The Journal of Experimental Biology 213, 1886-1900

doi:10.1242/jeb.026526

A model of visual–olfactory integration for odour localisation in free-flying fruit flies

Finlay J. Stewart[1], Dean A. Baker[2] and Barbara Webb[1,*]

[1]Institute of Perception, Action and Behaviour, School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh EH8 9AB, UK and [2]Department of Genetics, University of Cambridge, Downing Street, Cambridge CB1 3QA, UK

*Author for correspondence (bwebb@inf.ed.ac.uk)

Accepted 21 February 2010

SUMMARY

Flying fruit flies (*Drosophila melanogaster*) locate a concealed appetitive odour source most accurately in environments containing vertical visual contrasts. To investigate how visuomotor and olfactory responses may be integrated, we examine the free-flight behaviour of flies in three visual conditions, with and without food odour present. While odour localisation is facilitated by uniformly distributed vertical contrast as compared with purely horizontal contrast, localised vertical contrast also facilitates odour localisation, but only if the odour source is situated close to it. We implement a model of visuomotor control consisting of three parallel subsystems: an optomotor response stabilising the model fly's yaw orientation; a collision avoidance system to saccade away from looming obstacles; and a speed regulation system. This model reproduces many of the behaviours we observe in flies, including visually mediated 'rebound' turns following saccades. Using recordings of real odour plumes, we simulate the presence of an odorant in the arena, and investigate ways in which the olfactory input could modulate visuomotor control. We reproduce the experimental results by using the change in odour intensity to regulate the sensitivity of collision avoidance, resulting in visually mediated chemokinesis. Additionally, it is necessary to amplify the optomotor response whenever odour is present, increasing the model fly's tendency to steer towards features of the visual environment. We conclude that visual and olfactory responses of *Drosophila* are not independent, but that relatively simple interaction between these modalities can account for the observed visual dependence of odour source localisation.

Supplementary material available online at http://jeb.biologists.org/cgi/content/full/213/11/1886/DC1

Key words: *Drosophila melanogaster*, modelling, multimodal, olfaction, optomotor, vision.

INTRODUCTION

Animals interact with their environment using a number of sensory modalities. Consider, for example, an airborne fly searching for food. It could be alerted to the possibility of feeding by the odour of a fermenting banana, but in order to reach the fruit it may have to visually steer toward the yellow patch of its visual surround whilst avoiding obstacles such as branches. However, integrating multiple sensory modalities to produce appropriate behaviour is difficult, as they can give conflicting information. For example, a strong headwind may suggest to the fly, *via* its mechanosensory and proprioceptive systems, that it is moving forward, while its vision would suggest that it is in fact moving backwards.

In a tethered-flight experiment on *Drosophila*, Frye and Dickinson report that the response elicited by simultaneous presentation of visual and olfactory stimuli is approximated well by the sum of responses to each individual stimulus, and suggest that this 'addition' may simply represent simultaneous activation of different flight muscles controlled by separate sensory pathways (Frye and Dickinson, 2004). However, other experiments suggest more complex interactions. Free-flying Dipterans are only able to locate a concealed appetitive odour source when the walls of an experimental arena are patterned, and furthermore the patterns must contain vertical contrasts (Frye et al., 2003). Chow and Frye, in further tethered flight tests, show responses to optomotor stimuli (large-field rotation) that stabilise straight flight are enhanced by the presence of odour (Chow and Frye, 2008), and Duistermars and Frye find that reductions in the frequency and amplitude of saccades in response to odour are observed only when the fly has high contrast visual surrounds (Duistermars and Frye, 2008). Thus odour localisation, which one might at first assume to be a task that could be solved by purely chemotactic behaviour, appears to rely on features of the visual environment.

The purpose of the current study is to understand the interplay between these two modalities, and how this might lead to successful odour localisation in certain visual environments but not others. Even if the responses an animal makes to various stimuli are well characterised, it is very difficult to analytically extrapolate from these data to predict the behaviours that will result when the animal is allowed to move freely in its environment (Simon, 1982). This is because the animal's behaviour at one instant will shape its sensory experience of the next, in turn affecting its subsequent behaviour. In such a situation it is essential to use closed-loop models of the systems in question to test our hypotheses about the animal's behaviour.

A number of ways in which visual stimuli evoke flight responses have been identified in *Drosophila*. The optomotor response is a very well studied behaviour whereby a fly will attempt to rotate in the same direction as a rotating visual scene, thereby eliminating retinal slip (Götz, 1968; Tammero et al., 2004). Fruit flies also strongly avoid expanding visual patterns by initiating very rapid yaw turns termed saccades (Tammero and Dickinson, 2002a; Bender and Dickinson, 2006). Additionally, *Drosophila* regulate their flight velocity by reference to the visual environment (David, 1982; Mronz and Lehmann, 2008). Although rather less is known about flight responses to olfactory stimuli, it has been established that the sudden onset of an attractive odour results in increased forward velocity and a suppression of turning (Frye and Dickinson,

2004; Budick and Dickinson, 2006). It has also recently been found that *Drosophila*'s response to rotating visual stimuli strengthens in the presence of food odour, while the response to expanding patterns weakens (Chow and Frye, 2008).

We extend these findings by analysing the trajectories of free-flying *Drosophila* in different visual environments, with and without an odour source. Using these biological data, we develop a model of visuomotor behaviour based on the mechanisms described above. Finally, we add an odour plume to the simulated environment and investigate ways in which the olfactory signal could modulate behaviour. We show that altering just two parameters of the visuomotor model according to olfactory input is sufficient to account for the experimental findings. Thus, while the visual and olfactory responses of *Drosophila* cannot be considered independent, we propose a simple and parsimonious model of the interaction between these modalities that can explain flies' inability to locate an odour source in certain visual contexts.

MATERIALS AND METHODS

Free-flight arena

The experimental arena was cylindrical (100 cm diameter, 60 cm height) and made of 2 mm transparent acrylic (Fig. 1A). Lighting was provided with a ring of six evenly spaced 60 W incandescent lamps, and additional infrared (IR) illumination (undetectable by flies) provided by a single floodlight above the centre of the arena. The floor of the arena contained three holes, 25 cm from the centre of the arena and spaced 120 deg. apart (Fig. 1C), where small (1 ml capacity) vials were placed for odour trials. Vials were coloured black and positioned within the floor to minimise their visibility.

The experimental paradigm is closely based upon that of Frye and colleagues (Frye et al., 2003). The main differences between their methodology and ours are the IR floodlight (they used a ring of IR LEDs) and the presence of three vials, as opposed to just one.

Visual cues were provided by lining the arena with one of three 'wallpapers' made from white paper with opaque black cardboard shapes attached (Fig. 1B). The random chequerboard (CB) pattern was a grid of 43 mm squares (each subtending a visual angle of 5 deg. from the centre of the arena) that were either black or white with equal probability. The horizontal stripe (HS) pattern consisted of alternating black and white horizontal bands of width 43 mm, while the localised vertical contrast (LV) pattern was similar to the HS arena, except for the presence of a vertical white band 244 mm wide (or 28 deg., viewed from the centre of the arena) with a single black stripe of width 70 mm (8 deg.) running down its centre. Odours were then either aligned with the localised contrast (LV-near) or positioned 120 deg. away from it (LV-far) for odour localisation trials. All patterns contain strong contrasts throughout their extent in order to saturate the flies' contrast adaptation mechanisms. This prevents them from using any residual texture created by imperfections in the arena, which is a potential confounding phenomenon identified in previous free-flight studies (Frye and Dickinson, 2007).

3D tracking

Two cameras (Marlin F131B; Allied Vision Technologies, Ahrensburg, Germany) were mounted on the ceiling at either side of the floodlight, approximately 30 cm from the centre. The cameras had their IR cut-filters removed, giving some sensitivity to light in the near-IR range. Each camera was calibrated by means of eight LEDs mounted in known positions on the arena wall, allowing position and orientation to be computed, making it possible to determine the 3D location of an object imaged by them. The cameras acquire 50 frames s^{-1}, but because of bandwidth limitations it was not possible to image the whole arena in both cameras at this frame rate. Instead, a smaller region of interest was used for tracking flies in real-time, using a custom-written C++ program.

Tracking was achieved using an adaptive background subtraction algorithm to identify the fly's position in each image, and calculating the 3D position by triangulation. When possible, any missing data points in one camera were reconstructed from the other by linearly interpolating altitude. However, if gaps are too long or if both cameras lose the fly, reconstruction of the trajectory is not possible for that interval. In general, therefore, we do not have a single unbroken trajectory for each experimental trial, but rather a series of trajectory segments.

Experimental procedure

All experiments were performed on 2- to 4-day old mated female fruit flies, *Drosophila melanogaster* Meigen, from the same wild-type laboratory culture as used by Frye and colleagues (Frye et al., 2003). Flies were reared in a 12 h:12 h light:dark cycle, with experiments performed between 3 and 6 h after dawn. Flies were deprived of food but not water overnight (~19 h) to ensure motivation for foraging. An experimentally naive female was introduced to the arena by means of an aspirator, and allowed to fly within the arena for 10 s to remove any bias resulting from its position of introduction.

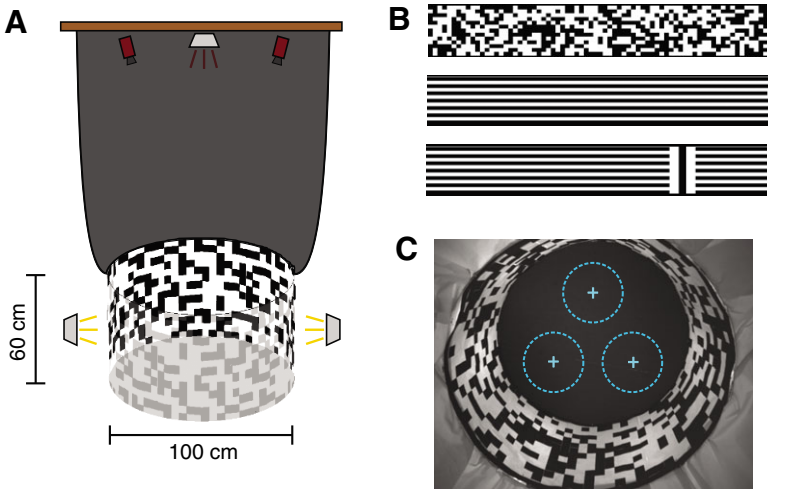

Fig. 1. Free-flight experiments. (A) Schematic representation of the free-flight arena. (B) The three wallpapers used to line the inside wall of the arena: chequerboard (CB, top), horizontal stripes (HS, middle) and localised vertical contrast (LV, bottom). (C) An image taken from one of the two tracking cameras, showing the positions of the vials and the extent of the zones used to quantify odour localisation behaviour. The curtain around the arena is black, but reflects near-IR light to which the camera is sensitive. The floor absorbs both visible and IR light.

Flight patterns were subsequently recorded for 4 min. Any recording of total length <30 s (due to the fly landing or flying up into the region of the cameras) was discarded.

In vision-only trials, all three vials contained water. In odour trials, two vials contained water and one balsamic vinegar (Carlo Magno, Nonantola, Italy). The vinegar vial was put in a slot 10 min before the start of an experiment to allow the odour to diffuse. After each replicate the odour vial was removed and the arena ventilated before repeating the procedure. During each day's 3 h session the odour always occupied the same slot. Data sets for odour conditions consist of an equal number of trials with the odour in each slot, to control for any bias that might exist in the arena.

Data analysis

Experimental trajectories were smoothed using Gaussian convolution with s.d.=28 ms to remove jitter and single-frame tracking errors. Trajectories were subsequently re-sampled at 20 ms intervals for further analysis. As the cameras have insufficient resolution to estimate the fly's orientation, we make the assumption that a fly is always aligned with its direction of motion. This approach has been adopted by a number of studies (Tammero and Dickinson, 2002b; Frye et al., 2003; Maimon et al., 2008). One might argue that this assumption is realistic because there is evidence to suggest that for a small fly like *Drosophila*, flight dynamics are dominated by frictional forces, meaning that side-slip is minimal (Hesselberg and Lehmann, 2007; Mronz and Lehmann, 2008). However, other authors have argued that significant side-slip could occur (Fry et al., 2003), and it has recently been directly demonstrated that *Drosophila* are capable of performing lateral flight manoeuvres (Ristroph et al., 2009). Thus, it should be recognised that our estimation of yaw orientation is an approximation.

Saccades are defined as any period when the fly's yaw velocity (calculated using the aforementioned orientation assumption) exceeds 450 deg. s^{-1}. This is a stricter criterion than the 300 deg. s^{-1} limit employed by previous studies (Tammero and Dickinson, 2002b; Frye et al., 2003), though it should be noted that they use a different smoothing procedure.

To quantify odour localisation (OL) performance, we define three cylindrical zones of radius 16 cm and height 60 cm centred on the odorant/water slots (Fig. 1C). We compare the time spent by the fly in the odour zone with that spent in the two water controls to compute an odour localisation index (OLI) as shown in Eqn 1:

$$\mathrm{OLI} = \frac{\mathrm{time}_{\mathrm{odour}}}{\mathrm{time}_{\mathrm{odour}} + \mathrm{time}_{\mathrm{control1}} + \mathrm{time}_{\mathrm{control2}}} . \quad (1)$$

Thus, the OLI ranges from 0 (perfect avoidance) to 1 (perfect localisation) with a chance baseline level of 0.33. Wilcoxon matched pairs (*T*) tests were used to determine significance of OL within one condition, while Mann–Whitney *U*-tests were used to compare OL performance between conditions.

Simulation

We modelled the fly's behaviour using a custom-written Java program. Simulated flies traversed an arena with dimensions and wall patterns identical to those used in the animal experiments, with the regions above and below the walls uniformly black. The model fly can 'see' the full 360 deg. azimuth and from –80 to +60 deg. elevation (Fig. 2), sampled at 1.8 deg. resolution. The pitch and roll orientation of the model fly's eyes is held constant, which is analogous to a real fly moving its neck to perfectly compensate for any changes in body orientation associated with flight manoeuvres. Blowflies are known to at least partially cancel out body movements by such a mechanism (Gilbert et al., 1995; van Hateren and Schilstra, 1999). The model fly also (for simplicity) maintains a constant altitude. As described in detail below, we model peripheral visual motion detection, then define wide-field filters for different patterns of optical flow, which in combination determine the speed and yaw of the model fly. These can be modulated by the presence of odour or rate of change of odour concentration.

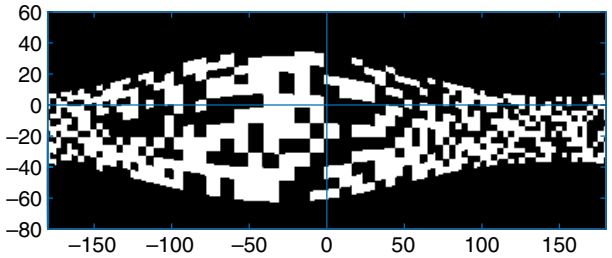

Fig. 2. Visual simulation. A sample of the model fly's retinal image in the CB arena. Angles in degrees.

Because of the difficulty of simulating the diffusion of odorants, we take an empirical approach to modelling the odour plume. A vial of ethanol was put in a mock-up of the experimental arena and the vapour was allowed to diffuse as in the fly experiments. Using a gantry robot equipped with a volatile organic compound (VOC) sensor (MiCS-5521; MicroChemical Systems, Corcelles, Switzerland), we repeatedly sampled the concentration of ethanol at each point on a 3D grid of resolution 10 cm in the arena (Fig. 3). This sensor was chosen for its rapid response in order to minimise the degree of temporal averaging in the readings, allowing us to capture the time-variant nature of the plume. Modelling the sensor as a first-order low-pass filter, we estimated its time constant to be approximately 160 ms (data not shown).

It is reasonable to assume that the diffusion of ethanol would be similar to vinegar because diffusion over distances of more than a few centimetres is predominantly turbulent (i.e. because of convective currents in the air) as opposed to molecular, so does not depend greatly on the molecular properties of the chemical (Murlis et al., 1992).

Each VOC measurement consisted of a 3.4 s duration recording at 20 Hz, i.e. 68 sensor readings. This was taken at every grid point in 15 different experimental trials at times ranging from 5 to 20 min after the introduction of odour to the arena. This gives a total of 15×68=1020 individual samples for each point in space. Data from mirror-symmetric positions about the *y*-axis were pooled, so in most positions there were in fact 2040 rather than 1020 readings to sample from. These data were then used in the simulation as follows. At every simulation time step, a grid point was chosen randomly using a 3D Gaussian distribution of s.d.=5 cm centred on the model fly's current position. This procedure was used to avoid the discontinuities between grid points that a strict nearest-neighbour scheme would create. One randomly selected reading recorded at the chosen grid point was then used as the odour intensity. This process was employed to emulate the high degree of variability that is found at any given point in the odour plume (Fig. 3).

Peripheral visual processing

Visual motion detection is achieved by using 'delay-and-correlate' elementary motion detectors (EMDs) (Hassenstein and Reichardt, 1956). These each take two photoreceptor inputs, which, based on physiological data (Stavenga, 2003), are positioned 5 deg. apart in

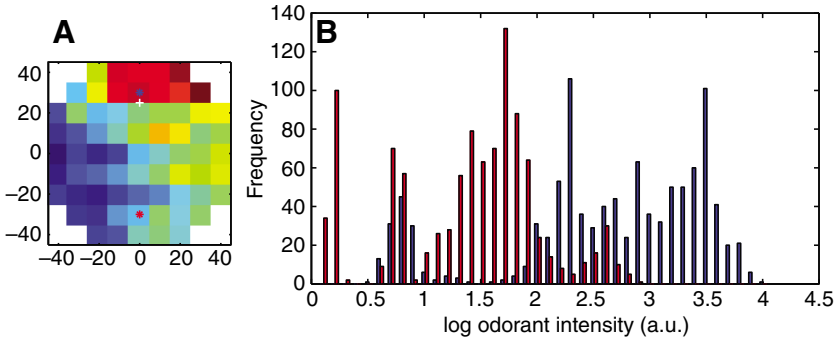

Fig. 3. Volatile organic compound (VOC) sensor measurements. (A) Map of mean log odorant intensity for the horizontal slice of altitude 35 cm (arbitrary heat scale). The white '+' indicates the odour source position. (B) Histogram showing the full set of readings for the points marked by the blue and red stars in A. Each dataset consists of a total of 1020 readings, in 15 batches of 68 readings (3.4 s recording at 20 Hz). Each batch is taken from a separate experimental trial, accounting for the multimodal nature of the distributions. This ensures that the plume simulation contains a realistically high degree of variability. a.u., arbitrary units.

visual space, with the spatial sensitivity of each modelled as a 2D Gaussian distribution of s.d.=1.49 deg. The photoreceptor signals are passed through first-order high-pass filters with a time constant (τ) of 10 s. This approach models slow photoreceptor adaptation and serves to remove any DC component of the signal. Subsequently, one of the signals is passed through a first-order low-pass filter (τ=40 ms) in order to introduce a delay, and is then correlated with the other (un-delayed) signal. This process is illustrated in Fig. 4. Note that EMDs do not respond to image velocity *per se*, but rather temporal frequency, i.e. the response depends on the spatial frequency of the retinal pattern as well as the speed at which it moves. Furthermore, their response is non-monotonic, peaking at an optimal temporal frequency. The time constant we have used for the low-pass filter is based on the finding that *Drosophila* respond optimally to moving patterns with a temporal frequency of ~7 Hz (Duistermars et al., 2007). The optimal frequency for a canonical EMD is given by $1/(2\pi\tau)$ (Zanker et al., 1999), corresponding to a τ of 23 ms. However, the addition of high-pass filters on the inputs tends to shift the tuning curve toward higher frequencies, so the higher value for τ was empirically chosen to counter this effect.

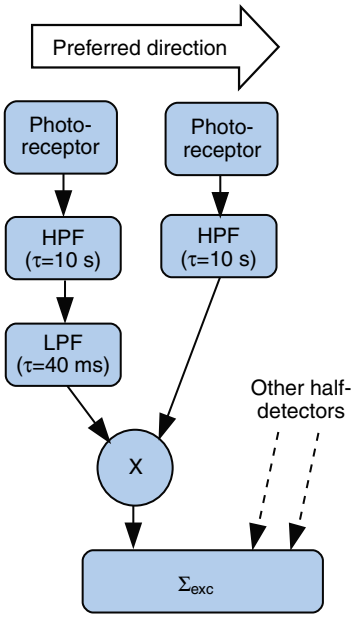

Fig. 4. Single delay-and-correlate 'half-detector'. The two photoreceptor inputs are 5 deg. apart in visual space, representing adjacent ommatidia. HPF denotes a first-order high-pass filter, and LPF a first-order low-pass filter. Motion in the preferred direction would elicit a positive response. The output of multiple half-EMDs is pooled by the component marked Σ_{exc}. Every half-detector like this one has a mirror-image counterpart (i.e. with the LPF on the other channel) whose output is passed to the corresponding inhibitory summation cell (not shown).

To achieve wide-field motion sensitivity, we pooled the outputs of an ensemble of EMDs as follows. Each EMD has two halves, an excitatory one that responds to movement in the preferred direction, and an inhibitory one that responds to movement in the opposite direction. All of the excitatory inputs (exc) are summed, as are the inhibitory ones (inh), and output of the whole ensemble is given by Eqn 2, where ε is a small constant and n is the number of EMDs in the ensemble:

$$\text{Output} = \frac{\sum_{i=1}^{n}\text{exc}_i - \sum_{i=1}^{n}\text{inh}_i}{\sum_{i=1}^{n}\text{exc}_i + \sum_{i=1}^{n}\text{inh}_i + n\varepsilon} . \quad (2)$$

This process achieves the phenomenon termed gain control (Borst et al., 1995; Single et al., 1997), whereby the response saturates to a temporal frequency-dependent value as the retinal size of the moving pattern increases, with the degree of pattern size invariance dependent on the constant 'leak' $n\varepsilon$. Each filter uses a different combination of EMD positions and orientations across the visual field. We chose filters that match patterns of retinal motion (or optic flow) created by either pure translational or pure rotational movements of the fly. Electrophysiological experiments on blowflies show that the optimal stimuli for wide-field motion-sensitive neurons bear a close resemblance to such patterns (Krapp and Hengstenberg, 1996; Krapp et al., 1998), and recent work indicates that analogous cells exist in *Drosophila* (Joesch et al., 2008). We define three filters that differ in the orientation of the pole of the optic flow pattern, and the parts of the visual surround to which they apply, which have different effects on the model fly motion as follows.

Optomotor response (OMR)

The OMR is perhaps the most straightforward visual reflex to model, as it is reasonable to assume that the appropriate optic flow filter is one corresponding to pure yaw rotation. We opted for a non-uniform distribution of EMDs, concentrating them in the frontal visual field (Fig. 5A). The resulting heightened sensitivity to frontal retinal slip increases the model fly's tendency to steer towards vertical contrasts in sparse environments (see Results).

The filter response is low-pass filtered, and passes through a leaky accumulator to remove noise by temporal averaging and is then used directly to control angular velocity. It has been noted (Srinivasan et al., 1999) that a simple OMR would cause a moving animal to steer towards obstacles, because the side on which objects are closer

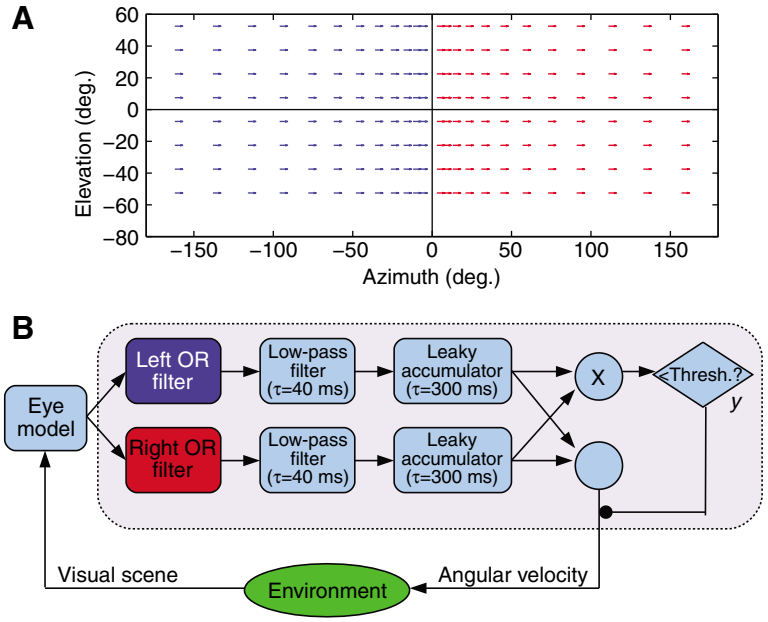

Fig. 5. Optomotor response (OMR). (A) Receptive field of the left (blue) and right (red) wide-field filters driving the optomotor response (OR). Each arrow represents a single EMD, with the endpoints denoting the positions of the two ommatidia giving input. The field corresponds to pure yaw rotation. (B) Control diagram of the OMR system. The dark blue and red boxes represent the filters shown in A, which detect optic flow patterns corresponding to yaw rotation in the two visual hemispheres. The filled circle represents an inhibition that entirely suppresses the signal on the line on which it impinges, which occurs if the left and right filters have strong but opposite sign responses. The dotted box represents the boundary of the OMR system.

will experience more rapid front-to-back motion, causing a turn to that side. However, such an effect is not normally observed in flying insects, suggesting the OMR is overridden in situations where the optic flow pattern more closely resembles translation than yaw rotation. Hence, we drive the optomotor response using two large-field filters, one for each visual hemisphere (Fig. 5A). The outputs of these are multiplied, and if the product is below a negative threshold value (i.e. left and right motion is in opposite directions, as expected for forward translation), the optomotor response is suppressed entirely. The whole process is illustrated in Fig. 5B.

Speed regulation (SR)

Flies clearly use visual information to modulate their flight speed, as their average speeds vary significantly in different visual surrounds (Tammero and Dickinson, 2002b), and they adjust their speed to compensate for motion of their visual environment (David, 1982; Mronz and Lehmann, 2008). A frontally oriented expansion field is a suitable filter to use for regulating velocity. Additionally, the ventral portions of the visual surround most reliably indicate flight speed (Neumann and Bülthoff, 2002) because the ground will generally be closer than the sky (thus giving a larger signal) and at a more constant distance than objects to the sides. Thus, the filter used is as shown in Fig. 6A. A simple proportional controller is used to attempt to keep the visual velocity at a set point, as shown in Fig. 6B. Although we do not model altitude regulation, a control system of this type has the property that flight speed increases proportionally to altitude (Franceschini, 2004).

Collision avoidance (CA)

It is known that expanding patterns trigger CA saccades directed away from the focus of expansion (Tammero and Dickinson, 2002a; Bender and Dickinson, 2006). One might therefore suppose that two expansion detectors centred at reasonably large azimuths to either side would be appropriate for CA. Indeed, this is essentially

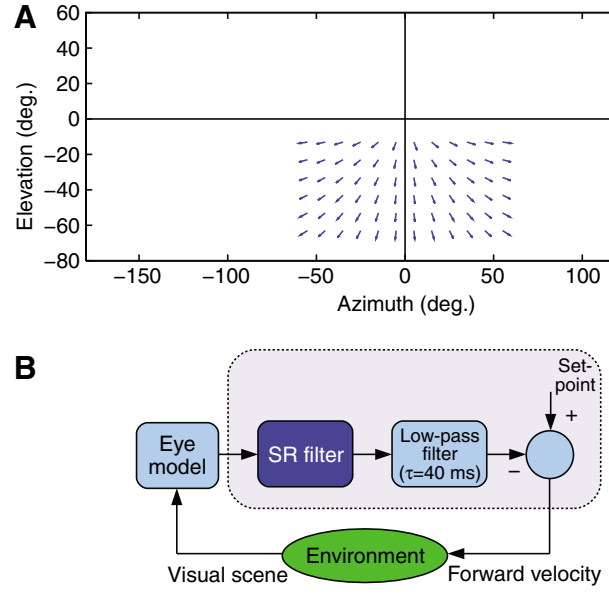

Fig. 6. Speed regulation (SR). (A) Receptive field for the filter driving speed regulation. (B) Control diagram for the SR system, which uses translational optic flow cues to control the model fly's forward velocity. The dark blue box represents the filter shown in A.

the approach taken by Reiser and Dickinson (Reiser and Dickinson, 2003). There are, however, some problems with this scheme – if the fly is flying straight, the focus of expansion will be frontally positioned, and only when the trajectory is substantially curved (or involves side-slip) would these 'matched filters' be a good to fit to the expansion stimulus. (Consequently, they will frequently cause the fly to saccade in the wrong direction in situations where horizontal motion cues must be relied upon, as illustrated in supplementary material Fig. S1.) A frontally centred expansion pattern split into two halves would resolve the issue described above, but suffers from the problem that a rightward rotation would cause the whole image to move left, promoting a rightward saccade and inhibiting a leftward one, and *vice versa*. Consequently we have chosen filters that represent a compromise between these two extremes (Fig. 7A). They have the same spatial extent to either side of the expansion pole, meaning that yaw rotation cancels out. However, they are only offset from the centre by 3 deg., resulting in a large overlap area. These filters could not be confined to EMDs from a single eye, because each eye can only see up to ~20 deg. into the contralateral visual field (Krapp and Hengstenberg, 1996). This does not represent a problem for the plausibility of the model, as binocularly sensitive cells have been identified, albeit in blowflies (Krapp et al., 2001; Farrow et al., 2006).

The full CA control scheme is shown in Fig. 7B. The left and right expansion filters feed into leaky accumulators, and if the signal exceeds a set threshold, a saccade in the opposite direction will be initiated and the model fly turns according Eqn 3, i.e. the weighted sum of two Gaussians (Fig. 8A), where t is in ms from the instant at which the saccade is triggered:

$$\text{Angular velocity}\ (t) = \text{amplitude} \times \left[0.7\exp\left(\frac{-(t-160)^2}{2\times 28^2}\right) + 0.3\exp\left(\frac{-(t-160)^2}{2\times 56^2}\right)\right]. \quad (3)$$

Thus the peak angular velocity is reached after 160 ms, when the term inside the square brackets is equal to 1. This angular velocity profile is based on our free-flight data.

The amplitude of the profile depends on the fly's flight speed, as we identified a relationship between these variables in our free-flight data (Fig. 8B). Thus amplitude (in deg. s^{-1}) is given by Eqn 4, where speed is in m s^{-1} and N denotes a random Gaussian variable. The parameters in this equation are determined by performing a linear regression on the free-flight data:

$$\text{Amplitude} = (1550 - 1106 \times \text{speed}) \times N(1,0.26)\,. \quad (4)$$

During a saccade a fly's forward velocity is reduced, possibly representing the effect of forward thrust being diverted into yaw torque. The model fly has its velocity adjusted during a saccade according to Eqn 5. Again, the scaling parameter is based on free-flight data (Fig. 8C):

$$\text{Adjusted speed}\ (t) = \text{speed} \times \left(1 - \frac{|\text{angular velocity}\ (t)|}{4000}\right). \quad (5)$$

Non-visual wall avoidance

If the model fly comes within 8 cm of the arena wall, and is not already performing a saccade, it will initiate an 'emergency' saccade away from the wall. While this is included for primarily practical reasons to prevent wall collisions by the simulated fly, there are a number of possible mechanisms the fly could be using for such emergency turns, e.g. detecting auditory reflections when it was within a few centimetres of the wall (Robert and Göpfert, 2002), performing escape when looming shapes reach a certain retinal size (Holmqvist and Srinivasan, 1991; Bender and Dickinson, 2006), or being stimulated by fine contrasts on the walls (e.g. the texture of the paper) that would only be resolved at very close range (and are not included in the simulated arena).

Olfactory response

We ignore the issue of odour identification altogether, since the chemical composition of the attractant is constant. Thus the model fly's olfactory input signal can be taken as a scalar value. Olfactory receptor neurons are known to have low-pass (Justus et al., 2004; Schuckel et al., 2008) as well as high-pass (i.e. adaptation) characteristics (Störtkuhl et al., 1999). Adaptation is also found at the first synapse of the olfactory system in the antennal lobe (Kazama and Wilson, 2008), and *Drosophila* larvae exhibit marked behavioural adaptation to the prolonged presentation of odour (Cobb and Domain, 2000). Indeed, it seems clear that the temporal

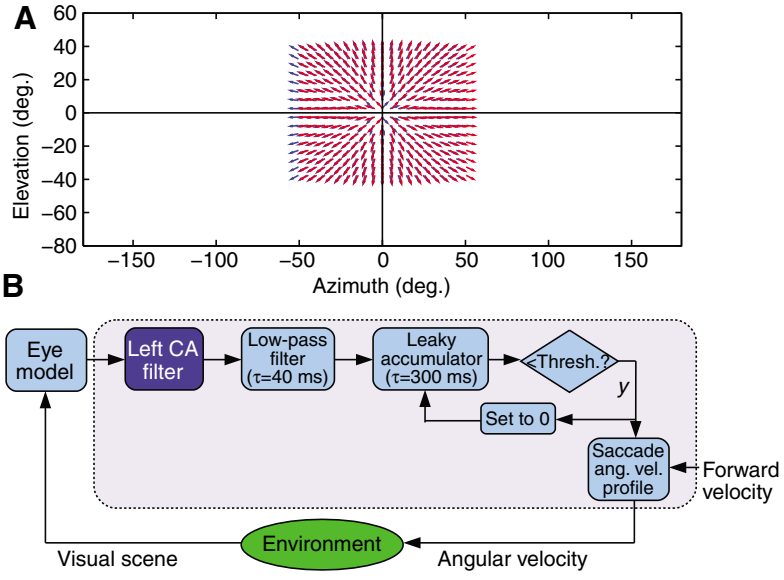

Fig. 7. Collision avoidance (CA). (A) The CA filters used in the model. Each covers ~105 deg. of azimuth but they are centred at ±3 deg. (elevation=0 deg.). (B) Control diagram for collision avoidance. Only the half of the system that triggers rightward saccades is shown for clarity; the other half has an identical configuration. The dark blue box represents the blue wide-field filter in A. The reset operation also applies to the other half of the system, i.e. a saccade in one direction sets both accumulators to 0. Thresh., threshold; ang. vel., angular velocity.

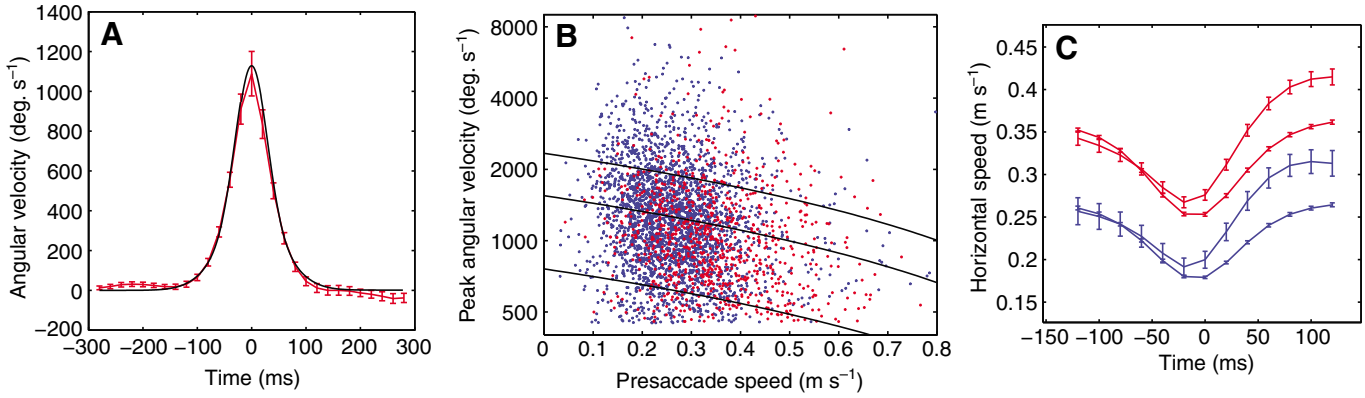

Fig. 8. Model fitting. (A) The mean angular velocity profile for saccades within 6 cm of the centre of the HS arena is shown in red. Because of the lack of vertical contrast and perspective cues, these saccades are minimally affected by visual feedback. Bars are 1 s.e.m. The sum-of-Gaussians fit is in black, with amplitude 1129 deg. s^{-1}. This amplitude was obtained using Eqn 4, based on the mean flight speed in the HS arena. (B) Saccade amplitude as a function of presaccadic speed. Blue is CB arena, red is HS. The black line shows the relationship between these variables defined by Eqn 4, with the dashed lines giving the 95% confidence interval of the Gaussian distribution. While the free-flight data are somewhat more variable than the model fit, it should be borne in mind that the more extreme data points may represent tracking errors. (C) Saccade-triggered horizontal speed profiles. Blue is CB arena, red is HS. The solid lines represent free-flight data, the dashed lines are predictions produced by Eqn 5 based on the real fly's angular velocity profile and presaccadic speed. The model is a rather crude approximation of the animal data, failing to capture the asymmetry created by the fly's rapid acceleration out of the saccade. Bars are 1 s.e.m.

derivative of odour intensity is more useful than the absolute value to an animal attempting to locate an odour source, given that baseline conditions may be highly variable. Fly larvae maintain a similar level of behaviour response over several orders of magnitude of odorant concentration. This phenomenon depends on the existence of two olfactory receptors which are sensitive to the same chemical but have very different intensity tuning curves (Kreher et al., 2008). Kreher et al. suggest that this non-linearity may be achieved by lateral connections in the antennal lobe, which have been documented in the adult fly (Olsen and Wilson, 2008).

To obtain a suitable olfactory signal (OD) we feed the raw odour level to two parallel first-order low-pass filters with different time constants. Subtracting the output of the slower filter from that of the faster filter gives an approximation of the temporal derivative, which we call OD′. Finally, we use an adaptive gain mechanism to keep the range of this signal constant despite changes in the overall amplitude of the odour signal. This is achieved by keeping an incremental estimate of the variance of the signal, with a time constant of 4 s. This is compared with a desired value for the variance, and the difference causes a compensatory change in the gain by means of a simple proportional controller. The output of the system, OD*, is given by OD′ multiplied by the adaptive gain parameter. This process is summarised in Fig. 9.

We tested several different ways in which the olfactory signal could interact with the visual control loops. As our aim is to determine what interactions best account for the experimental data, we will describe and discuss these mechanisms in the appropriate part of the Results. For each of these models of cross-modal interaction, there is some kind of gain parameter determining the strength with which olfaction affects visuomotor control. Clearly, it is essential to test a range of values for such parameters in order to properly investigate the suitability of each model. Thus, a systematic parameter search was used to ensure that models were not rejected unless we were convinced they could not reproduce the results. Details of these and other parameter values are given in the Appendix, while the complete Java source code for the model can be found http://www.ipab.inf.ed.ac.uk/cricketlab/olfaction_and_vision.html.

RESULTS

Behaviour in the absence of odour: experiments

Differences between behaviour in an environment rich in vertical contrast (the chequerboard pattern, CB) compared with one lacking vertical contrast (the horizontal stripe arena, HS) were broadly in line with previous studies (Tammero and Dickinson, 2002b; Frye et al., 2003; Frye and Dickinson, 2007). Flies approach the wall more closely and faster before saccading in the HS arena (Table 1), and while intersaccade intervals were similar across conditions, the intersaccadic distance was significantly higher in the HS arena due to the higher flight speed. Consistent with the suggestion that increasing flight force reduces manoeuvrability by constraining wing kinematics (Lehmann and Dickinson, 2001), there is a negative correlation between pre-saccade velocity and saccade amplitude in both the CB (ρ_{2539}=–0.112, P<0.0001) and the HS arena (ρ_{852}=–0.118, P=0.0003). Saccade amplitude is not significantly lower in the HS arena (96.1 *vs* 101.3 deg., $U_{24,24}$=208, P=0.099).

Straightness of flight (measured in terms of mean absolute angular velocity during intersaccadic segments) is similar across the two visual environments (CB: 34.0 deg. s^{-1} *vs* HS: 35.4 deg. s^{-1}, $U_{24,24}$=270, P=0.71). However, we find that flies veer away from the walls of the HS arena during intersaccadic flight as reported previously (Frye et al., 2003); there is a significant correlation between their angular velocity and the angle at which the arena wall is being approached (ρ_{246}=0.171, P=0.007; supplementary material Fig. S2). In contrast to the findings of Tammero and Dickinson

Table 1. Comparison of flight statistics

Statistic	CB	HS	*P*
Saccade distance to wall (cm)	31.1	25.9	**0.0001**
Intersaccadic speed (cm s^{-1})	26.7	38.1	**<0.0001**
Intersaccadic distance (cm)	21.0	29.9	**<0.0001**
Intersaccadic interval (ms)	809	826	0.76

P-values obtained using Mann–Whitney *U*-tests; significant results highlighted in bold.
There were 24 experimental replicates each for CB and HS.

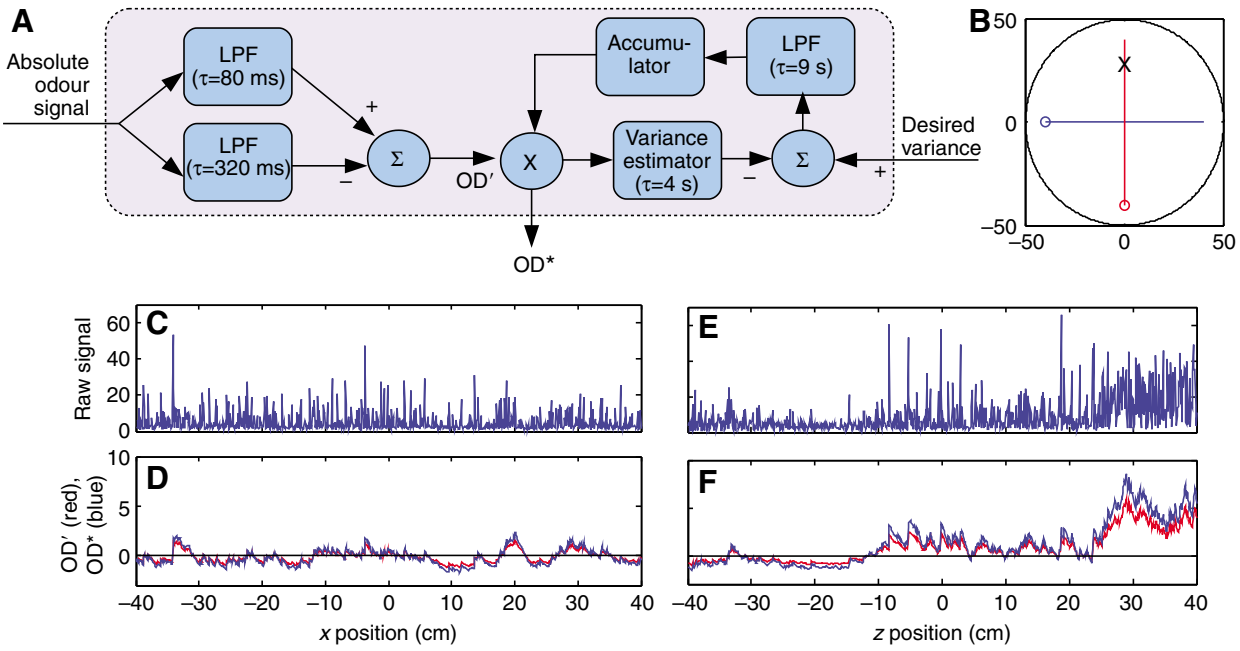

Fig. 9. Olfactory signal pre-processing. (A) Overview of the system. An estimate of the temporal derivative of the odour signal (OD′) is obtained by comparing two LPFs with different time constants. The gain of this signal is then adjusted by a slow feedback system which attempts to keep the variance equal to 1, giving us OD*. (B) Simulated trajectories for the plots in C and E (blue) and D and F (red), starting from the circle, with the odour source at the point marked X. (C) Raw input generated by the odour sampling method along the blue trajectory. Units are arbitrary. (D) Estimate of the temporal derivative of the odour signal in C both before (OD′, red) and after (OD*, blue) being subjected to adaptive gain. The values are similar because the gain is close to 1. Note the slight lag due to the use of LPFs. (E) Raw input generated from the odour sampling method along the red trajectory. (F) Estimate of the temporal derivative of the odour signal in E both before (OD′, red) and after (OD*, blue) being subjected to adaptive gain.

(Tammero and Dickinson, 2002b), no such effect is evident in the CB case (ρ_{624}=–0.012, P=0.76). Thus it appears that horizontal stripes (or rather, an absence of vertical contrasts) induce a bias in the fly's yaw stabilisation system (Frye and Dickinson, 2007).

One interesting issue in a system that combines saccades and an optomotor response is whether the rapid saccadic turn, which creates substantial visual rotation, triggers a compensatory optomotor response; or whether the response is suppressed, perhaps by some form of efference copy (Webb, 2004). We addressed this question by identifying saccades followed by long periods of uninterrupted intersaccadic flight. We notice a pronounced 'rebound', where the fly turns in the opposite direction to the saccade several hundred milliseconds after its peak (Fig. 10A). In itself, this observation could be dismissed as an artefact, given our uncertainty as to the flies' yaw orientation (see Materials and methods). However, the magnitude and time course of this rebound differ significantly between visual environments (at time=140 ms: $U_{24,24}$=452, P=0.0007). This suggests that the rebound effect is neither an artefact nor merely part of the saccade motor programme, but rather depends upon visual input, since the HS arena lacks vertical contrast and therefore gives few visual yaw cues. Further support for this account comes from the observation that in the HS arena there is a significant positive correlation between wall proximity and the size of the rebound, as measured at t=140 ms (ρ_{860}=0.103, P=0.0025) – because of perspective, there is a greater yaw cue created in the HS arena by apparent bulging of the wall as the fly gets closer to it. However, some caution in this interpretation is indicated by recent observations that side-slip reactions of flies (which we cannot distinguish in our experiments) might themselves be under some visual control (Sugiara and Dickinson, 2009).

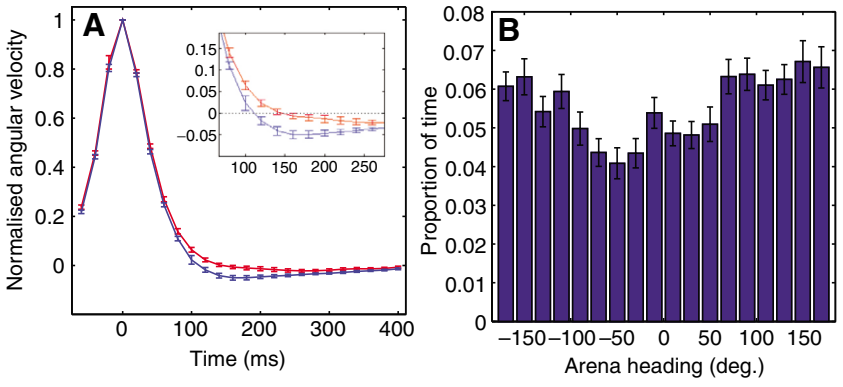

Fig. 10. Experimental results. (A) Saccade-triggered angular velocity profiles. Only data where no additional saccade occurred before t=500 ms are included. Blue is the CB arena, red is HS. Inset shows a zoomed view of the 'rebound' phase. Data are normalised according to the peak angular velocity to control for the slight difference in saccade amplitudes between arenas. Bars represent ±1 s.e.m. (B) Arena headings in the LV arena, with no odour present. The fly's heading is defined as the point on the arena wall towards which it is oriented at any given time. The vertical stripe is at the 0 deg. position. Clearly flies do not preferentially orient towards it. Bars represent ±1 s.e.m.

In an arena that has only localised vertical contrast (LV), one might expect that flies would be attracted to the vertical stripe, given that stripe fixation has previously been reliably observed in free-flight (e.g. Maimon et al., 2008) and tethered flight (e.g. Poggio and Reichardt, 1973; Duistermars and Frye, 2008). Surprisingly, we find no clear bias either towards or away from the vertical stripe (Fig. 10B; see also Fig. 12). Computing the OLI as if the vial next to the stripe contained odour gives a result of 0.322 (T_{42}=406, P=0.57), i.e. the fly does not spend its time preferentially in this region of the arena. Some possible explanations of this difference from previous studies are offered in the Discussion.

Behaviour in the absence of odour: simulation

We have tested our model with all three visual subsystems running in parallel and compare the resulting behaviour with the free-flight data. For 40 simulated seconds prior to each trial, the model fly was put in a new random (x,y) position and yaw orientation every 125 ms. This was to allow the slow adaptation of the photoreceptors to take place. The model was then initialised in a new random position with a velocity of 30 cm s^{-1}. The first five simulated seconds of flight are discarded, and the trial is only included if the model continues for a further 30 s without colliding with the wall. (Despite the emergency wall avoidance mechanism, collisions are still possible, e.g. if the resultant saccade is too small.) Each trial is ended after 40 s of recorded flight.

We find that all the differences seen between the CB and HS arenas are reproduced qualitatively by the model: in the HS arena, the model moves further and faster between saccades (Fig. 11, Fig. 13B,C), approaches the walls more closely (Fig. 12, Fig. 13A), and veers away from the walls between saccades (Fig. 11). There are, however, some quantitative discrepancies between the real and modelled flies: the model is generally less wall averse, and produces longer intersaccadic segments. Both these effects may be explained by the observation that flies make spontaneous saccades (Maye et al., 2007) in what is thought to be a sensory-independent optimal search behaviour (Reynolds and Frye, 2007). Not being triggered by visual expansion, these saccades could happen in any location. This would tend to increase the average distance from the wall at which saccades occur, and break up intersaccadic segments. As we cannot say whether such saccades are truly random as opposed to being under some unknown sensory or endogenous control, we have not included additional random saccades in our model.

It is possible to test the optomotor response in isolation, by attempting to reproduce the rebound effect we observed in the fly experiments. A model fly is placed in either the CB or HS arena, moving at the mean inter-saccadic velocity recorded in that arena. It is then made to perform a saccade (based on Eqn 4, but omitting the Gaussian noise) at the distance from the wall appropriate to that arena. The optomotor response system (Fig. 5) runs throughout the experiment in a closed-loop manner, with its output being summed with the open-loop saccade model. The results in Fig. 14 (cf. Fig. 10A) show the same pattern as in the fly experiments: the rebound is considerably larger and slightly quicker in the CB arena. Looking at the output of the OMR filters (also shown in Fig. 14) helps us to understand why. We see that the response is generally much larger in the CB case, because the vertical contrasts offer a far greater rotational cue than the bulging caused by perspective in the HS arena. The CB response is largest at moderate angular velocities, dipping at the peak of the saccade. This is because the wavelength of the pattern is of the order of tens of degrees, and therefore the temporal frequency experienced by the fly exceeds the optimal value of ~7 Hz. The HS pattern, effectively having a

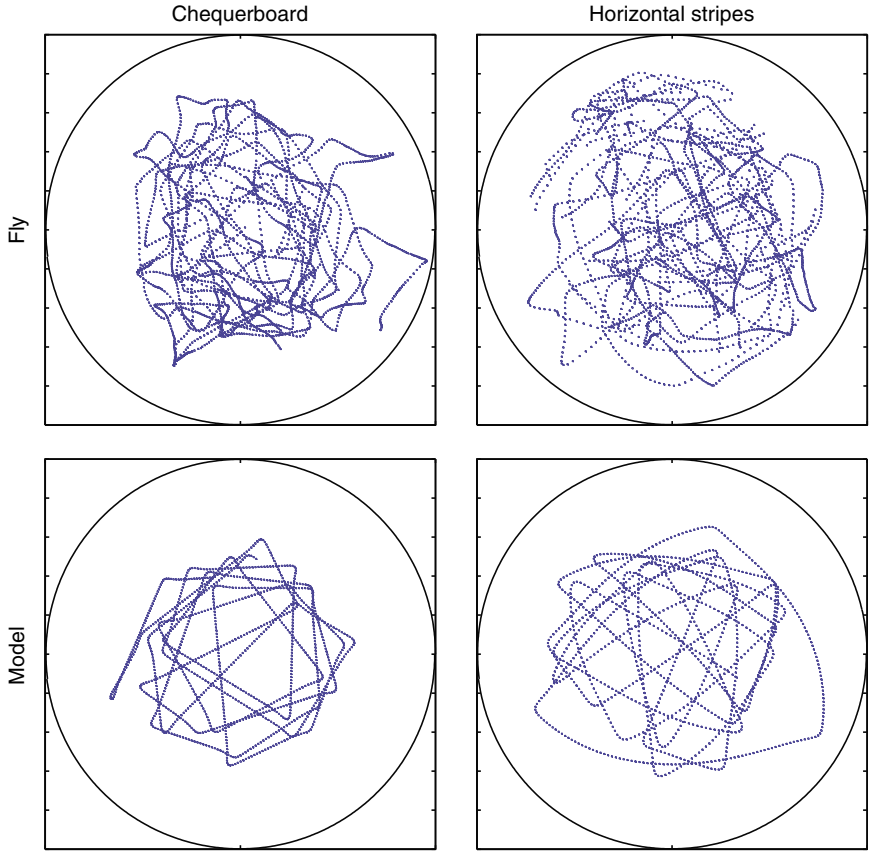

Fig. 11. Flight trajectories. Sample flights of real and modelled flies in the two visual environments. Dots represent positions every 30 ms, giving an impression of speed. In both the real and model data for the CB arena we can see slightly concave trajectories. This is due to the textured environment causing a large post-saccade optomotor rebound, meaning that the segments tend to curve in the opposite direction to a preceding saccade. Both the real and the model flies in the HS arena curve away from the walls, creating convex trajectories. In the case of the model at least, this phenomenon is due to an illusion whereby the lack of vertical contrast creates a rotation-type optic flow pattern although the fly is in fact translating, triggering the OMR.

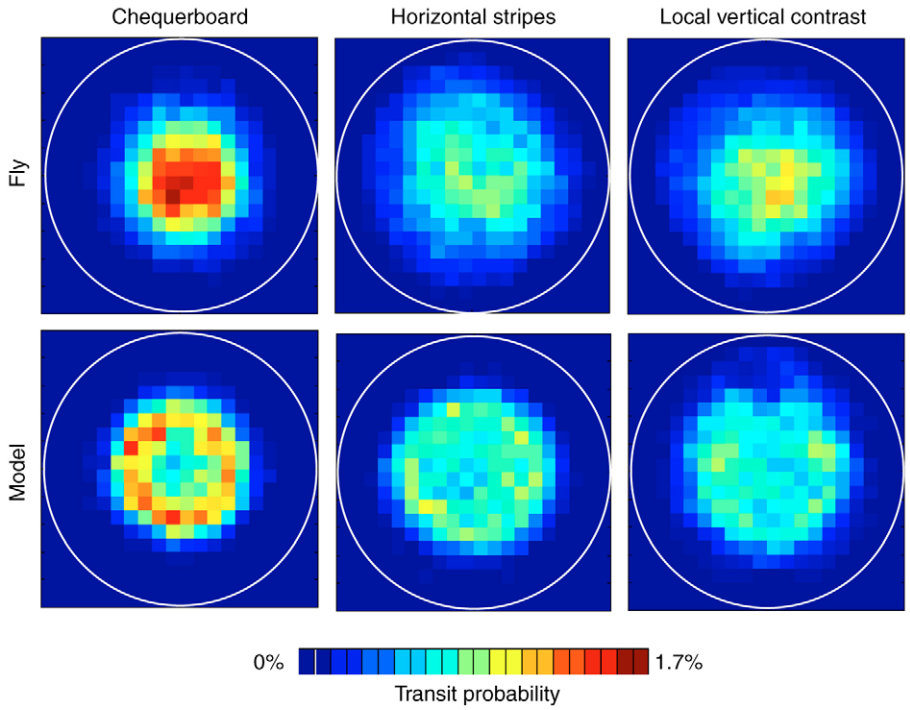

Fig. 12. Spatial distribution. Plots of transit probability (i.e. how frequently the fly passes through each grid square) for the real and modelled flies in the three arenas. The vertical contrast is at the 12 o'clock position in the LV arena. Both the real and model flies stay comparatively more centralised in the CB arena, venturing closer to the walls in HS arena. In the LV arena, neither the real nor the modelled flies are strongly attracted or repelled by the stripe.

wavelength of 360 deg., is not subject to this effect, and thus the signal peaks when the angular velocity is maximal. The amplitude of the model fly's rebound in both visual environments quite closely matches that observed in the fly, suggesting that no suppression of the OMR during or immediately following a saccade occurs in the fly. Instead, the decreasing response of EMDs to supraoptimal temporal frequencies appears sufficient to prevent instability in the OMR controller (Warzecha and Egelhaaf, 1996).

Importantly, the model shows neither a bias towards nor a bias away from the stripe in the LV arena (OLI=0.326, T_{24}=126. P=0.49), as is the case in the fly data. In tuning the model parameters, we found that the gain of the OMR system had a pronounced effect on the degree of target localisation, such that a strong OMR would bias the model to fly towards the area of local vertical contrast and a weak one would result in repulsion *via* the CA system (data not shown). The frontally concentrated OMR filter (Fig. 5A) was adopted to achieve the appropriate strength of target approach behaviour whilst maintaining a similar amplitude of optomotor rebound (i.e. response to visual yaw rotation spanning 360 deg.) as seen in the fly.

Odour localisation: experiments

As reported by Frye and colleagues (Frye et al., 2003), flies in the CB arena spend significantly more time in the vicinity of a vial filled with vinegar (an appetitive stimulus) than one filled with water (Fig. 15, OLI=0.467, T_{27}=29, P=0.0001). However, in contrast to this previous study, we find that odour localisation (OL) also takes place in a HS arena with no vertical contrasts (Fig. 15, OLI=0.378,

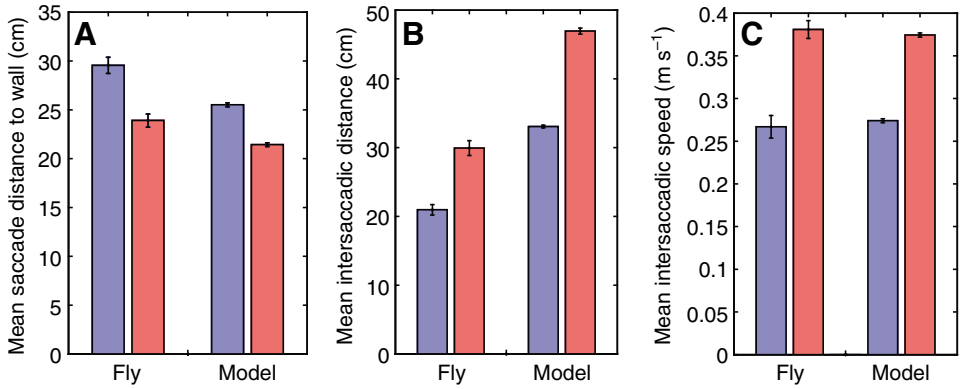

Fig. 13. Comparison of flight statistics. Blue bars are CB arena, pink are HS. (A) Mean distance from the wall at which saccades occur. (B) Mean length of intersaccadic segments. (C) Mean intersaccadic speed. Error bars are ±1 s.e.m.

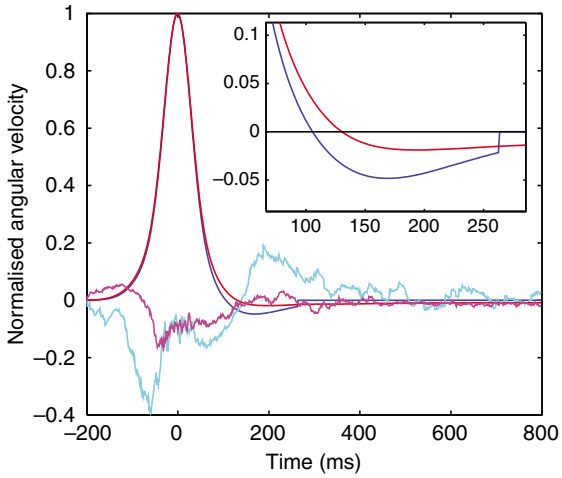

Fig. 14. Optomotor rebound in the modelled fly. Blue and red lines are normalised recordings of angular velocity in the CB and HS arenas, respectively, as in Fig. 10B. Inset shows a zoomed view of the 'rebound' phase. The discontinuity is caused by the optomotor response being suppressed by translational retinal motion (Fig. 5B). Cyan (CB) and magenta (HS) lines are recordings of OMR filter response (left and right summed), measured immediately downstream of the filters (Fig. 5B). Units are arbitrary, no normalisation occurs.

T_{27}=102, P=0.037), although OL performance is significantly stronger in the CB arena ($U_{27,27}$=237, P=0.028). Flies in the CB arena display a tendency to turn towards the wall during intersaccadic flight (supplementary material Fig. S2, ρ_{446}=–0.148, P=0.002), which they do not do in the absence of odour. The optomotor rebound effect is still exhibited, and is of a similar amplitude to that seen in the absence of odour (angular velocity at t=160 ms as a percentage of peak; CB control: –5.00% *vs* odour: –4.72%, $U_{27,24}$=300, P=0.65; HS: –0.61% *vs* –1.97%, $U_{27,24}$=379, P=0.30).

In the LV arena flies also locate the vinegar, but only when it is placed next to the stripe (Fig. 15; LV-near: OLI=0.512, T_{24}=28, P=0.0005; LV-far: OLI=0.398, T_{24}=92, P=0.098). These results can be compared to the finding that tethered *Drosophila* are unable to track an odour plume when presented only with a vertical object laterally offset from the plume (Duistermars and Frye, 2008). These authors argue from this that wide field stimuli are required for odour localisation, presumably assuming that using a stripe aligned with the odour direction would simply reveal stripe attraction rather than the potential sufficiency of this narrow-field stimulus. In our study, the ability to localise odour in the LV-near condition cannot simply be attributed to stripe attraction as this was not observed in the no odour trials. We conclude that wide-field vertical contrast is not necessarily required for odour localisation, but will discuss this further in the light of the simulation results below.

In all visual conditions the flies' mean altitude is significantly lower when odour is present, as was reported previously (Frye et al., 2003) (data not shown).

Odour localisation: simulation

Clearly, there are almost limitless ways in which olfaction could modulate flight control. We therefore employ the principle of Occam's razor, starting with the very simplest ways in which olfactory information could shape flight behaviour and only attempting more complex schemes if these can be shown to be insufficient to reproduce the experimental results.

Model 0: chemokinesis

Perhaps the simplest known algorithm that could produce odour localisation behaviour is undirected kinesis, a mechanism used, for instance, by *E. coli* to find food (Berg and Brown, 1972). This refers to a process whereby the animal will continue to move forward while odour intensity is rising, and turn randomly when it falls. The model is numbered 0 because it does not in fact involve any modulation of the visuomotor system by olfaction, but rather has the two systems operating entirely in parallel. Whenever OD* falls below a certain negative threshold value, a saccade is initiated with a 50/50 probability regarding the direction. Each odour-triggered saccade will inhibit visually triggered (CA) ones in just the same way as shown in Fig. 7B. Illustrative results are presented in Fig. 15, while supplementary material Fig. S3 shows the results of altering the saccade initiation threshold.

In all visual environments we see significant OL, clearly demonstrating that the simulated plume contains sufficient structure to allow its source to be found. Despite its simplicity, kinesis is evidently a viable strategy for odour location. However, this model does not match the behaviour of flies, as we would expect to see greater OL in the CB arena than in HS, and not to see significant OL in the LV-far condition. No value tested for the saccade initiation threshold is able to produce a significant difference in OLI between CB and HS arenas (supplementary material Fig. S3). Thus simple addition of independent olfactory and visual behaviours is not sufficient to explain fly behaviour, at least for this form of olfactory control.

It seems remarkable that such a simple control scheme can outperform flies in terms of OL, and indeed prompts the question of why flies do not make use of it. The reason for this may be that the algorithm implicitly assumes that motion of the model fly is never affected by air currents. An animal in the wild cannot generally assume that it is moving forward just because it is making the appropriate wing movements, and must therefore make use of visual input to obtain information about its movement relative to the world (Srinivasan et al., 1997). Similarly, odour plumes in natural environments are likely to be spatially and temporally more complex than in our arena set up. Consequently an animal evolving under such conditions would probably not find this simple solution to be as effective as it appears in the simulation.

Model 1: CA modulation

Clearly, olfaction must interact with the visual system to produce the environment-dependent odour localisation that we see in flies. An obvious way to produce kinesis-like behaviour *via* the visuomotor controller is to have the olfactory signal modulate the threshold for initiating CA saccades, such that saccades are inhibited by a rising odour intensity (positive OD*), and promoted by a falling one. Model 1 does just that, using OD* to continually modulate the threshold parameter in the CA system. Illustrative results are in Fig. 15, and supplementary material Fig. S4 shows the results of varying the gain with which OD* modulates the CA system.

This model produces the desired difference in OL performance between the CB and HS arenas. However, OLI is never significantly greater in LV-near than LV-far for any setting of the gain parameter, and furthermore the differences that are seen between these conditions are in the wrong direction (supplementary material Fig. S4). Thus, while this model can account for the OL effects seen in homogeneous visual environments, it fails to capture flies' behaviour in the presence of localised vertical contrast.

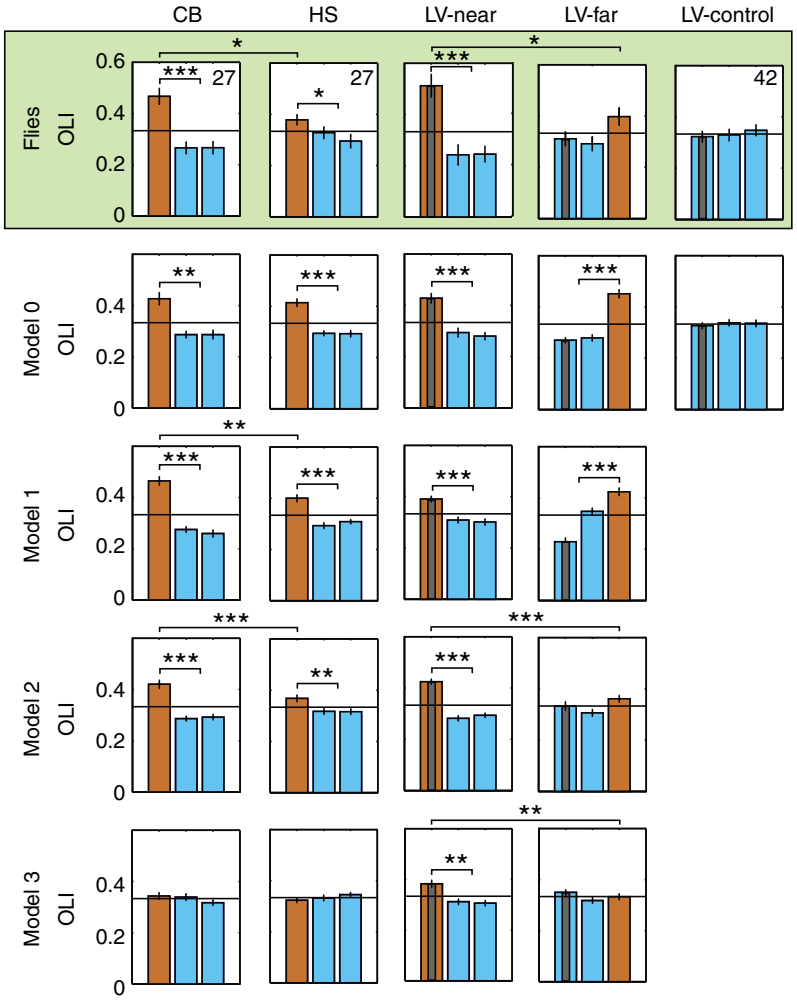

Fig. 15. Odour localisation (OL) performance of real and modelled flies. Animal data are highlighted in green. Bars show the odour localisation index (OLI) for the three vials in each test; brown bars are for vials containing vinegar, blue for ones containing water; error bars represent ±1 s.e.m. In the LV arena, the striped bar represents the vial aligned with the local vertical contrast. In the absence of odour all models are equivalent, so the LV no-odour control data are only shown once. For each model, the parameter settings that result in behaviour most closely matching that of the fly are used. Asterisks inside graphs refer to the significance of OL within that condition (i.e. odour vial *vs* mean of control vials), based on Wilcoxon matched pair tests, while those between graphs show significant differences in OLI between conditions, based on Mann–Whitney *U*-tests. Only Model 2 reproduces all of the effects seen in the animal data, i.e. significant OL in CB and LV-near, and significantly stronger OL performance in CB than in HS, and in LV-near than in LV-far. Unless otherwise specified, there were 24 experimental replicates. ***P<0.001, **P<0.01, *P<0.05.

Model 2: CA modulation + OMR boost

Recent work (Chow and Frye, 2008) indicates that flies increase the amplitude of responses to visual rotation in the presence of food odour, as well as decreasing their response to expansion stimuli. Model 2 attempts to incorporate the former effect, by boosting the gain of the OMR system whenever odour is present in the arena. Note that unlike the continuous modulation of the CA threshold based on OD*, this is a step change in OMR gain which occurs according to the presence or absence of food odour and persists for the entire experimental trial. Illustrative results are presented in Fig. 15.

Model 2 qualitatively reproduces all of the OL effects we observe in flight experiments. OL is stronger in CB than in HS (OLI=0.437 *vs* 361, $U_{24,24}$=110, P=0.0002), and unlike Model 1, we see OL occurring when the odour is in the same direction as the localised vertical contrast but not when it is positioned 120 deg. away from it (OLI=0.436 *vs* 332, $U_{24,24}$=41, P<0.0001). These effects are robust to small changes in the CA modulation gain or the OMR boost size (supplementary material Fig. S5). The step change to the OMR gain makes the model fly more inclined to approach the localised vertical contrast by increasing its tendency to steer towards it.

There are, however, some more subtle ways in which the model's behaviour is dissimilar to that of real flies. First, it displays no tendency to steer towards the wall of the CB arena when odour is introduced (ρ_{484}=–0.022, P=0.63). Second, the increased gain of the OMR system causes a larger optomotor rebound to occur in the presence of odour (peak amplitudes: CB control: –5.05% *vs* odour: –11.20%; HS: –2.08% *vs* –4.73%), which is not witnessed in flies. We return to these points in the Discussion.

Having demonstrated that Model 2 is sufficient to reproduce the main OL effects, we must finally show that it is necessary, i.e. that no simpler model of visual–olfactory interaction could reproduce the fly behaviour.

Model 3: OMR boost

Model 3 features only the odour-dependent alteration to the OMR system. Note that since this is a step change, no use is made of the actual odour signal intensity. Illustrative results are shown in Fig. 15. No significant OL is seen in either the CB (OLI=0.343, T_{24}=127, P=0.51) or the HS arena (OLI=0.323, T_{24}=113, P=0.29). This is to be expected, since these patterns are homogeneous and therefore provide no directional cues. If the model is unable to make use of the structure of the odour plume, then clearly it has no way of locating the odour source. We note that this would apply to any other purely step-change response, such as changing flight speed, or modulating CA, in the presence of odour, unless such a response was being switched on and off rapidly as the fly encounters 'patches' of odour above some threshold, which might occur more often nearer the odour source. However, from our measurements of odour concentration, the typical flight speed of the fly, and the plausible temporal responsiveness of olfactory interneurons, we do not think this 'patchy' experience of the odour is in fact characteristic of the arena conditions.

DISCUSSION

We have implemented a visuomotor model consisting of three parallel subsystems: the optomotor response (OMR), speed regulation (SR), and expansion-based collision avoidance (CA), demonstrating that this model can in simulation provide a reasonable approximation of the free-flight behaviour observed in *Drosophila*. We identified a visually dependent 'rebound' turning response immediately following saccades, which our model reproduces. We then simulated an odour plume, and investigated various ways the resultant olfactory input could modulate the functioning of the visuomotor system. We find that making just two simple alterations according to olfactory input is sufficient to reproduce the behaviour observed in flies. One of these effects takes the form of a step change to the OMR system gain whenever odour is present. This process could be thought of as the fly entering a different behavioural state or context in order to more effectively track and approach visual targets (Chow and Frye, 2008). The other component of this cross-modal system is a modulation of the CA system on a continuous basis according to an estimate of the temporal derivative of the odour signal. This results in visually mediated chemokinesis that promotes turning when the fly loses contact with the odour plume, as has been observed experimentally (Budick and Dickinson, 2006).

While we show that these elements of the model were necessary to reproduce the fly's behaviour, it is important to stress that these findings demonstrate the minimal level of interaction between the olfactory and visuomotor systems necessary to account for odour localisation in the experimental paradigm. One discrepancy between our model and real flies is that the latter turn towards the walls of the CB arena when odour is present. This effect can in fact be reproduced by disabling the suppression of the OMR system in response to translational optic flow (Fig. 5B) in an odour-dependent manner. However, it is our feeling that this slight improvement in behavioural match does not justify the associated increase in model complexity.

Furthermore, it may well be the case that additional mechanisms exist which could be critical for performance in situations experienced by flies in the wild, but which our experimental paradigm is incapable of detecting and thus we cannot sensibly model. For instance, though we see no overall change in intersaccadic velocity when odour is introduced (CB control: $26.7\,\mathrm{cm\,s^{-1}}$ *vs* odour: $27.3\,\mathrm{cm\,s^{-1}}$; $U_{27,24}$=294, P=0.57), there is considerable evidence that flies increase their flight speed upon encountering a food odour (Frye and Dickinson, 2004; Budick and Dickinson, 2006). Increasing velocity may serve to fight against a headwind in the wild, since flies would be most likely to detect an odour downwind of its source. On the other hand, its purpose may simply be to expedite progress within the plume.

In a similar vein, tethered-flight experiments (Chow and Frye, 2008) indicate that as well as strengthening the optomotor response, odour weakens responses to expanding stimuli. The latter phenomenon is not fully captured by our model, since the CA system is attenuated only by a rising odour intensity, not a constant elevated one. Also, we modulate the threshold for triggering a CA response, rather than the amplitude of the response *per se*, as we did not observe any difference in saccade sizes for flies in the odour conditions. This is perhaps more consistent with the observation reported in Frye and Dickinson that expansion response characteristics are independent of odour (Frye and Dickinson, 2004), under the assumption that the expansion stimulus used in that study was always well above threshold.

Attraction to long vertical objects is a well known reflex in *Drosophila* (Götz, 1987; Duistermars and Frye, 2008; Maimon et al., 2008), so the fact that we see no evidence of attraction to the single dark vertical stripe that appears in our localised vertical contrast condition is worthy of comment. We consider the most reasonable explanation for this difference in behaviour to be the presence of high contrast horizontal stripes in the rest of the arena, as the aforementioned studies use backgrounds with minimal texture or lower contrast vertical stripes. As noted previously, horizontally striped surfaces have a rather disruptive effect on flight control, eliciting considerably closer approaches than those to walls with no ostensible texture whatsoever (Frye et al., 2003; Frye and Dickinson, 2007). It is possible either that a tendency to fly towards the horizontal stripes counteracts their attraction to the vertical landmark, or that the target attraction system that normally causes approach to vertical landmarks is suppressed or interfered with by the horizontal stripe background. Experiments with two diametrically opposed vertical stripes against a horizontally striped background, analogous to 'Buridan's paradigm' for walking flies (Bülthoff et al., 1982; Strauss and Pichler, 1998), also fail to elicit any clear attraction (data not shown). Finally it is also possible that some other anomaly of our experimental set-up or tracking system explains why stripe attraction was not observed.

In our simulation, rather than attempt to model explicitly why 'failure' to approach the stripe occurs, we instead take the parsimonious decision not to implement a target fixation system at all, as it is not necessary to account for the phenomona we observed. Indeed, it has previously been shown that a single rotation-sensitive visuomotor system can achieve both optomotor stabilisation and target fixation in a sparsely textured environment (Poggio and Reichardt, 1973; Huber et al., 1999). This is consistent with our observation that simply increasing the gain of the OMR system elicits stripe direction-dependent localisation of odour in our model. However, it is perhaps more likely that flies do have an additional object detection and tracking system, such as those proposed in models from Higgins and Pant (Higgins and Pant, 2004) and Hennig and colleagues (Hennig et al., 2008). If such target attraction is somehow reduced in the horizontal stripe background, an increase in the strength of target attraction when odour is present might be an alternative explanation of our results: this would improve localisation when target and odour are aligned but will interfere with localisation when they are not. This might also explain the failure to orient reliably to odour when presented with a highly salient target stimulus at another location, as seen in the study by Duistermars and Frye (Duistermars and Frye, 2008).

Visual flight control has been the subject of a number of modelling studies. Many of these adopt a methodology which is biologically inspired but does not attempt to accurately reproduce the behaviour of any particular animal (e.g. Neumann and Bülthoff, 2002; Serres et al., 2006). A robotic implementation of a model closely based on *Drosophila* behaviour is described by Reiser and Dickinson (Reiser and Dickinson, 2003). This system achieves robust wall avoidance in a textured arena, but the model lacks the other visuomotor reflexes we have implemented, and is only investigated in a single visual environment.

The responses of individual lobula plate tangential cells (LPTCs) to visual stimuli have been modelled in an impressive level of detail (Lindemann et al., 2005). However, little is known about the processing between the optic lobes and the flight motor, making the task of closing the sensorimotor loop at this level of description very difficult. A recent study (Lindeman et al., 2008) attempts to use the modelled output of a particular LPTC (the HSE-neuron) to drive both optomotor and collision avoidance behaviours in a closed-loop simulation. Whilst producing realistic behaviour under certain

visual conditions, this model is found to lack robustness with respect to variations in the textural properties of the environment.

From the work discussed above, one could perhaps conclude that because of the redundancy of visual input, the precise configuration of optic flow filters is not critical for producing reasonable flight control given a richly textured environment. Unlike the previous studies, our model accounts for the specific differences in behaviour caused by changes in the visual setting.

Other than the simple high-level model presented by Frye and colleagues (Frye et al., 2003), ours is the first closed-loop model of odour localisation behaviour of an insect flying in still air to be proposed, to our knowledge. Pheromone tracking in moths is a subject which has received considerable attention, and several models have been put forward (Belanger and Arbas, 1998; Edwards et al., 2005; Pyk et al., 2006). However, these studies consider odour tracking in windy environments. In this situation, the wind provides a directional cue that the animal can follow to approach the odour source – information which is not available in our paradigm. Nevertheless, it would be worthwhile to extend our model to windy settings. This would very probably require the addition of visuomotor anemotaxis, perhaps using non-frontally aligned expansion filters to detect and compensate for side or backward motion due to wind. It would also be desirable to investigate whether our model can be generalised to the much larger distances over which moth chemotaxis typically occurs, or if a distinct algorithm is required in such instances.

One of the more surprising results of this study is that a model where visual and olfactory reflexes run in parallel with no interaction (Model 0) displays more robust odour localisation behaviour than the fly. We have hypothesised that the reason that evolution did not settle upon a mechanism of this sort is that it would fail if the fly's flight was perturbed by external forces, i.e. wind. Extending our model to operate in moving air would allow us to properly investigate this notion. If our hypothesis is correct, it is possible that visual–olfactory integration of the type implemented in our model represents a mechanism to overcome the general problem of goal-oriented navigation in a potentially non-static fluid. If this were the case, we might expect to find evidence for similar algorithms in a wide variety of airborne and aquatic animals.

APPENDIX

Model parameters

The models described in this study make use of a number of numerical parameters. To reduce the space of possible models and to prevent overfitting, it is clearly desirable to fix the values of as many of these parameters as possible. In some cases this could be achieved based on empirical data, either from previous studies or from our own experiments. Where this could not be done, we attempted to estimate a reasonable value and keep this constant throughout our investigation. Inevitably, some parameters remained which had to be manually tuned. Unfortunately, the complexity of the model is such that an exhaustive search of the space of these free parameters is not feasible.

Table A1 shows a complete list of modelling parameters. Where no units are given, they are arbitrary. Note: visual input ranges from –128 (black) to 127 (white). The olfactory adaptive gain system attempts to keep the variance of the olfactory signal at 1 (mean is 0).

LIST OF ABBREVIATIONS

CA	collision avoidance
CB	random chequerboard
EMD	elementary motion detector
HS	horizontally striped
IR	infrared
LPTC	lobula plate tangential cell
LV	localised vertical contrast
LV-far	LV arena with odour source offset from vertical contrast

Table A1. Modelling parameters

Based on literature	
Interommatidial angle	5 deg.
Ommatidial acceptance angle (Gaussian sensitivity profile)	s.d.=1.5 deg.
EMD LPF τ	40 ms
Based on free-flight experiments	
Fly altitude	36 cm
Saccade amplitude variability (Gaussian probability distribution)	s.d.=26%
Saccade amplitude/speed linear fit	offset=1550 deg. s^{-1}, gain=–1106 deg. sm^{-1}
Speed/angular velocity linear fit	offset=100%, gain=2.5×10^4% s deg.$^{-1}$
Estimated	
EMD HPF τ	10 s
Emergency saccade distance	8 cm
LPTC leak constant	1.2×10^4 per EMD input
Signal transduction LPF τ	40 ms
Olfactory adaptive gain variance estimator τ	4 s
Olfactory adaptive gain LPF τ	9 s
Tuned	
CA and OMR leaky accumulator τ	300 ms
OMR suppression threshold	–2.0
OMR gain	10.0 per ms
SR setpoint	0.021
SR gain	0.18 per ms
CA threshold	3.8
Olfactory pre-processing LPF τs	80 ms, 320 ms
Model 2 CA modulation gain	1.07
Model 2 OMR boost amplitude	141%

EMD, elementary motion detector; LPF, low-pass filter; HPF, high-pass filter; LPTC, lobula plate tangential cell; CA, collision avoidance; OMR, optomotor response; SR, speed regulation.

LV-near	LV arena with odour source aligned with the vertical contrast
OD	olfactory signal
OD′	estimated temporal derivative of OD
OD*	OD′ following adaptive gain control
OL	odour localisation
OLI	odour localisation index
OMR	optomotor response
SR	speed regulation
VOC	volatile organic compound

ACKNOWLEDGEMENTS

We would like to thank Lincoln Smith for allowing us to use the gantry robot and assisting us in operating it, Mark Frye for supplying the flies, and the reviewers for their helpful comments. This work was funded by the UK Engineering and Physical Sciences Research Council.

REFERENCES

Belanger, J. H. and Arbas, E. A. (1998). Behavioral strategies underlying pheromone-modulated flight in moths: lessons from simulation studies. *J. Comp. Physiol. A* **183**, 345-360.

Bender, J. A. and Dickinson, M. H. (2006). Visual stimulation of saccades in magnetically tethered *Drosophila*. *J. Exp. Biol.* **209**, 3170-3182.

Berg, H. C. and Brown, D. A. (1972). Chemotaxis in *Escherichia coli* analysed by three dimensional tracking. *Nature* **239**, 500-504.

Borst, A., Egelhaaf, M. and Haag, J. (1995). Mechanisms of dendritic integration underlying gain control in fly motion-sensitive interneurons. *J. Comp. Neurosci.* **2**, 5-18.

Budick, S. A. and Dickinson, M. H. (2006). Free-flight responses of *Drosophila melanogaster* to attractive odors. *J. Exp. Biol.* **209**, 3001-3017.

Bülthoff, H., Götz, K. G. and Herre, M. (1982). Recurrent inversion of visual orientation in the walking fly, *Drosophila melanogaster*. *J. Comp. Physiol.* **148**, 471-481.

Cobb, M. and Domain, I. (2000). Olfactory coding in a simple system: adaptation in *Drosophila* larvae. *Proc. R. Soc. Lond. B. Biol. Sci.* **267**, 2119-2125.

Chow, D. M. and Frye, M. A. (2008). Context-dependent olfactory enhancement of optomotor flight control in *Drosophila*. *J. Exp. Biol.* **211**, 2478-2485.

David, C. T. (1982). Compensation for height in the control of groundspeed by *Drosophila* in a new, 'barber's pole' wind tunnel. *J. Comp. Physiol. A* **147**, 485-493.

Duistermars, B. J. and Frye, M. A. (2008). Crossmodal visual input for odor tracking during fly flight. *Curr. Biol.* **18**, 270-275.

Duistermars, B. J., Chow, D. M., Condro, M. and Frye, M. A. (2007). The spatial, temporal and contrast properties of expansion and rotation flight optomotor responses in *Drosophila*. *J. Exp. Biol.* **210**, 3218-3227.

Edwards, S., Rutkowski, A. J., Quinn, R. D. and Willis, M. A. (2005). Moth-inspired plume tracking strategies in three-dimensions. *Proceedings of the 2005 International Conference on Robotics and Automation*, 1669-1674.

Farrow, K., Haag, J. and Borst, A. (2006). Nonlinear, binocular interactions underlying flow field selectivity of a motion-sensitive neuron. *Nat. Neurosci.* **9**, 1312-1320.

Franceschini, N. (2004). Visual guidance based on optic flow: a biorobotic approach. *J. Physiol. Paris* **98**, 281-292.

Fry, S. N., Sayaman, R. and Dickinson, M. H. (2003). The aerodynamics of free-flight maneuvres in *Drosophila*. *Science* **300**, 495-498.

Frye, M. A. and Dickinson, M. H. (2004). Motor output reflects the linear superposition of visual and olfactory inputs in *Drosophila*. *J. Exp. Biol.* **207**, 123-131.

Frye, M. A. and Dickinson, M. H. (2007). Visual edge orientation shapes free-flight behavior in *Drosophila*. *Fly* **1**, 153-154.

Frye, M. A., Tarsitano, M. and Dickinson, M. H. (2003). Odor localization requires visual feedback during free flight in *Drosophila melanogaster*. *J. Exp. Biol.* **206**, 843-855.

Gilbert, C., Gronenberg, W. and Strausfeld, N. J. (1995). Oculomotor control in calliphorid flies: head movements during activation and inhibition of neck motor neurons corroborate neuroanatomical predictions. *J. Comp. Neurol.* **361**, 285-297.

Götz, K. G. (1968). Flight contol in *Drosophila* by visual perception of motion. *Kybernetik* **4**, 199-208.

Götz, K. G. (1987). Course-Control, metabolism and wing interference during ultralong tethered flight in *Drosophila melanogaster*. *J. Exp. Biol.* **128**, 35-46.

Hassenstein, B. and Reichardt, W. E. (1956). Functional structure of a mechanism of perception of optical movement. *Proceedings of the First International Congress on Cybernetics, Namur*, 797-801.

Hennig, P., Möller, R. and Egelhaaf, M. (2008). Distributed dendritic processing facilitates object detection: A computational analysis of the visual system of the fly. *PloS ONE* **3**, e3092.

Hesselberg, T. and Lehmann, F. O. (2007). Turning behaviour depends on frictional damping in the fruit fly *Drosophila*. *J. Exp. Biol.* **210**, 4319-4334.

Higgins, C. M. and Pant, V. (2004). An elaborated model of fly small-target tracking. *Biol. Cybern.* **91**, 417-428.

Holmqvist, M. H. and Svinivasan, M. V. (1991). A visually evoked escape response of the housefly. *J. Comp. Physiol. A* **169**, 451-459.

Huber, S. A., Franz, M. O. and Bülthoff, H. H. (1999). On robots and flies: Modeling the visual orientation behavior of flies. *Rob. Auton. Syst.* **29**, 227-242.

Joesch, M., Plett, J., Borst, A. and Reiff, D. F. (2008). Response properties of motion-sensitive visual interneurons in the lobula plate of *Drosophila melanogaster*. *Curr. Biol.* **18**, 368-374.

Justus, K. A., Cardé, R. T. and French, A. S. (2004). Dynamic properties of antennal responses to pheromone in two moth species. *J. Neurophysiol.* **93**, 2233-2239.

Kazama, H. and Wilson, R. I. (2008). Homeostatic matching and nonlinear amplification at identified central synapses. *Neuron* **58**, 401-413.

Krapp, H. G. and Hengstenberg, R. (1996). Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* **384**, 463-466.

Krapp, H. G., Hengstenberg, B. and Hengstenberg, R. (1998). Dendritic structure and receptive-field organization of optic flow processing interneurons in the fly. *J. Neurophysiol.* **79**, 1902-1917.

Krapp, H. G., Hengstenberg, R. and Egelhaaf, M. (2001). Binocular contributions to optic flow processing in the fly visual system. *J. Neurophysiol.* **85**, 724-734.

Kreher, S. A., Mathew, D., Kim, J. and Carlson, J. R. (2008). Translation of a sensory input into behavioural output via an olfactory system. *Neuron* **59**, 110-124.

Lehmann, F. O. and Dickinson, M. H. (2001). The production of elevated flight force compromises manoeuvrability in the fruit fly *Drosophila melanogaster*. *J. Exp. Biol.* **204**, 627-635.

Lindemann, J. P., Kern, R., van Hateren, J. H., Ritter, H. and Egelhaaf, M. (2005). On the computations analyzing natural optic flow: Quantitative model analysis of the blowfly motion vision pathway. *J. Neurosci.* **25**, 6435-6448.

Lindemann, J. P., Weiss, H., Möller, R. and Egelhaaf, M. (2008). Saccadic flight strategy facilitates collision avoidance: closed-loop performance of a cyberfly. *Biol. Cybern.* **98**, 213-227.

Maimon, G., Straw, A. D. and Dickinson, M. H. (2008). A simple vision-based algorithm for decision making in flying *Drosophila*. *Curr. Biol.* **18**, 464-470.

Maye, A., Hsieh, C. H., Sugihara, G. and Brembs, B. (2007). Order in spontaneous behaviour. *PloS ONE,* **2**, e443.

Mronz, M. and Lehmann, F. O. (2008). The free-flight response of *Drosophila* to motion of the visual environment. *J. Exp. Biol.* **211**, 2026-2045.

Murlis, J., Elkinton, J. S. and Cardé, R. T. (1992). Odor plumes and how insects use them. *Annu. Rev. Entomol.* **37**, 505-532.

Neumann, T. R. and Bülthoff, H. H. (2002). Behavior-oriented vision for biomimetic flight control. *Proceedings EPSRC/BBSRC International Workshop on Biologically Inspired Robotics* **2002**, 196-203.

Olsen, S. R. and Wilson, R. I. (2008). Lateral presynaptic inhibition mediates gain control in an olfactory circuit. *Nature* **452**, 956-960.

Poggio, T. and Reichardt, W. (1973). A theory of the pattern induced flight orientation of the fly *Musca domestica*. *Kybernetik* **12**, 185-203.

Pyk, P. I., Badia, S. B., Bernardet, U., Knüsel, P., Carlsson, M., Gu, J., Chanie, E., Hansson, B. S., Pearce, T. C. and Verschure, P. F. M. J. (2006). An artificial moth: Chemical source localization using a robot based neuronal model of moth optomotor anemotactic search. *Auton. Robots* **20**, 197-213.

Reiser, M. B. and Dickinson, M. H. (2003). A test bed for insect-inspired robotic control. *Philos. Trans R. Soc. Lond. A* **361**, 2267-2285.

Reynolds, A. M. and Frye, M. A. (2007). Free-flight odor tracking in *Drosophila* is consistent with an optimal intermittent scale-free search. *PloS ONE* **2**, e354.

Ristroph, L., Berman, G. J., Bergou, A. J., Wang, Z. J. and Cohen, I. (2009). Automated hull reconstruction motion tracking (HRMT) applied to sideways maneuvers of free-flying insects. *J. Exp. Biol.* **212**, 1324-1335.

Robert, D. and Göpfert, M. C. (2002). Acoustic sensitivity of fly antennae. *J. Insect Physiol.* **48**, 189-196.

Schuckel, J., Meisner, S., Torkkeli, P. H. and French, A. S. (2008). Dynmaic properties of *Drosophila* olfactory electroantennograms. *J. Comp. Physiol. A* **194**, 483-489.

Serres, J., Ruffier, F., Viollet, S. and Franceschini, N. (2006). Toward optic flow regulation for wall-following and centring behaviours. *Int. J. Adv. Robotic Syst.* **3**, 147-154.

Simon, H. (1982). *The Sciences of the Artificial*, 2nd edn. Cambridge, MA: MIT Press.

Single, S., Haag, J. and Borst, A. (1997). Dendrtic computation of direction selectivity and gain control in visual interneurons. *J. Neurosci.* **17**, 6023-6030.

Srinivasan, M. V., Zhang, S. W. and Bidwell, N. J. (1997). Visually mediated odometry in honeybees. *J. Exp. Biol.* **200**, 2513-2522.

Srinivasan, M. V., Poteser, M. and Kral, K. (1999). Motion detection in insect orientation and navigation. *Vis. Res.* **39**, 2749-2766.

Stavenga, D. G. (2003). Angular and spectral sensitivity of fly photoreceptors. II. Dependence on facet lens F-number and rhabdomere type in *Drosophila*. *J. Comp. Physiol. A* **189**, 189-202.

Störtkuhl, K. F., Hovemann, B. T. and Carlson, J. R. (1999). Olfactory adaptation depends on the Trp Ca^{2+} channel in *Drosophila*. *J. Neurosci.* **19**, 4839-4846.

Strauss, R. and Pichler, J. (1998). Persistence of orientation toward a temporarily invisible landmark in *Drosophila melanogaster*. *J. Comp. Physiol. A* **182**, 411-423.

Sugiura, H. and Dickinson, M. H. (2009). The generation of forces and moments during visual-evoked steering maneuvers in flying *Drosophila*. *PLoS ONE* **4**, e4883.

Tammero, L. F. and Dickinson, M. H. (2002a). Collision-avoidance and landing responses are mediated by separate pathways in the fruit fly, *Drosophila melanogaster*. *J. Exp. Biol.* **205**, 2785-2798.

Tammero, L. F. and Dickinson, M. H. (2002b). The influence of visual landscape on the free flight behavior of the fruit fly *Drosophila melanogaster*. *J. Exp. Biol.* **205**, 327-343.

Tammero, L. F., Frye, M. A. and Dickinson, M. H. (2004). Spatial organization of visuomotor reflexes in *Drosophila*. *J. Exp. Biol.* **207**, 113-122.

van Hateren, J. H. and Schilstra, C. (1999). Blowfly flight and optic flow. II. Head movements during flight. *J. Exp. Biol.* **202**, 1491-1500.

Warzecha, A.-K. and Egelhaaf, M. (1996). Instrinsic properties of biological motion detectors prevent the optomotor control system from getting unstable. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **351**, 1579-1591.

Webb, B. (2004). Neural mechanisms for prediction: do insects have forward models? *Trends Neurosci.* **27**, 278-282.

Zanker, J. M., Srinivasan, M. V. and Egelhaaf, M. (1999). Speed tuning in elementary motion detectors of the correlation type. *Biol. Cybern.* **80**, 109-116.