

Direction Control for an Active Docking Behaviour Based on the Rotational Component of Log-Polar Optic Flow

Nick Barnes¹ and Giulio Sandini²

¹ Computer Vision and Machine Intelligence Lab (CVMIL),
Department of Computer Science and Software Engineering,
University of Melbourne, Victoria, 3010, AUSTRALIA
`nmb@cs.mu.oz.au`

² Laboratory for Integrated Advanced Robotics (LIRA-Lab),
Department of Communication Computer and Systems Science,
University of Genova, Via Opera Pia, 13, Genova, 16145, ITALY
`sandini@dist.unige.it`

Abstract. Docking is a fundamental requirement for a mobile robot in order to be able to interact with objects in its environment. In this paper we present an algorithm and implementation for a special case of the docking problem for ground-based robots. We require the robot to dock with a fixated environment point where only visual information is available. Specifically, camera pan/tilt information is unknown, as is the direction of motion with respect to the object and the robot's velocity. Further, camera calibration is unavailable. The aim is to minimise the difference between the camera optical axis and the robot heading direction. This constitutes a behaviour for controlling robot direction based on fixation. This paper presents a full mathematical derivation of the method and implementation used. In its most general form, the method requires partial segmentation of the optical flow field. The experiments presented, however, assume partial knowledge as to whether points are closer to the camera than the fixation point or further away. There are many scenarios in robotic navigation where such assumptions are typical working conditions. We examine two cases: convex objects; and distant background/floor. The solution presented uses only the rotational component of optical flow from a log-polar sensor. Results are presented with real image and ray-traced image sequences. The robot is controlled based on a single component of optical flow over a small portion of the image, and thus is suited to real-time implementation.

Keywords: Active vision and real-time vision, vision-guided mobile robots, and docking.

1 Introduction

Docking is a fundamental requirement for a mobile robot to interact with objects in its environment. In order to perform operations such manipulation (e.g.

autonomous fork-lifts [8]), or industrial assembly [12,9], a mobile robot must be able to dock. In this paper, we present the derivation and implementation of an active behaviour for controlling robot heading direction to support docking with a fixated environment point. Only visual information is required (i.e. pan/tilt information, and robot's heading direction and velocity are unknown). Further, camera calibration is not required. This method does not require knowledge of the object, other than a constraint on the distribution of the depth of points relative to the fixation point. Previously, we demonstrated that it was sufficient that the object be convex and centred in the image, or that the majority of background points be behind the fixated object [2]. This method can be used as a behaviour that is independent of fixation and high-level planning. The robot need only fixate on a point in the desired heading direction and invoke the behaviour, and then the robot will turn as required. This is elegant from an architectural viewpoint for general docking, but also facilitates systems where fixation is entirely separate from platform control. For example, consider a situation where fixation is controlled by a human operator with a camera attached to a head-set and no information is available about the head-set position.

This research will be integrated with our existing mobile robot system for circumnavigation [3], which uniquely identifies objects and moves around them. Integration of docking will facilitate close inspection and manipulation.

Fixation systems control camera direction to keep the projection of a scene point that is moving relative to the camera in a fixed position in the image. Fixation is fundamental to active vision, and as such there are many approaches available (e.g. [15,16]). Fixation can be used to facilitate perception of general motion [7]. Research with human subjects has shown that pedestrians can use fixation to gain information about their instantaneous heading direction [6].

Docking is a difficult problem, that is often handled with hardware solutions such as tactile sensors [13], or by using extensive knowledge about the visual properties of the object and camera [1,3,13]. Santos-Victor and Sandini [18] present an active approach to docking for a mobile robot that corrects heading direction, however, this assumes that the docking surface is planar.

2 The Log-Polar Sensor

Schwartz [19,20] derived an analytical formulation of biological vision systems based on experimental measures of the mapping from the retina to the visual cortex of monkeys. Visual data is transformed from the retinal plane in polar coordinates (ρ, θ) to log-polar Cartesian coordinates (ξ, γ) in the cortical plane. The relation can be expressed:

$$\xi = \log_a \frac{\rho}{\rho_0}, \gamma = q\eta, \quad (1)$$

where (ρ, η) are the polar coordinates of a point on the retinal plane and ρ_0 , q , and a are constants determined by the physical layout of the sensor. Thus,




Fig. 1. The log-polar sensor samples 64 evenly spaced angles, at 32 radii. The sensing elements increase in size toward the periphery. (a) sensor geometry; (b) image geometry.

sensing elements appear in a non-uniform distribution, with a high density at the central fovea, and continuously decreasing density toward the periphery.

A CMOS implementation simulating this type of sensor has been realised [14]. Figure 1 shows a log-polar image and its Cartesian reconstruction, illustrating the high-resolution at the fovea, and low resolution in the image periphery.

Jain [11] pointed out the advantages of using optical flow from a log-polar complex mapping for depth recovery from a translating camera with known motion parameters. The benefits of space-variant sensors for calculation of time-to-impact have also been demonstrated [14,21]. The method presented here demonstrates advantages for control of motion orthogonal to the image axis. The log-polar sensor parameterisation enables closed-loop robot heading direction control based directly on the rotational component of log-polar optical flow.

3 Theoretical Background

Consider a robot moving in three space, with a velocity vector $\mathbf{W} = (W_x, W_y, W_z)$ (see Figure 3). A camera mounted on the robot can move about all three axes with rotational velocities of $\boldsymbol{\omega} = (\theta, \phi, \psi)$, about (x, y, z) respectively.

Consider a point P on an object in the camera field of view, specified in camera coordinates (x, y, z) . Sandini and Tisarelli [17] derive the motion of P as rotational and translational components as follows. From the inverse perspective transform, the projected point on the image plane given a focal length F is:

$$P = [x, y] = \frac{F}{z}[x, y], \quad (2)$$

Differentiating (2) with respect to time, we may decompose the velocity vector into a component due to camera translation and due to camera rotation:




Fig. 2. (a) The coordinate system for a robot that moves in 3 space. (b) The coordinate system for the ground-based robot.

$$\begin{aligned} \mathbf{V} &= \mathbf{V}_t + \mathbf{V}_r \\ \mathbf{V}_t &= \left[\frac{xW_z - FW_x}{Z}, \frac{yW_z - FW_y}{Z} \right], \end{aligned} \quad (3)$$

$$\mathbf{V}_r = \left[\frac{xy\phi - [x^2 + F^2]\theta + y\psi}{F}, \frac{[y^2 + F^2]\phi - xy\theta - x\psi}{F} \right] \quad (4)$$

Now we derive equations for general motion observed by a log-polar sensor, following Tistarelli and Sandini [21]. The velocity in the image plane can be described in terms of radial and angular coordinates:

$$\dot{\rho} = \frac{xu + yv}{\rho} = u \cos \eta + v \sin \eta \quad (5)$$

$$\dot{\eta} = \frac{xv - yu}{\rho^2} = \frac{v \cos \eta - u \sin \eta}{\rho} \quad (6)$$

Substituting the motion equations (3) and (4):

$$\dot{\rho} = \left(\frac{xW_z - FW_x}{Z} + \frac{xy\psi - [x^2 + F^2]\theta + Fy\phi}{F} \right) \cos \eta + \left(\frac{yW_z - FW_y}{Z} + \frac{[y^2 + F^2]\psi - xy\theta - Fx\phi}{F} \right) \sin \eta \quad (7)$$

$$\dot{\eta} = \frac{1}{\rho} \left[\left(\frac{yW_z - FW_y}{Z} + \frac{xy\psi - [x^2 + F^2]\theta + Fy\phi}{F} \right) \cos \eta - \left(\frac{xW_z - FW_x}{Z} + \frac{xy\psi - [x^2 + F^2]\theta + Fy\phi}{F} \right) \sin \eta \right] \quad (8)$$

By substituting $x = \rho \cos \eta$ and $y = \rho \sin \eta$:

$$\dot{\rho} = \frac{1}{Z} [\rho W_z - F(W_x \cos \eta + W_y \sin \eta)] + \left(\frac{\rho^2}{F} \right) (\psi \sin \eta - \theta \cos \eta) \quad (9)$$

$$\dot{\eta} = \frac{F}{\rho} \left[\left(\frac{W_x}{Z} + \theta \right) \sin \eta + \left(\phi - \frac{W_y}{Z} \right) \cos \eta \right] - \phi. \quad (10)$$

However, the retinal sensor performs a logarithmic mapping as shown in Equation (1), thus we have [21]:

$$\begin{aligned}\dot{\xi} &= [\frac{1}{Z}[W_z - \frac{F}{\rho}(W_x \cos \frac{\gamma}{q} + W_y \sin \frac{\gamma}{q})] + (\frac{\rho}{F} + \frac{F}{\rho})(\phi \sin \frac{\gamma}{q} - \theta \cos \frac{\gamma}{q})] \log_a e, \\ \dot{\gamma} &= \frac{qF}{\rho} [(\frac{W_x}{Z} + \theta) \sin \frac{\gamma}{q} + (\phi - \frac{W_y}{Z}) \cos \frac{\gamma}{q}] - q\psi.\end{aligned}\quad (11)$$

Now let us consider the case of a ground-based robot that has a camera mounted on a pan-tilt platform, with no capacity for rotation about the optical axis (Figure 3). Thus, the robot moves with motion vector $\mathbf{W} = (W_x, W_z)$, and the camera has an angular velocity vector $\boldsymbol{\omega} = (\theta, \phi)$, given a combination of robot and pan/tilt platform motion. Consider also that the robot camera fixates independently on a target point in the environment. The pan and tilt velocities and absolute direction of the head are unknown, as is the absolute direction of the robot. Assume, without loss of generality, that the target object lies along the z axis, where the origin is a fixed to the robot along the optical axis.

We would like the robot to dock with the fixation point. To achieve this goal, the robot must adjust its heading direction such that W_x is zero. For a ground-based robot, the magnitude of W_y is not important. The issue of controlling the magnitude of velocity has been addressed previously [14,21].

If the robot is tracking a point in the environment, then from [21], we have:

$$\theta = \frac{-W_x}{D}, \phi = \frac{W_y}{D}, \quad (13)$$

where D is the distance to the point of fixation.

Substituting Equations (13) into Equations (11) and (12) to eliminate θ and ϕ , removing the redundant ψ , we obtain:

$$\dot{\xi} = [\frac{W_z}{Z} + \frac{1}{D}(\frac{\rho}{F} + \frac{F}{\rho}[1 - \frac{D}{Z}])](W_x \cos \frac{\gamma}{q} - W_y \sin \frac{\gamma}{q}) \log_a e, \quad (14)$$

$$\dot{\gamma} = \frac{qF}{\rho D}(1 - \frac{D}{Z})(W_y \cos \frac{\gamma}{q} - W_x \sin \frac{\gamma}{q}). \quad (15)$$

For Equation (15), if the object is not a long way above the ground relative to the robot-to-object distance, then the W_y residual from camera tilt will be small. Further, consider log-polar image region in which $\cos \frac{\gamma}{q}$ is close to zero, and $\sin \frac{\gamma}{q}$ is maximal, as shown in Figure 3 and specified in Equation (16).

$$\frac{\pi}{2} - k \bigcup \frac{3\pi}{2} + k, \quad (16)$$

where k is a constant specifying the width of the sensor region over which the mean is taken. Provided k is small, the coefficient of W_y is close to zero.

Thus, we may approximate $\dot{\gamma}$ in this region as:

$$\dot{\gamma} = \frac{qF}{\rho D}(\frac{D}{Z} - 1)W_x \sin \frac{\gamma}{q}. \quad (17)$$




Fig. 3. W_x dominates $\dot{\gamma}$ over the region shown in Equation (16).




Fig. 4. The intersection of the camera field of view and a sphere centred at the camera focal point with a radius the size of the fixation distance. The sign of $\dot{\gamma}$ is dependent on which side of this volume the world point appears.

From Equation 17, we can see that $\dot{\gamma}$ is directly proportional to W_x for this part of the image. The sign of $\dot{\gamma}$ is dependent on the sign of W_x and whether $D > Z$ for the given point. With no assumptions about scene formation $D > Z$ is unknown for any image point. Ideally, we could adjust the heading direction in either direction and check if $\dot{\gamma}$ reduces also. However, for the optical flow method used, the magnitude of $\dot{\gamma}$ can vary greatly for similar values of W_x , and so the second order derivative of motion cannot be determined with sufficient accuracy.

Note that in the region at the sides of the image, where $\cos \frac{\gamma}{q}$ is maximal, the effect on $\dot{\xi}$ of W_x will be maximal, however, for robot motion that is largely towards the object, $\dot{\xi}$ will be dominated by the expansion component, W_z , and so is less well suited than $\dot{\gamma}$ for direct use for control.

3.1 Geometric Constraints

As a log-polar sensor has high resolution at the fovea, and this resolution decreases toward the periphery, it is natural to fixate the object of interest at the centre of the image, i.e. the optical axis should point towards the object.

By definition, $D > Z$ for a given point P if the distance from the camera focal point is greater than that of the fixation point. By assuming the fixation point lies along the optical axis, we may define the region that is closer to the camera than the fixation point. It is the region contained within the intersection

of the cone defining camera visibility, and a sphere centred at the focal point, with radius equal to the fixation distance (see Figure 4). See [2] for an analysis.

In order to deduce robot heading angle relative to the optical axis we must be able to segment the visible scene into parts contained within this cone, or beyond it. A segmentation-based solution may be possible, however, in order to maintain processing speed, we assume constraints on scene geometry that allow $\dot{\gamma}$ to be used directly. Many assumptions are possible for particular situations, however, we wish to maintain the generality of the system. Two possible assumptions that have general application for ground-based mobile robots are:

1. The fixated object is convex and large in the image. Thus, most visible points are behind the fixation point from the robot's view point.
2. The fixation point is on the ground, such that all image points below it are ground points that are closer to the robot, and the majority of points above the object are background, and so behind the object. (This assumes the space between the object and the robot is not cluttered with other objects).

To clarify the theory above, Figure 5 shows optical flow for the motion discussed, and for several types of component motions.

4 Implementation

Many methods are available for calculating optical flow, with different strengths and weaknesses. See [4] for a comprehensive review of methods, and [10] for a review of the performance of these methods. Two major considerations drive our choice of method for calculating optical flow:

- Mobile robot docking is an on-line task, so any method must be capable of real-time performance; and,
- Our basic assumption on scene geometry does not facilitate the use of model-based methods.

The first consideration requires a fast method that avoids excessive computation. It should not deal with the whole image if a restricted part of the image will suffice. Although overall navigation performance must be robust to avoid unnecessary deviations, the optical flow method need not produce exact results. We avoided methods involving flow reconstruction, and instead used local flow calculation, relying on aggregation over part of the image to handle noise.

If knowledge is available about geometry of the docking surface, then methods can be applied that are based on assumptions such as that the surface is a first order function of image coordinates [22]. However, we wish to investigate a more general case where the only assumed knowledge is the approximate distance of points with respect to the fixation point.

The method chosen was that of Uras *et. al.* [23], which does not require assumptions about surface shape, is simple enough for fast implementation and is competitive in terms of accuracy [10]. Other fast local methods may also be appropriate. The method uses a local solution to the second order derivative:

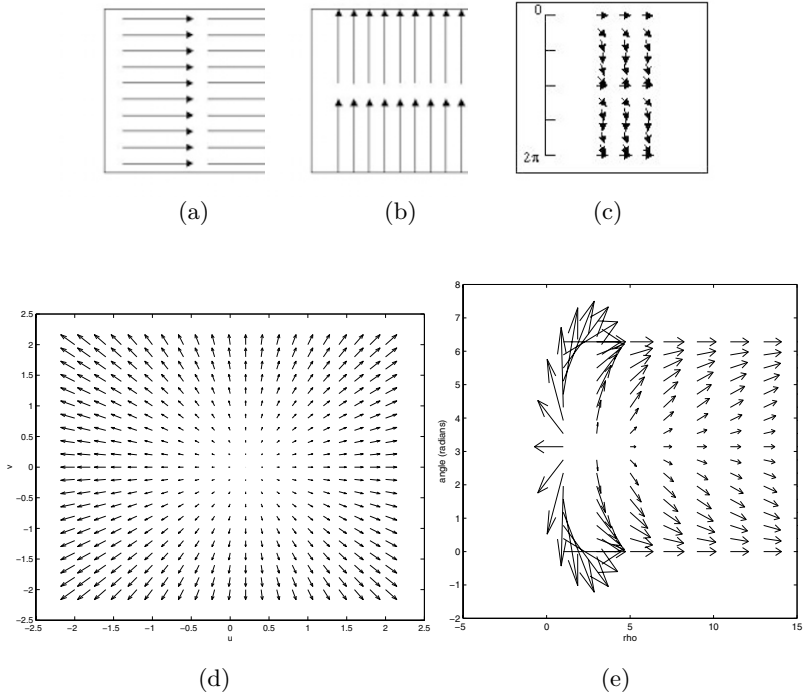


Fig. 5. Optical flow in polar coordinates. (a) The object moving closer to the camera. (b) Clockwise rotation in the plane. (c) Robot translation perpendicular to the image axis, while fixating on a point in front of the object. (d) Flow in Cartesian coordinates resulting from translation perpendicular to the image axis combined with motion toward the object, while fixating in front of the object. (e) In log-polar coordinates.

$$\begin{bmatrix} I_{xx} & I_{yx} \\ I_{xy} & I_{yy} \end{bmatrix} \begin{pmatrix} V_1 \\ V_2 \end{pmatrix} + \begin{pmatrix} I_{tx} \\ I_{ty} \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}. \quad (18)$$

In a log-polar image, values for optical flow can be calculated at each pixel. However, only angles where W_x dominates $\dot{\gamma}$ are used, as specified in Equation (16). Further, the log-polar sensor is a multi-scale device. The distance from the centre of the image at which the flow will have the most favourable signal to noise ratio is dependent on the flow scale. The flow was taken between two values ρ_u and ρ_l which were set manually dependent on geometry of the particular situation. Automated setting of scale is not addressed in this paper. We largely ignored noise in the flow calculation to facilitate speed, aiming instead for tolerance of noise in the input flow. We used the sign of the mean of $\dot{\gamma}$ (see Equation 17) to control heading direction, in a closed loop.

Note that there must be sufficient net optical flow for the signal-to-noise ratio to be adequate for robust control. For this to be the case, the depth between the

fixation point and the other image points must not be small with respect to the robot-to-object distance. For the second assumption, of a small object on the floor in a room, net flow will generally be sufficient given a distant background. However, as the robot moves close, the object may become large in the image so that little background is visible. Hardware final docking solutions may be appropriate at this stage (e.g. [1]). Alternatively, for the convex object assumption, the object must not be small with respect to the robot-to-object distance.

$$\mu = \sum_{\rho=\rho_l}^{\rho_u} \sum_{\gamma=\frac{\pi}{2}-k}^{\frac{\pi}{2}+k} \dot{\gamma} + \sum_{\rho=\rho_l}^{\rho_u} \sum_{\gamma=\frac{3\pi}{2}-k}^{\frac{3\pi}{2}+k} \dot{\gamma} \quad (19)$$

5 Results

We present experiments with real and simulated image sequences. All images are 256x256 Cartesian images that are subsampled into log-polar coordinates according to Equation (1). The simulated sequences show closed loop system performance. We have not yet completed the implementation on our mobile platform, however, we have taken real image sequences that confirm that real data does behave in the same manner as the simulated data. Thus, the system should perform correctly when the loop is closed.


5.1 Ray Traced Robot Simulations

Image sequences of textured objects were generated using the POV-Ray ray tracing package. These sequences generate noisy optical flow patterns as flow calculation incurs many of the difficulties that apply to real images. Simulation allow us to examine precise cases, enabling full evaluation of the mathematical theory. For example, it is difficult to ensure that the fixation point remains constant on the surface of a plane using real data.


For these experiments we used a robot simulator. The simulator takes an image every p msec, moving forward only with fixed velocity. The heading direction can be changed by applying an angular velocity. This is represented as an average value for the time interval. Thus, for an applied angular velocity of a radians per msec, the robot will turn ap radians during the interval.

Early simulated results have been presented previously [2]. The results presented here show more extensive trials of relevant situations.


Convex Object Assumption The algorithm was tested on two convex objects, using the assumption that all points are behind the fixation point. Figure 6(a) shows the first object, and (b) and (c) show the heading direction with respect to the optical axis against total distance travelled. The robot overshoots the zero heading direction in both cases. This is due partly to the fact that it is getting close to the object. It should be noted that the depth variation within the object is small in comparison with the distance from the object, the resulting flow is



(a)



(b)




(c)

Fig. 6. Docking with a stone-surfaced ellipsoid. (a) The object. (b) Docking when the initial direction of motion is at $\frac{\pi}{8}$ radians to the optical axis. (c) Initial direction of motion is at $\frac{\pi}{8}$ radians to the optical axis in the opposite direction.


also small, resulting in the noise seen in the experiment. Similar results were found for a convex polyhedral object with a brick-like surface.

Background assumption The second case is where the robot fixates on a small object in the foreground, and we assume that all points above the object are background (see Figure 7 (a)). The lower half of the image is floor, and a planar wall covers most of the upper half. This is a plausible setup for a ground-based robot attempting to dock with an object in a room with a background wall, and a flat floor. In this case, the system assessed only points in the upper part of the image. The floor could also have been used, however this configuration was chosen to be consistent with that used for the real image sequences. Figure 7(b) and (c) show plots of heading direction angle. In this case, convergence was faster and more stable, because the image points for which $\dot{\gamma}$ was calculated are further behind the object. Similar performance was also demonstrated with ray-traced images where all the points were in front of the fixation point.


Although theoretically, the algorithm should be able to correct heading direction with a planar surface, trials performed were unable to support this.



(a)



(b)



(c)

Fig. 7. Docking with an object with a distant background. (a) At the starting point, the object is small in the image, with a floor plane for all points below the object, and a background wall for all points behind the object. (b) Docking when motion begins at $\frac{\pi}{4}$ radians to the optical axis. (c) Docking when motion begins at $-\frac{\pi}{4}$ radians.

Mathematically there is a net rotational flow, however, in practice it appears this component is too small to be reliably extracted due to noise in this case.

5.2 Real Images

The object shown in Figure 8 object moved in a straight line along a rail. While it moved it was fixated by the LIRA head [5], using a colour-based binocular fixation method under development at the LIRA-lab. The images shown were taken from one of the cameras. Two types of image sequence were taken:

- The object moves from the right to the left of the image, toward the robot.
- The object moves from the left to the right of the image, toward the robot.

Figure 8 shows a sequence where the object moves at an angle such that in the coordinate system of Figure 3, W_x would be positive. With the object moving, motion is similar to when the robot moves, but the background is far away. This exemplifies the assumption that all above the object is background.




Fig. 8. A clown mask is fixated as it moves along a rail from right to left.

Figure 9 shows log-polar versions of the Cartesian images of Figure 8 (c) and (d), and the resulting log-polar optic flow. The scale of motion for this sequence is large, so the symmetric regions where $\dot{\gamma}$ was taken were towards the periphery of the image. Figure 10 shows the mean of $\dot{\gamma}$ for this sequence. Figure 11 shows a similar sequence, where the rail is at an angle with respect to the camera such that W_x would be negative. Figure 12 shows the mean of $\dot{\gamma}$.

These sequences show the sign of the mean changes with the direction of W_x . Due to slow recording of images, and minimum movement requirements between frames for the fixation algorithm, the changes between the images are large. More stable results can be expected with a higher sampling rate.

6 Conclusion

In this paper, we derived an algorithm which could be used to control heading direction for docking with an independently fixated object based on a class of assumptions about general scene properties. We demonstrated the algorithm's effectiveness on simulated and real images. This algorithm could be used as the basis of a docking behaviour, whereby a mobile robot need only fixate on a point and invoke the behaviour and it will move toward the point.

As this paper is early research in an interesting area, there is more work to be done. We are currently implementing this algorithm on a mobile robot to close the loop with real images. Further, as log-polar images are multi-scale, on-line automated determination of the optimal scale for heading control would be useful. Finally, this method is based on assumptions about the scene, even if these are high-level assumptions. Image segmentation into regions that are closer than the fixation point or otherwise would allow docking without scene knowledge.

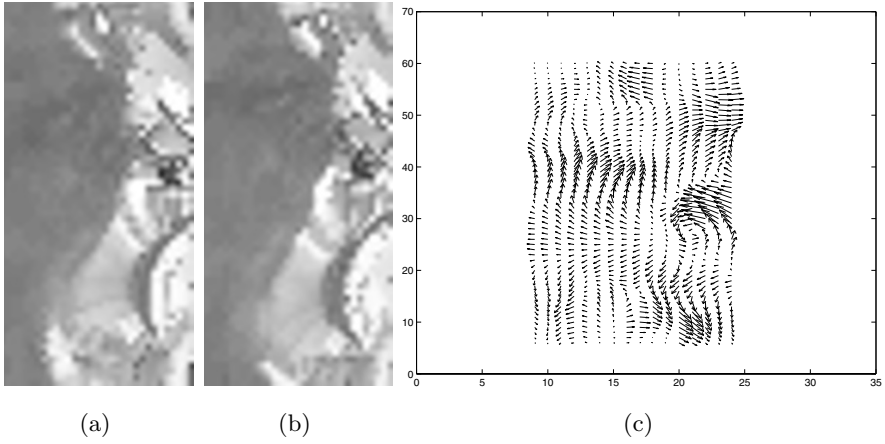


Fig. 9. Log-polar images of (c) and (d) in Figure 8, and the resulting optical flow.




Fig. 10. Mean value of γ for the image sequence of Figure 8.

Acknowledgements

The authors gratefully acknowledge Dr. Giorgio Metta for his assistance in the implementation reported in this paper. This work was supported in part by grants from The Queen's Trust for Young Australians, The Ian Potter Foundation, and an Australian Research Council (ARC) Small Grant, and in part by grants from the European Union: VIRGO and ROBVISION.

References

1. R C Arkin and D MacKenzie. Temporal coordination of perceptual algorithms for mobile robot navigation. *IEEE Trans on Robotics and Automation*, 10(3):276–286, June 1994.

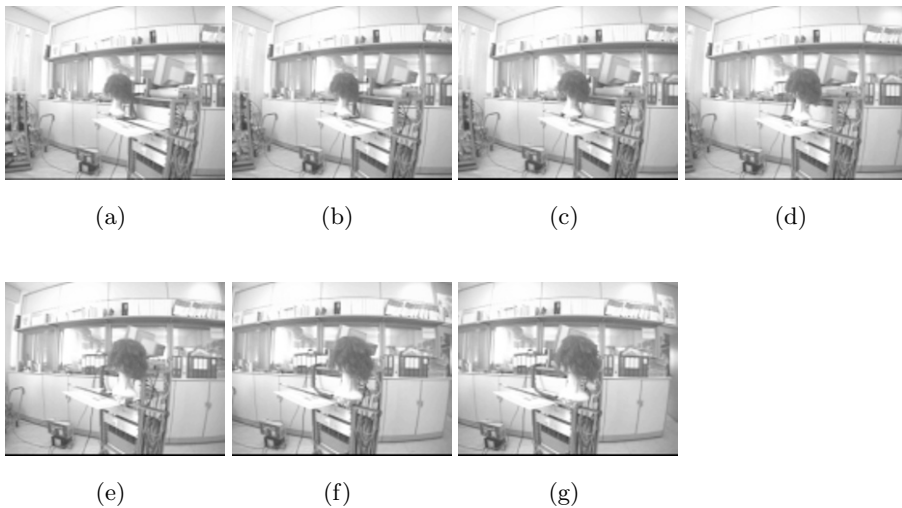


Fig. 11. The second image sequence, the motion is from left to right.




Fig. 12. Mean value of γ for the image sequence of Figure 11.

2. N Barnes and G Sandini. Active docking based on the rotational component of log-polar optic flow. In W-H Tsai and H-J Lee, editors, *ACCV Proceedings of the Asian Conference on Computer Vision*, pages 955–960, 2000.
3. N M Barnes and Z Q Liu. Vision guided circumnavigating autonomous robots. *Int. Journal of Pattern Recognition and Machine Intelligence*, 2000. in press.
4. S S Beauchemin and J L Barron. The computation of optical flow. *ACM Computing Surveys*, 27(3):433–467, 1996.
5. C Capurro, F Panerai, and E Grosso. The lira-lab head: mechanical design and control. Technical Report LIRA-TR 3/93, LIRA-Lab, DIST, Via Opera Pia, 13, University of Genova, Genova, Italy, August 1993.
6. J E Cutting, R F Wang, M Fluckiger, and B Baumberger. Human heading judgements and object-based motion information. *Vision Research*, 39:1079–1105, 1999.

7. C Fermuller. Navigational preliminaries. In *Active Perception*, pages 103–150. Lawrance Erlbaum Assoc., Hillsdale, NJ, 1993.
8. G Garibotto, S Masciangelo, M Ilic, and P Bassino. Service robotics in logistic automation: Robolift: Vision based autonomous navigation of a conventional fork-lift for pallet handling. In *8th International Conference on Advanced Robotics. Proceedings. ICAR '97*, pages 781–6, 1997.
9. A Hormann and U Rembold. Development of an advanced robot for autonomous assembly. In *IEEE Int. Conf. on Robotics and Automation*, pages 2452–7, 1991.
10. D J Fleet J L Barron and S S Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, 1994.
11. R C Jain, L Barlett, and N O'Brian. Motion stereo using egomotion complex logarithmic mapping. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, PAMI-9(2):356–369, Mar. 1987.
12. T C Leuth, U M Nassal, and U Rembold. Reliability and integrated capabilities of locomotion and manipulation for autonomous robot assembly. *Robotics and Automous Systems*, 14:185–198, 1995.
13. K Mandel and N A Duffie. On-line compensation of mobile robot docking errors. *IEEE Int. Journal of Robotics and Automation*, RA-3(6):591–598, Dec. 1987.
14. P Questa and G Sandini. Time to contact computation with a space-variant retinaline c-mos sensor. In *Proceedings of the International Conference on Intelligent Robots and Systems*, Osaka, Japan, 1996.
15. D Raviv and M Herman. A unified approach to camera fixation and vision based road following. *IEEE Trans. on Systems, Man and Cybernetics*, 24(8):1125–1141, Aug. 1994.
16. G Sandini, F Gandolfo, E Grosso, and M Tistarelli. Vision during action. In *Active Perception*, pages 151–190. Lawrance Erlbaum Assoc., Hillsdale, NJ, 1993.
17. G Sandini and M Tistarelli. Active tracking strategy for monocular depth inference over multiple frames. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(1):13–27, January 1990.
18. J Santos-Victor and G Sandini. Visual behaviours for docking. *Computer Vision and Image Understanding*, 67(3):223–28, Sept. 1997.
19. E L Schwartz. Spatial mapping in the primate sensory projection: Analytical structure and the relevance to perception. *Biological Cybernetics*, 25:181–194, 1977.
20. E L Schwartz. A quantitive model of the functional architecture of human striate cortex with application to visual illustration and cortical texture analysis. *Biological Cybernetics*, 37:63–76, 1980.
21. M Tistarelli and G Sandini. On the advantages of polar and log-polar mapping for direct estimation of time-to-impact from optical flow. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(4):401–410, April 1993.
22. H Tunley and D Young. First order optic flow from log-polar sampled images. In J-O Eklundh, editor, *European Conference on Computer Vision '94*, 1994.
23. S Uras, F Girosi, A Verri, and V Torre. A computational approach to motion perception. *Biological Cybernetics*, 60:79–87, 1988.