A Robust Docking Strategy for a Mobile Robot using Flow Field Divergence

Chris McCarthy and Nick Barnes
Vision Science, Technology and Applications Programme
National ICT Australia
Canberra, Australia.
cdmcc@rsise.anu.edu.au

***Abstract*— We present a robust strategy for docking a mobile robot in close proximity with an upright surface using optical flow field divergence. Unlike previous approaches, we achieve this without the need for explicit segmentation of the surface in the image, and using complete optical estimation (*i.e.* no affine models) in the control loop. A simple proportional control law is used to regulate the vehicle's velocity, using only the raw, unfiltered flow divergence as input. Central to the robustness of our approach is the derivation of a time-to-contact estimator that accounts for small rotations of the robot during ego-motion. We present both analytical and experimental results showing that through tracking of the focus of expansion to a looming surface, we may compensate for such rotations, thereby significantly improving the robustness of the time-to-contact estimate. This is demonstrated using an off-board natural image sequence, and in closed-loop control of a mobile robot.**

I. Introduction

For a mobile robot to interact with an object in its environment, it must be capable of docking in close proximity with the object's surface. Tasks such as plugging into a re-charging station, pallet lifting or transporting goods on a factory floor are common tasks that require some form of docking manoeuvre to be performed. Of particular importance is the control of the robot's deceleration to an eventual halt, close enough to the object for the interaction to take place. To achieve this, the robot must acquire a robust estimation of time-to-contact, and from this, control the velocity accordingly. The accuracy and robustness of the time-to-contact estimate is therefore crucial to the stability, and safety of the robot in performing this task.

For a single camera approaching an upright surface, a common method of estimating time-to-contact is to measure the image expansion induced by the apparent motion of the surface towards the camera. This can be obtained from the optical flow field divergence. This *looming effect* is characterised by flow vectors diverging from a single point in the image known as the focus of expansion (FOE). The use of visual motion to gauge time-to-contact is well supported by observations in biological vision. Srinivasan *et al.* [12] observe how honeybees use visual motion to decelerate and perform smooth graze landings. Lee [5] theorised that a human driver may visually control vehicle braking based on time-to-contact estimation obtained from image expansion.

While optical flow, and flow divergence, are commonly used to obtain time-to-contact estimates for obstacle avoidance [1], [8], few have applied optical flow to tasks requiring finer motion control such as docking. Some notable exceptions include Santos-Victor and Sandini [11], who apply an affine model of optical flow to obtain an approximation from normal flow vectors. In this work, time-to-contact is measured from the inverse of an affine flow parameter and used to control forward velocity while approaching a planar docking surface. Questa *et al.* [9] also use an affine approximation of flow, from which they measure divergence and calculate time-to-contact with a planar surface.

An important drawback of the application of affine flow models is that they require the explicit segmentation of the planar docking surface before flow is estimated. In contrast, methods for estimating general optical flow fields from local image regions, such as [6], require no *a priori* knowledge of scene structure, and therefore, no segmentation. In general, for systems such as road vehicles, optical flow is often used for other functions, such as a general sensor for salience to detect moving hazards over the whole scene, as well as for particular functions such as obstacle detection. Planar approximations are not adequate for this type of general use, and having multiple methods for calculating flow is implausible on restricted embedded hardware. In this paper we therefore develop time-to-contact estimation from general optical flow.

In much of the previous work with divergence-based time-to-contact estimation, divergence is measured at the same image location each frame. Ancona and Poggio [1], for example, use linear motion detectors to estimate time-to-contact at locations symmetrically placed about the image centre. Such strategies ignore the effect of FOE shifts on the divergence measure across the image. Robot egomotion is rarely precise, and even where only translation is intended, rotations will be present due to steering control adjustments, differing motor outputs, bumps, and noisy flow estimates. Such influences introduce instantaneous, frame-to-frame rotations of the robot, causing the optical axis to move with respect to the predominant direction of motion. As a result, the FOE is unlikely to be fixed with respect to the image centre. Therefore, to ensure consistency in time-to-contact estimates over time, we argue that divergence should be measured with respect to the FOE, and not the image centre.

Robustly estimating time-to-contact when the optical and translation axes are not physically aligned has been exam-

ined previously. Subbarao [13] considers time-to-contact with surfaces of arbitrary orientation, for a camera of arbitrary alignment with respect to the direction of motion. Subbarao shows that for non fronto-parallel surfaces, time-to-contact cannot be precisely computed, however, from the image deformation parameters, an upper and lower bound on the time-to-contact can be obtained. Subbarao, however, does not consider the effects of instantaneous rotations during egomotion, and therefore assumes the point of interest lies along the camera's optical axis. While a fixation-based strategy such as that used by Questa *et al.* [9] can keep the target point centred, a mobile robot is unable to achieve this without additional hardware support, which is not always available.

An alternative approach is to account for instantaneous rotations of the optical axis by tracking the location of the FOE for each frame. Van Leeuwen and Groen [14], [15] consider the use of FOE tracking to correct for the physical misalignment of the optical and translational axes as a result of the camera-robot configuration. However, while accounting for the constant, physical miss-alignment of these axes, they do not extend the use of FOE tracking explicitly to the removal of small frame-to-frame rotational effects during ego-motion, nor do they apply the time-to-contact measure to directly control the vehicle's velocity. In general, while previous work such as this has considered the use of FOE tracking for camera stabilisation during egomotion, no one has applied such an approach to tasks requiring fine motion control, nor provided a theoretical analysis supporting the advantages of such a strategy, and its potential use for control.

In this paper, we present a robust strategy for docking a mobile robot in close proximity with an upright surface using optical flow field divergence. We provide a theoretical justification for the constant tracking of the FOE as a means of accounting for shifts of the optical axis due to instantaneous rotations during ego-motion. We present off-board and on-board experiments demonstrating the application of this strategy to the task of docking. This strategy requires no explicit segmentation of the surface, uses complete gradient-based optical estimation, and employs only a simple proportional control law to regulate the robot's velocity.

II. Theoretical Background

Flow divergence is measured by examining the partial spatial derivatives of image velocity components in orthogonal directions. It is commonly defined as:

$$D = u_x + v_y, \tag{1}$$

where, for a given point in the image, u_x and v_y are the partial derivatives of their respective components of flow, in the x and y directions. Time-to-contact (τ) to a point along the optical axis of the camera can be measured from flow divergence, and is commonly defined as [2]:

$$\tau = \frac{Z}{T_r} = \frac{2}{D}, \tag{2}$$

where Z is the distance to the object along the axis of translation, and T_r is the velocity.

Flow divergence is invariant across the image plane if the surface plane is perpendicular to the camera's optical axis, and under this assumption, can be measured anywhere in the image. If precise fronto-parallel alignment with the surface is not maintained, further image deformation (curl [13]) occurs, causing the measured divergence to vary across the image.

Given instantaneous rotations during ego-motion, precise surface alignment is unlikely to exist. In the image domain, such effects are characterised by frame-to-frame shifts of the FOE, causing the divergence at any given image location to vary. As a result, (2) is unlikely to provide an accurate time-to-contact estimate in the presence of such rotations. To improve time-to-contact estimates during robot egomotion, a means of accounting for rotational effects is essential.

A. Compensating for rotation in time-to-contact estimation

The analysis presented in this section extends on the geometric modelling used by Santos-Victor and Sandini [11]. As in [11], we represent the docking surface as a plane in a camera centred coordinate system:

$$Z(X,Y) = Z_0 + aX + bY, \tag{3}$$

where Z_0 is the distance to the surface along the optical axis, (X,Y) is a point on the surface, and a and b give the slant and tilt with respect to the optical axis. By introducing the perspective projection equations, the surface plane can also be expressed as a function of the image coordinates, (x,y) [10]:

$$Z(x,y) = \frac{Z_0}{1 - a\frac{x}{f_x} - b\frac{y}{f_y}}. \tag{4}$$

Given a fixed camera with respect to the robot's direction of motion, we represent the translational velocity of the camera, T_c, as proportions of the predominant forward translational velocity, T_r, of the robot:

$$T_c = \begin{bmatrix} \alpha T_r & \beta T_r & \gamma T_r \end{bmatrix}. \tag{5}$$

The camera's angular velocity (ω_c) is given by:

$$\omega_c = \begin{bmatrix} \omega_x & \omega_y & \omega_z \end{bmatrix}, \tag{6}$$

where each component represents rotation about the axis indicated by its subscript. Figure 1 shows the geometric setup.

The optical flow induced by the apparent motion of the docking plane is defined by the well known equations [11]:

$$u(x,y) = f_x \left[\frac{\gamma T_r(\frac{x}{f_x} - \alpha)}{Z(x,y)} + \omega_x \frac{xy}{f_x f_y} - \omega_y(1 + \frac{x^2}{f_x^2}) + \omega_z \frac{y}{f_y} \right], \tag{7}$$

$$v(x,y) = f_y \left[\frac{\gamma T_r(\frac{y}{f_y} - \beta)}{Z(x,y)} + \omega_x(1 + \frac{y^2}{f_y^2}) - \omega_y \frac{xy}{f_x f_y} - \omega_z \frac{x}{f_x} \right], \tag{8}$$

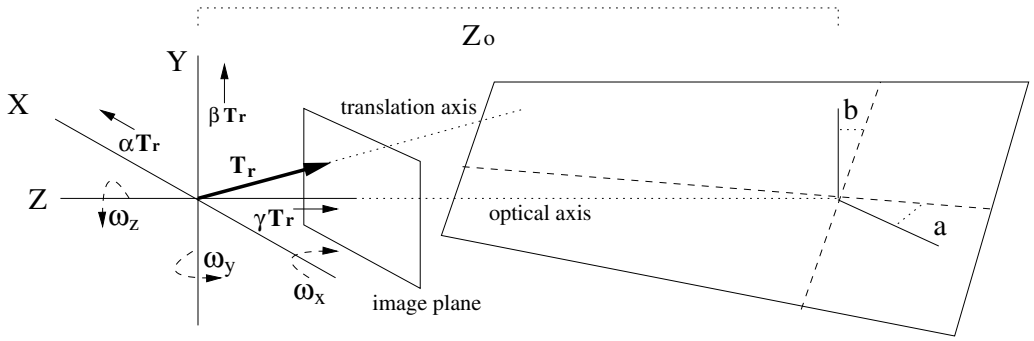

Fig. 1. Geometric setup

where $u(x,y)$ and $v(x,y)$ are the horizontal and vertical components of motion, and f_x and f_y are focal lengths expressed in pixels.

Let us now consider the effects of rotation, causing the FOE to shift. Let (x',y') be an arbitrary point in the image representing the FOE. We define the depth of the surface, $Z(x,y)$, with respect to the FOE:

$$Z(x,y) = \frac{Z(x',y')}{1 - a\frac{(x-x')}{f_x} - b\frac{(y-y')}{f_y}}. \tag{9}$$

Substituting (9) into Equations (7) and (8), we obtain:

$$u(x,y) = \frac{\gamma T_r(x - f_x\alpha)}{Z(x',y')}\left[1 - \frac{a(x-x')}{f_x} - \frac{b(y-y')}{f_y}\right] + \omega_x\frac{xy}{f_y} - \omega_y(f_x + \frac{x^2}{f_x}) + \omega_z\frac{y}{f_x}, \tag{10}$$

$$v(x,y) = \frac{\gamma T_r(y - f_y\beta)}{Z(x',y')}\left[1 - \frac{a(x-x')}{f_x} - \frac{b(y-y')}{f_y}\right] + \omega_x(f_y + \frac{y^2}{f_y}) - \omega_y\frac{xy}{f_x} - \omega_z\frac{x}{f_x}. \tag{11}$$

Given the optical flow at the FOE is zero, we can substitute for $x = x'$ and $y = y'$ in (10) and (11), and set their left sides both to zero. Solving then for ω_x and ω_y, we obtain:

$$\omega_x = \frac{f_y}{x'y'}\left(\frac{\gamma T_r}{Z(x',y')}(x' - f_x\alpha) + \omega_y(f_x + \frac{x'^2}{f_x}) + q\right), \tag{12}$$

$$\omega_y = \frac{\frac{T_r}{Z(x',y')}(x'y'\beta + f_yx' + f_xf_y\alpha + \frac{f_x\alpha y'^2}{f_y}) - r}{f_xf_y(1 + \frac{x'^2}{f_x^2} + \frac{y'^2}{f_y^2})}, \tag{13}$$

where

$$q = \omega_z\frac{y'}{f_y}$$
$$r = \omega_z(y' + \frac{y'^3}{f_y^2} - \frac{x^2y'}{f_xf_y}).$$

Taking the partial derivatives of $u(x,y)$ and $v(x,y)$ in their respective directions, and again substituting for $x = x'$, $y = y'$, we obtain the partial derivatives at the FOE, defined as:

$$u_x\Big|_{foe} = \frac{\gamma T_r}{Z(x',y')}\left[1 - a\left(\frac{x'}{f_x} + \alpha\right)\right] + \omega_x\frac{y'}{f_y} - \omega_y\frac{2x'}{f_x}, \tag{14}$$

$$v_y\Big|_{foe} = \frac{\gamma T_r}{Z(x',y')}\left[1 - b\left(\frac{y'}{f_y} + \beta\right)\right] + \omega_x\frac{2y'}{f_y} - \omega_y\frac{x'}{f_x}. \tag{15}$$

Summing these, we obtain the flow field divergence at the FOE (D_f):

$$D_f = \frac{-\gamma T_r}{Z(x',y')}\left[a\left(\frac{x'}{f_x} + \alpha\right) + b\left(\frac{y'}{f_y} + \beta\right) - 2\right] + 3\left(\frac{\omega_xy'}{f_y} - \frac{\omega_yx'}{f_x}\right), \tag{16}$$

and from this, we derive an equation for the time-to-contact of the surface point projecting to the FOE ($\tau_f = \frac{Z(x',y')}{T_r}$):

$$\tau_f = -\frac{\gamma}{D_f}\left[a\left(\frac{x'}{f_x} + \alpha\right) + b\left(\frac{y'}{f_y} + \beta\right) - 2\right] + \frac{3Z(x',y')}{D_fT_r}\left(\frac{\omega_xy'}{f_y} - \frac{\omega_yx'}{f_x}\right). \tag{17}$$

Using Equations (12) and (13), we substitute for ω_x and ω_y in (17) and thus remove both rotations from the equation:

$$\tau_f = -\frac{\gamma}{D_f}\Bigg[1 + a\left(\frac{x'}{f_x} + \alpha\right) + b\left(\frac{y'}{f_y} + \beta\right) - \frac{3}{\gamma x'}\Bigg(-f_x\alpha + \frac{(x'f_y + f_xf_y\alpha + x'y'\beta + \frac{y'^2f_x\alpha}{f_y})}{f_y(1 + \frac{x'^2}{f_x^2} + \frac{y'^2}{f_y^2})} + \frac{\omega_zy'T_r}{f_yZ(x',y')}\Big(f_y + \frac{y'^2}{f_y} - \frac{x'^2}{f_x} - 1\Big)\Bigg)\Bigg]. \tag{18}$$

Notably, this substitution introduces a term involving camera roll (ω_z). If required, techniques for roll removal such as [4] can also be applied without prior knowledge of the rotation.

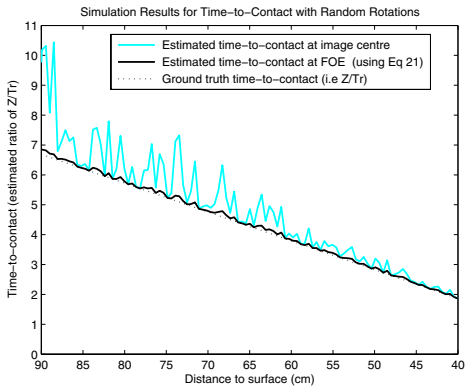

Fig. 2. Simulation results compare our FOE-based time-to-contact estimator (Equation (19)), with time-to-contact estimates obtained at the image centre using (2). For each sample, the robot's forward speed, and randomly chosen instantaneous rotational velocity ($-0.1 \leq \omega_y \leq 0.1$) were used to compute the corresponding horizontal flow. Ground truth gives exact time-to-contact, given the robot's velocity towards, and distance from the surface.

B. Time-to-contact for a ground-based mobile robot

Consider Equation (18) for the case of a mobile robot, moving on a ground plane towards a visible planar surface. Given a fixed, approximately forward facing camera, ω_z, will be negligible, and can therefore be set to zero. In addition, the camera orientation parameters: α, β and γ, can also be set to known values ($\alpha = \beta = 0, \gamma = 1$), thus reducing (18) to:

$$\tau_f = -\frac{1}{D_f}\left[1 + \frac{ax'}{f_x} + \frac{by'}{f_y} - \frac{3}{(\frac{x'^2}{f_x^2} + \frac{y'^2}{f_y^2} + 1)}\right]. \qquad (19)$$

The only potential remaining unknowns are the surface orientation parameters: a and b. If directional control maintains an approximate angle of approach, then bounds on these parameters, and hence τ_f, should exist. Even without precise knowledge of alignment, the FOE should still provide the best location from which to obtain the time-to-contact estimate given rotation is accounted for at this point. Notably, the existence of the FOE itself provides a natural constraint on the angle of approach. At extreme angles of approach, the FOE is unlikely to exist at all as the surface distance along the axis of motion becomes infinite. Keeping the FOE within the projected surface target area significantly narrows the range of allowable approach angles.

C. Simulation Results

To test the theory, a simulation was conducted, modelling the motion of a ground-based mobile robot, with camera, towards a planar, fronto-parallel surface. A 2D motion model was used, allowing only forward velocity and a single rotation in the ground plane. From this, a set of sample flow fields were obtained, and time-to-contact estimated using (19). For each consecutive sample, the distance to the surface was decremented by a constant amount, before a translational velocity, and randomly selected instantaneous rotational velocity were applied to the scene. The resulting motion vectors were then projected to the image plane, generating the expected flow resulting from the motion. The location of the FOE was then estimated, and time-to-contact computed.

Figure 2 gives the simulation results, showing a direct comparison of time-to-contact estimates obtained using the FOE-based time-to-contact estimator defined by (19), and time-to-contact estimates obtained from the measured divergence at the image centre (using (2)). A ground truth time-to-contact is also provided. This was computed from the robot's distance from the surface, and its known constant forward velocity towards the surface. It can be seen that the FOE-based time-to-contact measure almost perfectly reflects ground truth. Small discrepancies between the FOE-based measure and ground truth are the result of unavoidable quantisation errors in the image, disallowing the precise location of the FOE. Given infinite resolution, this discrepancy would disappear.

In contrast, time-to-contact measured along the optical axis exhibits significant fluctuation compared with that obtained at the FOE. Furthermore, the image centre always provides an over estimate of the time-to-contact. Notably, errors in time-to-contact are reduced as the distance to the surface approaches zero. This, however, is due to the robot's constant velocity towards the surface. As the translational flow increases, the effects of the robot's rotation are diminished. When the robot's velocity is decreasing in response to a looming surface, the translational flow will remain approximately constant, and so these fluctuations in time-to-contact at the image centre, can be expected throughout the approach.

In general, the derivation and simulation results presented in this section show that for a mobile robot with a single camera of known orientation with respect to the robot's axis of motion, rotations can be accounted for by estimating time-to-contact at the FOE. Conversely, estimating time-to-contact with no consideration for shifts of the FOE due to such rotations will result in inaccurate and unreliable time-to-contact estimates.

III. Implementation and Results

We now present two experiments showing the performance of the proposed FOE-based strategy for estimating time-to-contact from flow field divergence. We present results from an off-board experiment using a purpose-built image sequence, and results of the technique's integration into the control loop of a mobile robot for docking with an object surface.

In both experiments, the FOE was calculated using a simple algorithm that assumes the imaged surface occupies the entire viewing field. To obtain x', each row in the image was used to count the number of positive and negative horizontal flow components, which were then differenced, and averaged over all rows to locate the overall zero point for X. The algorithm was applied similarly to obtain y', using the signs of vertical components of flow. While more sophisticated algorithms exist, this was deemed suitable for these experiments.

A. Off-board time-to-contact estimation

A *looming wall sequence*, shown in Figure 3, was constructed to simulate the image expansion experienced when

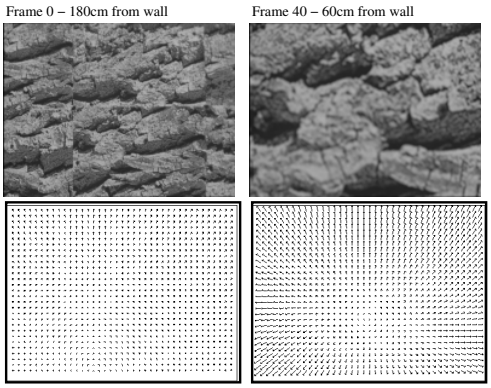

Fig. 3. Sample frames and ground truth optical flow fields from the looming wall image sequence.

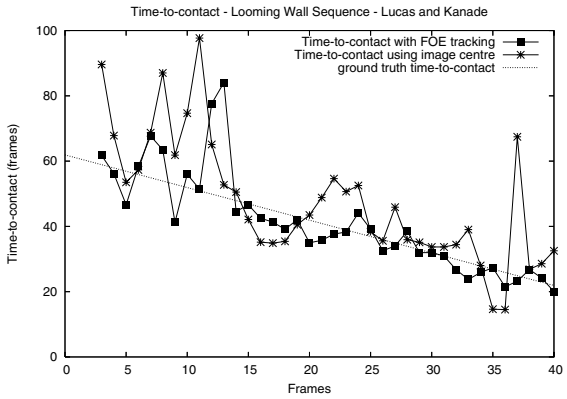

Fig. 4. Time-to-contact estimates for looming wall image sequence.

Fig. 5. Setup for on-board docking tests.

approaching a surface. The camera was moved 3cm per frame towards a heavily textured, approximately fronto-parallel wall. The camera was tilted slightly upwards while moving along the ground plane (to avoid viewing the ground plane).

Time-to-contact was estimated using four patches in the image, each centred on a distance of 12 pixels from the FOE, and each at 45 degrees from the horizontal and vertical axes that intersect at the FOE. For comparison, time-to-contact was also estimated by placing the four patches about the image centre.

Figure 4 shows time-to-contact estimates for each frame of the sequence (where time is measured in frames) for both strategies. A ground truth for the time-to-contact is also shown. Due to the unavailability of precise distance measures between the camera and surface, ground truth was obtained from the camera's known velocity, and a linear fit over time-to-contact measures obtained from ground truth flow fields.

From these results, a significant improvement in the consistency of time-to-contact estimates is achieved when divergence is calculated with respect to the FOE. Of particular note, the FOE-based time-to-contact strategy achieves a close match with ground truth from the fifteenth frame onward. In contrast, the image centre-based method consistently over-estimates time-to-contact, as well as exhibiting larger fluctuations than the FOE-based strategy across the sequence.

B. Robot docking experiment

The FOE-based time-to-contact measure was integrated into a simple closed-loop docking behaviour for velocity control of a mobile robot. In experiments conducted, a robot with a single forward facing camera approached a heavily textured, roughly fronto-parallel wall, attempting to safely stop as close to the wall as possible without collision (see Figure 5).

Velocity was governed by the proportional control law:

$$v_t = v_{t-1} + K_p(Dref - Dt), \tag{20}$$

where v_t is the current velocity, K_p is a proportional gain (tuned to 0.0275), D_t is the current flow divergence and D_{ref} is a reference divergence to servo to. In order to maintain a constant divergence during approach, the forward velocity of the robot must necessarily decrease.

Flow divergence was estimated using two 40×40 pixel image patches, each placed at 45 degrees on either side of the vertical axis passing through the FOE, centred on a distance of 25 pixels from the FOE. The patches were placed only above the FOE to avoid measuring divergence on the imaged ground plane. Optical flow was obtained using Lucas and Kanade's [6] gradient-based method, in combination with Fleet and Langley's recursive causal temporal filter [3], chosen on the basis of strong performances in a recent comparison of optical flow techniques for robot navigation tasks [7].

In each trial, the robot began approximately 1.5 metres from the wall at a speed of $0.4m/s$. At 1 metre from the wall, the docking behaviour was invoked manually. No directional control was used, however the robot was subject to small lateral drift and rotations. Using a calibrated overhead camera, the robot's course was tracked, and velocity updates(v_t) were logged.

Figures 6 and 7 show velocity-distance profiles, and the robot's plotted approach for four trials. Note that the plotted position of the robot is given with respect to the front of the robot. In all four trials, the FOE-based strategy docked in close proximity to the surface without collision.

The results show that close proximity stopping distances were achieved with surprisingly high consistency. The average stopping distance from the wall was 6cm, with the furtherest

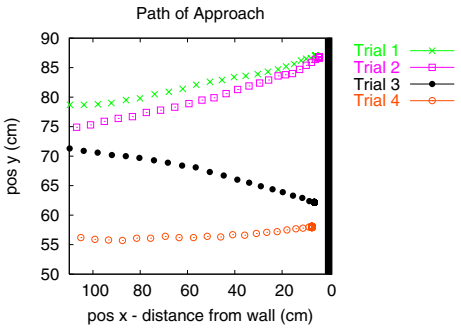

Fig. 6. Plotted position of robot for each docking trial.

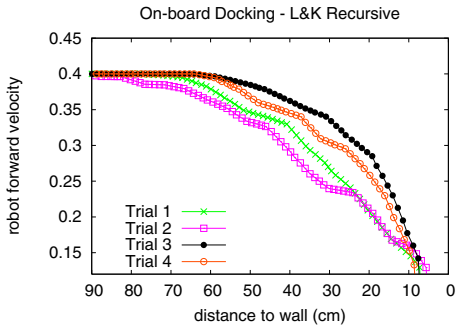

Fig. 7. On-board velocity-distance profiles for docking trials.

recorded being just 7cm. This consistency in stopping distance is encouraging when considering the simple control law used. This is also despite considerable variation in both the robot's initial starting position, and the extent (and direction) of its lateral drift in each trial, as is apparent in Figure 6.

An attempt was made to compare the FOE-based on-board control scheme with the same control scheme using an image-centre based divergence measure. The raw divergence estimates obtained at the image-centre, however, were found to be unworkable for the simple proportional control scheme used. A large range of tuning parameter values were explored.

The FOE-based docking strategy compares well with previous work in flow-based docking. The final stopping distances achieved are a significant improvement on Questa *et al.* [9] (approximately 15cm), and comparable with Santos-Victor and Sandini [11]. It is important to note that unlike previous work, we report highly consistent results over a set of trials. In addition, we obtain these results without the use of planar surface models, or filtering of the divergence estimates. Furthermore, only a a simple proportional control scheme is employed, with minimal tuning requirements. Additional post-filtering and/or a more sophisticated control scheme would be expected to improve performance further.

IV. Conclusion

This paper has presented a mobile robot docking strategy that utilises a time-to-contact estimation that is robust to noisy, instantaneous rotations induced by robot ego-motion. We have shown that through tracking the focus of expansion in the optical flow field, small rotations of the camera and miss-alignments of the optical and translational axes can be accounted for by calculating flow divergence with respect to the FOE. In this way, the effects of the rotation are effectively cancelled out, and improved accuracy and stability is achieved. Our results show a significant improvement in time-to-contact estimates when compared with common strategies that take no account of the shifting FOE during robot ego-motion. The accuracy and stability achieved using the FOE-based time-to-contact estimation strategy was demonstrated to be sufficient for fine motion control of a mobile robot for the task of docking. This was without any filtering of the control input signal, or complex control scheme.

References

[1] N. Ancona and T. Poggio, "Optical flow from 1d correlation: Application to a simple time-to-crash detector," in *Proceedings of the International Conference on Computer Vision*, 1993, pp. 209–14.

[2] D. Coombs, M. Herman, T. Hong, and M. Nashman, "Real-time obstacle avoidance using central flow divergence,and peripheral flow," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 1, pp. 49–59, 1998.

[3] D. J. Fleet and K. Langley, "Recursive filters for optical flow." *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 1, pp. 61–7, 1995.

[4] M. Hanada and Y. Ejima, "A model of human heading judgement in forward motion," *Vision Research*, vol. 40, no. 2, pp. 243–63, Jan. 2000.

[5] D. N. Lee, "A theory of visual control of braking based on information about time to collision," *Perception*, vol. 5, no. 4, pp. 437–59, 1976.

[6] B. Lucas and T. Kanade, "An iterative image registration technique with an application to stereo vision." in *Proceedings of DARPA Image Understanding Workshop*, 1984, pp. 121–30.

[7] C. McCarthy and N. Barnes, "Performance of optical flow techniques for indoor navigation with a mobile robot," in *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, 2004, pp. 5093–8.

[8] R. C. Nelson and J. Y. Alloimonos, "Obstacle avoidance using flow field divergence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 10, pp. 1102–6, 1989.

[9] P. Questa, E. Grossmann, and G. Sandini, "Camera self orientation and docking maneuver using normal flow," in *Proceedings of Spie - the International Society for Optical Engineering*, vol. 2488, Orlando, USA, 1995, pp. 274–83.

[10] J. Santos-Victor and G. Sandini, "Visual behaviors for docking, Tech. Rep. LIRA-TR 2/94, June 1994.

[11] ——, "Visual behaviors for docking," *Computer Vision and Image Understanding: CVIU*, vol. 67, no. 3, pp. 223–38, 1997.

[12] M. V. Srinivasan, S. W. Zhang, J. S. Chahl, E. Barth, and S. Venkatesh, "How honeybees make grazing landings on flat surfaces," *Biological Cybernetics*, vol. 83, pp. 171–83, 2000.

[13] M. Subbarao, "Bounds on time-to-collision and rotational component from first-order derivatives of image flow," in *Computer Vision, Graphics and Image Processing*, vol. 50, 1990, pp. 329 – 41.

[14] M. B. van Leeuwen and F. C. A. Groen, "Motion estimation wth a mobile camera for traffic applications," in *Proceedings of the IEEE Intelligent Vehicles Symposium 2000*, 2000, pp. 58–63.

[15] ——, "Motion interpretation for in-car vision systems," in *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2002, pp. 135–40.