plasticity, implying that the neural connectivity does not need to be explicitly encoded in the genome of the insect but rather can emerge during development. Finally, we illustrate that in our theory, the consistent presence of the eight-column organisation of head direction circuits across multiple insect species is not a chance artefact but instead can be explained by basic evolutionary principles.

1 Introduction

Insects exhibit an impressive ability to navigate the world, travelling long distances to migrate, find food or reach places of interest before returning to their nests (Müller and Wehner, 1988; Menzel et al., 1996; Heinze et al., 2013; Collett, 2019), a feat that requires them to keep track of their orientation across long journeys (Mappes and Homberg, 2004; Merlin et al., 2012; Warren et al., 2019; Collett, 2019; Beetz et al., 2022). This orientation tracking is achieved by the use of visual cues as well as integrating angular velocity signals over time to maintain a heading estimate relative to a starting angle (Seelig and Jayaraman, 2015; Taube, 2007), known as *heading integration*.

Electrophysiological and calcium imaging studies have shown that the neural population encoding the head direction in insects has a sinusoid-shaped activation pattern (Labhart, 1988, 2000; Loesel and Homberg, 2001; Pfeiffer et al., 2005; Pfeiffer and Homberg, 2007; Kinoshita et al., 2007; Heinze et al., 2009; Homberg et al., 2011; el Jundi et al., 2014; ElJundi et al., 2019). Furthermore, downstream neural populations also encode velocity signals as sinusoidal activations (Lyu et al., 2022).

Theoretical work has speculated that this recurring motif of sinusoidal activity patterns might be so prevalent because it enables easy elementwise vector addition, where vectors encoded as sinusoidal activity waveforms can be added together to give a sinusoidal waveform encoding the sum of the vectors (Touretzky et al., 1993; Wittmann and Schwegler, 1995; Vickerstaff and Paolo, 2005). This allows the encoded heading to be easily used by downstream circuitry to track the insect's position (Mittelstaedt, 1985; Wittmann and Schwegler, 1995; Vickerstaff and Paolo, 2005; Haferlach et al., 2007; Wessnitzer et al., 2008; Sakura et al., 2008; Stone et al., 2017). Further studies have shown that models closely aligned with biological data can indeed implement heading integration using sinusoidal activity

---

patterns (Pisokas et al., 2020; Turner-Evans et al., 2020), and that such circuits can be learned (Vafidis et al., 2022).

Here we show that enabling easy vector addition cannot be the unique driving factor for the presence of sinusoidal heading encodings, as many other circuits with different activity patterns can perform vector addition in the same way. This finding led us to question whether the sinusoidal activation patterns seen in insect navigation circuits are a coincidence, or if they might offer a particular functional advantage that was selected for during evolution.

To address this question, we consider the basic principles necessary for a circuit encoding direction. Of all the circuits fulfilling these requirements, the sinusoidal activity pattern offers the best resilience to noise for the encoded information. However, obtaining this activity requires a circuit with a specific connectivity pattern between neurons. Thus, our theory predicts that the heading integration circuit will have a sinusoidal connectivity pattern. We compare our predicted circuit with connectivity data for the desert locust (*Schistocerca gregaria*) and fruit fly (*Drosophila melanogaster*) using network analysis tools, showing a strong agreement. We then ask how an insect brain might develop such a circuit, finding that a simple Hebbian learning rule is sufficient. Finally, we combine ideas from replication dynamics with our theory, which leads us to the conclusion that the eight-column structure is a consequence of basic theoretical principles, rather than an evolutionary coincidence.

2 Results

2.1 A theoretical circuit for heading integration

We consider a population of N "compass neurons" with an activity that encodes the direction of the insect as an angular variable θ. We represent the activity of this population by a vector where each element corresponds to the activity of one neuron, $\mathbf{a}(\theta) = [\mathbf{a}_1(\theta), \mathbf{a}_2(\theta), ..., \mathbf{a}_N(\theta)]$. We take $N = 8$ neurons, consistent with data from many insect species which possess an eight-column organisation, with each column encoding a different direction (Honkanen et al., 2019; Stone et al., 2017; Pisokas et al., 2020).

Each neuron's activity is updated depending on its current firing rate and the inputs it receives, both from other neurons in the circuit and externally. We formulate this update rule as follows,

$$\dot{\mathbf{a}}(\theta) = -\mathbf{a}(\theta) + \phi\left[\mathbf{W}\mathbf{a}(\theta) + \mathbf{u}(t)\right], \tag{1}$$

where $\mathbf{W}$ is the circuit's weight matrix, representing the connections between neurons; ϕ is the neural activation function, that converts the total neural input into an output firing rate; and $\mathbf{u}(t)$ is the external input that encodes the insect's angular velocity, via the PEN population of neurons (Green et al., 2017; Turner-Evans et al., 2017; Sayre et al., 2021; Pisokas et al., 2020; Turner-Evans et al., 2020; Hulse et al., 2021).

To simplify our derivations we allow the neural activity values to be both positive and negative, interpreting these values as being relative to a baseline neural firing rate. Similarly, we allow the weights to be both positive and negative, a common simplification in computational models (Li et al., 2023; Cornford et al., 2020; Kriegeskorte and Golan, 2019). This simplification will be addressed in Section Comparing the predicted circuit with biological data where we compare our model with experimental data.

2.1.1 Mathematical principles for neural heading integration

A circuit capable of performing heading integration must fulfil the requirements outlined in Box *1*. The first two requirements allow us to establish the family of possible path integration circuits. We then use the principle of noise minimisation to determine which circuits perform best.

Circular topology: The activity should have the same topology as the variable it is encoding to prevent discontinuities. To encode heading, the activity should have the topology of a 1D circle.

Rotational Symmetry: The heading integration circuit should work similarly, irrespective of the direction in which the insect travels. There should not be a bias for any direction.

Noise minimisation: The circuit should minimise the noise of the neural representation so the insect can navigate as precisely as possible.

2.1.2 Constraints on the neural activity

The neural activity should encode the insect's head direction with a matching topology. Because the heading is a single angular variable, the topology of the activity space should be a 1D circle.

Furthermore, the symmetry requirement implies that rotating the heading of the insect should rotate the neural activity vector without changing its shape. Concretely, whether the insect is facing north or east, the activity of the neural population as a whole should be the same, but with the identity of the neuron with each activity value being different.

We can formalise the symmetry requirement by considering a head direction, θ, and a rotation by an integer multiple, k, of the angular spacing between neurons, $\Delta\theta = \frac{2\pi k}{N}$. In this case, individual neuron activities follow the relation

$$\mathbf{a}_n(\theta + \Delta\theta) = \mathbf{a}_{n-k \mod N}(\theta), \tag{2}$$

which enforces that the neural activity vector is circularly rotated as the head direction changes. This relation can be expressed in the Fourier domain where, by the shift property, the circular rotation becomes a multiplication by a complex exponential:

$$\mathcal{F}_f\left[\mathbf{a}(\theta + \Delta\theta)\right] = \mathcal{F}_f\left[\mathbf{a}(\theta)\right]e^{\frac{i2\pi kf}{N}}, \tag{3}$$

where $i = \sqrt{-1}$ and $f \in [0, ..., N-1]$ is the index of the spatial frequency, also called the harmonic. Activity with spatial frequency f has f "bumps" around the ring.

Since the complex exponential has unit norm, the magnitude of the Fourier components remains the same for any rotation, $\|\mathcal{F}[\mathbf{a}(\theta + \Delta\theta)]\| = \|\mathcal{F}[\mathbf{a}(\theta)]\| \ \forall \Delta\theta$. Therefore, the shape of the activity pattern, $\mathbf{a} \equiv \mathbf{a}(0)$, can also be fully specified by its Fourier domain representation, $\mathcal{F}[\mathbf{a}]$, and the phase of this activity profile around the network encodes the heading of the insect:

$$\mathcal{F}_f\left[\mathbf{a}(\theta)\right] = \mathcal{F}_f\left[\mathbf{a}\right]e^{i\theta f}. \tag{4}$$

Taking the inverse Fourier transform with the constraint that the neural activities must be real, we get the following form for the neural activity,

$$\mathbf{a}_n(\theta) = \sum_{f=0}^{N-1} \|\mathcal{F}_f\left[\mathbf{a}\right]\| \cos\left(\frac{2\pi fn}{N} + \theta f\right). \tag{5}$$

It should be noted that the phase offset, θf, of each cosine waveform scales with the spatial frequency. This is because higher frequency waveforms have shorter wavelengths (in terms of number of neurons), so need their phases to rotate more quickly to move at the same speed around the network.

To explain this in more detail, we consider the case where this scaling is not used and the phase offset is the same for all harmonics, shown in **_Figure 1_** top row. For $f = 1$ (red line in **_Figure 1_**), the waveform has one bump and a wavelength of 8 neurons, so a 180° rotation corresponds to the

waveform moving 4 neurons. But for $f = 2$ (blue line in **Figure 1**), the waveform has two bumps and a wavelength of 4 neurons, so a 180° rotation corresponds to moving only 2 neurons, while a 4 neuron shift would correspond to a full rotation. The fact that the waveforms for different frequencies move at different speeds through the network implies that, if multiple spatial frequencies are used to independently encode the heading, their positions would not be aligned and the combined waveform would change shape as different angles are encoded (purple line in **Figure 1**). This violates the rotational symmetry principle, as the activity would look significantly different when the insect is facing north compared to south.

To solve this problem of misaligned harmonics, we return to the case specified in **Equation 5** where the phase offset is scaled linearly with the spatial frequency so that all waveforms move at the same speed and rotational symmetry is ensured, which is shown in **Figure 1** bottom row. For a network with $N = 8$ neurons, a 180° rotation shifts the activity waveform by 4 neurons. This 4 neuron shift corresponds to a $\frac{4}{8}$ or 180° phase offset for the $f = 1$ waveform, which has a wavelength of 8 neurons, but a $\frac{4}{4}$ or 360° phase offset for the $f = 2$ waveform, which has a wavelength of 4 neurons. While this maintains rotational symmetry, it implies that the $f = 2$ waveform is the same whether the heading angle 0° or 180° is encoded. As a consequence, a higher harmonic waveform (having $f > 1$) cannot on its own specify the encoded angle. The activity in all harmonics must be considered simultaneously to decode the angle, but even then this decoding might not be unique in the presence of noise (see Section Ambiguities in multiple harmonics decoding with drift). The difficulties associated with encodings utilising multiple harmonics will be further addressed in the following sections.

Figure 1: **Encoding heading with multiple harmonics.** Each panel shows the activity profiles encoding a particular heading value. Curves $f = 1$, $f = 2$, and $f = 1 + f = 2$ respectively denote the waveform of the first harmonic, second harmonic, and the sum of the two. The vertical dashed ($f = 1$ peak) and dotted lines ($f = 2$ peaks) indicate the neurons which respond maximally for the first and second harmonic, respectively. **Top row:** Encoding the heading as the sum of two independent harmonics causes the waveform to change shape as the insect rotates, because the waveform for each harmonic can only rotate a distance of $\frac{360°}{f}$ as the insect rotates a full revolution to ensure an unambiguous representation. **Bottom row:** If all harmonics are aligned to rotate at the same speed, the combined waveform shape does not change. However, this alignment implies that higher harmonic waveforms cannot be uniquely mapped back to a heading: here the $f = 2$ encoding is the same for $\theta = 0°$ and 180°.

2.1.3 Constraints on heading integration circuits

The basic assumptions outlined earlier also constrain the possible heading integration circuits – such circuits should allow for neural activity with the required topology and rotational symmetry to stably exist and propagate. Together, these two principles require that the activity in the circuit should have a constant total magnitude. We consider that this constraint is enforced by the nonlinear neural activation function, ϕ, in **_Equation 1_**, as detailed in Section Path integration dynamics for heading and position.

If the network activity is at the desired level, and the external input $\mathbf{u}(t)$ is projected onto the ring attractor such that it does not alter the total network activity, then we can consider the network dynamics to be linear at this operating point. Therefore, while within the space of possible activities that corresponds to the ring attractor, our circuit dynamics are effectively described as

$$\dot{\mathbf{a}}(\theta) = -\mathbf{a}(\theta) + \mathbf{W}\mathbf{a}(\theta) + \mathbf{u}(t). \tag{6}$$

The rotational symmetry principle also applies to the network. For the same shaped activity waveform to be able to stably exist at any position around the network, the network connectivity should also be rotationally symmetric. For example, the connection strength between the neurons encoding the north and north-east directions should be the same as between those encoding south and south-west. Mathematically, this imposes that the weight matrix, $\mathbf{W}$, is circulant, specifically that $\mathbf{W}_{n,m} = \mathbf{W}_{n+k,m+k}$. This matrix is fully specified by its first row, called the *connectivity profile* and denoted ω, meaning that we can express the product of the matrix with the neural activity as

$$(\mathbf{W}\mathbf{a}(\theta))_k = \mathbf{W}_k\mathbf{a}(\theta) = \sum_{m=0}^{N-1} \mathbf{W}_{k,m}\mathbf{a}_m(\theta) = \sum_{m=0}^{N-1} \mathbf{W}_{0,m-k}\mathbf{a}_m(\theta) = \sum_{m=0}^{N-1} \omega_{m-k}\mathbf{a}_m(\theta), \tag{7}$$

which can be simplified in terms of the convolution operation,

$$\mathbf{W}\mathbf{a}(\theta) = \omega * \mathbf{a}(\theta). \tag{8}$$

Considering the case where the insect is not moving, $\mathbf{u}(t) = 0$, and the network activity is stable, $\dot{\mathbf{a}}(\theta) = 0$, we can combine **_Equation 6_** and **_Equation 8_** to get a relation for the stable network activity

$$\mathbf{a}(\theta) = \omega * \mathbf{a}(\theta). \tag{9}$$

In the Fourier domain, this simplifies into

$$\mathcal{F}_f[\mathbf{a}(\theta)] = \mathcal{F}_f[\omega]\,\mathcal{F}_f[\mathbf{a}(\theta)]. \tag{10}$$

As for the activity waveform analysis, we note that the Fourier transform is taken on the neural indices, not on the temporal domain.

Here we have N equations that have to be satisfied, since the Fourier transform of an N-dimensional vector is also N-dimensional. For each harmonic frequency f, **_Equation 10_** has two solutions:

- $\mathcal{F}_f[\omega] = 1$, which implies that activity with this spatial frequency is stable in the network, and therefore that it can encode the insect's heading.

- $\mathcal{F}_f[\mathbf{a}(\theta)] = 0$, meaning that the activity in this harmonic is zero, and thus nothing can be encoded in this frequency.

2.1.4 Minimising noise propagation in the circuit

The only constraint on the connectivity weights given by **_Equation 10_** is that, if a frequency is used for encoding then $\mathcal{F}_f[\omega] = 1$. There is no restriction on the weights for the inactive harmonics – they remain free parameters.

However, non-encoding channels can still propagate noise, which would be prevented by setting $\mathcal{F}_f[\omega] = 0$. To illustrate this, we consider white noise, denoted by ϵ, that is added to the neural activity. When this noisy activity evolves in accordance with the dynamics from **_Equation 8_**, we have

$$\omega * (\mathbf{a}(\theta) + \epsilon) = \omega * \mathbf{a}(\theta) + \omega * \epsilon, \tag{11}$$

where the term $\omega * \epsilon$ corresponds to noise and should therefore be dampened. We want to minimise the strength of that noise, which is quantified by its variance

$$\mathbf{Var}\left[\omega * \epsilon\right] = \mathbf{E}\left[\|\omega * \epsilon\|^2\right] = \|\omega\|^2 \mathbf{Var}\left[\epsilon\right]. \tag{12}$$

Hence the magnitude of the noise that passes from one time interval to the next is modulated by the magnitude of the weight vector. By Parseval's theorem,

$$\sum_f |\mathcal{F}_f\left[\omega\right]|^2 = \sum_n |\omega_n|^2 = \|\omega\|^2, \tag{13}$$

which implies that to minimise noise propagation in the network we should impose $\mathcal{F}_f\left[\omega\right] = 0$ for all harmonics, f, where $\mathcal{F}_f\left[\mathbf{a}(\theta)\right] = 0$. We show this in simulation in **Figure 2**, where only the first harmonic encodes information and we vary the weights of the other harmonics. The noise is indeed minimised if the weights for all non-encoding harmonics are set to zero.

This result establishes that all non-encoding harmonics in the network should be set to 0 to minimise noise propagation, and allows us to recover the circuit connectivity

$$\omega_n = \sum_{f \in \mathbf{F}} \cos\left(\frac{2\pi n}{N}f\right), \tag{14}$$

where $\mathbf{F}$ is the set of harmonics used to encode the head direction, $\mathbf{F} = \{f \in [1..N-1] : \|\mathcal{F}_f[\mathbf{a}]\| \neq 0\}$. The choice of $\mathbf{F}$ therefore determines both the harmonics used for encoding the angle in the activity, and the connectivity of the circuit that supports this activity.

We leverage the same logic to prove that the number of encoding channels does not affect the signal-to-noise ratio in the network. If we have c encoding channels, each with the same activity, the total activity will grow linearly with c. The noise will also grow with $\|\omega\|^2 = c$, meaning that the signal-to-noise ratio will remain constant.

We now return to the question of whether using one or multiple harmonics is better. As mentioned, the signal-to-noise ratio in the network remains constant as additional encoding channels, c, are used. But additional channels imply additional harmonics, which by Parseval's theorem require a higher total neural activity. If the total activity in the network is limited, the signal-to-noise ratio is in fact *decreased* when using more than one harmonic.

$$\text{SNR} \propto \frac{\|\mathbf{a}\|^2}{\|\omega\|^2} = \frac{\|\mathbf{a}\|^2}{c}. \tag{15}$$

Additionally, as detailed in Section Constraints on the neural activity and Section Ambiguities in multiple harmonics decoding with drift, circuits using multiple harmonics perform worse in high noise environments because they are not guaranteed to provide an unambiguous orientation encoding. Finally, as we will discuss in Section Learning rules and development and Section Convergence of Oja's rule with multiple harmonics, the use of multiple harmonics also complicates the circuit's development. We therefore only select circuits that use a single harmonic for further analysis, reducing the number of possibilities to N.

2.2 Determining the optimal circuit

As we consider networks with $N = 8$ neurons, consistent with multiple insect species (Honkanen et al., 2019; Stone et al., 2017; Pisokas et al., 2020), the possible single harmonic circuits are $f = \{1, 2, 3, 4, 5, 6, 7\}$, where the zeroth harmonic is discarded because it only represents the baseline neural activity. This gives the following activities and weights, derived from **Equation 5** and **Equation 14**

$$\mathbf{a}_n(\theta) = \cos\left(\frac{2\pi nf}{N} + \theta\right) \quad \omega_n = \cos\left(\frac{2\pi fn}{N}\right). \tag{16}$$

The circuits for the first 4 harmonics are plotted in **Figure 3**, but not all of these circuits are valid. In particular, circuits with even spatial frequencies have multiple neurons with identical activity values because they share a common divisor with $N = 8$, as explained in detail in Section Degenerate circuits.

Figure 2: **Increased number of harmonics introduces more noise.** Increasing the number of active harmonic frequencies increases the effect of errors in the network. **A:** Weight matrix profiles, ω, for networks with increasing numbers of harmonics. **B:** Normally distributed noise with zero mean and standard deviation 0.3 was added to the network activity, then the network state updated until it reached steady state. Networks with fewer harmonics better filtered out noise. **C:** Noise variance increases linearly as the number of active harmonics increases, as predicted by **Equation 12**. Multiple trials were performed with different random noise.

For $f = 4$ there are only 2 unique activity values, $\mathbf{a}_n = \pm \cos(\theta) \ \forall n$, so this circuit can only encode one dimension, not a circular topology (see **Figure 3** D). For $f = 2$ and $f = 6$ (not plotted), there are 4 unique activity values, allowing the angle to be properly encoded. However, these circuits are degenerate because they are composed of two independent subcircuits, each encoding one direction (see **Figure 3** B). Because the weights connecting the two subcircuits are all zero, the activity in each subcircuit is independent of the activity in the other, and so the overall activity cannot be constrained to have the required circular topology.

All circuits with odd harmonic frequencies are equivalent. For example, as shown in **Figure 3**, the connections in the circuit for $f = 3$ are the same as those for the $f = 1$ circuit after the neuron identities are permuted. As detailed in Section Equivalence under permutation, the frequencies $f = \{1, 3, 5, 7\}$ always give the same circuit because the odd frequencies are coprime with the number of neurons $N = 8$.

Since the activities and weights of all the non-degenerate circuits $f = \{1, 3, 5, 7\}$ are the same as the base harmonic $f = 1$ up to a permutation, we choose the lowest harmonic $f = 1$, which gives us the following activity and weights:

$$\mathbf{a}_n(\theta) = \cos\left(\frac{2\pi n}{N} + \theta\right)$$
$$\omega_n = \cos\left(\frac{2\pi n}{N}\right). \tag{17}$$

2.3 Comparing the predicted circuit with biological data

Our theory proposes an optimally noise-resistant circuit for heading integration, and its corresponding activity. The prediction that heading should be encoded as a sinusoidal activity bump is consistent with previous theoretical models (Touretzky et al., 1993; Wittmann and Schwegler, 1995; Hartmann and Wehner, 1995; Zhang, 1996; Vickerstaff and Paolo, 2005; Haferlach et al., 2007; Stone et al., 2017), as well as experimental evidence in both the locust and fruit fly (Heinze et al., 2009; Turner-Evans et al., 2017). We note, however, that data from the fruit fly shows a more concentrated activity bump than what would be expected from a perfect sinusoidal profile (Seelig and Jayaraman, 2015; Turner-Evans

Figure 3: **Circuits for encoding with different harmonics.** We plot four circuits corresponding to each of the individual harmonics $f = 1, 2, 3, 4$. Excitatory synapses are marked in red and inhibitory in blue. Neurons are shown in yellow with each having a black arrow that marks the direction to which it is tuned from *Equation 16*. The $f = 1$ circuit is the simplest and constitutes our baseline. For the other cases we plot the original connectivity in the upper row and the rearranged network in the lower row. We find that the $f = 2$ circuit consists of two independent subnetworks encoding orthogonal directions, the $f = 3$ case is identical to $f = 1$ after permuting neuron indices, and the $f = 4$ case results in two connected groups of neurons inhibiting each other, hence it can only encode one direction. As such, all cases either have a degenerate ring structure ($f = 2, 4$) or are equivalent to $f = 1$ after permutation.

et al., 2017), and that calcium imaging (which was used to measure the activity) can introduce biases in the activity measurements (Siegle et al., 2021; Huang et al., 2021). Thus the sinusoidal activity we model is an approximation of the true biological process rather than a perfect description.

Importantly, our theory proposes that the optimally noise-resilient heading integration circuit should have synaptic weights that follow a sinusoidal pattern, even though such weights are not necessary for producing sinusoidal activity as discussed in Section Minimising noise propagation in the circuit. This is the main prediction of our theory and we sought to validate it using connectivity data from the locust and fruit fly.

However, before we can directly compare our model with biological circuits, we must address a number of modelling simplifications:

- In the model we considered a population of eight neurons, but in insects there are eight neural columns, each with several neuron types (EPG, PEG, PEN, $\Delta 7$). Of these we model only the EPG (known as the compass neurons) which encode the integrated head direction.

- The synaptic connections between neurons in our model can be both positive and negative, while biological neurons follow Dale's law, meaning that a single neuron can either have only positive or only negative synapses.

- The neural activity in the model was centred around zero and could have both positive and negative activity (firing rates), while real neurons only have positive firing rates.

Therefore, we simplified the biological connectivity to produce an equivalent circuit that could be directly compared with our model prediction. The neural population in our theoretical model cor-

responds to the biological EPG neurons, as these encode the integrated heading. We considered the other three neurons types that are part of the compass circuit (PEG, PEN, and $\Delta 7$ – see Kakaria and de Bivort (2017); Pisokas et al. (2020)) as just implementing connections between EPG neurons in accordance with biological constraints. We counted the number of different paths between EPG neurons, accounting for the sign of the connections (whether the path passed through an inhibitory $\Delta 7$ neuron), and used the net path count as a proxy for connectivity strength. This process is explained in detail in Section Path counting.

We then computed the average connectivity profile, that is, how each neuron connected to its neighbours around the ring, and compared this profile to the closest fitting sinusoid. Because the neuron gains, absolute synaptic strength and biophysical properties of the neurons are unknown, the units of the net path count are not necessarily equivalent to our abstract connection strength. We therefore fit an arbitrarily scaled and shifted sinusoid to the connection counts:

$$\omega_{m-n} = \beta \cos\left(\frac{2\pi(m-n)}{N}\right) + \gamma, \tag{18}$$

where $m - n$ is the circular distance between two neurons while β and γ are constants that are fit to minimise the precision-weighted mean squared error compared with the experimental connectivity profile (see Section Fitting weights).

Analysing the data from Pisokas et al. (2020) we consider the shortest excitatory and inhibitory paths between EPG (also known as CL1a) neurons in the locust, which have lengths of 2 and 3, respectively. There are no direct connections of length 1, and all the paths of length ≥ 4 must pass through the same neuron type multiple times. This path counting analysis for the locust is shown in *Figure 4* A, and the procedure is detailed in Section Path counting. We find that the connectivity profile between neurons for the locust is very close to sinusoidal in shape, supporting our theoretical prediction.

For the fruit fly, we used data from Hulse et al. (2021) which provides synapse counts between pairs of neurons. These synapse counts were used as a proxy for connectivity strength, a view that has been validated in previous experiments (Liu et al., 2022; Barnes et al., 2022). After identifying the neurons and connections of interest, we grouped the neurons in eight columns following the logic presented in Pisokas et al. (2020), with detailed methodology explained in Section Data preprocessing. We repeated the path counting analysis for the fruit fly with synapse count data (*Figure 4* B), and found that while the data are noisy, the connectivity profile fits a single sinusoid pattern reasonably well. However, the high variability in the synapse counts makes our hypothesis difficult to differentiate from alternative shapes (see Section Fitting weights).

Taken together, this analysis shows that our theory is consistent with experimental data – using binary connectivity data for the desert locust and synapse-count data for the fruit fly.

Figure 4: **Comparison to biological data.** Biological network models with 4 distinct neural populations were simplified to equivalent networks with 1 population by counting paths of length 1, 2, and 3 between EPG neurons (which encode the integrated heading) and using the net signed path count as a proxy for connectivity strength. The average connectivity profile for each neuron to its neighbours around the ring was compared to the sinusoidal connectivity predicted by our theory. The locust network has no standard deviation because the data in Pisokas et al. (2020) are based on light microscopy which couldn't resolve variations between columns. Excitatory connections are shown in red and inhibitory connections in blue, while the strength of a connection is indicated by the line width.

2.4   Learning rules and development

Having validated the connectivity of our theoretical circuit by comparing it to experimental data, we ask whether our circuit lends itself to biological development. Specifically, we show that even though our circuit requires precise connection strengths, this connectivity can be developed naturally by a Hebbian learning rule. Our approach follows from previous research which has shown that simple Hebbian learning rules can lead to the emergence of circular line attractors in large neural populations (Stringer et al., 2002), and that a head direction circuit can emerge from a predictive rule (Vafidis et al., 2022). In contrast to this work, we focus only on the self-sustaining nature of the

heading integration circuit in insects and show that our proposed sinusoidal connectivity profile can emerge naturally.

Because the weight matrix is circulant, its eigenvalues are equal to the Fourier spectrum of its first row, which we derived to only have one nonzero value, for $f = 1$. The network therefore only has a single eigenvalue and so projects activity into a single dimension. This operation is similar to classical dimensionality reduction methods such as principal component analysis, which can be implemented by Hebbian-like learning rules (Dayan and Abbott, 2001). We thus analyse the effects of incorporating Oja's rule into our model, a classical variant of Hebbian learning where the synaptic strength between two neurons grows when both neurons are active simultaneously, and the total synaptic strength is regularised to prevent exploding weight growth,

$$\frac{d\mathbf{W}_{n,m}}{dt} = \eta \left(a_m(t)a_n(t) - a_n^2(t)\mathbf{W}_{n,m} \right),$$ (19)

where $\mathbf{W}_{n,m}$ is the synaptic connection strength from neuron m to neuron n, a_m and a_n are the pre- and postsynaptic activities respectively, and η sets the speed of the weight update dynamics relative to the activity dynamics.

For our analysis we assume an input $\mathbf{u}(t)$ which pushes the neural activity in every different direction (i.e. the insect faces all possible directions) and then integrate the weight updates over some long period of time,

$$\Delta\mathbf{W}_{n,m} = \eta \int_\theta \left[a_m(\theta)a_n(\theta) - a_n(\theta)^2\mathbf{W}_{n,m} \right] d\theta,$$ (20)

where θ is the integration space. We will assume that the insect moves uniformly over the full circle. Applying the activity from *Equation 17*, we can find the fixed point of this update rule, when $\Delta\mathbf{W}_{n,m} = 0$:

$$\mathbf{W}_{n,m} = \frac{\eta \int_\theta \cos\left(\theta + \frac{2\pi n}{N}\right)\cos\left(\theta + \frac{2\pi m}{N}\right)d\theta}{\eta \int_\theta \cos^2\left(\theta + \frac{2\pi n}{N}\right)d\theta} \quad = \cos\left(\frac{2\pi(n-m)}{N}\right).$$ (21)

Combined with *Equation 7* and *Equation 17*, this result means that if there is sinusoidal activity in the network, the weights will naturally converge to the optimal sinusoidal values by way of Oja's rule. This leads to the emergence of sinusoidal activity and weights from self-consistency: initial noisy sinusoidal activity will enforce weights that are close to a sinusoid, and those weights will filter out noise to make the activity even closer to a sinusoid, which will in turn make the weights more sinusoidal. This process will iteratively make the activity and the weights converge to the solution from *Equation 17*.

An important point in the use of Oja's rule is that it would tend to concentrate the activity in a single harmonic. In a linear network, the harmonics would compete during learning, leading to one single harmonic emerging and all others being suppressed, as shown in Section Linear neurons. For neurons with a nonlinear activation function, secondary harmonics would emerge, but would remain small under mild assumptions, as shown in Section Neurons with nonlinear activations. Oja's rule will still cause the weights to converge to approximately sinusoidal connectivity.

The finding that the weights will converge to a sinusoidal connectivity with learning has two interesting consequences from a biological standpoint: First, our circuit can emerge even when starting with only very coarse initial weights, without the need for high precision initial connectivity. Second, this simple plasticity rule allows the system to repair or recover from perturbations in its synapses as shown by simulations in *Figure 5*. However, this learning rule only applies to the weights that ensure stable, self-sustaining activity in the network. The network connectivity responsible for correctly integrating angular velocity inputs (given by the PEN to EPG connections in the fly) might require more elements than a purely Hebbian rule (Stringer et al., 2002), such as the addition of a predictive component (Vafidis et al., 2022).

Figure 5: **Recovery of synaptic weights with a Hebbian learning rule.** The synaptic weights converge to a sinusoidal pattern under Oja's rule when the network activity is dynamic and a sinusoidal input is provided. **A:** The weights start at zero and slowly converge to the prescribed sinusoidal profile, showing that this connectivity can emerge from scratch. **B:** The sinusoidal weights are perturbed by noise but learning ensures that the weight profile is corrected. In both cases the network's initial activity is corrupted with zero mean Gaussian noise. Noisy sinusoidal input is provided to rotate the activity bump around the network at a speed of $1/8$ neurons per timestep. The simulation runs for 100 periods. Parameters: $N = 8$, integration timestep $\Delta t = 0.01, \eta = 0.1, \|\mathbf{a}\| = 1, \sigma_W = 0.2, \sigma_a = 0.2, \sigma_u = 0.2$.

2.5 Evolution of the eight-column circuit

Having derived and experimentally validated the theoretical circuit, we now address another question: whether there might be a reason that insect head direction circuits have an eight-column architecture (Honkanen et al., 2019; Stone et al., 2017; Pisokas et al., 2020). The derivations leading to *Equation 17* are valid for other values of N, so there is no *a priori* reason to expect $N = 8$.

Our reasoning follows recent studies in genetics (Johnston et al., 2022; Dingle et al., 2018), which argue that an observed organism is more likely to have resulted from a simple than a more complex genome: evolution favours simplicity. We note that powers of two are easier to generate with replication dynamics than other numbers, because they just require each cell to divide a set number of times. Other numbers require that, at some point, two cells resulting from a division must behave differently, necessitating more complex signalling mechanisms and rendering this possibility less likely to have been developed by evolution without some other driving factor. We therefore expect N to be a power of two unless required to be otherwise.

As we show in Section Circuits with different neuron counts, not all numbers of neurons enable a working circuit. The circuits for $N = 2$ and $N = 4$ are degenerate – either producing a single dimensional encoding, or two disconnected circuits that do not enforce the required circular topology. $N = 8$ is the smallest power of two that could result in a non-degenerate circuit. This hints at the possibility that the eight-column architecture is not a chance evolutionary artefact, but rather that it is the genetically simplest circuit capable of performing heading integration.

3 Discussion

In this work we derived an optimal noise-minimising circuit for encoding heading and verified that this circuit matches experimental data from insects. Furthermore, we showed that such a circuit can be developed and maintained by a biological learning rule, and proposed a mathematical argument for the eight-column structure found in insect compass circuits. In this section, we discuss the implications and limitations of these contributions, and outline potential future work.

Heading integration circuits in insects have been extensively studied, with models ranging in complexity from simplified conceptual networks (Wittmann and Schwegler, 1995; Cope et al., 2017) to sophisticated models constrained by biological data and featuring multiple neuron types (Kakaria and de Bivort, 2017; Su et al., 2017; Kim et al., 2017; Pisokas et al., 2020; Lyu et al., 2022). Previous theoretical work has argued that a sinusoidal activity encoding is such a common motif in insect navigation because it facilitates elementwise vector addition (Wittmann and Schwegler, 1995). However, this cannot be the only reason because as we show there is a whole family of circuits with different encoding patterns that enable easy vector addition. By showing that sinusoidal activity emerges from the theoretically most noise-resilient heading integration circuit, and verifying that the corresponding circuit matches experimental data, we close this explanatory gap.

We also show that our proposed circuit can be developed by a simple Hebbian-based learning rule, and that the presence of the eight-column structure can be explained from the perspective of replication dynamics. Both results align with the idea that evolution should be biased towards structures that are easier to encode in the genome (Johnston et al., 2022; Alon, 2007) and learn (Zador, 2019). To the best of our knowledge, this is the first time that such arguments have been put forward in the context of a specific circuit with a specific function.

Our work still has some unaddressed limitations, in particular regarding the topology of the activity. The use of a circular topology to encode the head direction of the insect is valid only in 2D environments. But many species, including the fruit fly, can navigate in a 3D environment. We argue that even if these insects live in a 3D environment, the third dimension (up-down) is different from the other two due to gravity and the existence of a hard boundary (the ground). Further studies would be required to investigate the full effects of 3D motion.

We could also investigate circuits that integrate position, not only heading, which would require the activity to have a 2D plane topology instead of a circle. This would be particularly relevant for foraging insects such as bees or ants, whose ability to remember their position with respect to their nest has been the subject of many experimental and computational studies (Collett, 2019; Wehner and Srinivasan, 2003). The position integrating neurons in these insects are also predicted to have sinusoidal activation patterns (Wittmann and Schwegler, 1995; Vickerstaff and Paolo, 2005; Stone et al., 2017).

Finally, another interesting avenue for future work is to compare the encoding of direction in insects with that of mammals, which encode heading in a fundamentally different way that uses many more neurons. This raises a critical question: why would the circuit and encoding be different if navigation follows similar principles across species? We speculate that the difference might lie in the type of navigation that the two classes of animals use. Insects often rely on a reference system that is globally anchored to a certain point or phenomenon, whether it is their nests for ants and bees, the sunlight polarisation pattern for locusts, or the milky way for dung beetles. On the other hand, mammals such as rodents typically do not use global cues but rely on local landmarks that are context-dependent and only occur in specific locations. Therefore, mammals must employ a flexible encoding that can be updated as different environments are explored. Further investigation of this possibility would require a different set of principles than those selected here.

4 Materials and Methods

4.1 Path integration dynamics for heading and position

Here we make more precise statements about the dynamics of the circuit. The topology of the activity is a circular line attractor, so that any perturbation falls back into a circle and the position of the activity around the circle represents an angle. In our circuit with the dynamics from **Equation 1**, this circular line attractor is achieved by setting

$$\begin{aligned}
\|\phi'\left[\mathbf{W}\mathbf{a}\right]\| = 1 \quad & \forall \|\mathbf{a}\| = r \\
\|\phi'\left[\mathbf{W}\mathbf{a}\right]\| < 1 \quad & \forall \|\mathbf{a}\| > r \\
\|\phi'\left[\mathbf{W}\mathbf{a}\right]\| > 1 \quad & \forall \|\mathbf{a}\| < r.
\end{aligned} \tag{22}$$

where r is the radius of the attractor. Given that the dynamics are effectively linear we can apply a Fourier transform and write the original dynamics around the line attractor directly in the spatial Fourier basis,

$$\mathcal{F}_f[\dot{\mathbf{a}}] = \mathcal{F}_f[-\mathbf{a} + \omega * \mathbf{a} + \mathbf{u}(t)], \tag{23}$$

which we can expand to obtain

$$\mathcal{F}_f[\dot{\mathbf{a}}] = -\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\omega]\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\mathbf{u}(t)], \tag{24}$$

which only applies to the circular activity where the dynamics are linear. To incorporate the dynamics that force the activity to return to the cycle, we add a new term into the network dynamics which fulfils **Equation 22** – increasing $\|\mathcal{F}_f[\mathbf{a}]\|$ if $\|\mathbf{a}\| < r$ and decreasing it if $\|\mathbf{a}\| > r$. This increase or decrease can be incorporated as a simple scaling factor on the activity decay term

$$\mathcal{F}_f[\dot{\mathbf{a}}] = -\alpha(r - \|\mathbf{a}\|)\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\omega]\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\mathbf{u}(t)], \tag{25}$$

where $\alpha(r - \|\mathbf{a}\|)$ is a nonlinear smooth function with a single minimum and $\alpha(0) = 1$, which forces the activity to have the appropriate magnitude, $\|\mathbf{a}\| = r$.

As the activity always falls back to the circular line attractor, the heading integration is linear around the circle. This implies that any small movement of the animal or the perception of a sensory cue is first projected onto the circle, then linearly integrated. We therefore only consider how network inputs cause the activity to rotate around the network,

$$\mathcal{F}_f[\dot{\mathbf{a}}] = -\alpha(r - \|\mathbf{a}\|)\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\omega]\mathcal{F}_f[\mathbf{a}] + \mathcal{F}_f[\mathbf{u}(t)]_{\perp \mathbf{a}}, \tag{26}$$

where $\mathcal{F}_f[\mathbf{u}(t)]_{\perp \mathbf{a}}$ is the projection of the input signal perpendicular to the current activity pattern.

In more intuitive terms, the neurons have a saturating nonlinear activation function where they modulate their gain based on the total activity in the network. If the activity in the network is above the desired level, r, the gain is reduced and the activity decreases, and when the activity of the network is less than desired level, both the gain and the activity increase. Note that in this scenario transient deviations from the line attractor, which would induce nonlinear behaviour in the circuit dynamics, are tolerable. External inputs, $\mathbf{u}(t)$, could transiently modify the shape of the activity, producing activity shapes deviating from what the linear model can accommodate. For example, the shape of the bump attractor could be modified through nonlinearities while the insect attains high angular velocity (Turner-Evans et al., 2017).

Such nonlinear dynamics do not conflict with the theory developed here, which only requires linearity when the activity is projected onto the circular line attractor. In our framework, the linearity of integration *at the circular line attractor* is not a computational assumption, but rather it emerges from the principle of symmetry.

4.2 Ambiguities in multiple harmonics decoding with drift

We consider the case where multiple harmonics are used, and their phases have drifted from each other. We only focus on angular drift, rather than noise in the full activity, because as noted in Section Path integration dynamics for heading and position, any deviations in the overall activity level in the network will dissipate.

For example, we consider the harmonics f_1 and f_2. The activity is given by **Equation 5**

$$\mathbf{a}_n(\theta) = \cos\left(\frac{2\pi f_1 n}{N} + f_1\theta\right) + \cos\left(\frac{2\pi f_2 n}{N} + f_2\theta\right). \tag{27}$$

If the phase of the second harmonic drifts by $\delta\theta$,

$$\mathbf{a}_n^{\delta\theta}(\theta) = \cos\left(\frac{2\pi f_1 n}{N} + f_1\theta\right) + \cos\left(\frac{2\pi f_2 n}{N} + f_2\theta + \delta\theta\right) \tag{28}$$

We can calculate the alignment between the activity with drift and the activity without drift as a dot product of the activity vectors,

$$\frac{\langle \mathbf{a}(\theta), \mathbf{a}^{\delta\theta}(\theta)\rangle}{\|\mathbf{a}(\theta)\|\|\mathbf{a}^{\delta\theta}(\theta)\|} = \frac{1}{\|\mathbf{a}(\theta)\|^2}\sum_n \mathbf{a}_n(\theta)\mathbf{a}_n^{\delta\theta}(\theta) = 1 + \cos(\delta\theta). \tag{29}$$

However, there are other angles where the alignment is better. For example, we can consider an estimate angle, $\hat{\theta}$, with its corresponding activity, $\mathbf{a}(\hat{\theta})$, which gives

$$\frac{\langle \mathbf{a}(\hat{\theta}), \mathbf{a}^{\delta\theta}(\theta)\rangle}{\|\mathbf{a}(\hat{\theta})\|\|\mathbf{a}^{\delta\theta}(\theta)\|} = \frac{1}{\|\mathbf{a}(\theta)\|^2}\sum_n \mathbf{a}_n(\hat{\theta})\mathbf{a}_n^{\delta\theta}(\theta) = \cos(\theta - \hat{\theta}) + \cos(\theta - \hat{\theta} + \delta\theta). \tag{30}$$

For small $\delta\theta$, this is maximised when $\hat{\theta} = \theta + \delta\theta/2$. However, since θ is circular, if the drift is π then there are two possible positions given by $\hat{\theta} = \theta \pm \pi/2$.

This implies that an encoding using multiple harmonics does not necessarily offer a unique decoding in the presence of noise.

4.3 Equivalent circuits and degeneracies

4.3.1 Equivalence under permutation

The activity of neurons given by **Equation 5** implies that the preferred angle of neuron n is given by

$$\varphi_n = 2\pi\frac{fn \mod N}{N} \tag{31}$$

where in our case $N = 8$. **Table 1** shows $nf \mod (N)$ evaluated for all neurons in the network using different harmonics. We notice that for $f = \{1, 3, 5, 7\}$ all the numbers from zero to seven appear, while for $f = \{2, 6\}$ we only get the even numbers, for $f = 4$ we get only zero and four, and for $f = 8$ there is only zero.

The explanation is based on number theory. If f and N have the greatest common divisor $\gcd(N, f) = d$, then $nf \mod N = 0$ for $n = N/d$. This implies that the preferred angle of neuron $n = 0$ is the same as that of $n = N/d$. When N, f are co-prime $d = 1$, n goes from 0 to $N - 1$ without repeating any value. However, when $d > 1$, the neuron $n = N/d$ has the same tuning as the neuron $n = 0$, the neuron $n = N/d + 1$ has the same tuning as the neuron $n = 1$ and so on. In other words, the angular tuning of the neurons has a period of $\gcd(N, f)$.

This divides the possible circuits into the following groups:

- $f = \{1, 3, 5, 7\}$ which contain $N = 8$ angular tunings with values $[0, \pi/4, \pi/2, 3\pi/4, \pi, 5\pi/4, 3\pi/2, 7\pi/4]$ because every odd number is coprime with 8 and thus the $\gcd(N, f) = 1$.

- $f = \{2, 6\}$ which cycle through 4 possible directions $[0, \pi/2, \pi, 3\pi/2]$ because $\frac{N}{\gcd(N, f)} = 4$

- $f = 4$ which can only represent 2 possible directions $[0, \pi]$.

- $f = 8$ which can only represent the direction 0

	$n=0$	$n=1$	$n=2$	$n=3$	$n=4$	$n=5$	$n=6$	$n=7$	$\gcd(8,f)$
$f=1$	0	1	2	3	4	5	6	7	1
$f=2$	0	2	4	6	0	2	4	6	2
$f=3$	0	3	6	1	4	7	2	5	1
$f=4$	0	4	0	4	0	4	0	4	4
$f=5$	0	5	2	7	4	1	6	3	1
$f=6$	0	6	4	2	0	6	4	2	2
$f=7$	0	7	6	5	4	3	2	1	1
$f=8$	0	0	0	0	0	0	0	0	8

Table 1: Distribution of neural activity phases for different harmonics: We computed $fn \mod 8$ for $n = \{1,2,3,4,5,6,7\}$ and $f = \{1,2,3,4,5,6,7,8\}$ and found that all possible phases appear for any odd number. This happens because the $\mod 8$ operation imposes an abelian group structure, namely $\mathbb{Z}/8$ and any coprime with 8 will be a generator of the whole group. If we use instead $f = 2$ or $f = 6$ we have that we can divide the $N = 8$ and f by 2 and we get the abelian group $\mathbb{Z}/4$ which has four elements. The same argument applies to $f = 4$, leaving only two elements, and for $f = 8$ we get a single element.

4.3.2 Degenerate circuits

Given the groupings presented in the previous subsection, we notice that not all of them allow the encoding of a full range of angles, $(0, 360]$. Notably, for $f = 8$ the neurons only encode one angle, and for $f = 4$ only two complementary angles are encoded, 0 and π. This means that we cannot represent a full circle, because we need two different dimensions to do so, but in both of these cases we can only encode one.

For $f = 2, 6$ we obtain two groups of neurons, and in each group there are two neurons that encode two opposing directions, thus it is possible to cover a full circle. However, the circuit is degenerate, as shown in *Figure 3*. In this case there are no connections between even and odd neurons. Thus, we have two groups of neurons that are disconnected, and each group encodes only one direction (either north-south or east-west). While the encoding could work in principle, the circuit is decoupled, meaning that there is nothing in this circuit that prevents the activity from having activities outside the circular topology: as the two circuits are disconnected, a given value of activity in one group does not restrict the activity in the other.

4.3.3 Circuits with different neuron counts

We evaluate the viability of circuits with other values of N. From findings in the the previous section, we note that choosing N to be a prime number implies that all frequencies will be coprime with N, and thus that all the neurons will have a different tuning.

For $N = 2$, the two angles are $0, \pi$, in which case the circuit can only encode one dimension and thus cannot encode a full circle, as was the case for $N = 8, f = 4$ in Section Degenerate circuits.

For $N = 4$, if $f = 2$ the circuit can encode only the angles $0, \pi$ and we have the same case as for $N = 2$. But if $f = 1$ or $f = 3$, the neurons are tuned to the angles $0, \pi/2, \pi, 3\pi/2$ which can encode a circle. However, by looking at the connectivity matrix from *Equation 16*, we find that neurons $1, 3$ are connected to one another, but not to $2, 4$, as was the case for $N = 8, f = 2$ in Section Degenerate circuits. This implies that the circuit does not enforce a circular topology, and thus it does not work.

Notice that we could think of another structure where each neuron encodes a different direction. Intuitively, we would have a north-south neuron and an east-west neuron, as we would have in Cartesian coordinates. However, this would have the same problem as $N = 4$, where the two neurons would not interact and thus the topology would be wrong. Additionally, when the insect heads north, the north-south neuron would be very active, while when it heads south it would be very inactive. This means that the circuit as a whole would have a different firing rate depending on the direction, breaking the symmetry assumption.

4.4 Path counting

We counted the number of different paths between EPG neurons, accounting for the sign of the connections (whether the path passed through an inhibitory $\Delta 7$ neuron), and used the net path count

as a proxy for connectivity strength. The results of this analysis are shown in Section Comparing the predicted circuit with biological data, but we will detail the process here using the locust network from Pisokas et al. (2020) as an example, shown in *Figure 6*.

We consider the shortest excitatory and inhibitory pathways between EPG neurons, which have lengths of 2 and 3 respectively. There are no direct paths between EPG neurons in the locust circuit. In this case there are two paths of length 2 that implement self-excitatory connections for EPG_1:

- $EPG_1 \rightarrow PEG_1 \rightarrow EPG_1$

- $EPG_1 \rightarrow PEN_1 \rightarrow EPG_1$

And there is one connection of path length 2 that connects EPG_1 to each of its nearest neighbours:

- $EPG_1 \rightarrow PEN_1 \rightarrow EPG_2$

- $EPG_1 \rightarrow PEN_1 \rightarrow EPG_8$

For inhibitory connections there are four paths of length 3 that connect EPG_1 to the neuron on the opposite side of the ring, EPG_5. These are:

- $EPG_1 \rightarrow \Delta7_5 \rightarrow PEN_5 \rightarrow EPG_5$

- $EPG_1 \rightarrow \Delta7_5 \rightarrow PEG_5 \rightarrow EPG_5$

- $EPG_1 \rightarrow \Delta7_4 \rightarrow PEN_4 \rightarrow EPG_5$

- $EPG_1 \rightarrow \Delta7_6 \rightarrow PEN_6 \rightarrow EPG_5$

There are three connections of path length 3 connecting EPG_1 to EPG_4:

- $EPG_1 \rightarrow \Delta7_4 \rightarrow PEN_4 \rightarrow EPG_4$

- $EPG_1 \rightarrow \Delta7_4 \rightarrow PEG_4 \rightarrow EPG_4$

- $EPG_1 \rightarrow \Delta7_5 \rightarrow PEN_5 \rightarrow EPG_4$

Finally, there is one path of length 3 connecting EPG_1 to the neuron perpendicular to it around the ring, EPG_3:

- $EPG_1 \rightarrow \Delta7_4 \rightarrow PEN_4 \rightarrow EPG_3$

We then add these connections together to obtain the net path count profile for the network. Because the network is rotationally symmetric, the connections from EPG_1 are the same as the connections from EPG_2 but with the neurons indices incremented by 1. The generalised profile is shown in *Table 2*.

$EPG_n \rightarrow$	$n-4$	$n-3$	$n-2$	$n-1$	n	$n+1$	$n+2$	$n+3$
Excitatory path count	0	0	0	1	2	1	0	0
Inhibitory path count	4	3	1	0	0	0	1	3
Net path count	-4	-3	-1	1	2	1	-1	-3

Table 2: Net path count profile between EPG neurons in the locust circuit.

Figure 6: The full locust connectivity network from Pisokas et al. (2020) (left), and the connections with path lengths 2 and 3 from EPG$_1$ to all other EPG neurons in the network (right). Because the network is rotationally symmetric these path counts generalise to all EPG neurons as shown in **Table 2**.

4.5   Data preprocessing

The following section details our data preprocessing to produce a central complex circuit for the fruit fly similar to that in Pisokas et al. (2020) using synaptic counts from the fruit fly connectome data set (Hulse et al., 2021). Full details are shown in our available code.

We first identified the 6 neuron types which corresponded to the neurons of interest in the central complex (**Table 3**).

Neuron class	Number
PEN_a(PEN1)	20 neurons
PEN_b(PEN2)	22 neurons
PEG	18 neurons
EPG	46 neurons
EPGt	4 neurons
Delta7	42 neurons

Table 3: Neuron types and their numbers in the fruit fly connectome data set (Hulse et al., 2021).

We grouped the two PEN and two EPG populations together for further analysis.

From the connectome data we created a connectivity matrix (**Figure 7**) containing the number of synapses between all pairs of neurons of the above mentioned types. This contained 129473 synapses between 152 neurons in total.

Next, neurons in the same glomerulus were grouped together. The EPG neurons were divided between 18 glomeruli, L1-L9 and R1-R9. The PEN neurons were also divided between 16 glomeruli, L2-L9 and R2-R9. The PEG neurons were divided into 18 glomeruli with only 1 neuron in each, L1-L9 and R1-R9. The Delta7 neurons had 10 unique sub-types (**Table 4**) which were grouped into 8 glomeruli based on their left glomerulus index – i.e. L4R5_R and L4R6_R were grouped in L4.

Delta7 neuron sub-types	Glomerulus group
Delta7(PB15)_L1L9R8_R	Delta7 L1
Delta7(PB15)_L2R7_R	Delta7 L2
Delta7(PB15)_L3R6_R	Delta7 L3
Delta7(PB15)_L4R5_R	Delta7 L4
Delta7(PB15)_L4R6_R	Delta7 L4
Delta7(PB15)_L5R4_L	Delta7 L5
Delta7(PB15)_L6R3_L	Delta7 L6
Delta7(PB15)_L6R4_L	Delta7 L6
Delta7(PB15)_L7R2_L	Delta7 L7
Delta7(PB15)_L7R3_L	Delta7 L7
Delta7(PB15)_L8R1R9_L	Delta7 L8

Table 4: List of the Delta7 neuron sub-types in the fruit fly connectome data set (Hulse et al., 2021). These differ in the glomeruli their pre- and post-synaptic terminals innervate. We are referring to those altogether as Delta7 neurons in the present account.

After this grouping we had a connectivity matrix of 60 neurons in total (**Figure 8**).

For the EPG, PEN, and PEG neurons we then grouped glomeruli mirrored in the two hemispheres together – i.e. L1 and R1 were grouped into glomerulus 1. This resulted in a connectivity matrix of 34 neurons (**Figure 9**).

Finally, as in (Pisokas et al., 2020) we grouped the PEG9 and PEG1 neurons and EPG9 and EPG1 neurons together. The former because they both output to the same ellipsoid body segment, and the latter because they both receive common input. This resulted in a connectivity matrix with 32 neurons (**Figure 10**) which we used for the analysis in **Figure 4** as this network had the 8-fold rotational symmetry compatible with our theoretical model.

4.6   Fitting weights

Having calculated the paths or synaptic strengths between EPG neurons in the network, we next computed the average connectivity profile for the network – how each EPG neuron connected to its neighbours a certain distance around the ring. We then compared this connectivity profile to the closest fitting sinusoid to quantify to what degree our prediction of sinusoidal weights was consistent with biological data.

We use the vector

$$\hat{\omega}_d = \beta \cos\left(\frac{2\pi d}{N}\right) + \gamma, \tag{32}$$

where $d = ((m - n) + 4 \mod 8) - 4 \in [-4..3]$ is the signed circular distance from neuron m to neuron n.

In the networks from Pisokas et al. (2020) the weights are rotationally symmetric, so we minimise the mean squared error between the sinusoidal profile and the biological connectivity profile,

$$\beta^*, \gamma^* = \arg\min_{\beta,\gamma} \frac{1}{N} \sum_{d=-N/2}^{N/2-1} \left(\hat{\omega}_d(\beta,\gamma) - \omega_d\right)^2, \tag{33}$$

by using least squares. The network derived from the synaptic count data in Hulse et al. (2021) is not rotationally symmetric so each value in the average connectivity profile, ω_d, has a corresponding standard deviation, σ_d. In this case we minimise the precision-weighted mean squared error, which emphasises fitting connectivity profile values that are more consistently seen in the network:

$$\beta^*, \gamma^* = \arg\min_{\beta,\gamma} \frac{1}{N} \sum_{d=-N/2}^{N/2-1} \left(\frac{\hat{\omega}_d(\beta,\gamma) - \omega_d}{\sigma_d}\right)^2. \tag{34}$$

The result of fitting sinusoidal profiles to the data is shown in Section Comparing the predicted circuit with biological data. Since the fit for the fruit fly is not as clear as for the locust, we also compared

the observed weights in the fruit fly with Gaussian and von Mises curves. Since both the Gaussian and von Mises distributions also have a parameter specifying their width, this requires fitting an extra parameter. For the Gaussian we have

$$\beta_g e^{-\frac{n^2}{\sigma_g}} + \gamma_g, \tag{35}$$

where $\beta_g, \sigma_g, \gamma_g$ are parameters to fit. For the von Mises distribution we have

$$\beta_v e^{\kappa_v \cos(n)} + \gamma_v, \tag{36}$$

where $\beta_v, \kappa_v, \gamma_v$ are parameters to fit. For both the Gaussian and von Mises, we find a good agreement with the weights in *Figure 11*, and a lower root mean square error (RMSE) than with any combination of harmonics. However, we note that our sinusoidal model has two parameters instead of three, so to make a fair comparison we use the corrected Akaike Information Criterion (AICc) (Burnham and Anderson, 2004), which is given by

$$\text{AICc} = 2p - 2\ln(L) + \frac{2p^2 + 2p}{N - p - 1} = 2p + 2N\text{RMSE}^2 + \frac{2p^2 + 2p}{N - p - 1}, \tag{37}$$

where p is the number of distribution parameters, N the number of samples, and $\ln(L)$ is the log-likelihood, which in this case is the sum of squared errors, $\ln(L) = -N\text{RMSE}^2$. We find that the lowest AICc (corresponding to the best model) corresponds to the harmonic $f = 1$ (*Table 5*).

As a complementary approach to evaluate the shape of the distribution, we first fit the Gaussian and von Mises distributions to the best fit $f = 1$ curve. We then freeze the width parameters of the distributions (σ_g for the Gaussian and κ_v for the von Mises) and only optimise the amplitude and vertical offset parameters (β and γ) to fit the data. This approach limits the number of free parameters for the Gaussian and von Mises distributions to two, to match the sinusoid. The results are shown in *Figure 11* and *Table 5*. Both the fixed-width Gaussian and von Mises distributions are a slightly better fit to the data than the sinusoid, but the differences between the three curves are very small.

In simplifying the fruit fly connectome data, we assumed all synapses of different types were of equal weight, as no data to the contrary were available. Different synapse types having different strengths could introduce nonlinear distortions between our net synaptic path count and the true synaptic strength, which could in turn make the data a better or worse fit for a sinusoidal compared to a Gaussian profile. As such, we don't consider the 2% relative difference in RMSE between the $f = 1$ sinusoid and fixed-width Gaussian and von Mises distributions to be conclusive.

Overall, we find that the cosine weights that emerge from our derivations are a very close match for the locust, but less precise for the fly, where other functions fit slightly better. Given the limitations in using the currently available data to provide an exact estimate of synaptic strength (for the locust), and due to the high variability of the synaptic count (for the fruit fly), we consider that our theory is compatible with the observed data.

Figure 11: Gaussian and von Mises distributions provide a better fit than the weight profile with sinusoidal weights of frequency 1, but the former two distributions have an extra model parameter. When fixing the width parameters of these distributions to match the sinusoid, all three distributions provide a very similar fit. Note that the -4 and 4 neuron indices are the same and just duplicated for visualisation purposes.

Curve	RMSE	AICc
$f = 1$	0.3213	8.05
Gaussian	0.1941	12.60
von Mises	0.1854	12.55
Gaussian (fixed stddev)	0.3142	7.98
von Mises (fixed kappa)	0.3178	8.02

Table 5: Quality of fit for different possible weight structures. The Root Mean Square Error is weighted by the variance attached to each connection (in units of standard deviations). The corrected Akaike Information Criterion (AICc) is also shown for the weight profiles in **Figure 11**. The best fit after correcting for model size is the Gaussian with fixed standard deviation, but it is very closely matched by the von Mises with fixed kappa and $f = 1$.

4.7 Convergence of Oja's rule with multiple harmonics

4.7.1 Linear neurons

Integrating Oja's learning rule updates (**Equation 19**) over all angles in the position space gives us the following:

$$\Delta \omega = \mathbf{a} * \mathbf{a} - \|\mathbf{a}\|^2 \omega, \tag{38}$$

which we can transform into the Fourier domain:

$$\mathcal{F}_f[\Delta \omega] = \mathcal{F}_f[\mathbf{a}]^2 - \|\mathbf{a}\|^2 \mathcal{F}_f[\omega]. \tag{39}$$

The steady state solution when the weight update is 0 is:

$$\mathcal{F}_f[\omega] = \frac{\mathcal{F}_f[\mathbf{a}]^2}{\|\mathbf{a}\|^2}. \tag{40}$$

By Parseval's theorem, $\|\mathbf{a}\|^2 = \sum_f \|\mathcal{F}_f[\mathbf{a}]\|^2$ so the Fourier spectrum of the steady state weights is just the normalised Fourier spectrum of the input:

$$\mathcal{F}_f[\omega] = \frac{\mathcal{F}_f[\mathbf{a}]^2}{\sum_f \|\mathcal{F}_f[\mathbf{a}]\|^2}. \tag{41}$$

From this we can see that the stable L_1 norm of $\mathcal{F}[\omega]$ is 1,

$$\sum_f |\mathcal{F}_f[\omega]| = \frac{1}{\sum_f \|\mathcal{F}_f[\mathbf{a}]\|^2} \sum_f |\mathcal{F}_f[\mathbf{a}]^2| = 1. \tag{42}$$

If we combine this result with **Equation 10**, we find that in the case of a single encoding frequency, $\mathcal{F}_{f^*}[\mathbf{a}] > 0$, Oja's rule will result in a single stable harmonic with $\mathcal{F}_{f^*}[\omega] = 1$ because $\sum_f \mathcal{F}_f[\mathbf{a}]^2 = \mathcal{F}_{f^*}[\mathbf{a}]^2$.

If Oja's rule has a normalising factor added to account for the number of encoding harmonics used, $|\mathbf{F}|$,

$$\Delta\omega = \mathbf{a} * \mathbf{a} - \frac{1}{|\mathbf{F}|}\|\mathbf{a}\|^2\omega, \tag{43}$$

then the steady state solution for the weights becomes

$$\mathcal{F}_f[\omega] = |\mathbf{F}|\frac{\mathcal{F}_f[\mathbf{a}]^2}{\sum_f \mathcal{F}_f[\mathbf{a}]^2}, \tag{44}$$

where the L_1 norm of $\omega = |\mathbf{F}|$. In this case multiple harmonics could develop stable values of $\mathcal{F}_f[\omega] = 1$, but only if the activity magnitudes for these harmonics are identical. If any perturbation affects the activities, the harmonic which happens to have the larger activity will begin to dominate the other by the dynamics of **Equation 43**, until only one remains.

Therefore, the only case where Oja's rule results in a stable solution robust to perturbations is when only one harmonic is used.

4.7.2 Neurons with nonlinear activations

An analysis similar to the previous case applies when the firing rate of the neurons has a nonlinear activation. For a given activity profile, **Equation 41** still applies. Thus, for an activity defined by the weight profile $\mathbf{a}$, it will generate the weight profile

$$\omega = \frac{1}{\|\mathbf{a}\|} \sum_f |\mathcal{F}_f[\mathbf{a}]|^2, \tag{45}$$

where we set $\|\omega\| = 1$ for convenience of notation. Note that the squaring process will induce a weight distribution in which the larger harmonic will be enhanced more than the rest.

Having obtained the weights, the equation for the activity profile can be solved through the self-consistency setting in which $u(t) = 0$ and the activity is maintained,

$$\mathbf{a}_n(\theta) = \phi\left(\sum_{m=0}^{N} \omega_m \mathbf{a}_{n-m}(\theta)\right) \tag{46}$$

where ϕ is the nonlinear activation function. Given the rotational symmetry from our assumptions, it is enough to solve for $\mathbf{a}_n = \mathbf{a}_n(0)$, which gives us a set of $N/2$ equations which can then be combined with **Equation 45**, which gives us the following equation to solve

$$\begin{aligned}
\mathbf{a}_n &= \phi\left(\sum_{m=0}^{N}\left(\sum_{f=1}^{N}\frac{\mathcal{F}_f[\mathbf{a}]^2}{\|\mathbf{a}\|^2}\cos\left(\frac{2\pi(n-m)}{N}\right)\right)\mathbf{a}_{n-m}\right) \\
&= \phi\left(\frac{1}{\|\mathbf{a}\|^2}\sum_{m=0}^{N}\left(\sum_{f=1}^{N}\left[\sum_{k=0}^{N}\mathbf{a}_k\cos\left(\frac{2\pi kf}{N}\right)\right]^2\cos\left(\frac{2\pi(n-m)}{N}\right)\right)\mathbf{a}_{n-m}\right)
\end{aligned} \tag{47}$$

Since $\mathbf{a}_n = \mathbf{a}_{-n \mod N}$, we obtain $N/2$ equations for $N/2$ values of $\mathbf{a}_n$.

Note that there is no closed form solution in general, but some generic properties can be inferred. For example, as long as we have a single bump of activity, $\mathcal{F}_1 [\mathbf{a}] > |\mathcal{F}_f [\mathbf{a}]|, \ \forall f > 1$, therefore $\mathcal{F}_1 [\omega] > \mathcal{F}_f [\omega], \ \forall f > 1$. More specifically,

$$\mathcal{F}_f [\omega] = \frac{|\mathcal{F}_f [\mathbf{a}]|^2}{|\mathcal{F}_1 [\mathbf{a}]|^2} \mathcal{F}_1 [\omega]. \tag{48}$$

Figure 12 shows the results of the same simulation as Section Learning rules and development, but using neurons with a nonlinear (tanh) activation function, and distorting the network input to be a more square wave-like sinusoid. As predicted theoretically, Oja's rule still reinforces the dominant spatial frequency in the network, causing the weights to converge to the sinusoidal profile our theory predicts.

Figure 12: **Recovery of synaptic weights with a Hebbian learning rule.** The synaptic weights converge to a sinusoidal pattern under Oja's rule when the neurons have a nonlinear (tanh) activation function and a distorted (more square wave-like) sinusoidal input is provided. **A:** The weights start at zero and slowly converge to the prescribed sinusoidal profile, showing that this connectivity can emerge from scratch. **B:** The sinusoidal weights are perturbed by noise but learning ensures that the weight profile is corrected. Noise corrupted the initial activity, weights, and input with the same parameters as in **Figure 5**, the only difference being that these simulations were run for 2500 periods.

Figure 7: The connectivity matrix for the fruit fly containing all 152 neurons. On the horizontal axis are the names of the pre-synaptic neurons while on the vertical axis the names of the post-synaptic neurons. Neurons within each cell type are ordered by the glomerulus they innervate and arbitrarily within glomerulus.

Figure 8: The connectivity matrix for the fruit fly neurons grouped by innervated glomerulus so containing 60 groups. On the horizontal axis are the names of the pre-synaptic neurons while on the vertical axis the names of the post-synaptic neurons. Neuron groups are ordered by the glomerulus they innervate.

Figure 9: The connectivity matrix for the fruit fly neurons grouped by glomerulus and aggregated from both hemispheres, so containing 34 groups. On the horizontal axis are the names of the pre-synaptic neuron groups while on the vertical axis are the names of the post-synaptic neuron groups. Neuron groups are ordered by the glomerulus they innervate.

Figure 10: The connectivity matrix for the fruit fly neurons grouped by glomerulus and aggregated from both hemispheres, with PEG9 & PEG1 neurons and EPG9 & EGP1 neurons grouped, so containing 32 groups in total. On the horizontal axis are the names of the pre-synaptic neuron groups while on the vertical axis the names of the post-synaptic neuron groups. Neuron groups are ordered by the glomerulus they innervate.

Author contributions

P.V.A. and I.P. designed and conceptualised the project. D.D. performed the simulations. P.V.A. and D.D. performed the analytical calculations. P.V.A. and I.P. designed the data analysis and D.D. implemented it. All authors wrote the manuscript.

Data availability

The code used for performing all analysis and generating all figures is available at https://github.com/Dominic-DallOsto/insect-navigation-sinusoidal-optimality.

Acknowledgements

We would like to thank Benjamin F. Grewe for his support and helpful comments.

References

Alon U. Simplicity in biology. Nature. 2007; 446(7135):497–497. https://doi.org/10.1038/446497a.

Barnes CL, Bonnéry D, Cardona A. Synaptic counts approximate synaptic contact area in Drosophila. Plos one. 2022; 17(4):e0266064. https://doi.org/10.1371/journal.pone.0266064.

Beetz MJ, Kraus C, Franzke M, Dreyer D, Strube-Bloss MF, Rössler W, Warrant EJ, Merlin C, el Jundi B. Flight-Induced Compass Representation in the Monarch Butterfly Heading Network. Current Biology. 2022; 32(2):338–349.e5. https://doi.org/10.1016/j.cub.2021.11.009.

Burnham KP, Anderson DR. Model selection and multimodel inference. A practical information-theoretic approach. 2004; 2. https://doi.org/10.1007/b97636.

Collett TS. Path integration: how details of the honeybee waggle dance and the foraging strategies of desert ants might help in understanding its mechanisms. Journal of Experimental Biology. 2019 May; 222(11):jeb205187. https://doi.org/10.1242/jeb.205187.

Cope AJ, Sabo C, Vasilaki E, Barron AB, Marshall JAR. A Computational Model of the Integration of Landmarks and Motion in the Insect Central Complex. PLoS ONE. 2017 Feb; 12(2):e0172325. https://doi.org/10.1371/journal.pone.0172325.

Cornford J, Kalajdzievski D, Leite M, Lamarquette A, Kullmann DM, Richards B. Learning to live with Dale's principle: ANNs with separate excitatory and inhibitory units. bioRxiv. 2020; p. 2020–11. https://doi.org/10.1101/2020.11.02.364968.

Dayan P, Abbott LF. Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems. Computational Neuroscience, Cambridge, Mass: Massachusetts Institute of Technology Press; 2001.

Dingle K, Camargo CQ, Louis AA. Input–output maps are strongly biased towards simple outputs. Nature communications. 2018; 9(1):761. https://doi.org/10.1038/s41467-018-03101-6.

ElJundi B, Baird E, Byrne MJ, Dacke M. The brain behind straight-line orientation in dung beetles. Journal of Experimental Biology. 2019; 222. https://doi.org/10.1242/jeb.192450.

Green J, Adachi A, Shah KK, Hirokawa JD, Magani PS, Maimon G. A Neural Circuit Architecture for Angular Integration in *Drosophila*. Nature. 2017 May; 546(7656):101–106. https://doi.org/10.1038/nature22343.

Haferlach T, Wessnitzer J, Mangan M, Webb B. Evolving a neural model of insect path integration. Adaptive Behavior. 2007; 15(3):273–287. https://doi.org/10.1177/1059712307082080.

Hartmann G, Wehner R. The ant's path integration system: a neural architecture. Biological cybernetics. 1995; 73(6):483–497. https://doi.org/10.1007/s004220050204.

Heinze S, Florman J, Asokaraj S, el Jundi B, Reppert SM. Anatomical Basis of Sun Compass Navigation II: The Neuronal Composition of the Central Complex of the Monarch Butterfly. Journal of Comparative Neurology. 2013 Feb; 521(2):267–298. https://doi.org/10.1002/cne.23214.

Heinze S, Gotthardt S, Homberg U. Transformation of polarized light information in the central complex of the locust. The Journal of neuroscience: the official journal of the Society for Neuroscience. 2009 sep; 29(38):11783–93. https://doi.org/10.1523/JNEUROSCI.1870-09.2009.

Homberg U, Heinze S, Pfeiffer K, Kinoshita M, El Jundi B. Central neural coding of sky polarization in insects. Philosophical Transactions of the Royal Society B: Biological Sciences. 2011; 366(1565):680–687. https://doi.org/10.1098/rstb.2010.0199.

Honkanen A, Adden A, da Silva Freitas J, Heinze S. The insect central complex and the neural basis of navigational strategies. Journal of Experimental Biology. 2019; 222(Suppl_1):jeb188854. https://doi.org/10.1242/jeb.188854.

Huang L, Ledochowitsch P, Knoblich U, Lecoq J, Murphy GJ, Reid RC, de Vries SE, Koch C, Zeng H, Buice MA, et al. Relationship between simultaneously recorded spiking activity and fluorescence signal in GCaMP6 transgenic mice. eLife. 2021; 10:e51675. https://doi.org/10.7554/eLife.51675.

Hulse BK, Haberkern H, Franconville R, Turner-Evans DB, Takemura Sy, Wolff T, Noorman M, Dreher M, Dan C, Parekh R, Hermundstad AM, Rubin GM, Jayaraman V. A connectome of the Drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. eLife. 2021; 10. https://doi.org/10.7554/elife.66039.

Johnston IG, Dingle K, Greenbury SF, Camargo CQ, Doye JP, Ahnert SE, Louis AA. Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution. Proceedings of the National Academy of Sciences. 2022; 119(11):e2113883119. https://doi.org/10.1073/pnas.2113883119.

el Jundi B, Pfeiffer K, Heinze S, Homberg U. Integration of polarization and chromatic cues in the insect sky compass. Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology. 2014 Mar; 200(6):575–589. https://doi.org/10.1007/s00359-014-0890-6.

Kakaria KS, de Bivort BL. Ring Attractor Dynamics Emerge from a Spiking Model of the Entire Protocerebral Bridge. Frontiers in Behavioral Neuroscience. 2017 Feb; 11:8. https://doi.org/10.3389/fnbeh.2017.00008.

Kim SS, Rouault H, Druckmann S, Jayaraman V. Ring attractor dynamics in the Drosophila central brain. Science. 2017 may; 356(6340):849–853. https://doi.org/10.1126/science.aal4835.

Kinoshita M, Pfeiffer K, Homberg U. Spectral properties of identified polarized-light sensitive interneurons in the brain of the desert locust Schistocerca gregaria. The Journal of experimental biology. 2007; 210(Pt 8):1350–61. https://doi.org/10.1242/jeb.02744.

Kriegeskorte N, Golan T. Neural Network Models and Deep Learning. Current Biology. 2019 Apr; 29(7):R231–R236. https://doi.org/10.1016/j.cub.2019.02.034.

Labhart T. Polarization-opponent interneurons in the insect visual system. Nature. 1988; 331(6155):435–437. https://doi.org/10.1038/331435a0.

Labhart T. Polarization-Sensitive Interneurons in the Optic Lobe of the Desert Ant Cataglyphis bicolor. Naturwissenschaften. 2000; 87(3):133–136. https://doi.org/10.1007/s001140050691.

Li P, Cornford J, Ghosh A, Richards B. Learning better with Dale's Law: A Spectral Perspective. bioRxiv. 2023; https://doi.org/10.1101/2023.06.28.546924.

Liu TX, Davoudian PA, Lizbinski KM, Jeanne JM. Connectomic Features Underlying Diverse Synaptic Connection Strengths and Subcellular Computation. Current Biology. 2022 Feb; 32(3):559–569.e5. https://doi.org/10.1016/j.cub.2021.11.056.

Loesel R, Homberg U. Anatomy and physiology of neurons with processes in the accessory medulla of the cockroach Leucophaea maderae. Journal of Comparative Neurology. 2001; 439(2):193–207. https://doi.org/10.1002/cne.1342.

Lyu C, Abbott LF, Maimon G. Building an allocentric travelling direction signal via vector computation. Nature. 2022; 601(7891):92–97. https://doi.org/10.1038/s41586-021-04067-0.

Mappes M, Homberg U. Behavioral Analysis of Polarization Vision in Tethered Flying Locusts. Journal of Comparative Physiology A. 2004 Jan; 190(1):61–68. https://doi.org/10.1007/s00359-003-0473-4.

Menzel R, Geiger K, Chittka L, Joerges J, Kunze J, Müller U. The Knowledge Base of Bee Navigation. Journal of Experimental Biology. 1996; 199(1):141–146. https://doi.org/10.1242/jeb.199.1.141.

Merlin C, Heinze S, Reppert SM. Unraveling Navigational Strategies in Migratory Insects. Current Opinion in Neurobiology. 2012 Apr; 22(2):353–361. https://doi.org/10.1016/j.conb.2011.11.009.

Mittelstaedt H. Analytical Cybernetics of Spider Navigation. In: Barth FG, editor. *Neurobiology of Arachnids* Springer Berlin Heidelberg; 1985.p. 298–316. https://doi.org/10.1007/978-3-642-70348-5_15.

Müller M, Wehner R. Path integration in desert ants, *Cataglyphis fortis*. Proceedings of the National Academy of Sciences. 1988; 85(14):5287–5290. https://doi.org/10.1073/pnas.85.14.5287.

Pfeiffer K, Homberg U. Coding of Azimuthal Directions via Time-Compensated Combination of Celestial Compass Cues. Current Biology. 2007 jun; 17(11):960–965. https://doi.org/10.1016/J.CUB.2007.04.059.

Pfeiffer K, Kinoshita M, Homberg U. Polarization-sensitive and light-sensitive neurons in two parallel pathways passing through the anterior optic tubercle in the locust brain. Journal of neurophysiology. 2005; 94(6):3903–15. https://doi.org/10.1152/jn.00276.2005.

Pisokas I, Heinze S, Webb B. The Head Direction Circuit of Two Insect Species. eLife. 2020 Jul; 9:e53985. https://doi.org/10.7554/eLife.53985.

Sakura M, Lambrinos D, Labhart T. Polarized skylight navigation in insects: model and electrophysiology of e-vector coding by neurons in the central complex. Journal of neurophysiology. 2008; 99(2):667–82. https://doi.org/10.1152/jn.00784.2007.

Sayre ME, Templin R, Chavez J, Kempenaers J, Heinze S. A projectome of the bumblebee central complex. eLife. 2021; 10:e68911. https://doi.org/10.7554/elife.68911.

Seelig JD, Jayaraman V. Neural Dynamics for Landmark Orientation and Angular Path Integration. Nature. 2015 May; 521(7551):186–191. https://doi.org/10.1038/nature14446.

Siegle JH, Ledochowitsch P, Jia X, Millman DJ, Ocker GK, Caldejon S, Casal L, Cho A, Denman DJ, Durand S, et al. Reconciling functional differences in populations of neurons recorded with two-photon imaging and electrophysiology. eLife. 2021; 10:e69068. https://doi.org/10.7554/eLife.69068.

Stone T, Webb B, Adden A, Weddig NB, Honkanen A, Templin R, Wcislo W, Scimeca L, Warrant E, Heinze S. An Anatomically Constrained Model for Path Integration in the Bee Brain. Current Biology. 2017; 27(20):3069–3085. https://doi.org/10.1016/j.cub.2017.08.052.

Stringer S, Trappenberg T, Rolls E, Araujo I. Self-organizing continuous attractor networks and path integration: one-dimensional models of head direction cells. Network: Computation in Neural Systems. 2002; 13(2):217–242. https://doi.org/10.1080/net.13.2.217.242.

Su TS, Lee WJ, Huang YC, Wang CT, Lo CC. Coupled symmetric and asymmetric circuits underlying spatial orientation in fruit flies. Nature Communications. 2017; 8(1). https://doi.org/10.1038/s41467-017-00191-6.

Taube JS. The head direction signal: origins and sensory-motor integration. Annu Rev Neurosci. 2007; 30:181–207. https://doi.org/10.1146/annurev.neuro.29.051605.112854.

Touretzky DS, Redish AD, Wan HS. Neural Representation of Space Using Sinusoidal Arrays. Neural Computation. 1993 nov; 5(6):869–884. https://doi.org/10.1162/neco.1993.5.6.869.

Turner-Evans D, Wegener S, Rouault H, Franconville R, Wolff T, Seelig JD, Druckmann S, Jayaraman V. Angular Velocity Integration in a Fly Heading Circuit. eLife. 2017 May; 6:2112–2126. https://doi.org/10.7554/elife.23496.

Turner-Evans DB, Jensen KT, Ali S, Paterson T, Sheridan A, Ray RP, Wolff T, Lauritzen JS, Rubin GM, Bock DD, Jayaraman V. The Neuroanatomical Ultrastructure and Function of a Biological Ring Attractor. Neuron. 2020; 108(1):145–163.e10. https://doi.org/10.1016/j.neuron.2020.08.006.

Vafidis P, Owald D, D'Albis T, Kempter R. Learning accurate path integration in ring attractor models of the head direction system. eLife. 2022; 11:e69841. https://doi.org/10.7554/eLife.69841.

Vickerstaff RJ, Paolo EAD. Evolving neural models of path integration. Journal of Experimental Biology. 2005; 208(17):3349–3366. https://doi.org/10.1242/jeb.01772.

Warren TLTL, Giraldo YMYM, Dickinson MHMH. Celestial Navigation in *Drosophila*. Journal of Experimental Biology. 2019 Feb; 222:jeb186148. https://doi.org/10.1242/jeb.186148.

Wehner R, Srinivasan MV. Path Integration in Insects. In: Jeffery KJ, editor. *The Neurobiology of Spatial Behaviour* Oxford University Press; 2003.p. 9–30. https://doi.org/10.1093/acprof:oso/9780198515241.003.0001.

Wessnitzer J, Haferlach T, Mangan M, Webb B. Path Integration Using a Model of e-Vector Orientation Coding in the Insect Brain: Reply to Vickerstaff and Di Paolo. Adaptive Behavior. 2008; 16(4):277–280. https://doi.org/10.1177/1059712308094947.

Wittmann T, Schwegler H. Path integration - a network model. Biological Cybernetics. 1995; 73(6):569–575. https://doi.org/10.1007/bf00199549.

Zador AM. A critique of pure learning and what artificial neural networks can learn from animal brains. Nature communications. 2019; 10(1):3770. https://doi.org/10.1038/s41467-019-11786-6.

Zhang K. Representation of spatial orientation by the intrinsic dynamics of the head-direction cell ensemble: a theory. Journal of Neuroscience. 1996; 16(6):2112–2126. https://doi.org/10.1523/JNEUROSCI.16-06-02112.1996.