

Homing in Scale Space

David Churchill & Andrew Vardy
Department of Computer Science
Memorial University of Newfoundland
St. John's, Canada
www.cs.mun.ca/av

Abstract—Local visual homing is the process of determining the direction of movement required to return an agent to a goal location by comparing the current image with an image taken at the goal, known as the snapshot image. One way of accomplishing visual homing is by computing the correspondences between features and then analyzing the resulting flow field to determine the correct direction of motion. Typically, some strong assumptions need to be posited in order to compute the home direction from the flow field. For example, it is difficult to locally distinguish translation from rotation, so many authors assume rotation to be computable by other means (e.g. magnetic compass). In this paper we present a novel approach to visual homing using scale change information from Scale Invariant Feature Transforms (SIFT) which we use to compute landmark correspondences. The method described here is able to determine the direction of the goal in the robot's frame of reference, irrespective of the relative 3D orientation with the goal.

I. INTRODUCTION

Visual homing is the ability of an agent to return to a goal position by comparing the currently viewed image with an image captured at the goal, known as the snapshot image. It has been shown that insects such as bees and ants have the ability to visually home and that this is a crucial component in their overall navigational strategy [1]. Visual homing has been utilized in robotics as a means of executing learned paths [2], [3] and travelling between the nodes of a topological map [4], [5]. In this paper we propose a new visual homing method which is far less constrained than existing methods in that it can infer the direction of translation without any estimation of the direction of rotation. Thus, it does not require the current and snapshot images to be captured from the same 3D orientation.

Existing methods for visual based homing can be classified as holistic or correspondence-based [6]. Holistic methods rely on comparisons between images as a whole. An example of a holistic method is the method of Zeil et al. who posit a simple distance metric between images and implement homing as gradient descent in the space of this distance metric [7]. This method, while elegant in its simplicity, relies on the existence of a monotonic relationship between image distance and spatial distance. It also requires small exploratory movements of the robot in order to determine the gradient of the image distance function. Möller and Vardy described an alternative method based on gradient descent that removes the need for exploratory movements prior to computing a home vector [6]. Another holistic method is the so-called *warping method* of Franz et al. which searches for




Fig. 1. Ideal flow field for pure translation in a panoramic image [6].

the parameters of motion which make the warped snapshot image most similar to the current image. A warped snapshot image is generated by transforming the snapshot image *as if* the robot had actually moved according to the given motion parameters. To make this transformation tractable the assumption is made that all objects are equidistant from the goal. Despite the clearly unrealistic nature of this assumption, the warping method has been found to perform robustly in various indoor environments. In this paper we utilize the warping method to benchmark the performance of our algorithm. Correspondence methods detect features in the current image and attempt to correspond them with similar features in the snapshot image. The flow field that results is then interpreted to yield the direction of motion. Various features have been utilized, ranging in sophistication from raw image windows [8] to descriptors based on the Fourier-Mellin transform [9]. Recently, SIFT features have gained great popularity in many areas of computer vision and robotics due to the stability of their descriptor vectors with respect to changes in scaling, rotation, and illumination [10]. SIFT features have also been used in visual homing [11], [12], [13]. After finding the correspondences between features, the resulting correspondence vectors give the shift of the features in image space. If both the snapshot and current images are taken from the same orientation in a planar environment it is possible to compute the home direction analytically from a

single correspondence vector [8]. If the orientation is not the same, one can utilize some form of compass, or search for the change in orientation which would minimize the difference between the two images [7], [12].

The method described in this paper is similar to correspondence methods in that it relies upon finding correspondences between features. However, our interpretation of the resulting correspondences is markedly different. Consider the flow field for pure translation of an agent equipped with an omnidirectional camera. The field has a characteristic structure with foci of expansion and contraction separated by 180° (see Figure 1). If objects are distributed uniformly in the environment, half of them will appear to have expanded, while the remaining half will appear to contract. Typical correspondence methods consider how the features have shifted but not whether they have expanded or contracted. The problem is that in the presence of rotation it becomes much more difficult to determine the home direction from feature shifts. Hence, the two-stage process referred to above. However, whether a feature has changed in scale is independent of any change in orientation between the two views. We utilize the change in scale of corresponding SIFT features to determine the centre of the *region of contraction* which corresponds to the home direction.

In section 2 of this paper we will describe the Scale Invariant Feature Transform and how its parameters are related to our homing method. Section 3 defines many of the required concepts and tools to perform homing, as well introduces the method of homing in scale space. In section 4 we compare the performance of our method to the warping method. We discuss some future works related to our method in section 5, followed by some concluding remarks in section 6.

II. SCALE INVARIANT FEATURE TRANSFORM

The Scale Invariant Feature Transform, developed by Lowe [10] provides a robust keypoint descriptor which is able to match keypoints between images invariant of scale, rotation, or illumination. The process of determining a SIFT keypoint involves four stages. The first stage is the detection of scale space extrema. This stage involves the repeated convolution of the original image by Gaussian filters of increasing effective scale (values of σ) within each image octave. Each successive blur within the octave is subtracted from the previous to form a difference of Gaussians (DoG) image. Local maxima are detected in each scale of DoG space.

The second stage of the process involves the localization of the keypoint. In this step, keypoints are filtered for various reasons such as edge response, and contrast. The third stage of SIFT assigns an orientation to the keypoint based on the local gradient within the image. Lastly, the keypoint descriptor, which contains a number of orientation histograms relating to the gradients surrounding the keypoint, is constructed. This descriptor is stored with respect to both the scale and orientation of the keypoint, making it invariant to rotation, scale, as well as moderate 3D transformations.




Fig. 2. Robot pose diagram.

What we are left with is a SIFT keypoint storing the (x, y) location within an image, the scale σ and orientation ρ at which it was found, as well as the keypoint descriptor vector itself. We can describe a keypoint \mathbf{f} as follows:

$$\mathbf{f} = \{f_x, f_y, f_\sigma, f_\rho, \mathbf{f}_{kpd}\}.$$

A unique feature of our homing algorithm is its usage of f_σ .

III. HOMING IN SCALE SPACE

One requirement of our method is that the direction of translation be visible within the robot's field of view. Therefore we utilize panoramic images. Let **CV** represent the image taken at the location (**cv**) of the current panoramic view from the robot's perspective, and **SS** be the image taken at the goal location (**ss**).

Consider the diagram shown in Figure 2. If the robot has moved from position **ss** to position **cv**, the distance from the robot to feature **a** will have increased. This will be true of any feature on the same side of the perpendicular bisector of the line joining **ss** and **cv**. Similarly, the distance to the feature **b** will have decreased. We assume that this change in distance will be reflected in a corresponding change in the f_σ of the feature. Thus, we can classify features as either expanding or contracting. If there are sufficiently many features which are distributed evenly on either side of the dividing line then approximately half of them should experience expansion, while the other half should experience contraction. If we further assume that the features are distributed approximately uniformly throughout the environment then the home direction will be aligned with the centre of the region of contraction.

Locating the center of the region of contraction from **SS** to **CV** will allow us to determine the direction of translation that would move the robot towards the goal. We will use the change in scale information from SIFT feature correspondences to extract the center of the region of contraction, which coincides with the home direction under the assumption of uniform feature distribution.




Fig. 3. SIFT matched correspondences between **CV** (above) and **SS** (below). Correspondences in (a) show a scale decrease from **SS** to **CV**, thus having $\beta > 0$, indicating contraction. Conversely in (b) we see features which have $\beta < 0$, indicating expansion. Since these two regions should ideally be separated by π , they will be combined with a weighted average in order to more accurately compute the center of contraction.

The panoramic images used are w pixels wide by h pixels high. These images represent a complete viewing angle of 2π in the horizontal direction, as well as α radians in the vertical direction. Thus, each pixel represents a spacing of δ_x radians along the x-axis, and δ_y radians along the y-axis, computable by:

$$\delta_x = \frac{2\pi}{w} \quad \delta_y = \frac{\alpha}{h} \quad (1)$$

We therefore can convert our SIFT feature \mathbf{f} with location (f_x, f_y) within the images to angular coordinates $(f_{\theta_x}, f_{\theta_y})$ by

$$f_{\theta_x} = f_x \delta_x \quad f_{\theta_y} = f_y \delta_y \quad (2)$$

in order to facilitate proper directional calculations.

Determining the center of the region of expansion or contraction requires detecting whether a feature has grown or shrunk with respect to its size in the snapshot image. If we revisit our SIFT feature vector, not only does it give us the location of a feature of an image, but also the scale σ at which it was detected. Since σ is the level of Gaussian blur, the magnitude of scale is directly related to the feature's size. Through the application of SIFT, image features will be increasingly low-pass filtered as the scale parameter increases. Keypoints are detected at image coordinates where a feature is blurred out of existence. Intuitively, larger features will disappear at greater scales. Therefore, given a positive SIFT match between features \mathbf{f}_{ss} and \mathbf{f}_{cv} with scale values of σ_{ss} and σ_{cv} respectively, we can calculate:

$$\beta = \sigma_{ss} - \sigma_{cv}. \quad (3)$$

If $\beta > 0$ then the feature has shrunk from **SS** to **CV**, and conversely if $\beta < 0$ the feature has grown. We have many keypoint correspondences, so we must compute the center of these regions of expansion and contraction in order to find the home direction.

From the correspondences in Figure 3, we can see that the desks in the left half of the image appear to have grown larger while the filing cabinet seems to have shrunk. Since the cabinet represents the region of contraction in **CV**, this is the direction we wish to move. The key point of our method lies in this fact: no additional interpretation of the flow field is required. Merely the sign of β is enough to identify the change in feature size. The location of the corresponding keypoint within **SS** is not needed, since we are only concerned with

the features in **CV** which have contracted. It remains for us to accurately locate the center of this region of contraction. Since the relationship between the angular orientation of **CV** and **SS** is not needed, our method achieves complete invariance to changes in relative orientation between the two images. Also, since this method does not rely on any notion of an image horizon, it is invariant to changes in relative 3D orientation and elevation. This claim will be satisfied if the following conditions hold: (1) the camera's field of view encompasses the true direction of translation, (2) a significant number of correct correspondences are found, (3) the corresponded features are approximately uniformly distributed throughout the environment.

Let us denote a matched feature pair $\mathbf{m} = (\mathbf{f}_{ss}, \mathbf{f}_{cv})$. To calculate the center of a particular region, we partition our set of correspondences $M = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_n\}$ into M_{pos} and M_{neg} based on the sign of β . To determine the center of these partitioned regions with respect to the robot's heading, we use the angular mean of the data, that is: given any set of angles $\theta_1, \theta_2, \dots, \theta_n$:

$$\bar{\theta}(\theta_1, \theta_2, \dots, \theta_n) = \arctan \left(\frac{\sum_{i=1}^n \sin(\theta_i)}{\sum_{i=1}^n \cos(\theta_i)} \right). \quad (4)$$

We will denote the angular mean of our partitions as $\bar{\theta}_{pos}$ and $\bar{\theta}_{neg}$ respectively. We argued in section 1 that the regions of expansion and contraction are separated by π radians. We can use this fact to reduce the error in our calculation by allowing both the centers of expansion and contraction to contribute to the final result. Since both are always separated by a constant of π under ideal conditions, $\bar{\theta}_{pos} = \bar{\theta}_{neg} + \pi$.

We wish to allow both regions to contribute in such a way that a certain amount of confidence can be given to either set of data. It is often the case that $|M_{pos}|$ is significantly greater than $|M_{negs}|$, or vice versa. In an effort to assign confidence to a partition, we will use its cardinality to perform a weighted average of the mean of the data. This will shift the final calculation in the direction of the region with the most correspondences. We can compute our final home angle θ_{homing} as follows:

$$\bar{s} = |M_{pos}| \sin(\bar{\theta}_{pos}) + |M_{neg}| (\sin(\bar{\theta}_{neg}) + \pi) \quad (5)$$

$$\bar{c} = |M_{pos}| \cos(\bar{\theta}_{pos}) + |M_{neg}| (\cos(\bar{\theta}_{neg}) + \pi) \quad (6)$$

and finally:

$$\theta_{\text{homing}} = \text{atan2}(\bar{s}, \bar{c}). \quad (7)$$

This value for theta represents our final home vector with respect to the robot reference frame. Experimentally this weighted scheme has consistently shown to be more accurate than simply computing the unweighted average of the means of these regions. We summarize below our algorithm for determining the home direction:

1. Acquire an image **CV** from current robot location.
2. Perform SIFT feature matching on **SS** and **CV** to obtain a set of n matched feature pairs of the form $M = \{\mathbf{m}_1, \mathbf{m}_2, \dots, \mathbf{m}_n\}$.
3. Partition M into M_{pos} and M_{neg} where pos, neg denote the sign of β from equation 3.
4. Calculate the angular means $\bar{\theta}_{\text{pos}}$ and $\bar{\theta}_{\text{neg}}$ based on the values of f_{θ_x} from \mathbf{f}_{cv} (i.e. the angular location of the x coordinate of the keypoint from **cv**).
5. Calculate the weighted angular mean of both $\bar{\theta}_{\text{pos}}$ and $\bar{\theta}_{\text{neg}} + \pi$ based on their cardinality as shown in equations 5-7.
6. Move the robot in the direction of the computed angle, θ_{homing} .

IV. RESULTS & DISCUSSION

We have compared the performance of our method with that of the warping method [14]. Results from both methods were obtained using the panoramic image database described in [8] and on other databases of images collected at the University of Bielefeld¹. All databases contain images taken with a camera facing upwards at a hyperbolic mirror, at equal spacing and elevation within a capture grid. We used each location within the database as the goal location, then calculated the home direction from all other locations. The databases we used were:

Database	Description	Images	Spacing
A1originalH	Robot Lab	9 × 16	30cm
Chall1H	Hallway End	9 × 19	50cm
Chall2H	Hallway Entrance	7 × 19	50cm
Kitchen1H	Small Kitchen	11 × 8	10cm
Moeller1H	Living Room	21 × 10	10cm

The method for performance evaluation we used is the average angular error between correct home vectors and our computed home vectors. Our homing method takes images located at **cv** and **ss** and returns θ_{homing} , the angle we compute to be the homing angle. Since images in the database were taken at known locations, we can compute the ideal home vector angle as follows: given the (x, y) locations of current view **cv** and a goal snapshot **ss** on an evenly spaced capture grid, we can compute

$$\theta_{\text{ideal}}(\mathbf{ss}, \mathbf{cv}) = \text{atan2}(\mathbf{ss}_y - \mathbf{cv}_y, \mathbf{ss}_x - \mathbf{cv}_x) \quad (8)$$

¹The images databases used here are publicly available from the following URL: <http://www.ti.uni-bielefeld.de/html/research/avardy/>.

thus, the angular error $\text{AE}(\mathbf{ss}, \mathbf{cv})$ can be found by:

$$\text{AE}(\mathbf{ss}, \mathbf{cv}) = |\theta_{\text{ideal}} - \theta_{\text{homing}}| \quad (9)$$

We can obtain an overall average angular error as follows:

$$\text{AAE}(\mathbf{ss}) = \frac{1}{mn} \sum_{x=1}^m \sum_{y=1}^n \text{AE}(\mathbf{ss}, \mathbf{cv}_{xy}). \quad (10)$$

Finally, to obtain a measure of performance for the entire image database **db** we can use the total average angular error $\text{TAAE}(\mathbf{db})$, which computes the overall average of $\text{AAE}(\mathbf{ss})$ having computed the home direction to each possible location as the snapshot:

$$\text{TAAE}(\mathbf{db}) = \frac{1}{mn} \sum_{x=1}^m \sum_{y=1}^n \text{AAE}(\mathbf{ss}_{xy}). \quad (11)$$

Our method relies on having a number of SIFT feature matches in order to robustly compute the center of the regions of expansion and contraction. We require as many keypoints as possible for this process. Fortunately, SIFT feature matching offers a number of parameters which can be changed in order to maximize keypoint production, while still maintaining accurate results [10]. The values changed from those of Lowe's original implementation are as follows:

- 1) The number of scales at which keypoints are extracted is increased from 3 to 6 to increase the number of overall keypoints, while maintaining feasible running time
- 2) The peak threshold for the magnitude of difference of Gaussian values is decreased from 0.08 to 0.01 in order to choose more keypoints from areas of low contrast, since indoor environments often contain such areas
- 3) The ratio of scores from best to second best SIFT matching has been decreased from 0.6 to 0.8. As discussed in [10], this change results in a marginal decrease in match accuracy while dramatically increasing the number of matches.

Parameters for the warping method were selected to ensure fairness with respect to running time. We selected the following values for the parameters of the warping method search space: $\text{RhoMax} = 0.95$, $\text{RhoSteps} = 36$, $\text{AlphaSteps} = 36$, and $\text{PsiSteps} = 36$ [14]. On an Intel Core2 2.13GHz processor, this parameter selection resulted in an average execution time for the warping method which was 4.8% faster per snapshot than our scale space method. We consider this to be a fair metric for results comparison. Prior to initiating the homing process, we randomly rotate each image in the database horizontally by $\theta_r \in [0, 2\pi)$. This ensures that our results will be strictly rotation invariant. We will also show some sample results of shifting the pixels in the image vertically by a random amount $y_r \in [-15, 15]$ pixels to test whether our method is invariant to changes in elevation above the surface of movement. The remaining area of the image is filled with black pixels.

We see from the results in the top half of Figure 4 that the AAE for the warping method on the A1originalH database is 224% larger than that of scale space homing.




Fig. 4. Scale space homing (top left) of A1originalH database resulting in TAAE = 12.4°. Warping method (top right) of A1originalH database resulting in TAAE = 27.8°. Scale space homing (bottom left) of Kitchen1H database resulting in TAAE = 22.5°. Warping method (bottom right) of Kitchen1H database resulting in TAAE = 46.4°. Gray scale values range from black (0°) to white (90° for A1originalH, 130° for Kitchen1H), with darker values representing better results.

Our results also indicate that while the warping method has a more even performance throughout this database, our method has extremely high accuracy for areas in the center. While the warping method does outperform our method for some locations along the edges of the room, the overall performance for our approach is considerably better. The bottom half of Figure 4 shows the results from the Kitchen1H database, which, due to the existence of an object with many similar repeating features has yielded poor results for the warping method. For this database, the error for warping is 206% larger than that of scale space homing. The maximum AE for warping on Kitchen1H was 126°, while our method yielded a maximum AE of 60°. Over 22% of the individual snapshot results from warping yielded higher AE values than the maximum of 60° from scale space homing.

The following table shows the results from both homing in scale space, as well as homing using warping performed on the five databases:

Database	HiSS TAAE	Warping TAAE
A1originalH	12.4°	27.8°
Chall1H	14.3°	33.2°
Chall2H	22.2°	48.1°
Kitchen1H	22.5°	46.4°
Moeller1H	24.3°	35.4°

Although these tests show the TAAE is lower for each database using HiSS, we are also curious about how HiSS compares to warping on a case by case basis. We form two

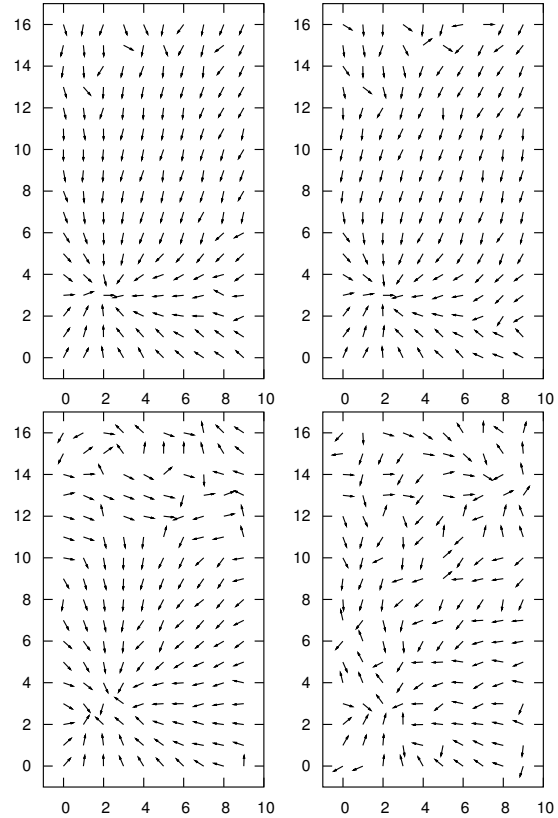


Fig. 5. Home vector images setting goal snapshot to (2,3). Top images show scale space homing with horizontal shift only (left, AAE=12.3°) and combined vertical shift (right, AAE=18.1°). Bottom images show warping method for horizontal shift only (left, AAE=39.2°) and combined vertical shift (right, AAE=59.4°)

vectors, \mathbf{h} (HiSS), \mathbf{w} (warping) consisting of the angular errors obtained by homing from all possible location within the database, to each possible home location. This allows us to perform a sign test [15] based on the median difference of $\mathbf{h} - \mathbf{w}$. In order to attempt to claim that our method performs better than warping, we use the alternative hypothesis that $\text{median}(\mathbf{h} - \mathbf{w}) < 0$. The following is a table of the median and p-values obtained from the sign test on this data:

Database	Median Difference	P-value
A1originalH	-0.0601545	2.2e-16
Chall1H	-0.130501	2.2e-16
Chall2H	-0.2333935	2.2e-16
Kitchen1H	-0.0991135	2.2e-16
Moeller1H	0	0.2085

so our hypothesis holds for all databases except Moeller1H, where the median value is actually 0 despite the large difference in mean angular error.

The results in Figure 5 show a sample location of (2,3) as our goal snapshot location. Homing was carried out to this location under both methods with horizontal shifting only (left), as well as the combined horizontal and vertical shift (right). From the images, it is clear that the warping method

suffers from the added vertical shift. Scale space homing does not suffer to the same degree and maintains an AE which is less than that of warping without any vertical shift.

These results show that our method outperforms the warping method for parameters which yield similar execution times. When we incorporate vertical shift into our database images to simulate variation in elevation, the performance gain shown by scale space homing is much greater.

V. FUTURE WORKS

Achieving success on the databases provided has created the need for further acquisition of images with varying parameters, such as changes in elevation and 3D orientation, as well as more natural features. A database similar to those mentioned in the results section could be taken at not only varied locations, but at varied heights and 3D orientations as well. The success of SIFT feature detection in natural environments [10] also lends a high probability that homing in scale space will be robust in outdoor, natural environments. Additionally, we could apply this method to airborne vehicles, performing homing in 3D by computing not just the one-dimensional angular mean, but the two-dimensional centroid of the region of contraction.

We will also plan to delve more deeply into methods for obtaining the center of the region of contraction when feature correspondences are few and/or erroneous. Although our weighted averaging scheme has shown good results, other more statistically sound methods may exist which yield even better results.

Scale space homing can also provide a measure of the likelihood that homing will succeed. Ideally, the regions of contraction and expansion are clearly separated and their centers are π radians apart. If these conditions do not hold then we may assume that homing will not be successful—perhaps because the images were taken from locations that are too far apart. We are currently investigating possible measures of homing success in terms of their correlation with measured homing success.

VI. CONCLUSION

We have described a method for performing visual homing using scale space data from Scale Invariant Feature Transforms. From our results, we have shown that this method of using SIFT scale change information to detect regions of expansion and contraction is robust with respect to arbitrary compass heading, as well as vertical shifts. By comparing ideal vectors to our computed vectors, we were able to show that the average angular error of homing in scale space is significantly less than that of the warping method, which has been widely used as a comparison in the field of local visual homing. From these observations we are able to conclude that homing in scale space is able to accurately determine goal directions while relying on fewer constraints than existing methods.

ACKNOWLEDGEMENTS

The authors wish to thank Ralf Möller for providing access to image databases collected at the University of Bielefeld and for the very helpful comments he made on an earlier draft of this paper.

REFERENCES

- [1] T. Collett and M. Collett, "Memory use in insect visual navigation," *Nature Reviews Neuroscience*, vol. 3, pp. 542–552, 2002.
- [2] A. Argyros, C. Bekris, S. Orphanoudakis, and L. Kavraki, "Robot homing by exploiting panoramic vision," *Journal of Autonomous Robots*, vol. 19, no. 1, pp. 7–25, 2005.
- [3] A. Vardy, "Long-range visual homing," in *Proceedings of the IEEE International Conference on Robotics and Biomimetics*. IEEE Xplore, 2006.
- [4] M. Franz, B. Schölkopf, H. Mallot, and H. Bülthoff, "Learning view graphs for robot navigation," *Autonomous Robots*, vol. 5, pp. 111–125, 1998.
- [5] W. Hübner and H. Mallot, "Metric embedding of view-graphs: A vision and odometry-based approach to cognitive mapping," *Autonomous Robots*, vol. 23, p. 183196, 2007.
- [6] R. Möller and A. Vardy, "Local visual homing by matched-filter descent in image distances," *Biological Cybernetics*, vol. 95, pp. 413–430, 2006.
- [7] J. Zeil, M. Hofmann, and J. Chahl, "Catchment areas of panoramic snapshots in outdoor scenes," *Journal of the Optical Society of America A*, vol. 20, no. 3, pp. 450–469, 2003.
- [8] A. Vardy and R. Möller, "Biologically plausible visual homing methods based on optical flow techniques," *Connection Science*, vol. 17, no. 1/2, pp. 47–90, 2005.
- [9] A. Rizzi, D. Duina, S. Inelli, and R. Cassinis, "Unsupervised matching of visual landmarks for robotic homing using Fourier-Mellin transform," in *Intelligent Autonomous Systems 6*, 2000, pp. 455–462.
- [10] D. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, 2004.
- [11] A. Briggs, Y. Li, D. Scharstein, and M. Wilder, "Robot navigation using 1d panoramic images," in *International Conference on Robotics and Automation*, 2006, pp. 2679–2685.
- [12] J. S. Pons, W. Hübner, J. Dahmen, and H. Mallot, "Vision-based robot homing in dynamic environments," in *13th IASTED International Conference on Robotics and Applications*, K. Schilling, Ed., 2007, pp. 293–298.
- [13] T. Goedemé, M. Nuttin, T. Tuytelaars, and L. Van Gool, "Omnidirectional vision based topological navigation," *International Journal of Computer Vision*, vol. 74, no. 3, pp. 219–236, 2007.
- [14] M. Franz, B. Schölkopf, H. Mallot, and H. Bülthoff, "Where did I take that snapshot? Scene-based homing by image matching," *Biological Cybernetics*, vol. 79, pp. 191–202, 1998.
- [15] J. Devore, *Probability and Statistics for Engineering and the Sciences*. Brooks/Cole, 1991.