

RESEARCH ARTICLE

How might ants use panoramic views for route navigation?

Andrew Philippides¹, Bart Baddeley¹, Ken Cheng² and Paul Graham^{3,*}

¹School of Informatics, University of Sussex, Falmer, Brighton BN1 9QJ, UK, ²Department of Brain, Behaviour and Evolution, Macquarie University, 2109 Sydney, NSW, Australia and ³School of Life Sciences, JMS Building, University of Sussex, Falmer, Brighton BN1 9QG, UK

*Author for correspondence (p.r.graham@sussex.ac.uk)

Accepted 13 October 2010

SUMMARY

Studies of insect navigation have demonstrated that insects possess an interesting and sophisticated repertoire of visual navigation behaviours. Ongoing research seeks to help us understand how these behaviours are controlled in natural complex environments. A necessary complement to behavioural studies is an understanding of the sensory ecology within which an animal behaves. To this end we have analysed ants' perspective views of a habitat within which desert ant navigation is well studied. Results from our analysis suggest that: parsimonious visual strategies for homing and route guidance are effective over behaviourally useful distances even in cluttered environments; that these strategies can function effectively using only the skyline heights as input; and that the simplicity and efficacy of using stored views as a visual compass makes it a viable and robust mechanism for route guidance.

Key words: insect navigation, view-based homing, snapshots, image difference function, route navigation, skyline.

INTRODUCTION

Many ants use fixed routes to travel between their nest and a profitable foraging ground, e.g. *Cataglyphis bicolor* (Santschi, 1913), *Formica rufa* (Rosengren and Fortelius, 1986), *Cataglyphis fortis* (Wehner et al., 1996), *Melophorus bagoti* (Kohler and Wehner, 2005). Experienced ants use visual landmarks to guide their routes and experiments show that portions of the route can be performed out of sequence and independently of path integration (Collett et al., 1992; Collett et al., 1998; Andel and Wehner, 2004; Kohler and Wehner, 2005). Interest in how this robust behaviour is produced by small-brained animals with their low-resolution visual system has intrigued biologists and roboticists. The subsequent cross-fertilisation of ideas has influenced the approach we have taken in this paper; a computational analysis of the visual information available to ants as they navigate routes through complex natural environments.

Laboratory and field studies using artificial landmarks at the nest (Wehner and R ber, 1979; Wehner et al., 1996; Akesson and Wehner, 2002; Narendra et al., 2007) or a feeder (Wolf and Wehner, 2000; Durier et al., 2003; Graham et al., 2004) have shown how visual landmark information can be used to guide the search for an important location. The basic mechanism is for a single view of the world to be stored at the goal location. The difference between the current view of the world and the view from the goal location can subsequently be used to drive the search for that goal, so-called 'view-based homing' or 'snapshot matching' (Cartwright and Collett, 1983). It is natural to ask whether view-based homing can also be used to guide long natural habitual routes.

Our knowledge of the mechanisms underpinning visually guided routes is less extensive than our understanding of view-based homing. A variety of experiments have highlighted how simple procedural rules can be used to associate directional information with visual landmarks and so guide small portions of a route (Collett

et al., 1992; Collett et al., 1998; Collett et al., 2001; Pratt et al., 2001; Graham and Collett, 2002; Graham et al., 2003; Harris et al., 2007; Collett, 2010). However, we do not have a general understanding of the mechanisms by which ants navigate habitual routes using information from natural visual panoramas. Similarly, we know little about how ants extract information from natural visual environments. It has long been suggested that in natural environments the skyline profile could provide a characteristic signature for a location (Wehner and R ber, 1979; Wehner et al., 1996) or provide easily identifiable discrete landmarks that can be associated with appropriate directions (Fourcassi , 1991; Fukushi, 2001). Recently, Graham and Cheng demonstrated that a skyline profile generated by an artificial arena can functionally mimic a natural panoramic scene even when colour cues, the distance distribution of objects and orientation relative to a celestial compass radically differ from the ants' familiar foraging locations (Graham and Cheng, 2009b). However, we do not know what information is extracted from a skyline and how this might be used for route guidance.

Our attempts to address these questions are fundamentally hampered by our lack of understanding of the ants' perspective of their environment. Here we begin the process of quantifying the visual information available to ants as they move through the world. Specifically, we have asked over what range can the comparison of the current view of the world with a remembered view of the world provide useful navigational information. This is a fundamental and powerful question because it relates to the issue of the level of world knowledge necessary for route performance.

Our approach follows that of Jochen Zeil and colleagues (Zeil et al., 2003; St rzl and Zeil, 2007), who captured sets of images using a panoramic imaging device within a natural environment of significance to behaving animals; in their case, ground-nesting solitary wasps. By measuring the difference between a reference

image and images from surrounding points, they defined an image difference function (IDF) and showed that, over a few metres, image differences increase monotonically with increasing distance. Insects can return to a goal by monitoring the difference between what they currently see and the stored reference image, then moving so that the difference decreases. Therefore, the presence of a smoothly increasing IDF showed that the information needed for view-based homing is available in unprocessed natural scenes and can be utilised over behaviourally significant distances. This result relies on the camera being aligned to an external frame of reference for all images. However, Zeil and colleagues also showed that the alignment of the camera when the reference image was taken (a proxy for heading information) can be robustly recovered at locations near the goal by comparing the reference image to rotated versions of the current image (Zeil et al., 2003). The orientation at which the current image best matches the reference image will be close to the orientation of the reference image. The implications for route following are stark: an insect with a visual system that is fixed relative to its body axis can recover a heading by rotating until it finds the best match between the current scene and a stored snapshot. Therefore the correct heading for a portion of a learnt route can be specified by a snapshot stored when the insect was previously moving in the correct direction along the route.

The study by Zeil and colleagues (Zeil et al., 2003) laid the foundations for a quantitative analysis of real visual environments with respect to navigation. Principally, the catchment area (Cartwright and Collett, 1987) of a panoramic view can be defined; after a certain distance from the goal the gradient of the IDF becomes flat, this distance being dependent on the depth structure of the world (Stürzl and Zeil, 2007). The general significance of these results is ensured by the parsimony of the analysis. Evaluating differences between current and reference images using an intentionally simplistic measure (the root mean square, r.m.s., pixel difference) means that usable information is available without the need for complex visual processing. Moreover, more sophisticated models of visual homing, which preserve retinotopic information, will generally be successful.

The work of Zeil and colleagues represents the first attempt to quantify the information that exists for visual homing in natural environments. We have taken a similar approach and measured the catchment areas of natural unprocessed scenes along with natural 1D skylines across a range of natural environments from open to cluttered. Specifically, we measured the distance from the goal over which the r.m.s. pixel difference between a reference image and route images increases smoothly. This analysis was performed using either full scenes or 1D skylines. Additionally, we have analysed the distance over which heading information can be usefully extracted when the difference between a reference image and rotated versions of the current image is used as a visual compass.

MATERIALS AND METHODS

Location

Images were collected from a field site 10 km south of Alice Springs, NT, Australia, where there are many colonies of the Australian desert ant *Melophorus bagoti* Lubbock. The navigational behaviour of this ant species has been well studied at this site (for a review, see Cheng et al., 2009).

Images

Panoramic images were collected from four transects through environments that were subjectively graded on a continuum from

open to cluttered. Although this ant is often referred to as a desert ant, its habitat contains lots of natural vegetation, with an abundance of grass tussocks, bushes and trees (Muser et al., 2005). The primary factor that correlated with the subjective assessment of clutter was the density of grass tussocks. For each transect, a straight line of approximately 30 m was pegged out and images were taken every metre using a GoPano panoramic lens (EyeSee360, Inc., Pittsburgh, PA, USA) with a Canon Powershot 720 digital camera (Canon UK, Reigate, Surrey, UK). At each location a thin piece of board was placed directly on the ground and levelled with a spirit level. The camera was positioned with the lens down to capture the panoramic scene from almost ground level and the images were assumed to be aligned to a common heading as they had been captured using a guideline, although this process could have introduced errors of a few degrees. Examples of images from the open and cluttered transects are shown in Fig. 1. The panoramic images were unwrapped with Photowarp[©] (EyeSee360, Inc.), cropped and resized to approximately 1 pixel deg⁻¹ which is of the order of the likely resolution of the *M. bagoti* compound eye. The resultant field of view was 360 deg by 90 deg (35 deg below and 55 deg above the horizon) and our analysis was performed with the portion of the image above the horizon.

Skyline extraction

Extracting the skyline was a two-stage process. Firstly the images were converted to a binary image by manually adjusting a threshold on the green channel so that as much foliage and ground was included (ON pixels) without parts of the sky also being classified as ON. Any flare or bleed interference from the edge of the lens

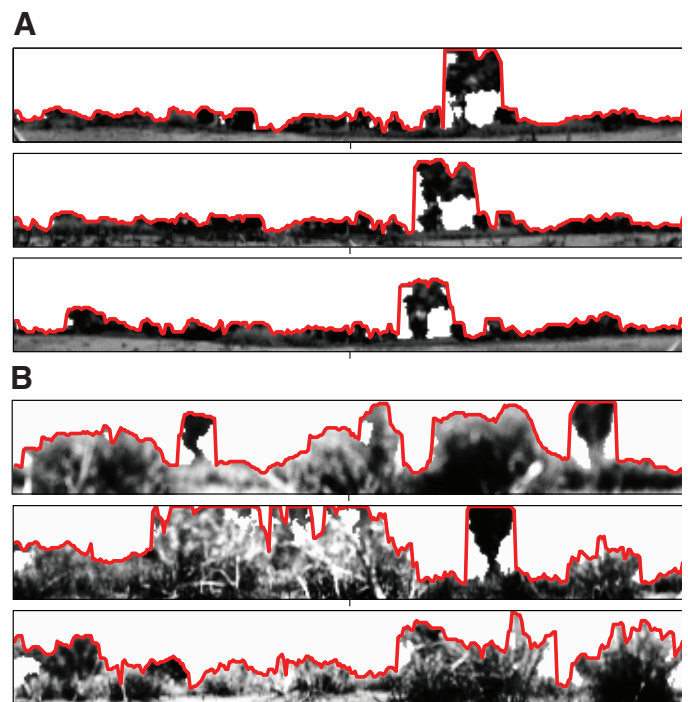


Fig. 1. Ant's eye views from open and cluttered environments. (A) The panoramic view from an ant's perspective of three locations on the most open of the transects. (B) Panoramic views from the most cluttered transect. Images are from locations 5 m apart and are shown after pre-processing, which involved histogram equalisation, sky detection and subsequent sky homogenisation. The thick red line on each image is the extracted skyline.

was removed manually. Subsequently, the first row of pixels was set to ON and any 'holes' in the binary map (OFF pixels surrounded by ON pixels) were filled in. Finally, any 'floating' objects (sets of ON pixels surrounded by OFF pixels) were removed, resulting in a binary 'template' image showing ground and foliage as ON and sky as OFF. The skyline was then defined as the height of the highest ON pixel in each azimuth. Skyline extraction is intentionally simple as we did not want the results to be dependent on sophisticated processing or optimised parameter choices. Possible variations in skyline extraction can be safely ignored as we tested variants of the skyline extraction process (e.g. taking the lowest ON pixel at each azimuth, not removing floating objects before selecting the highest ON pixel) and found a negligible effect on the overall results. It is likely that skyline extraction would be straightforward for ants. They have dichromatic vision with peak sensitivities in the UV and green range. A simple UV–green opponent channel would be perfectly suited to extracting the skyline (Möller, 2002).

Whole images

To evaluate properly the information available in an unprocessed scene we had to mitigate any influence of varying light levels or persistent light gradients from sun position, as they may have biased our recorded catchment areas. Firstly, contrast was normalised using histogram equalisation (histeq function in Matlab, MathWorks, Natick, MA, USA) of grayscale images, resulting in integer-valued pixels in the range 0–255. Subsequently, using the binary image template delineating sky from not-sky, we homogenised the sky to an intensity of 250 (several other sky intensities were tested – 50, 100, 150, 200 – with little effect on the overall results).

Catchment areas

For each transect we calculated two catchment areas for every image along the transect. The first was based on the IDF, the pixel-based r.m.s. difference between an image and a reference image when the two images are aligned to a common heading. The catchment area of an image is defined as the region within which an agent could return to the location where the reference image was taken by descending the gradient of the IDF (Zeil et al., 2003). The catchment area is defined as the number of consecutive locations spreading out from the reference image where the IDF gradient is positive, relative to the direction of movement, on either side of that location. In Figs 2, 3 and 4, we report the median radius of the catchment areas. This process was undertaken for whole images and scenes encoded as 1D skylines (e.g. Fig. 3A).

The second, the rotational catchment area, is determined from the rotational IDF (RIDF). The RIDF is calculated by evaluating the r.m.s. difference between a reference image and the current image rotated (*in silico*) in steps of 1 deg of azimuth, resulting in a 1×360 RIDF (e.g. Fig. 4A). The minimum value in the RIDF defines an orientation of the current image that gives the closest match with the reference image. In the vicinity of the reference image these orientations will be similar (Zeil et al., 2003). We defined a rotational catchment area as the region spreading out from the location of the reference image where the minimum in the RIDF is less than 45 deg from the true orientation of the reference image.

In a secondary analysis of the information available in the RIDF we applied a simple behaviourally plausible heuristic when calculating the rotational catchment areas. Starting from the goal locations we moved out along the transect. For each image we calculated the RIDF and from this extracted the three most prominent

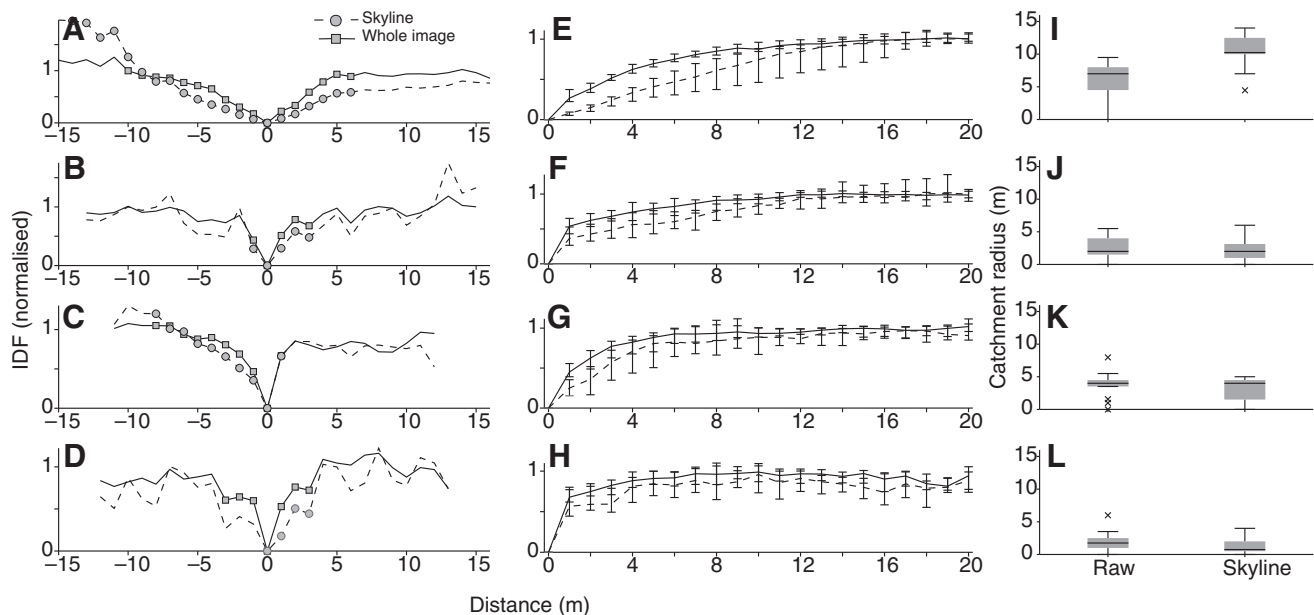


Fig. 2. Catchment areas based on image difference functions (IDFs). (A–D) IDF for a single reference image at the midpoint of each transect. The transects were subjectively ordered with A being the most open transect and D being the most cluttered. For each goal, the IDF was calculated for the whole image (solid line with squares) and the skyline only (dashed line with circles). For presentation, the IDFs are normalised to the 80th percentile. Points marked with symbols are considered to be within the catchment area of the reference image (see Materials and methods). (E–H) Median IDFs as a function of distance from the goal. For each transect, we collated the IDFs as a function of distance from the goal over all goal locations and illustrate the distributions by showing the median values, together with 25th and 75th percentiles as error bars, for the whole image (solid line) and skyline only (dashed line). $N > 7$ for all distances. For presentation, values are normalised to the maximum median IDF within each transect for the whole image and skyline, respectively. (I–L) Box plots showing the median radius of catchment areas for all positions on each route (box: 25th and 75th percentile; central line: median; whiskers: extent of the data; crosses: outliers, defined as points that are beyond the quartiles by more than 1.5 times the inter-quartile range). Data are shown for unprocessed images (left) and skyline (right).

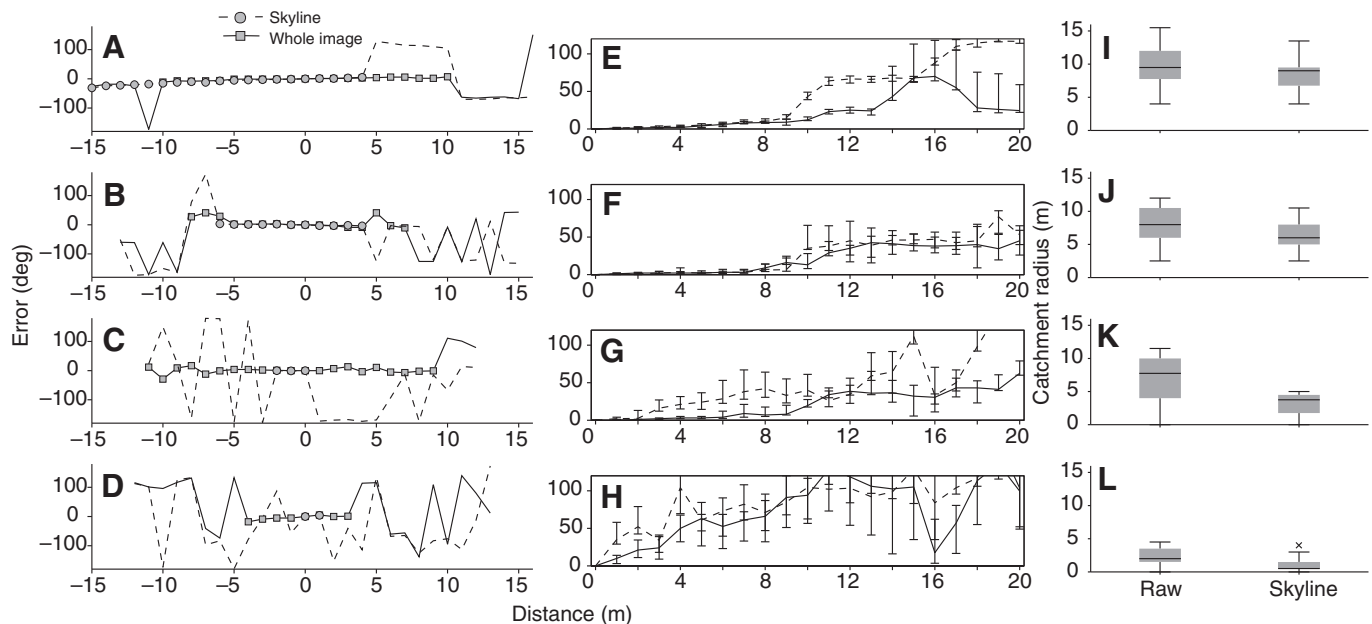


Fig. 3. Catchment areas based on rotational IDF. (A–D) Orientation errors after all images along a transect are rotated to find the best match with the reference image for whole images (squares) and skylines (circles). Points marked with symbols are considered to be within the catchment area of the reference image (see Materials and methods). (E–H) Median absolute orientation errors as a function of distance from the goal. We collated absolute orientation errors as a function of distance from the goal over all goal locations and illustrate the distributions of the errors by showing the median values, together with 25th and 75th percentiles as error bars, for the whole image (solid line) and skyline only (dashed line). $N > 7$ for all distances. (I–L) Box plots showing the median radius of catchment areas for all positions on each route (box: 25th and 75th percentile; central line: median; whiskers: extent of the data; crosses: outliers, defined as points that are beyond the quartiles by more than 1.5 times the inter-quartile range). Data are shown for unprocessed images (left) and skyline (right).

minima (Fig. 4A). Each minimum represents an orientation where there is a locally optimal match with the reference image. Rather than simply taking the best match, we took the minimum that is closest in orientation to the orientation at the previous location. This heuristic reduces aliasing by favouring RIDF minima close to the previous heading.

RESULTS

Gradient descent of IDF

We began our analysis of the information available in natural panoramic scenes by considering the IDF for all possible goal positions along our four transects. The IDF is generated by evaluating the pixel-wise r.m.s. difference between the reference image and all other images on the route. The catchment area of each goal image is estimated by looking for changes in the sign of the gradient of the IDF (e.g. Fig. 2A). Estimating the catchment area in this way gave us a direct measure of how useful a single stored snapshot would be for navigation in that environment. We performed this analysis with whole images that had been minimally processed to avoid systematic bias from varying light levels or gradients in sky intensities. We also measured catchment areas for panoramic images sparsely encoded as a 1D skyline profile representing the height of foliage against the sky for each azimuthal direction. For these image sets the IDF was generated by calculating the r.m.s. difference of skyline heights in the reference and route images.

Fig. 2 shows examples of IDF functions for whole images and skylines from a single goal at the midpoint of each transect (Fig. 2A–D), plots of how image difference relates to distance from the reference image (Fig. 2E–H) and also the distribution of catchment areas (Fig. 2I–L) for all possible goals along the

transect. The size of catchment areas is strongly dependent on the environment type and, as expected, catchment areas increase as the environments become more open. Nonetheless, even for the cluttered environment, where the appearance of the world can change very quickly (Fig. 1), there is still a gradient in the function relating median IDF to distance up to 6 m from the reference image (Fig. 2H). However, the median radius of the catchment areas is small (Fig. 2L). This is likely to be because of transient visual clutter and also because noise introduced during image collection is greater for this transect because of the increased difficulty in levelling the camera. When using precision gantry equipment (Zeil et al., 2003; Stürzl and Zeil, 2007) the IDFs from complex natural environments are always smooth. Therefore, we also looked at catchment radius size after smoothing the IDF (with a median filter of size 3). We assumed that a simple smoothing would be a biologically plausible mechanism, equivalent to ants performing a temporal averaging. In this case, the median radius increased from 2 to 4.5 m and from 1 to 4 m for raw images and skyline, respectively (Fig. 2L). Similarly with our second route the reported radii do not seem to match the smooth function in the median IDF versus distance (Fig. 2F). In this case we think a single anomalous image may be curtailing catchment areas. Again, a simple smoothing ameliorated this issue and catchment radius increased from 2 to 6.5 m and from 2 to 7 m for raw images and skyline, respectively (Fig. 2J).

These represent behaviourally significant distances when compared with typical foraging distances of 20 m for *M. bagoti* in this environment (Muser et al., 2005). Across all four transects there is little difference between using the whole image and using the skyline, suggesting that encoding the scene as a skyline does not discard too much useful information.

Using a goal image to retrieve orientation

The analysis of IDFs in natural environments shows that panoramic scenes contain information useful for view-based homing over a behaviourally relevant scale. However, if we consider in more detail what is actually required for route guidance, then perhaps stored views can be used more simply. Route guidance requires, in the main, an ant to decide in which direction to go rather than to home accurately to a specific goal or sub-goal. Therefore we investigated the range over which a stored view can be used to recover the orientation at which the reference image was taken. As ants are constrained by their morphology to travel in the direction of their long axis, the orientation of a goal image can serve as a proxy for route direction. At each point along a transect the image is rotated through 360 deg in steps of 1 deg and for each orientation the image is compared with the reference image. This creates a RIDF. The lowest value in this function represents the orientation of the current image that most closely matches the goal image. We consider this to be a useful match if the discrepancy between the orientations of the goal and current image is less than 45 deg. We define the rotational catchment area (RCA) as the region over which the discrepancy between the orientations of the goal and current image is less than 45 deg, which we consider to be a useful match.

Fig. 3 shows examples of RIDFs for whole images and skylines from a single goal on each transect (Fig. 3A–D), plots of how median rotational error relates to distance from the reference image (Fig. 3E–H) and also the median catchment areas (Fig. 3I–L) for all possible goals along the route. Again, the size of RCAs is strongly dependent on the environment type, with the largest catchment areas in open environments. In contrast to IDFs, whole images appear to out-perform skylines when used as a visual compass largely because aliasing is more likely when scenes are described only in 1D.

Improving performance with a simple heuristic

The RIDFs produced when comparing images often have a characteristic shape (Fig. 4A) with multiple minima at orientations where the two images match reasonably well. This can lead to aliasing and the retrieval of an inaccurate heading. To demonstrate how using a stored image as a visual compass lends itself to simple behavioural heuristics, we recalculated RCAs using at each location the nearest RIDF minimum to the orientation retrieved at the previous point (Fig. 4A). This reduces aliasing by favouring RIDF minima close to the previous heading, resulting in an increase in performance for all four transects (Fig. 4B–E and 4F–I). In addition to improving performance a heuristic like this would also reduce the processing requirement for an ant. Rather than analyse all possible orientations at each point she can scan either side of her current heading until perceiving a significant minima in the RIDF (see Baddeley et al., in press).

DISCUSSION

Catchment areas of natural scenes

We have analysed ants'-perspective views of natural visual environments for insights into the likely and viable mechanisms of visually guided route navigation. Our principal finding is that, in natural cluttered terrain, single panoramic views have useful catchment areas compared with the scale of natural routes, which for *M. bagoti* have been observed to be of the order of 20 m (Muser et al., 2005). This is true when we look at the IDFs for views that are aligned (Fig. 2); it is also true when we use a stored view as way of recovering an orientation (Fig. 3). At first glance it seems surprising that an environment such as transect 4, where the world changes significantly with movement (Fig. 1B), could contain any

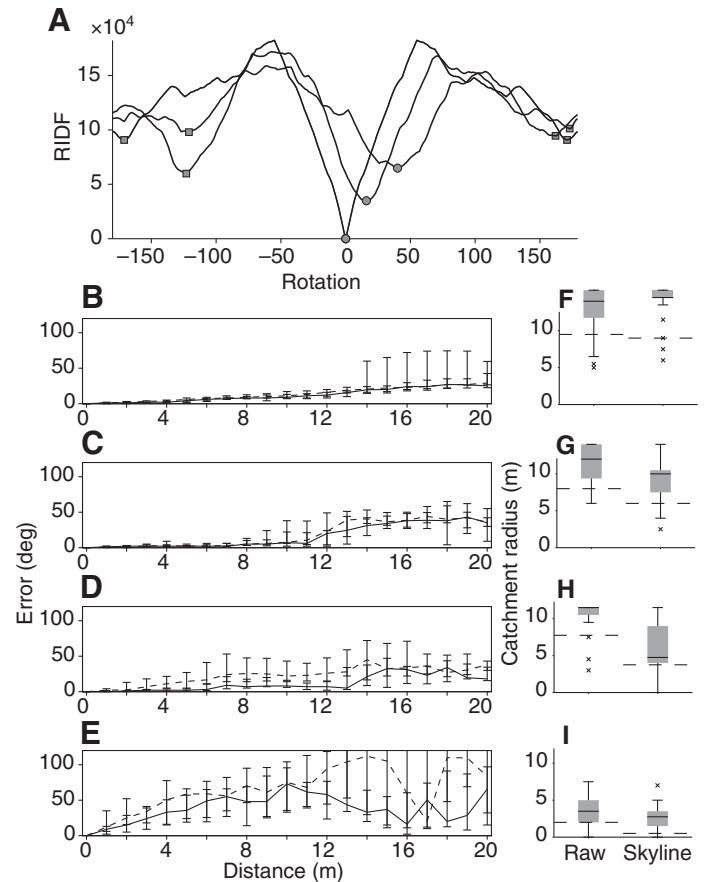


Fig. 4. Catchment areas based on augmented rotational IDFs (RIDFs). (A) RIDFs for a reference image (line where RIDF drops to 0) and two route locations (5 m apart). For each RIDF the three most prominent minima are marked. (B–E) Median absolute orientation errors as a function of distance from the goal when using the nearest minimum heuristic. We collated absolute orientation errors as a function of distance from the goal over all goal locations and illustrate the distributions of the errors by showing the median values, together with 25th and 75th percentiles as error bars, for the whole image (solid line) and skyline only (dashed line). $N > 7$ for all distances. (F–I) Box plots showing the median radius of catchment areas for all positions on each route when using the nearest minimum heuristic (box: 25th and 75th percentile; central line: median; whiskers: extent of the data; crosses: outliers, defined as points which are beyond the quartiles by more than 1.5 times the inter-quartile range). Data are shown for unprocessed images (left) and skyline (right). For each box the median rotational catchment area (RCA) without the nearest minimum heuristic is shown with a horizontal dashed line. For both whole image and skyline on all four transects the median RCA is improved when using the nearest minimum heuristic.

simple to use visual information. However, inspection of the panoramic scene from along these routes shows the information available. Fig. 5A shows the smooth and gradual change of the skyline in an open environment which underpins a smooth gradient in the IDF and accurate rotational matching (e.g. Fig. 2A,E,I and Fig. 3A,E,I). Fig. 5B shows the changing skyline for a cluttered route. Although the skyline changes rapidly, we still see sequences where features in the skyline persist and move slowly within the visual scene. These transiently stable features are enough to underpin gradients in the IDF (Fig. 2H).

Basten and Mallott (Basten and Mallott, 2010) have also investigated the properties of natural scenes from an ants' perspective by building a virtual simulation of a patch of desert

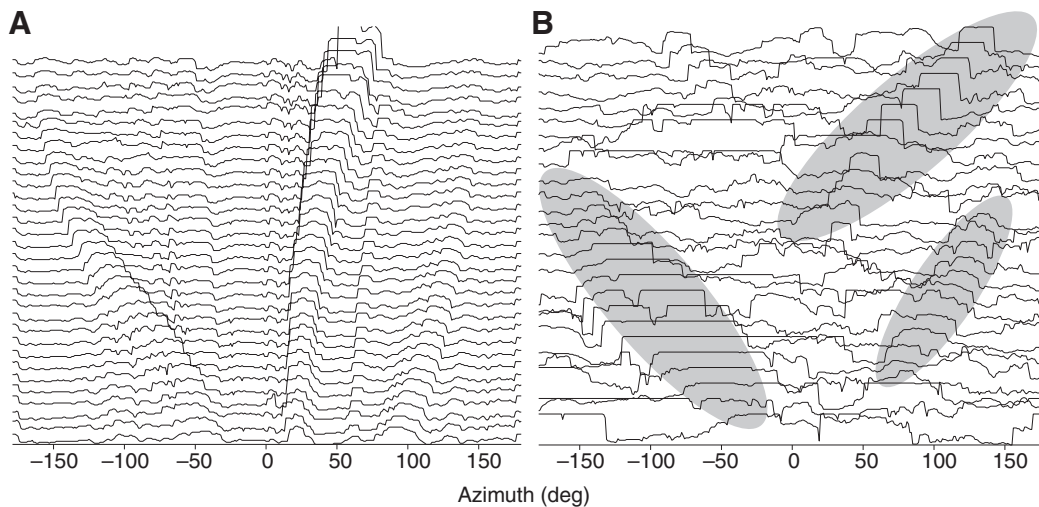


Fig. 5. Skylines along open and cluttered routes. Stacked skylines, showing the skylines for all locations where images were recorded on the open (A) and cluttered (B) transects. In B the highlighted regions show sequences of locations where a visual feature is persistent and moving smoothly across the visual field, thus providing information that could underpin view-based homing.

used for experiments by Kohler and Wehner (Kohler and Wehner, 2005). They were able to show the utility of a skyline code for uniquely defining a place. Their simulation, however, only included grass tussocks local to the region of interest and the reported catchment areas were around 2 m. A more realistic simulation of this particular semi-arid environment would have included medium (bushes) and large (trees) objects at a variety of distances (Muser, 2005). Zeil and Stürzl, in accordance with our Fig. 5, have shown that, in natural environments, a rich depth structure underpins robust view-based homing and directly influences catchment areas (Zeil and Stürzl, 2007). Therefore, it is likely that a more realistic world model would have given different results.

By analysing the information available in natural panoramic scenes we have shown that ants might not require a large set of views corresponding to a dense series of locations in order to control a habitual route. In the next section we discuss possible route guidance mechanisms that utilise remembered views.

Implications for route guidance mechanisms

We have shown that stored panoramic scenes from natural environments contain information that can be utilised over reasonable distances for descent in image difference (a proxy for view-based homing methods) and as a visual compass to recover an orientation. These two uses of stored images represent two very different mechanisms by which an ant could control a route. With the first, one would use a view-based homing algorithm to navigate to a sequence of views that are stored from points along the route. However, this intuitive strategy may not be as straightforward as it seems. Images need to be aligned to an external frame of reference requiring neural or behavioural mechanisms. Moreover, attempts to model route guidance as the sequential matching of a series of views (Smith et al., 2007; Smith et al., 2008; Vardy, 2006) have revealed non-trivial issues such as knowing when a sub-goal has been reached and reliably crossing boundaries into the catchment area of the next stored view. It may be that chaining snapshots is overly complicated for route guidance. Routes do not require the accuracy inherent in view-based homing (Collett et al., 1992). Rather, routes require simply the recall of headings appropriate to the current world location. Relevant to this, we have shown that stored natural views contain enough information to be used to recall route-appropriate headings (Zeil et al., 2003; Labrosse, 2006) without a global compass reference. Another attractive property of a mechanism that uses stored views to recall an orientation is that information from comparisons with multiple views can be sensibly pooled. For

instance, stored views that represent broadly similar directions could be simultaneously compared with the current view. Heading could then be set by some average, perhaps weighted by similarity, of the outputs across the multiple comparisons. Indeed, a mechanism like this, but with images stored at different headings, can even be used to home to a discrete location (Graham et al., 2010).

Using a stored image to retrieve an orientation is a neat way of controlling a route and, as described above, using this method it would be possible to combine results from simultaneous comparisons with multiple stored views. However, the problem of how to select the most appropriate views to define a route still remains. For instance, an arbitrarily chosen view along a route may contain objects that only appear transiently in the view sequence, which will reduce its effective catchment area. One possibility is that the views along a route segment could be averaged into a single view to remove noisy transient components to focus on the low frequency signal that should robustly define the correct orientation for that portion of the route. Given appropriate processing during route learning, it might be possible to learn a function that maps properties of a changing scene to appropriate headings along the route. Recent abstract models of route guidance based on experimental findings in ants (Harris et al., 2007; Collett, 2010) have shown how route control can be simplified when the route can be described by a smooth function that maps changes in (some aspect of) the visual scene onto a navigational instruction. Theoretical investigation into whether a moving average could extract a useful description of the views experienced along complex natural routes promises to be fascinating.

How to encode visual scenes

Inspired by recent behavioural findings with bees (Towne and Moscrip, 2008) and ants (Graham and Cheng, 2009a; Graham and Cheng, 2009b), we asked whether encoding panoramic scenes as 1D skylines influenced their catchment areas. Our results show that there is not a significant loss in performance between whole image and skyline when performing gradient descent on the IDF. For recovering orientation using the skyline, there is a small drop in performance because aliasing of prominent features is more likely when they are defined solely by their height, disregarding intensity and shape information. Despite this, there are compelling reasons why encoding panoramic scenes as skylines may be functionally successful. The skyline is an economical encoding that is probably easy for the insect visual system to perform (Möller, 2002) and is a robust parameterisation of a complex panorama which is

independent of lighting conditions and time of day. We have shown how simple behavioural heuristics can alleviate some aliasing problems, therefore rendering skyline encoding a plausible explanation of how ants parameterise natural scenes.

Insect navigation is a beautiful behaviour that continues to provide insight into insect perception and cognition. We believe a necessary complement to future behavioural studies of navigation is an understanding of the sensory ecology within which an animal behaves. To this end we have analysed ants' perspective views of a habitat within which desert ant navigation is well studied. Our analysis points to specific route guidance mechanisms that we can look for in future behavioural experiments.

ACKNOWLEDGEMENTS

We thank Matthew Collett for helpful discussions about this work. The work was supported by the Australian Research Council (DP0770300 to K.C.), The Royal Society, and the BBSRC and EPSRC through the Cognitive Foresight scheme (BBF0100521).

REFERENCES

- Åkesson, S. and Wehner, R. (2002). Visual navigation in desert ants *Cataglyphis fortis*: are snapshots coupled to a celestial system of reference? *J. Exp. Biol.* **205**, 1971-1978.
- Andel, D. and Wehner, R. (2004). Path integration in desert ants, *Cataglyphis*: how to make a homing ant run away from home. *Proc. Biol. Sci.* **271**, 1485-1489.
- Baddeley, B., Graham, P., Philippides, A. and Husbands, P. (in press). Holistic visual encoding of ant-like routes: navigation without waypoints. *Adapt. Behav.* Doi: 10.1177/1059712310395410.
- Basten, K. and Mallot, H. A. (2010). Simulated visual homing in desert ant natural environments: efficiency of skyline cues. *Biol. Cybern.* **102**, 413-425.
- Cartwright, B. A. and Collett, T. S. (1983). Landmark learning in bees. *J. Comp. Physiol. A*, **151**, 521-543.
- Cartwright, B. A. and Collett, T. S. (1987). Landmark maps for honeybees. *Biol. Cybern.* **57**, 85-93.
- Cheng, K., Narendra, A., Sommer, S. and Wehner, R. (2009). Traveling in clutter: Navigation in the Central Australian desert ant *Melophorus bagoti*. *Behav. Processes* **80**, 261-268.
- Collett, M. (2010). How desert ants use a visual landmark for guidance along a route. *Proc. Natl. Acad. Sci. USA* **107**, 11638-11643.
- Collett, M., Collett, T. S., Bisch, S. and Wehner, R. (1998). Local and global vectors in desert ant navigation. *Nature*, **394**, 269-271.
- Collett, T. S., Dillmann, E., Giger, A. and Wehner, R. (1992). Visual landmarks and route following in desert ants. *J. Comp. Physiol. A*, **170**, 435-442.
- Collett, T. S., Collett, M. and Wehner, R. (2001). The guidance of desert ants by extended landmarks. *J. Exp. Biol.* **204**, 1635-1639.
- Durier, V., Graham, P. and Collett, T. S. (2003). Snapshot memories and landmark guidance in wood ants. *Curr. Biol.* **13**, 1614-1618.
- Fourcassié, V. (1991). Landmark orientation in natural situations in the red wood ant *Formica lugubris* Zett (Hymenoptera: Formicidae). *Ethol. Ecol. Evol.* **3**, 89-99.
- Fukushi, T. (2001). Homing in wood ants, *Formica japonica*: use of the skyline panorama. *J. Exp. Biol.* **204**, 2063-2072.
- Graham, P. and Cheng, K. (2009a). Which portion of the natural panorama is used for view-based navigation in the Australian desert ant? *J. Comp. Physiol. A*, **195**, 681-689.
- Graham, P. and Cheng, K. (2009b). Ants use the panoramic skyline as a visual cue during navigation. *Curr. Biol.* **19**, R935-R937.
- Graham, P. and Collett, T. S. (2002). View-based navigation in insects: how wood ants (*Formica rufa* L.) look at and are guided by extended landmarks. *J. Exp. Biol.* **205**, 2499-2509.
- Graham, P., Fauria, K. and Collett, T. S. (2003). The influence of beacon-aiming on the routes of wood ants. *J. Exp. Biol.* **206**, 535-541.
- Graham, P., Durier, V. and Collett, T. S. (2004). The binding and recall of snapshot memories in wood ants (*Formica rufa* L.). *J. Exp. Biol.* **207**, 393-398.
- Graham, P., Philippides, A. and Baddeley, B. (2010). Animal Cognition: Multi-modal Interactions in Ant Learning. *Curr. Biol.* **15**, R639-R640.
- Harris, R. A., Graham, P. and Collett, T. S. (2007). Visual cues for the retrieval of landmark memories by navigating wood ants. *Curr. Biol.* **17**, 93-102.
- Kohler, M. and Wehner, R. (2005). Idiosyncratic route-based memories in desert ants, *Melophorus bagoti*: How do they interact with path-integration vectors? *Neurobiol. Learn. Mem.* **83**, 1-12.
- Labrosse, F. (2006). The visual compass: Performance and limitations of an appearance-based method. *J. Field. Robot.* **23**, 913-941.
- Möller, R. (2002). Insects that exploit UV-green contrast for landmark navigation. *J. Theor. Biol.* **214**, 619-632.
- Muser, B. A., Sommer, S. B., Wolf, H. A. and Wehner, R. (2005). Foraging ecology of the thermophilic Australian desert ant, *Melophorus bagoti*. *Aus. J. Zoo.* **53**, 301-311.
- Narendra, A., Si, A., Sulikowski, D. and Cheng, K. (2007). Learning, retention and coding of nest-associated visual cues by the Australian desert ant, *Melophorus bagoti*. *Behav. Ecol. Sociobiol.* **61**, 1543-1553.
- Pratt, S., Brooks, S. and Franks, N. (2001). The use of edges in visual navigation by the ant *Leptothorax albigipennis*. *Ethology* **107**, 1125-1136.
- Rosengren, R. and Fortelius, W. (1986). Ortstreue in foraging ants of the *Formica rufa* group-Hierarchy of orienting cues and long-term memory. *Insect. Soc.* **33**, 306-337.
- Santschi, F. (1913). Comment s'orientent les fourmis. *Rev. Suisse. Zool.* **21**, 347-425.
- Smith, L., Philippides, A., Graham, P., Baddeley, B. and Husbands, P. (2007). Linked local navigation for visual route guidance. *Adapt. Behav.* **15**, 257.
- Smith, L., Philippides, A., Graham, P. and Husbands, P. (2008). Linked local visual navigation and robustness to motor noise and route displacement. In *Animals to Animats 10, SAB 2008, LNCS*, Vol. 5040, pp. 179-188 (ed. M. Asada, J. C. Hallam, J.-A. Meyer and J. Tani). Berlin, Heidelberg: Springer-Verlag.
- Stürzl, W. and Zeil, J. (2007). Depth, contrast and view-based homing in outdoor scenes. *Biol. Cybern.* **96**, 519-531.
- Towne, W. F. and Moscrip, H. (2008). The connection between landscapes and the solar ephemeris in honeybees. *J. Exp. Biol.* **211**, 3729-3736.
- Vardy, A. (2006). Long-range visual homing. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics*, pp. 220-226. http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4141868
- Wehner, R. and Rüber, F. (1979). Visual spatial memory in desert ants, *Cataglyphis bicolor* (Hymenoptera: Formicidae). *Cell. Mol. Life. Sci.* **35**, 1569-1571.
- Wehner, R., Michel, B. and Antonsen, P. (1996). Visual navigation in insects: coupling of egocentric and geocentric information. *J. Exp. Biol.* **199**, 129-140.
- Wolf, H. and Wehner, R. (2000). Pinpointing food sources: olfactory and anemotactic orientation in desert ants. *J. Exp. Biol.* **203**, 857-868.
- Zeil, J., Hoffmann, M. I. and Chahl, J. S. (2003). Catchment areas of panoramic images in outdoor scenes. *J. Opt. Soc. Am. A*, **20**, 450-469.