

Published in final edited form as:

Neural Comput. 2022 September 09; 34(11): 2205–2231. doi:10.1162/neco_a_01540.

A neural model for insect steering applied to olfaction and path integration

Andrea Adden^{1,*}, Terrence C. Stewart², Barbara Webb³, Stanley Heinze^{1,4}

¹Vision Group, Department of Biology, Lund University, Lund, Sweden

²National Research Council of Canada, University of Waterloo Collaboration Centre, Waterloo, Canada

³Insect Robotics Group, School of Informatics, University of Edinburgh, Edinburgh, UK

⁴NanoLund, Lund University, Lund, Sweden

Abstract

Many animal behaviours require orientation and steering with respect to the environment. For insects, a key brain area involved in spatial orientation and navigation is the central complex. Activity in this neural circuit has been shown to track the insect's current heading relative to its environment, and has also been proposed to be the substrate of path integration. However, it remains unclear how the output of the central complex is integrated into motor commands. Central complex output neurons project to the lateral accessory lobes (LAL), from which descending neurons project to thoracic motor centres. Here, we present a computational model of a simple neural network that has been described anatomically and physiologically in the LALs of male silkworm moths, in the context of odour-mediated steering. We present and analyze two versions of this network, one rate-based and one based on spiking neurons. The modelled network consists of an inhibitory local interneuron and a bistable descending neuron ('flip-flop'), which both receive input in the LAL. The flip-flop neuron projects onto neck motor neurons to induce steering. We show that this simple computational model not only replicates the basic parameters of male silkworm moth behaviour in a simulated odour plume, but can also take input from a computational model of path integration in the central complex and use it to steer back to a point of origin. Furthermore, we find that increasing the level of detail within the model improves the realism of the model's behaviour, leading to the emergence of looping behaviour as an orientation strategy. Our results suggest that descending neurons originating in the LALs, such as flip-flop neurons, are sufficient to mediate multiple steering behaviours. This study is therefore a first step to close the gap between orientation circuits in the central complex and downstream motor centres.

Keywords

navigation; orientation; neural network

*Current address: Neural circuits and evolution lab, The Francis Crick Institute, London, UK

1 Introduction

Insects display an astonishing range of behaviours that include highly directed movements. For example, male moths navigate towards females emitting pheromones (Marsh et al., 1978; Obara, 1979) and female crickets move towards singing males (Simmons, 1988; Balakrishnan and Pollack, 1996). Other insects use visual cues to maintain a straight heading over short or long distances (dung beetle: (Dacke et al., 2003b; Baird et al., 2012); Monarch butterfly: (Mouritsen and Frost, 2002); Bogong moth: (Dreyer et al., 2018)), and can even rely purely on memory to navigate home (Srinivasan, 2015; Honkanen et al., 2019). While the cues used for navigation are different in these examples, they elicit very similar behaviours: upon encountering an appropriate stimulus, the animal chooses a direction with respect to that stimulus and begins moving in that direction. If the stimulus is temporarily lost, searching behaviour is initiated. Thus, the motor patterns elicited by different kinds of stimuli can be remarkably similar.

The spatial context for orientation and navigation is computed in the central complex (CX), the only unpaired and midline-spanning neuropil in the insect brain (Fig. 1; (Heinze, 2017)). In recent years, progress has been made in understanding and modelling this ‘compass system’ of insects. Neurons in the CX integrate external inputs with self-generated angular velocity cues, thus providing a reliable internal representation of the animal’s heading (Green et al., 2017; Kakaria and de Bivort, 2017). An extended model of the CX network furthermore showed that the CX is a possible substrate for path integration (Stone et al., 2017), i.e., continuously integrating velocity to maintain an estimate of the direction and distance to a reference location. This model also demonstrates how CX output can serve directly as a steering command: the summed activity of columnar output neurons in each hemisphere is compared and any imbalance between the two hemispheres should produce a turn towards the relevant side, while a balanced output results in straight movement. However, the model does not postulate a biologically plausible mechanism to achieve steering, and instead summarises the entire steering system as one theoretical motor command. Despite our increasingly complete understanding of CX networks, a question that remains unanswered is how CX output is translated into motor control; i.e. how it is transmitted to thoracic motor centres to influence behavioural decisions.

The answer likely lies in the lateral accessory lobes (LAL), a paired neuropil flanking the CX. The LAL forms part of the pheromone-processing pathway that has been described in detail in the silk-worm moth (Namiki et al., 2014; Kymre et al., 2021). While the CX integrates information from modalities such as vision and mechanoreception, it does not appear to process olfactory information. Odours, including pheromones, are integrated in higher-order neuropils including the lateral horn, mushroom bodies and finally the LAL. The LAL has been described as a pre-motor centre, as several types of descending neurons that project to thoracic motor centres have post-synaptic endings in the LAL (Namiki et al., 2018a). It has been suggested that interactions between the two LALs are sufficient to mediate search behaviours triggered by e.g. pheromone input, but that goal-oriented behaviours require CX input into the LAL network (Namiki and Kanzaki, 2016; Álvarez-Salvado et al., 2018). Indeed, CX output neurons project to the LALs, but how they interact with the LAL circuitry is currently unknown. Additionally, how descending neurons

encode motor commands on a population level is currently not well understood, although multiple recent studies have been able to dissect single neural circuits that underlie specific behaviours (e.g. (Schnell et al., 2017; Cande et al., 2018; Namiki et al., 2018b)).

One such behaviour, which has been examined in detail, is the pheromone-following behaviour of silkworm moths. Male silkworm moths display a highly stereotyped behavioural sequence when following a female's pheromone plume. Upon first contact with the plume, the moth responds with a 'surge', that is a straight movement towards the source of the odour. When the odour plume is lost, several ways of re-acquiring the plume have been described; most notably 'casting', during which the moths walk in a zig-zag pattern until it finds a new odour pocket, and looping downwind (review: (Cardé and Willis, 2008)). Early studies have identified several descending neuron types whose activity correlates with turning behaviour when a male moth orients in a pheromone plume (Mishima and Kanzaki, 1999). Among these, the most notable are "flip-flop" neurons, which are bistable neurons that switch between a high-activity and a low-activity state in response to a trigger stimulus (Olberg, 1983; Kanzaki et al., 1994; Mishima and Kanzaki, 1999). That is, the same stimulus can cause the neuron to increase or decrease its firing, depending on whether it is in the low or high activity state, respectively, when that stimulus occurs.

These neurons have post-synaptic terminals in the LALs, and their axons descend through the ventral nerve cord and synapse onto neck motor neurons, which in turn activate neck muscles that control head movements (Kanzaki and Mishima, 1996; Mishima and Kanzaki, 1998). Thus, if the left-descending flip-flop neuron is in its high-activity state, the left neck motor neuron and the left neck muscle are also active, causing the head and consequently the moth to turn left. Although this network has been described in the context of pheromone following, other studies have shown that flip-flop neurons can also be triggered by light flashes (Olberg, 1983) and sound (Zorovic and Hedwig, 2011). It therefore seems likely that flip-flop neuron mediated steering may constitute a general form of targeted steering, independent of the stimulus modality that drives the behaviour (Steinbeck et al., 2020).

In this study, we aim to evaluate whether a basic flip-flop neuron network can produce naturalistic steering in (a) a simulated odour plume and (b) when presented with compass input via a neural model of the central complex. To this end, we present a rate-based and a spiking computational model of a simple flip-flop network. The two models presented here allow us to compare two different implementations of the same network, where both models follow the same connectivity pattern. The rate-based model uses continuous-valued sigmoid neurons, while the spiking model uses leaky integrate-and-fire (LIF) spiking neurons. Both models are effective at navigating in the simulated olfactory and visual tasks, however the spiking model produces more realistic trajectories. Comparing the behaviour of our models to behavioural data from male silkworm moths, we find that this simple flip-flop based neural circuit is sufficient to replicate the basic characteristics of the moths' paths. Furthermore, we describe looping behaviour as an emerging orientation strategy when sensory input is directionally ambiguous. Finally, we demonstrate for the first time that the flip-flop network can work as a general steering network when combined with a computational model of the CX (Stone et al., 2017). This study is therefore a step towards

closing the gap between higher processing centres in the brain that make navigational decisions, such as the CX, and the thoracic motor circuits that ultimately move the insect.

Materials and methods

The neurons modelled here were physiologically and anatomically described in (Olberg, 1983; Kanzaki et al., 1992; Mishima and Kanzaki, 1999; Kanzaki et al., 1994). We infer input and output regions of neurons from their anatomical appearance, i.e. smooth terminals are assumed to be inputs, while varicose terminals are assumed to be outputs. Two neurons are assumed to be connected if the input region of one neuron overlaps with the second neuron's output region. The model connections are furthermore based on the network proposed in (Mishima and Kanzaki, 1999) with some small modifications.

The network consists of two pairs of neurons: One flip-flop neuron (FF) and one protocerebral bilateral neuron (PBN; Fig 1A) per hemisphere. Both cell types receive input directly from the plume, or from the output neurons of the central complex (CPU1/PFL neurons) when connected to the path integration network (PI). PBN neurons were proposed to provide bilateral inhibition between the two LALs (Mishima and Kanzaki, 1999; Kanzaki et al., 2005) and are therefore modelled to inhibit the contralateral FF neuron. FF neurons have excitatory connections directly onto the contralateral motor, based on the finding that FF neuron activity correlates with neck motor neuron activity (Kanzaki and Mishima, 1996).

We present two different implementations of this network, which we call 'rate-based model' and 'spiking model'. The rate-based model uses continuous-valued sigmoid neurons, which is the most common neural model used in insect brain network models. However, while such neurons are easier to work with and build models from, they abstract away many of the underlying biological details. To complement this, we also present a spiking model built from leaky integrate-and-fire (LIF) neurons, bringing richer temporal dynamics to the model. While still abstracted, we will show that the spiking model leads to more realistic behavior.

The general approach for creating the spiking model is known as the Neural Engineering Framework (NEF; Eliasmith and Anderson (2003)). This allows us to take any particular desired dynamics (such as the dynamics of a flip-flop neuron) and construct it using more basic components (such as a group of spiking LIF neurons). In particular, here we approximate one FF neuron using 100 recurrently-connected spiking LIF neurons. The NEF treats this as an optimization problem and finds the ideal connection weights among those components, such that the overall system produces behavior that is as close as possible to the desired rate-based description of a FF neuron. Our interpretation of the model is that this is one flip-flop neuron with 100 internal components, and the flip-flop behaviour arises out of these internal interactions.

Both models are described in detail below. The source code is available at https://github.com/stanleyheinze/insect_steering. The spiking model was implemented using the software toolkit Nengo (Bekolay et al., 2014).

Models

Rate-based model—For the rate-based model, we use continuous-valued sigmoid neurons, with one addition described below for the FF neurons. If the total input to the neuron is J , then the output r from the neuron is generally defined as

$$r = 1/(1 + e^{-aJ - b}) \quad (1)$$

where a and b are the gain and bias constants for the neuron, respectively. The total input to the neuron is the weighted sum of the rates of all incoming connections

$$J_j = \sum w_{ij}r_i \quad (2)$$

The weights w are either 0 (no connection), 1 (excitation), or -1 (inhibition), with gaussian noise of standard deviation 0.01 added when the model is created. We also add random noise $\epsilon \hookrightarrow N(0, \sigma^2)$ with $\sigma^2 = 0.02$ to the input to each neuron on every time step. These two sources of randomness are meant to give some individual variation to the models. For both PBN and FF neurons, we also add the neurons' previous output to its own input, to allow for the sustained activity that has been described for these neurons. This gives them the following equation, where r_t is the output for time step t , which is what we use in our model.

$$r_t = 1/(1 + e^{-a(J + r_{t-1} + \epsilon) - b}) \quad (3)$$

For the FF neurons, we added a mechanism to produce flip-flop behaviour, where an input stimulus will switch a neuron from a high state to a low state and vice versa, depending on the state of the neuron preceding the stimulus. Since sigmoid neurons by themselves are too simple to produce this behaviour, we added a feature to the rate-based model where if the current output r is large (>0.8) and the input is large (>0.5), then the output of this FF neuron is reduced by 0.5 and the opposite FF is increased by 0.5. This produces the required flip-flop behaviour (Fig S1), but does not postulate a plausible mechanism whereby this behaviour is produced. We present a more realistic mechanism in the next section on the spiking model.

Spiking model—As described above, each neuron in the spiking model is represented by a group of 100 LIF neurons. The NEF approach to doing this is to start by writing out the particular desired dynamics as a differential equation $\frac{dx}{dt} = f(x, u)$ where x is the internal state and u is the input. In the case of our flip-flop neuron, we can write the flip-flop mechanism as given in Equation 4.

$$\frac{dx}{dt} = \begin{cases} 1 - x, & \text{if } x > 0.5 \text{ and } u < 0.2 \text{ (if FF is high and no input, stay high)} \\ 0 - x, & \text{if } x > 0.5 \text{ and } u > 0.2 \text{ (if FF is high and input, go low)} \\ 0 - x, & \text{if } x < 0.5 \text{ and } u < 0.2 \text{ (if FF is low and no input, stay low)} \\ 1 - x, & \text{if } x < 0.5 \text{ and } u > 0.2 \text{ (if FF is low and input, go high)} \end{cases} \quad (4)$$

This differential equation will give similar behaviour to the sigmoid FF neuron described above. In particular, note that if a differential equation is outputting a value of the form $T - x$, then the value x will move towards T until $T = x$. This means that in the first and last cases above, the state of the FF neuron will go to 1 (i.e. the FF is high), and in the second and third cases, the FF neuron will go to 0 (low). This gives us a differential equation approximation of the sigmoid FF.

The next step is to train a feed-forward single-hidden-layer neural network that approximates $y = \tau f(x, u) + x$. Importantly, we use LIF neurons in the hidden layer, but we have no non-linearity at all in the input and output layers. These hidden layer neurons will eventually become the recurrently connected internal components that will approximate a FF neuron. To train this network, the training data is a set of randomly generated x values and the y values are generated using Equation 4. The τ parameter is the time constant of the post-synaptic currents that will be used in the next step, which we choose here to be 100 ms. While any neural network training approach can be used (e.g. back-propagation of error), here we take the simple approach of randomly generating the first layer of weights and using regularized leastsquares minimization (i.e. ridge regression) to find the second layer of weights. This is a fast approach which works for any neuron model and avoids the difficulties of applying back-propagation to spiking neurons (Eliasmith and Anderson, 2003).

Given this network, we now connect the output of the network *back to its own input*, and include an exponential synapse model ($h(t) = e^{-t/\tau}/\tau$ for $t > 0$; i.e. every time a neuron spikes, the current it produces in the neurons it is connected to follows an exponential decay). Surprisingly, the resulting recurrent neural network will approximate the desired $\frac{dx}{dt} = f(x, u)$. To prove this, we note that the synapse will have the effect of convolving the output of the network with $h(t)$. Since this output is also the input to the network, we have $x(t) = (\tau f(x(t)) + x(t)) * h(t)$. Taking the Laplace transform, $X(s) = (\tau F(s) + X(s))H(s)$. Since the Laplace transform of the synapse model is $H(s) = 1/(1 + s\tau)$, this gives:

$$\begin{aligned} X(s) &= (\tau F(s) + X(s))/(1 + s\tau) \\ X(s)(1 + s\tau) &= \tau F(s) + X(s) \\ X(s) + s\tau X(s) &= \tau F(s) + X(s) \\ s\tau X(s) &= F(s) \end{aligned}$$

Finally, converting back to the time domain, we get $\frac{dx}{dt} = f(x, u)$. That is, this recurrently connected set of LIF neurons will approximate the desired flip-flop dynamics. Furthermore, since our feed-forward network has no non-linearity at the input or output, we can convert the system into a single pool of 100 LIF neurons whose recurrent connection weights are given by $W = W_{out} W_{in}$, where W_{in} is the weights from the input to the hidden layer and W_{out} is the weights from the hidden layer to the output.

This overall method (Eliasmith and Anderson, 2003) lets us take any desired dynamics (written as a differential equation) and convert it into a pool of basic components (here we use 100 LIF neurons) that are connected to each other using exponential synapses. As the

number of components increases, the accuracy of this approximation will improve, up to the limit of how well the feed-forward neural network approximates $y = \tau f(x, u) + x$. We can think of the resulting system as defining a spiking attractor network whose dynamics are governed by that function.

Note that we are not suggesting the flip-flop neurons must have recurrent connections, but rather that these neurons have a sufficiently complex internal state that such feedback is necessary to capture their properties. The resulting spiking model produces a notably different time course of response compared to the rate-based model, although the key qualitative characteristics are maintained. One crucial difference between the rate-based and spiking models is in how synchronization between the two FF cells is achieved. In the rate-based model, we have explicitly added a rule that forces the contralateral flip-flop to its up-state when an input pushes the ipsilateral flip-flop to the down-state, and vice versa. This is not necessary in the spiking model, as the internal components of a FF neuron can successfully switch the activity state of the neuron. However, there is no way for the internal components of one FF neuron to directly affect the contralateral FF neuron. There is, however, an inhibitory connection between the PBN and the contralateral FF, so this allows the system to push the contralateral FF to a low state when the ipsilateral FF and PBN get input from the sensor. This allows for some asynchrony to emerge in the spiking model, rather than forcing it as in the rate-based model. Note that this type of asynchronous flip-flopping behaviour can be most easily observed when using artificial input (Fig. S1A). It becomes less obvious when the agent is exposed to a naturalistic stimulus, such as the simulated odour plume (Fig. 1). In this more complex scenario, it is evident that the spiking model is only an approximation of the ideal equations given, and as such it is possible for both FFs to be high or low at the same time.

Experimental situation 1: Following an odour plume

Our model is directly inspired by the flip-flop neurons that have been implicated in pheromone tracking in moths, hence we first evaluate its ability to control the behaviour of a simulated agent in an odour plume (Fig 2A). To simulate a realistic plume, we use an efficient model of odor dispersion in a turbulent medium (Farrell et al., 2002), as implemented with the python-based module *pompy* (<https://github.com/InsectRobotics/pompy>, by Matthew Graham, Insect Robotics Group, Edinburgh University). This model was validated against measured gypsy moth plume data (Jones, 1983) and generates meandering plumes (Farrell et al., 2002), while still being efficient enough to be run at a very fine time scale (5 ms) to capture the high-frequency changes observed in moth plumes (≈ 0.1 s; Mafra-Neto and Cardé (1994); Vickers et al. (2001); Levakova et al. (2018)). The plume was dispersed by a weak constant wind (2 m/s) flowing from the direction of the odour plume source towards the agent's starting point. The agent was equipped with two fronto-lateral sensors which were designed to mimic antennae. The antenna size and angle (away from the centre line) from each other was adjusted to match real moths (Loudon and Koehl (2000); see table 1 for moth parameters). We set a maximum antenna sensitivity, above which the response is saturated, i.e. any higher concentration does not elicit a stronger response. We then chose a maximum response value to correspond to this level of concentration, and scaled the response linearly between this value and zero for lower

concentrations. These two parameters were adjusted manually. At each point in time, the network processes the input and makes a steering decision, which is applied to the centre of mass of the agent and scaled by the maximum rotation speed (table 1), analogous to (Stone et al., 2017). The agent's walking speed was determined by an acceleration and a drag parameter, analogous to (Stone et al., 2017) (see table 1). The agent's acceleration setting in the simulation was adjusted to match the actual walking speed of moths (Loudon and Koehl, 2000). The resulting forward speed was not influenced by the headwind, but agents were given a weak tendency to turn upwind in the absence of plume input, in accordance with observations in silkworm moths (Cardé and Willis, 2008). Unless otherwise specified, all parameters were tuned manually to match the behaviour of real moths as closely as possible. Certain arbitrary parameters of the models, such as the maximum output value for the antennae and the maximum rotation speed, were set such that the resulting tracks approximated the tracks of silkworm moths as reported in (Ando et al., 2013).

Ando et al. (2013) presented a robot that was steered through an odor plume by an on-board moth walking on a track ball. The moth's movement on the trackball was translated into wheel speeds for the robot. The authors presented trajectories for both the moth-controlled robot and the moth only in the odor plume. Here, we adjusted the parameters of this model such that the agent's trajectories were similar to the moths' trajectories presented in (Ando et al., 2013), based not only on the trajectories, but also quantitatively on the turn duration, turn angle and turn velocity.

Using their robot, Ando and colleagues performed further experiments where they added a bias to the turning of the robot. That is, a constant signal was added to the left (or right) wheel, while the moth was controlling the robot. This causes the moth to drift to the edge of the pheromone plume, but they show the moth is able to compensate for this and continue to follow the plume. We performed the same experiment in our simulation by adding a constant bias (in rad/s) to the moth rotation, causing the simulated agent to have an extra tendency to turn in a particular direction. Ando and colleagues show that moths' plume-following behaviour is robust to this sort of manipulation, and that the resulting paths tend to follow along the edge of the plume.

Behavioural measurement thresholds were defined in accordance with (Ando et al., 2013), to allow for comparing the models to data from male silkworm moths. A turn was identified if the agent's turn duration was larger than 0.5 s, the turn angular velocity was larger than 5 deg/s, and the turn angle was larger than 30° (Fig 2B, C; Ando et al. (2013)). Loops were detected based on their high rotation rate, above 30 deg/s for at least 5 s. An experiment was considered successful if the agent arrived within 5 cm of the goal. All data was analysed in Python 3.5.5.

Experimental situation 2: Path integration

A second motivation for our model was to understand how output from the central complex is translated to steering behaviour. Specifically, we connected our steering system to a previously developed CX path integrator model using sigmoid neurons (Stone et al., 2017). This model is based on the neuroanatomy of the central complex, and we take the output neurons from the model (CPU1/PFL; figure 6A) and project them to the lateral accessory

lobes in our model, where they may interact with the flip-flop descending neurons as well as the protocerebral bilateral neurons modelled here. Since there are eight CPU1 cells per hemisphere whose summed activity is thought to activate the motor, but only one flip-flop neuron in our model, the activity of all 8 neurons was summed by projecting onto the same flip-flop cell.

In the plume experiments, the value of each sensor was between 0 and 1, so that the absolute difference between the two sensors could also fall between 0 and 1. However, the output from the path integrator had a narrower spread (0.5 to 1). While our models worked with that smaller difference, we also re-scaled the path integrator output to a scale of 0 to 1 to test whether this would improve the models' behaviours.

To evaluate the behaviour in this situation, we used exactly the same simulator as before (pompy), but removed the odour plume. We then caused the agent to take a random exploratory path by setting its rotation rate to be a Gaussian white noise process with $\sigma = 0.1 \text{ rad/s}$ while moving forward at a constant speed (see table 1), for 15 seconds in total. The agent then attempted to return directly home using the CX output combined with our steering model. The simulation was continued for another 40 seconds. An experiment was considered successful if the agent arrived within 5 cm of the starting location.

Results

In order to determine how realistic the behaviour of our model is, we compared the simulated tracks quantitatively to data from silkworm moths (*Bombyx mori*, originally published in (Ando et al., 2013), Fig. 3). The total turn duration of the rate-based model falls within the standard deviation of real moth data but was slightly higher for the spiking model (Fig. 3A). The total turn angle of both models fell within the standard deviation of real moth data (Fig. 3C), while the mean turn velocity of both models was 10-15 % lower compared to real moths (Fig. 3B). However, both models performed well with respect to finding the origin of the plume, with a success probability of 0.84-1.0 (Fig. 3D). The tracks of both models display both cross-wind zigzagging and straight surges, as well as loops (Fig. 3E, F), with the rate-based model having straighter paths than the spiking model. Overall, when comparing the rate-based and spiking models to moth data, we find that both models replicate real moth data reasonably well.

A further way of testing how well our models replicate real moth behaviour was to add a turn bias to the simulation. When given a turn bias, silkworm moths were shown to track the edges of the odour plume instead of the centre (Ando et al., 2013). This was also the case for our models (Fig. 4). When analysing the angle between the current position of the agent and the source of the odour plume, we found that the models shift away from the centre of the plume already at a turn bias of 1 rad/s (Fig. 4A, B). With increasing turn bias, the models' success rates decrease, and at a turn bias of 3 rad/s, no agent simulated by the rate-based model reaches the goal. The spiking model is more robust but starts failing at a turn bias of 5 rad/s (data not shown). While these results are more difficult to compare quantitatively to real moth data owing to the different ways of implementing the turn bias, it is clear that

qualitatively the models behave very similarly to real moths, as they track the edge of the plume rather than the centre when given a turn bias.

Interestingly, looping behaviour emerges from the spiking model, but is only rarely observed in the rate-based model (Fig. 5A, B). In fact, we were unable to find ways of modifying the rate-based model to create these tight loops while still being able to successfully follow the plume (Fig. S3, S4). It should be noted that, while it is possible to observe clean flip-flop behaviour in ideal circumstances (Fig S1A), with more realistic input the model exhibits a much richer array of behaviours (Fig S1B), including times where both flip-flops are active, and times when neither are active. This is due to the fact that the spiking model uses internal components to approximate the ideal flip-flop algorithm, but this approximation is imperfect. Thus, the spiking model may be able to get into a transient stable state where both flip-flops are on, but one is slightly stronger, causing a long turn in one direction. Alternatively, there may be other low-level differences between the components used for the two sides of the model, leading to subtle asymmetries in the model. With this in mind, further examination of the looping behaviour is warranted. Since moths have been described to perform loops when they lose the odour plume (Baker and Haynes, 1996), we analysed the sensor values during looping. Loops were detected automatically based on their high rotation rates, allowing us to divide the trajectory into looping and non-looping stretches (Fig. 5A). Analysing the difference in value between the two sensors during looping, we find that looping occurs proportionally more often when the two sensors have similar or equal values (Fig. 5C-E), suggesting that looping emerges from the model when the sensor data is non-directional. Furthermore, the average difference in sensor values is consistently lower during a loop than during a turn, and both terminate when the difference between left and right sensor increases (Fig. 5C). During a loop, FF neurons switch activity state less and their output tends to remain stable on one side as compared to a turn (Fig. 5E).

To investigate what a possible advantage of looping could be, we compared the spiking FF network to a control agent in which the sensors were connected directly to the motor and which showed no looping behaviour. We removed the tendency to turn upwind in the absence of sensory input, to simulate a situation in which the agent solely relies on the odour plume, and tested both agents in successively sparser plumes. Under these conditions, the spiking FF agent had a significantly higher success rate than the control brain (Fig. 5F). Surprisingly, at very low plume concentrations, the control brain had a high success rate. We suspected that this may be a side effect of the agent pointing towards the odour source at the start of the simulation, and simply walking straight in the absence of sensor input. Indeed, when we changed the starting angle such that the agent faced 90 degrees to the right of the odour source, the success rate of the control brain was reduced to almost zero while the spiking FF agent remained successful at finding the source of the plume. This indicates that looping is indeed an orientation strategy that, in the absence of all other sensory information, can aid in locating the target.

Having established that the simple flip-flop networks are able to reliably replicate several characteristics of male silkworm moth behaviour, we then proceeded to test whether the models could also take CX output as input signals. For this experiment, we used a computational model of the CX that computes path integration in an anatomically

constrained network (Fig. 6A; (Stone et al., 2017)). In short, path integration is a computation that combines, at each time step, the current heading of the animal (represented in TB1/delta7 neurons, Fig. 6A) with its forward speed (represented in TN/LNO neurons, Fig. 6A). The resulting vector (updated in the CPU4/PFN memory loop, Fig. 6A) points in the direction of the path's origin, and the length of the vector represents the distance of the animal from that origin. Thus, an animal that continuously updates this vector during an outbound path has the possibility to return back to its origin in a straight line. Once the animal decides to return, the desired heading (encoded in the CPU4/PFN vector) is compared to the current heading (TB1/delta7 neurons), and mismatches between the two are transferred to an unspecified motor via CPU1/PFL neurons. As CPU1 neurons project from the CX to the LAL, it is plausible that they interact with flip-flop neurons there. Importantly, steering signals are encoded as an imbalance between the summed activity of all CPU1 neurons in the right hemisphere and those in the left hemisphere. We therefore used the summed CPU1 activity as input to the flip-flop network. Using the same model parameters as for odor-plume experiments, both the rate-based and the spiking model steer an agent back to its origin based on path integrator output (Fig. 6B), albeit with a lower success rate than the ideal path integrator (Fig. 6C). We assessed the accuracy of homing by analysing the orientation of the agent relative to the origin, where 0° indicates that the agent is perfectly tracking along the straight line between the end of the outbound path and its origin (Fig. 6D). Without any connections to the steering model, the ideal path integrator peaks at an orientation of 0° and has a standard deviation of 15.4° . The spiking model also peaks at 0° but has a wider standard deviation of 67.9° . Interestingly, the rate-based model has one peak at 0° , a standard deviation of 50.2° , and two additional peaks at $\pm 90^\circ$.

Due to this odd distribution of orientations, and due to the relatively low success rate for the two models (0.7 and 0.57, respectively), we examined different ways of connecting the path integrator to our steering system. There is no a priori reason for expecting an output value of 0.5 from the path integrator to mean the same thing as a 0.5 from the odour detection system. However, we do not want to postulate complex neural mechanisms between these two systems. Two simple things to adjust are the gain of this connection (which would correspond to increasing the number of synapses, or moving the synapse closer to the spike initiation zone) and the bias current (which would correspond to changing the threshold at which the neuron will fire). While neither of these on their own significantly improved behaviour, we found that adjusting both gain and bias such that a path integrator value of 0.5 is mapped to a 0 input to the steering system and a value of 1.0 stays at 1.0 (and intermediate values are linearly interpolated between these) greatly improved performance while keeping the same qualitative effects (Fig. 6). When re-scaling the path integrator output to a scale between 0 and 1, we find that both models are significantly more successful as well as more accurate in tracking along a straight line back to the origin (Fig. 6C-F). At a sensor difference of 0.5 or above, the rate-based model's percentage of successful runs increases from 0.7 to 0.78, and the spiking model increases from 0.57 to 0.92 (Fig. 6C). Thus, the flip-flop circuit can be used to steer using multiple input sources, with minimal modifications.

Discussion

We modelled a simple flip-flop network based on neurons that have been described in detail in the silkworm moth. The computational model has only two pairs of neurons: the flip-flop neurons (FF), which are bistable neurons, and the PBNs, which provide inhibition between the two flip-flops and thereby synchronise the two hemispheres. Note that in our model, there are no other connections between the two hemispheres, thus this synchronisation is entirely mediated by the PBNs. This network was modelled both as a rate-based model and as a spiking model. Surprisingly, both models were able to replicate the behaviour of real moths in an odor plume reliably, despite their simplicity. We found that additional neurons were not necessary to produce the behaviour presented in this paper (Fig. S2).

We also tested whether this simple steering network could serve as an interface between the CX and downstream motor centres by combining it with the CX path integrator network (Stone et al., 2017). We could show that both the rate-based and the spiking model can take input from the path integrator and use it for steering towards a target, and that the efficiency of steering depends on the scaling of the input into the system.

In the following, we will discuss our findings with respect to the models, their behaviours and the predictions and conclusions we can draw from these experiments.

Rate-based vs. spiking model

When comparing the rate-based and the spiking models, both produced similar overall plume-following behaviours, despite the difference in complexity. This reflects that rate-based models are, to a certain degree, a valid way of modelling neural networks, despite the many simplifications they involve. However, the more subtle behaviours of the system seem to be more realistic in the spiking version, in particular for looping behaviour and for the influence of turning bias. Interestingly, when a turning bias is introduced (Fig. 4B), the spiking model produces the same edge-following behaviour observed in the original experiment (Ando et al., 2013), while the rate-based model exhibits some shift, but with a wider spread. Overall, this indicates that the spiking model produces more realistic behaviour.

A major open question is how flip-flopping can be achieved in a real neuron. In the spiking model presented here, we combine low-level components, which consist of voltage build-up and spikes, to approximate the flip-flop behaviour. While this produces more realistic behaviour than the rate-based approach, more details could be added. In particular, various models of neural bistability currently exist (e.g. (Camperi and Wang, 1998; Gruber et al., 2003)) that might serve as the basis for a more accurate model of flip-flop neurons.

Model behaviour

When evaluating the behaviour of both models, the turn angle, turn angular velocity and turn duration agree well with values reported for male silkworm moths and fall within the mean \pm one standard deviation of moths (Ando et al., 2013). The high positive deviation of the spiking model's turn duration (Figure 3A) can be explained by this model's tendency to loop, leading to more and longer turns. The rate-based model also produced looping

behaviour, but at a much lower rate than the spiking model. Loops occurred when the difference between the two sensors' values was very low and the directional information of the sensory input was therefore ambiguous. We could show that looping conferred a distinct advantage in sparser odor plumes, leading to significantly higher success rates as compared to a non-looping control agent. This suggests that looping is a basic orientation strategy that allows an insect to sample its entire local environment, and may be especially useful if the insect orients relative to an external cue, or towards the source of an intermittent cue such as an odor plume. As an orientation strategy, looping has been described in several insects: For example, fruit moths perform a loop as a search mechanism when they lose the odor plume (Baker and Haynes, 1996); dung beetles perform a circular dance on their dung ball to take a 'snapshot' of skylight cues before rolling their dung ball in a straight line relative to these cues (Baird et al., 2012; el Jundi et al., 2016), and perform another dance after losing their bearing; and desert ants use a similar strategy to learn their visual surroundings during learning walks (Fleischmann et al., 2017). Looping therefore appears to be a robust strategy to sample and learn the local sensory environment, as well as to re-acquire a sensory signal that has been lost. Note that this orientation strategy would not be expected in simple taxis behaviours, in which the animal navigates purposefully towards or away from a sensory cue. In this situation, a very small difference between the two sensor values would be expected to elicit a straight walk, with the aim of keeping the difference as small as possible.

Central complex output can be used for steering

In addition to steering towards the source of an odor plume, we have shown that the models also steer well when getting input from the CX path integrator (PI) network published in Stone et al. (2017) (Stone et al., 2017). The output from this network is in essence a steering signal that represents the difference between the intended heading and the current heading. This signal is asymmetric between the right and left hemisphere, depending on whether the agent needs to correct to the right or to the left.

Our steering models can take this input and steer the agent towards its point of origin, using the same parameters that were used for odor-based steering. This was surprising, considering that the sensory input experienced in an odor plume is quite different from the input provided by the PI. Odor plume input is intermittent and varies at a high temporal frequency, whereas PI output is constant, changes smoothly without sudden jumps, and ideally varies within a relatively small range around the intended heading. Re-scaling the PI output to match the range given by the odor plume resulted in successful control behaviour without necessitating fundamental changes to the model. Biologically, this re-scaling could be achieved by the presence of interneurons and/or neuromodulators in the inputs to the flip-flop system, and facilitate adapting to new sensory environments.

Our results suggest that the flip-flop system can take and integrate input signals from multiple modalities, including the CX, to generate downstream steering commands. One important limitation is that the multimodal input signals need to be directional, that is, there must be an imbalance between the signals in the right and left hemisphere. Here, we tested input signals that are derived from visual cues, but there is no reason to assume that the input should be restricted to olfactory and visual cues alone (Steinbeck et al., 2020). Silkworm

moth flip-flop neurons are known to switch state in response to odor cues as well as light flashes, and bistable neurons of a similar morphology were found to switch state in response to auditory input in the cricket (Zorovic and Hedwig, 2011). Our data therefore supports the idea that this neural network may work as a general purpose steering network at least in the context of targeted orientation behaviours (excluding taxis behaviours).

Purpose of the flip-flop system

One important question is whether the flip-flop neurons convey any actual advantage to the steering circuit. Indeed, the path integrator model without a flip-flop steering component works perfectly well ((Stone et al., 2017); see also the “ideal” model in Fig. 6). Similarly, pheromone plumes could be followed by directly connecting the left and right sensor data to downstream motor neurons, bypassing the flip-flop network altogether (Fig. 5F, S3). However, we have shown that one advantage of using the flip-flop network when following an odor plume is the emergence of looping behaviours, which increase the robustness of the behaviour to disturbances such as intermittent loss of sensory input. Furthermore, the flip-flop models produce paths that cover a wider part of the available space, while still successfully finding the plume source. We thus conclude that one purpose of the flip-flop system is to support reliable goal-directed behaviours while also facilitating exploratory behaviours that cause the insect to vary its position and orientation, rather than following a direct path.

Predictions

Our analysis generates several testable predictions. First and foremost, CPU1 neurons that project from the CX to the LAL are expected to have either direct or indirect excitatory synaptic connections with flip-flop neurons. To our knowledge, only two similar connections have been described so far: In the fruit fly, CPU1 neurons (PFL in *Drosophila* nomenclature) were shown to synapse onto bilateral LAL interneurons (Franconville et al., 2018) (see also (Hulse et al., 2021), as well as the ipsilaterally-descending LAL neuron DNa02 (Rayshubskiy et al., 2020). However, whether CX output neurons also project onto contralaterally-descending neurons in the LAL remains unknown. Finding the interaction sites between CX output neurons and LAL descending neurons, such as the flip-flop neurons modelled here, will be an important step towards understanding how the CX controls behaviour.

Secondly, our findings support that the flip-flop network does not only underlie olfactory steering, but that it can be a multimodal steering network. If this is correct across insects, we would expect flip-flop neurons to switch state in response to any stimulus that elicits targeted locomotion (excluding taxis behaviours). Here, we will discuss three examples of targeted locomotion: straight line orientation, migration, and path integration.

Dung beetles perform short-distance, straight-line orientation when rolling their ball away from the dung pile (Dacke et al., 2003a). To keep their path straight, they rely on skylight cues such as the position of the Sun, the polarization pattern of the sky, and the sky’s spectral gradient (Dacke et al., 2003b, 2004, 2013; el Jundi et al., 2014). These cues are integrated in the CX to generate a current heading (el Jundi et al., 2016), which can be used

to steer the animal along its straight path and adjust for deviations. We therefore expect that dung beetle flip-flop neurons should respond to a change of the skylight cue, such as a sudden rotation of the polarization pattern, with a state change.

When it comes to long-distance migration, the Monarch butterfly and the Bogong moth are well-known insect models for diurnal and nocturnal migration, respectively (Reppert et al., 2016; Warrant et al., 2016). The Monarch butterfly uses a time-compensated Sun compass, as well as the geomagnetic field to migrate from its breeding grounds in North America to overwintering regions in central Mexico (Mouritsen and Frost, 2002; Guerra et al., 2014). The Bogong moth uses the geomagnetic field in combination with visual landmarks to migrate from its breeding grounds in southern Queensland and western New South Wales (Australia) to its overwintering sites in the Australian Alps (Dreyer et al., 2018). The CX path integration network processes Sun compass information and has been proposed to be a possible substrate for computing long-distance migration (Heinze and Reppert, 2011; Honkanen et al., 2019), thus making it likely that the resulting steering commands are passed on to LAL descending neurons. We would expect flip-flop neurons in the Monarch butterfly and the Bogong moth to switch state in response to sudden changes in the skylight cues or landmark configuration that they use to orient. Furthermore, many flying insects use optic flow for flight control, including moths (Fry et al., 2009; Weir et al., 2014; Stöckl and Kelber, 2019; Mauss and Borst, 2020). One might therefore expect flip-flop neurons to also respond to a change in the rotational component of optic flow with a state change.

Finally, path-integrating ants and bees are obvious targets for measuring flip-flop neuron responses, considering that we use the PI network as an input for our steering models. However, since the flip-flop neurons are not driven directly by sensory input that can be controlled in an experimental situation, but rather by a memory state, it is more difficult to test how flip-flop neurons respond during homing. One could however test optic flow cues and compass cues separately, to dissect how the different components of the path integrator drive the flip-flop neurons. Alternatively, it may be possible to perform extracellular tetrode recordings from flip-flop neurons during natural homing on a track ball (Dahmen et al., 2017).

Conclusions and outlook

Of course, motor control mediated by descending neurons is much more complex than the simple model presented here. In the silkworm moth, several other neuron types were described to play a role in pheromone-mediated steering, which can be added to the model for increased complexity and biological relevance. Additionally, complex motor patterns are most likely mediated by not just one cell type, but by a population code across a number of descending neurons (Namiki et al., 2018a). Creating larger and more integrated models of this form is a useful tool for a more complete understanding of these complex interactions.

We believe that the model we have developed and presented here is one small step towards understanding the connection between the heading direction system in the CX and downstream motor centers. Importantly, the approach we have taken to develop this model is flexible and suitable for a wide range of model features. To date, this is the

most complex insect-based model developed using the Nengo neural modelling software. However, Nengo has also been used for a wide variety of mammal-based models, including Spaun, a large-scale model of the human brain (Eliasmith et al., 2012). While modelling insect brains offers different challenges than mammalian brains, we believe our work has shown that this sort of large-scale model is possible, and can lead to more realistic behaviour than some traditional modelling approaches.

Developing a more complete model cannot be done by a single group of researchers. We have made our model freely available at https://github.com/stanleyheinze/insect_steering, and we hope that a community of researchers can, over time, add neuron types and neural systems to increase the complexity of the model and advance our understanding of this general steering system in insects.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We would like to thank the organizers and participants of the Nengo summer school 2016, in particular Ben Morcos and Xuan Choo, for contributing to the initial moth model and for helpful discussions. We are also grateful to Noriyasu Ando for sharing his data with us.

References

- Álvarez-Salvado E, Licata AM, Connor EG, McHugh MK, King BM, Stavropoulos N, Victor JD, Crimaldi JP, Nagel KI. Elementary sensory-motor transformations underlying olfactory navigation in walking fruit-flies. *eLife*. 2018; 7: 1–38.
- Ando N, Emoto S, Kanzaki R. Odour-tracking capability of a silkmoth driving a mobile robot with turning bias and time delay. *Bioinspiration and Biomimetics*. 2013; 8 (1)
- Baird E, Byrne MJ, Smolka J, Warrant EJ, Dacke M. The dung beetle dance: An orientation behaviour? *PLoS ONE*. 2012; 7 (1) 1–6.
- Baker TC, Haynes KF. Pheromone-mediated optomotor anemotaxis and altitude control exhibited by male oriental fruit moths in the field. *Physiological Entomology*. 1996; 21 (1) 20–32.
- Balakrishnan R, Pollack GS. Recognition of courtship song in the field cricket, *Teleogryllus oceanicus*. *Animal Behaviour*. 1996; 51 (2) 353–366.
- Bekolay T, Bergstra J, Hunsberger E, DeWolf T, Stewart TC, Rasmussen D, Choo X, Voelker AR, Eliasmith C. Nengo: a Python tool for building large-scale functional brain models. *Frontiers in Neuroinformatics*. 2014; January. 7: 1–13.
- Camperi M, Wang XJ. A model of visuospatial working memory in prefrontal cortex: Recurrent network and cellular bistability. *Journal of Computational Neuroscience*. 1998; 5 (4) 383–405. [PubMed: 9877021]
- Cande J, Namiki S, Qiu J, Korff W, Card GM, Shaevitz JW, Stern DL, Berman GJ. Optogenetic dissection of descending behavioral control in *Drosophila*. *eLife*. 2018; 7: 1–23.
- Cardé RT, Willis MA. Navigational strategies used by insects to find distant, wind-borne sources of odor. *Journal of Chemical Ecology*. 2008; 34 (7) 854–866. [PubMed: 18581182]
- Dacke M, Baird E, Byrne M, Scholtz CH, Warrant EJ. Dung Beetles Use the Milky Way for Orientation. *Current Biology*. 2013; 23 (4) 298–300. [PubMed: 23352694]
- Dacke M, Byrne MJ, Scholtz CH, Warrant EJ. Lunar orientation in a beetle. *Proceedings of the Royal Society B: Biological Sciences*. 2004; 271 (1537) 361–365.
- Dacke M, Nilsson D-E, Scholtz CH, Byrne M, Warrant EJ. Insect orientation to polarized moonlight. *Nature*. 2003a; 424 (6944) 33.

- Dacke M, Nordström P, Scholtz CH. Twilight orientation to polarised light in the crepuscular dung beetle *Scarabaeus zambesianus*. *The Journal of Experimental Biology*. 2003b; 206 (9) 1535–1543. [PubMed: 12654892]
- Dahmen H, Wahl VL, Pfeffer SE, Mallot HA, Wittlinger M. Naturalistic path integration of *Cataglyphis* desert ants on an air-cushioned lightweight spherical treadmill. *Journal of Experimental Biology*. 2017; 220 (4) 634–644. [PubMed: 28202651]
- Dreyer D, Frost B, Mouritsen H, Gunther A, Green K, Whitehouse M, Johnsen S, Heinze S, Warrant E. The Earth's Magnetic Field and Visual Landmarks Steer Migratory Flight Behavior in the Nocturnal Australian Bogong Moth. *Current Biology*. 2018; 28: 1–7. [PubMed: 29249662]
- el Jundi B, Foster JJ, Khaldy L, Byrne MJ, Dacke M, Baird E. A Snapshot-Based Mechanism for Celestial Orientation. *Current Biology*. 2016; 26 (11) 1456–1462. [PubMed: 27185557]
- el Jundi B, Smolka J, Baird E, Byrne MJ, Dacke M. Diurnal dung beetles use the intensity gradient and the polarization pattern of the sky for orientation. *The Journal of experimental biology*. 2014; 217 (13) 2422–2429. [PubMed: 24737763]
- Eliasmith, C, Anderson, CH. *Neural engineering: Computation, representation, and dynamics in neurobiological systems*. MIT Press; Cambridge, MA: 2003.
- Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T, Tang Y, Rasmussen D. A Large-Scale Model of the Functioning Brain. *Science*. 2012; 338 (6111) 1202–1205. [PubMed: 23197532]
- Farrell JA, Murlis J, Long X, Li W, Cardé RT. Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes. *Environmental Fluid Mechanics*. 2002; 2 (1-2) 143–169.
- Fleischmann PN, Grob R, Wehner R, Rössler W. Species-specific differences in the fine structure of learning walk elements in *Cataglyphis* ants. *Journal of Experimental Biology*. 2017; 220 (13) 2426–2435. [PubMed: 28679795]
- Franconville R, Beron C, Jayaraman V. Building a functional connectome of the drosophila central complex. *eLife*. 2018; 7: 1–25.
- Fry SN, Rohrseitz N, Straw AD, Dickinson MH. Visual control of flight speed in *Drosophila melanogaster*. *The Journal of experimental biology*. 2009; 212: 1120–30. [PubMed: 19329746]
- Green J, Adachi A, Shah KK, Hirokawa JD, Magani PS, Maimon G. A neural circuit architecture for angular integration in *Drosophila*. *Nature*. 2017; 546 (7656) 101–106. [PubMed: 28538731]
- Gruber AJ, Solla SA, Surmeier DJ, Houk JC. Modulation of striatal single units by expected reward: A spiny neuron model displaying dopamine-induced bistability. *Journal of Neurophysiology*. 2003; 90 (2) 1095–1114. [PubMed: 12649314]
- Guerra PA, Gegear RJ, Reppert SM. A magnetic compass aids monarch butterfly migration. *Nature Communications*. 2014; 5
- Heinze S. Unraveling the neural basis of insect navigation. *Current Opinion in Insect Science*. 2017; 24 (Figure 1) 58–67. [PubMed: 29208224]
- Heinze S, Reppert SM. Sun Compass Integration of Skylight Cues in Migratory Monarch Butterflies - supplemental information. *Neuron*. 2011; 69 (2) 345–358. [PubMed: 21262471]
- 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
- Hulse BK, Haberkern H, Franconville R, Turner-Evans DB, Takemura SY, Wolff T, Noorman M, Dreher M, Dan C, Parekh R, Hermundstad AM, et al. A connectome of the drosophila central complex reveals network motifs suitable for flexible navigation and context-dependent action selection. *eLife*. 2021; 10: 1–180.
- Jones CD. On the structure of instantaneous plumes in the atmosphere. *Journal of Hazardous Materials*. 1983; 7 (2) 87–112.
- Kakaria KS, de Bivort BL. Ring Attractor Dynamics Emerge from a Spiking Model of the Entire Protocerebral Bridge. *Frontiers in Behavioral Neuroscience*. 2017; 11
- Kanzaki R, Ikeda A, Shibuya T. Morphological and physiological properties of pheromone-triggered flipflopping descending interneurons of the male silkworm moth, *Bombyx mori*. *Journal of Comparative Physiology A*. 1994; 175: 851.
- Kanzaki R, Mishima T. Pheromone-Triggered 'Flipflopping' Neural Signals Correlate with Activities of Neck Motor Neurons of a Male Moth, *Bombyx mori*. *Zoological Science*. 1996; 13 (1) 79–87.

- Kanzaki R, Nagasawa S, Shimoyama I. Neural Basis of Odor-source Searching Behavior in Insect Brain Systems Evaluated with a Mobile Robot. *Chemical Senses*. 2005; 30 (suppl 1) 285–286.
- Kanzaki R, Sugi N, Shibuya T. Self-generated Zigzag Turning of *Bombyx mori* Males During Pheromone-mediated Upwind Walking. *Zoological Science*. 1992; 9 (3) 515–527.
- Kymre JH, Liu XL, Ian E, Berge CN, Wang GR, Berg BG, Zhao XC, Chu X. Distinct protocerebral neuropils associated with attractive and aversive female-produced odorants in the male moth brain. *eLife*. 2021; 10: 1–27.
- Levakova M, Kostal L, Monsempès C, Jacob V, Lucas P. Moth olfactory receptor neurons adjust their encoding efficiency to temporal statistics of pheromone fluctuations. *PLoS Computational Biology*. 2018; 14 (11) 1–17.
- Loudon C, Koehl Ma. Sniffing by a silkworm moth: wing fanning enhances air penetration through and pheromone interception by antennae. *The Journal of experimental biology*. 2000; 203 (Pt 19) 2977–2990. [PubMed: 10976034]
- Mafra-Neto A, Cardé RT. Fine-scale structure of pheromone plumes modulates upwind orientation of flying moths. *Letters to Nature*. 1994; May. 369: 142–144.
- Marsh D, Kennedy JS, Ludlow AR. An analysis of anemotactic zigzagging flight in male moths stimulated by pheromone. *Physiological Entomology*. 1978; 3: 221–240.
- Mauss AS, Borst A. Optic flow-based course control in insects. *Current Opinion in Neurobiology*. 2020; 60: 21–27. [PubMed: 31810007]
- Mishima T, Kanzaki R. Coordination of flipflopping neural signals and head turning during pheromone-mediated walking in a male silkworm moth *Bombyx mori*. *J Comp Physiol A*. 1998; 183: 273–282.
- Mishima T, Kanzaki R. Physiological and morphological characterization of olfactory descending interneurons of the male silkworm moth, *Bombyx mori*. *Journal of Comparative Physiology A*. 1999; 184 (2) 143–160.
- Mouritsen H, Frost BJ. Virtual migration in tethered flying monarch butterflies reveals their orientation mechanisms. *PNAS*. 2002; 99 (15) 10162–10166. [PubMed: 12107283]
- Namiki S, Dickinson MH, Wong AM, Korff W, Card GM. The functional organization of descending sensory-motor pathways in *drosophila*. *eLife*. 2018a; 7: 1–50.
- Namiki S, Iwabuchi S, Pansopha Kono P, Kanzaki R. Information flow through neural circuits for pheromone orientation. *Nature communications*. 2014; 5 5919
- Namiki S, Kanzaki R. The neurobiological basis of orientation in insects: insights from the silkmoth mating dance. *Current opinion in insect science*. 2016. February. 16–26.
- Namiki S, Wada S, Kanzaki R. Descending neurons from the lateral accessory lobe and posterior slope in the brain of the silkmoth *Bombyx mori*. *Scientific Reports*. 2018b; 8 (1) 1–26. [PubMed: 29311619]
- Obara Y. *Bombyx mori* Mating Dance: an Essential in Locating the Female. *Applied Entomology and Zoology*. 1979; 14 (1) 130–132.
- Olberg RM. Pheromone-triggered flip-flopping interneurons in the ventral nerve cord of the silkworm moth, *Bombyx mori*. *Journal of Comparative Physiology*. 1983; 152 (3) 297–307.
- Rayshubskiy A, Holtz S, D’Alessandro I, Li A, Vanderbeck Q, Haber I, Gibb P, Wilson R. Neural circuit mechanisms for steering control in walking *Drosophila*. *bioRxiv*. 2020. 2020.04.04.024703
- Reppert SM, Guerra PA, Merlin C. Neurobiology of Monarch Butterfly Migration. *Annual Review of Entomology*. 2016; 61 (1) 25–42.
- Schnell B, Ros IG, Dickinson MH. A Descending Neuron Correlated with the Rapid Steering Maneuvers of Flying *Drosophila*. *Current Biology*. 2017; 27 (8) 1200–1205. [PubMed: 28392112]
- Simmons LW. The calling song of the field cricket, *Gryllus bimaculatus* (De Geer): constraints on transmission and its role in intermale competition and female choice. *Animal Behaviour*. 1988; 36 (2) 380–394.
- Srinivasan MV. Where paths meet and cross: navigation by path integration in the desert ant and the honeybee. *Journal of Comparative Physiology A*. 2015; 201 (6) 533–546.
- Steinbeck F, Adden A, Graham P. Connecting brain to behaviour: A role for general purpose steering circuits in insect orientation? *Journal of Experimental Biology*. 2020; 223 (4)

- Stöckl AL, Kelber A. Fuelling on the wing: sensory ecology of hawkmoth foraging. *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*. 2019; 205 (3) 399–413. [PubMed: 30880349]
- 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. e11 [PubMed: 28988858]
- Vickers NJ, Christensen TA, Baker TC, Hildebrand JG. Odour-plume dynamics influence the brain's olfactory code. *Nature*. 2001; March. 410
- Warrant E, Frost B, Green K, Mouritsen H, Dreyer D, Adden A, Brauburger K, Heinze S. The australian bogong moth *Agrotis infusa*: A long-distance nocturnal navigator. *Frontiers in Behavioral Neuroscience*. 2016; April. 10
- Weir PT, Schnell B, Dickinson MH. Central complex neurons exhibit behaviorally gated responses to visual motion in *Drosophila*. *Journal of Neurophysiology*. 2014; 111 (1) 62–71. [PubMed: 24108792]
- Zorovic M, Hedwig B. Processing of species-specific auditory patterns in the cricket brain by ascending, local, and descending neurons during standing and walking. *Journal of Neurophysiology*. 2011; 105 (5) 2181–2194. [PubMed: 21346206]

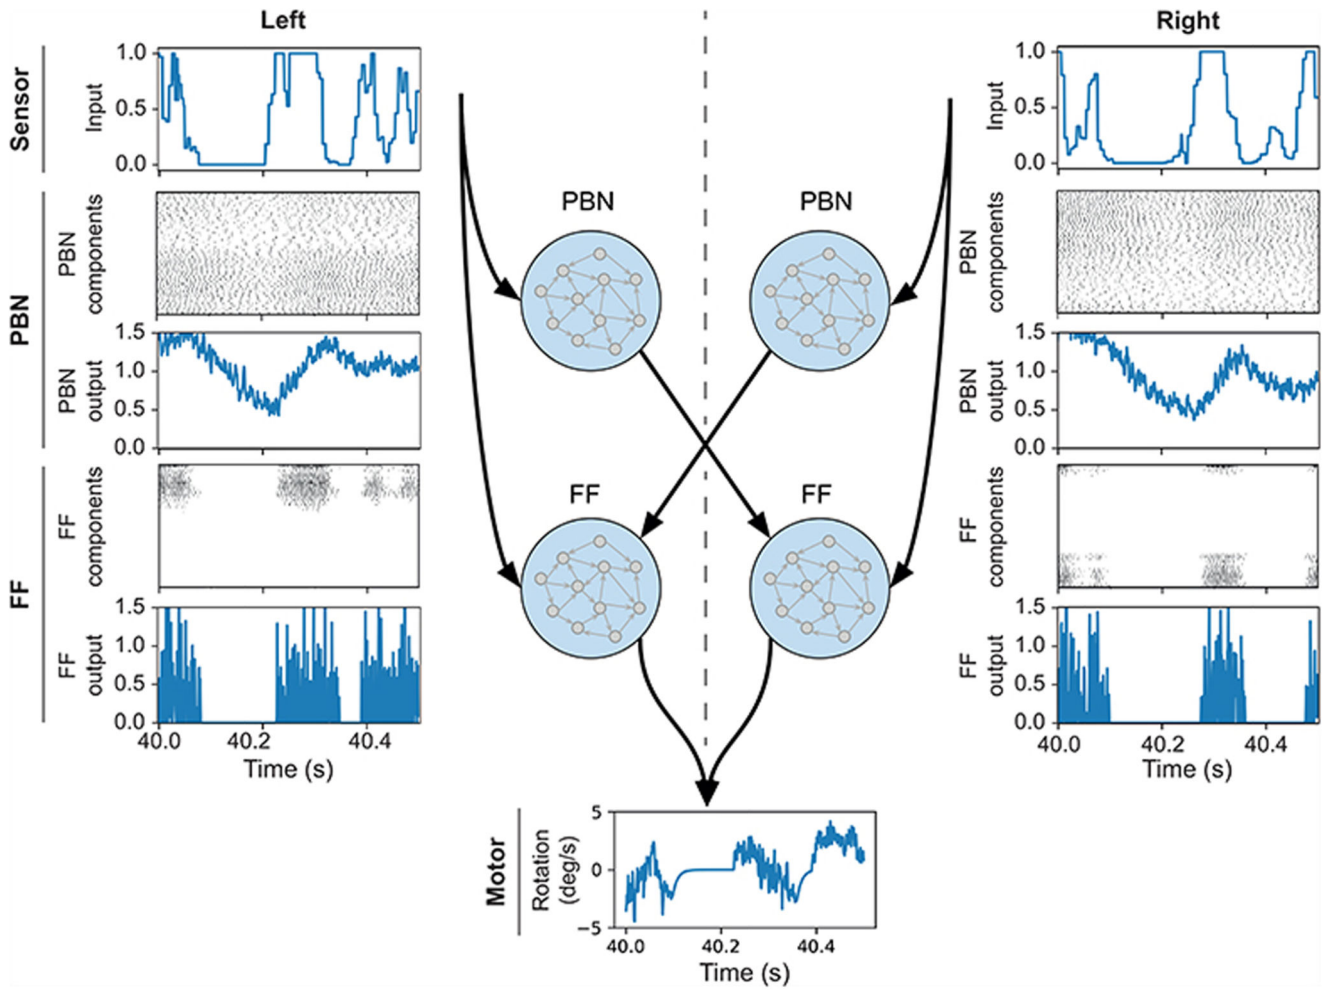


Figure 1.

The flip-flop network. Protocerebral bilateral neurons (PBN) and flip-flop neurons (FF) both get input directly from the ipsilateral sensor. PBNs inhibit the contralateral FF. FF neurons activate the motor. For the spiking model, each neuron is comprised of 100 spiking components, recurrently connected so as to achieve an approximation of the rate-mode behaviour. The output from each neuron is a weighted sum of the output of its components. The final output controls the turning rate of the moth, and is computed as the left FF output minus the right FF output times a fixed scaling factor. Activity shown here is for the spiking network. See Fig. S1 for a comparison of activity propagation through the rate-based vs. spiking model given artificial input. See supplementary video for an animated version of this figure.

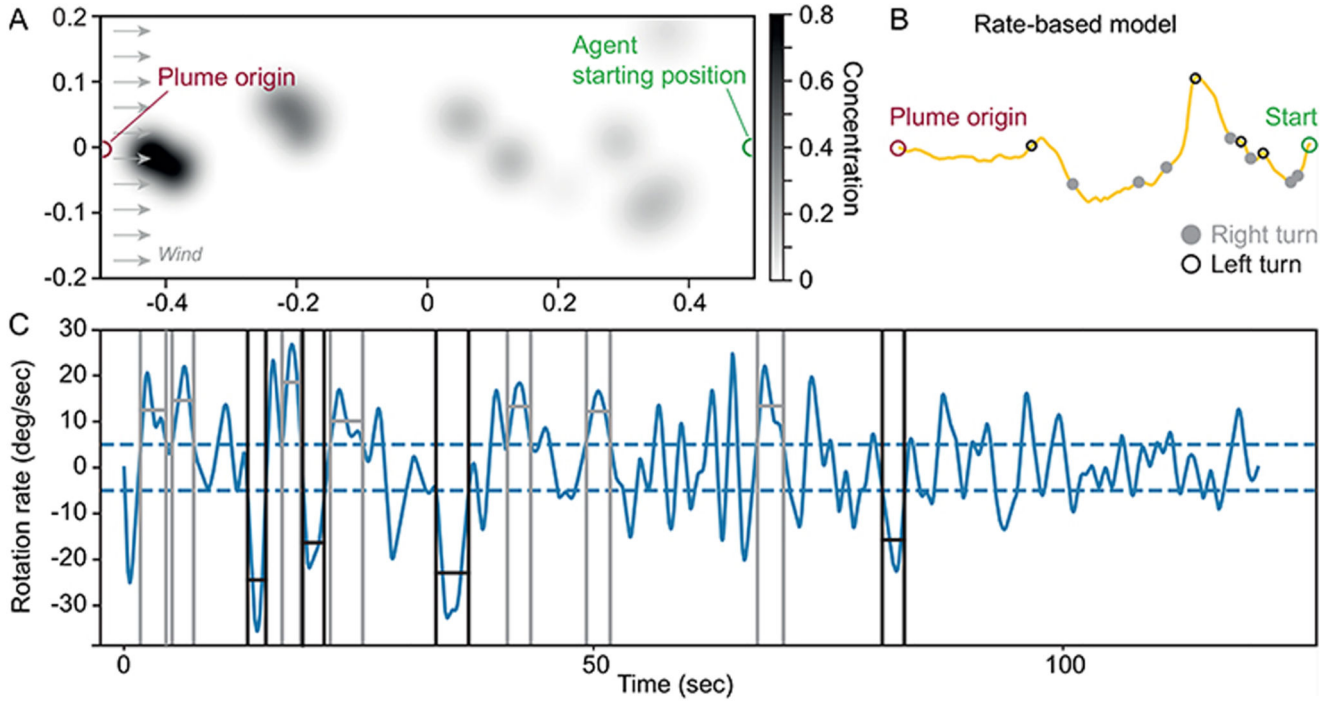


Figure 2. Overview over the odor plume experiment.

A: The pheromone plume is released from the source and dispersed by wind. Shown is one frame of the plume, with the grey value reflecting the pheromone concentration on a scale from 0 to 1, where 0 translates into no input to the agent's sensor and 1 into maximal sensory input. In the absence of input, the agent turns upwind. B: Typical path of the rate-based model, with right turns indicated by filled grey circles and left turns indicated by empty black circles. C: The rotation rate defines whether a rotation is classified as a turn. Detected turns are marked (right turn = grey, left turn = black). Definition as in (Ando et al., 2013).

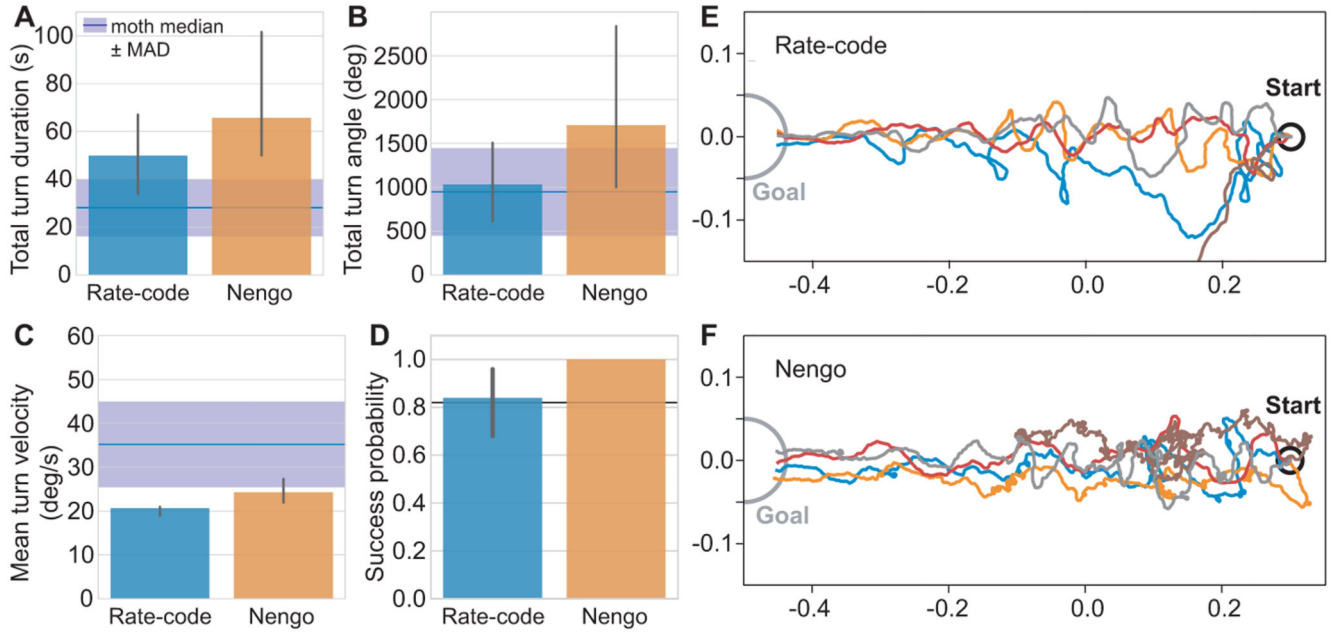


Figure 3. Behaviour of both models in a simulated odor plume.

A-D: Comparison of the rate-based and spiking model to data from silkworm moths. Bar plots represent the median, with the error bar giving the bootstrapped 95 % confidence interval. The moth mean \pm standard deviation (black line and grey area) and the moth median \pm median absolute deviation (dark blue line and blue area) are given for comparison. Moth data reproduced with permission from Ando et al. (2013). E-F: Example trajectories for the rate-based (E) and spiking (F) models, with different colours denoting individual trials. For additional paths, see Fig. S2

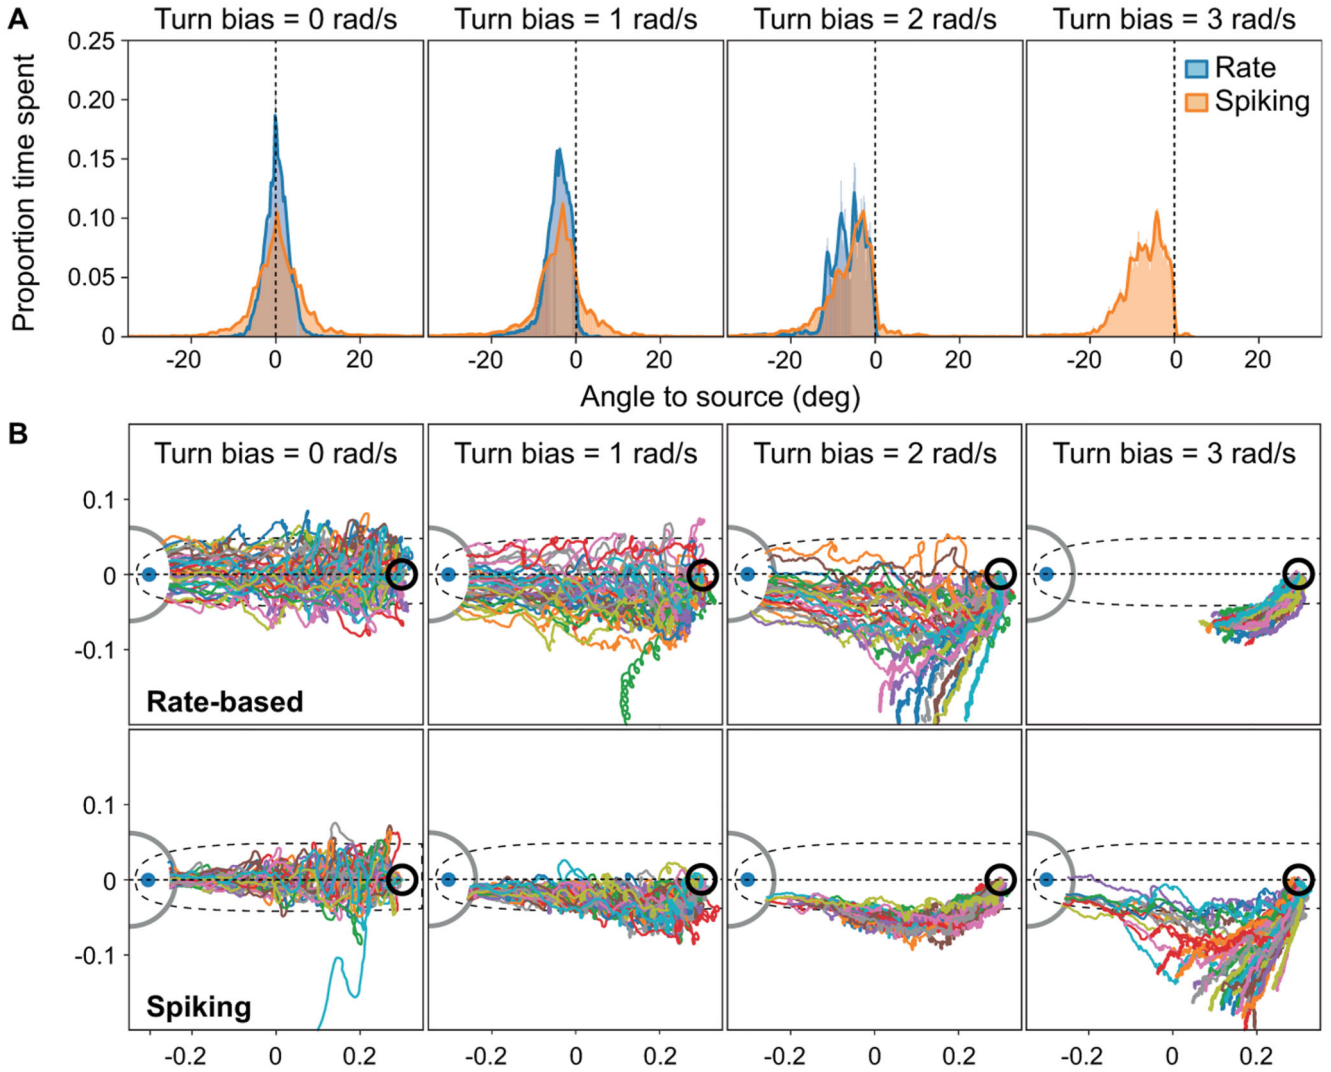


Figure 4. Turn bias.

A: Proportion of time the agent spent at a specific angle relative to the source of the odor plume. Only successful trials were included. When a turn bias was added to the left motor, the agents shifted away from the center of the plume towards the plume’s edge. At a turn bias of 3, the rate-based model failed consistently. B: Paths of the rate-based and spiking models with increasing turn bias. $N = 40$ per condition. Dotted line = center of the odor plume; dashed outline = area of odour concentration of at least 10 percent of the maximum detectable level, averaged across 3000 time steps; blue dot = plume origin; grey circle = area around plume origin that needs to be reached in order to count as success.

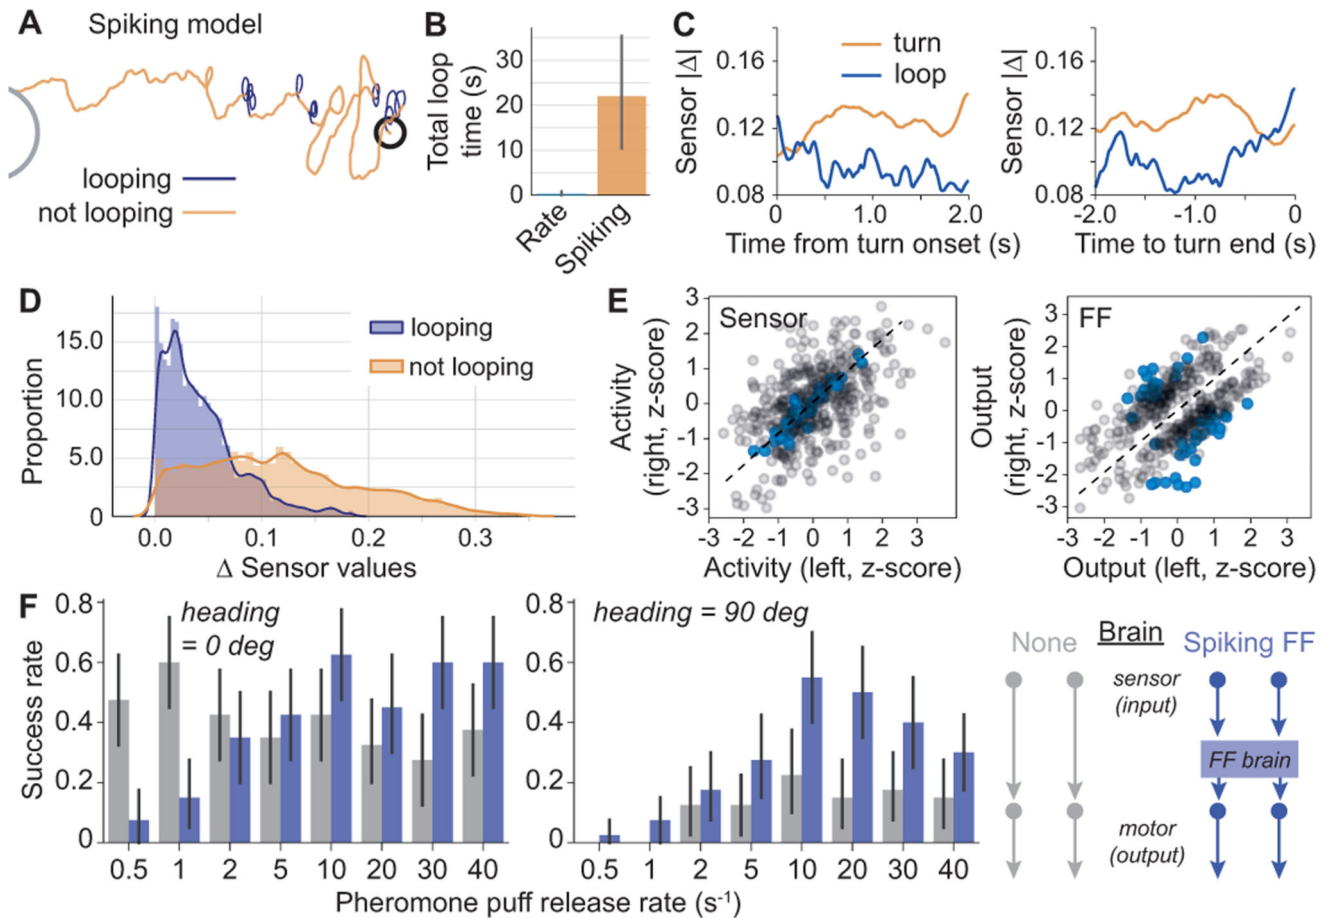


Figure 5. Looping behaviour of the spiking model.

A: Example track of a spiking model agent orienting in an odor plume, with loops highlighted in purple. B: During tracks lasting on average 253 seconds, the spiking model spends over 20 s looping, while loops almost never occur in the rate-based model. C, D: The difference in sensor input is consistently lower during a loop than during a turn. E: Z-score of the sensor activity and FF neuron output during loops (blue) vs. turns (grey). Sensor activity is substantially more similar during a loop than during a turn, while FF neuron output tends to remain on one side. F: Comparison of the spiking FF brain to a control brain in which the sensors are directly connected to the motor, at starting angles of 0 degrees and 90 degrees relative to the origin of the plume. The spiking FF model is significantly more robust to increasingly sparse plumes, with a high success rate down to a puff release rate of 10 puffs per second, independent of starting angle.

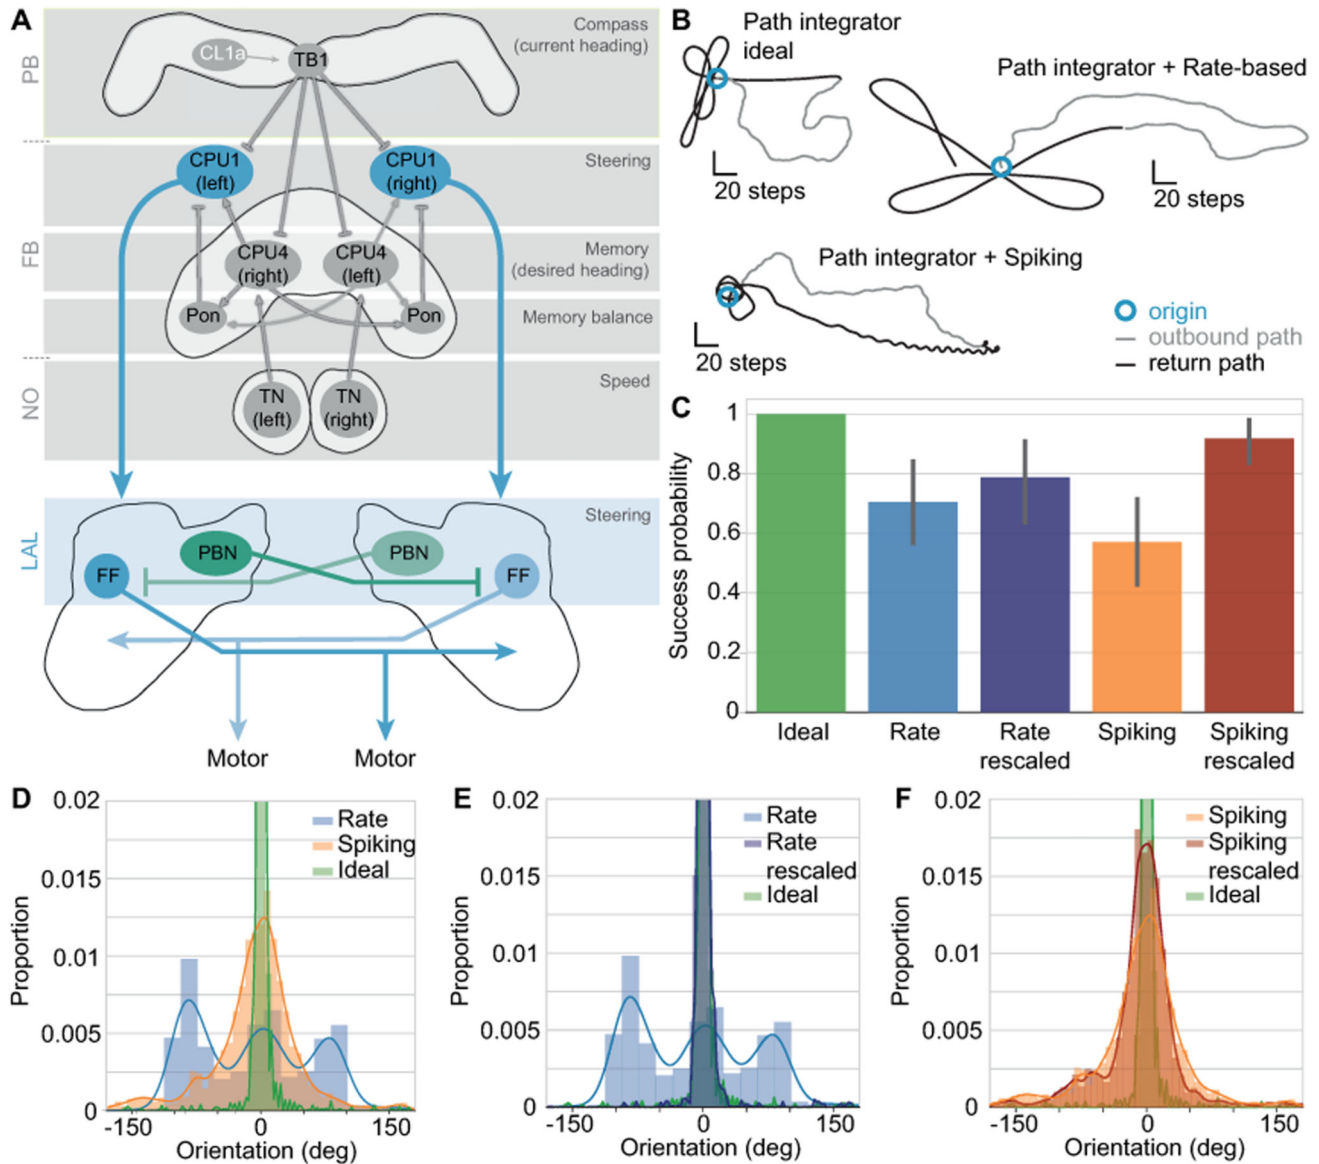


Figure 6. Integrating the flip-flop models with the central complex path integrator.

A: Schematic of the integrated network. The path integrator receives compass input via CL1a neurons, and speed input via TN neurons. TB1 neurons compute the current heading direction. CPU4 and pontine (Pon) neurons form a memory loop which integrates the current heading direction with the current speed, resulting in a vector that points to the path's origin. When the agent wants to return home, CPU1 neurons compare the current heading (represented in TB1 neurons) to the desired heading given by the CPU4 home vector. CPU1 neuron output is then fed directly into the flip-flop networks, with the same input driving both the inhibitory protocerebral bilateral neurons (PBN) and the flip-flop neurons (FF). B: Example paths of the path integrator without a steering network (ideal), with the rate-based model, and with the spiking model. C: Success rate of the integrated models compared to the ideal path integrator. The success probability of both models increases when the CPU1

output is rescaled to a range of 0 to 1. Error bars represent the bootstrapped 95 % confidence interval around the median. D: Proportion of time the models spend at a certain angle relative to the path's origin. E: With rescaled input, the distribution of the rate-based model becomes almost identical to that of the ideal path integrator. F: With rescaled input, the distribution of the spiking model has a smaller spread around 0° and becomes more similar to the ideal path integrator. $N = 50$ for all path integrator experiments. Only successful trials were considered in D-F.

Table 1

Simulation parameters.

Parameter	Value
Number of LIF components to implement one neuron	100
Random noise added to neurons	$N(0,0.02)$
Wind speed	2m/s
Plume puff release rate	50 Hz
Plume puff initial radius	0.1 m
Antenna size	7.5 mm
Antenna angle	45°
Antenna maximum sensitivity	100
Antenna maximum output value	3.0
Moth maximum rotation speed	8.0 rad/s
Moth acceleration	$0.2m/s^2$
Moth drag	0.5 m/s
Moth turn-into-wind rate	0.1 rad/s