

# Adaptive Behavior

<http://adb.sagepub.com/>

---

## What is the relationship between visual environment and the form of ant learning-walks? An in silico investigation of insect navigation

Alex DM Dewar, Andrew Philippides and Paul Graham  
*Adaptive Behavior* published online 10 February 2014  
DOI: 10.1177/1059712313516132

The online version of this article can be found at:  
<http://adb.sagepub.com/content/early/2014/02/10/1059712313516132>

---

Published by:



<http://www.sagepublications.com>

On behalf of:

# ISAB

International Society of Adaptive Behavior

Additional services and information for *Adaptive Behavior* can be found at:

**Email Alerts:** <http://adb.sagepub.com/cgi/alerts>

**Subscriptions:** <http://adb.sagepub.com/subscriptions>

**Reprints:** <http://www.sagepub.com/journalsReprints.nav>

**Permissions:** <http://www.sagepub.com/journalsPermissions.nav>

**Citations:** <http://adb.sagepub.com/content/early/2014/02/10/1059712313516132.refs.html>

>> [OnlineFirst Version of Record](#) - Feb 10, 2014

[What is This?](#)

# What is the relationship between visual environment and the form of ant learning-walks? An in silico investigation of insect navigation

Alex DM Dewar<sup>1</sup>, Andrew Philippides<sup>2</sup> and Paul Graham<sup>1</sup>

Adaptive Behavior  
1–17

© The Author(s) 2014

Reprints and permissions:

sagepub.co.uk/journalsPermissions.nav

DOI: 10.1177/1059712313516132

adb.sagepub.com



## Abstract

The learning walks of ants are an excellent opportunity to study the interaction between brain, body and environment from which adaptive behaviour emerges. Learning walks are a behaviour with the specific function of storing visual information around a goal in order to simplify the computational problem of visual homing, that is, navigation back to a goal. However, it is not known at present why learning walks take the stereotypical shapes they do. Here we investigate how learning-walk form, visual surroundings and the interaction between the two affect homing performance in a range of virtual worlds when using a simple view-based homing algorithm. We show that the ideal form for a learning walk is environment-specific. We also demonstrate that the distant panorama and small objects at an intermediate distance, particularly when the panorama is obscured, are important aspects of the visual environment both when determining the ideal learning walk and when using stored views to navigate. Implications are discussed in the context of behavioural research into the learning walks of ants.

## Keywords

Insect navigation, learning walks, view-based homing, visual navigation in ants

## 1 Introduction

The ability to navigate between important locations is found in almost all mobile organisms, and, because it is so widespread, navigation provides a valuable benchmark for understanding general principles of cognitive organization across taxa (e.g. Cheng, 2010). For a researcher, navigational behaviours are particularly amenable to study. There is a clearly defined aim (i.e. ‘move towards the goal’) and animals’ internal ‘state’ is evident in directions of motion. This close link between control system, body, behaviour and world sits well with the ‘embodied’ view of cognition, where body, brain and environment are treated as a single system, with behaviour as an emergent property (Brooks, 1999; Pfeifer & Scheir, 1999). Thus, even apparently complex cognitive systems can often be reduced to a set of simple, task-specific competencies (Brooks, 1999). Indeed, such a basic ‘toolkit’ of abilities is thought to underlie navigation in ants (see Wehner, 2008). Here we leverage the completeness of the ant navigation model to investigate the relationship between visual homing and active learning.

Visual homing, which is employed by many species (Wang & Spelke, 2002; Wiener et al., 2011), involves an

agent storing knowledge of how the world appears from a goal location so that future returns can be guided by comparisons between the current view and the stored view(s). How these stored views are learned initially is poorly understood (reviews: Collett, Graham, Harris, & Hempel de Ibarra, 2006; Zeil, 2011), though in insects, learning is thought to take place mostly during the stereotypical orientation behaviours performed on leaving the nest or a food-source (Collett, 1995; Collett & Lehrer, 1993; Jander, 1997; Zeil, 1993a, 1993b); in ants these are known as ‘learning walks’. The first time an ant leaves the nest or food-source is when she will spend longest performing her learning walk; the duration declines thereafter (Graham & Collett, 2006; Nicholson, Judd, Cartwright, & Collett, 1999; Wehner, Meier, & Zollikofer, 2004). A large change in the visual

<sup>1</sup>School of Life Sciences, University of Sussex, Falmer, UK

<sup>2</sup>Department of Informatics, University of Sussex, Falmer, UK

### Corresponding author:

Alex DM Dewar, School of Life Sciences, University of Sussex, Falmer, BN1 9QG, UK.

Email: a.dewar@sussex.ac.uk

surroundings, as with the introduction of a prominent new object, results in a resurgence of these behaviours (Müller & Wehner, 2010). During learning walks and flights, it is not known whether learning takes place continuously or at discrete points (Collett & Lehrer, 1993), though the latter is suggested by the presence of multiple discrete points within a learning walk where the ant faces the nest, something which is seen in the learning walks of multiple species. For instance, *Ocymyrmex robustior* ants intermittently pause and turn to face the (invisible) nest entrance during their learning walks (Müller & Wehner, 2010). We know that these learning behaviours are almost certainly tailored to the surroundings (Wystrach, personal communication, 4 November 2013; see also Wei & Dyer, 2009; Wei, Rafalko, & Dyer, 2002), although the details are not known. Here, we are interested in the theoretical question of whether there are optimal places to store discrete nest-oriented views and how this depends on visual environment.

In order to understand how learning walks provide suitable visual information, we need to understand how visual information can be used for homing. The mechanisms underlying visual homing can be simple. For instance, view-based navigation (i.e. navigation to the location where a single view was stored) can be carried out by performing gradient descent on the output of a basic image difference function (IDF) between the stored goal view and aligned views from nearby locations (Philippides, Baddeley, Cheng, & Graham, 2011; Zeil, Hofmann, & Chahl, 2003), as, even in natural scenes, image difference increases smoothly with distance (Zeil et al., 2003). An alternative strategy is to use a stored view to recall a heading, by making comparisons between the goal views and the current view *at all possible rotations*, then heading in the direction for which the difference is at a minimum: this is the rotational image difference function (RIDF) (Philippides et al., 2011; Wystrach, Mangan, Philippides, & Graham, 2013; Zeil et al., 2003). With the RIDF, combining the information from multiple views into a single heading is computationally simple (e.g. with a weighted average), making it more appropriate for modelling learning walks with multiple ‘nest saccades’, such as those in Müller and Wehner (2010). Moreover, the types of movement required by the RIDF algorithm (forward locomotion and sampling multiple directional views) match up well with the ways in which ants move while navigating. A forager can sample the world over a range of directions either by stopping and turning on the spot, before heading in a new direction (Philippides et al., 2011; Wystrach, Philippides, Cheng, & Graham, submitted) or as a by-product of her typically sinuous path (Lent, Graham, & Collett, in press). Evidence is mounting that such scanning is a mechanism for obtaining a heading based on visual information and is accordingly performed far more frequently by ants placed in unfamiliar environments.

We investigated learning walks with a series of in silico experiments examining the relationship between the locations where visual information is acquired and subsequent homing success. The aims of this study are, first, to ask what distinguishes a ‘good’ from a ‘bad’ learning walk, as defined by how easy or difficult it renders subsequent homing efforts, and, second, to better understand the interaction between learning walks and environment. In this way we will shed light on the navigational strategies used by ants in practice.

## 2 Methods

### 2.1 Virtual reality system

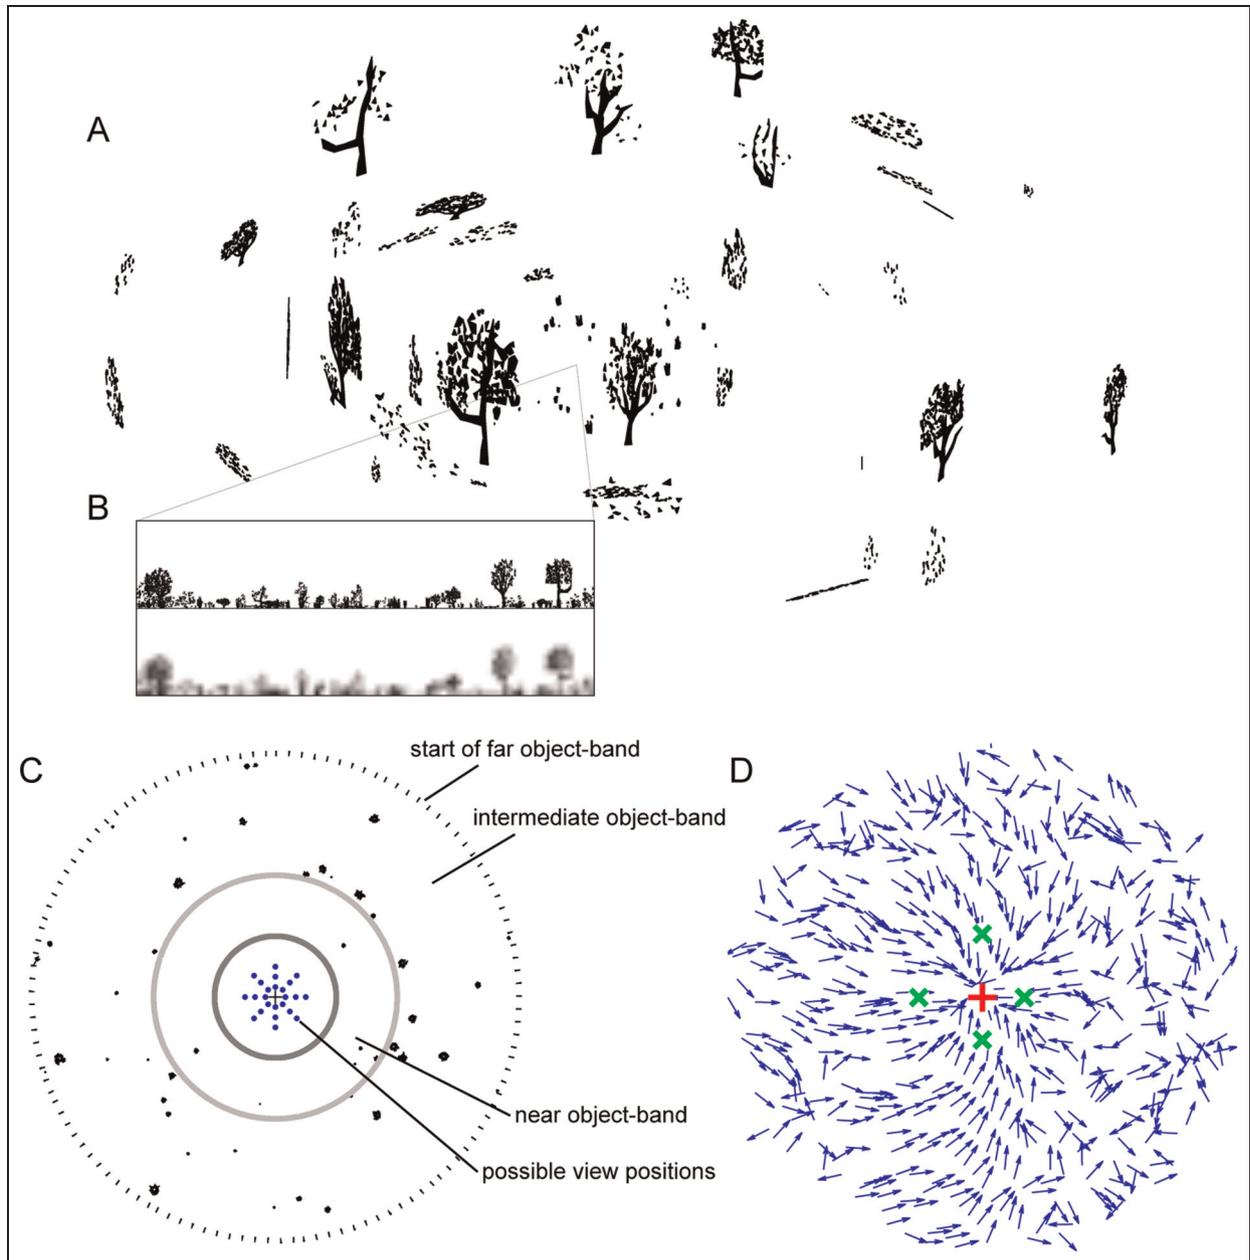
The virtual reality (VR) system we use is the same as that presented in Baddeley, Graham, Husbands, and Philippides (2012). Briefly, the generated environments, or ‘worlds’, are designed to resemble the visually sparse, semi-arid habitat of *Melophorus bagoti* and so consist of a random assortment of tussocks, bushes and trees (Figure 1A). Matlab® (MathWorks, Natick, MA, USA) is used throughout, both to generate the worlds and to perform testing and analysis.

Views recorded from worlds are panoramic, as is (almost) the case for real ants’ eyes, covering 360° in azimuth and 68° in elevation. Acuity is reduced to an ‘ant level’ by local averaging so that each pixel is equivalent to 4° of visual angle (Schwarz, Narendra, & Zeil, 2011; Zollikofer, Wehner, & Fukushi, 1995). The visual ‘objects’ (tussocks, bushes and trees) are generated from sets of pre-defined triangles in random configurations. They are also subject to random rotation and reflection. Both their relative size and distribution within worlds are roughly to scale; accordingly, we describe distances in this paper in terms of metres. The tussocks are rendered in 3d, whereas distant bushes and trees were 2d objects placed in the 3d world as they were sufficiently far from the portion of the world where testing was performed that full 3d was redundant. All the images we use are greyscale.

In all worlds, the centre (0, 0) is the ‘nest entrance’ or goal. The ‘far object-band’, a random assortment of bushes and trees, begins at a radius of 12 m from the nest entrance. The number of such objects varies randomly: for each of 99 iterations of a loop there was a probability of 0.5 for whether the object was placed or not. Tussocks are placed within the 12 m radius. Although their exact placement varies between experiments, they are never placed within 3 m of the nest entrance (Figure 1C). For each world we use images from a grid of 9845 positions,  $x_j$ , which correspond to locations ~6 cm apart within a radius of 3 m.

### 2.2 Navigation algorithm

**2.2.1 From image differences to a new heading.** Any two images can be compared with an IDF in order to obtain a measure for the degree of difference. The IDF here is



**Figure 1.** One example world, from various perspectives, with homing performance shown for a particular SVP. A: Example world. B: Panoramic views from the centre of the world. High resolution (top); resolution at a level comparable to that of an ant's eye, as used in the homing algorithm (bottom). C: Bird's-eye view of the centre portion of the world. Outside of the outermost circle is the far object-band (not shown). The nest entrance is denoted with a cross in the centre and possible snapshot positions are indicated by dots. D: An example of successful homing within the world. The SVP used is indicated by the  $\times$ s and the  $+$  in the centre indicates the nest entrance.

r.m.s. pixel difference (Zeil et al., 2003), chosen for its simplicity. Other possibilities, such as mutual information (Kim, Szenher, & Webb, 2009), exist but had a negligible difference on results when tested (presumably because our images have uniform contrast).

For a position  $x_1$  at rotation  $\phi$  in world  $W$ , the view will be  $V(W, x_1, \phi)$ , with the pixel in column  $m$  and row  $n$  denoted as  $V(W, x_1, \phi)_{m,n}$ . The r.m.s. image difference between two views is:

$$d(W, x_1, \phi, x_2, \theta) = \sqrt{\frac{\sum_{m=1}^w \sum_{n=1}^h (V(W, x_1, \phi)_{m,n} - V(W, x_2, \theta)_{m,n})^2}{wh}} \quad (1)$$

where  $w$  is image width (here: 90 px) and  $h$  is image height (here: 17 px).

For each world, views from 24 different positions ( $s_i$ ) facing the nest entrance are set as potential snapshots. These view positions are arranged on a radial grid centred on the nest, at three different radii (0.5 m, 1 m and 1.5 m) and eight different (evenly spaced) angles (Figure 1C). A set of view positions (SVP) uses four of the possible 24, with the further constraint that only one view can be on each radial arm of the grid. This gives a total of 5670 potential SVPs. The choice of four stored views follows preliminary tests that showed that homing performance plateaus with increasing number of views after four, and aligns with Graham, Philippides, and Baddeley (2010).

To get a heading from a stored view, at position  $s_i$ , we use the RIDF (Philippides et al., 2011; Zeil et al., 2003), which involves making multiple IDF comparisons between the stored and current views with the current view rotated incrementally through  $360^\circ$ . This yields a range of difference values, with the angle at which the minimum occurs indicating the best-matching direction:

$$r(W, s_i, \mathbf{x}) = \min_{\theta} d(W, s_i, \mathbf{x}, \theta) \quad (2)$$

$$\hat{h}(W, s_i, \mathbf{x}) = \arg \min_{\theta} d(W, s_i, \mathbf{x}, \theta) \quad (3)$$

where  $\theta \in \{4^\circ, 8^\circ, \dots, 360^\circ\}$ ,  $r(W, s_i, \mathbf{x})$  is the RIDF minimum value and  $\hat{h}(W, s_i, \mathbf{x})$  is the corresponding heading. Note that we lose the dependency on  $\phi$  as stored views are always oriented toward the nest.

To use an SVP  $S$  for homing, RIDFs comparing the current view  $V(W, \mathbf{x}_j, \theta)$  with each of the snapshot views, at position  $s_i$ , are calculated giving  $r(W, s_i, \mathbf{x}_j)$  and  $\hat{h}(W, s_i, \mathbf{x}_j)$  for each  $s_i \in S$ . The four headings  $\hat{h}(W, s_i, \mathbf{x}_j)$ , each an estimate of the direction to the nest, are then combined by weighting items by their ‘goodness of match’,  $w_i$ , where

$$r_{\min}(W, S, \mathbf{x}_j) = \min_{s_i \in S} r(W, s_i, \mathbf{x}_j)$$

$$w(W, S, s_i, \mathbf{x}_j) = \frac{r_{\min}(W, S, \mathbf{x}_j)}{r(W, s_i, \mathbf{x}_j)}$$

Hence, by definition, a weight of one is given to the best-matching view, with other views’ contributions to the heading weighted according to their ‘goodness of match’ relative to this benchmark. The weighted circular mean is then calculated as follows (circular statistics toolbox for Matlab: Berens, 2009):

$$\hat{H}(W, S, \mathbf{x}_j) = \arg \left( \sum_{s_i \in S} w(W, S, s_i, \mathbf{x}_j) \cdot \exp(\hat{h}(W, s_i, \mathbf{x}_j) \cdot \iota) \right)$$

where  $\iota$  is the imaginary unit. Other combinations of weightings were trialled with no qualitative effects on results.

### 2.3 Calculating the error on the homeward component

A measure commonly used for determining performance on a homing task is the average homeward

component (AHC; see Batschelet, 1981), which varies from zero, for an estimated heading  $\geq 90^\circ$  from the true heading, to one, for a completely accurate estimated heading. Here, so as to have an error rather than a performance measure, we use a slight variant, the error on the homeward component (EHC), which we define as one minus the AHC. Hence, the error  $E(W, S, \mathbf{x}_j)$  at a given position in the world  $\mathbf{x}_j$  for a heading  $\hat{H}(S, \mathbf{x}_j)$  is

$$\varepsilon(W, S, \mathbf{x}_j) = |\hat{H}(W, S, \mathbf{x}_j) - H(\mathbf{x}_j)| \quad (4)$$

$E(W, S, \mathbf{x}_j) =$

$$\begin{cases} 1 - \cos(\varepsilon(W, S, \mathbf{x}_j)) & \text{for } \varepsilon(W, S, \mathbf{x}_j) < 90^\circ \\ 1 & \text{for } \varepsilon(W, S, \mathbf{x}_j) \geq 90^\circ \end{cases} \quad (5)$$

where  $\varepsilon(W, S, \mathbf{x}_j)$  is the angular difference between estimated heading,  $\hat{H}(W, S, \mathbf{x}_j)$ , and the true heading to the nest,  $H(\mathbf{x}_j)$ . This measure accordingly gives  $0 \leq E(W, S, \mathbf{x}_j) < 1$  when  $\varepsilon(W, S, \mathbf{x}_j) < 90^\circ$ , and  $E(W, S, \mathbf{x}_j) = 1$  otherwise. The value in doing this is that all ‘bad’ headings (i.e. those pointing away from the goal) are given a value of one, and we then distinguish only between ‘good’ headings.

The mean error over all positions  $\mathbf{x}_j$  for a set of SVP  $P$  in a world  $W$  is then

$$\bar{E}(W, P) = \frac{1}{9845 \cdot |P|} \sum_{S \in P} \sum_{j=1}^{9845} E(W, S, \mathbf{x}_j) \quad (6)$$

## 2.4 Experiments

**2.4.1 Initial experiments.** Ten randomly generated worlds were used for the initial experiments. In these worlds, there were six tussocks in the near object-band and 40 in the intermediate object-band. Placement in the near object-band was pseudo-random: the experimenter selected from among random configurations ones where the tussocks were relatively evenly spaced around the nest entrance. This was done to avoid generating unrealistically asymmetric environments that might lead to a strong directional bias. The tussocks in the intermediate object-band were placed completely at random. The far object-band contained 2d objects as described in Section 2.1.

**2.4.2 Varying the environment.** The purpose of these experiments was to investigate the relationship between components of the world, SVPs and homing performance. To this end, we divided worlds into three ‘object-bands’ (near: 3–6 m, intermediate: 6–12 m and far: >12 m) with different ‘types’ of cue. The near object-band contained one large tussock (height = 1.11 m) and the far object-band contained the horizon as before. In order to observe the effect of changing the ‘level of clutter’, we generated eight intermediate

object-bands that varied in number of tussocks (0, 33, 54, 75, 96, 117, 138 or 159;  $M_{\text{height}} = 0.550$ ,  $SD_{\text{height}} = 0.323$ ). To see the effect of objects in each object-band, we compared worlds in which objects in one band were rotated at angles of  $0^\circ$ ,  $90^\circ$ ,  $180^\circ$  and  $270^\circ$ . By choosing worlds that differ only in the rotation of one object-band, we can effectively ‘pit’ that band against the other two.

The logic of these experiments is perhaps best illustrated with an example. Imagine we have worlds  $A$ ,  $B$  and  $C$ , with  $B$  differing from  $A$  only in that the near object-band is rotated by  $+90^\circ$  and  $C$  differing from  $A$  only in that the intermediate and far object-bands are rotated by  $-90^\circ$ . If we have the set of the 100 ‘best’ SVPs for  $A$ ,  $G$ , the mean homing error for  $G$  in  $A$  is  $\bar{E}(A, G)$  (equation (6)). The change in error when the near band is rotated is then given by  $\epsilon_N = \bar{E}(A, G) - \bar{E}(B, G)$ , and when both the intermediate and far object-bands are rotated, by  $\epsilon_{I\&F} = \bar{E}(A, G) - \bar{E}(C, G)$ . Based on which of the error changes is greater for these two conditions, we can see whether the near object-band contributes more to which SVPs perform best than the intermediate and far object-bands combined (i.e.  $\epsilon_N > \epsilon_{I\&F}$ ) or vice versa ( $\epsilon_{I\&F} > \epsilon_N$ ).

**2.4.3 Descriptive measures for SVPs.** To analyse performance as a function of SVP, we used a number of simple measures.

**Spread of stored views.** For an SVP  $S$  the spread of stored views  $s(S)$  is

$$s(S) = \frac{\sum_{s_i \in S} \sum_{s_j \neq s_i} \|s_i - s_j\|}{n(n-1)}$$

where  $n$  is the number of stored views in the SVP (four in our case) and  $s_i$  is the coordinates of the  $i$ th view in  $S$ .

**View dissimilarity.** This is the mean of the r.m.s. differences between each of the images in the SVP,

$$D(W, S) = \frac{\sum_{s_i \in S} \sum_{s_j \neq s_i} d(W, s_i, 0, s_j, 0)}{n(n-1)}$$

where  $n$  is the number of stored views and  $d(W, s_i, 0, s_j, 0)$  is the r.m.s. difference (equation (1)) between the stored views at  $s_i$  and  $s_j$  in a world  $W$ .

**Nest distance.** This is the median distance of the stored views from the nest entrance.

**‘Surroundedness’.** This is the circular standard deviation (Batschelet, 1981) of the angles of the view positions,  $s_i$ , in the SVP relative to the nest. An SVP where the nest entrance has views distributed all around it will

score higher than one where the views are closer together.

**‘Oppositeness’.** This was defined as the number of pairs of stored views that were in line with each other and with the nest entrance; in other words, a straight line could be drawn connecting these three points. As four stored views were used in each SVP and only one can be on each radial arm, this means the possible values for this measure are zero, one and two.

### 3 Results

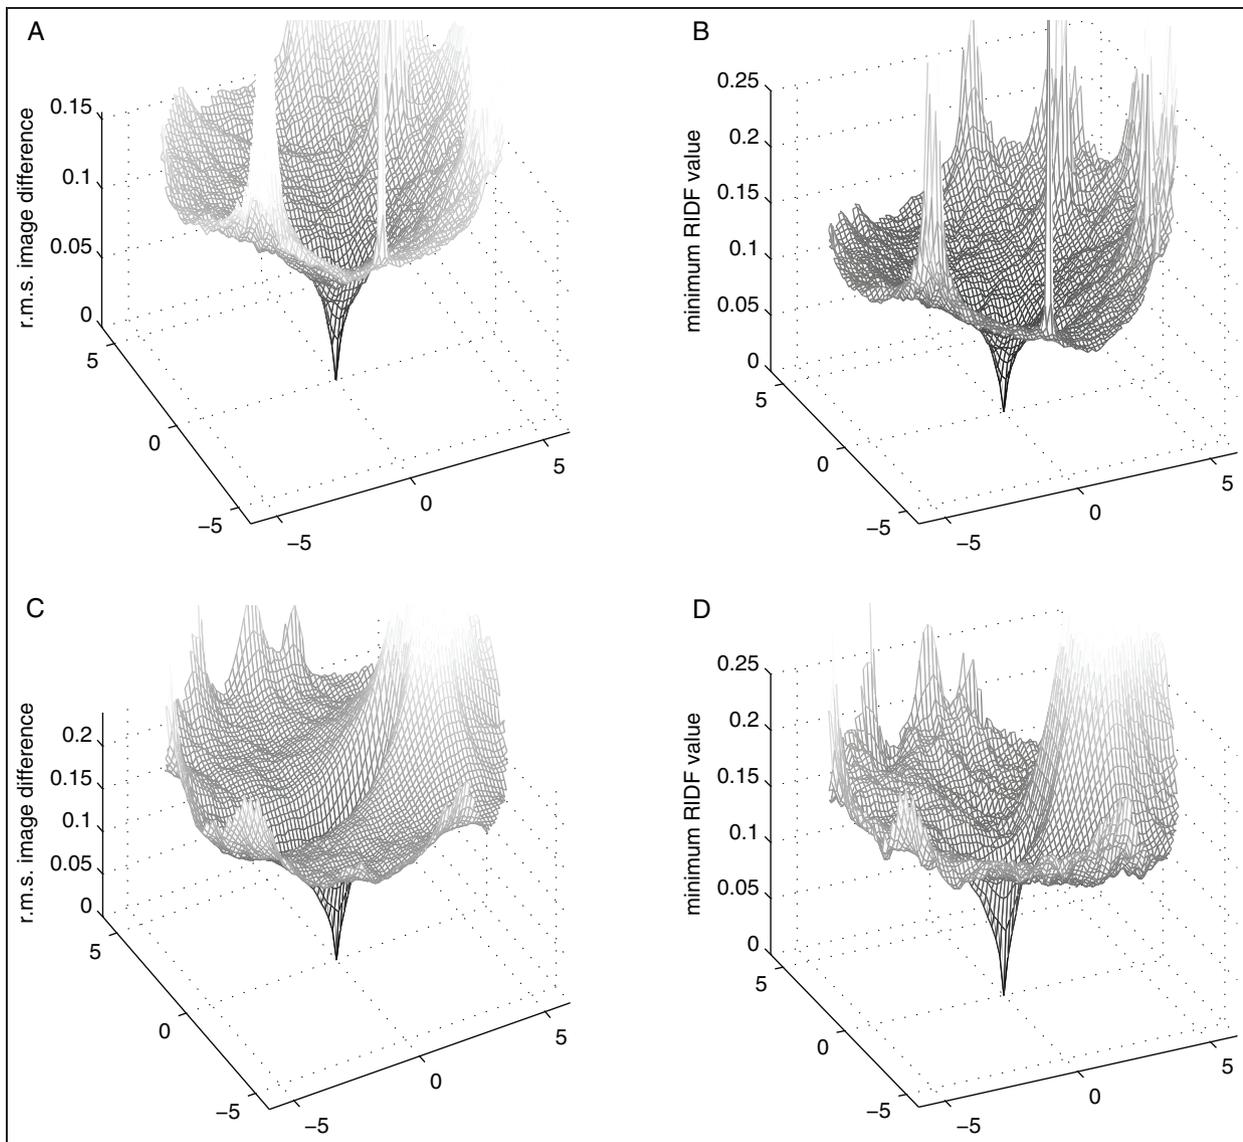
The aim of this study is to investigate what comprises a ‘good’ learning walk and to ask how this relates to the visual environment. The ultimate goal is to better understand what visual-homing strategies are used by ants. To do this, we implemented a simple, biologically plausible visual-homing method in VR worlds. We experimented with sets of views, asking how good a given set of views is for homing from locations throughout the environment. We also looked at how the visual surroundings affect the optimal placement of views. As we are investigating the form of learning walks, we care about where views are taken relative to the nest entrance, and we have dubbed this information a ‘set of view positions’. Note that the term refers only to view *locations*: the same SVP in two different worlds will contain different images as each world looks different.

#### 3.1 Natural image statistics of artificial worlds

We first verified that the image data drawn from these virtual worlds were qualitatively similar to ‘real-world’ data. Graphs of IDFs (Figure 2A,C) and RIDFs (Figure 2B,D) were found to resemble those generated from natural images (Philippides et al., 2011; Zeil et al., 2003). Specifically, there are smooth gradients in image difference as one moves away from the location of a goal image. The IDFs also give ‘catchment areas’, generally for several metres around the goal, within which an agent would always be able to return to the goal. Also, the goal image can be compared to rotated versions of images at other locations in order to recapture the orientation at the goal image over similar distances. Information from multiple locations can then be used to navigate back to the goal (Section 2; Figure 1D). This suggests we have a meaningful simulation and homing algorithm for investigating how learning walks depend on visual environment.

#### 3.2 Are SVPs world-specific?

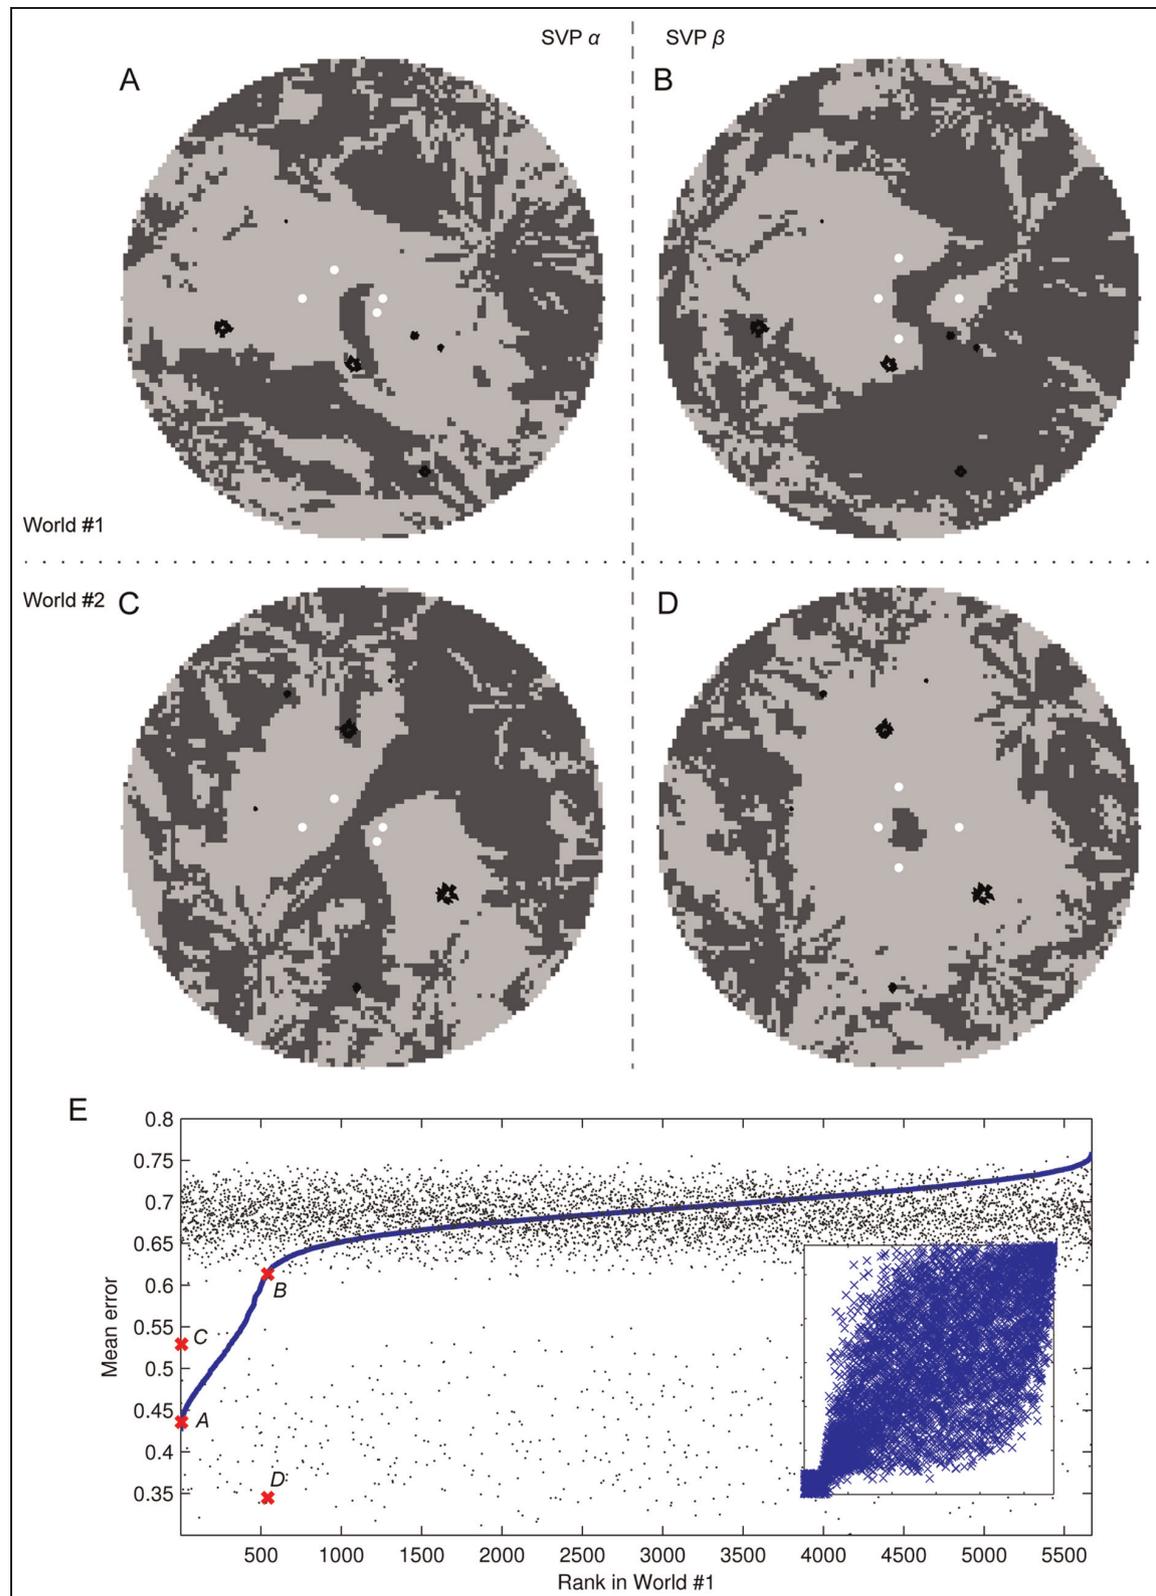
After verifying that a biologically plausible and parsimonious algorithm produces reliable homing in our environments (Figures 1 to 3), we turned to the question of how optimal view locations (SVPs) relate to



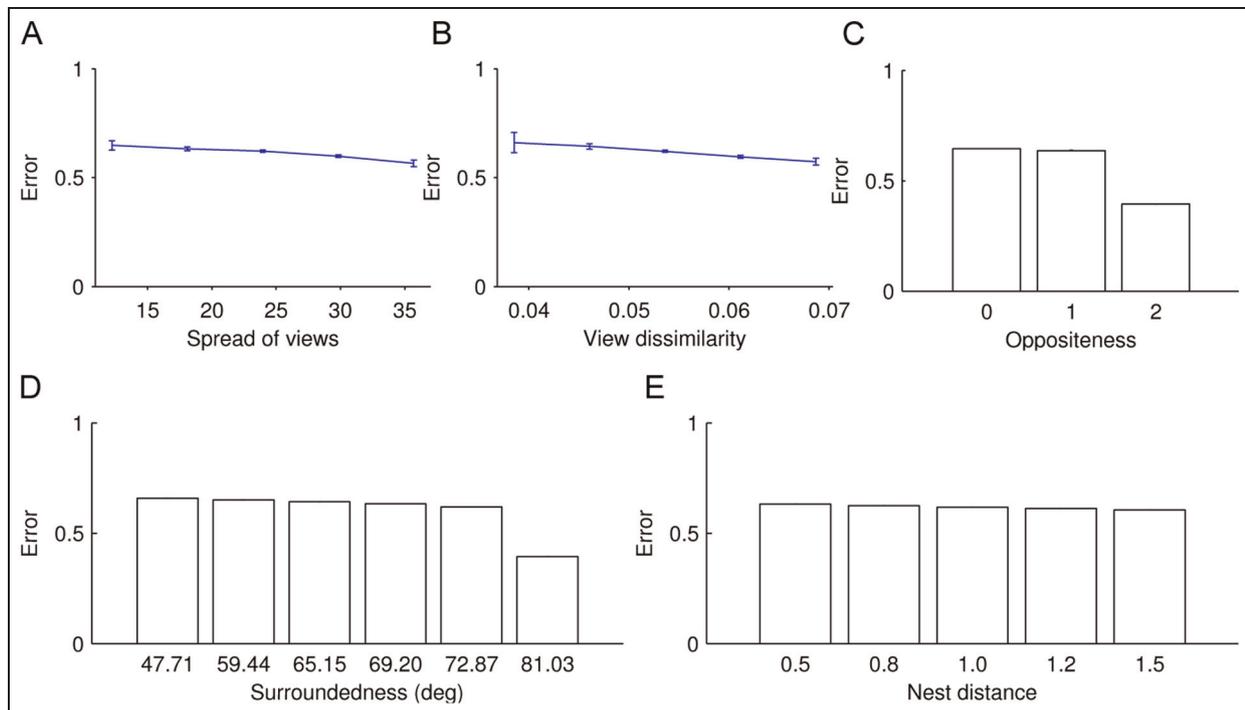
**Figure 2.** A comparison of IDF and RIDF minimum surfaces for two different worlds. A and B represent one world; C and D, another more complex world. The goal view was from the nest entrance in each of the worlds, that is, (0, 0). As the RIDF is the best match for all rotations, at any given point in space the score will be less than or equal to that for the IDF (i.e. with no rotation). Note that the RIDF graphs indicate the quality of the best match and not the best-matching direction. The best directions, however, were almost always within a few degrees of the orientation of the reference view, supporting the use of the visual compass for homing. For B, almost all headings aligned with the reference image (i.e. to 0°; mean error = 12.1°; SD = 30.3°), although for D there was more scatter (mean error = 69.9°; SD = 49.5°), presumably owing to the greater complexity of the environment.

visual environment. For two worlds we compared SVPs ranked by mean error (equation (6)). We found that, whilst an SVP that performed well in one world would not necessarily perform well in another, some SVPs performed reasonably well across worlds, albeit with much variability. This is shown in Figure 3, where homing success (defined as  $\varepsilon < 45^\circ$ ) is indicated by light grey and failure by dark grey. It can be seen that an SVP,  $\alpha$ , that performs well in World #1 (Figure 3A) does not perform well in World #2 (Figure 3C). A second SVP,  $\beta$ ,

however, performs poorly in World #1 (Figure 3B) but well in World #2 (Figure 3D). Note that placement of tussocks influences what regions of the world the agent is able to home from successfully: in cases where tussocks are directly in between the location and the nest entrance, performance is often poorer. We show that ranks are significantly different (ranks for SVPs in World #1 vs ranks in World #2: Figure 3E). As well as these specific examples (Worlds #1 and #2, SVPs  $\alpha$  and  $\beta$ ), we also compared ranks for all SVPs across 10 test



**Figure 3.** Is there an ideal SVP for all worlds? A–D: ‘Success/fail’ plots (success: light grey; fail: dark grey) for a range of locations in two worlds using two SVPs. Homing was deemed successful for locations where heading error was less than or equal to  $45^\circ$ . A and B are the results for one world; C and D are the results for another. The same SVP was used in A and C, and another was used in B and D. The positions of views are shown in white, and those of tussocks in black. E: error vs rank for World #1 (line; A and B) and the error for the same SVP in World #2 (dots; C and D). The positions on the graph for the different ‘tests’ (A–D) are labelled. Inset is a scatter plot of how the rank of SVP differs between the two worlds. The difference in median ranks of SVPs between worlds was significant: Friedman test,  $\chi^2(9, N = 51022) = 61.26, p < 10^{-9}$ . Kendall’s  $W$  was 0.59, indicating strong differences between worlds.



**Figure 4.** How does the form of an SVP relate to performance? For SVPs we looked at how their mean error (EHC) varied as a function of the descriptive measures described in Section 2.4.3. A–E: SVP performance versus ‘spread of views’, ‘view dissimilarity’, ‘oppositeness’, ‘nest distance’ and ‘surroundedness’, respectively;  $N(\text{worlds}) = 10$ ;  $N(\text{SVP}) = 5670$ . Error bars indicate standard error. Data in A and B are grouped into five unevenly sized bins because of uneven distribution of possible values for the measures.

worlds (Section 2.4.1). This indicates that SVP performance depends on the environment to some extent.

An interesting side point is that we can also see patches of unsuccessful homing centred on the nest entrance, which might seem counterintuitive. The problem is that once the agent is closer to the nest entrance than the view positions, *all* views will give a good match and the algorithm will not give a clear heading. Of course, this would not be an issue for ‘real’ ants, as close to the nest entrance other cues, such as conspecifics or carbon dioxide emanating from the nest (Buehlmann, Hansson, & Knaden, 2012), can be used, with systematic search as a fallback (e.g. Wehner & Srinivasan, 1981). We make no claim that visually guided homing alone suffices for navigation, although it is certainly a, if not *the*, critical part of the ant’s navigational repertoire (Cheng, Narendra, Sommer, & Wehner, 2009; Wehner, 2008).

We have shown that SVP success can be affected by the visual environment, although it seems intuitively sensible that there will be properties of SVPs that make them more likely to be successful, regardless of the visual environment; this is examined in the next section.

### 3.3 Properties of SVPs which perform well across worlds

To investigate how the form of SVPs relates to performance, we categorized SVPs with the measures described in Section 2.4.3. The aim was to test our

intuitions about what properties of an SVP might best facilitate homing. We initially investigated whether high-performing SVPs were frequently just rotated versions of one another, but this was not borne out (Monte Carlo simulation; data not shown). This suggests that the measures that best predict SVP performance will likely relate to the broad rather than the fine structure of SVPs.

The first property we examined was what we termed ‘spread of views’, or mean Euclidean distance between views (Figure 4A). One might expect that the best performance would be obtained from SVPs with more separated views, presumably giving a greater range of information about the environment; indeed, there was a clear trend in this direction. There is a similar pattern for ‘view dissimilarity’ (Figure 4B), a measure that should correlate strongly with spread of views. We next looked at ‘surroundedness’, or azimuthal spacing of view positions (Figure 4E). It might also be expected that SVPs that are more angularly spaced will also perform better, and, although the most spread-out group ( $S(\text{SVP}) = 81.03^\circ$ ) had a markedly lower mean error than the others ( $\sim 25\%$  less), the others did not differ from one another. Another, similar, measure is median distance of view positions from the nest (‘nest distance’). By the same reasoning as before, it seems plausible that an SVP with views further from the nest entrance will perform better. Yet there was only a very slight effect for nest distance (Figure 4D), with less

than 2% difference in mean error between the nearest and furthest groups. However, we note that in ants, the mechanism for orienting toward the nest is likely to be path integration and thus have an error dependent on distance. Among these similar measures, ‘spread of views’ seems the most useful. Another measure investigated is what we called ‘oppositeness’ (Figure 4C), or the number of pairs of views that are opposite each other (with the nest entrance in the middle); this one is perhaps slightly less intuitive but was found to be a reliable predictor of homing success. As we are using four views, the possible values are zero, one and two. For oppositeness (Figure 4C), there was only a 1% difference between mean error scores for SVPs with zero or one pairs of ‘opposite’ views, but a ~27% improvement in error scores between one and two pairs.

Of course, these measures are not independent of one another: views with a large angular distance between them are likely to be further apart in Euclidean space, too. The extremes of performance for different measures are probably being driven by many of the same SVPs. The purpose, then, of this section is not an exhaustive investigation of SVP properties, but rather an examination of whether there are any simple, easy-to-calculate measures that are predictive of the homing success for a given SVP.

Overall, our various measures show that views should be spread around the nest, though, interestingly, the distance of the views from the nest is not too important.

### 3.4 What aspects of a world ‘resonate’ with particular SVPs?

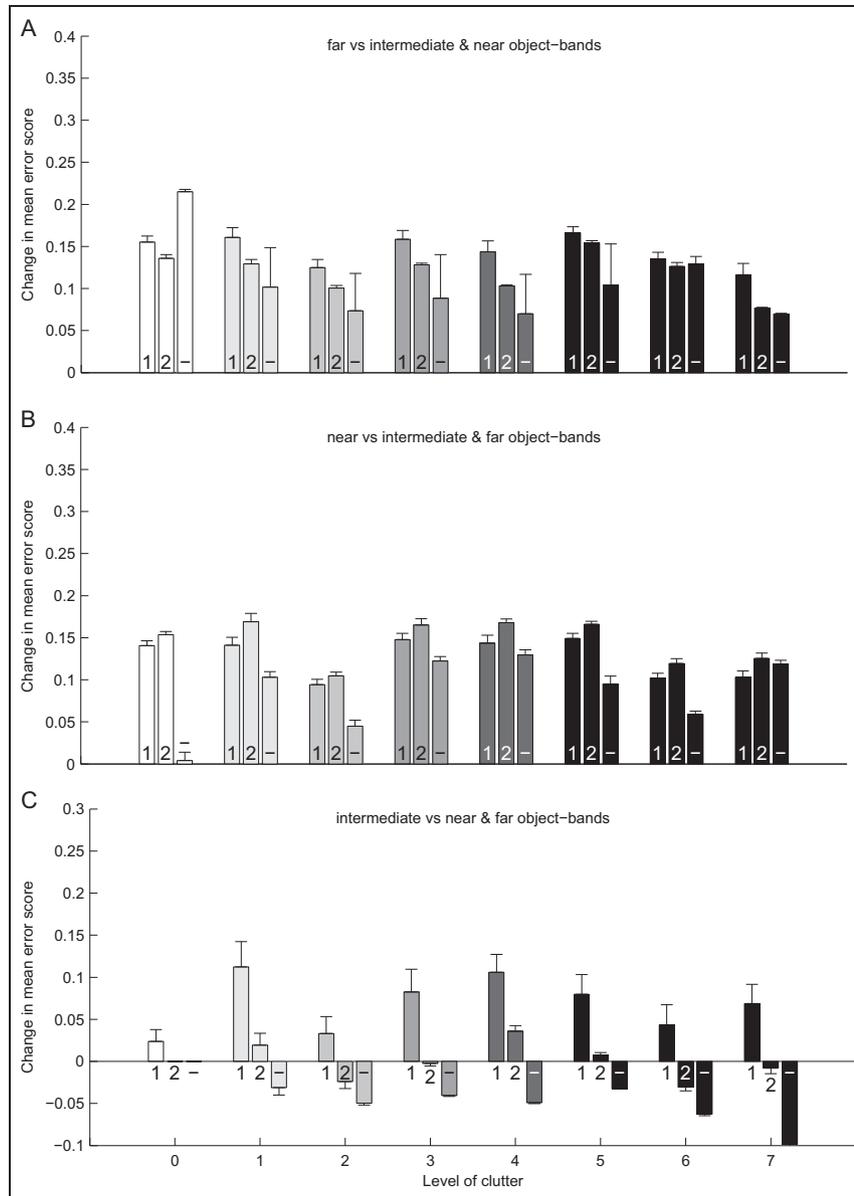
**3.4.1 Comparing SVPs between worlds.** We have shown above that how an SVP performs is partly determined by its form, independently of the world in which it finds itself, although the structure of the world is also important. We next examine which parts of visual worlds affect the performance of SVPs. To do this, we defined three different ‘object-bands’, which together describe the worlds used in this set of experiments: the near object-band, comprised of a single, large tussock, which from the nest entrance always appeared higher than the horizon; the intermediate object-band, containing small tussocks; and the far object-band, which contained 2d bushes and trees (Section 2.4.1). Another manipulation involved varying the number of tussocks in the intermediate object-band (in eight even levels,  $0 \leq N \leq 159$ ); we call this ‘level of clutter’. Each object-band could be rotated independently of the others or removed entirely. We generated worlds for all combinations of clutter level and object-band rotations. Comparisons between relevant worlds were then used to estimate the relative contribution of each object-band to determining which SVPs perform well in that world.

As an aside, we should explain why, with an agent that does not segregate its visual environment, we feel it necessary to run trials in worlds which we have ourselves divided into ‘object-bands’. This division is simply so that we can ask how the homing mediated by particular SVPs is impacted by certain large-scale alterations to the environment. Of course, from the perspective of the algorithm, there is *no* division between object-bands and none contributes to homing performance independently of the others; for us, views are an undifferentiated mass of pixels, not a set of extracted and labelled visual features.

The aim of these experiments was to see what effect changes to the visual environment have on which SVPs perform well. To do this we took the 100 best-performing SVPs for a ‘reference world’ and assessed how their performance changed with the rotation of one or two object-bands. There were four trials for each condition (a rotation of 0°, 90°, 180° or 270°) with the 100 SVPs changing depending on which performed best in the relevant reference world. Small changes in error indicate that the SVPs are not markedly better suited to the reference world or test world and so suggest that similar learning walks would be effective in both and that the visual environment is not so important in determining which SVPs perform well. When the score changes, however, we can compare the changes in mean error between test worlds where one (‘Rotate 1’) or two (‘Rotate 2’) object-bands are rotated to see which properties of the environment ‘resonate’ with the well-performing SVPs. Note that we are looking at view positions (SVPs) not view images (to which we turn in the next section).

We first compared the effect of rotating the far object-band with that of rotating the near and intermediate object-bands together (Figure 5A). Which of the error changes is more, when rotating the far (‘Rotate 1’) or intermediate and near object-bands (‘Rotate 2’), indicates which of these rotations has the greatest effect on homing performance for the best-performing SVPs. An error change of zero would indicate that the SVPs perform equally well in reference and test worlds, or, in other words, that as ‘strategies’ they are as viable in either world and thus that there is no value in tailoring a learning walk to the object-band(s) under examination. As shown in Figure 5, the error change was significantly less when rotating the near and intermediate object-bands than the far object-band, which shows that the far object-band is more important in determining the SVPs that were good for the reference world. This is presumably because the far object-band is the furthest away and therefore less subject to visual translation, making it a reliable cue over space.

Next we compared the near with the intermediate and far object-bands together (Figure 5B). If rotation of the near object-band had a greater effect than



**Figure 5.** How is SVP performance influenced by changes to the visual environment? The 100 best SVPs for a ‘reference world’ were implemented in new versions of the world with one band of visual objects rotated or removed; we analysed SVP performance when aligned with the focal object-band or aligned with the other two object-bands. Bar charts show the change in mean error score (EHC), where a positive value represents a decrease in performance. The analysis was performed for seven levels of clutter in the middle band of landmarks, shown from white (low clutter) to black (high clutter). A–C: the focal visual object-bands were far (A), near (B) and intermediate (C). For a given level of clutter, changes in mean error are reported for three conditions: Rotate 1 (1); Rotate 2 (2) or Absent (–), where Rotate 1 means that SVP stay aligned with the focal object-band, Rotate 2 means SVP are aligned with the other two object-bands, and Absent means that the focal object-band is removed. A: Rotate 1 has higher error than Rotate 2 (sign test; 8/8;  $p < 0.005$ ); Rotate 2 has a higher error than Absent (sign test; 7/8;  $p < 0.01$ ). B: Rotate 2 has higher error than Rotate 1 (sign test; 8/8;  $p < 0.005$ ); Rotate 1 has a higher error than Absent (sign test; 7/8;  $p < 0.01$ ). C: Rotate 1 has higher error than Rotate 2 (sign test; 8/8;  $p < 0.005$ ); Rotate 2 has a higher error than Absent (sign test; 7/8;  $p < 0.01$ ).

rotation of the intermediate and far object-bands on the performance of the best-performing SVPs, we would expect a greater error change for the ‘Rotate 1’ than the ‘Rotate 2’ condition. However, the opposite was observed, indicating that the intermediate and far object-bands together are a significantly bigger determinant of which SVPs perform best.

We then compared the intermediate with the near and far object-bands together (Figure 5C). We would expect a greater error change when rotating the intermediate object-band (‘Rotate 1’) than the near and far object-bands (‘Rotate 2’) if the intermediate is more important. This is indeed what was observed, indicating the greater importance of the intermediate

object-band in determining optimal SVPs. This is perhaps surprising given that the intermediate object-band is comprised solely of small tussocks which might be thought to act as ‘clutter’, acting only to obscure information in the far object-band; however, this result shows that sufficient information is discernible from these small tussocks for reliable homing.

To examine this further, we looked at the effect of removing object-bands completely (‘Absent’) on the best-performing SVPs. This is because the previous manipulations (the rotations) in effect change the world in two ways: the object-band being rotated has been removed from its original location *and* it has been placed in a new one. Hence, it is not surprising that in all three sets of experiments the ‘Absent’ condition had a significantly lower error change than both the ‘Rotate 1’ and the ‘Rotate 2’ conditions, with one exception: the ‘near vs ...’ trials with no clutter (clutter level 0). For this latter case, the error change was greater instead, presumably because, without the far object-band or the intermediate object-band, the homing is driven solely by the near object-band, which covers only a small portion of the visual field, so giving a high level of visual mismatch. The influence of removing an object-band is particularly pronounced for the ‘intermediate vs ...’ trials, where removal of the intermediate band led to a *decrease* in error, reflecting better homing without than with the intermediate object-band. Yet we also know that the intermediate object-band is important, as shown by the rotation trials. This indicates that the far object-band is the most reliable cue, but if it is obscured by the intermediate object-band this in turn becomes the most important.

Lastly, we varied the ‘level of clutter’ in order to see if the increase in number of objects, with a concomitant decrease in the visibility of the distant panorama, affects the direction of these trends or homing performance generally. It appears not to have a big impact on the direction of trends, although the mean error for the reference world did increase with increasing levels of clutter (data not shown), intuitively suggesting that high-clutter environments are more difficult to navigate within.

Thus, the far and intermediate object-bands have a substantial effect, whereas, perhaps counterintuitively, a prominent nearby object does not. Note that this does not imply that the near object-band necessarily contains no useful information; simply that SVPs should be tailored to the intermediate and far object-bands, but not the nearby object. We next turn to the question of what happens to homing performance when object-bands are rotated *after* views have been learned from the unrotated world.

**3.4.2 Comparing sets of view images between worlds.** In the previous section we investigated how SVP performance

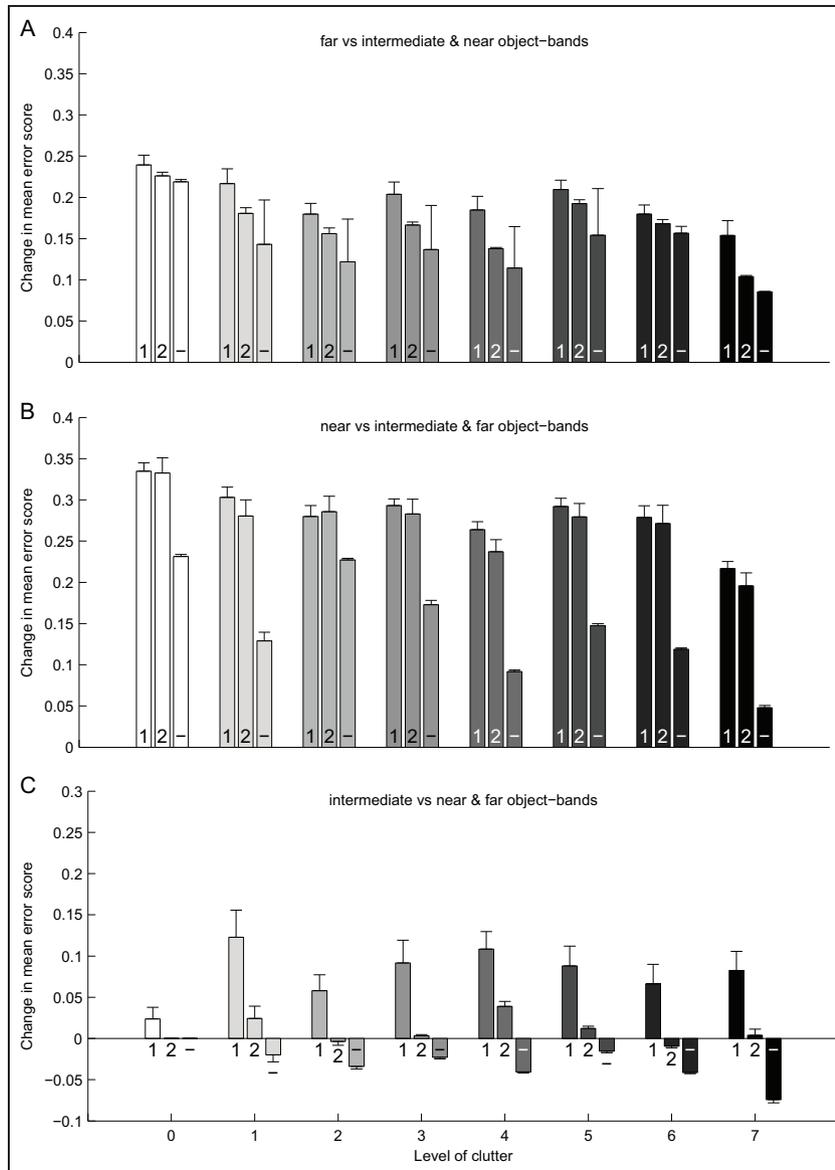
varies between worlds differing along certain specific parameters. Note that it was the set of positions and *not* the image content of views being investigated; the same SVP contains different images in different worlds (as the worlds look different). We next looked at what happens if the image content of views is conserved between reference and test worlds; in other words, we looked at sets of view images (SVIs) rather than SVPs. The ‘real-world’ analogue of these experiments would be observing the effect on an experienced forager’s homing after the environment is suddenly altered.

These experiments were mostly of the same form as those in the previous section, but with one crucial difference: whereas previously the same SVPs were used across worlds, but with the SVIs varying (i.e. because each world looks different), here, the same SVIs, drawn from the reference world, were used throughout. This enables us to investigate whether specific view sets are robust to large changes in the visual world. In this instance, the 100 best-performing SVPs from the reference world were used as the baseline.

For this analysis the pattern of results was similar to the SVP experiments (Figure 5) for the ‘far vs ...’ and ‘intermediate vs ...’ conditions (Figure 6A,C), but was in the opposite direction for the ‘near vs ...’ conditions (Figure 6B): in other words, there was greater increase in error for ‘Rotate 1’ than ‘Rotate 2’. Though the intermediate object-band ‘wins out’ from among these pairwise comparisons, the far object-band is still important, as indicated by the relatively greater error changes for ‘Rotate 1’ in the ‘far vs ...’ compared with ‘Rotate 1’ in the ‘intermediate vs ...’ conditions. Hence, for this set of experiments, as before, it is the movement of objects at intermediate and far distances that have the greatest effect on homing performance. However, there is also a (slight) effect for the near object-band, indicating that moving prominent nearby objects will substantially impact upon homing performance. This suggests that, though the learning walks of ants near to a prominent object may be similar regardless of where the object is, nonetheless, moving it may trigger a fresh bout of learning walks (Müller & Wehner, 2010).

### 3.5 The relation of image difference to SVP and SVI performance

We also examined whether the pattern of error changes, for all of the conditions mentioned above, is predicted by the r.m.s. difference between images from the ‘reference’ and ‘test’ worlds (see Figures 7 and 8). There could conceivably be a relation between image difference and homing performance, as we know the former indicates how different two scenes are and we would presume that on average a similar ‘test’ world should have smaller scene-differences and a small increase in error change compared to the reference world. Yet the



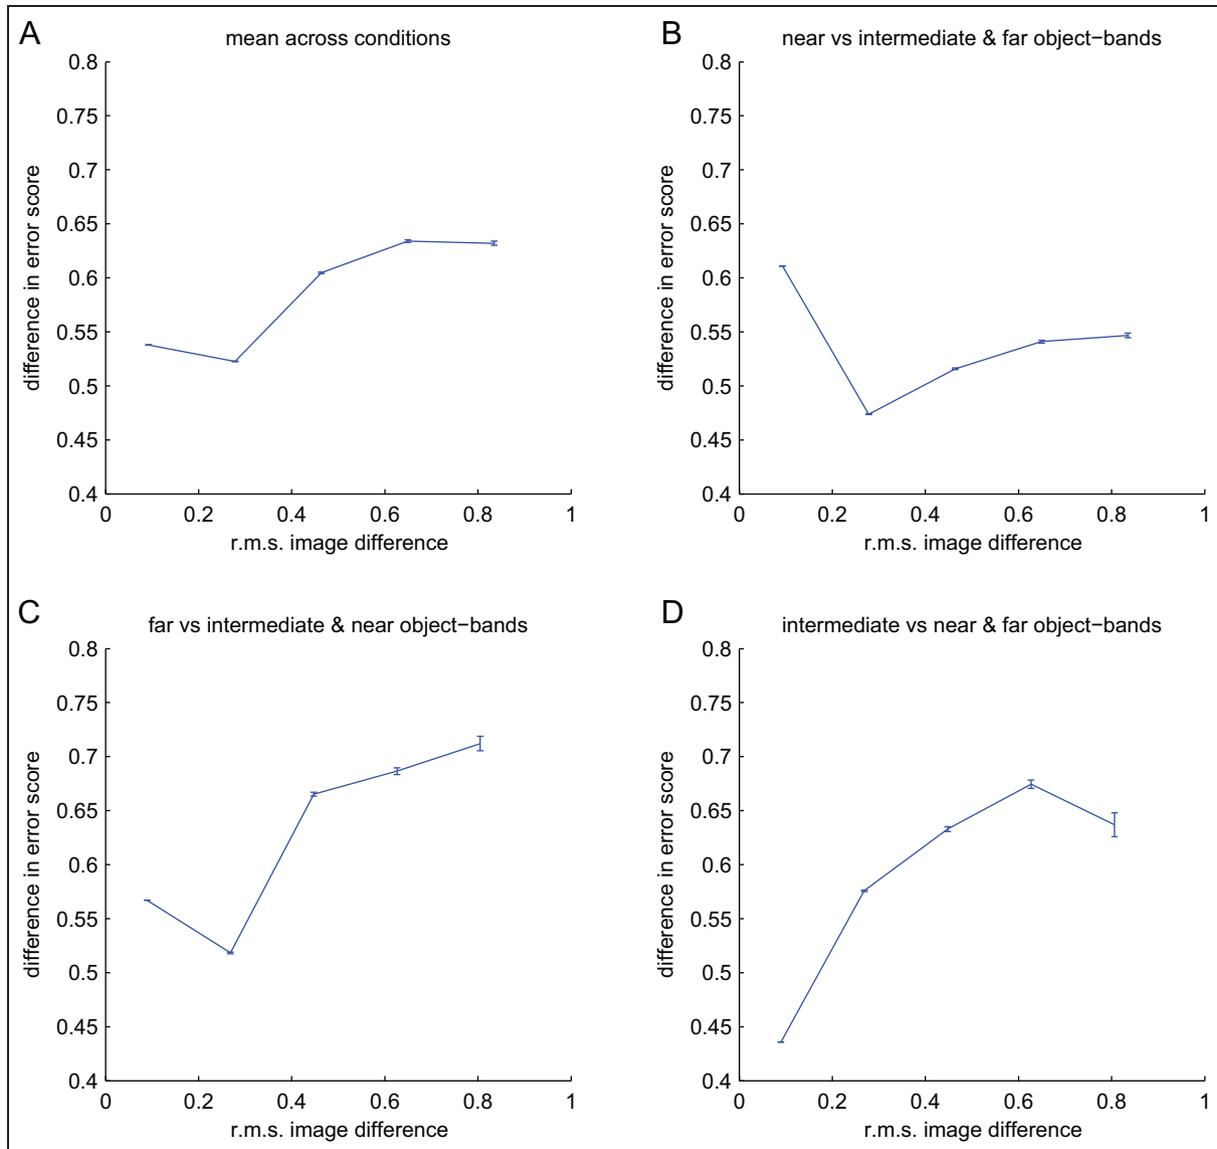
**Figure 6.** How is SVI performance influenced by changes to the visual environment? Conventions as in Figure 5, though data refer to SVI rather than SVP performance across conditions. Hence, views drawn from the reference world were used as the basis for homing in the test worlds. A–C: the visual object-bands were far (A), near (B) and intermediate (C). The three conditions were as before (see Figure 5): Rotate 1 (1), Rotate 2 (2) and Absent (–). A: Rotate 1 has higher error than Rotate 2 (sign test; 8/8;  $p < 0.005$ ); Rotate 2 has a higher error than Absent (sign test; 8/8;  $p < 0.005$ ). B: Rotate 1 has higher error than Rotate 2 (sign test; 7/8;  $p < 0.01$ ); Rotate 2 has a higher error than Absent (sign test; 8/8;  $p < 0.005$ ). C: Rotate 1 has higher error than Rotate 2 (sign test; 8/8;  $p < 0.005$ ); Rotate 2 has a higher error than Absent (sign test 8/8;  $p < 0.005$ ).

trend was weak, with a large variance. For view-based homing there must be a function relating image difference and error. These results show how the particulars of homing success as a function of image differences are dependent, in non-trivial ways, on the distribution of objects in the world.

#### 4 Discussion

The purpose of this study was to examine how homing performance is affected by learning-walk form, the

visual environment and the relation between the two. To do this we used a simple view-based homing algorithm in VR worlds, which were composed of 3d tussocks and 2d bushes and trees. Stored views were drawn from a number of positions, fixed relative to the nest entrance, and subsets of these views (SVIs) were used for homing. We investigated how different combinations of views perform relative to one another and how this differs between worlds, in order to see which properties of SVPs and worlds are important, and how they depend upon one another.

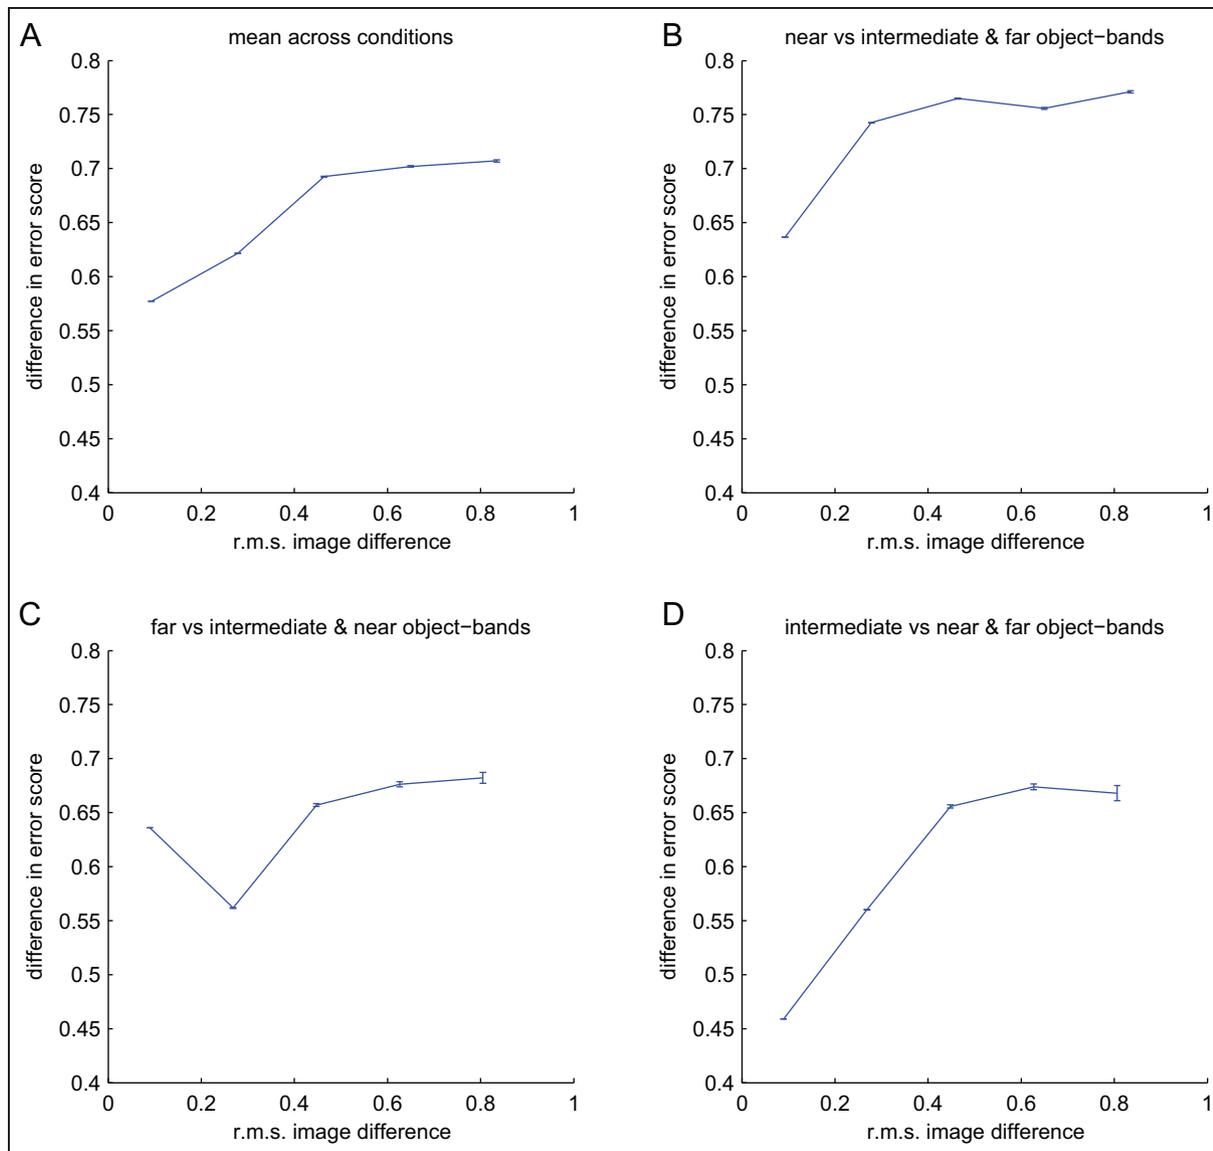


**Figure 7.** Relationship between changes in error score and image difference for the 100 best-performing SVPs. For all the data in Figure 5 we look at the relationship between change in error score and the r.m.s. image difference. Image differences are calculated for all 9845 locations in the worlds before and after the rotation of one or two object-bands, or the removal of one, with data divided between five bins of equal width on the basis of image difference. Error bars give standard error. Note that there is not a clear positive trend for any of the conditions, indicating that image difference alone is not predictive of relative homing success.

Our major findings were that certain intuitive properties of SVPs, such as ‘spread of views’ and ‘view dissimilarity’, are partially predictive of homing performance irrespective of the visual environment. However, no single SVP performed well in all worlds, indicating that it is also advantageous to tailor learning walks to the surroundings. Further investigation found that a large, salient object adjacent to the nest is likely *not* to be the most important aspect of a visually rich environment. The distant panorama (our ‘far object-band’) will generally be important, as will objects at an intermediate distance from the nest when there is no clear line of sight to the horizon. This seemingly counterintuitive finding is in line with results that show ants

failing to extract and utilize information from a salient beacon-like object (Wystrach, Beugnon, & Cheng, 2011; Wystrach & Graham, 2012). A second intuition was that our results might be predicted by simple image differences. For instance, differences in error scores following experimental manipulation may correlate with r.m.s. differences between images in the original and transformed worlds, but the relationship was weak (Figures 7 and 8).

Of course, visual homing is bound to be more complex in real ants, as, among other things, they do not have a completely panoramic field of view (though it is broad: e.g. for *M. bagoti*,  $\sim 150^\circ$  per eye: Schwarz et al., 2011). Additionally, *M. bagoti* ants have larger anterior



**Figure 8.** Relationship between changes in error score and image difference for the 100 best-performing SVIs (i.e. with view images drawn from the reference world; see Figure 7). For all the data in Figure 6 we look at the relationship between change in error score and the r.m.s. image difference. Image differences are calculated by comparing all 9845 locations before and after the rotation of one or two object-bands (or removal of one). Error bars give standard error. Again, there is not a strong positive relationship between image difference and relative homing success.

than posterior facet diameters (cf. *Cataglyphis bicolor*: Zollikofer et al., 1995), which could equate to a 50% difference in resolution (Schwarz et al., 2011). We also know that ants must employ early visual filtering to distinguish sky from skyline, among other things, using a green–UV colour-opponent system (Möller, 2002). These factors might change the fine detail of the results, though the broad pattern would remain the same.

#### 4.1 Modelling navigation with ant-like constraints

The value of simulation work is that one can close the loop between behaviour, environment and computational algorithm, because one has control over all three.

As we are interested in biological questions, we have tried to constrain these three things so that they are ant-like. Although the relationship between simulation studies and biology is strong (Webb, 2000), especially with respect to insect navigation, there are few modelling studies where the constraints on the system are ‘ant-like’. Since the original homing models (see Cartwright & Collett, 1983) describe agents at a fixed orientation in space, many follow-up works have maintained this constraint and thus there are few ant-specific models of visual homing. Möller (2012) presents a neural network-based model of ant navigation that relies on predicting changes to the environment from visual translation. It makes use of ‘scanning’ in order to find a heading, as

does the model of Baddeley et al. (2012). Basten and Mallot (2010) investigated homing performance in an ant-like world using either pixel intensity or skyline height for the stored views and either average landmark vector (ALV) (Lambrinos, Möller, Labhart, Pfeifer, & Wehner, 2000) or gradient ascent algorithms to inform the agent's movement, but relied on only one stored view at a time, and images were aligned as with 'snapshot' (sensu Cartwright & Collett, 1983) models.

In order to maintain ant-like constraints, we employed an RIDF-based homing algorithm (Graham et al., 2010; Narendra, Gourmaud, & Zeil, 2013; Philippides et al., 2011) across a range of virtual environments (shown to be realistic: see Figure 2), primed with views drawn from different possible SVPs. The algorithm performed well, with reliable homing achieved across a range of worlds and of possible SVPs. This work thus represents the first attempt to investigate the active learning of views for homing with biologically realistic, ant-like constraints on visual environment, motor system and behaviour.

#### 4.2 Relation to biological data

Our simulation studies make a number of predictions about real-life learning walks. First, we predict that a single large landmark near the nest entrance will alter learning walks substantially *only* in a relatively featureless environment. Second, we expect the distant panorama to act as a reliable cue, as it is relatively invariant with motion of the forager. Yet although there is potentially a high level of information on the horizon, it will often be obscured by visual 'clutter' (such as tussocks) at intermediate distances from the nest. In this case, we would expect the clutter *itself* to be a useful cue; there will be a small effect of visual translation when compared to very proximal objects and if there is a clear line of sight it can be reliable over space.

Unfortunately, there is very little in the way of detailed learning-walk data in the literature with which to compare our results. One exception is found in Müller and Wehner (2010), who present several *Oecomyrmex robustior* learning walks, which the ants executed following placement of a large black cylinder near the nest entrance. The authors show two learning walks, roughly spiral-shaped, but asymmetric and substantially different from each other. All fixations were directed at the invisible nest entrance; the cylinder, despite being highly salient in an otherwise featureless environment, attracted no fixations, even though its placement triggered the learning walks in the first place. This fits with the idea that important visual cues can be useful to homing without being fixated. However, the ants' prior experience (number and direction of previous foraging bouts) is not specified, so its potential influence on these learning walks is unknown. The other important work is by Nicholson et al. (1999),

who trained ants to a feeder next to which were placed one, two or no black cylinders. In this case, however, the ants fixated the cylinders rather than the goal. The discrepancy in results could be due to the difference in context (nest entrance vs feeder) or in species (*O. robustior* vs *F. rufa*; for species differences in homing behaviours, see Schwarz & Cheng, 2010).

As the ultimate aim of this project is to understand how learning walks function in ants, further behavioural work will be extremely informative. Of course, detailed recordings, across species and ecologies, will prove useful to many interested in the natural history of learning walks. However, are for an in-depth, mechanistic understanding of the phenomenon systematic manipulations to environments are needed, with observations of how this changes learning walks. We hope that the results here will provide a 'starting-off point' for such investigations and will stimulate discussion about the simple computational mechanisms that could underlie this critically important behaviour.

#### Funding

AD and PG are both funded by BBSRC (grant codes: BB/F015925/1 and BB/H013644).

AP is funded by EPSRC (code: EP/I031758/1) and grant FP7 ICT 308943.

#### References

- Baddeley, B., Graham, P., Husbands, P., & Philippides, A. (2012). A model of ant route navigation driven by scene familiarity. *PLoS Computational Biology*, 8(1), 1–16.
- Basten, K., & Mallot, H. A. (2010). Simulated visual homing in desert ant natural environments: Efficiency of skyline cues. *Biological Cybernetics*, 102(5), 413–425.
- Batschelet, E. (1981). *Circular statistics in biology*. London: Academic Press.
- Berens, P. (2009). CircStat: A MATLAB toolbox for circular statistics. *Journal of Statistical Software*, 31(10), 1–21.
- Brooks, R. A. (1999). *Cambrian intelligence*. Cambridge, MA: MIT Press.
- Buehlmann, C., Hansson, B. S., & Knaden, M. (2012). Path integration controls nest-plume following in desert ants. *Current Biology*, 22(7), 645–649.
- Cartwright, B. A., & Collett, T. S. (1983). Landmark learning in bees: Experiments and models. *Journal of Comparative Physiology*, 151, 521–543.
- Cheng, K. (2010). Common principles shared by spatial and other kinds of cognition. In F. L. Dolins, & R. W. Mitchell (Eds.), *Spatial cognition, spatial perception: Mapping the self and space*. Cambridge: Cambridge University Press.
- Cheng, K., Narendra, A., Sommer, S., & Wehner, R. (2009). Traveling in clutter: Navigation in the Central Australian desert ant *Melophorus bagoti*. *Behavioural Processes*, 80, 261–268.
- Collett, T. S. (1995). Making learning easy: The acquisition of visual information during the orientation flights of social wasps. *Journal of Comparative Physiology A*, 177, 737–747.
- Collett, T. S., Graham, P., Harris, R. A., & Hempel de Ibarra, N. (2006). Navigational memories in ants and bees:

- Memory retrieval when selecting and following routes. *Advances in the Study of Behaviour*, 36, 123–172.
- Collett, T. S., & Lehrer, M. (1993). Looking and learning: A spatial pattern in the orientation flight of the wasp *Vespa vulgaris*. *Proceedings: Biological Sciences*, 252(1334), 129–134.
- Graham, P., & Collett, T. S. (2006). Bi-directional route learning in wood ants. *The Journal of Experimental Biology*, 209, 3677–3684.
- Graham, P., Philippides, A., & Baddeley, B. (2010). Animal cognition: Multi-modal interactions in ant learning. *Current Biology*, 20(15), 639–640.
- Jander, R. (1997). Comparative psychology of invertebrates: The field and laboratory study of insect behaviour. In G. Greenberg, & E. Tobach (Eds.), *Research in developmental and comparative psychology* (pp. 79–99). New York, NY: Garland Publishing, Inc.
- Kim, P., Szenher, M. D., & Webb, B. (2009). Entropy-based visual homing. In *International Conference on Mechatronics and Automation* (pp. 3601–3606).
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., & Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, 30, 39–64.
- Lent, D. D., Graham, P., & Collett, T. S. (In press). Phase-dependent visual control of the zigzag paths of navigating wood ants. *Current Biology*.
- Möller, R. (2002). Insects could exploit UV-green contrast for landmark navigation. *Journal of Theoretical Biology*, 214(4), 619–631.
- Möller, R. (2012). A model of ant navigation based on visual prediction. *Journal of Theoretical Biology*, 305, 118–130.
- Müller, M., & Wehner, R. (2010). Path integration provides a scaffold for landmark learning in desert ants. *Current Biology*, 20(15), 1368–1371.
- Narendra, A., Gourmaud, S., & Zeil, J. (2013). Mapping the navigational knowledge of individually foraging ants, *Myrmecia croslandi*. *Proceedings of the Royal Society B*, 280(1765), 1–9.
- Nicholson, D. J., Judd, S. P. D., Cartwright, B. A., & Collett, T. S. (1999). Learning walks and landmark guidance in wood ants (*Formica rufa*). *The Journal of Experimental Biology*, 202(13), 1831–1838.
- Pfeifer, R., & Scheir, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
- Philippides, A., Baddeley, B., Cheng, K., & Graham, P. (2011). How might ants use panoramic views for route navigation? *The Journal of Experimental Biology*, 214, 445–451.
- Schwarz, S., & Cheng, K. (2010). Visual associative learning in two desert ant species. *Behavioral Ecology and Sociobiology*, 64(12), 2033–2041.
- Schwarz, S., Narendra, A., & Zeil, J. (2011). The properties of the visual system in the Australian desert ant *Melophorus bagoti*. *Arthropod Structure and Development*, 40, 128–134.
- Wang, R., & Spelke, E. S. (2002). Human spatial representation: Insights from animals. *Trends in Cognitive Sciences*, 6(9), 376.
- Webb, B. (2000). What does robotics offer animal behaviour? *Animal Behaviour*, 60(5), 545–558.
- Wehner, R. (2008). The architecture of the desert ant's navigational toolkit (Hymenoptera: Formicidae). *Myrmecological News*, 12, 85–96.
- Wehner, R., Meier, C., & Zollikofer, C. P. E. (2004). The ontogeny of foraging behaviour in desert ants, *Cataglyphis bicolor*. *Ecological Entomology*, 29, 240–250.
- Wehner, R., & Srinivasan, M. V. (1981). Searching behaviour of desert ants, genus *Cataglyphis* (Formicidae, Hymenoptera). *Journal of Comparative Physiology*, 142(3), 315–338.
- Wei, C. A., & Dyer, F. C. (2009). Investing in learning: Why do honeybees, *Apis mellifera*, vary the durations of learning flights? *Animal Behaviour*, 77(5), 1165–1177.
- Wei, C. A., Rafalko, S. L., & Dyer, F. C. (2002). Deciding to learn: Modulation of learning flights in honeybees, *Apis mellifera*. *Journal of Comparative Physiology A*, 188(9), 725–737.
- Wiener, J., Shettleworth, S., Bingman, V. P., Cheng, K., Healy, S., Jacobs, L. F., ... Newcombe, N. S. (2011). Animal navigation: A synthesis. In R. Menzel, & J. Fischer (Eds.), *Animal thinking: Contemporary issues in comparative cognition* (vol. 8, pp. 51–78). Cambridge, MA: MIT Press.
- Wystrach, A., Beugnon, G., & Cheng, K. (2011). Landmarks or panoramas: What do navigating ants attend to for guidance? *Frontiers in Zoology*, 8(1), 21.
- Wystrach, A., & Graham, P. (2012). What can we learn from studies of insect navigation? *Animal Behaviour*, 84(1), 13–20.
- Wystrach, A., Mangan, M., Philippides, A., & Graham, P. (2013). Snapshots in ants? New interpretations of paradigmatic experiments. *The Journal of Experimental Biology*, 216, 1766–1770.
- Zeil, J. (1993a). Orientation flights of solitary wasps (*Cerceris*; Sphecidae; Hymenoptera). I. Description of flight. *Journal of Comparative Physiology A*, 172(2), 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. *Journal of Comparative Physiology A*, 172, 207–222.
- Zeil, J. (2011). Visual homing: An insect perspective. *Current Opinion in Neurobiology*, 22, 1–9.
- Zeil, J., Hofmann, M. I., & Chahl, J. S. (2003). Catchment areas of panoramic snapshots in outdoor scenes. *Journal of the Optical Society of America A*, 20(3), 450–469.
- Zollikofer, C. P. E., Wehner, R., & Fukushi, T. (1995). Optical scaling in conspecific *Cataglyphis* ants. *The Journal of Experimental Biology*, 198, 1637–1646.

## About the Authors



**Alex Dewar** is a PhD candidate, working in the centre for computational neuroscience and robotics (CCNR) at the University of Sussex. He is interested in how computationally simple heuristics can in combination lead to complex behaviours, with an especial emphasis on strategies for visually guided navigation.



**Andrew Philippides** is a senior lecturer in informatics. His main research focus is best described as computational neuroethology: to combine traditional biology experiments with modelling and robotics. Current research topics include visual navigation in insects and robots, neuromodulation in (real and artificial) neural networks and agent-based modelling.



**Paul Graham** is a reader in biology and member of the CCNR. His research involves behavioural studies of insect navigation and he is interested in how modelling and robotics can be used as tools for understanding biological systems.