Proceedings of the 2015
IEEE Conference on Robotics and Biomimetics
Zhuhai, China, December 6-9, 2015

Bio-inspired Visual Guidance: from Insect Homing to UAS Navigation

Aymeric Denuelle[1,2] and Mandyam V. Srinivasan[1,3]

Abstract— **In the context of autonomous navigation, vision has demonstrated its potential in replacing bulkier sensors and achieving multi-tasking, operating as the primary sensor for guidance as well as simultaneously fulfilling other dedicated missions. Vision sensors are of particular interest in aerial navigation where their versatility and light weight simplify the avionics and increase the available payload of unmanned aerial systems (UAS). Biological studies provide us with insights into the visual cues by which insects are able to navigate effectively despite their limited brain resources and low resolution eyes, and enable the development of simple, yet efficient navigation techniques for UAS guidance. This paper reviews biological models of insect homing, their subsequent applications into the robotics field, and the extent to which these models have been applied to the guidance of UAS. Suitable methods for autonomous aerial navigation in natural environments are described and the remaining challenges and opportunities for bio-inspired techniques of visual guidance are discussed.**

I. Introduction

Unmanned aerial systems (UAS) are gaining popularity in many civilian applications, being deployed in search and rescue situations, structure inspection, environmental surveying, precision agriculture, etc. The need for advanced autonomous capabilities, along with aiming at downsizing the onboard sensor suite to increase UAS' autonomy, have led researchers to look into alternative solutions to current autopilots which rely primarily on global satellite navigation systems such as GPS. Low-cost, lightweight, multi-purpose imaging sensors have promoted the use of vision as the primary sensor for navigation. Although presenting many advantages, vision algorithms remain computationally intensive. Insects, such as bees, wasps and ants, use visual cues and display remarkable navigational skills despite their limited computational capabilities. Based on biological observations, in particular by looking at homing strategies in insects, more parsimonious, bio-inspired, methods have been developed for vision-based local guidance and long-range visual homing.

Although flying insects have been shown to rely on optic flow (OF) for attitude control and navigation [1], we do not discuss in this survey the standard OF techniques applied to UAS navigation (for a review, see [2]), which often require advanced processing (calibration, optimisations). Additionally, due their high computational burden and as they do not appear to be biologically plausible, visual self-localisation methods based on techniques such as SLAM (simultaneous localisation and mapping) [3] are outside the scope of this review. Indeed, insects do not seem to use maps for navigation; they show evidence of storing multiple routes but not their spatial relationships [4].

The remainder of this paper presents bio-inspired models for visual local homing and their existing applications in Section II. Section III reviews long-range homing schemes and their application to robot navigation. Special attention is given to UAS applications in Section IV; Section V presents the remaining challenges and Section VI concludes the review.

Authors are with the Queensland Brain Institute[1] and School of Information Technology and Electrical Engineering[3], The University of Queensland, St Lucia, Australia, and the Autonomous Systems Program[2], CSIRO, Pullenvale, Australia. Corresponding author: `a.denuelle@uq.edu.au`

The research described here was supported partly by the ARC Centre of Excellence in Vision Science (CE0561903), by Boeing Defence Australia Grant SMP-BRT-11-044, ARC Linkage Grant LP130100483, ARC Discovery Grant DP140100896, a Queensland Premier's Fellowship, and an ARC Distinguished Outstanding Researcher Award (DP140100914).

II. Bio-inspired Local Visual Homing Models

In this section, we review biologically plausible techniques that have been applied in the field of robot navigation in order to perform visual local homing (within one's visual field), inspired by different strategies that insects use to return to a known location (usually a food source or nest). The homing underlying mechanism is the capacity for relating currently experienced visual information with that stored in memory. Techniques developed for visual homing can be broadly separated into landmark-based and view-based methods. The first category relies on the extraction of particular features or the identification of objects (landmarks) to compute the homing vector while the second one uses the image as a whole without extracting any features.

A. Landmark-based Models

Landmark-based methods are subdivided into two categories: template-based and parameter-based methods. The latter simply stores a few parameters to describe a location and its landmarks, while template-based methods make use of image sub-regions to characterize landmarks. Although the template concept was predominant in the last decades to explain insect navigation abilities, it is still not clear whether homing insects use templates, parameters or both [5].

1) Parameter-based models: [6] developed an image-based local homing system to allow a robot to navigate using panoramic images. Their scheme makes use of one-dimensional image strips sampled along the horizon, used to extract characteristic points (landmarks). Biological observations indeed suggest that insects, such as ants, use skyline encoding to parameterise natural scenes [7]. The differences in angular positions of these landmarks between the current and the reference images are used to determine the home vector. However, as only vertical edges show up prominently as characteristic points, this method is best suited to indoor

978-1-4673-9675-2/15/$31.00 © 2015 IEEE

environments. This approach implicitly employs the isotropic distance distribution (IDD) assumption, also used in many template-based models [8], [9], that landmark distances can vary, but the distribution of distances is independent of the viewing direction.

A simpler version of this algorithm is presented in [10], where it is assumed that all landmarks are at the same distance from the snapshot location. The authors sampled a ring at the horizon to obtain a 1D image, and built a matched filter which predicts the displacement field of the landmarks, based on the equal distance (ED) assumption. The current view is then warped according to this displacement field for various (unit) directions of displacement of the robotic platform from the home location, and compared with the reference snapshot by using the dot product to measure the degree of match. The best-fitting direction gives an approximation of the home vector, which is then used to guide the platform in a test arena. Although this is an exhaustive search procedure (and hence could become computationally expensive when using images of greater dimension), it provides the robot with an estimate of its orientation, the home direction as well as the distance from home (correlated to the image distance). Even if the ED assumption for the environment does not seem very realistic, the resulting homing performance is not severely affected and, close to the goal location, this scheme outperforms algorithms that assume an isotropic landmark distribution. Indeed, the error in the computed home direction using the ED assumption decreases when approaching the goal location (unlike in the case of the IDD assumption). Additionally, this scheme does not require an external compass, because the orientation is included in the search space of the warping method.

The homing algorithm proposed by [11] can be regarded as an extension of the warping algorithm [10] to the frequency domain. This scheme uses 1D images obtained by sampling panoramic images at the horizon. By approximating these 1D images in the Fourier domain (using sine-cosine and amplitude-phase representations), the authors derive an efficient homing scheme, using a mobile robot within an office environment, where the calculation of one home vector is approximately 100 times faster when compared with the warping method. Indeed, their algorithm only requires a small number of Fourier coefficients to achieve a good homing performance. In order to increase the homing accuracy and working distance, the authors proposed a coarse-to-fine homing strategy: start with a small number of Fourier coefficients (low frequencies) and then increase this number as the distance to the goal location decreases.

[12] developed the average landmark vector (ALV) model, in which a unit landmark vector, pointing from the current location towards the landmark, is assigned to each visual landmark feature. The landmark vectors are averaged to get the ALV of the goal location. By simply subtracting the stored ALV (computed at the goal location) from this ALV, an approximation of the home vector is obtained. Consequently, this model is really efficient in terms of computational resources, as only the goal ALV needs to be stored and the home vector is easily obtained by subtraction of two AL vectors instead of requiring an image matching process. [12] and [13] demonstrated successful implementations of this parsimonious ALV model for guiding a mobile robot to its goal location.

2) Template-based models: [8] was the first to propose a novel view-based homing scheme, termed the 'snapshot' model, to explain honeybees' behaviour when searching for a goal location (hive). In this approach, landmarks in the current panoramic image are segmented, and each landmark is tagged with an elementary movement vector (EMV) that is directed toward the landmark if it appears smaller than in the reference image, or away from it if it appears larger. The home vector is obtained as a weighted sum of these EMVs, with the weights proportional to the degree of mismatch.

[14] later introduced the term 'catchment area' to describe the region within which a bee can successfully return to its goal location. The size of this catchment area depends on the density and size of landmarks present in the environment. Their model proposes that bees actually use two snapshots to reach a location, one which excludes landmarks in the vicinity of the goal, and another which includes them. The discrepancy between the filtered snapshot (without nearby landmarks) and the current retinal image (also filtered in the same way) gradually increases as the bee moves away from the reference location. Using such filtered snapshots, a model bee can find its way back from a relatively long distance and then switch to the unfiltered snapshot to pinpoint the goal location by maximizing the fit between the unfiltered retinal image and the memorized unfiltered snapshot. Realistically, the use of filtered snapshots requires the ability to measure the distances to objects in order to filter the snapshot according to the distance. Bees can obtain this information by means of motion parallax.

In the snapshot model, to achieve the matching process, both retinal image and memorized snapshot are assumed to be oriented in the same absolute direction – an assumption which has been verified in the homing behaviour of insects. Indeed, [15] show that bees face the same magnetic compass direction when capturing a snapshot or looking for the goal location. Ants seem to navigate in a similar way by keeping their body axis aligned so that the landmark is always in front of the eye [16].

Most variants of the snapshot model implemented in simulation and in physical robots use 1D panoramic images of landmarks generated from raw 2D images [9]. [17] presented a novel approach operating on 2D images, and using image corners as features. They compared their scheme against an implementation of the ALV model on real-world images and demonstrated superior performance in terms of return ratio (number of successful homing runs to total number of homing trials) and size of the catchment area, which was more than doubled because it tracked image features in two dimensions.

[18] extended the warping algorithm [10] by introducing novel 2D warping methods which perform considerably better with a small additional computational effort. However,

from a biological perspective, it is unlikely that the warping method is implemented in the brain of an insect. Indeed, [19] showed that the 1D scheme already requires about 850 000 operations, which is approximately the total number of neurons that worker honeybees possess [20].

Contrary to previously described models, [21] divided the panoramic image into sub-areas centered around local features and used correlation techniques to determine best corresponding local views (used as landmarks) within the image. In their approach, when the robot reaches a place of interest, it moves around this new goal location to learn how to reach it from different neighbouring locations, by associating home vectors with the images acquired at each of these locations. This process creates an enlarged catchment area for the goal location. The angular differences between the local views and the sub-areas within the learned image provide a similarity measure between the current image and the reference snapshot taken at a goal location. Then the best matching panoramic image is chosen among the learned ones to select the homing vector that will get the robot to its goal. The authors argued that the recognition of these sub-areas around selected features facilitates a more robust scene recognition than classical global correlation performed without feature extraction. This approach better tolerates a lack of landmarks or the misinterpretation of a few of them, however it requires the robot to actively explore its environment around the goal location before deciding on the best homing direction.

B. View-based Methods

As demonstrated in [22], panoramic images can be used to perform homing in outdoor natural scenes because image differences increase smoothly with distance from a reference location. However, the catchment areas shrink when navigating in cluttered environments as the distance to objects decreases. Although the image difference function is sensitive to illumination variations and the consequent effects of shadows, it can be made robust to these effects by simple operations such as local contrast normalisation [23].

The homing strategy presented in [22] does not require any image matching but instead rests upon a gradient descent scheme that uses image differences (GDID). In order to measure the gradient of the image distance at the current image location, three sample locations (2D navigation) and their associated image distances need to be used. In addition, these image distances must be taken in three non-collinear directions (usually the three locations form a right-angled triangle in 2D space, or a small tetrahedron in 3D space). The home vector is then simply the negative of this gradient vector. This scheme requires the use of an external compass (so that all image distances are computed from views oriented in the same direction), however it can be replaced by minimising the image distances over rotation [22].

When it comes to applying the GDID scheme to robot navigation, the sampling of these three image distances requires exploratory movements which can be unfeasible. Indeed, they must involve sharp turns, in order to best estimate the gradient. Furthermore, this technique increases the overall length of the homeward trajectory. In addition, as the relative vectors between these sample locations need to be known, the resulting gradient estimation is affected by odometry errors.

In order to avoid these exploratory movements, [24] developed a novel model based on matched-filter descent in image distances. The core idea is to estimate the gradient of the image distances by using two predicted images for movement along two perpendicular directions from the current location. These two image predictions are derived from the current image by applying two 'template' optic flow fields that correspond to the two perpendicular translations in the horizontal plane. Each of these templates consists of a typical pattern of optic flow vectors that displays a focus of expansion in the direction of movement, a focus of contraction in the opposite direction, and a region with horizontal flow vectors between them. These templates are computed using the equal distance assumption [10].

Similarly, [25] used 2D warped panoramic images to perform visual homing for a robot in a laboratory environment. Their scheme makes use of raw colour images and only performs simple image processing operations. The authors compare images by using the Euclidean distance in the RGB space. A gradient descent scheme is employed to perform local homing, where exploratory movements are avoided by synthesising the corresponding images from the current view. Their experiments showed good repeatability and performance despite a systematic bias in the goal location (signalled by a sudden change in gradient orientation).

Both descent and matched-filter descent in image differences (respectively DID and MFDID) performances depend on the spatial structure of the image distance function. This function must rise smoothly and monotonically with spatial distance in order to apply the concept of gradient descent. Indoor experiments conducted in [24] showed that MFDID consistently outperformed DID. This result may appear surprising as DID actually samples the image distance function whereas MFDID only approximates it. In addition, MFDID uses the equal distance assumption which is likely to be untrue in many environments. However, image distances were measured at positions 30cm distant from the current location in case of DID while the gradient was estimated using two synthetic images taken at infinitesimally close locations. Hence, the gradient obtained by MFDID is more sensitive to the fine local structure of the approximated image distance function. The authors concluded that DID could match the performance of the MFDID scheme by drastically decreasing the step size of the exploratory movements.

When navigating to a goal location in environments with an anisotropic distribution of landmarks, home vectors produced by GDID deviate from the true home direction. [26] applied Newton's method to the MFDID scheme in order to reduce these deviations, and demonstrated, on indoor image databases, substantially improved performance over MFDID.

[27] introduced the concept of dynamic snapshot matching (based on optic flow amplitudes), as opposed to the use

of static snapshots. This may be better applicable to flying insects in many situations, as they are more likely to be in constant motion. It describes a means by which honeybees can make use of statically camouflaged landmarks (with the same contrast and texture as the background) by generating motion parallax [28]. In effect, this new matching scheme makes use of the depth structure of the environment, thus offering a robust strategy for navigation. Additionally, this concept of dynamic snapshot matching can be applied to the detection and description of scene changes [29], which is useful in long-range homing to distinguish whether one has entered a new visual locale.

III. Long-range Visual Homing

In this section, long-range visual homing refers to the ability of an agent to reach a goal location which can be outside its visual field, by following a visual route.

From their snapshot model, [14] described a longer range navigation scheme where the model bee makes use of a goal-centred map, built from a stack of snapshots, each of which is associated with a home vector. The use of multiple snapshots is also supported by ant experiments [16].

[30] presented a novel approach for navigation in large-scale environments by augmenting the ALV homing scheme [12]. Their scheme for learning a route between two distinct locations consists of three components: foraging, visual route learning and visual navigation. First, the agent keeps track of its random foraging journey from the starting location by performing path integration until it reaches its goal location. A global home vector is then stored. On the first homeward run (navigating using path integration), an AL vector, pointing towards the goal, is calculated at each time-step and is then stored as soon as it significantly changes from the one computed at the previous time-step. Indeed, a discontinuity in the ALV space is representative of the appearance or disappearance of an object from the visual field, signifying that the agent has reached a new visual locale. Consequently, an ordered series of AL vectors is accumulated during the homeward journey. On its second outbound journey, the previously stored global home vector is now used to determine the direction of the goal location. The agent navigates using path integration and, similarly, accumulates a second ordered series of AL vectors (or waypoints). From then on, visual navigation uses only the ALVs for navigation in either direction. The standard ALV homing scheme is used to home to the first waypoint (AL vector). Once the difference between the current ALV and the intermediate goal location approaches zero, the agent sets the next waypoint in the series as the goal and repeats the process until it gets to its home location. Even if the simulated environment used in this approach does not cope with real issues such as object segmentation (the landmarks used were black cylinders) and noise in visual input, the authors demonstrate a basic strategy which guides an agent with continuous sensory feedback through successive waypoints in a large-scale environment.

In photo-realistic simulations, [31] used a view-based method where learning a route involves storing a new snapshot whenever the angular difference between the direction computed by the visual homing scheme (MFDID, [24]) and that suggested by odometry exceeds a certain threshold angle, thus building a sequence of snapshots, each of which is stored with an associated odometry motion vector. In this scheme, route following relies on the stored odometry vectors, and visual homing (MFDID) acts to correct for errors to reach each intermediate goal location. Arrival at a waypoint is declared when the image difference drops below a certain threshold. Even if high values for this threshold lead to shorter routes (as arrivals at waypoints are detected earlier), the author chose a smaller, conservative, value. Indeed, if the threshold is set too high, the robot can switch from one waypoint to the next prematurely, and lose its route. Performance was shown to rely upon a certain degree of stability with regard to the environmental conditions (illumination for instance) during the learning and route following stages.

To avoid the environment-specific tuning required by methods using thresholds for waypoints selection and retrieval, [32] developed a non-threshold-based framework to link local image-based homing methods into a route. Their linked local navigation (LLN) scheme assumes that an agent can perform a single journey to a goal location, using path integration, for instance, while constructing the route waypoints by storing views of its environment as the number of perceived landmarks changes. Similarly, when navigating, a waypoint is assumed to be reached if the number of landmarks changes. Rather than formally counting landmarks, the agent is supposed to notice when a feature appears or disappears. This binary and significant event specifies a visual locale and appears to be biologically plausible [33]. The LLN scheme still uses the ALV model to perform local navigation (even if it could be replaced by any image-based algorithm that extracts features). Its success requires waypoints to be spread out in time, as failures are more likely to occur if waypoints are close to each other (consequently involving more boundary crossing between locales, leading to more failures caused by perceptual aliasing, for instance).

Although feature extraction can be fast, it often requires assumptions to be made about the type of features and the structure of the environment. In addition, visual landmarks are not usually obvious in natural scenes. Consequently, [34] proposed to use 2D raw colour panoramic images to perform short and long-range navigation. Local homing is based on the gradient scheme from [25] previously described, and longer range navigation is achieved by successively homing to a series of intermediate goal locations (way-images). Contrary to [25], instead of monitoring a sudden change in gradient orientation to detect that a way-image has been reached (a situation which does not occur anymore in most cases due to a faster implementation and subsequent higher frame rate), their scheme simply keeps track of whether the robot has been behind the way-image (monitoring when the commanded absolute turning angle, difference between the gradient direction and the current orientation, rapidly increases). The author claims a significantly faster processing speed (on 2D images) compared to the efficient visual homing scheme done in the frequency domain by [11]. His method reduces aliasing issues (encountered in 1D images), it does not require feature extraction and only performs simple image processing operations. Long-range navigation is achieved as a succession of short-range homing steps, demonstrating good repeatability although the remaining weakness is one of reliably detecting of when each waypoint is passed.

As the need to robustly determine that a waypoint has been reached during navigation still remains a non-trivial problem, [35] proposed a novel scheme for navigating without waypoints. A route is no longer defined as a series of discrete waypoints but is learned more holistically instead. In this framework, a classifier is used in order to predict whether a given view is on or off the learned route. Route following is achieved by scanning the environment and moving in the direction which seems to be part of the route. Haar-like features are extracted from positive views (forward-facing views belonging to the route) and negative ones (left and right-facing views not part of the route) and used to train a boosted classifier. The authors demonstrated the possibility of using a simple view-classification strategy to learn a non-trivial route. A classifier provides a compact way of storing information for route recognition and also a measure of the expected uncertainty of the classification. The simple interaction between a behavioural strategy (that is to say visual scanning, a behaviour that is supported by observations that ants will often turn rapidly on the spot [36]) and learned information, provides a robust scheme for route following. It also demonstrates that routes can be represented holistically, and that route recapitulation can be described as a recognition problem rather than a recall one, using familiarity with, rather than similarity to, a particular reference snapshot.

IV. Homing Models Applied to UAS Navigation

Many visual homing schemes have been tested in simulation or indoor environments using mobile terrestrial robots in 2D. Although most of these strategies are inspired by flying insects, relatively few examples of application to UAS can be found in the literature. In this section, the discussion is restricted to the bio-inspired techniques developed for airborne vehicles.

[37] extracted relative positions of three artificial landmarks (red cylinders) to guide a blimp-type flying robot in an indoor environment. Their model consists in controlling only yaw and height (roll and pitch are controlled passively due to gravity and the blimp-nature of the platform) by keeping the centre-most landmark centered horizontally, and vertically, respectively, in the image. The moving direction is then decided using an area-action relation based on the geometrical relation between landmarks, home location and the robot. Although it does not require any absolute compass, their scheme requires the three landmarks to be seen at all times. We do not extend our discussion to visual servoing techniques as most of them use complex artificial markers (concentric circles, H-shaped landing pad, etc.), unlikely to be found in natural environments and used by insects.

In the UAS domain, hovering is a straightforward application of local visual homing. The snapshot concept is used in [38] for the control of hover of a quadrotor by using a template matching technique to compute optic flow between the stored visual snapshot and the current image. One of the limitations of the snapshot hover is that a measure of scale is required to bootstrap the algorithm, provided here by integrating the image loom. Snapshot matching schemes eliminate the long-term drift in hover that would occur if optic flow measurements were integrated. Another example using snapshot matching is the bio-inspired technique described in [39], [40] for 3D hover in natural environments.

Inspired by the fly's visual micro-scanning movements, [41] developed a small-scale artificial compound eye, which estimates displacement by measuring angular positions of contrasting features, and mounted it on a tethered robot flying indoors to enable hover and initial position retrieval after perturbations. Another insight from biology is that insect employ omnidirectional vision. In the computation of visual egomotion, the use of panoramic vision makes translation easily distinguishable from rotation [42]. This observation has led to the development of bio-inspired miniature optic flow sensors for the guidance of ultra-lightweight UAS [43]–[45]. Successful goal-directed navigation has been demonstrated with a blimp-like platform by using bio-inspired elementary motion detectors (EMDs) for course stabilisation and visual odometry [46].

[47] proposed an implementation of a view-based navigation strategy to allow a quadrotor to follow a route defined by a concatenation of stored image waypoints. The localisation process consists of finding the image which best fits the current view in the visual memory. The matching is performed by computing a cross correlation score on features extracted using the Harris corner detector. Although the authors presented a method for the autonomous navigation of a UAS using a single camera in an indoor environment, the image matching is done on a ground station with an average computational time of 94ms/image to match a mean of 73 robust features for each frame.

V. Discussion

To achieve local visual homing, the ALV algorithm [12] appears to be the most parsimonious scheme among the landmark-based models that have been applied to autonomous mobile robots. Nevertheless, as this method only considers the angular direction to the visual landmarks, it requires robust feature extraction and is strongly dependent on a good compass accuracy. View-based methods appear to be more suited for outdoor navigation as they do not require the extraction of features, which can be a significant problem in natural, unstructured environments. Image differences increase monotonically and smoothly from the home location [22], making methods like gradient descent schemes suitable for outdoor navigation. However, one remaining crucial issue is the strong reliance on the ability to align current and stored views, which is a non-trivial problem for autonomous flying robots moving in a 3D environment. Moreover, to obtain a good estimate of the direction and magnitude of the gradient, exploratory movements are needed unless one uses the equal distance assumption, which is not very realistic for natural scenes. As the computation of the gradient requires the execution of sharp turns, these periodic exploratory movements disrupt the rendering of a smooth and efficient trajectory when performing visual homing onboard UAS in an unstructured 3D environment.

Visual homing models inspired by biological studies perform well in small-scale environments but tend to degrade significantly with the increasing scale of the environment. During long-range homing, when an agent uses waypoints in order to reach its home location, one unsolved problem is that of crossing robustly from one visual locale to the next. To date, in order to detect the reaching of a waypoint, the extant methods rely strongly either on an environment-dependent image discontinuity threshold [30], [31], or on the odometric system [32], [34]. Although [35] described a navigation scheme without waypoints, their method requires the extraction of features to train a classifier. This may not always be feasible in the context of performing homing in unstructured 3D environments. Additionally, future work requires the ability to detect a navigational error (misdetection of the arrival at a waypoint). An indication of error can be obtained by comparing the movement vector provided by visual homing with that suggested by visual odometry and adopting of an appropriate corrective behaviour. For instance, desert ants seem to abandon landmark navigation when it leads them in a direction which differs from that suggested by path integration [48].

To summarise, the applicability and implementation of a view-based method for UAS navigation in unstructured 3D environments, which can deal with image alignment errors and avoids undesirable exploratory movements at every time step, is a continuing challenge for explaining visual homing in animals, as well as devising biologically inspired homing algorithms. Additionally, a practical and reliable long-range visual homing scheme for autonomous aerial navigation in outdoor natural scenes, which is less environment-dependent, and robust to cumulative errors in odometry, still remains to be developed.

VI. Conclusion

This paper provides a review of bio-inspired visual homing models applied to the field of autonomous navigation. Vision-based techniques, enhanced by biological insights into the navigational strategies of insects, are of particular interest for UAS guidance where payload limitations and GPS-independent autonomy represent an even higher stake than in traditional mobile robotics. In natural environments, view-based methods have the advantage of not requiring feature extraction (which can be costly and challenging) but often depends on accurate image alignment. Future work would require the development of novel approaches for UAS navigation, either making use of sparse, yet robust route descriptions and taking advantage of efficient, bio-inspired, holistic local homing strategies to successfully reach intermediary waypoints en route to the goal location, or directly by using holistic route representations to recognise and follow the correct homeward journey.

References

[1] M. V. Srinivasan, "Honeybees as a Model for the Study of Visually Guided Flight, Navigation, and Biologically Inspired Robotics," *Physiological Reviews*, vol. 91, no. 2, pp. 413–460, 2011.

[2] H. Chao, Y. Gu, and M. Napolitano, "A Survey of Optical Flow Techniques for Robotics Navigation Applications," *Journal of Intelligent & Robotic Systems*, vol. 73, no. 1-4, pp. 361–372, Jan. 2014.

[3] J. Fuentes-Pacheco, J. Ruiz-Ascencio, and J. M. Rendón-Mancha, "Visual simultaneous localization and mapping: a survey," *Artificial Intelligence Review*, vol. 43, no. 1, pp. 55–81, Jan. 2015.

[4] F. Dyer, "Spatial memory and navigation by honeybees on the scale of the foraging range," *The Journal of Experimental Biology*, vol. 199, no. 1, pp. 147–154, 1996.

[5] R. Möller, "Do Insects Use Templates or Parameters for Landmark Navigation?" *Journal of Theoretical Biology*, vol. 210, no. 1, pp. 33–45, May 2001.

[6] J. Hong, X. Tan, B. Pinette, R. Weiss, and E. Riseman, "Image-Based Homing," *IEEE Control Systems*, vol. 12, no. 1, pp. 38 –45, Feb. 1992.

[7] A. Philippides, B. Baddeley, K. Cheng, and P. Graham, "How might ants use panoramic views for route navigation?" *Journal of Experimental Biology*, vol. 214, no. 3, pp. 445–451, Feb. 2011.

[8] B. A. Cartwright and T. S. Collett, "Landmark Learning in Bees," *Journal of Comparative Physiology*, vol. 151, no. 4, pp. 521–543, 1983.

[9] T. Röfer, "Controlling a Wheelchair with Image-based Homing," in *Proceedings of the AISB Symp. on Spatial Reason. in Mobile Robots and Animals, Manchester University, UK*, 1997.

[10] M. O. Franz, B. Schölkopf, H. A. Mallot, and H. H. Bülthoff, "Where did I take that snapshot? Scene-based homing by image matching," *Biological Cybernetics*, vol. 79, no. 3, pp. 191–202, Oct. 1998.

[11] W. Stürzl and H. Mallot, "Efficient visual homing based on Fourier transformed panoramic images," *Robotics and Autonomous Systems*, vol. 54, no. 4, pp. 300–313, Apr. 2006.

[12] D. Lambrinos, R. Möller, T. Labhart, R. Pfeifer, and R. Wehner, "A mobile robot employing insect strategies for navigation," *Robotics and Autonomous Systems*, vol. 30, no. 1–2, pp. 39–64, Jan. 2000.

[13] R. Möller, "Insect visual homing strategies in a robot with analog processing," *Biological Cybernetics*, vol. 83, no. 3, pp. 231–243, Aug. 2000.

[14] B. A. Cartwright and T. S. Collett, "Landmark Maps for Honeybees," *Biological Cybernetics*, vol. 57, no. 1-2, pp. 85–93, 1987.

[15] T. S. Collett and J. Baron, "Biological compasses and the coordinate frame of landmark memories in honeybees," *Published online: 10 March 1994*, vol. 368, no. 6467, pp. 137–140, Mar. 1994.

[16] S. P. D. Judd and T. S. Collett, "Multiple stored views and landmark guidance in ants," *Nature*, vol. 392, no. 6677, pp. 710–714, Apr. 1998.

[17] A. Vardy and F. Oppacher, "Low-Level Visual Homing," in *Advances in Artificial Life*, ser. Lecture Notes in Computer Science, W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, and J. T. Kim, Eds. Springer Berlin Heidelberg, Jan. 2003, no. 2801, pp. 875–884.

[18] R. Möller, "Local visual homing by warping of two-dimensional images," *Robotics and Autonomous Systems*, vol. 57, no. 1, pp. 87–101, Jan. 2009.

[19] A. Vardy and R. Möller, "Biologically plausible visual homing methods based on optical flow techniques," *Connection Science*, vol. 17, no. 1-2, pp. 47–89, 2005.

[20] H. Hertel and U. Maronde, "Processing of Visual Information in the Honeybee Brain," *Neurobiology and Behavior of Honeybees*, pp. 141–157, 1987.

[21] P. Gaussier, C. Joulain, J. Banquet, S. Leprêtre, and A. Revel, "The visual homing problem: An example of robotics/biology cross fertilization," *Robotics and Autonomous Systems*, vol. 30, no. 1–2, pp. 155–180, Jan. 2000.

[22] J. Zeil, M. I. Hofmann, and J. S. Chahl, "Catchment areas of panoramic snapshots in outdoor scenes," *Journal of the Optical Society of America A*, vol. 20, no. 3, pp. 450–469, Mar. 2003.

[23] W. Stürzl and J. Zeil, "Depth, contrast and view-based homing in outdoor scenes," *Biological Cybernetics*, vol. 96, no. 5, pp. 519–531, May 2007.

[24] R. Möller and A. Vardy, "Local visual homing by matched-filter descent in image distances," *Biological Cybernetics*, vol. 95, no. 5, pp. 413–430, Nov. 2006.

[25] D. Binding and F. Labrosse, "Visual local navigation using warped panoramic images," in *Proceedings of Towards Autonomous Robotic Systems. University of Surrey, Guildford, UK*, 2006, pp. 19–26.

[26] R. Möller, A. Vardy, S. Kreft, and S. Ruwisch, "Visual homing in environments with anisotropic landmark distribution," *Autonomous Robots*, vol. 23, no. 3, pp. 231–245, Oct. 2007.

[27] L. Dittmar, "Static and dynamic snapshots for goal localization in insects?" *Communicative & Integrative Biology*, vol. 4, no. 1, pp. 17–20, 2011.

[28] L. Dittmar, W. Stürzl, E. Baird, N. Boeddeker, and M. Egelhaaf, "Goal seeking in honeybees: matching of optic flow snapshots?" *The Journal of Experimental Biology*, vol. 213, no. 17, pp. 2913–2923, Sep. 2010.

[29] N. Nourani-Vatani, P. V. K. Borges, J. M. Roberts, and M. V. Srinivasan, "On the Use of Optical Flow for Scene Change Detection and Description," *Journal of Intelligent & Robotic Systems*, vol. 74, no. 3-4, pp. 817–846, Jun. 2014.

[30] L. Smith, A. Philippides, and P. Husbands, "Navigation in Large-Scale Environments Using an Augmented Model of Visual Homing," in *From Animals to Animats 9*, ser. Lecture Notes in Computer Science, S. Nolfi, G. Baldassarre, R. Calabretta, J. C. T. Hallam, D. Marocco, J.-A. Meyer, O. Miglino, and D. Parisi, Eds. Springer Berlin Heidelberg, Jan. 2006, no. 4095, pp. 251–262.

[31] A. Vardy, "Long-Range Visual Homing," in *IEEE International Conference on Robotics and Biomimetics, 2006. ROBIO'06*, Dec. 2006, pp. 220 –226.

[32] L. Smith, A. Philippides, P. Graham, B. Baddeley, and P. Husbands, "Linked Local Navigation for Visual Route Guidance," *Adaptive Behavior*, vol. 15, no. 3, pp. 257–271, Jan. 2007.

[33] R. Möller, "Insects Could Exploit UV–Green Contrast for Landmark Navigation," *Journal of Theoretical Biology*, vol. 214, no. 4, pp. 619–631, Feb. 2002.

[34] F. Labrosse, "Short and long-range visual navigation using warped panoramic images," *Robotics and Autonomous Systems*, vol. 55, no. 9, pp. 675–684, Sep. 2007.

[35] B. Baddeley, P. Graham, A. Philippides, and P. Husbands, "Holistic visual encoding of ant-like routes: Navigation without waypoints," *Adaptive Behavior*, vol. 19, no. 1, pp. 3–15, Jan. 2011.

[36] A. Wystrach, A. Philippides, A. Aurejac, K. Cheng, and P. Graham, "Visual scanning behaviours and their role in the navigation of the Australian desert ant *Melophorus bagoti*," *Journal of Comparative Physiology A*, vol. 200, no. 7, pp. 615–626, Jul. 2014.

[37] H. Kobayashi, K. Kikuchi, K. Ochi, and Y. Onogi, "Navigation Strategies Referring to Insect Homing in Flying Robots," in *IEEE International Conference on Robotics and Automation, 2001. Proceedings 2001 ICRA*, vol. 2, 2001, pp. 1695 – 1700 vol.2.

[38] P. Li, M. Garratt, and A. Lambert, "Monocular Snapshot-based Sensing and Control of Hover, Takeoff, and Landing for a Low-cost Quadrotor: Monocular Snapshot-based Sensing and Control," *Journal of Field Robotics*, Mar. 2015.

[39] A. Denuelle, S. Thurrowgood, F. Kendoul, and M. V. Srinivasan, "A view-based method for local homing of unmanned rotorcraft," in *Automation, Robotics and Applications (ICARA), 2015 6th International Conference on*. IEEE, 2015, pp. 443–449.

[40] A. Denuelle, S. Thurrowgood, R. Strydom, F. Kendoul, and M. V. Srinivasan, "Biologically-inspired Visual Stabilization of a Rotorcraft UAV in Unknown Outdoor Environments," in *Unmanned Aircraft Systems (ICUAS), 2015 International Conference on*. IEEE, 2015, pp. 1084–1093.

[41] F. Colonnier, A. Manecy, R. Juston, H. Mallot, R. Leitel, D. Floreano, and S. Viollet, "A small-scale hyperacute compound eye featuring active eye tremor: application to visual stabilization, target tracking, and short-range odometry," *Bioinspiration & Biomimetics*, vol. 10, no. 2, p. 026002, Feb. 2015.

[42] R. C. Nelson and J. Aloimonos, "Finding Motion Parameters from Spherical Motion Fields (Or the Advantages of Having Eyes in the Back of Your Head)," *Biological Cybernetics*, vol. 58, no. 4, pp. 261–273, 1988.

[43] W. E. Green, P. Y. Oh, and G. Barrows, "Flying Insect Inspired Vision for Autonomous Aerial Robot Maneuvers in Near-Earth Environments," in *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 2347–2352.

[44] J.-C. Zufferey and D. Floreano, "Fly-Inspired Visual Steering of an Ultralight Indoor Aircraft," *IEEE Transactions on Robotics*, vol. 22, no. 1, pp. 137–146, Feb. 2006.

[45] P.-E. J. Duhamel, C. O. Perez-Arancibia, G. L. Barrows, and R. J. Wood, "Biologically Inspired Optical-Flow Sensing for Altitude Control of Flapping-Wing Microrobots," *IEEE/ASME Transactions on Mechatronics*, vol. 18, no. 2, pp. 556–568, Apr. 2013.

[46] F. Iida, "Biologically inspired visual odometer for navigation of a flying robot," *Robotics and Autonomous Systems*, vol. 44, no. 3-4, pp. 201–208, Sep. 2003.

[47] J. Courbon, Y. Mezouar, N. Guenard, and P. Martinet, "Visual navigation of a quadrotor aerial vehicle," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. IROS 2009*, Oct. 2009, pp. 5315 –5320.

[48] M. Collett, T. Collett, S. Bisch, and R. Wehner, "Local and global vectors in desert ant navigation," *Nature*, vol. 394, pp. 269–272, 1998.