

What can be learnt from analysing insect orientation flights using probabilistic SLAM?

Bartholomew Baddeley · Andrew Philippides ·
Paul Graham · Natalie Hempel de Ibarra ·
Thomas Collett · Phillip Husbands

Received: 10 February 2009 / Accepted: 7 July 2009 / Published online: 29 July 2009
© Springer-Verlag 2009

Abstract In this paper, we provide an analysis of orientation flights in bumblebees, employing a novel technique based on simultaneous localisation and mapping (SLAM) a probabilistic approach from autonomous robotics. We use SLAM to determine what bumblebees might learn about the locations of objects in the world through the arcing behaviours that are typical of these flights. Our results indicate that while the bees are clearly influenced by the presence of a conspicuous landmark, there is little evidence that they structure their flights to specifically learn about the position of the landmark.

Keywords Orientation flights · SLAM · Navigation · Bumblebees

1 Introduction

1.1 Orientation flights in bees and wasps

When bees and wasps first leave their nests or a productive food source they perform a series of stereotypical flight manoeuvres referred to as orientation flights (von Frisch 1967; Becker 1958; Vollbehre 1975; Zeil and Kelber 1991; Zeil 1993a; Capaldi et al. 2000), or Turn Back and Look

(TBL) flights (Lehrer 1991, 1993; Lehrer and Collett 1994)¹. Immediately following take off they turn to face the place of interest and fly a series of arcs of increasing distance roughly centred on the goal. It is known that disruption of these flights affects the bees' ability to successfully relocate the goal and that as the bees become more familiar with their visual environment they spend less time performing these orienting behaviours until eventually they fly almost directly away (Lehrer 1993). Experienced bees will however re-initiate orientation flights if they have trouble finding the goal (Iersel and Assem 1964; Wehner 1981).

Many previous authors have made the observation that the sideways arcing movements that are typical of these flights would be suitable for measuring depth through motion parallax (Cheng et al. 1987; Lehrer 1991; Lehrer and Collett 1994; Collett 1995; Collett and Zeil 1997; Zeil 1993a,b). Lehrer and Collett (1994) investigated what depth cues bees learned during their approach to a feeder and during their TBL flights when leaving. Bees were trained to collect a reward from a feeder that was located in a fixed position relative to a cylinder. In three experimental conditions, bees were able to view the cylinder: only on their approach to the feeder; only on their departure; or, during both approach and departure. Probe tests demonstrated that bees that viewed the cylinder only on arrival learned its apparent size whereas bees that viewed the cylinder on departure learned the absolute distance. Bees that were able to view the cylinder on arrival and departure learned both apparent size and absolute distance, with absolute distance providing the primary cue during the

B. Baddeley (✉) · A. Philippides · P. Graham ·
T. Collett · P. Husbands
Department of Informatics, Centre for Computational
Neuroscience and Robotics, University of Sussex,
Brighton, UK
e-mail: bartbaddeley@googlemail.com

N. H. de Ibarra
School of Psychology, University of Exeter, Exeter, UK

¹ Jander (1997) characterises orientation flights as having two distinct phases. The first phase is generally near the nest and composed of small slow movements and is followed by larger spirals around the nest. For the remainder of this paper any reference to orientation flights will relate to the small scale TBL behaviours that the bees perform on first leaving their nest.

initial phase of learning and absolute size becoming more important later on.

These observations beg the question of how visual information is used during these flights. The results discussed above clearly implicate TBL flights in learning about the absolute depth of landmarks. Here we investigate whether flights are structured to optimally extract depth information from the environment. Holding such metric information might enable the development of a metric representation of the spatial relationships between a discrete set of salient features or landmarks within a geo-centric or world-based frame of reference, a map. Whether or not insects solve spatial problems in this way has generated a great deal of controversy over the years and is still regarded as an open question (Menzel et al. 2005; Gallistel 1989; Gould 1986; Wehner and Menzel 1990; Wehner et al. 2006). Much of the debate centres on the cognitive demands and the utility of learning such map-like representations given the short lifespan and limited neural resources available to insects. In contrast, we approach the issue by examining whether the stereotypical arcing manoeuvres observed during TBL flights are optimised for learning the metric information required for a map-like representation of the world.

TBL flights potentially provide cues to depth in at least two ways. First, the movements produce patterns of optic flow across the whole visual field that allow landmarks to be ranked according to depth (Cheng et al. 1987; Zeil 1993b). Alternatively, the positions of a discrete set of visual features or landmarks could be learned without the direct use of optic flow. By tracking the path of individual features across the retina and combining this information with information about ego-motion, an estimate of the position of each of the features could gradually be built up. Each time a feature is viewed, the estimate of its position could be updated and improved providing an increasingly coherent spatial representation of the world. Maintaining such a representation would allow the bee to re-orient upon viewing familiar landmarks in the vicinity of the goal. This is consistent with the observation that experiencing a prominent landmark en route narrows the search distribution of returning bees (Srinivasan et al. 1997).

1.2 The Bayesian brain

There is growing evidence that many aspects of sensorimotor behaviour in animals can be closely modelled in terms of Bayesian inference and estimation, a framework that can handle the inherent uncertainty and noise in the natural world. For instance, many authors argue that a framework based on Bayesian inference is highly suited to modelling and understanding vision, as it is capable of dealing with the complexities and ambiguities of natural images while accounting for fundamental perceptual tasks such as recognition (e.g. Yuille and Kersten 2006; Lee and Mumford 2003). More

specifically, Cheng et al. (2007) review numerous instances of integration of spatial cues in diverse species, including insects, and show that the ways in which the cues are combined closely follows a Bayesian interpretation of perception. Courville et al. (2006) show how a Bayesian approach provides a principled interpretation of conditioning in animals in a changing world, specifically the finding that surprising events provoke animals to learn faster. Kording and Wolpert (2006) show that action selection in human motor behaviour is close to that predicted by Bayesian decision theory, which defines optimal choice in a world characterised by uncertainty. The success of such models has led some to speculate that brains possess neural circuitry that does something very close to estimating probability distributions (Knill and Pouget 2004); indeed Rao has demonstrated that small recurrent networks of noisy integrate-and-fire neurons can perform approximate Bayesian inference (Rao 2004). For further details of such approaches see the wide ranging review in Griffiths et al. (2008).

The success of such Bayesian interpretations suggest that *if* the bees in the study described here are learning the position of the landmark as part of some metric representation of the environment, it is reasonable to hypothesise that they are doing so in a way that can be closely modelled in terms of Bayesian estimation. If the flights are structured in such a way as to learn positional information, a probabilistic analysis should reveal changes in uncertainty estimates associated with the landmarks. As described in the next section, we have adapted SLAM, a powerful probabilistic framework developed in mobile robotics, to provide quantitative tools to study bee flights within a Bayesian framework. Because the probabilistic robotics framework, which was specifically developed to deal with closed loop sensorimotor behaviour, is so general, it can be readily adapted to the study of insect behaviour.

1.3 Probabilistic robotics and SLAM

Over the past fifteen years or so, approaches based on probabilistic inference have become prominent in mobile robotics, particularly in map building and navigation. The mapping problem involves a period of exploration during which a robot builds a map of its environment that can then be used for accurate navigation. A related problem is that of localisation—the ability of a robot to determine where it is, relative to a map, from its sensor readings. Most work in this area has concentrated on the simultaneous localisation and mapping (SLAM) problem (Smith et al. 1990; Dissanayake et al. 2001; Durrant-Whyte and Bailey 2006), also sometimes known as the concurrent mapping and localisation problem (Thrun et al. 1998). This requires a mobile robot, when placed at an unknown spot in an unknown environment, to incrementally construct a consistent map of the environment at

the same time as determining its location in the map. A great deal of progress has been made and for certain types of environments good solutions to the problem have been found (Durrant-Whyte and Bailey 2006; Milford and Wyeth 2008; Cummins and Newman 2008). Nearly all these solutions rely on probabilistic models of the robot and its environment, and employ probabilistic inference in building maps from the robot’s sensor readings.

The success of the probabilistic approach stems from the fact that the mapping problem is inherently uncertain and robot sensors are noisy, as is robot movement. The probabilistic approaches embrace these characteristics of the problem rather than ignoring them or trying to hide them. The overall approach is to use recursive Bayesian methods to build up estimates of unknown probability density functions over time using incoming sensor measurements and a mathematical process model as outlined below. In probabilistic terms the SLAM problem requires the following probability distribution to be computed for every time t .

$$P(\mathbf{x}_t, |Z^t, U^t, \mathbf{x}_0) \tag{1}$$

where the vector \mathbf{x}_t represents the system state describing both the robot (robot position and orientation etc.) and pertinent information about the environment (usually landmark locations, but more complex metric relationships describing various aspects of the geometrical layout of the environment are also sometimes used), Z^t and U^t represent all sensor readings, z_t , and motor controls, u_t , from time $t = 0$ until the present: $Z^t = \{z_0, z_1, \dots, z_t\}$, $U^t = \{u_0, u_1, \dots, u_t\}$. Thus equation 1 is the conditional probability density of the system state, given the recorded sensor inputs and the motor controls, along with the initial position of the robot, \mathbf{x}_0 . Using Bayes Theorem it is possible to recast this distribution in terms of the following recursive equation (Thrun et al. 2005):

$$P(\mathbf{x}_t|Z^t, U^t, \mathbf{x}_0) = \eta P(z_t|\mathbf{x}_t) \int P(\mathbf{x}_t|\mathbf{x}_{t-1}, u_t) \dots P(\mathbf{x}_{t-1}|Z^{t-1}, U^{t-1}, \mathbf{x}_0) d\mathbf{x}_{t-1} \tag{2}$$

where η is a normalising constant. This equation involves a probabilistic motion model, $P(\mathbf{x}_t|\mathbf{x}_{t-1}, u_t)$, which assumes a Markov process in which \mathbf{x}_t depends only on the previous position, \mathbf{x}_{t-1} , and the applied motor control, u_t , as well as a probabilistic sensor observation model, $P(z_t|\mathbf{x}_t)$, which describes the probability of making sensor observations when the robot location and the map are known. Solutions to the SLAM problem involve finding appropriate representations for the motion and sensor observation models such that the various probability densities can be calculated via efficient recursive procedures. The most common approaches use an extended Kalman filter (EKF) (Guivant and Nebot 2001) or a more general particle filter (Montemerlo et al. 2002).

1.4 Outline of the paper

In this paper, we use the SLAM framework in a different way. By running a SLAM simulation with the recorded flight trajectories as the movements that are made and using a model of a bee’s visual system as a sensor observation model together with a very general probabilistic motion model, we can examine which parts of the local environment would be mapped most effectively. Here, probabilistic SLAM becomes a tool for analysing noisy behavioural data to investigate whether the structure of TBL flights is consistent with bees learning a metric representation. To use this approach we simulate the learning of a map given particular flight structures. The uncertainty values associated with entries in the map provide a direct measure of the efficacy of the flight, or parts of the flight, for learning about that position in the map. Note that this analysis does not depend on bees using such maps. It simply provides a method for analysing the informational content of the bee’s movements.

We have recorded trajectories of bumblebees performing orientation flights on leaving their nest in the presence of a single conspicuous landmark in order to examine what the bees might be attending to and learning. We analyse these flights in terms of generated optic flow and their potential for learning metric information about the positions of visual features. If bees are attempting to learn metric information about the landmark, we would expect the structure of the flights to show some bias towards learning about the depth structure of the world at or near to the position of the landmark.

In the next section, we describe the SLAM approach to learning and how we acquired the flight data to which we have applied our analyses. Finally, we report our results and conclude with a discussion of their implications.

2 Methods

2.1 A probabilistic solution to the SLAM problem

The essence of the SLAM approach for a visually guided agent is to estimate the current state of the system (i.e. positions etc. of the agent and all landmarks in the world) using a two-step procedure which importantly also attempts to quantify the uncertainty of each estimate. First, the state estimate, together with a covariance matrix that reflects the uncertainty in the estimate, is propagated forward in time using a process or movement model that defines how the agent’s state (position, velocity etc.) changes in response to a control input. Following this, a (noisy) measurement of all visible objects is made and used to update the state and covariance estimates once again. Here, both update steps are implemented through an EKF, the most common approach for non-linear state estimation.

The process and sensor models are therefore crucial and their accuracy determines how the uncertainty in the system evolves. In the first stage, uncertainty in the agent’s position will increase, due to inaccuracies and noise in the process model while state and covariance estimates of object positions are unchanged as they are unaffected by agent movement. In contrast, the measurement phase acts to reduce the uncertainty in the entire state estimate, with reduction in uncertainty being determined by the accuracy of the measurements.

2.1.1 The state and covariance

The state \hat{x} of the system is a vector containing both the state of the agent \hat{x}_v and the positions of all tracked features \hat{y}_i . In our current implementation, we chose to track a single feature in each of our simulations. By changing the position of this one feature we are able to examine how the recorded flight trajectories affect learning of all possible fictive object positions within a region of interest (ROI) defined as the area viewed by the camera. The state vector is paired with a covariance matrix, P , partitioned as follows:

$$\hat{x} = \begin{pmatrix} \hat{x}_v \\ \hat{y}_1 \\ \hat{y}_2 \\ \vdots \end{pmatrix}, \quad P = \begin{pmatrix} P_{xx} & P_{xy1} & P_{xy2} & \dots \\ P_{y1x} & P_{y1y1} & P_{y1y2} & \dots \\ P_{y2x} & P_{y2y1} & P_{y2y2} & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (3)$$

where P_{ij} represents the covariance between the state variables i and j . The choice of representation and the frame of reference of the state are both arbitrary.

2.1.2 The process model

The state and covariance are updated following a movement of duration Δt using the following equations:

$$\hat{x}_{(t+1)} = f(\hat{x}_{(t)}, u_t, \Delta t) \quad (4)$$

$$P_{(t+1)} = \frac{\partial f}{\partial x} P_{(t)} \frac{\partial f^T}{\partial x} + Q_t \quad (5)$$

where, f is a differentiable state transition function, that maps the state estimate, \hat{x}_t , and control inputs, u_t , at time t , into a subsequent state \hat{x}_{t+1} . A matrix describing the process noise Q_t , together with the Jacobian of the state transition function ($\frac{\partial f}{\partial x}$) is used to update the covariance matrix P .

The process noise is intended to account for any un-modelled movements and is given by:

$$Q_t = \frac{\partial f}{\partial u_t} U_t \frac{\partial f^T}{\partial x_t} \quad (6)$$

where U_t is the diagonal covariance matrix of u_t .

2.1.3 The measurement model

As with the process model, in the EKF framework the measurement model need not be linear but must be differentiable. The measurement model provides a model of the sensor array and allows the prediction of the expected sensor activations given the current state estimate. For example, assuming a very simple ray tracing model of the visual system whereby a measurement can be made if there is a direct line of sight to the measured feature. Then the measurement model provides a prediction of the direction to the feature given the current estimates of the agent’s position and orientation and the position of the feature. In the case of a simple ray tracing model of vision, this calculation involves straightforward trigonometry.

As well as a point prediction of the measurement h_i , the measurement model allows us to calculate the innovation covariance matrix S_i . The innovation covariance matrix S_i represents the expected uncertainty in measurement h_i and is given by:

$$S_i = \frac{\partial h_i}{\partial x_v} P_{xx} \frac{\partial h_i^T}{\partial x_v} + \frac{\partial h_i}{\partial x_v} P_{xyi} \frac{\partial h_i^T}{\partial y_i} + \frac{\partial h_i}{\partial y_i} P_{yix} \frac{\partial h_i^T}{\partial x_v} + \frac{\partial h_i}{\partial y_i} P_{yiyi} \frac{\partial h_i^T}{\partial y_i} + R \quad (7)$$

where, $\frac{\partial h_i}{\partial x_v}$ and $\frac{\partial h_i}{\partial y_i}$ are the Jacobian matrices of the measurement model with respect to x_v and y_i , respectively, P_{xx} , P_{xyi} , P_{yix} and P_{yiyi} are sub-matrices of P , and R is the measurement noise covariance that describes the accuracy of the measurements.

The accuracy with which measurements can be inferred is limited by the accuracy with which estimates of the bee’s position and orientation can be made using the video recordings. Therefore, in this instance the measurement noise reflects both the limited acuity of the bee’s visual system and the limited accuracy with which positional information about the bee can be determined from the video data.

2.1.4 Updating the state

Following a measurement z_i , the Kalman gain, W , is calculated and used to update the state and covariance estimates using:

$$W = P \frac{\partial h_i^T}{\partial x} S^{-1} \quad (8)$$

$$\hat{x}_{t+1} = \hat{x}_t + W(z_i - h_i) \quad (9)$$

$$P_{t+1} = P_t - W S W^T \quad (10)$$

3 The probabilistic framework applied to orientation flights

As we stated earlier we do not suggest that bees necessarily form the sort of maps that are used in SLAM. However, the approach provides a tool to quantify the information about the environment that a bee might extract on a given flight. By replacing the measurement model that describes a camera's optics with one that describes the optics of a bee, and similarly, replacing the process model with a motor model that describes the bee's movements we can construct a closed-loop system for investigating active vision. The SLAM simulation using simple models of the bees' sensory and motor capabilities was run in parallel with the transcribed video data.

For clarity, it is worth restating that in our simulations we only ever track a single feature and that the position of this feature need not coincide with the true position of the landmark. By changing the position of a fictive feature and inferring where on the bee's retina its image would fall we are able to examine how the recorded flight trajectories affect learning. In the uncertainty maps of Figs. 5 and 7 each pixel represents the final uncertainty in the positional estimate of a fictive landmark at that pixel location following the simulated flight, given that the landmark was visible.

For the sensor model we assume that the bee is able to observe a given feature if it falls within the field of view and there is a direct line of sight to it. We further assume that the measurement noise and therefore the visual acuity is fixed. For the majority of our analyses we assume a functionally uniform retina, meaning that measurements are treated equally irrespective of where on the retina they are made. In a later analysis we look at the effect of restricting the input to the frontal retina. The motor model that we employ assumes that we do not have access to the control input. Instead it is assumed that all forces and torques that act on the bee are small and normally distributed around zero (Davison and Murray 2002) so at each time-step the velocity is assumed to remain constant while the covariance of the state estimate increases. This very simple and general motor model is sufficient to implement SLAM and allow our analysis of the orientation flights.

Setting the correct level for the process and measurement noise is key to the successful implementation of probabilistic SLAM. The measurement noise should represent as closely as possible the true accuracy of the measurements that are made of the world and the process noise should reflect, again as closely as is possible, the accuracy of the process model. Setting these values too high will result in slow convergence of the map estimates resulting in poor performance. Setting them too low can result in catastrophic failure as the system converges too quickly to an incorrect solution.

In order to set the value of the process noise covariance, Q , we assumed a zero mean impulse model and measured the squared difference between the model predictions and the observed movements. The variance of the measurement noise, R , needs to be able to account for both the limited visual acuity of the bees [$\approx 5^\circ$ (Spaethe et al. 2001)] and the limited accuracy with which it is possible to determine the position and orientation of the bees from the video recordings ($\pm 5^\circ$). R was therefore set to approximately 5 degrees.

One aspect of the SLAM problem that we do not address is the issue of initialising the map estimates. In a standard SLAM implementation operating in the real world deciding when and how to initialise a new feature into the map is a non-trivial problem. Since it is not possible to determine the distance to a feature given a single view of it, it is difficult to initialise the map entry to a sensible value. Most approaches have a separate initialisation procedure whereby a feature will be tracked for a few frames prior to the feature actually being added to the map. In our simulations we choose to initialise the map estimates at their correct value and to set the uncertainty in the landmark estimate to be high (100) and the uncertainty in the bee estimate to be low (1.0). This is intended to reflect that the bee knows the position of itself relative to the nest when it exits the nest but does not know the position of the landmark.

Using this probabilistic framework we can follow orientation flights, tracking how the structure of a flight reduces the uncertainty in the bee's estimate of its own position and of visual features within the environment. In Fig. 1a, these uncertainties are represented as ellipses around the landmark and the bee's positions. As the flight proceeds the uncertainty in the positional estimate of the bee increases while the uncertainty in the landmark positional estimate decreases as measurements of the landmark are made. This can be seen between $t = 0.78$ and $t = 1.5$ s, following measurements of the landmark which collapse the uncertainty perpendicular to the direction of viewing, and between $t = 2.96$ and $t = 3.7$ s, where the bee moves perpendicularly to the line of sight to the landmark, reducing uncertainty in the direction parallel to the line of sight.

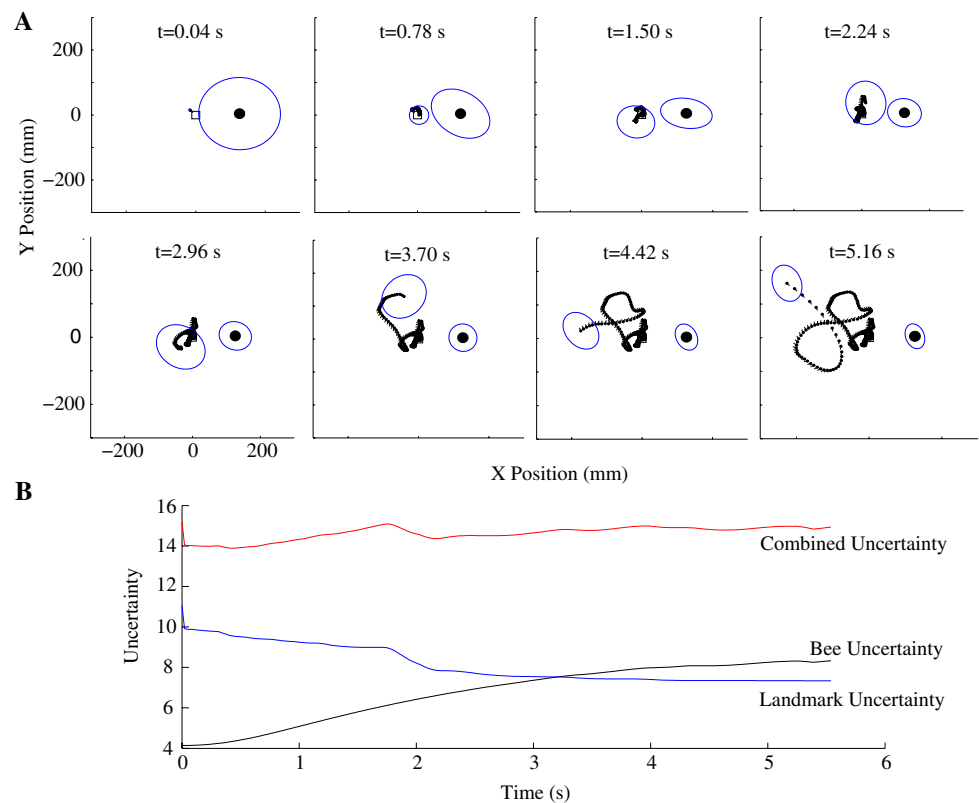
3.1 Measuring uncertainty

As Fig. 1b shows, the total uncertainty in the system, measured in terms of the entropy, H , of the covariance matrix P ;

$$H = \frac{1}{2} \log(2\pi^d |P|) \quad (11)$$

is divided between the bee and the landmark estimates, where d is the dimension of the state estimate and $|P|$ is the determinant of the covariance matrix.

Fig. 1 a An example of how the estimated uncertainty in a SLAM based simulation of a real flight evolves during the course of a flight. Uncertainty in the landmark positional estimate and the estimate of the bee's position are represented by ellipses and shown at several instances throughout a sample orientation flight. All distances are measured in mm and the timestep was 0.02 s. The nest (*open square*) is at (0,0) and the landmark (*filled circle*) is at (125,0). **b** The *bottom panel* shows the evolution in uncertainty over the course of the flight shown in **a**. The *top line* shows how the uncertainty of the whole system develops, the *middle line* shows the uncertainty in the estimate of the landmark position and the *bottom line* shows the bee's positional uncertainty.



The goal in a standard robotics implementation of SLAM is to reduce the uncertainty in the whole system. However, for our purposes, the total uncertainty might not be appropriate. We therefore look at how both the landmark's and the observer's uncertainties evolve during the course of a flight in order to determine the correct measure to use in our analyses.

The bee's positional uncertainty increases monotonically and smoothly during the flight (bottom line Fig. 1b). The only deviation from this pattern occurs towards the later portions of the flight when close inspection reveals occasional small reductions in uncertainty. The increase in uncertainty reflects any noise effects or inaccuracies in our process/motor model and the small reductions in uncertainty relate to sections of the flight when the bee is able to re-orient itself relative to the landmark.

The estimated landmark uncertainty shows a different pattern, starting high and reducing at a fairly constant rate. The reduction in the landmark uncertainty due to each measurement is determined by four things, namely the prior level of uncertainty in the landmark estimate, the level of uncertainty of the bee's positional estimate, the accuracy of the measurement (determined by the measurement noise), and finally on the movements that are executed by the bee. The initial rapid reduction in uncertainty is mainly due to the bee being very certain of its own position and to the relatively large move-

ment across the retina of the landmark due to its close proximity. As the flight progresses the reductions in uncertainty become smaller. This is due to the increased uncertainty in the bee's own positional estimate and to the smaller relative movements that occur at greater distances.

If bees do map salient visual features in a way that can be modelled using SLAM, and use such maps to re-orient when features come back into view, then it is the estimated landmark uncertainty that is the correct measure of learning efficacy to use in our analyses. The reasoning is as follows. Following a learning flight the bee leaves the immediate nest locale for its foraging flight. During this flight its positional uncertainty will increase due to inaccuracies in estimating ego-motion. In contrast the map entries for any landmarks learned during the learning flight will remain unchanged while out of view. If the learning was effective then the estimated uncertainty will be highly correlated with the accuracy of the map entry. In this way the uncertainty in a landmark estimate should give a measure of how effectively the bee has been able to learn the position of the landmark and consequently how well the bee will be able to re-orient when the landmark comes back into view. Given this reasoning, it is the uncertainty in the landmark estimate that proves relevant for our analyses, since this measure reflects how suitable a given set of movements were for learning the position of a particular landmark. It is this measure of uncertainty in the

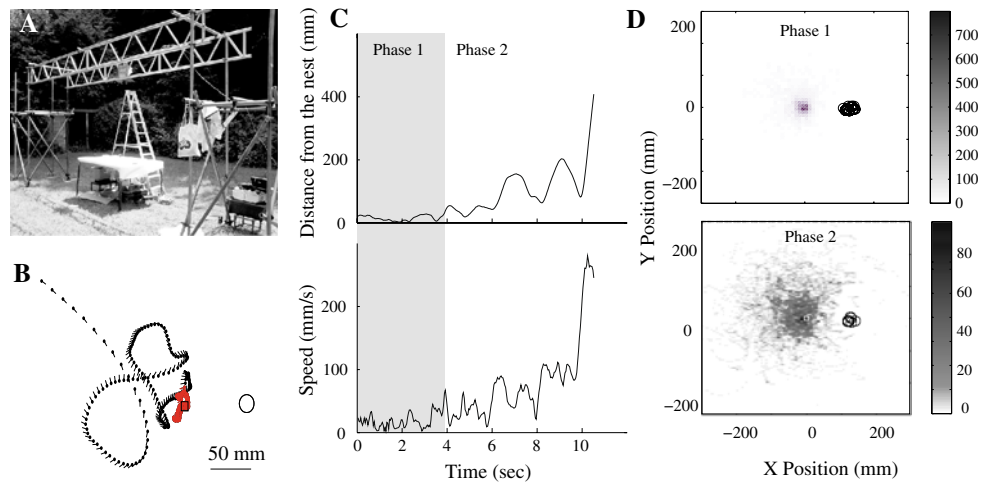


Fig. 2 **a** A commercially available bumblebee hive is positioned underneath a featureless but textured 180×150 cm board containing a single conspicuous landmark in the form of a small black cylinder. The bees enter and exit through an inconspicuous entrance that is positioned a small distance from the landmark. A camera was positioned on a frame and pointed down on the board in order to record the position and orientation of the bees during the initial phase of their foraging trips. **b** An example of a bumblebee flight recorded with this set-up. The ball and stick icons represent the head position and body axis of the bee at 20ms intervals. The black circle represents the position of the 5 cm diameter landmark. All distances are measured in mm. The nest (*open*

square) is at (0,0) and the landmark (*open circle*) is at (125,0). **c** The upper and lower traces represent the bee's distance away from the nest entrance and the bee's speed, respectively. These parameters were used (see Sect. 2) to determine the end of the first phase of bees' flights (*red markers* in **b**). **d** Density plots showing the distribution of bee positions across all 37 flights for the first (*top*) and second (*bottom*) phases. The scale indicates the overall number of time-steps for which a bee was present at a given position in the ROI. The nest (*open square*) is at (0,0) and the landmarks (*open circles*) were at approximately (125,0). All distances are measured in mm.

positional estimate of the landmark we use throughout the rest of this paper.

3.2 Acquiring behavioural data

The data consist of 37 separate recordings of bees exiting an inconspicuous nest entrance that is positioned on a featureless but textured 180×150 cm board containing a single conspicuous landmark in the form of a small black cylinder. A camcorder (Sony HDR HC7E) was suspended from scaffolding oriented along a WSW and ENE axis 2 m above the table (Fig. 2a). A separate tape-deck or second camcorder fed by the camcorder above the nest-hole served as a data recorder. The camcorder recorded at 25 frames per s. Frames were split to avoid interlacing and to improve the temporal resolution to 20 ms. The information from digital tapes was transferred to a hard drive using Adobe Premiere Pro.

Orientation flights are performed over the first 5–10 departures of a foraging bee from its nest. We recorded the arcing behaviours that the bees performed up to the point when they flew out of the camera's field of view. The field of view was $1,020 \times 698$ pixels corresponding to an approximate ROI of 1,000 by 700 mm. Results figures throughout use this ROI. Software written in MatLab extracted a bee's horizontal coordinates in terms of the centre of mass of the bee's image. The program also determined the compass orientation in which

the bee faced. The orientation of the body axis is given by the major axis of the bee's image with the head end identified from both the shape of the image and the direction of movement between frames. The program allowed the computed values to be checked and when necessary adjusted by hand. The accuracy of the positional data was ± 1 mm and the orientation data was $\pm 5^\circ$. A typical example of a flight is shown in Fig. 2b. We only have horizontal positional information as height was not monitored, however, we observed that the bees fly close to the table for most of each learning flight. Height is only gained towards the end of the recorded flight segments.

Most orientation flights follow a similar pattern. An initial phase where bees are very close to the nest entrance and moving slowly is followed by the bee moving away from the nest entrance and commencing the characteristic arcing behaviour. To ensure our results were not dominated by the initial phase where the bee is very close to the nest, we performed our analyses on the whole flights and on each phase separately. As there were some occasions when the bee continued the first phase behaviour of slow flight close to the nest after a brief foray away from the nest, we could not simply decide the phases based on distance to the nest. We therefore defined the end of the first phase as the point where the distance from the nest multiplied by the squared velocity exceeded 500. This ad hoc approach divided the flights into

distinct behavioural sections, in which the second phase was not dominated by the bee adopting positions close to the nest. Figure 2c shows the velocity and distance traces for the flight shown in Fig. 2b. The first phase is marked by red icons in Fig. 2b and a grey area in 2c.

3.3 Which parts of the world do the bees view?

For our analyses we need to know the extent of the bumblebees' horizontal visual field. For each frame of the video we took the bee's position and horizontal body axis orientation and used this to infer which parts of the environment were in view, which were being occluded by the bee's own body, and which parts were being viewed binocularly. Measurements are not available for the frontal binocular region and for the posterior occluded region in bumblebees and we have taken values based on the honeybee (Seidl and Kaiser, 1981) adjusted for the larger body of the bumblebee. We therefore use a value of 30° for the frontal binocular region and 55° for the posterior occluded region.

Each pixel in the image was scored according to the following criteria; 0 points for occlusion and 1 point for being in view and 2 points for pixel positions that were viewed binocularly². A running total was maintained for each pixel over the course of a flight. The final totals were used to create a 2D frequency histogram that showed which parts of the environment were viewed most often. We performed similar analyses using only the frontal 20° as the region in view and also without treating binocular regions differently. These analyses produced qualitatively similar results to those presented and are therefore not reported here.

3.4 What optic flow is generated by flights?

From the bee's position and velocity we can infer the optic flow that would be generated by a visual feature at a given position. For each time-step (t) we calculated the change in retinal position, from time $t - 1$, for hypothetical visual features at every position within our ROI. Since any optic flow produced by the bee's own rotations will be uniform across the entire retina it will provide no cues to depth. High speed and high magnification recordings show that wasps segregate translational and rotational movements during orientation flights. They punctuate spells of translation when they hold a constant body orientation with brief saccadic turns (Zeil et al. 2007). With this in mind we calculated both the total optic flow and the optic flow that resulted from purely translational movements. We then considered two different ways of measuring the optic flow that would be produced by

each point in space, given the recorded flight trajectories. We first considered which points in the world consistently generated perceptible optic flow. As a measure of consistency we keep a cumulative score for each position representing the total number of frames where a landmark at that location would generate optic flow less than 100°s^{-1} . This upper bound is set by the blur velocity of the bee visual system, which is dependent on the spatial and temporal resolution of the compound eye (Land 1999). We set no lower bound on what constitutes behaviourally relevant optic flow since any signal that can be measured should provide some information. We also measured the total magnitude of optic flow by summing the optic flow signal across all time-steps in which the optic flow was less than 100°s^{-1} . This resulted in four different measures of optic flow.

4 Results

4.1 Visual consequences of flight structure

In accordance with previous descriptions of the orientation flights of flying insects, bumblebees in this set of flights tended to face the landmark and nest whilst gradually backing away from them performing arcs of increasing radius. Consequently, the area of the environment viewed most often is a triangular region emanating from the nest roughly symmetrical about the landmark (Fig. 3). The results for three typical flights, having short, medium, and long durations, and summary data for all flights show the same general structure. Additionally, we observed that orientation flights seemed to be composed of two phases (Fig. 2c, d). In an initial phase, the bees remained close to the nest, moved slowly often rotating as if to scan the entire visual field. During the second phase bees flew a series of arcs, with both speed and distance from the nest gradually increasing. As these phases could be functionally distinct, here and throughout we have augmented our analysis of the entire flights by also examining these phases independently. When considering which parts of the environment are viewed most often, the results are qualitatively similar for both flight phases (Fig. 3b, c). The major difference comes from the increased arc lengths later in the flights when views tend to be more focussed on the region between nest and landmark.

The bees flights are thus structured so that the area between nest and landmark is viewed most often. If however, viewing this area were all the bees were trying to achieve then it is unlikely that we would see any other structure in the data. We next examine the optic flow generated by the flights.

Figure 4 shows two different measures of optic flow, total magnitude and overall consistency, for movements with and without rotations. When translational and rotational movements are considered together the two measures of optic flow

² In a noiseless system there is no benefit to making multiple measurements of the same point in space, the situation changes however when we consider noisy measurements.

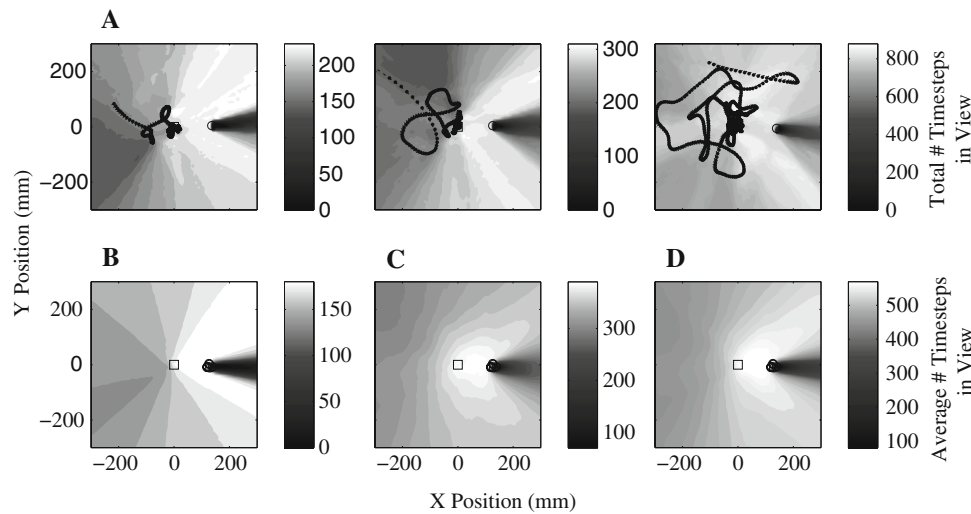
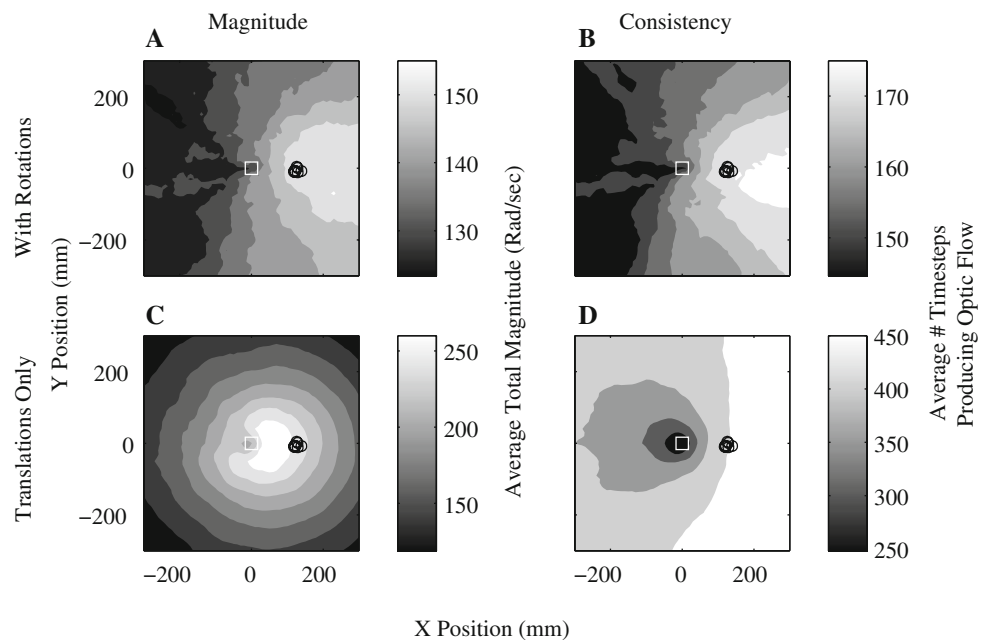


Fig. 3 a Density plots of the frequency with which each location in the environment is viewed during single short, medium and long orientation flights (left to right, respectively). Bee position and orientation are shown every 20 ms as ball and stick. All distances are measured in mm. The nest (open square) is at (0,0) and the landmark (open circle) is at (125,0). Lighter regions are viewed most often. Scales vary due

to sample size differences. **b–d** Combined data for 37 flights showing first phase, second phase and entire flight, respectively (see Sect. 2 for description of phases). Landmark positions (black circles) varied slightly between recordings. Again lighter regions are viewed most often

Fig. 4 Four different measures of optic flow. Combined data for 37 flights. **a, b** Magnitude and consistency of optic flow for combined rotational and translational movements. **c, d** Magnitude and consistency of optic flow for purely translational movements (see text for details). All distances are measured in mm. The nest (open square) is at (0,0) and the landmark (open circle) is at (125,0). Landmark positions varied slightly between recordings



produce broadly similar distributions (Fig. 4a, b). We observe a clear bias towards the area containing the landmark, indicating that this region produces both a consistently perceptible and large optic flow signal. However, when we consider only translationally induced optic flow as might be produced during the translational phase of a body saccade, the two measures produce qualitatively different distributions (Fig. 4c, d). Measuring the magnitude of the signal results in a symmetrical distribution centred on a region midway between the

nest and the landmark (Fig. 4c). In contrast, measuring the consistency of the signal results in a minimum at the nest position, with regions behind the landmark consistently generating perceptible flow.

Considering the implications of these results for how well each flight is structured for learning the landmark position begs the questions of what optic flow bees use, how they use it and how it might be extracted. Zeil et al. (1996) observed that wasps seem to keep the image of the area around the

nest as stationary as possible during their learning flights. This is broadly consistent with Fig. 4a and b. In a subsequent paper, Voss and Zeil (1998) show how behavioural routines that combine rotations and translations could facilitate the extraction of depth information by performing movements that produce informative patterns of optic flow. The patterns of optic flow produced by our flights are dependent on which measure we use. This highlights the difficulty of interpreting optic flow: is high or low optic flow useful? Determining and justifying what constitutes behaviourally relevant optic flow and the difficulty in deciding how to measure its informational content provides motivation for our use of probabilistic SLAM.

4.2 Analysis of flight structure using probabilistic SLAM

There are two important features of our probabilistic SLAM model. First, as the agent learns the position of a visual feature in the environment it also maintains an estimate of the uncertainty of this estimate. This uncertainty value can provide a quantitative measure of the utility of a particular flight structure for learning the position of a visual landmark. If the learning is effective then uncertainty should be highly correlated with accuracy and so the uncertainty in a landmark estimate should give a measure of how accurately the bee will be able to re-orient when the landmark comes back into view. Our analysis involved, for all flights, running the model multiple times to simulate how learning would have progressed for all possible landmark positions. Comparison of the final uncertainty values for different fictive landmark

positions, tells us where in the environment a visual landmark would be learnt most accurately for a given flight.

The second key feature of our approach is that by using a closed-loop model, we capture the sensory consequences of movements. The impact that particular movements have on positional estimates is automatically incorporated. We do not need to make assumptions about which movements might be useful to the agent but instead only consider the effects of a flight structure on uncertainty reduction. This means that we do not need to classify movements as, for instance, peering or pivoting, we simply consider the utility of that movement in terms of reducing positional uncertainty (e.g. see Fig. 1).

Figure 5 shows a SLAM analysis of the three example flights together with an analysis of the combined data from all flights in our dataset. We did not normalise the individual uncertainty maps prior to combining them since this would have biased our results in favour of the shorter flights during which the arcing behaviours were less prominent. Moreover, the uncertainty maps for individual flights had a similar range of values so we simply calculated the mean value across all flights for each position within our ROI. The SLAM analysis suggests that the region around the nest entrance would be the region that would be most effectively learned given the recorded flights and our assumptions about the bees' sensor and motor capabilities. Again, we see no observable bias towards the position of the landmark. This analysis therefore also argues against the hypothesis that the bees are attempting to specifically learn about the position of the landmark by structuring their flight to extract depth information in an optimal or even near-optimal way.

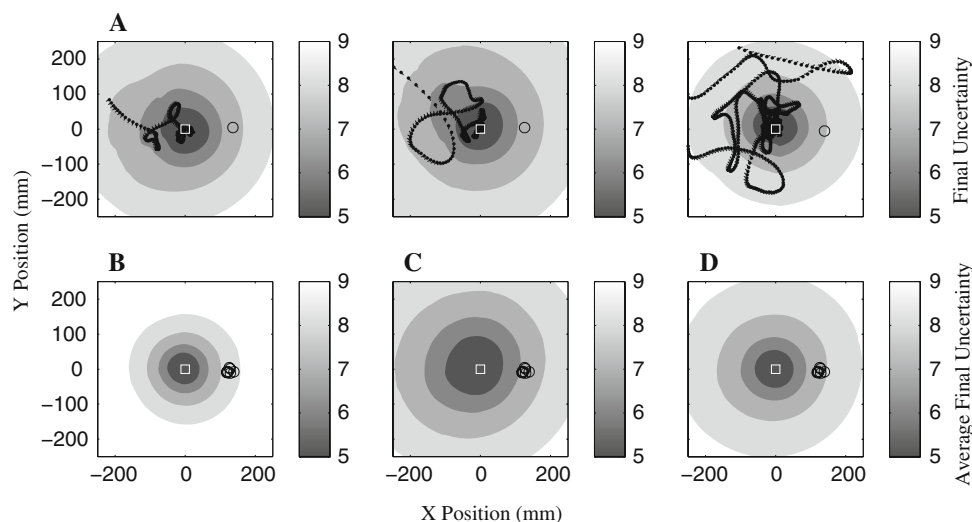



Fig. 5 a SLAM analysis of single flights. The contour plots show the final uncertainty in a SLAM based map for fictive landmarks placed at each point in the ROI. The position of the nest entrance is indicated by the white square and the true position of the landmark is shown by the black circle. b–d Combined data for 37 flights showing first phase, sec-

ond phase and entire flight, respectively (see Sect. 2 for description of phases). All distances are measured in mm. The nest (*open square*) is at (0,0) and the landmark (*open circle*) is at (125,0). Landmark positions varied slightly between recordings

Fig. 6 SLAM analysis of artificial flights showing the effects of viewing direction and flight structure. The *top row* shows a straight flight, the *middle row* a spiral centred on the nest, and the *bottom row* a spiral centred on the landmark. The *left column* shows the result of each flight when facing the direction of movement, the *middle column* when facing the nest, and the *right column* when facing the landmark. All distances are measured in mm. The nest (*open square*) is at (0,0) and the landmark (*open circle*) is at (125,0)



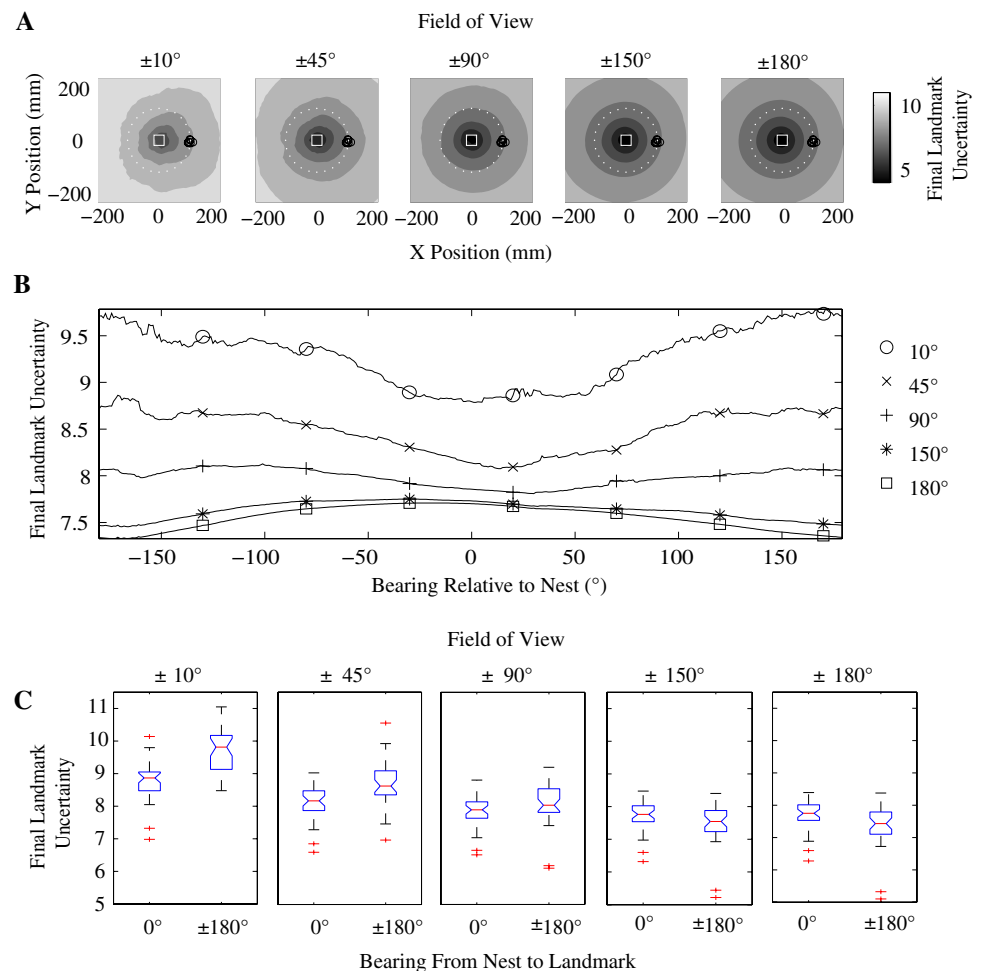
To help us interpret the uncertainty maps of Fig. 5 we also looked at the uncertainty maps that would be produced by artificial flights. The artificial flights were generated by taking a recorded flight and changing the viewing direction and direction of movements, while maintaining the same speed profile. Figure 6 shows the uncertainty maps generated for three illustrative artificial trajectories, a straight flight, and nest or landmark centred spirals. To assess the impact of viewing direction we simulated three looking directions for each trajectory: in the direction of movement and fixating the nest or landmark. There are several points of note. First, the starting position has a strong influence on the final uncertainty map. In the straight flight when the agent faces forward (top left), despite the fact that the start is viewed less often than the end of the flight, the minimum is clearly at the start of the trajectory. The two other straight flights (top middle and right) show the more subtle effect of changing the viewing direction. It is possible to induce an asymmetry in the map by varying the viewing direction, however, the minimum clearly remains centred on the starting position. Second, it is possible to shift the focus of the uncertainty map. When the spiral flight is landmark rather than nest-centred, the uncertainty distribution changes and the region containing low uncertainties is clearly biased toward the landmark. Overall, it is evident that the path of the flight has a much stronger influence on the final uncertainty maps than viewing direction and it would be straightforward to

design a flight where the landmark location was learnt most accurately.

4.3 Varying the field of view

Anecdotal evidence suggests that during visuo-motor behaviour the frontal retina is important. In a final analysis we look at varying the field of view between $\pm 10^\circ$ up to a full 360° to investigate what effect this has on the uncertainty maps in our SLAM simulations. Figure 7 shows a summary of the results for five different fields of view. The first thing to note is that irrespective of the field of view the uncertainty is always lowest for the area centred on the nest. However, when we look in more detail we see that a smaller field of view induces asymmetries in the resulting uncertainty map. This asymmetry is made clearer in Fig. 7b which shows the final uncertainty for fictive landmarks along a circle centred on the nest (shown in white in Fig. 7a). For narrow fields of view, $\pm 10^\circ$ and $\pm 45^\circ$, there is a clear asymmetry with a minimum centred on the true bearing to the landmark from the nest. For wider fields of view the reverse pattern is found, with the minimum at a bearing of 180° relative to the true bearing to the landmark from the nest. The difference in the mean uncertainty (Fig. 7c) for fictive landmarks at 0° and 180° bearings is very small but significantly different (*t* test, $P < 0.05$) for all fields of view. The slight decrease in uncertainty for the position opposite to the true landmark position

Fig. 7 **a** Contour plots showing the final uncertainty maps for five different fields of view generated by combining all 37 recorded flights. The position of the nest entrance is indicated by the *white square* and the positions of the landmark are shown by the *black circles*. The *large white circle* in each plot indicates a line of equal distance from the nest that passes through the landmark. **b** Final uncertainty for different fields of view along lines of equal distance from the nest. The true position of the landmark is at a bearing of 0° . **c** Box plots describing the final landmark uncertainty for fictive landmarks at bearings relative to the nest of 0° and $\pm 180^\circ$ for different fields of view. The boxes have lines at the lower quartile, median, and upper quartile values. The whiskers are lines extending from each end of the boxes to show the extent of the rest of the data, excluding outliers (*crosses*). The difference in the means at bearings relative to the nest of 0° and $\pm 180^\circ$ are significantly different (*t* test, $P < 0.05$) for all fields of view



is due to the bee backing away from the landmark and therefore flying closer to this point. The asymmetry that we see for narrower fields of view is probably due to the bees maintaining the landmark in the frontal part of their visual field as they back away from the nest.

5 Discussion

Orientation flights have been proposed to aid learning in several ways. For instance, it could be that the flight simply keeps the bee in the vicinity of the nest, thus enabling it to sample the visual environment in this region many times. A second hypothesis is that the flights facilitate the acquisition of 2D images of the world (snapshots), possibly at the ends of the arcs (Collett and Lehrer 1993; Collett 1995).

The first hypothesis is clearly compatible with our data although it fails to explain the arcing behaviour that is so typical of these flights. It is not clear to us how we could test the second hypothesis with our current approach. We therefore leave open the possibility that the structure and function of

the flights is related to learning snapshots. Our approach can however be applied to examining a third hypothesis, namely that the arcing structure of the flights is optimised for measuring the distance to conspicuous landmarks in the immediate vicinity of the nest (Cheng et al. 1987; Lehrer 1991, 1993; Zeil 1993b; Collett and Zeil 1997).

The aim of our study was to determine to what degree the structure of orientation flights is consistent with bees learning a metric representation of their environment. To do this we asked whether the flights were, in some well defined way, optimised for learning the position of a single conspicuous landmark. The key analyses assessed which parts of the world consistently generate optic flow and which parts would be learnt effectively through a map-based feature tracking framework.

Our analysis of optic flow proved problematic. It is not clear how we can relate any of our measures of optic flow to the accuracy with which a bee might learn the position of a landmark, thus motivating our use of SLAM.

The SLAM analysis of flights shows little or no bias towards the actual landmark position. Instead, for our flight paths, positions of fictive landmarks near the nest entrance

were consistently learnt more accurately than those near the actual landmark position. We conclude therefore that although bees are clearly influenced by the landmark, in that they preferentially maintain it in their field of view, the focus of learning appears to be only very slightly influenced by the landmark position.

That learning is centred on the nest is perhaps unsurprising. That there appears to be only a very subtle bias towards the position of the landmark argues against the hypothesis that the flights are specifically structured so as to learn about the position of the landmark. When we restrict the field of view of the simulated bee (Fig. 7) the actual landmark position now has a lower uncertainty than other positions at the same distance from the nest. While the difference in uncertainty is consistent, it is small and we do not yet know if bees preferentially use their frontal field for landmark guidance and so it is hard to assess the behavioural significance of this result. Moreover, even in the case of the highly restricted field of view, positions close to the nest have by far the lowest uncertainty.

Learning flights in ground-nesting bees and wasps are known to have a distinct 3D structure, with the insects gaining height and horizontal distance from the nest at roughly the same rate (Zeil 1993a; Zeil et al. 1996). However, for the flights we observed, the bees maintained a fairly constant height throughout the portion of the flight that we were able to record. Given a 3D trajectory, a SLAM analysis could equally well be performed in three dimensions. If we included height we would expect to see a similar pattern of results to the ones we report here, although we might expect to see higher uncertainty in the height estimate due to the lower variation in movements in this dimension. Interestingly, bees with a nest on a vertical surface exhibit a more oscillatory vertical component during their orientation flights. Rudolf Jander observed Euglossine bees performing almost circular orientation flights outside their arboreal nest entrance (Jander 1997). It would be interesting to apply a SLAM analysis in 3D to these flights to see whether learning was focussed on the nest in the third dimension in the same way as we have observed for movements in the horizontal plane.

In this paper, we have introduced an approach derived from probabilistic SLAM to analyse bee orientation flights and investigate whether they are structured to efficiently learn the positions of the landmarks that are available in a given terrain. Our analysis focussed on the accuracy with which fictive objects at different locations would be localised. We showed that our flights are not optimised to learn about the position of a prominent landmark, but are more suited to learn about objects near the nest. One of the benefits of this approach is that we can vary the parameters of the sensor and motor models to investigate different aspects of a system in a closed-loop. This also allows the incorporation of new information about the motor patterns and sensory

capabilities of bees or, indeed, other species, as and when they become available. We contend that this approach provides a powerful new tool for the study of active spatial learning.

Acknowledgments Financial support came from the EPSRC and BBSRC under the Cognitive Foresight scheme and PG was supported by a Leverhulme Fellowship. We thank two referees for their constructive advice.

References

- Becker L (1958) Untersuchungen über das Heimfindevermögen der Bienen. *Z Vergl Physiol* 41: 1–25
- Capaldi EA, Smith AD, Osborne JL, Fahr-Bach SE, Farris SM, Reynolds DR, Edwards AS, Martin A, Robinson GE, Poppy GM, Riley JR (2000) Ontogeny of orientation flight in the honeybee revealed by harmonic radar. *Nature* 403:537–540
- Cheng K, Collett TS, Pickhard A, Wehner R (1987) The use of visual landmarks by honeybees: Bees weight landmarks according to their distance to the goal. *J Comp Physiol A* 161:469–475
- Cheng K, Shettleworth S, Huttenlocher J, Rieser J (2007) Bayesian integration of spatial information. *Psychol Bull* 133(4):625–637
- Collett T (1995) Making learning easy: the acquisition of visual information during the orientation flights of social wasps. *J Comp Physiol A* 177:737–747
- Collett T, Lehrer M (1993) Looking and Learning: A Spatial Pattern in the Orientation Flight of the Wasp *Vespa vulgaris*. *Proc R Soc* 252(1334):129–134
- Collett T, Zeil J (1997) The selection and use of landmarks by insects. In: Lehrer M (ed) *Orientation and communication in arthropods*. Birkhauser, Basel, pp 41–65
- Courville A, Daw N, Touretzky D (2006) Bayesian theories of conditioning in a changing world. *Trends Cognit Sci* 10:294–300
- Cummins M, Newman P (2008) FAB-MAP: Probabilistic Localization and Mapping in the Space of Appearances. *Int J Robot Res* 27(6):647–665
- Davison A, Murray D (2002) Simultaneous Localisation and Map-Building Using active Vision. *IEEE Trans PAMI* 24(7):865–880
- Dissanayake M, Newman P, Clark S, Durrant-Whyte H, Csorba M (2001) A solution to the simultaneous localization and map building (SLAM) problem. *IEEE Trans Robot Autom* 17(3):229–241
- Durrant-Whyte H, Bailey T (2006) Simultaneous localisation and mapping (SLAM): part i the essential algorithms. *Robot Autom Mag* 13:99–110
- Gallistel CR (1989) Animal cognition: The representation of space, time, and number. *Annu Rev Psychol* 40:155–189
- Gould JL (1986) The Locale Map of Honey Bees: Do Insects Have Cognitive Maps? *Science* 232(4752):861–863
- Griffiths T, Kemp C, Tenenbaum J (2008) Bayesian models of cognition. The Cambridge handbook of computational cognitive modeling. Chap. 3. Cambridge University Press, London
- Guivant J, Nebot E (2001) Optimization of the simultaneous localization and map-building algorithm for real-time implementation. *IEEE Trans Robot Autom* 17(3):242–257
- Jander R (1997) Macroevolution of a fixed action pattern for learning: the exploration flights of bees and wasps. In: Greenberg G, Tobach E (eds) *Comparative psychology of invertebrates: the field and laboratory study of insect behavior*. Garland, New York, pp 79–99
- Knill D, Pouget A (2004) The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci* 27(12):712–719
- Kording K, Wolpert D (2006) Bayesian decision theory in sensorimotor control. *Trends Cognit Sci* 10:319–326

- Land M (1999) Motion and vision: why animals move their eyes. *J Comp Physiol* 185:341–352
- Lee T, Mumford D (2003) Hierarchical Bayesian inference in the visual cortex. *J Opt Soc Am A* 20(7):1434–1448
- Lehrer M (1991) Bees which turn back and look. *Naturwissenschaften* 78:274–276
- Lehrer M (1993) Why do bees turn back and look? *J Comp Physiol A* 172:549–563
- Lehrer M, Collett TS (1994) Approaching and departing bees learn different cues to the distance of a landmark. *J Comp Physiol A* 175:171–177
- Menzel R, Greggers U, Smith A, Berger S, Brandt R, Brunke S, Bundrock G, Hülse S, Plümpe T, Schaupp F, Schüttler E, Stach S, Stindt J, Stollhoff N, Watzl S (2005) Honey bees navigate according to a map-like spatial memory. *PNAS* 102(8):3040–3045
- Milford M, Wyeth G (2008) Mapping a suburb with a single camera using a biologically inspired slam system. *IEEE Trans Robot* 24(5):1038–1053
- Montemerlo M, Thrun S, Koller D, Wegbreit B (2002) Fast-SLAM: A factored solution to the simultaneous localization and mapping problem. *Proc AAAI Natl Conf Artif Intell* 2002:593–598
- Rao R (2004) Hierarchical Bayesian inference in networks of spiking neurons. *Adv. neural information processing systems (NIPS 04)*, vol 17. MIT Press, Cambridge
- Seidl R, Kaiser W (1981) Visual field size, binocular domain and ommatidial array of the compound eyes of honey bees. *J comp Physiol A* 143:17–26
- Smith R, Self M, Cheeseman P (1990) Estimating uncertain spatial relationships in robotics. In: Cox IJ, Wilfon GT (eds) *Autonomous Robot Vehicles*. Springer, Heidelberg, pp 167–193
- Spaethe J, Tautz J, Chittka L (2001) Visual constraints in foraging bumblebees: Flower size and color affect search time and flight behavior. *PNAS* 98(7):3898–3903
- Srinivasan M, Zhang S, Bidwell N (1997) Visually mediated odometry in honeybees. *J Exp Biol* 200:2513–2522
- Thrun S, Fox D, Burgard W (1998) A probabilistic approach to concurrent mapping and localization for mobile robots. *Mach Learn* 31(1):29–53
- Thrun S, Burgard W, Fox D (2005) *Probabilistic robotics*. MIT Press, Cambridge
- Vollbehre J (1975) Zur Orientierung junger Honigbienen bei ihrem ersten Orientierungsflug. *Zool Jb allg Zool Physiol* 79:33–69
- van Iersel J, van den Assem J (1964) Aspects of orientation in the digger wasp *Bembix rostrata*. *Anim Behav (Suppl I)*:145–162
- von Frisch K (1967) *The dance language and orientation of bees*. Belknap/Harvard University Press, Cambridge
- Voss R, Zeil J (1998) Active vision in insects: an analysis of object-directed zig-zag flights in a ground-nesting wasp (*Odynerus spinipes*, Eumenidae). *J Comp Physiol A* 182:377–387
- Wehner R (1981) Spatial vision in arthropods. *Handbook of sensory physiology VII/6C*. Springer, Heidelberg 287–616
- Wehner R, Menzel R (1990) Do insects have cognitive maps? *Annu Rev Neurosci* 13:403–414
- Wehner R, Boyer M, Loertscher F, Sommer S, Menzi U (2006) Ant navigation: one-way routes rather than maps
- Yuille A, Kersten D (2006) Vision as Bayesian inference: analysis by synthesis? *Trends Cognit Sci* 10:301–308
- Zeil J (1993a) Orientation flights of solitary wasps (Cerceris; Sphecidae; Hymenoptera). I Description of flight. *J Comp Physiol A* 172:189–205
- Zeil J (1993b) Orientation flights of solitary wasps (Cerceris; Sphecidae; Hymenoptera). II Similarities between orientation and return flights and the use of motion parallax. *J Comp Physiol A* 172:207–222
- Zeil J, Kelber A (1991) Orientation flights in ground nesting wasps and bees share a common organization. *Verh Dtsch Zool Ges* 84:371–372
- Zeil J, Kelber A, Voss R (1996) Structure and function of learning flights in bees and wasps. *J Exp Biol* 199:245–252
- Zeil J, Boeddecker N, Hemmi J, Stürzl W (2007) Going wild: towards an ecology of visual information processing. In: North G, Greenspan R (eds.) *Invertebrate neurobiology*. Cold Spring Harbor Monograph Series 49. Chap. 15