# 3D Motion and Structure from Planar Passive Navigation

by
Enrico De Micheli & Alessandro Verri

Technical Report 92-22
October 1992

Department of Computer Science
The University of British Columbia
Vancouver, B. C. V6T 1Z2
Canada

email: DEMICHELI@genova.infn.it

# 3D Motion and Structure from Planar Passive Navigation

Enrico De Micheli (*) & Alessandro Verri (†)

Technical report 92-22
October 1992

(*) Department of Computer Science - U.B.C., Vancouver, B.C., V6T 1W5 Canada.[1]
(†) International Computer Science Institute, 1947 Center Street, Berkeley CA 94704, USA.[2]

## Abstract

The work presented in this report is based on the observation that if a viewing camera is appropriately mounted on a vehicle which moves on a planar horizontal surface parallel to the instantaneous direction of translation, then the optical flow obtained from the moving camera depends on two parameters only; the angular velocity around an axis orthogonal to the planar surface and the ratio between the viewed depth and the translational speed (i.e., *generalized time-to-collision*). Elementary error analysis shows that the angular velocity can be robustly estimated by averaging the horizontal component of the optical flow along the vertical line through the center of the image. Once the angular velocity has been recovered, the generalized time-to-collision, which measures depth in time units, can be computed from one component only of the optical flow. It is found that the accuracy of depth estimation from the vertical component of the optical flow is more accurate, increases with the distance from the horizontal line through the center of the image, and is almost independent of the angular velocity. Therefore, in the case of planar motion, the computation of the two-dimensional (2D) optical flow over the entire 2D image plane is hardly necessary. A number of experiments on synthetic and real images support the presented analysis. Since the precision with which the viewing camera is positioned on the moving vehicle is not critical, it is concluded that the proposed method is likely to be very useful in applications like autonomous robot navigation.

[1]On leave from Istituto di Cibernetica e Biofisica – CNR, Via Dodecaneso 33, Genova, Italy.
[2]On leave from Dipartimento di Fisica, Università di Genova, Via Dodecaneso 33, Genova, Italy.

# 1 Introduction

In many dynamical applications, like autonomous robot navigation and car driving, useful visual information can be obtained from a camera mounted on the mobile vehicle. The vehicle motion throughout the surrounding environment produces spatial and temporal changes in the viewed image which can be used for the reconstruction of both the vehicle motion and the three-dimensional (3D) structure of the scene in various ways [1-10]. Many of the proposed techniques consist of two steps. In the first step, the optical flow, that is, the apparent motion of the image brightness pattern on the image plane, is computed. In the second step, the equations which relate optical flow to the viewed motion and structure are solved. Both of these steps are computationally rather expensive and their implementation in real systems seems to require the use of special hardware.

The key point of the paper is the observation that, if the motion of the vehicle is restricted on a planar surface, the optical flow equations can be drastically simplified. It is first shown that if the image plane of the camera mounted on the vehicle is orthogonal to the planar surface and the optical axis is parallel to the instantaneous direction of translation, then the angular velocity is the only motion parameter which is left to be computed. The optical flow equations become linear and the viewed motion and structure can be easily reconstructed. An elementary error analysis shows that the angular velocity can be optimally recovered from only the horizontal component of optical flow along the vertical line which goes through the optical center of the image plane. Once the angular velocity has been determined, depth can be estimated by means of one component only of the optical flow. The vertical component of the optical flow is found to be better suited than the horizontal component. Since one-dimensional (1D) optical flow computed over lines of the image plane provides reliable information on the viewed motion and structure, the computation of the full 2D optical flow over the entire image plane is unnecessary and conventional hardware may be sufficient to develop a simple optical flow based module for planar passive navigation.

The report is organized as follows. In section 2 the geometry of the problem is described. Section 3 discusses how to compute motion and structure parameters from the simplified optical flow equations. In Section 4 experimental results on synthetic data and real images are presented. Finally, Section 5 summarizes the obtained results.

# 2 Planar navigation

In this section the general problem of passive navigation is restricted to the case in which the viewing system moves on a planar surface. No assumption is made on the spatial structure of the viewed scene, and the motion is otherwise arbitrary.

Let us first establish some basic notations. Let $S$ denote the viewing system which

moves throughout a static environment and $(X, Y, Z)$ the coordinates of a point $P$ in a system of orthogonal axes X, Y and Z attached to $S$. Let $O$, the origin of the system, be the center of projection and $f$ the focal length. If the normal vector of the image plane is parallel to the Z-axis, then $x = fX/Z$, $y = fY/Z$, and $z = f$ are the coordinates of $p$, the perspective projection of $P$ in the image plane with respect to $O$.

Ideally, the optical flow at $p$ can be thought of as the perspective projection onto the image plane of the $3D$ velocity field $\mathbf{V}$ at $P$. In the case of arbitrary rigid motion, $\mathbf{V}$ can be described in terms of $\mathbf{T} = (T_X, T_Y, T_Z)$ and $\mathbf{\Omega} = (\Omega_X, \Omega_Y, \Omega_Z)$, the translational and angular velocity between $S$ and the environment respectively. The components $u$ and $v$ of the optical flow along the X- and Y-axis can then be written as [1]

$$\begin{aligned} u &= \frac{T_X f - T_Z x}{Z} + \frac{(x^2 + f^2)\Omega_Y - y(x\Omega_X + f\Omega_Z)}{f} \\ v &= \frac{T_Y f - T_Z y}{Z} - \frac{(y^2 + f^2)\Omega_X - x(y\Omega_Y + f\Omega_Z)}{f} \end{aligned} \tag{1}$$

In this general setting, the problem of recovering the viewed motion and structure from optical flow at $n$ points can be stated as the problem of solving $n$ pairs of equations like Eqs. 1 in $6 + n$ unknowns, that is, the six motion parameters, $T_X$, $T_Y$, $T_Z$, $\Omega_X$, $\Omega_Y$, and $\Omega_Z$, and the depth $Z$ at each of the $n$ points. This problem is clearly nonlinear and, independent of the employed method, reliable solutions seem to require the integration of optical flow estimates from rather different locations [10,11].

If the motion of $S$ is restricted to a planar surface $\pi$ the optical flow equations can be greatly simplified. The hypothesis of planar motion is clearly relevant to many natural and cultural scenarios. Two further assumptions will be needed: *(i)* the image plane of the viewing system is orthogonal to $\pi$, and *(ii)* the optical axis is parallel to the instantaneous direction of translation.

Before commenting on the meaning of *(i)* and *(ii)*, let us first look at their impact on the optical flow equations. Since the motion is constrained to the planar surface $\pi$, the angular velocity $\mathbf{\Omega}$ is completely described by a vector in the direction orthogonal to $\pi$. By virtue of *(i)*, this direction is parallel to the Y-axis and $\mathbf{\Omega}$ can be simply written as

$$\mathbf{\Omega} = (0, \Omega, 0). \tag{2}$$

The translational component of motion, $\mathbf{T}$, by means of *(ii)*, lies entirely along the Z-axis, or

$$\mathbf{T} = (0, 0, T). \tag{3}$$

The substitution of the expressions 2 and 3 into Eqs. 1, if $\tau = -Z/T$, yields

$$
\begin{aligned}
u &= \Omega f + x\left(\frac{\Omega}{f}x + \frac{1}{\tau}\right) \\
v &= y\left(\frac{\Omega}{f}x + \frac{1}{\tau}\right)
\end{aligned}
\tag{4}
$$

which, in matrix notation, can be written as

$$
\begin{pmatrix} u \\ v \end{pmatrix} = M \begin{pmatrix} \Omega \\ 1/\tau \end{pmatrix}
$$

with

$$
M = \begin{pmatrix} f + x^2/f & x \\ xy/f & y \end{pmatrix}
$$

Eqs. 4 must be compared with Eqs. 1. In the case of Eqs. 1, the problem was nonlinear and consisted of $2n$ equations in $6 + n$ unknowns. Instead, in the case of Eqs. 4, the problem is linear and consists of 2 equations in 2 unknowns, $\Omega$ and $1/\tau$. Note that $\tau$, the time required to cover the distance $Z$ at the speed $|T|$, reflects the well known ambiguity between absolute distance and speed of a viewed point and can be thought of as the depth $Z$ in time units. In what follows we will refer to $\tau$ as the *generalized time-to-collision*. Since the translational parameter is hidden in the denominator of $\tau$, $\Omega$ remains the only motion parameter to be computed.

The analysis of the existence and uniqueness of the solution to Eqs. 4 is elementary. With the exception of the points for which $|M|$, the determinant of $M$, vanishes, Eqs. 4 can always be solved to uniquely determine the viewed motion and structure. The solution takes the form

$$
\Omega = \frac{uy - vx}{fy} \tag{5}
$$

and

$$
\frac{1}{\tau} = \frac{f^2 v - x(uy - vx)}{f^2 y}. \tag{6}
$$

Since $|M| = fy$, it follows that Eqs. 4 cannot be solved at the points with $y = 0$, that is, on the horizontal line $l_h$ which goes through the optical center.

Finally, let us comment on the constraints *(i)* and *(ii)* which, in addition to the hypothesis of planar motion, underlie the presented analysis. First the constraints *(i)* and *(ii)* can be thought of as defining a reference frame in which the problems of motion and structure recovery in the case of planar motion can be uncoupled and stated with a minimal number of unknowns. Interestingly, the constraints *(i)* and *(ii)* are naturally met by a viewing camera mounted in the standard upright position on a mobile vehicle. Independent of the trajectory of the vehicle, the image plane of

the camera is orthogonal to the ground plane and the direction of gaze is parallel to the instantaneous direction of translation. Let us now discuss in some detail the use of Eqs. 5 and 6 for the computation of motion and structure.

# 3 1D optical flow

In this Section the stability of the computation of the viewed motion and structure from Eqs. 5 and 6 is studied. In agreement with the simple camera model described in section 2, the errors on the horizontal and vertical optical flow components are assumed to be equal and constant over the image plane. Let us first consider the estimation of $\Omega$ from Eq. 5.

## 3.1 Angular velocity

By differentiating Eq. 5, if $\Delta u$ and $\Delta v$ are the errors in the computation of $u$ and $v$, the error $\Delta\Omega$ on the angular velocity can be written as

$$\Delta\Omega = \frac{1}{f}\left(\Delta u + \left|\frac{x}{y}\right|\Delta v\right) \tag{7}$$

From Eq. 7 it can be concluded that $\Delta\Omega$ is minimum at the points with $x = 0$, that is, the points which lie on the vertical line $l_v$ through the center of projection. For $x = 0$, Eq. 5 reads

$$\Omega = \frac{u}{f} \tag{8}$$

from which it follows that $u$ is constant over $l_v$. From Eq. 8 it can easily be seen that in order to optimally estimate the angular velocity it is sufficient *(a)* to compute the horizontal component of the optical flow along the vertical line $l_v$ and *(b)* to estimate $\Omega$ as

$$\Omega = \frac{< u >}{f} \tag{9}$$

where $< u >$ is the average of $u$ along $l_v$. It is remarkable that the very simple Eq. 9 is true for arbitrary (planar) motion and independent of the 3D structure of the scene.

## 3.2 3D structure

Let us now look at the estimate of the 3D structure of the viewed scene by means of Eq.6. By substituting the expression for $\Omega$ of Eq.5 into Eq.6, the equation for the generalized time-to-collision $\tau$ reads

$$\frac{1}{\tau} = \frac{v}{y} - \frac{x}{f}\Omega \tag{10}$$

Therefore, once the angular velocity has been recovered, the 3D structure of the scene can be reconstructed by computing only the vertical component of optical flow. As Eq. 10 indicates, for $\Omega = 0$, all the structure information on the viewed scene is in the ratio $v/y$. In the presence of rotation, $v$ changes an amount equal to $xy\Omega/f$ (see Eq.4), and the second term in the right-hand-side of Eq. 10 compensates exactly for that change. For most typical values of $f$ and $T \neq 0$, the term $|x\Omega/f|$ is usually much smaller than $|v/y|$. This agrees with the intuitive fact that a rotation around a vertical axis has a small (and purely perspective) effect on the vertical component of the apparent motion.

The stability analysis of Eq. 10 is straightforward. By differentiating Eq. 10, the error on the inverse of the estimated generalized time-to-collision, $\Delta(1/\tau)$, is

$$\Delta\left(\frac{1}{\tau}\right) = \left|\frac{1}{y}\right| \Delta v + \left|\frac{x}{f}\right| \Delta\Omega \tag{11}$$

Thus, the larger the distance from the horizontal line $l_h$ (i.e., the larger $y$), the smaller the error on the depth computed from Eq. 10. From Eq. 11 it follows that the error increases with $|x|$, but since $|v/y|$ is usually much larger than $|x\Omega/f|$ this dependence should be hardly noticeable. In section 4 these predictions will be checked on synthetic and real data.

## 3.3 Symmetry breaking

The previous analysis has shown that, in the case of planar navigation, the two components of the optical flow and the horizontal and vertical directions of the image plane are not symmetric. This somewhat surprising result has a simple explanation. Clearly, the a priori symmetry between the horizontal and vertical direction is broken by the assumption of planar motion and by the particular positioning of the viewing camera on the moving vehicle. The first asymmetric finding is that Eqs. 1 cannot be solved along the horizontal line $l_h$ of the image plane. The second asymmetric result is that the horizontal component of the optical flow is better suited than the vertical component for the computation of the angular velocity. Let us now show that the two components of the optical flow cannot be equivalently used for depth estimation. Let us substitute Eq.5 in Eq. 6 and write $1/\tau$ as a function of $\Omega$ and $u$. After simple calculations, it is easy to obtain

$$\frac{1}{\tau} = \frac{u}{x} - \left(\frac{x}{f} + \frac{f}{x}\right)\Omega \tag{12}$$

Eq. 12 is only superficially similar to Eq. 10. Since a rotation around a vertical axis has a large effect on the horizontal component of the apparent motion, the second

term on the right-hand-side of Eq. 12 is now also important. The stability analysis of Eq. 12 yields

$$\Delta\left(\frac{1}{\tau}\right) = \left|\frac{1}{x}\right| \Delta u + \left|\frac{x^2 + f^2}{fx}\right| \Delta\Omega \tag{13}$$

From Eq. 13 it follows that the accuracy in the reconstruction of the 3D structure on the horizontal component of the optical flow depends rather critically from the horizontal location of the image point (the error is unbounded for $x \rightarrow 0$). Since the vertical location is also critical (the optical flow equations cannot be solved along $l_h$), the use of the horizontal component of the optical flow for depth estimation is unlikely to be effective. Let us now corroborate the analysis presented in this section by means of some experimental results.

# 4 Experimental results

In this section the stability analysis of the previous section is first tested on synthetically generated flows corrupted by noise. Then, results of experiments on sequences of real images are reported and discussed.

## 4.1 Synthetic data

Let us first study the effect of the uncertainty in the knowledge of the geometry of the problem. The previous analysis was based on some structural assumptions on the planarity of the observed motion, position of the viewing camera with respect to the environment and, knowledge of a few intrinsic parameters of the camera (like focal length and location of the optical center). In practice, the observed motion is not exactly planar and the position of the viewing camera and the values of the intrinsic parameters are not known precisely. Therefore, it is useful to study the robustness with which Eq. 8 and 10 can be solved in the presence of noise and uncertainty in the intrinsic parameters. The two optical flows of Figure 1*a* and *b* have been synthetically produced by simulating the motion of a slanted planar surface which is moving toward the image plane. The image plane is assumed to be $256 \times 256$ pixels and the focal length 500 pixels. In both cases the planar surface is translating along the $Z$-axis with $T = 1cm/sec.$ Since the distance between the planar surface and the image plane, in the same units, is $50cm$, we have that $\tau = 50sec.$ Figure 1*a* depicts the case of pure translation. To better simulate motion in the real world small amounts of angular velocity were added along the $X$- and $Y$-axis ($.02°/sec$ and $.03°/sec$ respectively). This is the reason why the focus of expansion in Figure 1*a* is not exactly in the center of the optical flow. In Figure 1*b* the planar surface is also rotating with $\Omega = .3°/sec.$ In this case noisy motion was simulated by adding a small component of angular velocity around the $X$-axis ($.02°/sec$). Finally, both

| True values | $\Omega\ (x=0)$ | $\Omega\ (x=2)$ | $\Omega\ (x=4)$ | $\Omega\ (x=8)$ |
|---|---|---|---|---|
| .03 *deg/frame* | $.0297 \pm .0001$ | $.0327 \pm .0001$ | $.0350 \pm .0001$ | $.0404 \pm .0002$ |
| .3 *deg/frame* | $.299 \pm .001$ | $.303 \pm .001$ | $.305 \pm .001$ | $.311 \pm .001$ |

Table 1: Stability of the estimation of the angular velocity by means of Eq. 9. The first and second row refer to the optical flow of Figure 1*a* and *b* respectively. The first column reproduces the true values of $\Omega$, the second the estimates obtained by using the vertical line which goes exactly through the center of the image plane, while the third, fourth, and fifth columns the results which were obtained by using a vertical line 2, 4, and 8 pixels away from the true vertical line respectively. The angular velocity is averaged over 256 estimates because the noise added to the flow is uncorrelated.

the flows were corrupted by additive random noise (see the legend of Figure 1 for details).

In Figure 1*c* and *d* the 3D reconstruction obtained over the horizontal line drawn in Figs. 1*a* and *b* is shown. The filled circles reproduce the depth estimates obtained through Eq. 10 from the vertical component of optical flow, while the open triangles the estimates obtained through Eq. 12 from the horizontal component of optical flow. In both cases the angular velocity was computed from the vertical line 2 pixels away from $l_v$. Notice that the results displayed in Figure 1*c* and *d* are in very good agreement with the theoretical analysis of the previous section. While the depth estimates from the vertical component of the optical flow are almost independent of both the angular velocity and the horizontal location of the image point, the accuracy in the depth estimates from the horizontal component is much lower, becomes meaningless in the proximity of $x = 0$, and decreases noticeably when $\Omega$ increases.

Table 1 shows the stability of the estimation of $\Omega$ computed from Eq. 9 under the assumption that the location of the optical center is known within a certain margin of error. The first and second row of Table 1 correspond to the case of Figure 1*a* and *b* respectively. The first column of Table 1 reproduces the true values, while the second column shows the estimates of $\Omega$ obtained from $l_v$, the vertical line which goes exactly through the center of the image plane. The third, fourth, and fifth columns contain the results obtained from a vertical line which is 2, 4, and 8 pixels away from $l_v$ respectively (the values in the third column were used in the examples of Figure 1). From Table 1 it can easily be seen that the estimation of the angular velocity which can be obtained from Eq. 9 is rather stable, accurate, and nearly independent of the amount of rotational component. Equivalent results were obtained in many other similar experiments.

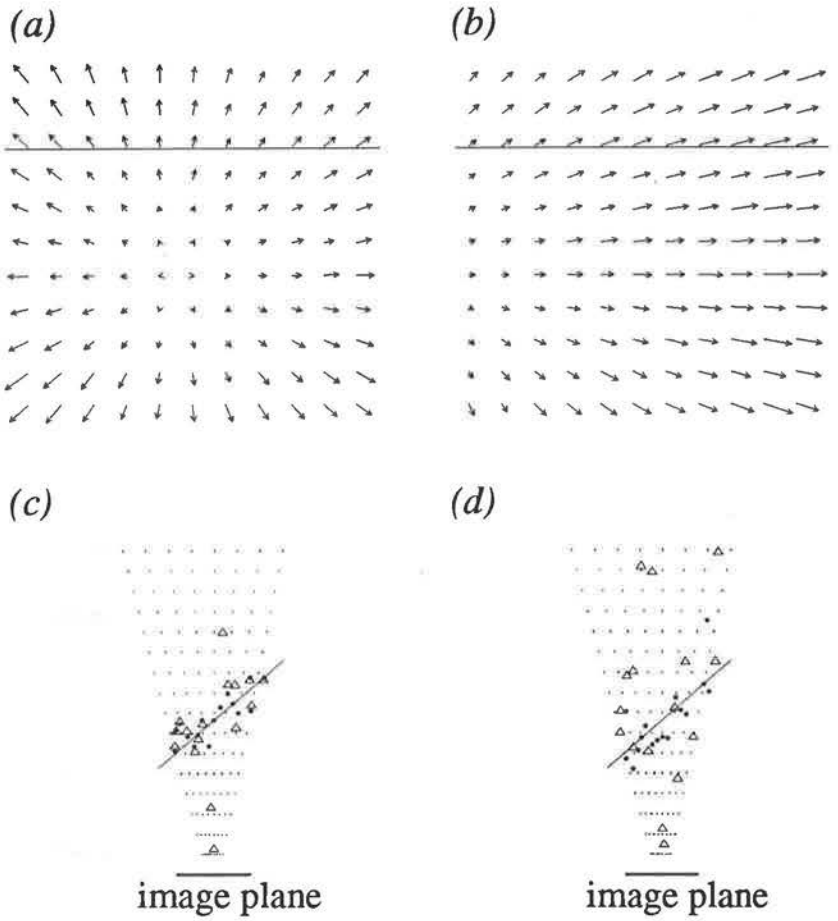


Figure 1: Depth estimates from noise corrupted synthetic data. ($a$) and ($b$) Synthetically generated optical flows which reproduce the motion of a slanted planar surface which is moving toward the image plane. See the text for numerical details. The flow vectors, which were corrupted by zero mean white Gaussian noise with standard deviation equal to 10% of the maximum value, are subsampled for better visualization. ($c$) Top view of the subsampled 3D reconstruction obtained by computing depth by means of Eq. 10 (filled circles) and Eq. 12 (open triangles) along the line of ($a$) and ($b$) respectively. The slanted straight lines mark the true location of the moving plane. The mean of the relative errors in the reconstruction of depth were $\epsilon_v = .10$ and $\epsilon_u = .59$ for the estimates obtained from the vertical and horizontal component respectively for ($c$), and $\epsilon_v = .12$ and $\epsilon_u = .84$ for ($d$). The estimates of the angular velocity which were used are shown in the third row of Table 1.

## 4.2 Real images

Let us now present an experiment on a sequence of real images. Figure 2*a* shows a frame of a sequence of 17 frames in which the viewing camera is translating and rotating toward the viewed scene on a rail. The camera motion was controlled by a linear stepping motor and the motion parameters were $\Omega = 0.05°/frame$ and $T = 5.08cm/frame$. The scene consists of a piece of furniture on the left part of the image, a table in the bottom and bottom-right part, and a wall in the background. The sharp depth discontinuity between the wall and the piece of furniture was nearly $2.9m$, while the distance between the wall and the camera, at the first frame of the sequence, was nearly $7.8m$. The presented experiment may appear deceptively simple. In fact, in a sufficiently small time interval, due to the constraint of planar motion and the particular position of the viewing camera on the moving vehicle, the experiment describes a typical motion.

Figure 2*b* show the optical flow associated with the third frame of the sequence and computed according to the algorithm described in [12]. Table 2 shows the angular velocity estimated over the sequence by means of Eq. 8 and assuming the optical center in the image center. As it can easily be seen from Table 2 the estimates are consistently very good throughout the entire sequence. Since the method runs at a few hertzs on a SPARC workstation (computation of the horizontal component of the optical flow over the vertical axis included), it is possible to measure the amount of the observed rotation, or control the heading of a mobile vehicle without the need of dedicated hardware.

Figure 2*c* and *d* show the 3D reconstructions which were obtained by using the vertical component $v$ of the optical flow along the two horizontal lines superimposed to the flow of Figure 2*b* at the third and tenth frame respectively. Correctly, the depth discontinuity is detected at the lower horizontal lines (filled circles) but not at the higher horizontal lines (open circles). Similar estimates were obtained with the remaining frames. Qualitatively, the results are in very good agreement with the distances directly measured in the scene. Quantitatively, the depth estimate is not always accurate mainly because the method which has been used for the computation of optical flow does not perform equally well over different ranges of image motion and texture. Similar results have been obtained in two other image sequences of comparable structural complexity.

# 5 Conclusion

In this report, the problem of passive navigation has been restricted to a particular but interesting case. A simple analysis has shown that if the moving vehicle is moving on a planar surface, an appropriate positioning of the viewing camera with respect to the vehicle makes it possible to recover all the relevant visual information by computing 1D optical flow. Apart from the assumption of planar motion, the presented analysis was not based on further hypotheses on the viewed motion, nor on

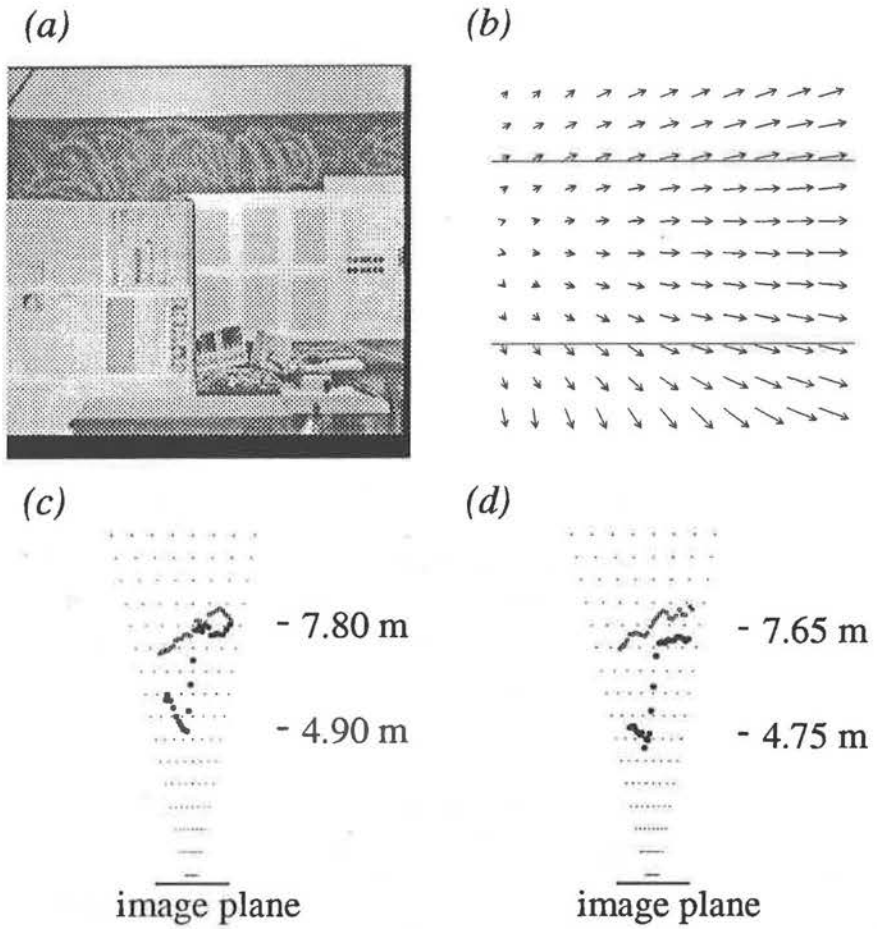


Figure 2: The recovery of motion and structure from real images. ($a$) One frame of a sequence in which the viewing camera was translating along the optical axis and rotating around a vertical axis orthogonal to the ground plane. The motion parameters were $\mathbf{T} = (0, 0, 5.08) cm/frame$ and $\mathbf{\Omega} = (0, .05, 0)°/frame$. The sequence consisted of 17 frames grabbed by means of a Panasonic BL202 camera and a Datacube Digimax board. Each image consists of $480 \times 512$ pixels. The focal length was 745 pixels. The background wall was at $7.8m$ from the camera and the piece of furniture to the left was $4.9m$ at the beginning of the sequence. ($b$) The subsampled optical flow relative to the frame of ($a$) computed according to the algorithm described in [12]. ($c$) Top view of the subsampled 3D reconstruction of ($a$) from the vertical component $v$ of the optical flow of the third frame of the sequence along the lower line of ($b$) (filled circles) and the upper line of ($b$) (open circles). ($d$) Top view of the subsampled 3D reconstruction of ($a$) from the vertical component $v$ of the optical flow of the tenth frame of the sequence along the same lines of ($c$). The angular velocity estimates are shown in Table 2.

| frame | $\Omega$ [*deg/frame*] |
|---|---|
| 3 | .0494 $\pm$ .0003 |
| 4 | .0490 $\pm$ .0002 |
| 5 | .0493 $\pm$ .0003 |
| 6 | .0500 $\pm$ .0003 |
| 7 | .0502 $\pm$ .0002 |
| 8 | .0494 $\pm$ .0002 |
| 9 | .0497 $\pm$ .0002 |
| 10 | .0481 $\pm$ .0002 |
| 11 | .0501 $\pm$ .0001 |
| 12 | .0497 $\pm$ .0002 |
| 13 | .0494 $\pm$ .0001 |
| 14 | .0479 $\pm$ .0001 |
| 15 | .0508 $\pm$ .0001 |

Table 2: Estimates of the angular velocity over the sequence of Figure 2. Even if the optical flow was computed at each point of the image plane, the angular velocity was estimated by averaging the horizontal component of the flow at only 20 locations equally spaced along the vertical line $l_v$. A finer sampling would not improve the accuracy because the flow estimates are spatially correlated [12]. The errors are the standard errors in the assumption that the 20 flow estimates are uncorrelated.

the 3D structure of the viewed scene. Experiments on synthetic data and sequences of real images indicate that 1D optical flow [13] can be used effectively for the recovery of the viewed motion and structure. The horizontal component of the optical flow along the vertical line through the center of the image is well-suited for the estimation of the angular velocity, while the vertical component of the optical flow is better suited for the estimation of the generalized time-to-collision, i.e., the ratio between the viewed depth and the translational speed. Therefore, in the case of planar motion, the computation of the 2D optical flow over the entire image plane can be probably avoided and it should be possible to develop a simple optical flow based method for passive navigation on conventional hardware.

**Acknowledgments.** We wish to thank Rod Barman, Stewart Kingdon, and Dan Razzell for their kind help. Fruitful discussions with J. Little, E. Bandari, J. Weber, and J. Malik are gratefully acknowledged. This research was partially supported by the "Progetto Sensori ed elaborazione dell'informazione per la guida automatica di robot" of the Consiglio Nazionale delle Ricerche, and project A-1 of the Institute for Robotics and Intelligent Systems, in the Network of Center for Excellence in Canada.

## References


1. Longuet-Higgins, H.C., and Prazdny, K., 1980. The interpretation of a moving retinal image. *Proc. R. Soc. London B* **208**, 385-397.

2. Prazdny, K., 1981. Egomotion and relative depth from optical flow. *Biological Cybernetics* **36**, 87-102.

3. Tsai and Huang, 1984. Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces. *IEEE Trans. Patt. Anal. Mach. Intell.* **6**, 13-27.

4. Adiv, G., 1985. Determining 3D motion and structure from optical flow generated by several moving objects. *IEEE Trans. Patt. Anal. Mach. Intell.* **7**, 384-401.

5. Rieger, J.H. and Lawton, D.T., 1985. Processing differential image motion. *J. Opt. Soc. Amer. A* **2**, 354-359.

6. Horn and Negahdaparipour, 1987. Direct Passive Navigation: Analytical Solution for Planes. *IEEE Trans. Patt. Anal. Mach. Intell.* **9**, 168-176.

7. Faugeras, O.D., Lustman, F., and Toscani, G., 1987. Motion and Structure from Motion from Points and Line Matches. *Proc. 1st Int. Conf. Comput. Vision*, London, 25-34.

8. Heeger, D.J. and Jepson, A.D., 1990. Simple method for computing 3D motion and depth. *Proc. 3rd Intern. Conf. Comput. Vision* Osaka, 96-100.

9. Tistarelli, M., and Sandini, G., 1992. Dynamics Aspects in Active Vision. *CVGIP: Image Understanding* **56**, 108-129.

10. Heeger, D.J. and Jepson, A.D., 1992. Subspace Methods for Recovering Rigid Motion I: Algorithm and Implementation. *Int. J. Comput. Vision* **7**, 95-117.

11. Adiv, G., 1989. Inherent ambiguities in recovering 3D motion and structure from a noisy flow field. *IEEE Trans. Patt. Anal. Mach. Intell.* **11**, 477-489.

12. Campani, M. and Verri, A., 1990. Computing optical flow from an overconstrained system of linear algebraic equations. *Proc. 3rd Intern. Conf. Comput. Vision* Osaka, 22-26.

13. De Micheli, E., Uras S. and Torre, V., 1993. The Accuracy of the Computation of Optical Flow and of the Recovery of Motion Parameters. *IEEE Trans. Patt. Anal. Mach. Intell.* **15-5**, 1-14.