

Neural circuit mechanisms for steering control in walking *Drosophila*

Aleksandr Rayshubskiy¹, Stephen L. Holtz¹, Isabel D'Alessandro¹, Anna A. Li¹, Quinn X. Vanderbeck¹, Isabel S. Haber¹, Peter W. Gibb¹, and Rachel I. Wilson^{1,*}

¹Department of Neurobiology, Harvard Medical School, Boston, MA, USA

*e-mail: rachel_wilson@hms.harvard.edu

Abstract: Navigation can be directed toward distant targets represented within the brain's spatial maps; alternatively, navigation can be directed toward objects in the local environment. Here we identify neurons in the *Drosophila* brain that integrate these two types of navigation drives. These neurons send axonal projections to the ventral nerve cord, and their activity predicts and influences steering during walking. Meanwhile, their dendrites integrate steering signals from the compass in the brain's spatial memory center, as well as stimulus-directed steering signals from multimodal sensory pathways that bypass the compass. Using a computational model, we show how the specific connectivity of this network can generate steering behavior directed toward internal (remembered) goals, and we show how environmental cues can dynamically alter the balance of stimulus- and memory-directed steering. Our results suggest a framework where motor dynamics emerge from the integration of parallel feedback loops that drive steering toward internal versus external goals.

Introduction

Navigation is sometimes directed toward remembered targets or distant locations (memory-directed navigation). At other times, navigation is directed toward objects observed in the immediate environment (stimulus-directed navigation); here, the value of the attractive object may have been learned, but the navigational route to the object does not need to be learned, and this makes stimulus-directed navigation fundamentally different from memory-directed navigation, where the target is not currently observable. Navigating in a complex environment may involve switching between these different navigation strategies¹. In mammals, memory-directed navigation depends largely on hippocampal pathways, whereas stimulus-directed navigation depends more on striatal pathways²⁻⁴. Ultimately, both pathways must somehow communicate with motor systems, including the descending neurons in the mammalian brainstem which send steering commands to the spinal cord⁵.

Drosophila and other arthropods can also switch between memory- and stimulus-directed navigation⁶⁻⁸. In arthropods, the locus of spatial memory is the central complex, which drives steering toward a remembered spatial goal. So-called "compass neurons" in the central complex encode the organism's current heading relative to environmental cues⁹⁻¹¹, analogous to head direction cells in mammals. When compass neurons are genetically silenced, *Drosophila* can no longer steer toward a remembered spatial goal; remarkably, however, they can still steer directly toward an observable object^{12,13}. These results suggest that memory-directed navigation and stimulus-directed navigation are mediated by distinct pathways in *Drosophila*, just as in mammals. It has been proposed that these pathways are merged in a region called the lateral accessory lobe or LAL¹⁴ (**Fig. 1a**). The LAL receives input from visual, olfactory, and auditory brain regions, in addition to input from the central complex¹⁴. Meanwhile, the LAL sends descending projections to the ventral nerve cord¹⁴⁻¹⁹, the insect analog of the spinal cord. LAL neuron activity is known to correlate with steering and influence steering in crickets, silkmoths, and cockroaches^{16-18,20}. However, the pathway from the compass neurons to LAL descending neurons has not been elucidated in any species. Moreover, it is not known how memory-directed navigation and stimulus-directed navigation are integrated.

Here, we identify descending neurons that carry information from the *Drosophila* brain to the limbs to control steering during walking. We show that these neurons are two synapses downstream from compass neurons, and they participate in memory-directed steering by the compass system. We show that same descending neurons also receive multimodal sensory input that bypasses the compass system, and they participate in stimulus-directed steering. We combine electrophysiology data and connectomics data to generate a computational model

of the feedback loops that integrate external sensory stimuli and internal maps of space to control steering during navigation. This model can explain how the brain transitions seamlessly between the pursuit of remembered goals and observable goals.

Results

Our first goal was to identify descending neurons (DNs) in the *Drosophila* brain that participate in steering during walking. We focused on DNs that receive synaptic input in the LAL and project to all the leg-zones of the ventral cord. A recent survey¹⁵ identified four *Drosophila* DN types that fulfill these criteria. We focused on two of these (DNa01 and DNa02; **Fig. 1a**) because these neurons are morphologically similar to cricket neurons that are active during steering^{18,21,22}. In *Drosophila*, DNa01 and DNa02 can be targeted using selective split-Gal4 lines that drive expression in a single cell per brain hemisphere¹⁵, and a study using one of these lines has already shown that DNa01 is active in the ventral cord during steering²³.

Descending neurons that predict different features of steering

Using these split-Gal4 lines to target our electrodes, we performed whole cell recordings from either DNa01 or DNa02 in flies walking on a spherical treadmill (**Fig. 1b**). For each cell, we then quantified the relationship between neural activity and behavioral dynamics. Specifically, we computed the linear filters that describe the relationship between firing rate and behavior. These filters tell us how behavior changes after a brief increase in firing rate – i.e., a firing rate “impulse”. To compute these filters, we cross-correlated firing rate with the fly’s rotational velocity, sideways velocity, or translational velocity. We then normalized each filter by the autocorrelation in the cell’s firing rate (**Extended Data Fig. 1**).

These filters showed that DNa01 and DNa02 firing rate increases were typically followed by relatively large changes in rotational and sideways velocity (**Fig. 1c,d**). By comparison, DNa01 and DNa02 firing rate increases were not followed by similarly large changes in translational velocity (**Fig. 1e**). It should be noted that rotational and sideways velocity are themselves strongly correlated, meaning that steering often involves rotation as well as side-stepping, whereas steering is only weakly correlated with translational velocity^{22,43} (**Extended Data Fig. 1**)^{24,25 20,21,20,21}. Our results imply that DNa01 and DNa02 are mainly related to steering, rather than translational movement.

We also found several interesting functional differences between DNa01 and DNa02. First, steering filters (i.e., rotational and sideways velocity filters) were larger for DNa02 (**Fig. 1c,d**). This means that that an impulse change in firing rate predicts a larger change in steering. Moreover, steering filters for DNa02 were more accurately predictive of behavior – i.e., when we convolved each filter with the cell’s spike train, the DNa02 filters produced a better prediction of the fly’s subsequent steering, as compared to the DNa01 filters (**Fig. 1f**). Finally, DNa02 steering filters were biphasic, whereas DNa01 steering filters were monophasic (**Fig. 1c,d**). A biphasic filter converts a sustained input into a transient output, while a monophasic filter converts a sustained input into a sustained output. Thus, DNa02 predicts more transient changes in steering.

To more directly compare the firing rate dynamics of the two cell types, we made simultaneous recordings from DNa01 and DNa02 on the same side of the brain (**Fig. 2a**) using a combination of transgenes that targets both cells (**Extended Data Fig. 2**). These dual recordings showed that spike bursts were more transient in DNa02 than in DNa01 (**Fig. 6a,b, Extended Data Fig. 1**). Recall that, even we normalize for the firing rate dynamics of each cell type, the behavioral impulse response is more transient for DNa02 than for DNa01 (**Fig. 1c,d**). For both these reasons – a distinction in firing rate dynamics, and a distinction in the behavioral impulse response – it seems that DNa02 is associated with relatively transient steering, whereas DNa01 is associated with more sustained steering.

To explore this idea, we computed rotational velocity filters for each cell in these dual recordings. We convolved each cell's firing rate with its filter to obtain the rotational velocity prediction for each cell. We found that DNa02 generally predicted large, transient fluctuations in rotational velocity, whereas DNa01 generally predicted smaller, more sustained changes in rotational velocity in the same experiment (**Fig. 2c**). This finding further supports the idea that increases in DNa02 firing rates are associated with transient, high-gain steering, whereas increases in DNa01 firing rates are associated with sustained, low-gain steering.

Next, we asked how rotational velocity might be predicted by the simultaneous firing rates of both cell types. Specifically, we compared neural activity with behavior 150 ms later, to account for the average time lag between neural activity and the subsequent change in the fly's rotational velocity (**Fig. 1c**). We observed that rotational velocity was a relatively steep function of DNa02 activity, but a comparatively shallower function of DNa01 activity (**Fig. 2d-f**). This analysis supports the conclusion that there is a high-gain relationship between DNa02 activity and steering, whereas there is a lower-gain relationship between DNa01 activity and steering.

Translational velocity was not strongly correlated with neural activity in these dual recordings (**Extended Data Fig. 3**), as we might expect. Interestingly, however, when DNa01 was more hyperpolarized than DNa02, the fly was generally moving backward (**Extended Data Fig. 4**). During backward walking, right-left differences in DNa02 activity were linearly related to steering, just as during forward walking (**Extended Data Fig. 4**).

To determine how steering depends on the combined activity of DNa02 neurons on both sides of the brain, we made dual recordings from DNa02 neurons on the right and left (**Fig. 3a**). We found that rotational velocity was consistently related to the difference in right-left firing rates (**Fig. 3b**). This relationship was essentially linear through its entire dynamic range, and was consistent across paired recordings (**Fig. 3c**). We obtained similar results in dual recordings from DNa01 neurons on the right and left (**Extended Data Fig. 5**).

We then extended our linear filter analysis to right-left DNa02 pairs. We computed each cell's linear filter, and we then convolved each filter with the cell's spike train to predict the fly's rotational velocity. By adding the predictions from simultaneously recorded right-left pair, we found that we could accurately predict turns in both directions (**Fig. 3d,e**).

To summarize, we find that different DNs predict distinct features of steering behavior. DNa02 activity precedes transient, high-gain steering, whereas DNa01 precedes sustained, low-gain steering. DNa02 is overall more strongly predictive of steering, and in this cell type, steering velocity is linearly related to the right-left difference in firing rate.

Compass-directed steering

Having identified steering-related DNs, we then investigated the brain circuits that provide input to these DNs. We were particularly interested in whether these particular steering DNs are functionally downstream from the compass neurons of the central complex (E-PG neurons)⁹. Steering is correlated with the activity of some central complex neurons^{26,27}, but the central complex contains no DNs¹⁵, and no previous studies have demonstrated a functional connection from the central complex to DNs.

To investigate whether DNa02 is functionally downstream from the compass, we made whole-cell recordings from DNa02 while monitoring and manipulating compass neurons. To track the dynamics of the compass in real time, we imaged GCaMP6f in the entire compass neuron ensemble. Meanwhile, we microstimulated central complex neurons called P-EN1 neurons¹³ by expressing the ionotropic ATP receptor P2X₂ in these neurons under Gal4/UAS control²⁸ and puffing ATP onto their dendrites (**Fig. 4a**). P-EN1 neurons function to relay information about the fly's angular velocity to compass neurons, and so microstimulating P-EN1 neurons causes the compass to register a fictive behavioral turn.

As expected from previous work, we found that there is a single bump of activity in the compass neuron ensemble which rotates as the fly steers in darkness^{9,10}. The position of the bump constitutes a working memory of the fly's heading direction. Again consistent with previous work, we found that briefly microstimulating P-EN1 neurons causes this bump to jump to a new location¹³. As a negative control, we confirmed that bump jumps rarely follow ATP injection when the Gal4 transgene is omitted (**Fig. 4b**).

The bump jump should create a mismatch between the fly's internal steering goal and its remembered compass heading. Accordingly, the fly typically executes a compensatory behavioral turn shortly after the bump jump¹³. This turn brings the compass bump back to its initial location (**Fig. 4c,d**). Note that the behavioral turn can be in either direction: a leftward turn causes the bump to return clockwise, while a rightward turn causes the bump to return counter-clockwise (as viewed from the posterior side of the brain; **Fig. 4a,d**). On average, we found that the bump returned to its initial location in 40% of trials.

Throughout each experiment, we performed a whole-cell recording from DNa02. Notably, we found that DNa02 cells often fired a burst of spikes just before the fly performed its compensatory behavioral turn to bring the compass bump back to its initial location (**Fig. 4e**). To quantify this effect, we identified the time point in each trial where the bump's return speed was maximal. We used this time point to align the data across trials to account for the fact that the bump's return had different latencies on different trials (**Fig. 4c,d**). Because we always recorded from DNa02 cells on the left, we focused on the trials where we expect a leftward behavioral turn – i.e., trials where the bump returned via a clockwise path. In these trials, we found that the bump jump was typically followed by increased DNa02 firing on the left, and then a leftward steering maneuver, and then a clockwise return of the bump (**Fig. 4f**). Moreover, we found that DNa02 could predict much of the trial-to-trial variability in the magnitude of compensatory behavioral turns (**Fig. 4g**).

Trials where the bump returned in the opposite direction (counter-clockwise) provide a negative control. In these trials, the fly's behavioral turn was typically rightward rather than leftward (**Fig. 4d,h**). Accordingly, we found no DNa02 firing rate increase on the left side (**Fig. 4h**). This shows that DNa02 is not activated bilaterally and thus nonspecifically. Rather, DNa02 is recruited specifically on the side that predicts the direction of the behavioral turn.

In summary, DNa02 activity can predict the magnitude and direction of compass-directed steering maneuvers. Thus, DNa02 is likely to be functionally downstream from the central complex. This would mean that DNa02 is a component of the pathway that transforms spatial maps in working memory into steering maneuvers.

Stimulus-directed steering

Next, we asked whether DNa02 is functionally engaged in stimulus-directed steering, in addition to compass-directed steering. Stimulus-directed steering does not require compass neurons^{12,13}. Thus, stimulus-directed steering and compass-directed steering likely involve separate, parallel pathways. These pathways might merge in the brain, or alternatively they might not merge until a later stage of motor control, in the ventral nerve cord.

To generate an attractive sensory stimulus, we expressed the channelrhodopsin variant CsChrimson in olfactory receptor neurons under the control of an *Orco-LexA* transgene. We then stimulated the right or left antennae independently using two thin optical fibers. This fictive odor stimulus produced steering toward the stimulated antenna (**Fig. 5a**). This finding is consistent with previous studies showing that flies will turn toward a lateralized odor, irrespective of the composition of the odor or which olfactory glomeruli are stimulated^{29,30}. Notably, these lateralized fictive odor stimuli also produced asymmetric responses in DNa02, with higher firing rates on the ipsilateral side (**Fig. 5a**). We confirmed there was no neural or behavioral response when the *LexA* transgene was omitted (**Extended Data Fig. 6**).

To generate an aversive stimulus, we expressed CsChrimson in the antennal thermoreceptor neurons that are excited by heat, under the control *Gr28b.d-LexA*. As before, we stimulated the right or left antenna independently. These fictive heat stimuli drove flies to steer away from the stimulated antenna (**Fig. 5b**). These stimuli again produced responses in DNa02, but now with higher firing rates on the side contralateral to the stimulus (**Fig. 5b**). Thus, DNa02 activity was higher on the side ipsilateral to the attractive stimulus, but contralateral to the aversive stimulus. In other words, DNa02 encoded the laterality of the stimulus-evoked action, not the laterality of the stimulus itself.

On average, DNa02 was activated about 150 ms before the stimulus-evoked turn (**Fig. 5c**). Moreover, DNa02 could predict much of the trial-to-trial variability in the magnitude of stimulus-evoked turns (**Fig. 5d**). Together, these findings suggest that DNa02 is the part of the neural circuitry that detects a lateralized sensory stimulus and orients the body toward that stimulus. We obtained similar results for DNa01 neurons (**Extended Data Fig. 7**), supporting the idea that DNa01 and DNa02 often function together. We also confirmed that there is a similar relationship between DNa02 firing rate and behavior for spontaneous steering, compass-directed steering, and stimulus-directed steering (**Extended Data Fig. 8**), suggesting that DNa02 plays a similar role in all these behaviors.

Parallel steering circuits that converge onto descending neurons

Our findings raise the question of whether DNa02 neurons are a site of integration for compass signals and sensory signals. Alternatively, the integration of these pathways might occur upstream from DNa02. To investigate this issue, we set out to identify DNa02 in an electron microscopy (EM) dataset comprising the entire adult female *Drosophila* brain³¹. After we traced all the primary neurites in the relevant tract, we found one cell that was an excellent match to DNa02 (**Fig. 6a,b**). Next, we reconstructed the backbones of several hundred cells directly presynaptic to DNa02. We then selected 28 presynaptic cells for full reconstruction (**Fig. 6c**), with the goal of sampling inputs from many different brain regions. For each presynaptic cell type, we reconstructed multiple individual exemplars (**Supplementary Table 1**). Finally, as our manuscript was in preparation, the Janelia FlyEM Project released a machine-annotated EM connectome of the dorsal-right portion of another adult female *Drosophila* brain (the ‘hemibrain’ dataset)³². We identified FlyEM cell 1140245595 as an excellent match to DNa02, allowing us to compare the inputs to DNa02 in two independent adult female brain samples.

Central complex. We found that DNa02 receives a major, direct input from one central complex cell type, which is called PFL3 in the FlyEM annotation (**Fig. 6d**). PFL3 neurons receive synaptic input in the protocerebral bridge (including direct input from compass neurons). PFL3 neurons receive additional synaptic input in the fan-shaped body. Meanwhile, PFL3 neuron axons project to the LAL, where they form many synapses onto DNa02 (**Supplementary Table 1**). These results demonstrate that DNa02 is only two synapses downstream from compass neurons. The anatomy of PFL3 neurons suggests that they encode the relationship between the fly’s heading direction (as represented by compass neurons) and the fly’s internal goal (which might be represented by neurons of the fan-shaped body), and they influence steering via DNa02. The FlyEM annotation identifies PFL3 as a major input to DNa02 – specifically, the second-ranked among all DNa02 presynaptic cell types, in terms of synapse numbers.

Visual system. We reconstructed two cell types that project to DNa02 from the medial anterior optic tubercle (**Fig. 6e**). This structure is the largest optic glomerulus in the brain, and the only one with a retinotopic projection from the visual system³³. The optic lobe inputs to this structure are required for steering toward visual objects³⁴. Thus, the cells we reconstructed might participate in visual object-directed steering. We also found a cell type (LC9) that projects directly to DNa02 from the lobula of the optic lobe (**Fig. 6f**). FlyEM annotations support the conclusion that DNa02 receives many direct inputs from the optic lobe and medial

anterior optic tubercle.

Superior protocerebrum. We reconstructed two cell types that project to DNa02 from the superior medial protocerebrum and crepine (**Fig. 6g**). These areas are both major outputs of the mushroom body and lateral horn^{35,36}. Therefore, these cells might carry olfactory and thermosensory signals to DNa02 neurons. Indeed, in silkworms, cells with similar morphologies are odor-responsive^{37,38}. FlyEM annotations suggest that there are many similar cell types presynaptic to DNa02.

Ventrolateral protocerebrum. We reconstructed one cell type that projects to DNa02 from the anterior ventrolateral protocerebrum and the posterior ventrolateral protocerebrum (**Fig. 6h**). These brain regions receive auditory/mechanosensory input, as well as visual input. Both auditory and visual responses have been described in the cricket neurons that resemble DNa02^{18,21,22}. FlyEM annotations suggest that there are many neurons like these which are presynaptic to DNa02.

Contralateral LAL. We found that DNa02 receives direct input from the contralateral LAL, via LAL bilateral neurons (**Fig. 6i**). FlyEM annotations support this conclusion, and indeed similar cells have been described in silkworms¹⁴. These cells may compare input to the right and left LAL.

Other DNs. We found that DNa02 receives direct input from other DNs (**Fig. 6j**). Indeed, FlyEM annotations suggest that many DNs connect to DNa02. DN→DN synapses may coordinate DNs with related functions.

In summary, we find that DNa02 receives disynaptic input from compass neurons. Meanwhile, DNa02 also receives input from sensory pathways that bypass the compass system. This finding can explain why stimulus-directed steering does not require compass neurons^{12,13}.

Transforming compass signals into steering drives

The mechanisms that link the brain's spatial maps with motor control systems are not understood in any species. Our EM reconstructions show that, in the *Drosophila* brain, PFL3 neurons constitute a remarkably direct link between the brain's internal compass and steering descending neurons. This finding motivated us to investigate the anatomy of PFL3 neurons in detail, to determine how these neurons might transform compass readouts into steering drives. We found similar PFL3 connectivity in the full fly brain dataset³¹ and in the recent 'hemibrain' dataset³², implying that the PFL3 wiring diagram is fairly stereotyped across brains.

In this wiring diagram (**Fig. 7a**), each PFL3 neuron receives selective input from specific compass neurons, and it sends selective output to either the right or left copy of DNa02. Based on our electrophysiological recordings, we expect the right-left difference in DN firing rates to predict the fly's rotational velocity. Moreover, based on previously published work, we know how the fly's rotational velocity is used to continuously update the compass^{9,10,39}. We synthesized all these results to construct a computational model of the brain-behavior feedback loop (**Fig. 7a**).

For simplicity, in the initial form of this model, the only inputs to PFL3 neurons were from compass neurons. Each PFL3 neuron simply rectified its compass input and transmitted it to a DN, which then summed those PFL3 inputs. The right-left difference in DN firing rate then became a steering signal which continuously updated the fly's heading direction at each time step of the simulation, as well as continuously updating the compass.

This model exhibited a striking property: it steered the fly toward a specific stable heading, regardless of the fly's initial heading (**Fig. 7b**). This occurred because the network always converged to a stable state where left

and right DNs received equal input (**Fig. 7c**, top). Thus, this network transformed the output of the compass (a ring attractor¹¹) into a point attractor – i.e., a system with only one stable state.

These point attractor dynamics depended critically on the wiring of PFL3 neurons. Specifically, there are two special PFL3 neurons whose preferred heading is not shared by another PFL3 neuron projecting to the opposite DN (**Fig. 7a**). As a result, when the “compass bump” touches either of these special PFL3 neurons, the fly turns. When we omitted these two PFL3 neurons from the network, then all headings were equally stable (**Fig. 7b**).

To generate changes in goal heading over time, we needed to add a mechanism that can shift the stable state of the model network. To do this, we augmented the model by giving each PFL3 neuron a synaptic input in the fan-shaped body (FB). Most of the FB input to PFL3 neurons originates with FB local neurons, according to the FlyEM ‘hemibrain’ annotation, but little is known about FB local neurons. Therefore, we simply modeled FB local neurons as providing a static spatial pattern of inhibition to PFL3 neurons. When we switched on the local neurons in the left FB, we found this caused a shift in the network’s stable state (**Fig. 7c**, bottom). As a result, a new stable heading direction emerged (**Fig. 7d**). Conversely, when we switched on the local neurons in the right FB, or the middle of the FB, these inputs again produced new and different stable headings (**Fig. 7d**). Thus, a simple change in FB input can change the system’s policy for translating heading into steering. In essence, the static spatial pattern of FB neuron activity specifies the “policy” that translates a spatial state (heading) into an action (steering).

Finally, we incorporated stimulus-directed steering into the model, by giving DNs a lateralized sensory input that bypasses the compass (**Fig. 7e**). Partway through the simulation, after the fly’s heading had stabilized, we simulated the appearance of an object in the environment (**Fig. 7f,g**). When the object was on the right, the sensory input to the right DN was larger, and *vice versa*. DNs summed the signals they received from the sensory stimulus pathway and the compass pathway. As expected, the fly steered in the direction of the object as long as it was present. When the object disappeared, the compass pathway automatically caused the fly to turn back to its original goal heading (**Fig. 7f,g**). Importantly, the fly’s behavior in this model reproduces the behavior of a real fly confronted with a sequence of observable and remembered stimuli. Specifically, a real fly can be cued to head toward an unobserved (remembered) goal, and then cued to detour away from that path by a transient “distractor” object; when the distractor disappears, a fly typically turns back toward its original remembered goal, but only if its compass system is intact⁶. Our model provides a mechanistic explanation for this behavioral sequence. The conceptual solution is to generate steering commands based on convergent parallel control loops. When the stimulus-driven loop is activated, it can override the compass loop. Then, when the stimulus-driven loop is deactivated, the compass loop regains control and can produce the proper compensatory steering signal to turn the fly back toward its original remembered goal.

Contributions of single descending neuron types to steering behavior

Our results indicate that at least four DNs are recruited during steering maneuvers – i.e., two copies of DNa02, and two copies of DNa01. The population of steering-related DNs may indeed be larger than this. In the context of this DN ensemble, do single DNs make any measurable contribution to steering behavior? A previous study has shown that bilateral optogenetic stimulation of either DNa01 or DNa02 promotes walking⁴⁰, but the effects of unilateral stimulation have not yet been investigated.

To address this question, we used *hs-FLP* to stochastically express a channelrhodopsin variant (ReaChR) in DNa02 neurons. This produced either unilateral expression or bilateral expression (**Fig. 8a**). As each fly walked on a spherical treadmill, we illuminated its head from above to activate the ReaChR+ neuron(s). This illumination was symmetric on the two sides of the head.

In the flies with unilateral expression, we found that light evoked a small average steering bias in the ipsilateral

direction – i.e., the direction of the ReaChR+ neuron (**Fig. 8b**). To determine if this steering bias was significant, we used control sibling flies with bilateral ReaChR expression in DNa02. Specifically, each fly with bilateral expression was randomly assigned a label (“right expression” or “left expression”). Using these flies, we then computed the average rotational velocity in the “ipsilateral” direction, just as we did for the flies with true unilateral expression. This simulation was repeated independently for 1000 bootstrap replicates (resampling with replacement) to obtain the 95% confidence interval of the mean, under the null hypothesis that there is only a random relationship between DNa02 expression and steering behavior. We found that the mean of the data (with unilateral expression) was outside the 95% confidence interval of the simulation outcomes (**Fig. 8b,c**). Thus, in flies with unilateral expression, the turning bias was significantly larger than we would expect by chance.

In these experiments, symmetric illumination from above also caused changes in translational velocity. These effects were likely due to visual startle, because they were indistinguishable in flies with unilateral expression, bilateral expression, and no expression (data not shown). Due to the confounding effect of visual startle, we could not use this dataset to investigate how optogenetic activation of DNa02 affects the fly’s translational velocity.

Next, to investigate the effect of bilateral DN silencing, we expressed an inhibitory opsin (GtACR1) in either DNa01 neurons or DNa02 neurons (**Fig. 8d**). As a genetic control, we also used an “empty” split-Gal4 line. As each fly walked in a small arena, we switched an overhead light on and off periodically (2 min on, 2 min off). We analyzed behavior as a function of genotype (Gal4/control) and light (on/off) using a set of two-factor ANOVAs, corrected for multiple comparisons. A significant (genotype×light) interaction would be evidence for a behavioral effect of DN silencing.

First, we analyzed the body movements that are correlated with DNa01 and DNa02 activity – namely, rotational and sideways movements (**Fig. 8e**). When DNa01 neurons were silenced, sideways speed was significantly reduced. Surprisingly, however, rotational speed was unchanged when either DNa01 or DNa02 neurons were silenced.

Recall that DNa01 and DNa02 predict rotational movements about as well as they predict sideways movements (**Fig. 1f**). Moreover, unilateral activation of DNa02 influences ipsilateral rotational movements (**Fig. 8b,c**). Why, then, is the body’s rotational movement unaltered when these DNs are bilaterally silenced? Importantly, body rotation can be achieved via many different leg maneuvers^{24,41}. For example, in some body-rotation events, the stance duration of the inner back leg is prolonged relative to the other legs – in essence, the fly pivots on its inner back leg – whereas this maneuver is absent in other body-rotation events (**Fig. 8f**). Perhaps DN silencing interferes with a specific leg maneuver, rather than body rotation *per se*.

To investigate this idea, we identified five leg maneuvers that often accompany body rotations in control flies (**Extended Data Fig. 9**). We then examined whether each leg maneuver was altered by DN silencing. We found that silencing DNa01 produced a significant defect in stance prolongation in the inner back leg – i.e., it reduced the fly’s tendency to pivot on the inner back leg during a body rotation (**Fig. 8g**). We observed the same result when we silenced DNa02, although here the result fell just short of our cutoff for statistical significance. Thus, silencing DNa01 (and possibly DNa02) causes a defect in inner-back-leg-stance-prolongation, but this defect is evidently compensated via other mechanisms to achieve a normal rotational movement of the body.

Finally, we noticed that silencing either DNa01 or DNa02 altered the fly’s normal leg movement rhythm, regardless of whether the fly was walking straight or turning. Specifically, step frequency decreased, while leg swing distance increased (**Fig. 8h**). Silencing DNa01 also decreased the fly’s translational velocity. Notably, all these results phenocopy the consequences of proprioceptor ablation⁴². These results are consistent with the idea that DNs modulate walking by controlling the gain of proprioceptive reflex loops^{42,43}. Specifically, we propose

that lateralized DN activity produces lateral asymmetries in proprioceptive reflexes, leading to asymmetric leg movements, and thus turning.

Discussion

Ensemble codes for motor control

In this study, we showed that steering maneuvers during walking recruit at least two types of descending neurons in the *Drosophila* brain. Both DN types correlate with rotational velocity and sideways velocity in the ipsilateral direction (Fig. 5). That said, the two DN types also have non-identical functions. Whereas DNa02 predicts large, rapid steering events, DNa01 predicts smaller and slower steering events (Figs. 5 and 6). In addition, we found that silencing DNa01 decreases sideways speed and translational velocity, whereas silencing DNa02 does not, possibly because the consequences of DNa02 silencing are more completely masked by compensatory mechanisms. Notably, DNa01 and DNa02 have distinct axonal projection patterns, suggesting they have non-identical downstream targets. Whereas DNa01 terminates in the medial part of the leg neuropil, DNa02 projects through the entire medial-lateral extent¹⁵ (**Fig. 1a**). Given the myotopic map of motorneuron dendrites within each leg neuropil⁴³, DNa01 may influence motorneurons controlling proximal leg muscles, whereas DNa02 may influence motorneurons controlling the entire leg.

We hypothesize that additional DNs also participate in steering during walking, because silencing either DNa01 or DNa02 produces only small defects in steering. In the future, to better understand the contributions of different steering DNs, it will be important to compare DN spike trains with simultaneous leg kinematic measurements. It will also be important to understand the circuits downstream from these DNs in the ventral nerve cord. One intriguing possibility is that lateralized activity in these DNs generates lateral asymmetries in leg proprioceptive reflexes, leading to asymmetric leg movements, and thus turning. Consistent with this idea, we found that silencing steering DNs can phenocopy the consequences of proprioceptor ablation⁴². Indeed, there is already evidence from other insect species that DNs can modulate walking by controlling the gain of proprioceptive reflex loops^{42,43}.

In the future, it should be possible to compare the brain circuits upstream from different steering DNs. In particular, steering in flight involves specific DNs which are distinct from those we have identified here⁴⁴. Indeed, the DNs that target the wings and the legs are largely distinct populations. This raises the question of how much upstream brain circuitry is shared for flight control and walking control, and what happens to these systems when they are not currently in use.

In this regard, it is perhaps relevant that we observed clear responses to lateralized sensory stimuli in steering DNs even when the fly was not walking (**Extended Data Fig. 10**). Similarly, a recent study found that *Drosophila* DNs that evoke landing during flight are still responsive to visual stimuli when the fly is not flying⁴⁵. We might think of these residual DN responses as a latent bias toward action, or a preparation for an action that never ultimately occurs⁴⁶. In this regard, it is relevant that mammalian corticospinal neurons can be active in the absence of movement, either in the period preceding a movement⁴⁷ or when a movement is being simply observed⁴⁸. These latent corticospinal signals may represent movement preparation, which may additionally involve neurons which are not actually active during movement⁴⁹. The discovery of latent motor activity in the *Drosophila* brain provides a potential starting point for future mechanistic investigations of brain dynamics during motor preparation.

Integrating compass-directed steering and stimulus-directed steering

Recent studies in *Drosophila* suggest there are mechanisms downstream from compass neurons that compare the current heading with some stored goal heading, in order to keep the fly on course toward its goal^{12,13}. These

studies found that, when a fly is steadily locomoting straight toward an unseen target, the pattern of activity in compass neurons is constant – in other words, the brain’s “compass needle” is steady. If the fly (or the compass) is pushed away from this steady orientation, the fly typically makes a corrective turn to regain its original goal heading, thereby bringing the compass back in its originally steady position. These findings have raised the question of how the goal is represented in the brain, and how it is compared to the fly’s current heading.

Our results provide insight into the mechanisms of goal-oriented steering by the compass system. We found that the magnitude and direction of compass-directed steering maneuvers are predicted by DNa02 neurons. Moreover, there is only one network layer interposed between DNa02 and compass neurons. This layer is composed of a cell type (PFL3). Using EM reconstruction of PFL3 neurons, we were able to map the wiring pattern of PFL3 neurons onto DNa02 neurons. When we instantiated this wiring pattern in a computational model, we found that – in the absence of any sensory stimulus – the network steers the fly in a specific heading direction, regardless of the fly’s initial heading. Essentially, the model network is a ring attractor (the compass) embedded within a point attractor – a system with only one stable state. The model’s point attractor dynamics depend on the specific wiring pattern of PFL3 neurons. These point attractor dynamics also depend on the fact that the fly’s steering behavior is fed back to the compass to generate an updated heading representation, causing an update in the pattern of PFL3 activity, DN activity, and behavior.

There are likely mechanisms to rotate the brain’s internal heading goal as circumstances warrant. We conjecture that the goal heading is rotated by shifting the spatial pattern of FB inputs to PFL3 neurons. These FB inputs would specify the policy that translates state (heading) into action (steering). The FB receives projections from higher brain regions involved in reinforcement learning – specifically, the mushroom body and its associated regions in the superior protocerebrum⁵⁰ – which may convey information about the relative value of different policies.

Our model can also explain how the fly can hold a remembered goal in working memory, while also making a detour toward a novel object to investigate it. This type of behavior was described in a classic study⁶ where a fly was cued to head toward an unseen (remembered) target. As the fly was walking straight toward the remembered target, a new object was made to appear to the side of the fly’s path, causing the fly to make a detour toward the object. When the object then disappeared, the fly was observed to turn back toward its original remembered target, implying the remembered target was retained in working memory; this memory was shown to require an intact compass system. Our model explains this behavior as arising straightforwardly from the concerted effects of two parallel feedback loops which converge onto DNs (**Fig. 7e**) – a stimulus-directed feedback loop which bypasses the compass, and a memory-directed feedback loop which engages the compass. When the stimulus-driven loop is activated by an object in the environment, it can override the compass loop. When that object disappears, the compass system regains control of steering. Because the compass operates continuously during the detour, it contains an up-to-date representation of the current heading, and it can generate the precise steering maneuver required to re-orient toward the remembered goal. This same process would operate when the fly detours to avoid an obstacle or threat along its path.

Recent work has just identified DNs in the mammalian brainstem that influence steering⁵. Given this, it should be possible to work backward from mammalian steering DNs to find points of convergence between memory- and stimulus-directed steering pathways. In mammals, these pathways are thought to involve the hippocampus and striatum, respectively². Ultimately, these pathways must be integrated to produce sequences of navigation behaviors. This problem has generally been considered from the perspective of cognitive neuroscience; our aim here has been to approach this type of problem from the perspective of motor systems. Once we understand the neural code for motor control at the level of motor systems (e.g., descending neurons), it is possible to re-frame the problem of controlling behavior as a problem of controlling descending neural dynamics. Here we show how sequences of neural and behavioral dynamics during navigation can emerge from the integration of parallel feedback loops which converge at the level of the brain’s motor systems.

Author Contributions: A.R. and R.I.W. designed the study. A.R. built all experimental devices and performed all experiments. S.L.H. performed the analyses in Figs. 1, 2c, and 3d-e, as well as Extended Data Figs. 1, 4c, and 10d. I.D. performed analyses and visualizations of EM reconstructions in Fig. 6. A.A.L., Q.X.V., I.S.H., and P.W.G. reconstructed neurons from EM data in Fig. 6. R.I.W. performed computational modeling in Fig. 7. A.R. performed all other analyses. A.R. and R.I.W. wrote the manuscript with input from S.L.H. and I.D.

Acknowledgments: We are grateful to Gwyneth Card, Michael Dickinson, Hiro Namiki, and Wyatt Korff for sharing descending neuron morphology data and split-Gal4 drivers (*SS00730* and *SS00731*) pre-publication. Luke D. Lavis and Jonathan B. Grimm provided SiR110-HaloTag dye pre-publication. Julian Ng, Sebastian Cachero, and Gregory Jefferis shared *pJFRC81-td3-Halo7::CAAX(attP18)* flies pre-publication. Tzumin Lee shared *Orco-LexA* flies, Marco Gallio shared *Gr28b.d-LexA* flies, Gero Miesenböck shared *UAS-P2X₂* flies, Tom Clandinin shared *hsFLP.70* flies, and Mike Crickmore shared *UAS-GtACR1::eYFP(VK00005)* flies. Jonathan Green provided excellent advice on central complex stimulation. Tom Kazimiers helped transfer our EM reconstructions to the FAFB-v14 community workspace. Mert Erginkaya, Saba Ali, Kelli Fairbanks, Tansy Yang, Emily Tenshaw, Markus Pleijzier, Imaan Tamimi, Eugenia Chiappe, Vivek Jayaraman, Barry Dickson, Gwyneth Card, and Greg Jefferis contributed to reconstructing neurons in Fig. 4. Douglas Hayden provided statistical advice. Richard Mann and members of the Wilson lab provided feedback on the manuscript. A.R. is supported by National Research Service Award 5F31DC015701. S.L.H. is supported by a National Research Service Award 5F31NS106982. This work is supported by NIH grants R01DC008174, R01NS101157, and U19NS104655. R.I.W. is an HHMI Investigator.

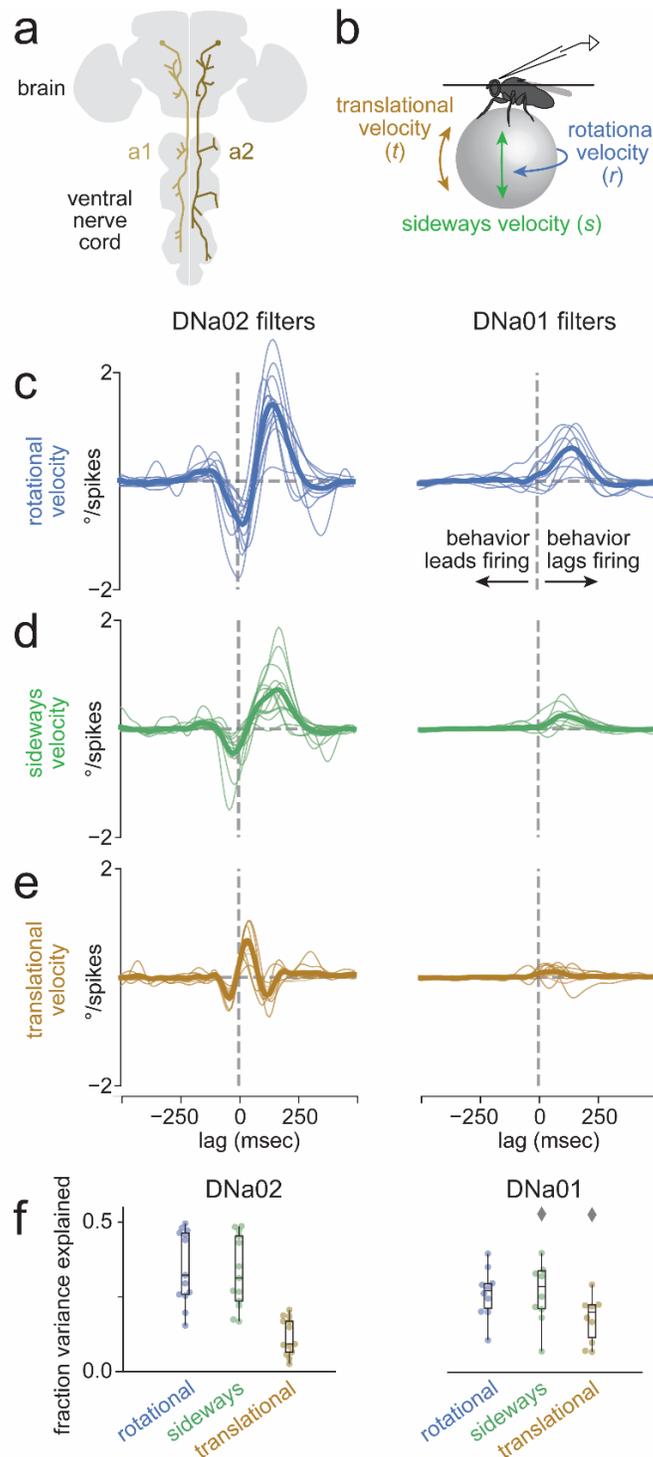


Fig. 1: Specific descending neurons that predict steering behavior.

a, Schematic morphologies of DNa01 and DNa02. These two descending neurons have similar dendritic arbors that extend into several brain regions, including the LAL.

b, Schematic showing the three axes of movement on the spherical treadmill. **c**, Rotational velocity filters for DNa02 and DNa01. These filters describe the average rotational velocity impulse response, given a delta function (unit impulse) in firing rate. Thin lines are filters for individual flies ($n=13$ for DNa02, $n=10$ for DNa01), thick lines are averages. For positive lags, behavior follows firing; for negative lags, behavior leads firing. As expected, the filter is ~ 0 for negative lags. Extended Data Fig. 1 shows the reverse analysis (behavior \rightarrow neuron filters).

d-e, Same for sideways and translational velocity.

f, Variance explained by each filter type. Data points are flies. Boxplots show interquartile range (IQR), whiskers show range of data (except for diamonds, observations $>1.5 \times \text{IQR}$).

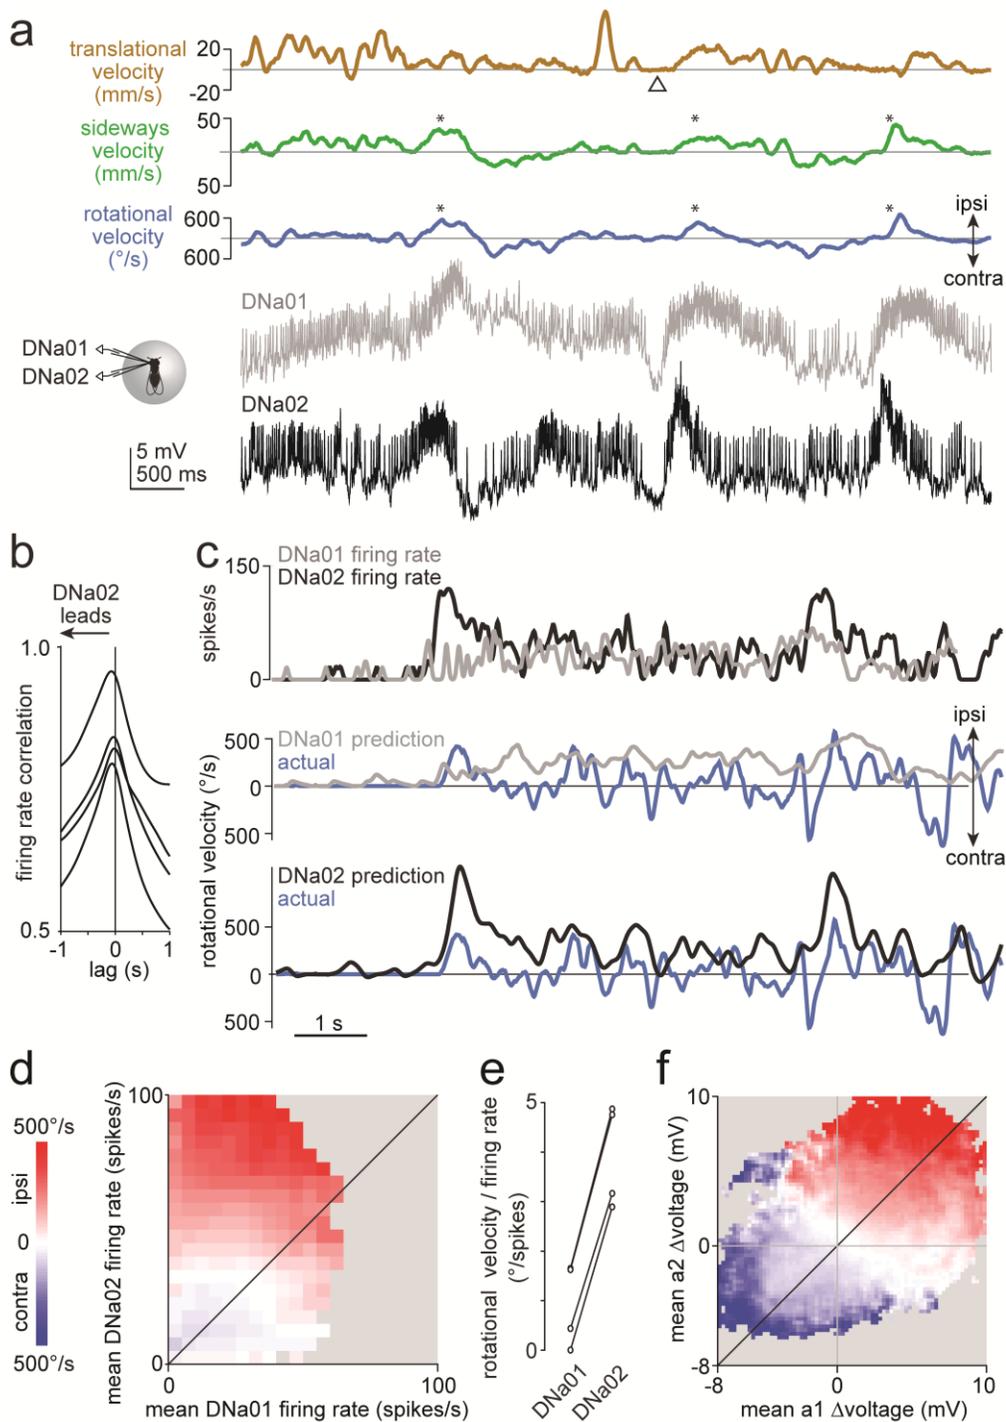


Fig. 2: DNa01 and DNa02 predict steering with different kinetics and different gain.

a, Example dual recording from DNa01 and DNa02, both on the left. Both cells are depolarized just before the fly steers in the ipsilateral direction (asterisks). Firing rate increases tend to be more transient for DNa02 than for DNa01 (see also Extended Data Fig. 1). Note that both cells are hyperpolarized when the fly briefly stops moving (Δ).

b, Correlation between DNa01 and DNa02 firing rates. Each line is a different paired recording (4 flies).

c, We computed the rotational velocity filter for each cell, and then convolved each cell's filter with its firing rate to generate rotational velocity predictions. Whereas DNa02 over-predicts large rapid steering events, DNa01 under-predicts these events.

d, Colormap shows binned and averaged rotational velocity for each value of DNa01 and DNa02 firing rates for an example paired recording. Note the plot is not symmetrical about the line of unity; instead, rotational velocity is more sensitive to DNa02 than to DNa01.

e, Slope m of the linear regression fitting rotational velocity r to each cell's firing rate f ($r=mf+b$). The difference between cell types is statistically significant ($p=10^{-4}$, paired t-test, $n=4$ flies).

f, Same as (**d**) but for voltage changes rather than firing rate changes.

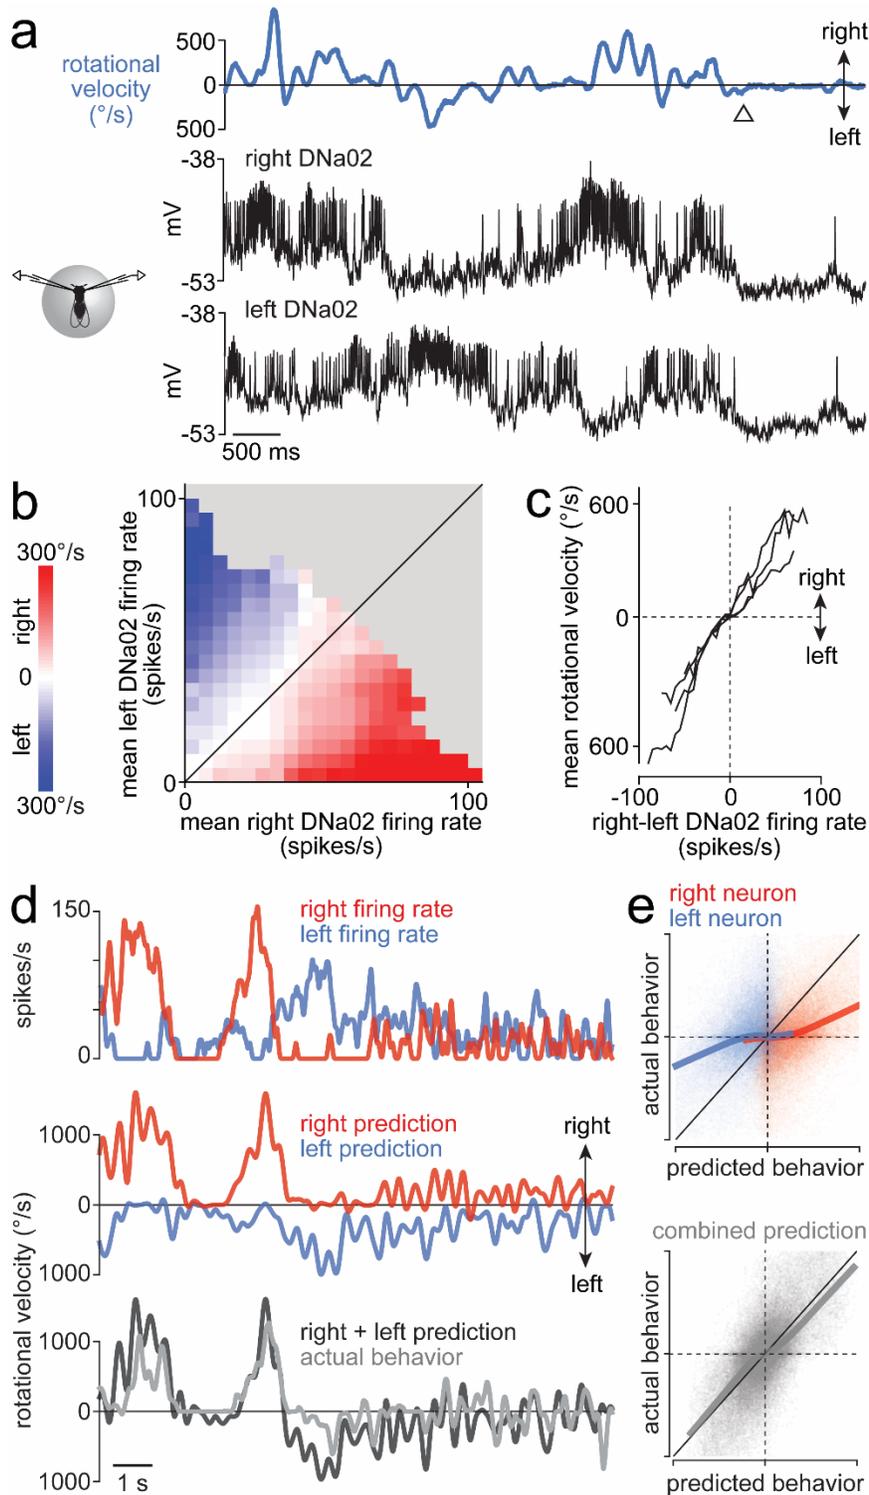


Fig. 3: Steering is linearly related to right-left differences in DN firing rates.

a, Example bilateral recording from DNa02. The neurons are anti-correlated when the fly is turning. Both neurons are hyperpolarized when the fly stops moving (Δ).

b, Colormap shows binned and averaged rotational velocity for each value of bilateral DNa02 firing rates for an example paired recording.

c, Mean rotational velocity for each value of the bilateral firing rate difference. Each line is a different paired recording (n=4 flies).

d, We computed the rotational velocity filter for each cell, and then convolved each cell's filter with its firing rate to predict rotational velocity. Combining the predictions of the two cells (with equal weighting) generates a good prediction.

e, Top: predicted rotational velocity versus actual rotational velocity, for both single-cell predictions for an example experiment. Bottom: the dual-cell prediction for this experiment. Thick lines are LOWESS fits. Each axis ranges from 800 $^{\circ}/s$ rightward to 800 $^{\circ}/s$ leftward, and dashed lines denote 0 $^{\circ}/s$.

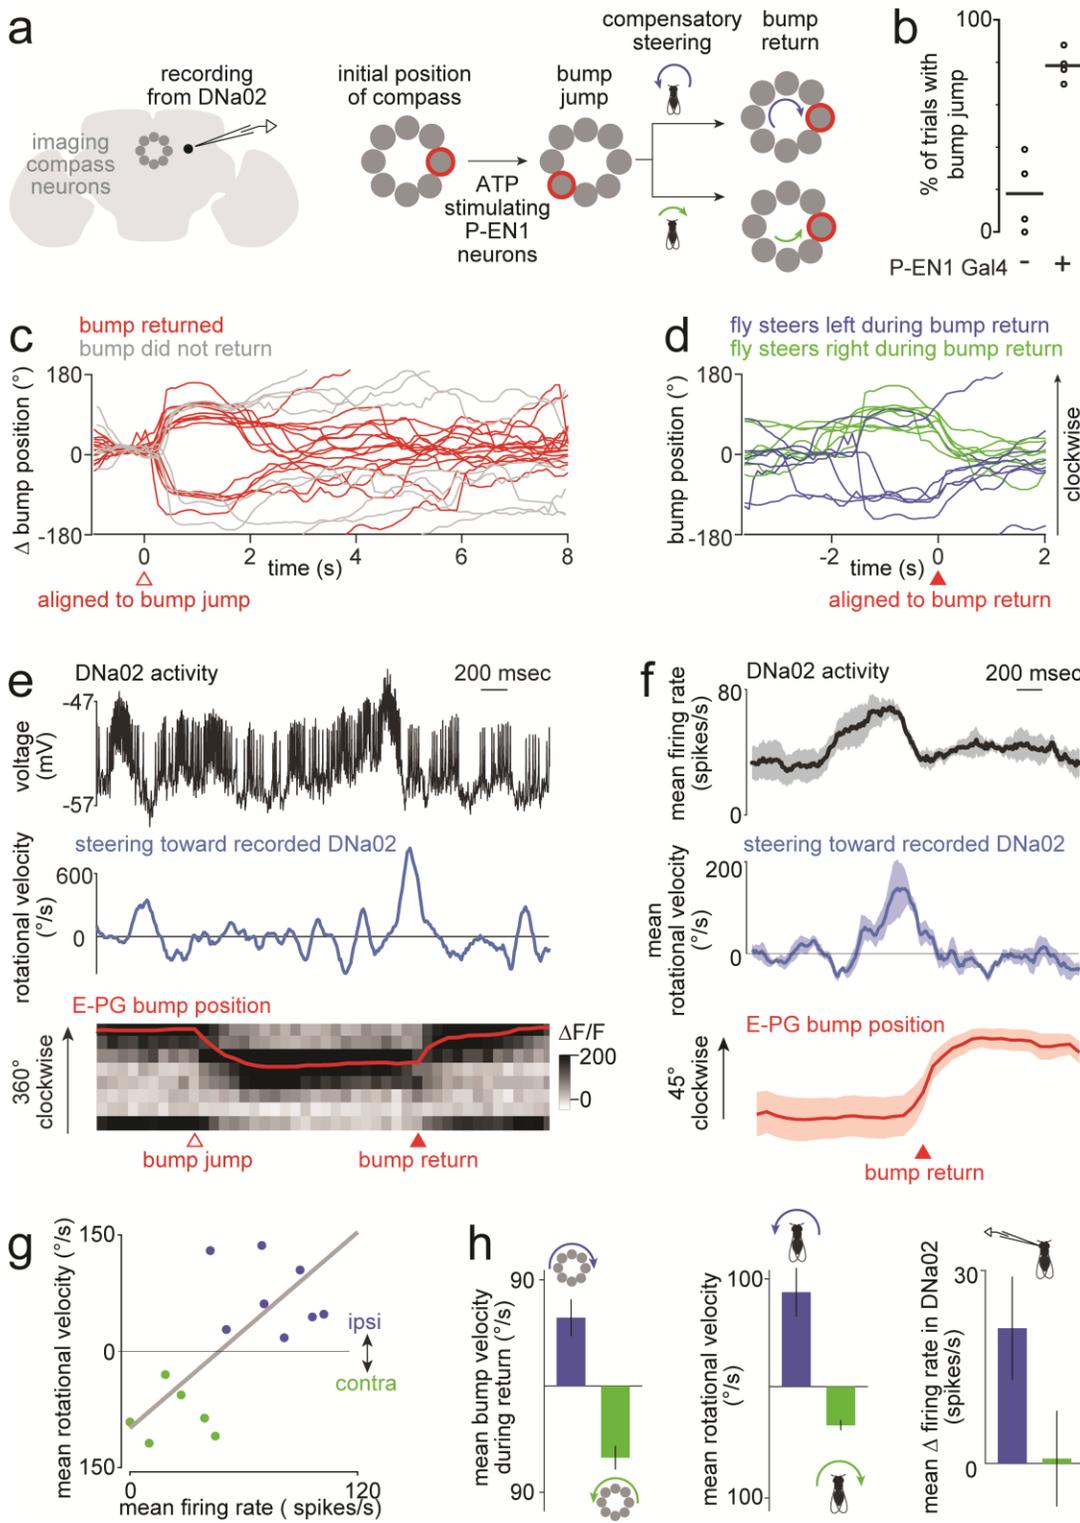


Fig. 4: DNa02 participates in compass-directed steering.

a, A bump of activity rotates around the ensemble of E-PG neurons (compass neurons) as the fly turns. We microstimulate P-EN1 neurons with ATP to evoke a bump jump. This drives compensatory steering which returns the bump to its initial location.

b, Trials with a bump jump after the ATP puff. Each point is a fly (each genotype $n=4$). In controls lacking P-EN1 Gal4, ATP only occasionally preceded bump jumps, which is likely coincidental, as the bump often moves spontaneously.

c, Change in bump position after ATP. The bump returns to its initial position in many trials (red) but not all (gray) in this experiment.

d, Red trials from **c** aligned to the time of maximal bump return speed, color-coded by steering direction during bump return.

e, Example trial. Top: DNa02 activity. Middle: fly's rotational velocity toward the left side, i.e., the side of the recorded DNa02 neuron. Bottom: grayscale E-PG $\Delta F/F$ over time, where each row is a 45 $^{\circ}$ -sector of the compass, and red is bump position. ATP causes a bump jump (Δ). Then, the left copy of DNa02 bursts, the fly turns left, and the bump returns via a clockwise path (\blacktriangle).

f, Trials where the bump returned to its initial location via a clockwise path were aligned to the time of peak

bump speed. Trials were averaged within a fly, and then across flies (mean \pm SEM across flies, $n=4$ flies).

g, In an example experiment., rotational velocity correlates with DNa02 firing rate on a trial-to-trial basis ($R^2=0.51$, $p=3 \times 10^{-3}$, two-tailed t-test). Trials are color-coded as in **d**. In all other experiments, p was also <0.05 .

h, Data sorted by the direction of the bump's return (mean \pm SEM across flies, $n=4$ flies). Whereas clockwise (blue) bump returns were typically preceded by leftward turning, counter-clockwise (green) bump returns were preceded by rightward turning, as expected. On average, the left copy of DNa02 was only excited on trials where bump moved clockwise, meaning the fly was turning left.

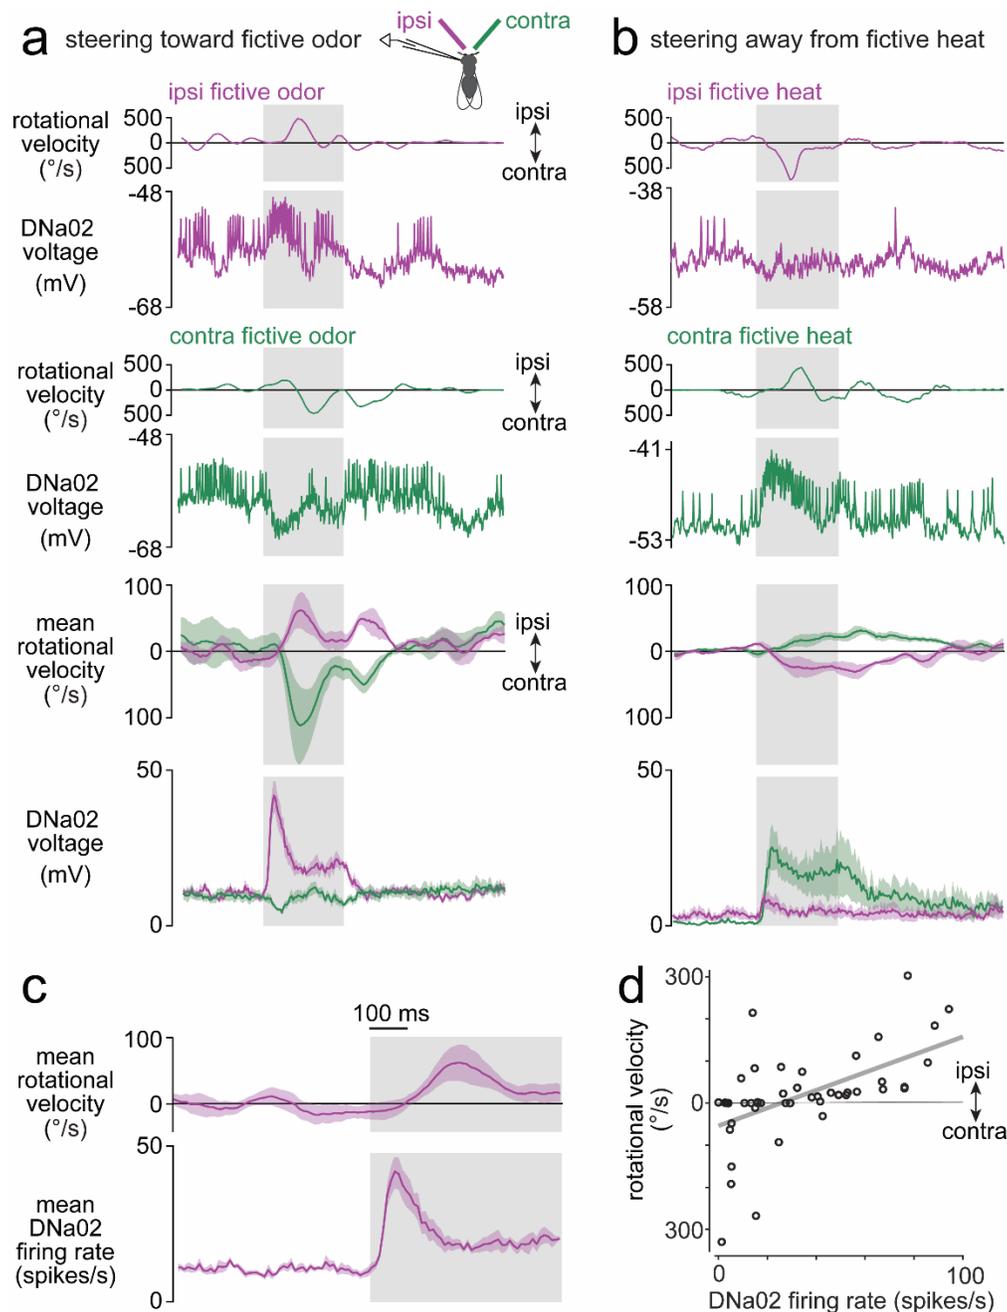


Fig. 5: DNa02 participates in stimulus-directed steering.

a, CsChrimson was expressed in most olfactory receptor neurons. As a fly walked on a spherical treadmill, a fiber-optic filament illuminated either the right or left antennae alternately. Ipsi- and contralateral stimuli are defined relative to the recorded neuron. Top: two example trials, one ipsi and one contra, showing the fly's rotational velocity and DNa02 activity on each trial. Bottom: mean \pm SEM across flies, $n=4$ flies. Gray shading shows the 500-ms period of fictive odor.

b, Same as **a** but for fictive heat. CsChrimson was expressed in heat-activated neurons of the antenna. Fictive heat drives behavioral turning away from the stimulus, rather than toward it.

c, Mean data on an expanded timescale to show that DNa02 firing rate increases precede turning toward ipsilateral fictive odor (mean \pm SEM across flies, $n=4$ flies).

d, Trial-to-trial variability in an example experiment. Each datapoint is a trial where fictive odor was presented on the ipsilateral side; gray line is linear fit ($R^2=0.33$, $p=10^{-5}$, two-tailed t-test). In other experiments, R^2 ranged from 0.16 to 0.40, with p always <0.005 .

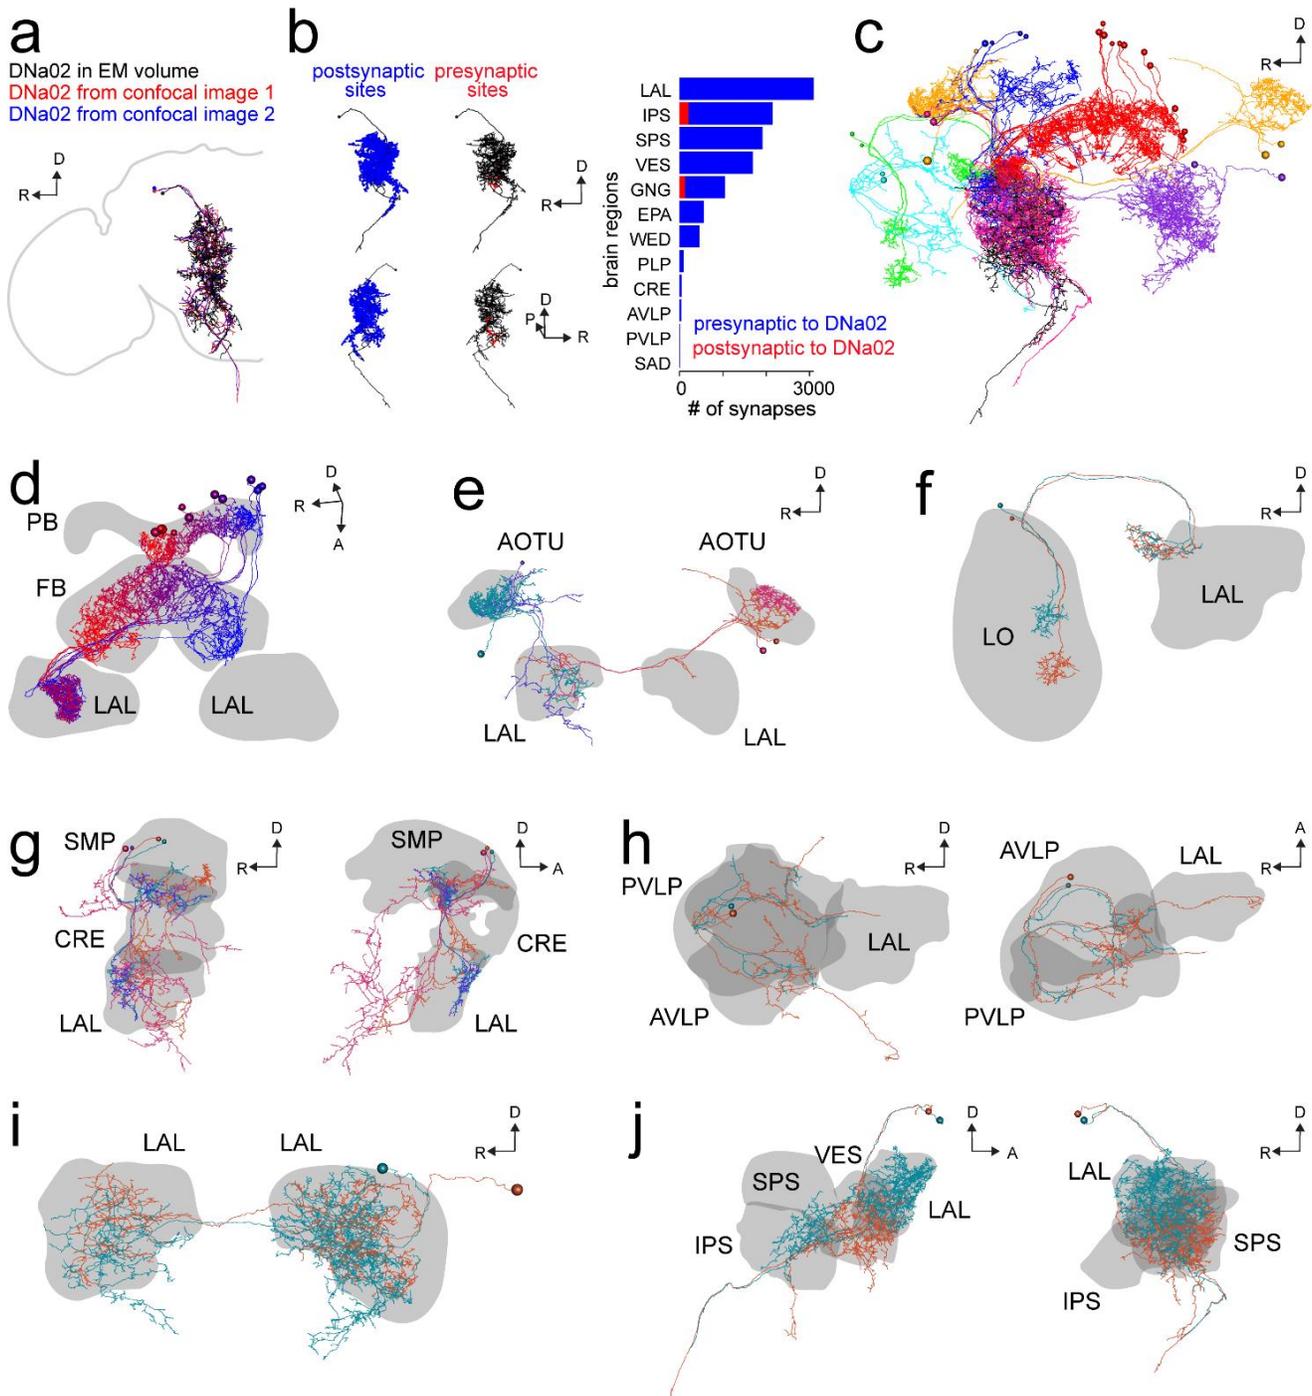


Fig. 6: Pathways for internal compass cues and external sensory cues converging onto DNa02.

a, The cell we identified as DNa02 in the EM dataset (anterior view, right brain outlined), overlaid with the DNa02 morphology we traced from confocal data. Scale bars are 50 μm (here and throughout this figure).

b, All pre- and postsynaptic sites in DNa02, shown from two angles (left) and tabulated by brain region (right).

c, All DNa02 input neurons selected for full EM reconstruction. Cell types are shown in different colors, with DNa02 in black.

d, Central complex output neurons (PFL3 neurons). Inputs to DNa02 on the right arise primarily from the left protocerebral bridge (PB). These cells have dendrite-like arbors in the PB and fan-shaped body (FB). Cells are color-coded by position in the PB and FB.

e, Four input neurons with dendrites in the anterior optic tubercle (AOTU). Two cross the midline (magenta/orange).

f, Two input neurons with dendrites in the lobula (LO). These are lobula columnar type 9 cells (LC9).

g, Four input neurons arising from the superior medial protocerebrum (SMP) and crepine (CRE). Two project to the anterior LAL (blue/green), while the others project posteriorly. Two views are shown.

h, Two input neurons arising from the posterior/anterior ventrolateral protocerebrum (PVLP/AVLP). Two views are shown.

i, Two input neurons arising from the contralateral LAL.

j, Two DNs that provide direct input to DNa02. Both arborize in the LAL, vest (VES), and superior/inferior posterior slope (SPS/IPS). Two views are shown. **Supplementary Table 1** lists synapses per connection and neuron names archived at virtuallyflybrain.org.

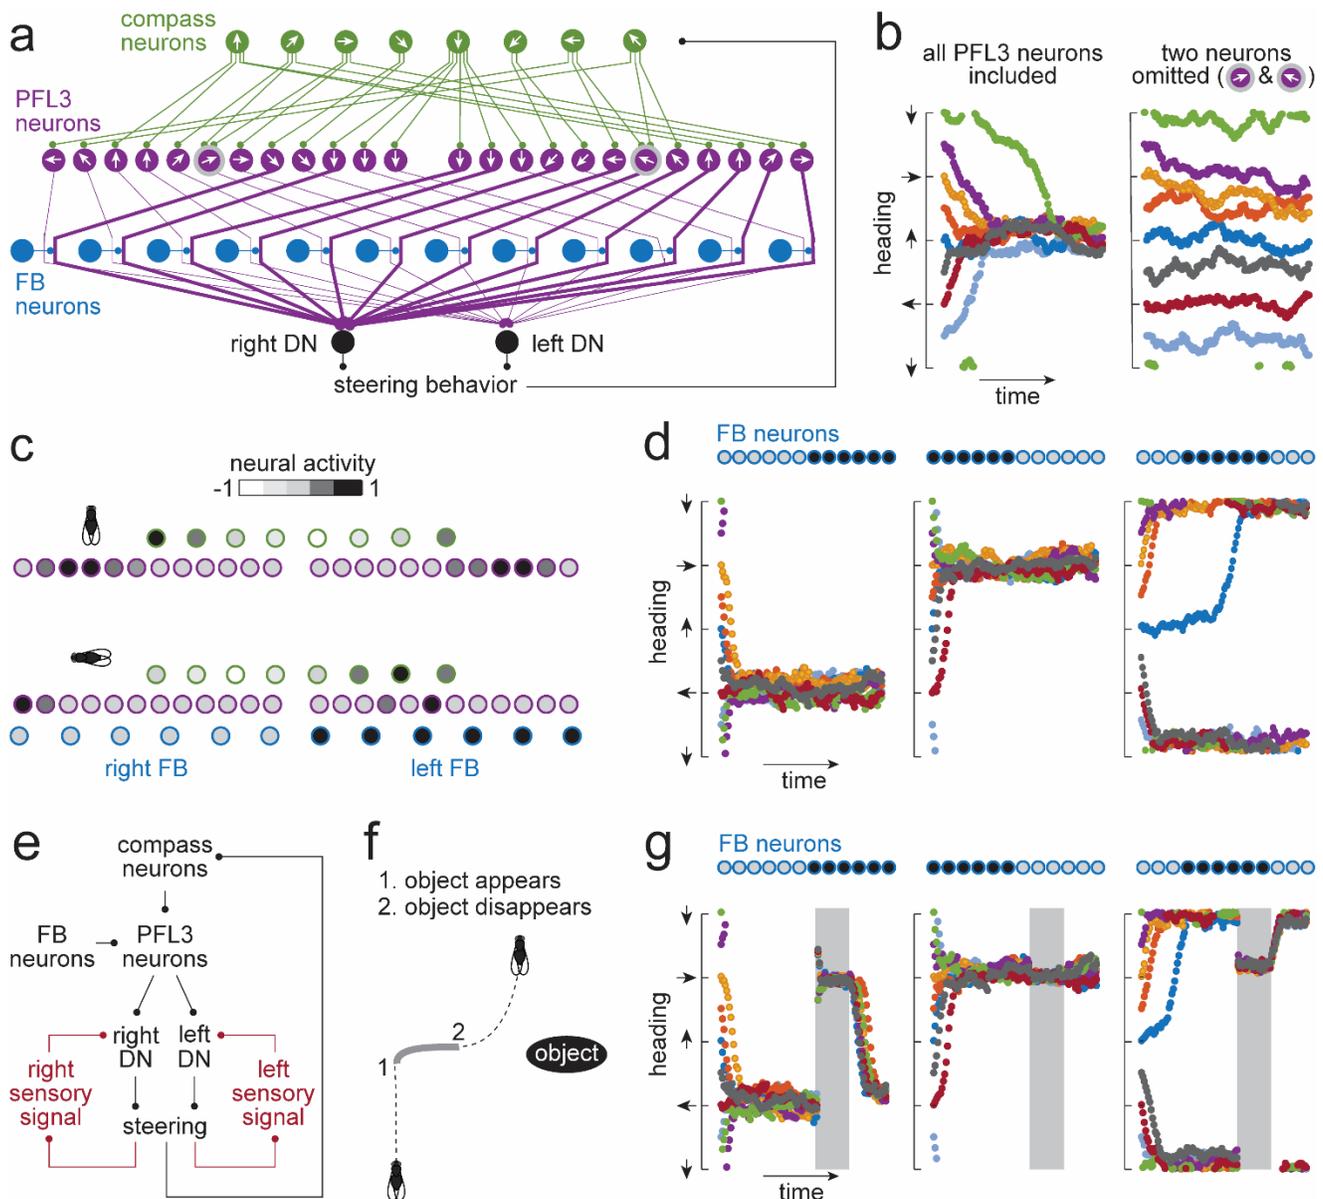


Fig. 7: A computational model of integrative steering control.

a, Architecture of the model. Compass neurons project to PFL3 neurons. The preferred heading of each compass neuron and PFL3 neuron is denoted by an arrow. Each PFL3 neuron also receives an input from an FB local neuron. Each PFL3 neuron sums its compass input and its FB input; this sum is then rectified and transmitted to its postsynaptic DN. PFL3 neurons projecting to right and left DNs are shown as thick and thin lines, respectively. The right/left difference in DN activity becomes the steering command which is used to update the fly's heading direction at the next time step of the simulation, thereby closing the feedback loop. A small amount of noise is added to the steering command to emulate the noisiness of real neural networks. Note that "right" and "left" refer to the fly's right and left, and the wiring diagram is shown as if viewed from the anterior face of the brain, following convention⁵¹.

b, Heading versus time (with arbitrary time units), for 9 initial headings. (FB local neurons are not included in this initial version of the model.) The network always steers the fly toward the same heading ("North"), regardless of the fly's initial heading. When two specific PFL3 neurons are omitted (indicated with gray rings in **a**), then all headings are equally stable.

c, Top: the system's stable state with no FB neurons. Bottom: the system's stable state after we add 12 FB local neurons to the model, and we selectively activate these neurons on the left side of the FB.

d, Heading versus time for three different spatial patterns of FB neuron activity. In each panel, results are shown for 9 initial headings. The activity patterns of the FB local neurons are shown above each panel.

e, Architecture of the full model. Compass neurons, PFL3 neurons, and FB local neurons are wired as before (see **a**), but here DNs also receive sensory input that bypasses the compass.

f, Schematic summary of the "detour" behavior paradigm⁶. As the fly is navigating toward a remembered spatial goal, an object appears along the fly's path, causing the fly to approach it. When the object disappears, the fly reverts to its remembered goal.

g, Model results. Same as **d**, but now an object appears during the time denoted with gray shading. We positioned the object to the East of the fly, and so the fly heads East when it is present.

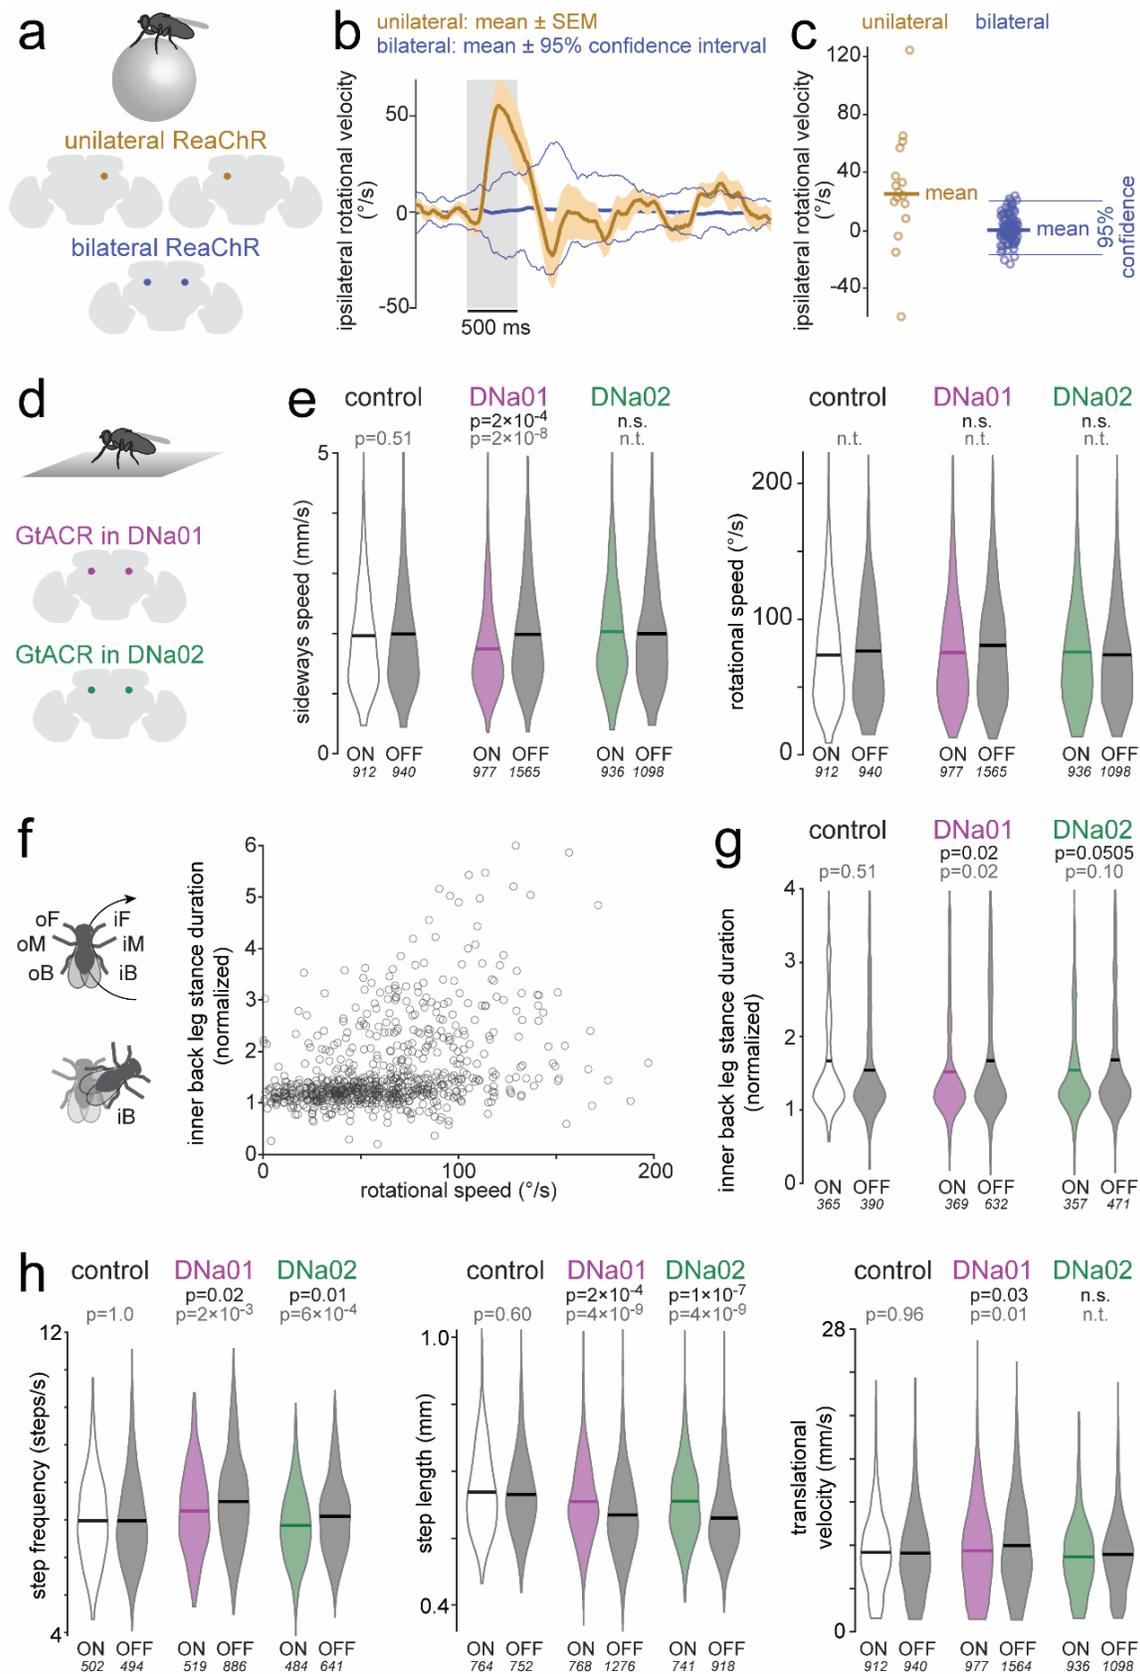


Fig. 8: Behavioral results of directly activating and silencing DNs.

a, As flies walked on a spherical treadmill, they were illuminated repeatedly for 500-ms epochs. ReaChR was expressed uni- or bilaterally in DNa02.

b, Ipsilateral rotational velocity, mean of 16/17 flies (uni/bilateral ReaChR). The mean of the unilateral data (thick gold line) lies outside the 95% confidence interval (thin blue lines) of the distribution of outcomes we obtain when we randomly assign bilateral

(control) flies to the “right” or “left” expression group.

c, Ipsilateral rotational velocity. Each data point is one fly (or simulated fly), averaged across the illumination epoch.

d, GtACR1 was expressed bilaterally in DNa01 or DNa02. Flies were illuminated repeatedly for 2-min epochs as they walked in an arena.

e, Sideways and rotational speed distributions for light on/off in each genotype. The p-values in black denote ANOVA genotype×light interactions after Bonferroni-Holm correction; the p-values in gray denote post hoc Tukey tests comparing light on/off within a genotype; n.s. = not significant, n.t. = not tested (because genotype×light interaction was not significant). The number of points (500-ms time windows) in each distribution is shown in italics; see Methods for fly numbers.

f, Stance duration of inner back (iB) leg, normalized to other legs (oF, oM, oB, iF) was measured in 500-ms time windows (n=755 windows) and plotted versus rotational speed. For some windows where rotational speed is high, iB stance duration is prolonged i.e., the fly pivots on its iB leg.

g, Distribution of normalized iB stance durations, for 500-ms time windows where rotational speed exceeded a threshold of 20 %/s for ≥ 100 ms.

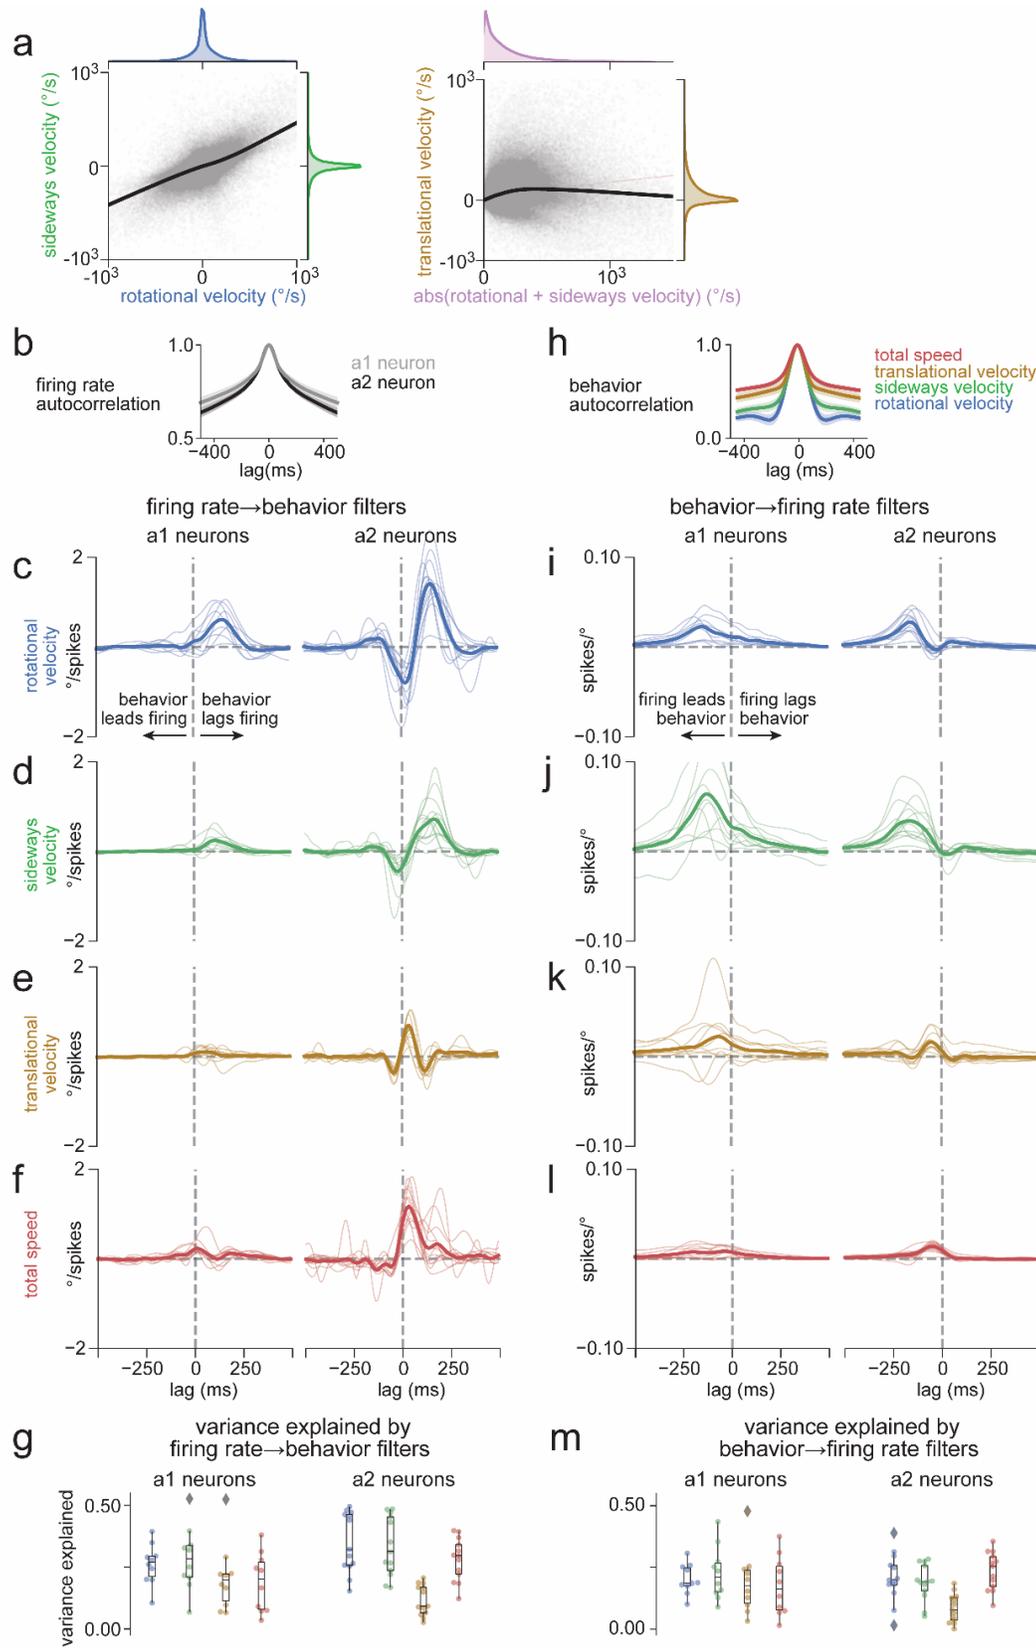
h, Step frequency, step length, and translational velocity distributions.

References

1. Eichenbaum, H. The role of the hippocampus in navigation is memory. *J Neurophysiol* **117**, 1785-1796 (2017).
2. Packard, M.G., Hirsh, R. & White, N.M. Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. *J Neurosci* **9**, 1465-1472 (1989).
3. Packard, M.G. & McGaugh, J.L. Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems. *Behav Neurosci* **106**, 439-446 (1992).
4. Packard, M.G. & McGaugh, J.L. Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiol Learn Mem* **65**, 65-72 (1996).
5. Cregg, J.M., *et al.* Brainstem neurons that command mammalian locomotor asymmetries. *Nat Neurosci* **23**, 730-740 (2020).
6. Neuser, K., Triphan, T., Mronz, M., Poeck, B. & Strauss, R. Analysis of a spatial orientation memory in *Drosophila*. *Nature* **453**, 1244-1247 (2008).
7. Bregy, P., Sommer, S. & Wehner, R. Nest-mark orientation versus vector navigation in desert ants. *J Exp Biol* **211**, 1868-1873 (2008).
8. Murakami, H., Tomaru, T. & Gunji, Y.P. Interaction between path integration and visual orientation during the homing run of fiddler crabs. *R Soc Open Sci* **4**, 170954 (2017).
9. Seelig, J.D. & Jayaraman, V. Neural dynamics for landmark orientation and angular path integration. *Nature* **521**, 186-191 (2015).
10. Green, J., *et al.* A neural circuit architecture for angular integration in *Drosophila*. *Nature* **546**, 101-106 (2017).
11. Kim, S.S., Rouault, H., Druckmann, S. & Jayaraman, V. Ring attractor dynamics in the *Drosophila* central brain. *Science* **356**, 849-853 (2017).
12. Giraldo, Y.M., *et al.* Sun navigation requires compass neurons in *Drosophila*. *Curr. Biol.* **28**, 2845-2852 e2844 (2018).
13. Green, J., Vijayan, V., Mussells Pires, P., Adachi, A. & Maimon, G. A neural heading estimate is compared with an internal goal to guide oriented navigation. *Nat. Neurosci.* **22**, 1460-1468 (2019).
14. Namiki, S. & Kanzaki, R. Comparative neuroanatomy of the lateral accessory lobe in the insect brain. *Front. Physiol.* **7** (2016).
15. Namiki, S., Dickinson, M.H., Wong, A.M., Korff, W. & Card, G.M. The functional organization of descending sensory-motor pathways in *Drosophila*. *eLife* **7**, e34272 (2018).
16. Olberg, R.M. Pheromone-triggered flip-flopping interneurons in the ventral nerve cord of the silkworm moth, *Bombyx mori*. *J. Comp. Physiol.* **152**, 297-307 (1983).
17. Zorović, M. & Hedwig, B. Descending brain neurons in the cricket *Gryllus bimaculatus* (de Geer): auditory responses and impact on walking. *J. Comp. Physiol. [A]*. **199**, 25-34 (2013).
18. Staudacher, E. & Schildberger, K. Gating of sensory responses of descending brain neurones during walking in crickets. *J. Exp. Biol.* **201**, 559 (1998).
19. Bidaye, S.S., Machacek, C., Wu, Y. & Dickson, B.J. Neuronal control of *Drosophila* walking direction. *Science* **344**, 97 (2014).
20. Harley, C.M. & Ritzmann, R.E. Electrolytic lesions within central complex neuropils of the cockroach brain affect negotiation of barriers. *J. Exp. Biol.* **213**, 2851 (2010).
21. Brodfuehrer, P.D. & Hoy, R.R. Ultrasound sensitive neurons in the cricket brain. *J. Comp. Physiol. [A]*. **166**, 651-662 (1990).
22. Zorović, M. & Hedwig, B. Processing of species-specific auditory patterns in the cricket brain by ascending, local, and descending neurons during standing and walking. *J. Neurophysiol.* **105**, 2181-2194 (2011).
23. Chen, C.L., *et al.* Imaging neural activity in the ventral nerve cord of behaving adult *Drosophila*. *Nat. Commun.* **9**, 4390 (2018).

24. DeAngelis, B.D., Zavatone-Veth, J.A. & Clark, D.A. The manifold structure of limb coordination in walking *Drosophila*. *eLife* **8** (2019).
25. Katsov, A.Y., Freifeld, L., Horowitz, M., Kuehn, S. & Clandinin, T.R. Dynamic structure of locomotor behavior in walking fruit flies. *eLife* **6** (2017).
26. Martin, J.P., Guo, P., Mu, L., Harley, C.M. & Ritzmann, R.E. Central-complex control of movement in the freely walking cockroach. *Curr. Biol.* **25**, 2795-2803 (2015).
27. Shiozaki, H.M., Ohta, K. & Kazama, H. A multi-regional network encoding heading and steering maneuvers in *Drosophila*. *Neuron* **106**, 126-141 e125 (2020).
28. Lima, S.Q. & Miesenböck, G. Remote control of behavior through genetically targeted photostimulation of neurons. *Cell* **121**, 141-152 (2005).
29. Borst, A. & Heisenberg, M. Osmotropotaxis in *Drosophila melanogaster*. *J. Comp. Physiol. [A]*. **147**, 479-484 (1982).
30. Gaudry, Q., Hong, E.J., Kain, J., de Bivort, B. & Wilson, R.I. Asymmetric neurotransmitter release at primary afferent synapses enables rapid odor lateralization in *Drosophila*. *Nature* **493**, 424-428 (2013).
31. Zheng, Z., *et al.* A complete electron microscopy volume of the brain of adult *Drosophila melanogaster*. *Cell* **174**, 730-743 (2018).
32. Xu, C.S., *et al.* A connectome of the adult *Drosophila* central brain. *bioRxiv* **2020.01.21.911859** (2020).
33. Wu, M., *et al.* Visual projection neurons in the *Drosophila* lobula link feature detection to distinct behavioral programs. *eLife* **5** (2016).
34. Ribeiro, I.M.A., *et al.* Visual projection neurons mediating directed courtship in *Drosophila*. *Cell* **174**, 607-621 e618 (2018).
35. Aso, Y., *et al.* The neuronal architecture of the mushroom body provides a logic for associative learning. *eLife* **3**, e04577 (2014).
36. Frechter, S., *et al.* Functional and anatomical specificity in a higher olfactory centre. *eLife* **8**, e44590 (2019).
37. Kanzaki, R., Arbas, E.A. & Hildebrand, J.G. Physiology and morphology of protocerebral olfactory neurons in the male moth *Manduca sexta*. *J. Comp. Physiol. [A]*. **168**, 281-298 (1991).
38. Namiki, S., Iwabuchi, S., Pansophia Kono, P. & Kanzaki, R. Information flow through neural circuits for pheromone orientation. *Nature Commun.* **5**, 5919 (2014).
39. Turner-Evans, D., *et al.* Angular velocity integration in a fly heading circuit. *eLife* **6** (2017).
40. Cande, J., *et al.* Optogenetic dissection of descending behavioral control in *Drosophila*. *eLife* **7**, e34275 (2018).
41. Strauss, R. & Heisenberg, M. Coordination of legs during straight walking and turning in *Drosophila melanogaster*. *J Comp Physiol A* **167**, 403-412 (1990).
42. Mendes, C.S., Bartos, I., Akay, T., Marka, S. & Mann, R.S. Quantification of gait parameters in freely walking wild type and sensory deprived *Drosophila melanogaster*. *eLife* **2**, e00231 (2013).
43. Brierley, D.J., Rathore, K., VijayRaghavan, K. & Williams, D.W. Developmental origins and architecture of *Drosophila* leg motoneurons. *J. Comp. Neurol.* **520**, 1629-1649 (2012).
44. Schnell, B., Ros, I.G. & Dickinson, M.H. A descending neuron correlated with the rapid steering maneuvers of flying *Drosophila*. *Curr. Biol.* **27**, 1200-1205 (2017).
45. Ache, J.M., Namiki, S., Lee, A., Branson, K. & Card, G.M. State-dependent decoupling of sensory and motor circuits underlies behavioral flexibility in *Drosophila*. *Nat. Neurosci.* **22**, 1132-1139 (2019).
46. Kien, J. Neuronal activity during spontaneous walking—I. Starting and stopping. *Comp. Biochem. Physiol., Part A Mol. Integr. Physiol.* **95**, 607-621 (1990).
47. Tanji, J. & Evarts, E.V. Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J. Neurophysiol.* **39**, 1062-1068 (1976).
48. Kraskov, A., *et al.* Corticospinal mirror neurons. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **369**, 20130174 (2014).
49. Churchland, M.M., Cunningham, J.P., Kaufman, M.T., Ryu, S.I. & Shenoy, K.V. Cortical preparatory activity: Representation of movement or first cog in a dynamical machine? *Neuron* **68**, 387-400 (2010).

50. Aso, Y., *et al.* Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. *eLife* **3**, e04580 (2014).
51. Wolff, T., Iyer, N.A. & Rubin, G.M. Neuroarchitecture and neuroanatomy of the *Drosophila* central complex: A GAL4-based dissection of protocerebral bridge neurons and circuits. *J. Comp. Neurol.* **523**, 997-1037 (2015).



Extended Data Figure 1: Forward and reverse linear filters for DNa01 and DNa02 neurons

a, Rotational velocity and sideways velocity are strongly correlated (left), whereas translational velocity is not strongly correlated with steering speed (defined as the absolute value of rotational velocity plus sideways velocity, both in units of $^{\circ}/s$). Black lines are LOWESS fits; margins show kernel density estimates (n=20 flies)

b, Firing rate autocorrelation for DNa01 and DNa02.

c, Firing-rate-to-rotational-velocity filters (f \rightarrow r), for DNa01 and DNa02.

d, Firing-rate-to-sideways-velocity filters (f \rightarrow s), for DNa01 and DNa02.

e, Firing-rate-to-translational-velocity filters (f \rightarrow t), for DNa01 and DNa02.

f, Firing-rate-to-total-speed filters (f \rightarrow total speed), for DNa01 and DNa02, where total speed = $|r| + |s| + |t|$. As DNa02 predicts large changes in r and s, it is not surprising that it is also predicts large changes in total speed. Filters in (c-f) can be convolved with neural firing rates to predict behavior, where $b(t) = H(t) * f(t)$. Filters in (c-e) are reproduced from Fig. 1.

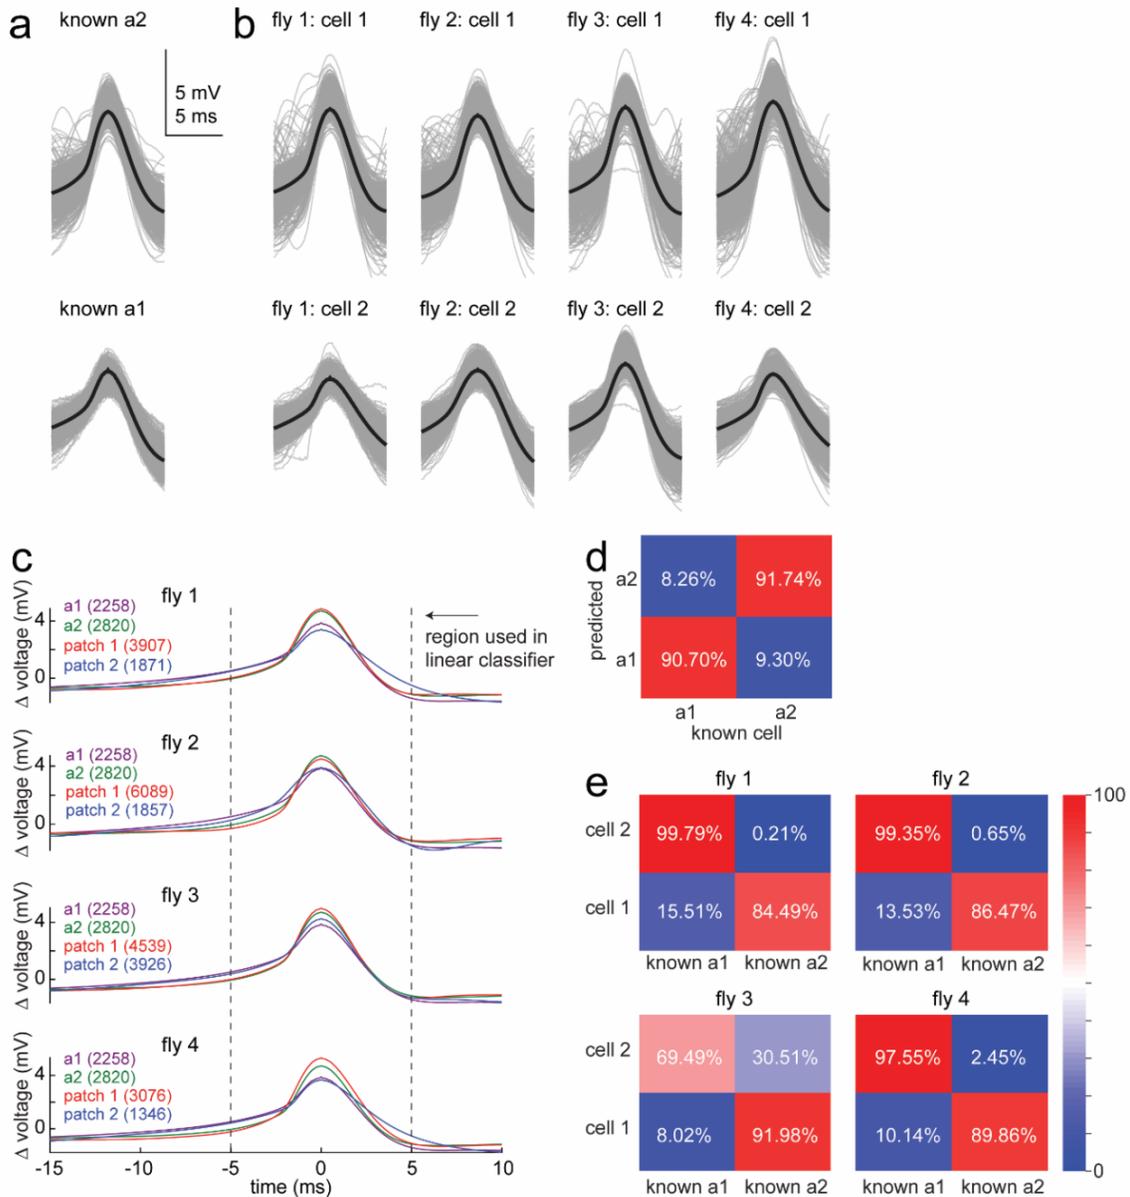
g, Variance explained by each forward filter, for DNa01 and DNa02 neurons. Colors denote variables in (c-f).

h, Behavior autocorrelation for each kinematic variable.

(i-l) Same as in (c-f), but in the reverse direction (r \rightarrow f, s \rightarrow f, etc.). These filters can be convolved with a behavioral variable to predict neural firing rates, where $f(t) = H(t) * b(t)$. Whereas the forward filters in

(c-f) are normalized for the neuron's autocorrelation function, these reverse filters are normalized for behavior autocorrelation.

m, Variance explained by each reverse filter, for each kinematic variable, for DNa01 and DNa02 neurons.



Extended Data Figure 2: Discriminating DNa01 from DNa02 in dual recordings

We combined the split-Gal4 line that targets DNa01 (*SS00731*) with a Gal4 line that targets DNa02 (*R75C10-GAL4*). With a *10XUAS-IVS-mCD8::GFP(attP40)*, this combination of driver lines labeled only two somata in the vicinity of DNa01 and DNa02. Here we show that we can discriminate DNa01 from DNa02 in dual recordings from this genotype.

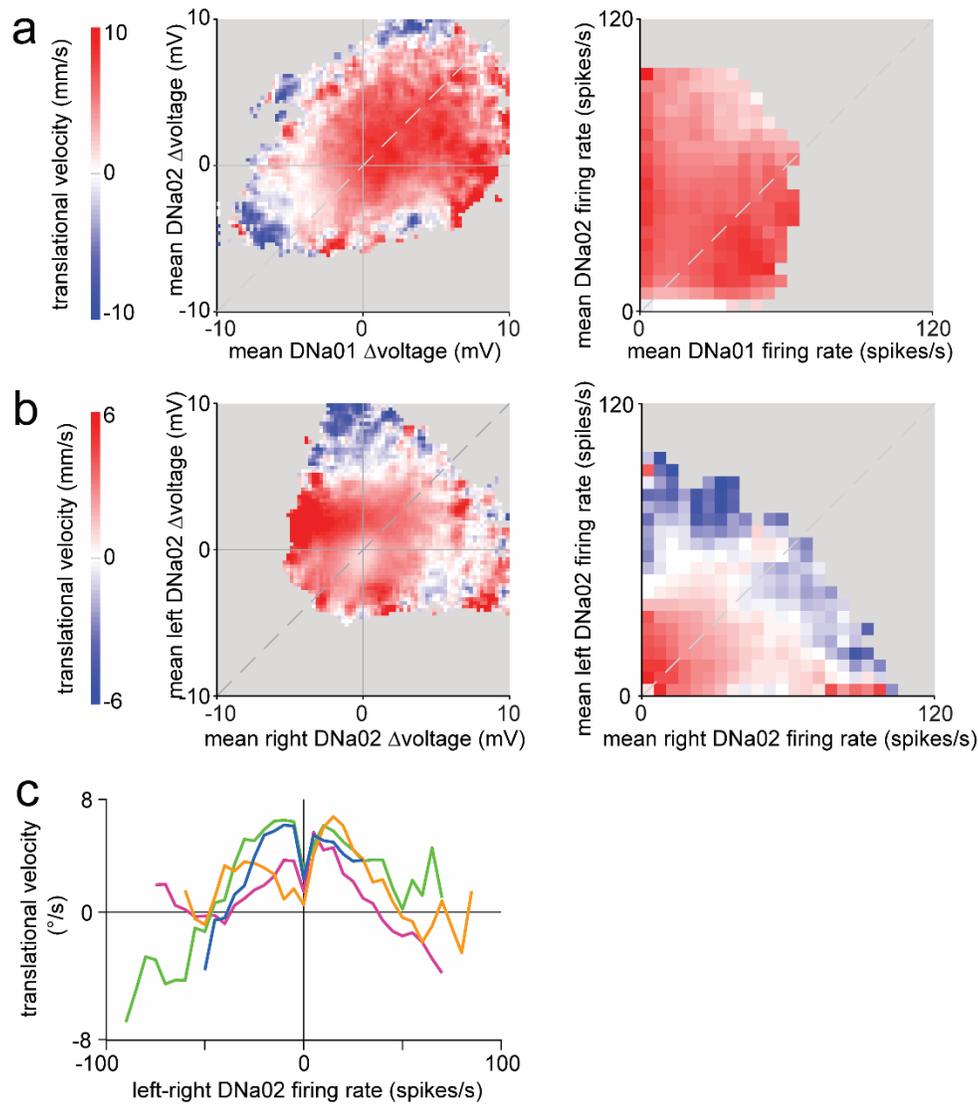
a, Top: spike waveforms from a known DNa02 cell, recorded in the specific split-Gal4 line (*SS00730*). Bottom: spike waveforms from a known DNa01 cell, recorded in a different specific split-Gal4 line (*SS00731*). Individual spikes are gray, averages in black.

b, Spike waveforms from dual DNa01/DNa02 recordings in 4 flies.

c, Magenta and green lines show overlaid averages from (a). Red and blue show overlaid averages from (b). Dashed lines show the part of the waveform used in classification analysis. Values in parentheses show the number of spikes in each average.

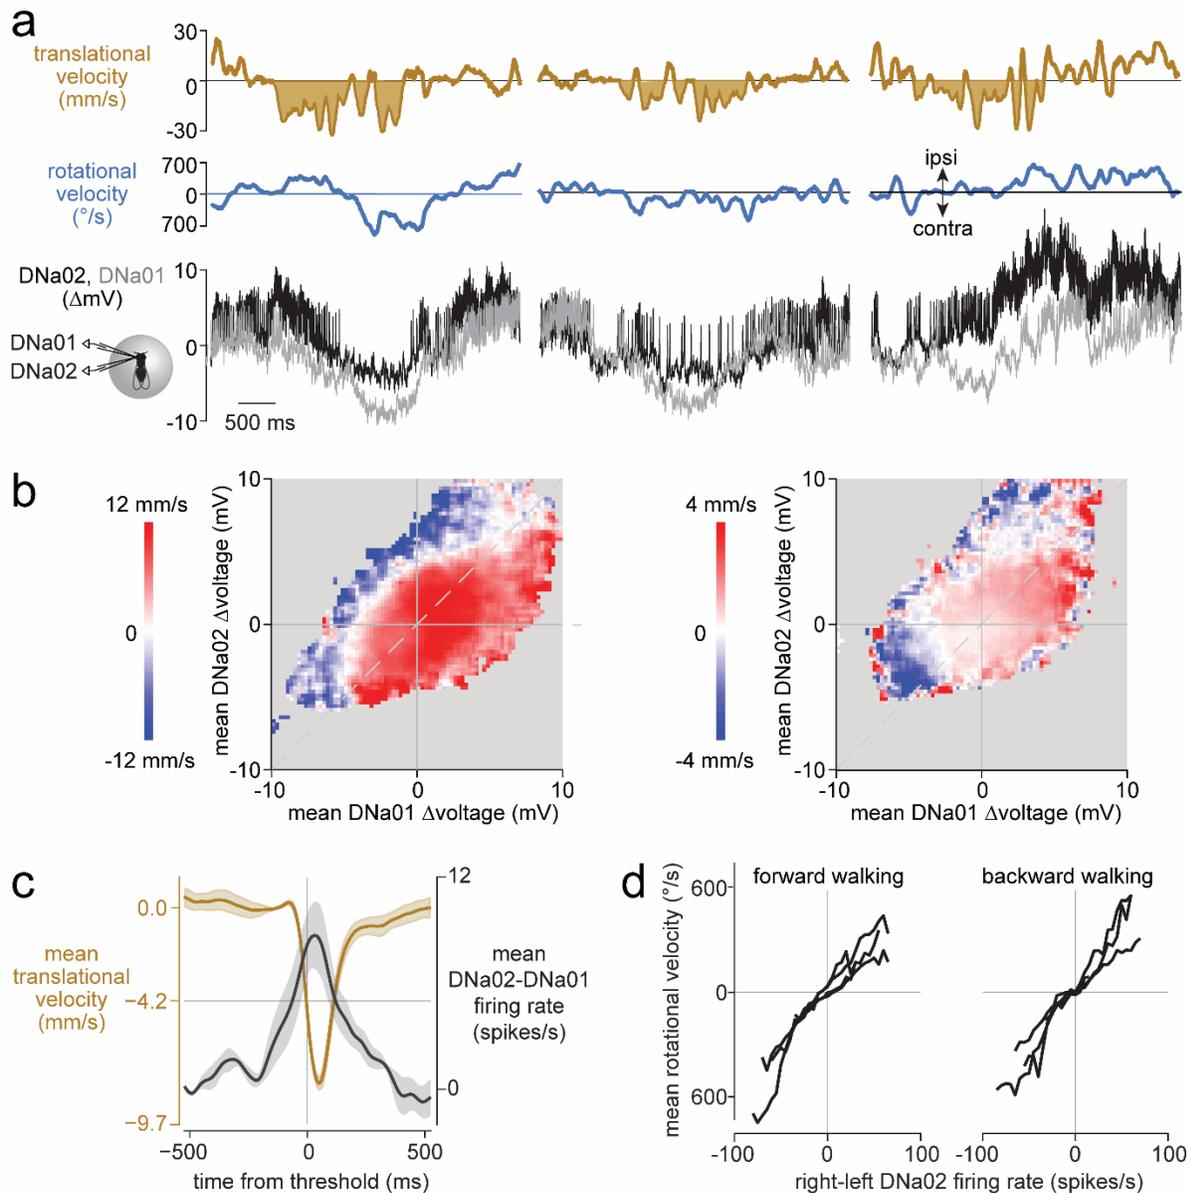
d, A linear classifier was trained to separate cells from known DNa01 and DNa02 recordings, using the data shown in (A). In essence, the classifier separates spike waveforms based on their shapes and amplitudes. This classifier was then tested using another pair of known DNa01 and DNa02 recordings. Values shown are the percentage of spikes assigned to each cell type. Because most of the DNa01 spikes were assigned to the DNa01 category, and most of the DNa02 spikes were assigned to the DNa02 category, we can conclude that this classifier produces the correct identification of the two cells.

e, The same classifier was then applied to dual DNa01/DNa02 recordings from all 4 flies, in order to determine which cell was which. In every case, the classification was essentially unambiguous. Moreover, in every case, the classifier assigned the DNa01 label to the cell with the smaller and deeper soma, consistent with our observation that the soma of DNa01 is somewhat smaller and deeper than the soma of DNa02.



Extended Data Figure 3: Translational velocity and neural activity in paired recordings

- a**, From an example DNa01/DNa02 paired recording, colormaps show binned and averaged translational velocity for each paired value of Δ voltage (left) or firing rate (right). Both DNa01 and DNa02 were recorded in the left hemisphere. When (Δ voltage_{DNa02} \gg Δ voltage_{DNa01}), the fly is typically moving backward. When (firing rate_{DNa02} \gg firing rate_{DNa01}), the fly is also often moving backward, but forward movement is still more common overall, and so the net effect is that translational velocity is small but still positive when (firing rate_{DNa02} \gg firing rate_{DNa01}). Note that when we condition our analysis on behavior rather than neural activity, we do see that backward walking is associated with a large firing rate differential (Extended Data Fig. 7c).
- b**, From an example DNa02/DNa02 paired recording, colormaps show binned and averaged translational velocity for each paired value of Δ voltage (left) or firing rate (right). When both left and right DNa02 cells are firing zero spikes, the fly is typically stopped (see also Fig. 2). When either copy of DNa02 is firing at a high rate, the fly is often moving backward; this is likely because extremely fast turns are often associated with backward movement.
- c**, Mean translational velocity for each left-right DNa02 firing rate difference. Each line is a different fly ($n=4$). Large left-right differences are associated with backward movement; again, this is likely because extremely fast turns are often associated with backward movement.



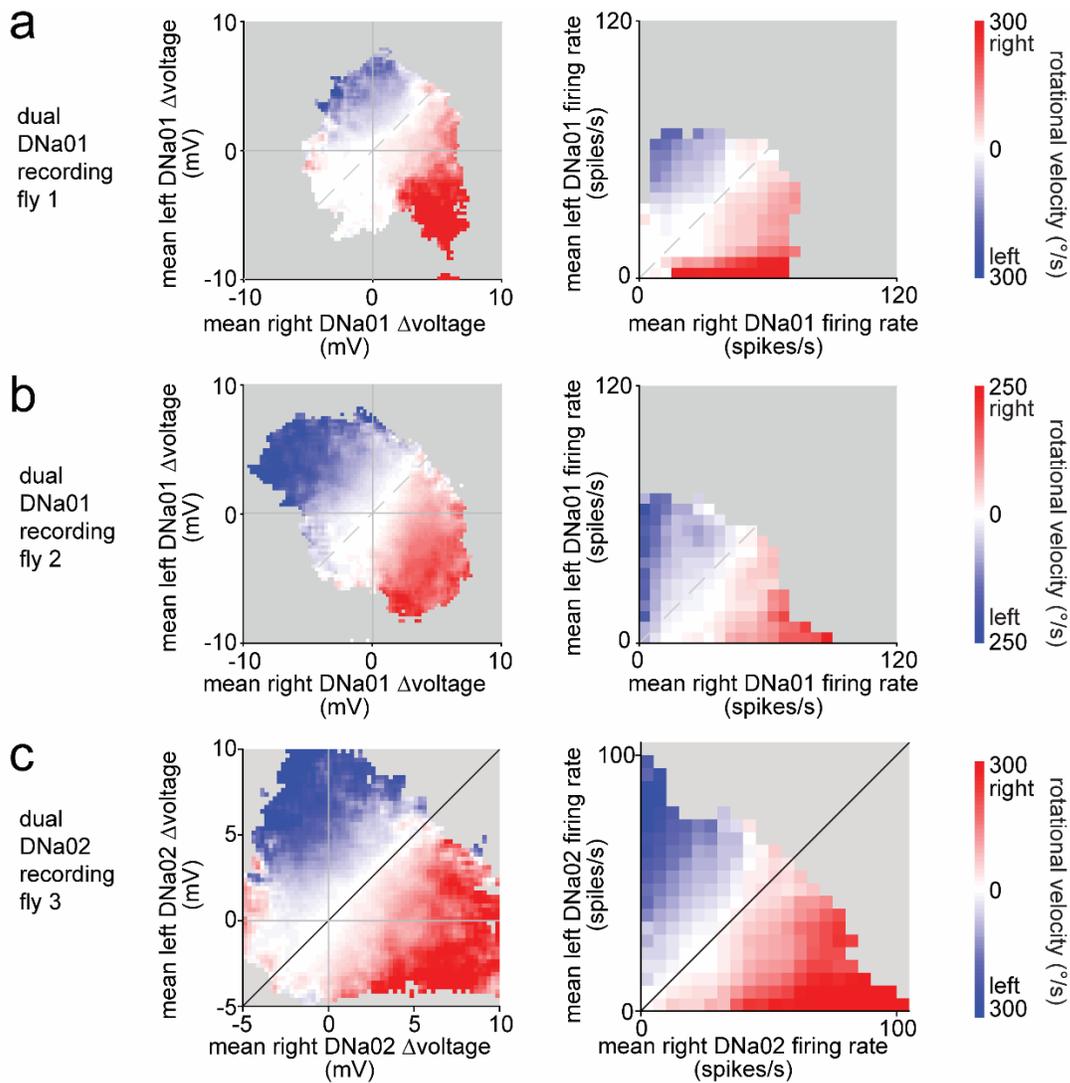
Extended Data Figure 4: Backward walking

a, Three examples of dual recordings from DNa01 and DNa02. When DNa02 is much less hyperpolarized than DNa01 (or much more depolarized), the fly is generally moving backward. Flies typically moved backward $\sim 15\%$ of the time.

b, Colormap shows binned and averaged translational velocity for each value of DNa01 and DNa02 membrane voltage, for three example paired recordings. Voltages are expressed as changes from each cell's mean.

c, Mean difference in firing rates (DNa02-DNa01), \pm SEM across flies ($n=4$ flies), aligned to the onset of backward walking. Overlaid is mean translational velocity \pm SEM across flies; horizontal line represents the threshold used to trigger onset of backward walking.

d, Binned and averaged rotational velocity versus bilateral firing rate difference for epochs of forward walking and backward walking. Each line is a different paired recording ($n=4$ flies).

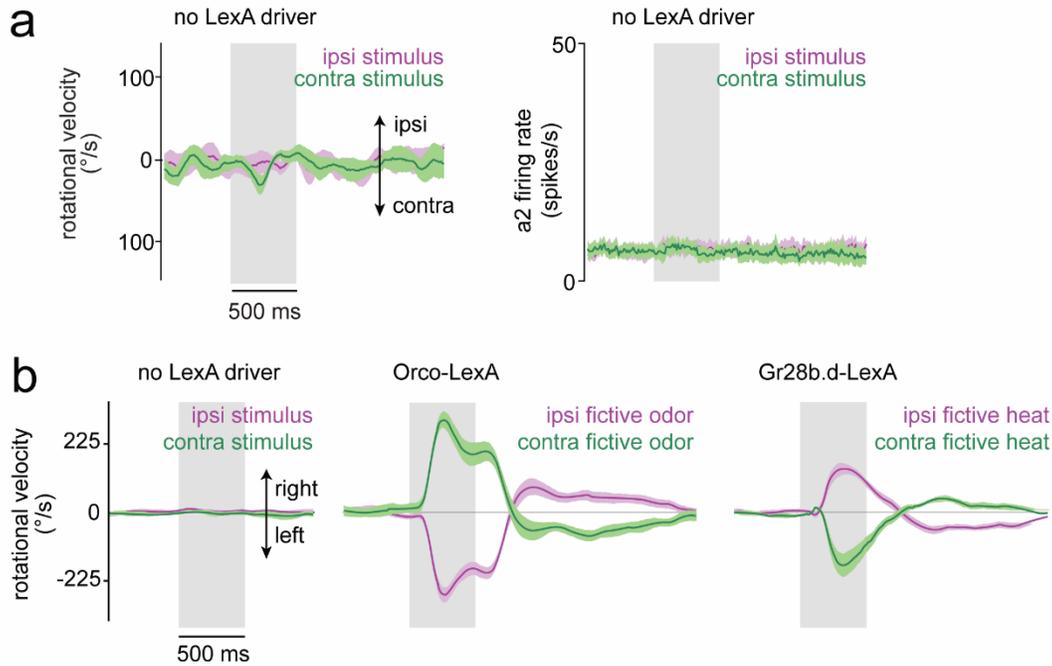


Extended Data Figure 5: Dual recordings from the right and left copies of DNa01.

a, An example dual recording from DNa01 on both sides of the brain. Colormap shows binned and averaged rotational velocity for each value of bilateral DNa01 Δ voltage (left) and bilateral DNa01 firing rates (right). Note that rotational velocity is related to the right-left firing rate difference in this DNa01 paired recording. This is similar to our results in DNa02 paired recordings (Fig. 3).

b, Same but for another example dual DNa01/DNa01 recording.

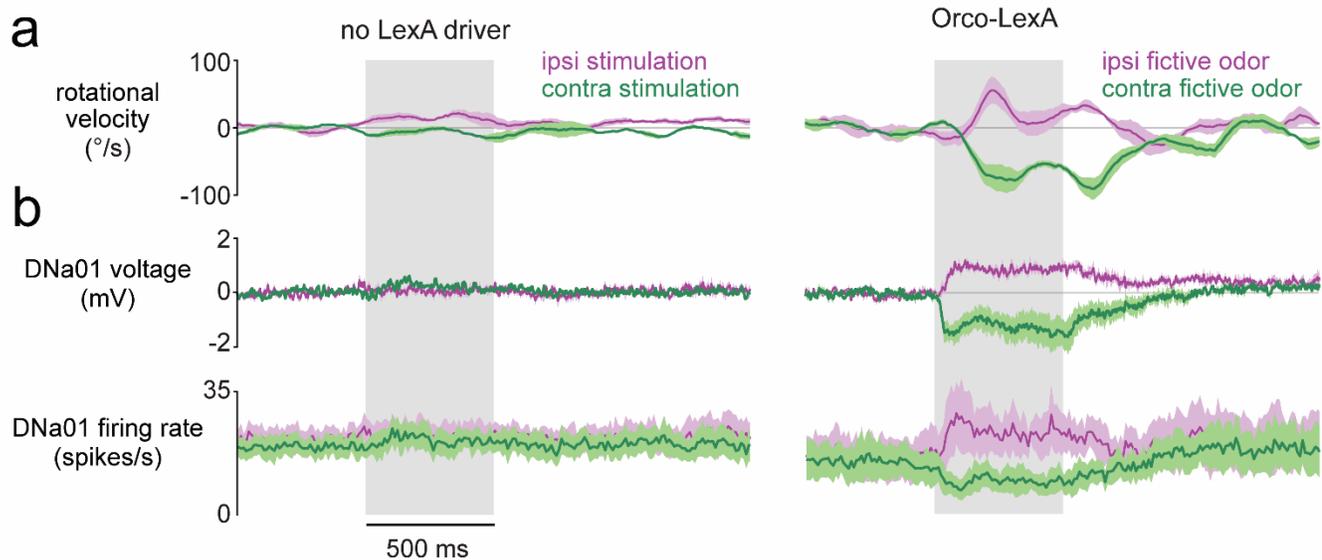
c, An example dual DNa02/DNa02 recording, shown for comparison. The right-hand plot is reproduced from Fig. 3b.



Extended Data Figure 6: Genetic controls for fictive sensory stimuli and behavioral responses in intact flies

a, Rotational velocity (left) and DNa02 firing rate (right) in genetic control flies where no LexA transgene was present (mean \pm SEM across flies, n=4 flies). These flies were treated just like those in Fig. 1, meaning that each antenna was illuminated in the same way. There is essentially no steering behavior or DNa02 response. This result confirms that the steering behavior and DNa02 responses we describe in Fig. 1 are not due to the visual or thermal effects of the fiber optic illumination *per se*. The lack of visual responses is likely related to the fact that the fiber optic filament is very small (50 μ m diameter), and the illumination from the fiber is partially blocked by the antenna, which is positioned very close to the fiber.

b, Behavioral responses in flies where no dissection was performed, i.e. we did not open the head capsule. Left: genetic controls lacking a *LexA* driver (n=6 flies). Middle: fictive odor (*Orco-LexA*, n=9). Right: fictive heat (*Gr28b.d-LexA*, n=9 flies). Plots show mean \pm SEM across flies. Note that behavioral responses are larger and less variable in these intact flies. The reduced behavioral performance of the flies in Fig. 1 is likely due to local removal of the peri-neural sheath and/or local disruption of the neuropil surrounding DNa02 somata during the patching procedure. Nonetheless, the qualitative features of behavior in Fig. 1 are comparable to intact flies.

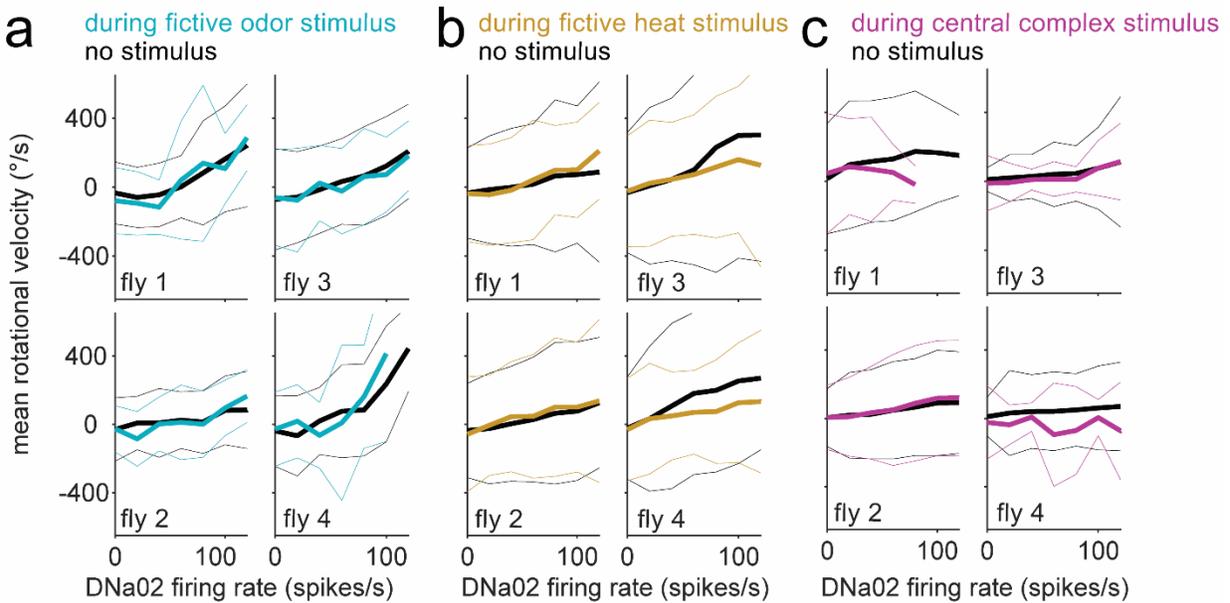


Extended Data Figure 7: Lateralized fictive odor stimuli produce laterally asymmetric responses in DNa01.

a, Rotational velocity during fictive odor presentations.

b, DNa01 activity in the same experiments.

Each plot shows mean \pm SEM across flies (n=4 genetic controls with no LexA driver, n=3 flies harboring *Orco-LexA*). On average, DNa01 activity is weaker than DNa02 activity during odor-evoked turning behavior (compare to Fig. 5), but qualitatively DNa01 and DNa02 have similar relationships to odor-evoked turning.



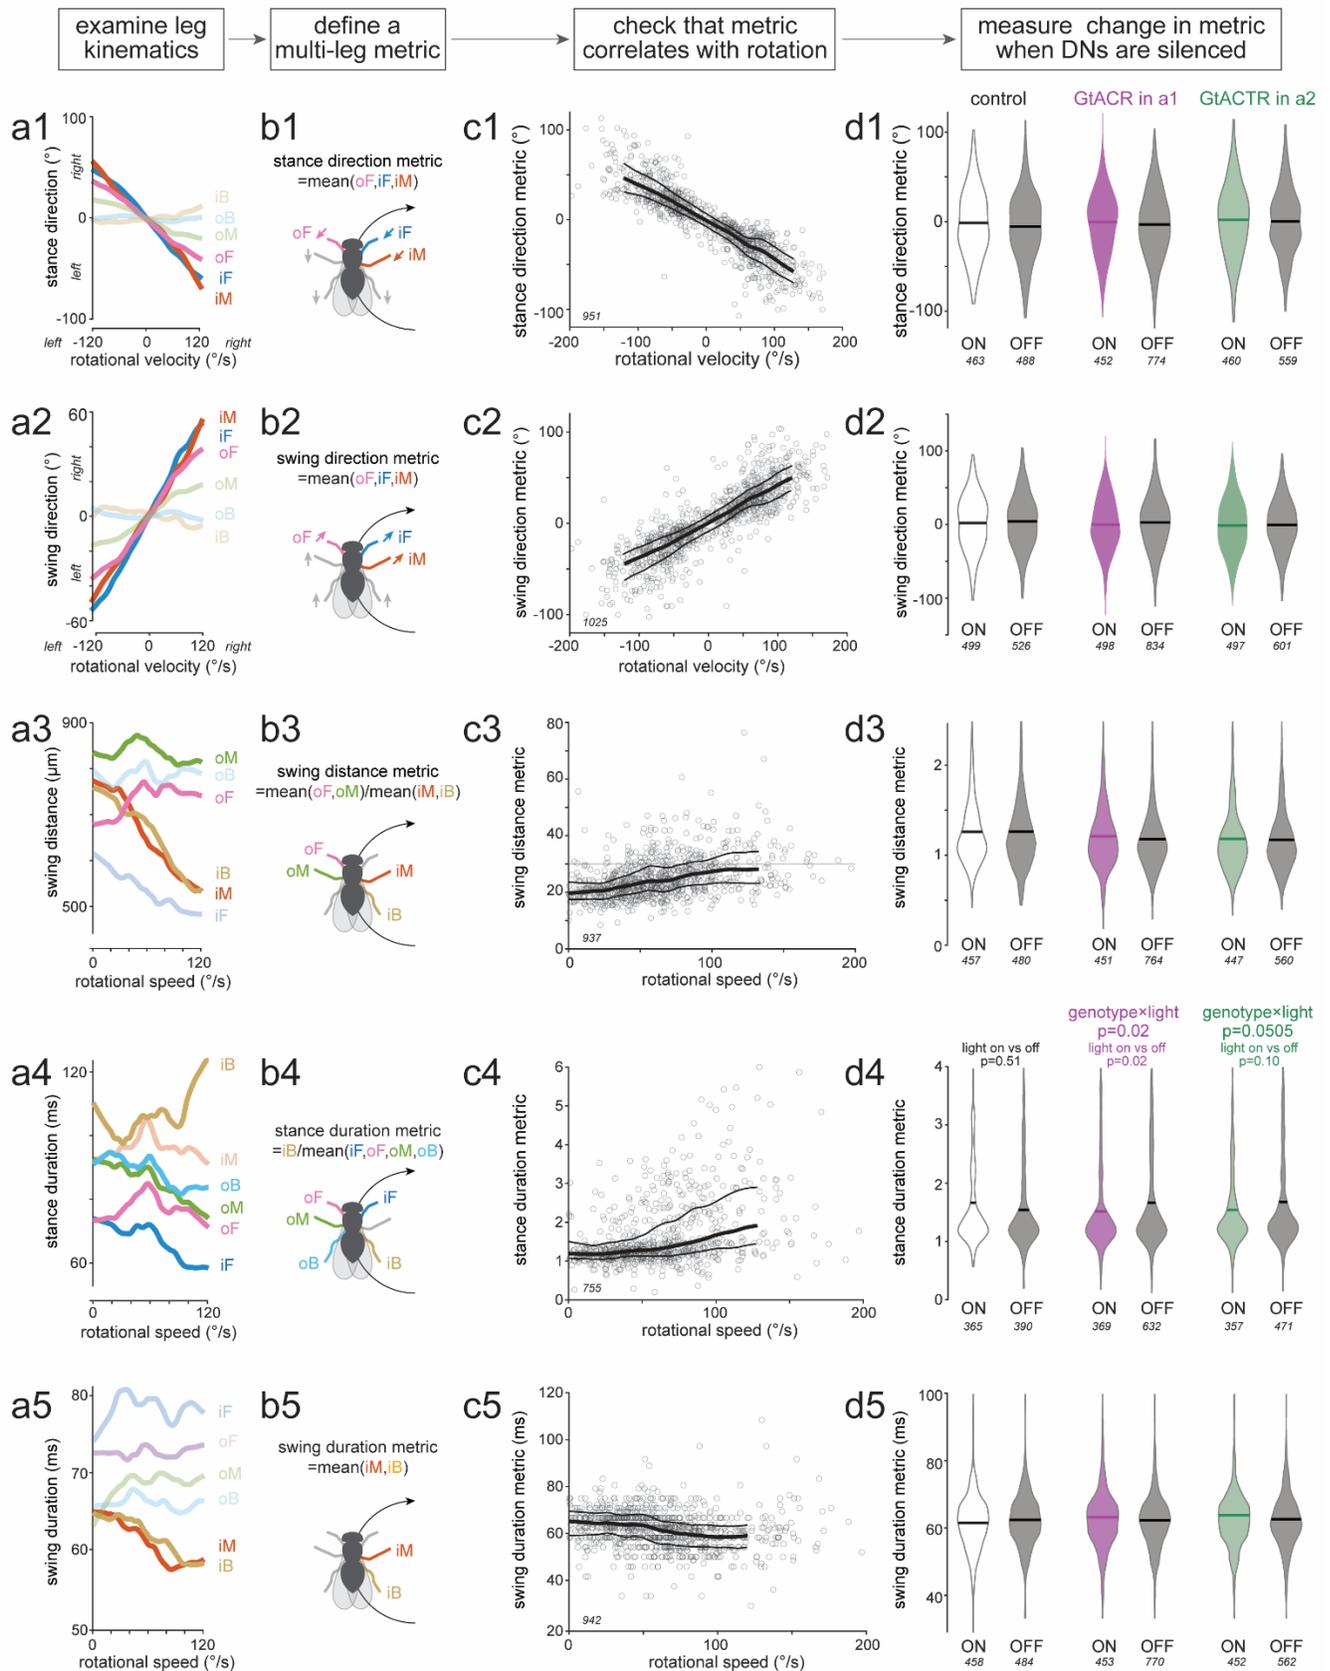
Extended Data Figure 8: Comparing neuron/behavior curves during stimulus-evoked steering and spontaneous steering

a, In each fictive odor experiment (Fig. 5), we measured rotational velocity and DNa02 firing rate in 50-ms time windows. We separated windows when the stimulus was present from windows when the stimulus was absent. Both ipsi- and contralateral stimulus periods were included. We then binned the data by firing rate and averaged rotational velocity values within each bin (thick line). Thin lines show 95% confidence intervals around the mean. Each panel is a different fly.

b, Same but for fictive heat experiments.

c, Same but for central complex stimulation experiments. Here, the “stimulus” epoch is defined as the 750 ms window prior to the maximum bump return speed. A few of these plots show a shallower slope (e.g., flies 2 and 3), possibly because these flies were in a different behavioral state.

Note that the neuron-behavior curve is similar in the two conditions in each experiment. This implies that DNa02 has the same quantitative relationship to steering behavior, regardless of whether steering is evoked by an olfactory stimulus, a thermal stimulus, a compass steering drive, or a spontaneous steering drive.

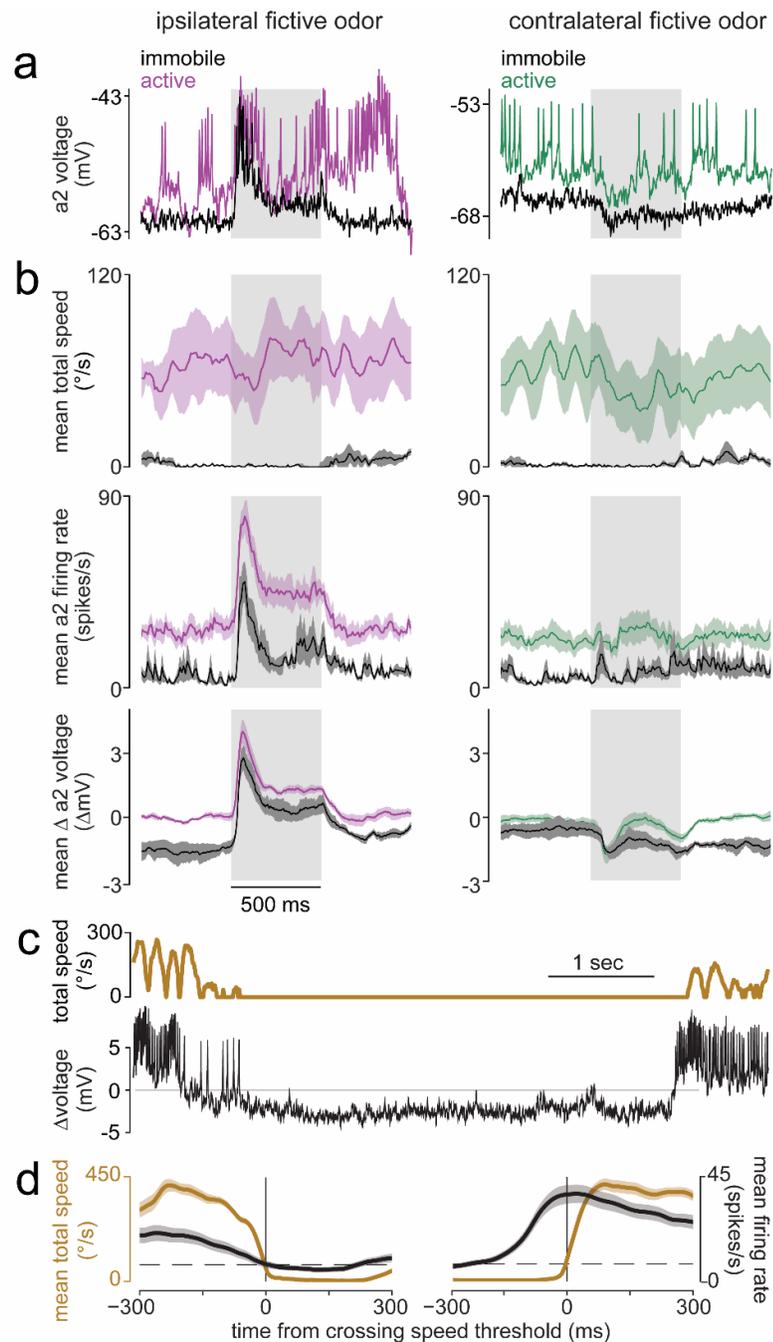


Extended Data Figure 9: Leg movements associated with body rotation

a, We selected 5 leg-kinematic variables likely to be associated with body rotation. In control flies, we computed the mean value of each variable, for each leg, binned by rotational velocity (for directional variables) or rotational speed (for nondirectional variables).

- b**, We defined 5 multi-leg metrics, each based on one leg-kinematic variable. For each metric, the relevant legs are labeled and color-coded. Gray legs were not used because, for these legs, the variable in question did not show a reliable correlation with the fly's rotational movement (a). Gray legs in (b) are shown in faded colors in (a).
- c**, Next, we plotted each metric versus rotational velocity or rotational speed, to verify that the expected relationship was observed. We refined our definitions in order to optimize the signal-to-noise ratio of this relationship. Each datapoint is a time window (with *n* values in italics near the origin of each plot). Thick lines show LOWESS fits; thin lines show 95% confidence intervals; only a small number of datapoints have rotational speeds $>120^\circ/\text{s}$, so fits are not shown above this ceiling. Importantly, we fixed the number of metrics we tested, and also the definitions of those metrics, based purely on data from control flies, before we examined any of these metrics in our experimental genotypes.
- d**, Finally, we asked how these metrics change when DNa01 or DNa02 neurons are silenced. We identified all body rotation events in every experiment by searching forward in time for moments when rotational speed crosses a threshold of $20^\circ/\text{s}$. We extracted a 500-ms time window around each threshold-crossing. Events were constrained to be non-overlapping. Body rotation events were tabulated separately for light-on and light-off. The number of time windows in each distribution is shown in italics. To determine if any metric was changed when DNa01 neurons were silenced, we performed 5 two-factor ANOVAs and looked for a significant interaction between genotype (control/DNa01) and light (on/off); these p-values were Bonferroni-Holm corrected to avoid type 1 errors ($m=5$ tests). Then, to determine if any metric was changed when DNa02 neurons were silenced, we performed 5 two-factor ANOVAs and looked for a significant interaction between genotype (control/DNa02) and light (on/off); these p-values were again Bonferroni-Holm corrected ($m=5$ tests); nonsignificant p-values are not shown. Then, in cases where these p-values were significant (**d4**, $p=0.02$, DNa01 vs control) or nearly significant (**d4**, $p=0.505$, DNa02 vs control), we followed up by performing a post hoc Tukey test to determine if there was a significant difference between light on/off in each genotype. Interestingly, the only metric to show an effect was the stance duration metric (**d4**). In control flies, a prolonged stance duration is observed in some but not all body rotation events (**c4**). This implies that body rotation can be driven by an inner-back-leg pivoting maneuver, but body rotation does not absolutely require this maneuver. This idea is consistent with our finding that rotational speeds are overall unchanged when DNa01 or DNa02 is silenced (Fig. 8e).

Note that the data for the stance duration metric (**a4-d4**) are reproduced in Fig. 8f,g, where this metric is called “inner back leg stance duration (normalized)”, i.e. iB pivoting.



Extended Data Figure 10: Neural activity during immobility.

a, Examples of DNa02 activity during fictive odor presentation when the fly was active versus immobile. Note that, when the fly is immobile, ipsilateral odor still evokes depolarization and spiking. (This phenomenon is also visible in Fig. 5d, where there are multiple trials where the fly is immobile and thus rotational velocity is exactly zero, but DNa02 firing rate is somewhat high.) Note also that contralateral odor can still evoke hyperpolarization during immobility. More examples of DNa01/DNa02 hyperpolarization during immobility can be seen in Figs. 2a and 3a.

b, Summary data (mean \pm SEM across flies, $n=4$ flies) for active versus immobile trials. Odor produces a similar *change* in neural activity for trials where the fly is behaviorally active versus immobile. However, the entire dynamic range of neural activity is shifted downward during immobility, so that the peak firing rate (and peak depolarization) is lower than what we observe during behaviorally active trials. We defined immobility as essentially zero total speed (i.e., the sum of speed values in all three axes of ball rotation) during the 500 msec prior to stimulus onset and the entire period of stimulus presentation.

c, Example of a fly transitioning from activity to immobility, and then back to activity again.

d, Mean total speed and DNa02 firing rate during (left) the transition to immobility and (right) the transition to activity (\pm SEM across flies, $n=7$ flies). Transitions were detected by setting a threshold for total speed (dashed line). The firing rate change precedes the onset of movement by ~ 250 ms. Total speed is defined as rotational speed + sideways speed + translational speed.

Supplementary Table 1: Neurons presynaptic to DNa02 reconstructed from the full-brain EM volume.

Neuron Name	Synapses per Connection	Neuron from Figure 4
PFL3_1	15	4d (1)
PFL3_2	20	4d (2)
PFL3_3	10	4d (3)
PFL3_4	17	4d (4)
PFL3_5	20	4d (5)
PFL3_6	19	4d (6)
PFL3_7	16	4d (7)
PFL3_8	17	4d (8)
PFL3_9	9	4d (9)
PFL3_10	9	4d (10)
PFL3_11	27	4d (11)
PFL3_12	20	4d (12)
AOTUtoL_1	17	4e, green
AOTUtoL_2	37	4e, blue
AOTUtoL_3	10	4e, orange
AOTUtoL_4	4	4e, magenta
LC9_1	4	4f, green
LC9_2	5	4f, orange
SMPtoL1_1	6	4g, green
SMPtoL1_2	3	4g, blue
SMPtoL2_1	58	4g, orange
SMPtoL2_2	19	4g, magenta
AVLPtoL1_1	8	4h, green
AVLPtoL1_2	20	4h, orange
LtoL1_1	32	4i, green
LtoL1_2	41	4i, orange
DN_1	172	4j, green
DN_2	124	4j, orange

ONLINE METHODS

Fly husbandry and genotypes

We used female flies 1-2 days post-eclosion, except where specified below. We used virgins to minimize egg-laying, which tended to stop movement of the spherical treadmill. Flies were housed in a 25° incubator at 50-70% humidity. Unless otherwise noted, flies were cultured on molasses food (B7 recipe, Archon Scientific) and kept on a 12/12 light/dark cycle. Flies for the experiments in **Fig. 5 and 8** (and **Extended Data Figs. 6, 7, 8a-b, 9, 10a,b**) were cultured on molasses food and then transferred 0-1 days after eclosion onto a vial of standard cornmeal agar fly food supplemented with rehydrated potato flakes (Formula 4-24, Carolina Biological Supply) and 100 µL of all-*trans*-retinal (35 mM in ethanol; R2500, Sigma-Aldrich). These vials were covered with aluminum foil to prevent photoconversion of all-*trans*-retinal by ambient light.

The split-Gal4 lines⁵² targeting DNs¹⁵ were *SS00730* for DNa02 (*R75C10-p65ADZp* in *attP40*; *R87D07-ZpGdbd* in *attP2*) and *SS00731* for DNa01 (*R22C05-p65ADZp* in *attP40*; *R56G08-ZpGdbd* in *attP2*); these were obtained from Gwyneth Card (Janelia Research Campus). Each of these split-Gal4 lines typically drives expression in one cell per brain hemisphere, although *SS00731* occasionally drives expression in two cells per hemisphere. *R75C10-Gal4(attP2)*^{53,54} and *R60D05-LexA(attP40)*⁵² were obtained from the Bloomington Drosophila Stock Center (BDSC). *VT032906-Gal4(attP2)*⁵⁵ was obtained from the Vienna Drosophila Resource Center. *Orco-LexA* (on chromosome 3)⁵⁵ was obtained from Tzumin Lee (Janelia Research Campus). *Gr28b.d-LexA* (on chromosome 2) was obtained from Marco Gallio (Northwestern University)⁵⁶. *hsFLP.122*⁵⁷⁻⁵⁹ was obtained from Thomas Clandinin (Stanford University). *UAS-GtACR1::eYFP(VK00005)* was obtained from Michael Crickmore (Boston Children's Hospital); this reagent has not been published previously. *pJFRC81-td3-Halo7::CAAX(attP18)* was obtained from Greg Jefferis via Luke Lavis; this specific reagent has not been published previously, although general methods its construction have been published⁶⁰. *UAS-P2X2* (chromosome 3)²⁸ was obtained from Gero Miesenböck (Oxford University). Finally, the following stocks were obtained from the BDSC: *10XUAS-IVS-mCD8::GFP(su(Hw)attP8)*⁵², *10XUAS-IVS-mCD8::GFP(attP40)*⁵², *UAS(FRT.mCherry)ReaChR(VK00005)*⁶¹, *p65.AD.Uw(attP40)*; *GAL4.DBD.Uw(attP2)*⁶², *13XLexAop2-IVS-GCaMP6f-p10(su(Hw)attP5)*⁶³, *13XLexAop2-IVS-CsChrimson::mVenus(attP40)*⁶⁴, *13XLexAop2-IVS-CsChrimson::mVenus(attP18)*⁶⁴, *13XLexAop2-IVS-CsChrimson::mVenus(attP2)*⁶⁴.

Genotypes of fly stocks used in each figure are as follows:

DNa02 single/dual recording (Figs. 1 and 3; Extended Data Figs. 1, 3b-c, 4d, 5c, 10c-d)

y¹, w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R75C10-p65ADZp (attP40) / +; R87D07-ZpGdbd (attP2) / +

and

w¹¹¹⁸, 13XLexAop2-IVS-CsChrimson::mVenus(attP18) / w; 10XUAS-IVS-mCD8::GFP(attP40) / R75C10-P65ADZP(attP40); Orco-LexA / R87D07-ZpGdbd(attP2)

DNa01 single/dual recording (Fig. 1; Extended Data Figs. 1, 5a-b, 7)

y¹, w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R22C05-p65ADZp(attP40) / +; R56G08-ZpGdbd(attP2) / +

DNa01/DNa02 dual recording (Fig. 2; Extended Data Figs. 2, 3a, 4a-c)

y¹, w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R22C05-p65ADZp(attP40) / 10XUAS-IVS-mCD8::GFP(attP40); R56G08-ZpGdbd(attP2) / R75C10-GAL4(attP2)

fictive odor with DNa02 recording (Fig. 5a,c,d; Extended Data Figs. 6b, 8a, 10a-b)

w¹¹¹⁸, 13XLexAop2-IVS-CsChrimson::mVenus(attP18) / w; 10XUAS-IVS-mCD8::GFP(attP40) / R75C10-

P65ADZP(attP40); Orco-LexA / R87D07-ZpGdbd(attP2)

fictive heat with DNa02 recording (Fig. 5b; Extended Data Figs. 6b, 8b)

y^1 , w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R75C10-P65ADZP(attP40) / Gr28b.d-LexA; R87D07-ZpGdbd(attP2) / 13XLexAop2-IVS-CsChrimson::mVenus(attP2)

fictive odor with DNa01 recording (Extended Data Fig. 7)

y^1 , w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R22C05-p65ADZp(attP40) / Orco-LexA; R56G08-ZpGdbd(attP2) / 13XLexAop2-IVS-CsChrimson::mVenus(attP2)

no-LexA control for fictive odor and fictive heat, DNa02 recordings (Extended Data Fig. 6)

y^1 , w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R75C10-P65ADZP(attP40) / +; R87D07-ZpGdbd(attP2) / 13XLexAop2-IVS-CsChrimson::mVenus(attP2)

no-LexA control for fictive odor and fictive heat, DNa01 recordings (Extended Data Fig. 7)

y^1 , w, 10XUAS-IVS-mCD8::GFP(su(Hw)attP8) / w; R22C05-p65ADZp(attP40) / +; R56G08-ZpGdbd(attP2) / 13XLexAop2-IVS-CsChrimson::mVenus(attP2)

calcium imaging and DNa02 recording during central complex stimulation (Fig. 4; Extended Data Fig. 8c)

w, pJFRC81-td3-Halo7::CAAX(attP18) / w; R75C10-P65ADZP(attP40) / R60D05-LexA(attP40), 13XLexAop2-IVS-GCaMP6f-p10(su(Hw)attP5); R87D07-ZpGdbd(attP2) / UAS-P2X2, VT032906-Gal4(attP2)

central complex-evoked turning, genetic control lacking P-EN1 Gal4 (Fig. 4b)

w, pJFRC81-td3-Halo7::CAAX(attP18) / w; R75C10-P65ADZP(attP40) / R60D05-LexA(attP40), 13XLexAop2-IVS-GCaMP6f-p10(su(Hw)attP5); R87D07-ZpGdbd(attP2) / UAS-P2X2

Unilateral optogenetic activation of DNa02 (Fig. 8a-c)

y, w, hsFLP.122 / w; R75C10-P65ADZP(attP40) / +; R87D07-ZpGdbd(attP2) / UAS-FRT.mCherry.FRT.ReachR(VK00005)

Bilateral optogenetic silencing of DNa01 (Fig. 8d-h, Extended Data Fig. 9)

w; R22C05-p65ADZp(attP40) / +; R56G08-ZpGdbd(attP2) / UAS-GtACR1::eYFP(VK00005)

Bilateral optogenetic silencing of DNa02 (Fig. 8d-h, Extended Data Fig. 9)

w; R75C10-P65ADZP(attP40) / +; R87D07-ZpGdbd(attP2) / UAS-GtACR1::eYFP(VK00005)

Genetic control for optogenetic silencing (“empty” split-Gal4) (Fig. 8d-h, Extended Data Fig. 9)

w; p65.AD.Uw(attP40) / +; GAL4.DBD.Uw(attP2) / UAS-GtACR1::eYFP(VK00005)

Fly mounting and dissection for electrophysiology and calcium imaging

On the day of the experiment, a fly was cold-anesthetized and then inserted into a hole in a stainless-steel, machine milled platform (Harvard Medical School Research Instrumentation Core). We secured the fly to the platform using UV-curable glue (AA3972 Loctite) applied on the thorax and eyes.

We then extended the proboscis with forceps and waxed it in place to prevent brain motion (Electra Waxer, Almore); prior to this step, a manual manipulator (MX160R Siskiyou) was used to lower a shield over the legs to protect them from the tip of the waxer. We then covered the brain with extracellular saline composed of: 103 mM NaCl, 3 mM KCl, 5 mM TES, 8 mM trehalose, 10 mM glucose, 26 mM NaHCO₃, 1 mM NaH₂PO₄, 1.5 mM CaCl₂, and 4 mM MgCl₂ (osmolarity 270-275 mOsm). The saline was bubbled during the dissection with

95% O₂ and 5% CO₂ and reached an equilibrium pH of 7.3. The top of the cuticle was removed, followed by air sacs and fat globules around the patching site. Muscle 16 was severed to prevent brain motion. The frontal head air sac, which extends ventrally under the brain, was also removed to improve illumination from below the brain. Finally, we mechanically removed the peri-neural sheath over the patching site. (For behavior-only experiments in **Fig. 8d-h** and **Extended Data Fig. 6b**, all these dissection steps were omitted.)

Finally, we transferred the fly to the spherical treadmill, where we used a motorized three-axis manipulator (MT3-Z8, Thorlabs) to adjust the fly relative to the ball. Two cameras (BFLY-PGE-13E4M, FLIR) were used to visualize the fly during this adjustment.

Spherical treadmill and behavioral measurements

We used a sensor-based setup for tracking the movements of the spherical treadmill⁶⁵. Briefly, a hollow HDPE sphere (6.35 mm diameter) was floated on pressurized air, in a plenum mounted on a manipulator to allow alignment of the center of the ball with two motion sensors located at a distance from the ball. The ball was painted with a random dot pattern to make it easier for motion sensors to detect movement. We positioned two cameras 90° apart to track the motion of the ball in three dimensions (rotational, sideways, translational velocity), and we took the fly's fictive movements in each direction as equal and opposite to the ball's movements. These motion cameras were custom built by housing the motion sensor array (ADNS-9800, JACK Enterprises) in a custom housing with a lens (M2514-MP2, Computar). Data from the sensor was read out and translated into an analog signal using a digital-to-analog converter (MCP4725, SparkFun) and an Arduino device (Due, Arduino). Analog data were digitized at 10 kHz with a multi-channel NI DAQ board (PCIe-6353, National Instruments), controlled via the Matlab Data Acquisition Toolbox interface (MathWorks). Analog data were then smoothed using a 2nd-order lowpass Butterworth filter and downsampled to 100 Hz.

Motion sensor outputs were calibrated to obtain the conversion from sensor outputs to ball velocities. We performed this calibration by rotating a ball using a planetary gear motor with an encoder (Actobotics, Part 3638298). The motor was rotated at 6 different speeds in all 3 axes of motion, and linear regression was used to calculate the mapping between the output of the motion sensors and the rotation of the ball.

We noticed that steering responses to sensory stimuli were larger and more consistent in intact flies, as compared to flies used for electrophysiology experiments (**Extended Data Fig. 6b**). The reduced performance of flies used for electrophysiology experiments is likely due to local removal of the peri-neural sheath and/or local disruption of the neuropil surrounding the targeted somata. Nonetheless, the qualitative features of behavior are comparable in intact flies (**Extended Data Fig. 6b**) versus flies used for electrophysiology (**Fig. 5**).

Electrophysiology

For patch clamp experiments, we illuminated the brain with an infrared fiber (SXF30-850, Smart Vision Lights) mounted under the fly tethering platform in a sagittal plane at ~20° from the horizontal and we visualized the brain with a camera (BFLY-PGE-13E4M, Point Gray) mounted on a fluorescence microscope (BX51, Olympus).

Before attempting to obtain seals, we first used a large-diameter glass pipette filled with extracellular saline (4-5 MΩ), to clear the area around the soma of interest by applying positive pressure and also gently sucking away any cells lying on top of the soma of interest. Patch pipettes (9-11 MΩ) were pulled the day of the recording and filled with internal solution containing 140 mM KOH, 140 mM aspartic acid, 10 mM HEPES, 1 mM EGTA, 1 mM KCl, 4 mM MgATP, 0.5 mM Na₃GTP, and 13 mM biocytin hydrazide (pH adjusted to 7.2 ± 0.1, osmolarity adjusted to 265 ± 3 mOsm). Whole-cell patch-clamp recordings were performed from fluorescently

labeled somata in current-clamp mode using an Axopatch 200B amplifier. Data were low-pass filtered at 5 kHz, digitized at 10 kHz by a 16 bit A/D converter (National Instruments, BNC 2090-A), and acquired using the Matlab Data Acquisition Toolbox (MathWorks). Recordings were stable for at least 15 min and up to 2 hrs.

We always recorded from the neuron in the left hemisphere, except when we recorded from the same cell type bilaterally. For bilateral recordings, we performed the cleaning and air-sac clearing steps on both sides of the brain. For dual recording experiments of DNa01 and DNa02, we performed recordings on the left side of the brain.

We targeted DNa01 using *SS00731*, and we targeted DNa02 using *SS00730*¹⁵; the only exception was that, in dual DNa01 and DNa02 recordings, we targeted DNa01 using the *SS00731*, and we targeted DNa02 using *R75C10-Gal4*. In these flies, we could identify DNa01 and DNa02 based on the depths of their somata (DNa02 is more superficial), the sizes of their somata (DNa02 is larger), and their spike waveforms (DNa02 spikes are larger). Finally, to confirm the discriminability of DNa02 and DNa01 spikes, we quantified the performance of automated spike waveform classification (**Extended Data Fig. 2**).

Calcium imaging and electrophysiology during central complex stimulation

In the experiments for **Fig. 4** (and **Extended Data Fig. 8c**), the pitch angle of fly's head was carefully adjusted during the dissection procedure. This allowed patch-pipette access to the anterior side of the brain (where DNa02 somata reside) as well as pressure-pipette access to the protocerebral bridge on the posterior side of the brain (where P-EN1 dendrites reside).

To position the pressure-ejection pipette in the protocerebral bridge, and to obtain a patch recording from an DNa02 somata, we visualized the brain and the patch pipette using an infrared fiber (SXF30-850, Smart Vision Lights) mounted under the tethering platform in a sagittal plane at ~20° from the horizontal and a near-infrared camera (GS-U3-41C6NIR FLIR) mounted on the widefield viewing port of a 2-photon microscope (Bergamo II, Thorlabs).

To obtain a recording from DNa02, we needed to fluorescently label DNa02 somata. However, we needed to label DNa02 without also labeling P-EN1 neurons (which also expressed Gal4), because P-EN1 labeling could have contaminated our GCaMP6f signal in the ellipsoid body. We therefore used a chemogenetic approach⁶⁰. Specifically, we drove expression of HaloTag under *Gal4/UAS* control, and we then applied SiR110-HaloTag dye⁶⁶ (a gift from Luke Lavis) to the preparation. The dye was prepared as a stock solution (500 μM in DMSO), and then 1 μL of the stock was dissolved in 500 μL of extracellular saline. Most of the saline in the recording chamber was removed and the dye solution was added to the chamber. After a 15 min incubation, the chamber was rinsed and re-filled with regular extracellular saline. We used a TexasRed filter cube (49017, Chroma) to visualize labeled DNa02 somata. Because the perineural sheath was removed selectively in the vicinity of DNa02 somata, the bath-applied SiR110-HaloTag dye bound mainly to DNa02 somata, with essentially no binding to the ellipsoid body neuropil.

For GCaMP6f imaging of E-PG neuron dendrites in the ellipsoid body, we used a volumetric, galvo-resonant scanning 2-photon microscope (Bergamo II, Thorlabs) equipped with a 20×, 1.0 n.a. objective (XLUMPLFLN20XW, Olympus) and GaAsP detectors (Hamamatsu). We used ScanImage 2018 software (Vidrio Technologies) to control the microscope. Two-photon excitation was provided by a Chameleon Ultra II Ti:Sapphire femtosecond pulsed laser with pre-compensation (Vision-S, Coherent). To image GCaMP6f, we tuned the laser to 940 nm. The objective was translated vertically using a scanner with a 400-μm travel range (P-725KHDS PIFOC, Physik Instrumente). The ellipsoid body was volumetrically imaged with 12 z-planes of 256×64 pixels separated by 4-5 μm. We acquired ~12 volumes/s.

For P-EN1 stimulation, we followed a published procedure¹³. We prepared a solution of ATP (A7699, Sigma, 0.5 mM) and Alexa594 dye (A33082, FisherScientific, 20 μ M,) in extracellular saline. This solution was used to fill a glass pipette slightly smaller than a patch pipette. We used a red emission filter cube (49017, Chroma) to visually locate the pipette via the microscope eye pieces. The tip of the pipette was placed into the protocerebral bridge, where the dendrites of P-EN1 neurons are located, generally close to glomerulus 4⁵¹. To eject ATP and dye, we delivered a 50-100 ms pulse of pressure to the back of the pipette (10-20 p.s.i.) using a pneumatic device gated by a voltage pulse (PV820; World Precision Instruments). We confirmed ejection by observing a bolus of dye appear in the center of the protocerebral bridge neuropil, and the resulting stimulus-locked rotation of the E-PG bump in the ellipsoid body.

In a given experiment, we noticed that the bump tended to jump to the same location in every trial. We used that fact to try to obtain large bump jumps. Specifically, we manually timed the ATP injection in each trial to maximize the size of the resulting bump jump. We ejected ATP every 15-25 s for a total of 87 ± 10 stimulations per fly.

As a negative control, we performed the same type of experiment in flies lacking the P-EN1 Gal4 line, and we confirmed that the ATP puff only rarely preceded bump jumps in this genotype (**Fig. 4b**). This negative result demonstrates that the bump jump in the experimental genotype is primarily driven by P2X₂ expression in P-EN1 neurons. Bump jumps in control flies are likely coincidental, as the bump is often moving in a typical E-PG imaging experiment. In control flies, in the few trials where the bump jumped, the bump returned to its initial location 35% of the time. These returns are not unexpected, because the bump's location is thought to be continuously compared to an internal goal location stored downstream from E-PG neurons, and any deviation from the fly's angular goal is corrected via compensatory turning maneuvers.

Note that, in these experiments, E-PG neurons expressed GCaMP6f (under the control of *60D05-LexA*), P-EN1 neurons expressed P2X₂ and HaloTag (under the control of *VT032906-Gal4*), and DNa02 neurons also expressed P2X₂ and HaloTag (under the control of *SS00730*). We took several steps to verify that ATP directly stimulated P-EN1 neurons but not DNa02 neurons. First, we used dye to check that the ATP bolus was confined to the protocerebral bridge. Second, we confirmed that DNa02 neurons were not depolarized during the period when P-EN1 neurons were depolarized and the bump therefore jumped, i.e. in the first 0.5 s after ATP ejection (mean DNa02 voltage change in the 0.5 s after ATP ejection was 0.12 ± 0.27 mV in experimental flies and 0.14 ± 0.27 mV in genetic controls lacking the P-EN1 Gal4 driver). Rather, DNa02 neurons were only excited just before the fly made a compensatory steering maneuver, which usually occurred 1.0-1.5 s after the bump jump. Third, we verified that DNa02 was only recruited when the bump returned to its initial location via a clockwise path; if ATP had been directly exciting DNa02, then the side where DNa02 was active should not be contingent on the bump's path.

Fictive odor and fictive heat stimuli

In **Fig. 5** (and **Extended Data Figs. 6, 7, 8a-b, 10a,b**), we used the LexA system to express *LexAOp-CsChrimson::mVenus* under the control of either *Orco-LexA*⁵⁵ or *Gr28b.d-LexA*⁵⁶. To illuminate the antennae, we fixed a pair of 50- μ m fiber optic cannulas (10 mm long, FG050UGA, Thorlabs) to the underside of the recording platform, so that each fiber was directed at one antenna. Alignment of each cannula was confirmed in every experiment by looking for uniform and symmetric illumination of each antenna³⁰. A 660 nm light source was fiber-coupled to each cannula (M660F1, Thorlabs). Stimulation lasted for 0.5 s, with an inter-stimulus interval of 11 s.

EM reconstruction

Identifying DNa02. We traced neurons in the female adult fly brain (FAFB) dataset³¹, using the CATMAID toolkit⁶⁷. To identify DNa02, we began by tracing all the neurites with large cross-sections in the tract where DNa02's primary neurite resides, in the vicinity of the ventral LAL. All these primary neurites were reviewed by two experts, who agreed that only one neuron resembled DNa02. We then fully reconstructed the skeleton of this DNa02-like neuron. Next, we confirmed that this neuron bore an excellent resemblance to skeletons we traced from two 3D confocal images of the DNa02-specific split-Gal4 line (*SS00730*) driving CD8::GFP expression (**Fig. 6a**). The first of these images was provided to us by Gwyneth Card¹⁵. The second image we obtained independently, via immunostaining and confocal imaging using standard procedures. We registered both 3D images to a common template brain⁶⁸ (JFRC2013) using the neuropil counterstain (anti-Bruchpilot immunostaining). Image registration was carried out as described previously⁶⁹ using the CMTK registration suite (<https://www.nitrc.org/projects/cmtk>). Single neurons were manually traced from 3D confocal images using the Simple Neurite Tracer plugin⁷⁰ in ImageJ. For comparison with light-level data, the EM skeleton was also registered onto the JFRC2013 template brain.

Inputs to DNa02. Once we identified DNa02 in the FAFB dataset, we annotated all the presynaptic sites (1,497) and postsynaptic sites (10,842) within the brain domain of this neuron. We then partially reconstructed the skeletons of 658 neurons presynaptic to DNa02 and 25 neurons postsynaptic to DNa02. Finally, we selected 28 presynaptic neurons for complete skeletonization. Each of these 28 neurons has been reconstructed by one expert and independently reviewed by a second expert.

Analysis of EM data. Analyses of FAFB skeletons and synapses were performed in R using open-source software. RCatmaid⁶⁸ was used to interface with CATMAID servers. All other neuron analyses and visualizations were performed with natverse⁶⁸ functionality or custom-written code. Synapses and connected skeletons were retrieved from CATMAID using catnat (<https://jefferislab.github.io/catnat/>). In order to group synapses by brain region (**Fig. 6b**), we used the FAFB neuropil meshes corresponding to standard brain region definitions⁷¹. To identify brain regions containing presynaptic neurons (**Fig. 6d-j**), we calculated the cable length within each neuropil mesh (3D bounding box). The 2-4 brain regions containing the most cable were hand-traced in Adobe Illustrator and added to **Fig. 6d-j**.

Identifying DNa02 in the FlyEM dataset. We identified the FlyEM 'hemibrain' neuron 1140245595 as DNa02 based on its similarity to DNa02 in the FAFB dataset. Specifically, we considered (1) cell morphology, (2) the number of pre- and postsynaptic sites in the cell, and (3) the number of synaptic inputs to the cell from clearly identifiable presynaptic cell types (e.g., PFL3 neurons, LC9 neurons, and other cell types not shown). By all these criteria, the match was very good.

Computational modeling

After reconstructing all 12 PFL3 neurons that project to DNa02 on the right side of the brain, we determined the protocerebral bridge (PB) glomerulus and FB column containing each PFL3 neuron's dendritic arbors. We then used the FlyEM Project hemibrain v.1.0.0 dataset to determine that this pattern is fairly stereotyped across brains and mirror-symmetric on the two sides of the brain. The 'hemibrain' dataset confirmed that PFL3 neurons receive direct input from E-PG neurons in the PB, as well as a major input from Delta7 neurons. We assume that Delta7 inputs simply reinforce the effect of E-PG inputs, and so we do not explicitly model Delta7 neurons.

We modeled E-PG neurons ("compass neurons") as 8 units with uniformly-spaced preferred headings (**Fig. 7a**). At each time step, the E-PG neuron whose preferred heading most closely matched that of the fly's current heading (θ_t) was assigned a firing rate of 1, and E-PG neurons with progressively more distant heading preferences were assigned firing rates of 0.5, 0, -0.5, and -1. In an actual brain, the preferred heading of an E-PG neuron is not hard-wired^{9,72,73}, so we arbitrarily assigned a preferred direction of "South" (π rad) to the E-

PG neurons in the ventral wedge of the ellipsoid body. The mapping from E-PG neurons to PB glomeruli was based on published anatomical data⁵¹. On each side of the brain, there are 11 PFL3 neurons that each arborize in a single PB glomerulus; we assigned one E-PG input neuron to each of these uni-glomerular PFL3 neurons. Moreover, on each side of the brain, there is one PFL3 neuron that arborizes in two PB glomeruli; we modeled the E-PG input to each of these bi-glomerular PFL3 neurons by taking the mean of two adjacent E-PG firing rates. All these connectivity patterns are shown explicitly in **Fig. 7a**. In this diagram, the spatial relationships among PFL3 neurons are intended to resemble the depiction in **Fig. 6d**: the 12 model PFL3 somata on the fly's right side represent the neurons that receive input from the right PB ($i = 1$ to 12), and the 12 somata on the fly's left side represent the neurons that receive input from the left PB ($i = 12$ to 14).

In the initial version of the model (**Fig. 7b**), there was no FB input to PFL3 neurons. Each PFL3 neuron firing rate (P_i) was simply a rectified version of its E-PG input. Each DN summed its PFL3 inputs

$$D_R = \sum_{i \in \{8,10,13,14,16,18,19,20,21,22,23,24\}} P_i \quad \text{Equation 1}$$

$$D_L = \sum_{i \in \{1,2,3,4,5,6,7,9,11,12,15,17\}} P_i \quad \text{Equation 2}$$

where D_R and D_L are the firing rates of the right and left DN. The right-left difference in DN firing rates was used to update θ after adding Gaussian noise:

$$\theta_{t+1} = \theta_t + a(D_R - D_L + \varepsilon) \pmod{2\pi \text{ rad}} \quad \text{Equation 3}$$

where $a=0.1$ and $\varepsilon \sim \text{Normal}(\mu = 0, \sigma^2 = 0.36)$. The new heading was then used to update the E-PG input to PFL3 neurons in the next time step of the simulation. This model assumes that PFL3 neurons are excitatory, but if we instead assume they are inhibitory, the model behaves similarly. When there is no FB input to PFL3 neurons, the model always converges to the ‘‘North’’ heading (0 rad; **Fig. 7b,c**). When noise pushes the heading away from this stable point, the result is a right/left asymmetry in PFL3 input to DNs which steers the heading back to its stable point. In a noiseless model, South (π rad) is also stable.

In the next version of the model (**Fig. 7d**), we included FB input to PFL3 neurons. The FlyEM hemibrain dataset identifies FB local neurons as the major FB inputs to PFL3 neurons. Because essentially nothing is known about FB local neurons, we modeled their population activity as a simple time-invariant spatial pattern. Fig. 7d shows the effect of three static FB local neuron firing rate patterns: $\{0, 0, 0, 0, 0, 0, 1, 1, 1, 1, 1, 1\}$, $\{1, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0\}$, and $\{0, 0, 0, 1, 1, 1, 1, 1, 1, 0, 0, 0\}$. We modeled the FB input to PFL3 neurons as inhibitory (with a weight of -1) because many FB local neurons express Gal4 under the control of a VGlut-Gal4 driver (flycircuit.tw), and glutamate is often inhibitory in the *Drosophila* brain⁷⁴⁻⁷⁷. There are 12 FB local neurons in the model. We noticed that each PFL3 dendrite in the FB overlaps fairly precisely with the dendrite of one other PFL3 neuron in the FB (e.g., $i = 1$ and 8 overlap; $i = 2$ and 10 overlap; etc.; **Fig. 7a**), so we modeled these specific pairs of PFL3 neurons as postsynaptic to the same FB local neuron input. Each PFL3 neuron summed its E-PG and FB input, and rectified the result:

$$P_i = \max(E_j - F_k, 0) \quad \text{Equation 4}$$

where E_j and F_k are the firing rates of the presynaptic E-PG neuron and FB local neuron, respectively.

In the final version of the model (**Fig. 7g**), we kept the PFL3→network as it was before (Equations 1 and 2), but we also added a direct sensory input to DNs which bypasses the compass. Partway through each simulation, we simulated the appearance of an object to the East of the fly ($\pi/2$ rad) by creating a difference in right-left DN firing rates of $(\pi/2 - \theta_t)$:

$$\theta_{t+1} = \theta_t + a(D_R - D_L + \varepsilon + b(\pi/2 - \theta_t)) \pmod{2\pi \text{ rad}} \quad \text{Equation 5}$$

where $b=8$. We chose a value of b that was large enough so that the stimulus-directed steering drive effectively overrode the effect of PFL3 neurons when the object was present. Lower values of b produced a heading which

was intermediate between the compass-directed goal heading (North) and the stimulus-directed goal heading (East). The object did not change in intensity as the fly approached it.

Optogenetic activation of DNa02

For the experiments in **Fig. 8a-c**, we used 2-3 day old females. Flies expressed FLP under the control of a heat-shock promoter, and they harbored *UAS-FRT.mCherry.FRT.ReachR.citrine* and *SS00730* transgenes. These flies were allowed to develop at room temperature ($\sim 21^{\circ}\text{C}$), as we found that this was sufficient to produce stochastic labeling. We took virgin flies from this culture and we fixed them into our standard platform. However, dissection was performed prior to making behavioral measurements, i.e. the head cuticle remained intact. After the flies acclimated to the spherical treadmill, orange light (617 nm, M617L3, Thorlabs) was delivered from above the head, so that the left and right sides of the head were illuminated symmetrically. The power density at the fly's head was $68 \mu\text{W}/\text{mm}^2$. Light was switched on for 0.5 s every 11 s for a total of 150 trials. The experimenter remained blind to the fly's behavior throughout these trials. At the end of these trials, the dorsal cuticle was dissected away and ReaChR expression was scored. In a total of 73 experiments, we observed bilateral expression in 47 brains, left-only expression in 9 brains, right-only expression in 8 brains. and no expression in 8 brains. After ReaChR expression was scored, flies were divided into data analysis categories, and behavioral data was analyzed accordingly.

Optogenetic inactivation of DNa01 or DNa02

For the experiments in **Fig. 8d-h and Extended Data Fig. 9**, we used 2-3 day old females, comprising 7 control flies (empty-splitGal4), 11 flies where DNa01 expressed GtACR1 (*SS00731*), and 6 flies where DNa02 expressed GtACR1 (*SS00730*). Flies were kept on potato food infused with all-*trans* retinal (35 mM in ethanol; R2500, Sigma-Aldrich) for 24 hrs prior to behavioral testing. Flies were cold-anesthetized and placed in round arenas (38 mm in diameter) with transparent floors (clear acrylic, McMaster-Carr, 8536K134) and custom dome-shaped tops which were 3D-printed with white resin to promote walking⁷⁸ and illuminated with diffuse overhead infrared light (850 nm ring LEDs, Alexnld). Flies were imaged from below using a camera (FLIR, BlackFly BFS-U3-13Y3M-C) equipped with a macro-zoom lens (Rainbow, S6X11). The lens was fitted with a 720 nm long pass filter (Hoya Filter, R72 Infrared) to block room light and optogenetic stimulation light. The video image was 1024 by 1280 pixels and was acquired at 60 Hz. Optogenetic stimulation was delivered using a green LED (Thorlabs, M530L4) positioned under the arena, and focused on the area of the arena using an aspheric condenser lens (Thorlabs, ACL25416U). Each experiment lasted 2 hrs, and during this time, the green LED was continuously cycled between on (2 min) and off (2 min). Videos were streamed to disk in 2-GB (27-s) chunks, while light on/off status was recorded using a custom Matlab program. In general, flies moved more with the light on, and this effect was quantitatively indistinguishable in the three genotypes. Therefore, we decided to focus on time windows where the fly was moving, in order to ask whether specific movement features are affected by DN silencing. For this reason, in the analyses below, we discarded data chunks where the fly's time-averaged translational velocity was $<4.0 \text{ mm/s}$.

Data processing and analysis

Data alignment

Behavioral, stimulus, and imaging data were aligned using triggers acquired on the same NI-DAQ board.

Data analysis – spherical treadmill data preprocessing

Kinematic data used in computing linear filters, signal autocorrelations, LOWESS-fit scatter-plots, and event-triggered analysis were processed as follows. First, the zero-point of each kinematic signal was corrected for

small artifactual offsets by finding the median of periods of inactivity, defined as regions with instantaneous difference of less than 0.025°/s. The offset value during these periods was largely consistent an entire experiment, so a global subtraction was used to remove this offset for each experiment.

We then lightly smoothed each kinematic variable by convolving with a 50-ms Gaussian kernel bounded from -3.5σ to $+3.5\sigma$, and re-normalized. In order to remove remaining high frequency noise while preserving large signal excursions, we then used Gaussian process smoothing. This smoothed the data by fitting the time series to a Gaussian random walk process and taking the maximum a posteriori probability (MAP) estimate for each time point, as found with the L-BFGS-B optimizer. The model was implemented with PyMC3⁷⁹. Briefly, the model is as follows:

$$z_i \sim \text{Normal}(z_{i-1} + \mu, (1 - \alpha) * \sigma^2) \quad \text{Equation 6}$$

$$y_i \sim \text{Normal}(z_i, \alpha * \sigma^2) \quad \text{Equation 7}$$

Where each time point y_i was modeled as a draw from a normal distribution parameterized with mean $z_i + \mu$ and standard deviation $\alpha * \sigma^2$, where α acts as a factor to assign signal variance to a moving average or normally distributed noise. All variables except α were fit using the approximate inference software; μ was given a large initial prior standard deviation (100000, or 5-6 times the largest value present in the data) and it quickly converged to a stable value. We set α to 0.2 to assign 20% of signal variance to noise and the remaining 80% to the random walk. All time series were scrutinized to ensure proper baseline correction and smoothing.

Total speed was calculated by taking the absolute value of rotational velocity, sideways velocity, and translational velocity, and then summing these values:

$$\text{total speed} = |r| + |s| + |t| \quad \text{Equation 8}$$

where r , s , and t are all expressed in units of °/s (i.e. sideways and translational velocity were not yet converted into units of mm/s). This metric quantifies the overall level of fly movement, irrespective of the direction of movement.

Data analysis - electrophysiology data preprocessing

To compute mean changes in membrane voltage relative to baseline (mean Δ voltage) in **Fig. 2f** (and **Extended Data Figs. 3a,b, 4b, 5, 10b**), we first removed spikes by median-filtering the raw voltage trace. To compute firing rates in **Figs. 5 4f-h, 2b,d,e, and 3b,c** (and **Extended Data Figs. 3, 4d, 6a, 7b, 10b**), we first detected spikes by using a relative prominence metric in the Matlab findpeaks function. We then counted spikes in 10-ms non-overlapping bins, and smoothed with an exponential filter using the smoothts function in Matlab with a 30-ms window size to generate a continuous waveform of firing rate versus time.

For **Figs. 1, 2c, and 3d,e** (and **Extended Data Figs. 1, 4c, 10d**), firing rate was calculated as follows. First, data were conservatively band-pass zero-phase filtered using a first-order Butterworth filter with natural (3 dB) frequency of 100 Hz. Filtered time series were then normalized to a 500-ms rolling estimate of the median absolute deviation (MAD). MAD is defined as the median of absolute deviations from the median of the entire dataset. We used a corrected version of this estimate to maintain asymptotically Normal consistency by multiplying by 1.4826. This aids spike detection when firing rates are high. Spikes were then detected using the relative prominence metric in the find_peaks function in the SciPy Python library, where the prominence of a given spike was defined over a 10-s period. Prominence values were selected for each experiment and neuron. While requiring threshold tuning for some experiments, this method was robust in all DNa01 and DNa02 recordings. To estimate the firing rate, detected spikes were then binned at 1.25 ms and smoothed by convolving with a 2.5-ms Gaussian kernel as described above. All steps of firing rate calculation were scrutinized to ensure proper behavior by plotting periods known to be difficult or easy to estimate (i.e. small inter-spike-interval or high MAD).

To compute autocorrelations (**Extended Data Fig. 1b,h**) and linear filters (**Figs. 1, 2c, 3d,e**, and **Extended Data Fig. 1c-g,i-m**), and for event-triggered analyses (**Extended Data Figs. 4c, 10d**), we processed electrophysiology data as follows. We first estimate an initial “offset” by taking the median of the first 30-60 s. We then median-filtered the entire time series using a kernel length of 35 ms. This window length was empirically selected to optimally preserve low-frequency voltage modulations while excluding most of the spike waveform. We then removed high frequency artifacts remaining after median filtering by convolving with a 5-ms Gaussian kernel (bounded from -3.5σ to $+3.5\sigma$ and re-normalized). The resulting time series was then detrended by subtracting the average of linear fits to voltage data over every 30 s and 120 s, and then adding back the initial offset. All time series were scrutinized to ensure that median filtering and detrending were successful.

To compute neuron-behavior relationships, firing rates and membrane voltage values were downsampled to the same sampling rate as the spherical treadmill data. This was done by applying an upsampling by a factor of 4, applying a linear-phase finite impulse response filter, and then downsampling.

Data analysis - Signal autocorrelation analysis

Kinematic variable autocorrelation and neuron firing rate autocorrelation were calculated using the correlate function in the python library SciPy (**Extended Data Fig. 1b,h**) or Matlab’s xcorr (**Fig. 2b**).

Data analysis - Linear filter analysis

Linear filters (**Figs. 1, 2c, 3d,e**, and **Extended Data Fig. 1c-g,i-m**) were calculated using the sampling rate of the spherical treadmill data, meaning that electrophysiology data were downsampled as described in the preprocessing section. We calculated these linear filters by treating one signal as input and the other as output. We found the fast Fourier transform (FFT) of 1-sample overlapping 4-s windows of both input and output, and then computed

$$F(\Omega) = (input^*(\Omega) \times output(\Omega)) / (input^*(\Omega) \times input(\Omega)) \quad \text{Equation 9}$$

where $input^*$ represents the complex conjugate of $input$. The denominator of this equation is the input power spectrum, and the numerator is the cross-correlation between input and output in the frequency domain. We then applied the inverse of the fast Fourier transform to F to yield the linear filter relationship between input and output. For filter analyses, all behavioral variables were represented in units of $^\circ/s$ (i.e., sideways velocity and translational velocity were not converted from $^\circ/s$ to mm/s). This was done so that all filters had the same units ($^\circ/s$ in **Fig. 1** and **Extended Data Fig. 1c-g**, or $s/^\circ$ in **Extended Data Fig. 1i-m**), allowing direct comparison of different filter amplitudes.

Our input and output signals have little high-frequency content, and so division in the frequency domain introduces some noise into the linear filter. Moreover, each class of input-output relationship has its own characteristic frequency content. For this reason, we low-pass filtered the linear filters to remove this noise differently for each class. Behavior→Neuron filters were filtered with a 0th order Slepian window of 6 Hz bandwidth, whereas Neuron→Behavior filters were low pass filtered using a 0th order Slepian window of 15 Hz. These windows maximize energy in the central lobe and were empirically selected to avoid distortion of spectral content within the window.

To generate behavioral predictions, we convolved each cell’s firing rate estimate with its Neuron→Behavior filters. To find the variance explained by each filter, we linearly regressed the predicted output against the actual output and took the R^2 value. Since each experiment differed in signal-to-noise, and low pass filtering was applied globally, the length of filter used in prediction was optimized within valid bounds for each neuron (100 ms to 4 s). In general, we found that a shorter filter was optimal for rotational velocity and sideways velocity, whereas a longer filter was optimal for translational velocity and total speed, consistent with the

differing timescales of the autocorrelograms of these variables (**Extended Data Fig. 4h**). For generating a two-cell behavioral prediction from DNa02 dual recordings, we simply summed the predictions of both filters (**Fig. 3d,e**).

Data analysis - transitions between movement and immobility

To generate **Extended Data Fig. 10d**, we defined total speed as $|r|+|s| + |t|$, as noted above, where r , s , and t are velocities expressed in units of $^{\circ}/s$. We applied a threshold of $75^{\circ}/s$ to the fly's total speed in each experiment in the increasing direction (to detect "starts") and also the decreasing direction (to detect "stops"). We then discarded starts where the fly was not consistently stopped before starting (i.e., in the window 750 ms before threshold crossing, 90% of data points were not less than or equal to $\frac{1}{2}$ the threshold value), and we also discarded starts where the fly was not consistently moving after threshold crossing (i.e., in the window 750 ms after threshold crossing, 90% of the data points were not greater than the threshold value). Conversely, we also discarded stops where the fly was not moving consistently before stopping (i.e., in the window 750 ms before threshold crossing, 90% of the data points were not above the threshold value), and we also discarded stops where the fly was not consistently stopped after the threshold crossing (i.e. in the window after threshold crossing, 90% of the data points were not less than or equal to $\frac{1}{2}$ the threshold value). Our conclusions based on this analysis were generally insensitive to the length of the window chosen. Finally, we aligned events by the time of threshold crossing, and we averaged data across events within each experiment before averaging across experiments.

Data analysis – transitions to backward walking

To generate **Extended Data Fig. 4c**, a threshold of $-75^{\circ}/s$ was applied to translational velocity in the decreasing direction. Events were accepted if they maintained the threshold value for 90% of a 75-ms window after threshold crossing, and if translational velocity was also $>40^{\circ}/s$ in 90% of the 200 ms prior to threshold crossing. We aligned events by the time of threshold crossing, and we averaged data across events within each experiment before averaging across experiments.

Data analysis – colormaps of behavioral data in 2-neuron space

To generate the colormaps in **Figs. 2d,f and 3b** (and **Extended Data Figs. 3a,b, 4b, 5**), we shifted the neural data forward by 150 ms to account for the delay in behavior relative to neural data. We then divided both neural data and behavioral data into 50-ms nonoverlapping time windows. We computed the average firing rate or membrane voltage in each window; these values were then used to bin the behavioral data in 2-dimensional firing rate space or 2-dimensional membrane voltage space.

Data analysis - central complex stimulation

For the imaging data in **Fig. 4**, imaging planes through the ellipsoid body (EB) were resliced to obtain coronal sections. Using custom Matlab code, we manually divided the EB into 8 equal wedge-shaped sectors. We then calculated $\Delta F/F$ for each EB sector, defining F as the average over time of the lower half of the raw fluorescence values for that sector.

We estimated the position of the bump of activity in E-PG neurons by computing a population vector average (PVA)⁹. Our experimental design required that we focus on trials where the bump was clearly visible and stable, jumped after the ATP puff, and then returned to its initial position; our goal was to study DNa02 neuron responses during this sequence of events. We therefore discarded trials where the bump amplitude faded during the trial – specifically, where PVA magnitude dropped below the 7th percentile (for that experiment) for more than 1 s during the trial. We then discarded trials where the bump was already moving during the 1 s prior to the

ATP puff, meaning that the standard deviation of the bump's position during that period was >1.5 sectors. At this point, we computed the "initial position" of the bump in each trial as its mean position during the 1 s prior to the ATP puff. Next, we discarded trials where the bump did not jump in response to the ATP puff, i.e. where the bump's position did not move by at least 0.5 sectors. We also discarded trials where the bump did not return to dwell within 0.5 sectors of its initial position (i.e., its position before the bump jump) for at least 0.5 s. Our results were not especially sensitive to the precise values of any of these data inclusion criteria.

We defined the bump's return period as the time window beginning with the maximum excursion of the bump (from its initial position) and ending with the bump's return to a position within 0.5 sectors of its initial position. We found the time point during each trial where the bump's return speed was maximal. We then aligned trials by this time point before averaging data across trials and experiments. We confirmed that counterclockwise bump movements are generally associated with rightward behavioral turns, while clockwise bump movements are associated with leftward behavioral turns, as described previously³⁹.

Data analysis - optogenetic inactivation of DNa01 or DNa02

We removed 27-s video chunks where the fly moved less than one body length. We then cropped each chunk, with the field of view centered around the center of the fly, and we used these videos as inputs to the DeepLabCut⁸⁰ annotation software. In 1300 frames from 12 flies, we manually annotated 8 body parts in each frame: the tarsi of each leg, the posterior tip of the abdomen, and the center of the anterior edge of the head. Frames were chosen for manual annotation to maximize diversity of walking angles in the field of view. Training was performed on 95% of this 1300-frame dataset, and the remaining 5% was used to visually evaluate model performance by comparing the locations of manually labeled points and auto-labeled points. When the model mislabeled frames, it was generally because of an unusual walking angle or lighting condition, and so we added frames with these conditions to our training set and recomputed the detection model to improve performance in these situations. We also excluded stretches of data when the fly was grooming, jumping, or walking on the ceiling of the arena, because all these events reduced the model's performance. We labeled a leg "in stance" if the smoothed instantaneous velocity of the limb was <8 mm/s (averaged over 3 consecutive frames); otherwise the leg was labeled "in swing".

Sideways speed and rotational speed (Fig. 8e) Taking data from all three genotypes, we divided each video chunk into non-overlapping 500-ms windows. We then computed the vector between the head and the abdomen tip. We used this vector to compute the fly's translational velocity, sideways speed, and rotational speed, for each frame, averaged over all frames in the window. We discarded windows where the fly's average translational velocity was <2.0 mm/s. To determine whether there was a significant effect of DNa01 silencing on sideways speed or rotational speed, we performed a two-factor ANOVA with genotypes (DNa01/control) and light (on/off) as factors, and we examined the p-values associated with genotype \times light interaction, after p-values were corrected for multiple comparisons (Bonferroni-Holm correction, $m=2$ tests). We used the same procedure to determine whether there was a significant effect of DNa02 silencing.

Leg movements associated with body rotation events (Fig. 8f,g and Extended Data Fig. 9) We analyzed five metrics intended to capture the multi-leg kinematic features associated with body rotations. These metrics were based on previous descriptions of leg movements during body rotations^{24,41}. Importantly, we determined the number of metrics we would test, and we fixed the exact definitions of those metrics, based on data from control genotypes alone. Our goal was to identify leg-kinematic features associated with body rotation in control flies, in order to subsequently determine if any of these features were altered when DNs were silenced. To analyze leg-kinematic features associated with body rotation, our first step was to detect body-rotation events by searching forward in time for moments where the fly's rotational speed exceeded a threshold of $20^\circ/s$ and stayed above that threshold for at least 0.1 s. We then extracted a 500-ms window of data starting 100 ms prior to the threshold crossing. (We explored several window sizes, and chose 500 ms because it produced the

strongest relationship between the multi-leg metrics and rotational movement, considering all five metrics taken together.). Next, for each leg, we detected every complete swing epoch and every complete stance epoch in the window. Then, in each window, we computed the following metrics (iF = inner front, iM = inner middle, iB = inner back, oF = outer front, oM = outer middle, oB = outer back):

1. *Stance direction* (**Extended Data Fig. 9**, row 1): The vector pointing backward along the body's long axis was taken as 0° . We measured each leg's movement direction (in body-centric coordinates) during each stance. When the fly is walking forward, leg stance directions are close to 0° (i.e., every leg is moving almost straight in the backward direction relative to the body). During left turns, the legs in stance tend to move rightward ($0^\circ < \theta < 180^\circ$), while conversely, during right turns, the legs in stance tend to move leftward ($0^\circ > \theta > -180^\circ$); see **Extended Data Fig. 9b₁**. The stance direction metric was defined as the mean of the stance directions of the oF, iF, and iM legs (averaged over all epochs in the window).
2. *Swing direction* (**Extended Data Fig. 9**, row 2) The vector pointing forward along the body's long axis was taken as 0° . We measured each leg's movement direction (in body-centric coordinates) during each swing. When the fly is walking forward, leg swing directions are close to 0° (i.e., every leg is moving almost straight in the forward direction relative to the body). During left turns, the legs tend to swing leftward ($0^\circ < \theta < -180^\circ$), while conversely, during right turns, the legs tend to swing rightward ($0^\circ > \theta > 180^\circ$); see **Extended Data Fig. 9b₂**. The swing direction metric was defined as the mean of the stance directions of the oF, iF, and iM legs (averaged over all epochs in the window).
3. *Swing distance* (**Extended Data Fig. 9**, row 3) We measured the distance each leg moved during its swing epochs. The swing distance metric was defined as the mean of the swing distances of the OF and OM legs (averaged across swing epochs), divided by the mean swing distances of the iM and iB legs (averaged across epochs in the window).
4. *Stance duration* (**Extended Data Fig. 9**, row 4 and **Fig. 8f,g**) We measured the time each leg spent in each of its stance epochs. The stance duration metric was defined as the stance duration of the iB leg (averaged over stance epochs), divided by the mean stance durations of the iF, oF, oM, and oB legs (averaged across epochs in the window).
5. *Swing duration* (**Extended Data Fig. 9**, row 5) We measured the time each leg spent in each of its swing epochs. The swing duration metric was defined as the mean swing duration of the iM and iB legs (averaged across epochs in the window).

To determine whether there was a significant effect of DNa01 silencing on any metric, we performed a two-factor ANOVA with genotypes (DNa01/control) and light (on/off) as factors, and we examined the p-values associated with genotype \times light interaction, after p-values were corrected for multiple comparisons (Bonferroni-Holm correction, $m=5$ tests). We used the same procedure to determine whether there was a significant effect of DNa02 silencing.

Step frequency, step length, and translational velocity (Fig. 8h) We divided each video chunk into non-overlapping 500-ms windows. We discarded windows where the fly's time-averaged translational velocity was <2.0 mm/s. Step frequency was defined as the mean frequency of stance onset for all six legs, averaged across all complete stride epochs within that window, for all legs. Step length was defined as the mean stride distance of all six legs, averaged across all complete stride epochs within that window, for all legs. Translational velocity was defined as the velocity of the vector connecting the abdomen tip to the head, in the direction of that vector, for each frame, and then averaged over the window. To determine whether there was a significant effect of DNa01 silencing on any of these three variables (step frequency, step length, translational velocity), we performed a two-factor ANOVA with genotypes (DNa01/control) and light (on/off) as factors, and we examined the p-values associated with genotype \times light interaction, after p-values were corrected for multiple comparisons (Bonferroni-Holm correction, $m=3$ tests). We used the same procedure to determine whether there was a significant effect of DNa02 silencing.

For the leg-centric metrics above, if any leg included in a given metric did not complete at least one of the relevant epochs during a particular time window, then that window was omitted. Moreover, for the ANOVAs

described above, if a fly did not contribute at least 5 time windows to the dataset for that particular metric, then the fly was omitted. The number of flies that contributed data to each panel was as follows (control/DNa01/DNa02 genotypes): **Fig. 8e** (7/11/6); **Fig. 8f,g** (5/8/6); **Fig. 8h** – step frequency (6/8/6), step length (6/10/6), translational velocity (7/11/6). For **Extended Data Fig. 9**, the values are: row 1 (6/9/6); row 2 (6/9/6); row 3 (6/9/6); row 4 (5/8/6); row 5 (6/9/6).

Post hoc statistical tests: When any ANOVA yielded a significant genotype×light interaction, we followed up by performing a post hoc Tukey test to determine whether, for each genotype, there was a significant effect of light on/off. Finally, to confirm that it was reasonable to treat each time window as an independent datapoint, we performed a separate two-factor ANOVA with fly ID and light as factors, and we verified that variance due to fly ID is much smaller than total variance.

Data inclusion

In analysis of electrophysiology experiments, we excluded cells where the recording lasted ≤ 15 minutes. Late in a recording, the membrane voltage sometimes became depolarized, which we interpret as a sign of poor recording quality; we therefore discarded any extended epoch where the membrane voltage was more depolarized than -33 mV.

In **Fig. 5** (and **Extended Data Figs. 8, 10**), flies were excluded that displayed highly asymmetric stimulus-evoked steering behavior. For example, we excluded a fly if on average, she displayed significant stimulation-evoked turns to the right but not to the left. These cases are likely due to asymmetric positioning of the fiber optic filaments near the antennae. This excluded 2/6 flies for fictive odor experiments and 3/7 flies for fictive warming experiments.

In **Fig. 4**, flies were excluded if there were < 6 trials that passed the checks described above (see *Data analysis – central complex stimulation*). This excluded 4 of 8 recordings.

In single-cell recordings (**Figs. 1, 4, 5; Extended Data Figs. 1b-m, 7b, 8, and 10**), we excluded 1 DNa01 cell (out of 8 total) and 1 DNa02 cell (out of 11 total) where the recording quality fluctuated substantially over time.

DATA AND CODE AVAILABILITY

EM reconstructions will be archived at virtualflybrain.org. Code will be available at <https://github.com/wilson-lab>. Raw data are available upon reasonable request.

REFERENCES FOR ONLINE METHODS

52. Pfeiffer, B., *et al.* Refinement of tools for targeted gene expression in *Drosophila*. *Genetics* **186**, 735-755 (2010).
53. Pfeiffer, B.D., *et al.* Tools for neuroanatomy and neurogenetics in *Drosophila*. *Proc. Natl. Acad. Sci. U. S. A.* **105**, 9715-9720 (2008).
54. Jenett, A., *et al.* A GAL4-driver line resource for *Drosophila* neurobiology. *Cell Rep.* **2**, 991-1001 (2012).
55. Lai, S.L. & Lee, T. Genetic mosaic with dual binary transcriptional systems in *Drosophila*. *Nat. Neurosci.* **9**, 703-709 (2006).
56. Frank, D.D., Jouandet, G.C., Kearney, P.J., Macpherson, L.J. & Gallio, M. Temperature representation in the *Drosophila* brain. *Nature* **519**, 358 (2015).
57. Chou, T.-b. & Perrimon, N. The autosomal FLP-DFS technique for generating germline mosaics in *Drosophila melanogaster*. *Genetics* **144**, 1673-1679 (1996).

58. Golic, K.G. & Lindquist, S. The FLP recombinase of yeast catalyzes site-specific recombination in the *Drosophila* genome. *Cell* **59**, 499-509 (1989).
59. Bloomington *Drosophila* Stock Center. Notes on X-linked insertions of heat shock-FLP constructs. [^https://bdsc.indiana.edu/stocks/recombinases/hsflp_notes.html](https://bdsc.indiana.edu/stocks/recombinases/hsflp_notes.html), accessed April 1, 2020
60. Sutcliffe, B., *et al.* Second-generation *Drosophila* chemical tags: Sensitivity, versatility, and speed. *Genetics* **205**, 1399-1408 (2017).
61. Inagaki, H.K., *et al.* Optogenetic control of *Drosophila* using a red-shifted channelrhodopsin reveals experience-dependent influences on courtship. *Nature Methods* **11**, 325-332 (2014).
62. Hampel, S., Franconville, R., Simpson, J.H. & Seeds, A.M. A neural command circuit for grooming movement control. *eLife* **4**, e08758 (2015).
63. Chen, T.W., *et al.* Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295-300 (2013).
64. Klapoetke, N.C., *et al.* Independent optical excitation of distinct neural populations. *Nat. Methods* **11**, 338-346 (2014).
65. Seelig, J.D., *et al.* Two-photon calcium imaging from head-fixed *Drosophila* during optomotor walking behavior. *Nat. Methods* **7**, 535-540 (2010).
66. Zheng, Q., *et al.* Rational design of fluorogenic and spontaneously blinking labels for super-resolution imaging. *ACS Cent. Sci.* **5**, 1602-1613 (2019).
67. Saalfeld, S., Cardona, A., Hartenstein, V. & Tomancak, P. CATMAID: collaborative annotation toolkit for massive amounts of image data. *Bioinformatics* **25**, 1984-1986 (2009).
68. Bates, A.S., *et al.* The natverse, a versatile toolbox for combining and analysing neuroanatomical data. *eLife* **9**, e53350 (2020).
69. Cachero, S., Ostrovsky, A.D., Yu, J.Y., Dickson, B.J. & Jefferis, G.S. Sexual dimorphism in the fly brain. *Curr. Biol.* **20**, 1589-1601 (2010).
70. Longair, M.H., Baker, D.A. & Armstrong, J.D. Simple Neurite Tracer: Open Source software for reconstruction, visualization and analysis of neuronal processes. *Bioinformatics* (2011).
71. Ito, K., *et al.* A systematic nomenclature for the insect brain. *Neuron* **81**, 755-765 (2014).
72. Fisher, Y.E., Lu, J., D'Alessandro, I. & Wilson, R.I. Sensorimotor experience remaps visual input to a heading direction network. *Nature* **576**, 121-125 (2019).
73. Kim, S.S., Hermundstad, A.M., Romani, S., Abbott, L.F. & Jayaraman, V. Generation of stable heading representations in diverse visual scenes. *Nature* **576**, 126-131 (2019).
74. Liu, W.W. & Wilson, R.I. Glutamate is an inhibitory neurotransmitter in the *Drosophila* olfactory system. *Proc Natl Acad Sci U S A* **110**, 10294-10299 (2013).
75. Collins, B., Kane, E.A., Reeves, D.C., Akabas, M.H. & Blau, J. Balance of activity between LN(v)s and glutamatergic dorsal clock neurons promotes robust circadian rhythms in *Drosophila*. *Neuron* **74**, 706-718 (2012).
76. Hamasaka, Y., *et al.* Glutamate and its metabotropic receptor in *Drosophila* clock neuron circuits. *J Comp Neurol* **505**, 32-45 (2007).
77. McCarthy, E.V., *et al.* Synchronized bilateral synaptic inputs to *Drosophila melanogaster* neuropeptidergic rest/arousal neurons. *J Neurosci* **31**, 8181-8193 (2011).
78. Simon, J.C. & Dickinson, M.H. A new chamber for studying the behavior of *Drosophila*. *PLoS One* **5**, e8793 (2010).
79. Salvatier, J., Wiecki, T.V. & Fonnesebeck, C. Probabilistic programming in Python using PyMC3. *PeerJ Computer Science* **2**, e55 (2016).
80. Mathis, A., *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* **21**, 1281-1289 (2018).