The active camera as a projective pointing device *

Andrew J. Davison, Ian D. Reid and David W. Murray
Department of Engineering Science, University of Oxford,
Parks Road, Oxford, OX1 3PJ, U.K.
email: [ajd,ian,dwm]@robots.oxford.ac.uk

Abstract

This paper demonstrates an approach which exploits an active camera as a projective pointing mechanism. The optical centre of a static camera is notionally substituted by the centre of rotation of the active camera, as is the image plane by a frontal plane, a plane perpendicular to the optical axis of the active camera in its resting direction. Algorithms devised for 3D motion and 3D structure recovery using a single passive camera become immediately applicable to the active camera without need for reformulation. Furthermore, because the active camera can access a panoramic field of view, instabilities which may arise when the field of view is small, or because the shared field of view between successive views after movement is small, are lessened. Two quite different applications of the idea are presented. In the first, the homography between a planar surface in the scene and the frontal plane is recovered and used to recover scene trajectories. In the second, the essential matrix between points in two frontal-plane "views" is recovered and used to determine the motion of a mobile vehicle.

1 Introduction

One of the arguments in favour of an active approach to visual sensing is that the burden of representing the environment is lightened by allowing the world to be its own best memory. Having control over its gaze direction allows the visual system to take a first or another look at parts of the scene that might have been neglected or forgotten. Although representational burdens are lightened, one might ask what burden remains? Certainly some memory and some representation are required if the system is not to be in state of perpetual and unsophisticated surprise.

In this paper we demonstrate an approach which regards the active camera as a *projective pointing mechanism.* The centre of projection is not (necessarily) the optical centre of the camera, but rather is the centre of rotation of the camera. Similarly the image plane is not that of the camera, but a notional fronto-parallel plane — that is, perpendicular to the optical axis of the camera in its resting direction.

The significant advantage of this representation is that algorithms devised for 3D motion and 3D structure recovery using a single passive camera become immediately applicable to the active camera without need for reformulation. Moreover,

*Supported by Grant GR/J65372 from the EPSRC. AJD is funded by an EPSRC research studentship. IDR is a supported by a Glasstone Fellowship from the University of Oxford.

BMVC 1995 doi:10.5244/C.9.45

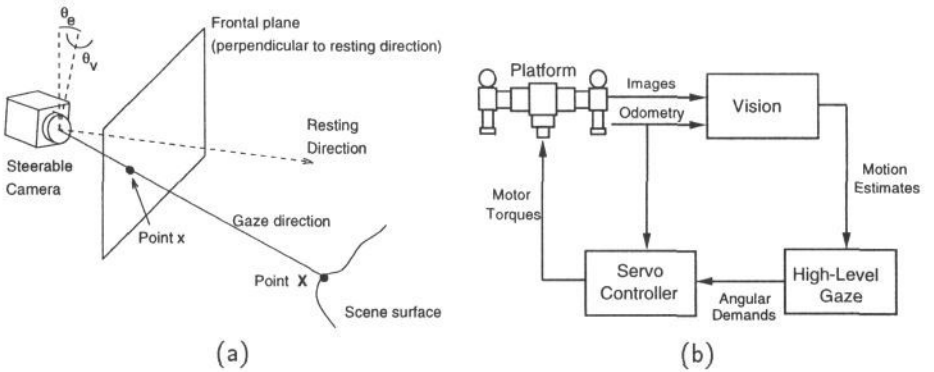

Figure 1: (a) The geometrical relationship between the directable camera, the frontal plane and the scene point. (b) The overall system architecture of the visuo-control loop in our active camera system.

the instabilities in such algorithms which may arise when the field of view is small are lessened because the active camera provides a panoramic but nevertheless undistorted field of view.

We illustrate the idea with two demonstrations:

- in §2, the calibration of a scene plane by finding the homography between scene plane and fronto-parallel plane and its subsequent use in recovering planar trajectories; and
- in §3, the recovery of ego-motion of a mobile vehicle using an active camera mounted upon it by applying the 8-point algorithm [1, 2].

The camera and scene geometries under consideration throughout the paper are sketched in Figure 1(a). We assume a single directable camera which rotates about an elevation (up-down) axis and a vergence [1] (left-right) axis. In our configuration these axes are perpendicular, but more relevant is that the two axis angles (θ_e, θ_v) together define a gazing direction which intersects a notional "frontal plane" at point $\mathbf{x}$. The frontal plane is fixed perpendicular to the resting gaze direction defined by $\theta_e = \theta_v = 0$, and lies an arbitrary distance in front of the rotation centre of the camera. For our camera platform, at zero vergence the camera optic axis is perpendicular to the elevation axis and the inverse kinematics give, taking the arbitrary distance as unity,

$$\mathbf{x} = (\tan\theta_v / \cos\theta_e,\ \tan\theta_e,\ 1)^\mathsf{T} .$$

The gazing direction goes on strike a surface in the scene at point $\mathbf{X}$.

The principal elements in our active vision system — camera platform, vision system, gaze controller and servo controller — are shown in Figure 1(b).

[1] The camera is one of a stereo pair. Although the second camera is not used we retain the vergence label.

2 Active recovery of a plane-plane homography

2.1 Theory

In the first example, the scene is assumed planar. Any pair of corresponding points on the frontal plane and scene plane must therefore be related by an homography

$$\mathbf{X} \equiv [\mathtt{M}]\mathbf{x} \ , \tag{1}$$

where $\equiv$ denotes projective equality and where the vectors are expressed in homogeneous coordinates $\mathbf{x} = (x, y, 1)^\top$ and $\mathbf{X} = (X, Y, 1)^\top$. $[\mathtt{M}]$ is a 3×3 matrix which, because scale is arbitrary, has only eight degrees of freedom and thus can be recovered by establishing the correspondence between (at least) four known points [3] or four lines [4] in the two planes.

To establish the calibration using points, the active camera is directed to view in succession the calibration points whose scene coordinates $\mathbf{X}_i$ are known. The fronto-parallel plane position $\mathbf{x}_i$ is derived from the joints angles $(\theta_e, \theta_a)_i$. Each point correspondence provides an equation

$$\begin{pmatrix} \lambda_i X_i \\ \lambda_i Y_i \\ \lambda_i \end{pmatrix} = \begin{bmatrix} M_{11} & M_{12} & M_{13} \\ M_{21} & M_{22} & M_{23} \\ M_{31} & M_{32} & 1 \end{bmatrix} \begin{pmatrix} x_i \\ y_i \\ 1 \end{pmatrix} \tag{2}$$

where the explicit unknown scale restores actual equality, and where to constrain the degrees of freedom we set $M_{33} = 1$. Eliminating λ_i leaves two equations contributing to the system

$$\begin{bmatrix} \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_i & y_i & 1 & 0 & 0 & 0 & -X_i x_i & -X_i y_i \\ 0 & 0 & 0 & x_i & y_i & 1 & -X_i y_i & -Y_i y_i \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \begin{pmatrix} M_{11} \\ M_{12} \\ M_{13} \\ \vdots \\ M_{32} \end{pmatrix} = \begin{pmatrix} \vdots \\ X_i \\ Y_i \\ \vdots \end{pmatrix} \ . \tag{3}$$

or abbreviated $[\mathtt{A}]\mathbf{m} = \mathbf{b}$. For four or more points, the system can be solved in the least-squares sense as $\mathbf{m} = [\mathbf{A}^\top\mathbf{A}]^{-1}[\mathbf{A}^\top]\mathbf{b}$.

If the active camera then tracks an object, generating a sequence of joint angles $\theta_e(t), \theta_v(t)$, the frontal plane trajectory $\mathbf{x}(t)$ can be computed, and hence the ground plane trajectory recovered as $\mathbf{X}(t) = [\mathtt{M}]\mathbf{x}(t)$.

2.2 Experiment

2.2.1 Calibration by recovering the homography

The view obtained in the resting $\theta_e = \theta_v = 0$ direction from the window of the tower block is given in Figure 2(a). Overlaid on the image are the four points used for calibration. In the scene these points $\mathbf{X}_i$ are the corners of a nominally rectangular lawn measured as $(X, Y)_A = (0.0, 0.0)$, $(X, Y)_B = (21.1, 0.0)$, $(X, Y)_C = (21.1, 7.9)$, $(X, Y)_D = (0.0, 7.9)$ in the ground plane, where each unit is 1 metre. Centering the active camera on these points gave head joints

angles of $(\theta_e, \theta_v)_A = (1.24°, 21.51°)$, $(\theta_e, \theta_v)_B = (-3.62°, 4.66°)$, $(\theta_e, \theta_v)_C = (-0.48°, -0.67°)$, $(\theta_e, \theta_v)_D = (3.53°, 16.40°)$, giving

$$[\mathtt{M}] = \begin{bmatrix} 0.228 & 0.601 & -73.281 \\ 0.135 & -0.479 & 51.757 \\ 0.002 & 0.013 & 1.000 \end{bmatrix} .$$

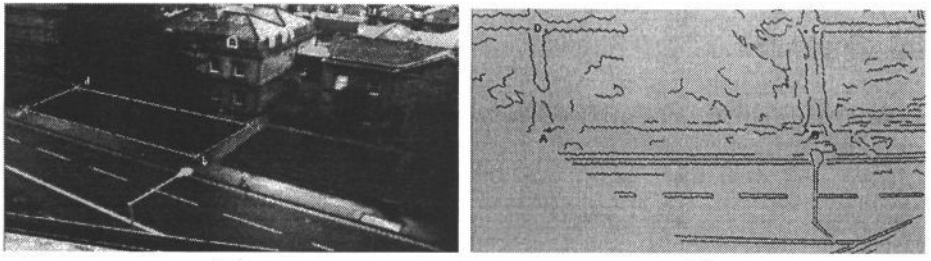

(a) (b)

Figure 2: (a) The four points used for computing the homography [M]. (b) The effect of the transformation on static features.

2.2.2 Tracking to recover the observer trajectory

To track, edge-normal components $\mathbf{v}$ of the motion field $\dot{\mathbf{r}}$ are derived at 25Hz from the central portion of the image[5, 6]. After subtracting the motion due to head rotation (which is depth independent), any static background will have $\mathbf{v}_s \sim \mathbf{0}$ and can be excluded from further consideration. Foreground motion regions are grown by spatially grouping the non-background vectors. Because of the small size of the fovea ($\sim 6°$ field of view in both x and y), multiple moving regions found are assumed to arise from the same moving object, and so the motion and location of all segmented regions are combined to yield a mean position and mean velocity estimate, $\langle \mathbf{r} \rangle, \langle \dot{\mathbf{r}} \rangle$. These demands are filtered using a constant velocity filter, and then "prompt" demands are predicted over the latency in the visual feedback loop, and sent via the inverse kinematics to the servo-controller.

Figure 3: The optical flow as detected, segmented and fitted to the foveal imagery while tracking a person. The stills are every fourth frame from a 25 Hz sequence.

An example of the output from the foveal optical flow process during tracking a person is shown in Figures 3, and the joint angle trajectory $\theta_v(t)$ and $\theta_e(t)$ captured over a period of one minute when that person left the University's Computing Centre, turned left along the pavement for some 30m, and then doubled back, cutting across the lawn, is shown in Figure 4(a) and (b). This movement, along

with several others, is shown as a trajectory $\mathbf{x}(t)$ in the frontal plane in Figure 4(c), and the tracks transformed to $\mathbf{X}(t)$ in the ground plane are shown in (d).

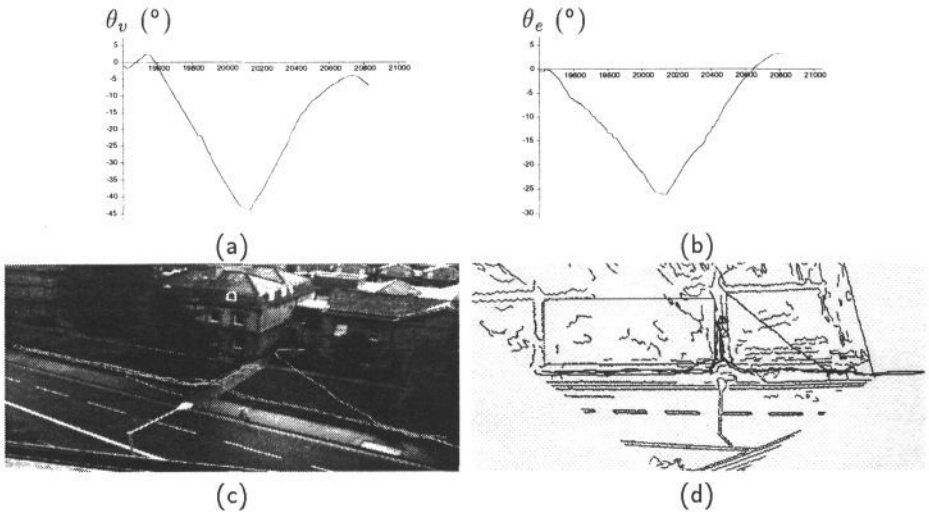

Figure 4: Recorded joint angles during one trajectory (a), (b). Recovered trajectories $\mathbf{x}(t)$ in frontal plane (c) and transformed to $\mathbf{X}(t)$ in ground plane (d).

3 Active realization of the 8-point algorithm

The second example of the exploitation of the active camera as a projective pointing device is an implementation in the frontal plane of the 8-point method of Longuet-Higgins [1] and Tsai and Huang [2]. In a conventional implementation, at least eight point correspondences between a pair of images obtained from a moving camera enable the calculation of the camera's motion through a scene and the scene structure up to the depth/speed scaling ambiguity. For our purposes, replacing the single passive camera with the logically identical combination of active pointing device and frontal plane allows us to use the algorithm to find the egomotion of a mobile vehicle on which the active camera is mounted.

As before, we use one of the cameras of a stereo platform, here mounted on a mobile vehicle as shown in Figure 5. For convenience, the resting direction $\theta_e = \theta_v = 0$ is approximately lined up with the front-back axis of the vehicle so that the frontal plane is nominally parallel to the front face of the vehicle.

The raw measurements used to generate points on the frontal plane are the angles of elevation and vergence necessary to point the active camera directly at a world feature. The active device is therefore *calibrated*: we know explicitly the relationship between image positions and ray angles. This means we can relate correspondences between points in successive positions of the frontal plane by the essential matrix rather than the fundamental matrix.

As explained in § 3.1, in an initial position of the vehicle the active camera fixates in turn upon a set of features X_i, for each of which the platform angles

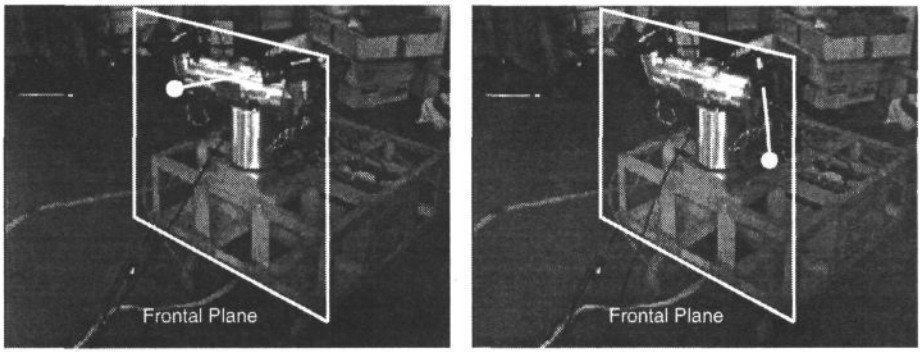

Figure 5: The active platform mounted on a vehicle. The frontal plane is perpendicular to the $\theta_e = \theta_v = 0$ direction, which is lined up with the axis of vehicle. (a) and (b) show how two points on the plane correspond to camera angles.

$(\theta_e, \theta_v)_i$ are measured and frontal plane coordinates $\mathbf{x}_i$ derived. The robot is now driven to a second position. Let $[\mathtt{R}]$ and $\mathbf{t}$ be the rotation and translation that relate the coordinate frames attached to the vehicle in positions 1 and 2, so that a point X in the scene has descriptions in the two frames related by

$$\mathbf{X} = [\mathtt{R}]\mathbf{X}' + \mathbf{t} \ .$$

The active camera now seeks to fixate in turn on the same set of scene features as before, generating a new set of joint angles $(\theta_e, \theta_v)'_i$, and new coordinates in the frontal plane $\mathbf{x}'_i$. Now corresponding coordinates are related by

$$\mathbf{x}^\top [\mathtt{E}]\mathbf{x}' = 0 \ ,$$

where the essential matrix is $[\mathtt{E}] = [\mathtt{t}_\times][\mathtt{R}]$, in which the antisymmetric matrix $[\mathtt{t}_\times]$ is made from the components of the translation vector:

$$[\mathtt{t}_\times] = \begin{bmatrix} 0 & -t_3 & t_2 \\ t_3 & 0 & -t_1 \\ -t_2 & t_1 & 0 \end{bmatrix} \ .$$

Weng *et al.* [7] give a method of computing and decomposing the essential matrix into a rotation matrix $[\mathtt{R}]$ and the unit vector in the direction of $\mathbf{t}$. In brief, omitting certain sign-checks, this comprises computing the essential matrix by minimizing $||[\mathtt{A}]\mathbf{e}||$ where

$$[\mathtt{A}] = \begin{bmatrix} \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x'_i x_i & x'_i y_i & x'_i & y'_i x_i & y'_i y_i & y'_i & x_i & y_i & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{bmatrix} \ .$$

Vector $\mathbf{e} = (e_1 \ldots e_9)^\top$ is found as the unit eigenvector associated with the smallest

eigenvalue of $[\mathtt{A}]^\top[\mathtt{A}]$, and then

$$[\mathtt{E}] = \begin{bmatrix} e_1 & e_4 & e_7 \\ e_2 & e_5 & e_8 \\ e_3 & e_6 & e_9 \end{bmatrix} .$$

The unit vector $\hat{\mathbf{t}}$ is found by minimizing $||[\mathtt{E}]^\top\hat{\mathbf{t}}||$. The minimization can be achieved by singular value decomposition of $[\mathtt{E}]^\top$ into $[\mathtt{E}]^\top = [\mathtt{U}][\mathtt{W}][\mathtt{V}]^\top$. The unit vector will be the right singular vector $\mathbf{V}_i$ associated with the smallest singular value, which should be zero or close to zero. The rotation matrix $[\mathtt{R}]$ is determined by minimizing $||[\mathtt{R}]^\top[-\mathbf{t}]_\times - [\mathtt{E}]^\top||$.

3.1 Experiments

To use the 8-point algorithm we need to establish correspondence between the frontal plane coordinates of points seen from two different vehicle positions. A method is needed whereby features in the 3D world can be reliably fixated by the active head from two potentially widely spaced viewpoints. Here we use correlation which proves to be sufficiently robust to significant vehicle motions. The choice of correlation mask size is taken as 25 × 25 pixels — small enough to provide good localization, but large enough to provide distinctive information content.

With the vehicle at an initial position, the active camera is directed so that the chosen feature is within the field of view, and a 25 × 25 pixel patch including it is selected by hand in the image. Using correlation the system then performs a closed-loop "lock on" procedure to move the active head until the feature is fixated in the centre of the view; this uses correlation and approximate camera calibration to iteratively move the head to the required angles. The platform angles θ_e and θ_v are used to calculate the frontal plane coordinates for the feature, and the now centralized image patch is saved.

When around 15–20 features are found from the first vehicle position, the vehicle is driven to a new position. One by one the features are re-fixated. Using whatever knowledge of the vehicle motion is available, the camera is pointed in an estimated direction[2] for each object and the image is searched for the patch giving the best correlation with that saved from the first view. Once it is found in the image, the closed-loop search re-fixates the camera on the feature. The joint angles enable the new frontal plane coordinates to be calculated. Figure 6 shows examples of fixated matches. With a suitable threshold set for a successful correlation and for the motions tried, the occurrence of false matches was infrequent, despite the additional problem that our platform geometry introduces cyclotorsion.

When the set of correspondences is as complete as possible, the frontal plane coordinates are used to determine the essential matrix and hence the vehicle motion. We report experiments with the four different vehicle motions shown in Figure 7 and also describe an initial experiment to investigate the benefits in accuracy afforded by the wide viewing range of the active camera.

Experiment 1: translation and rotation in the ground plane. In the first experiment the vehicle was moved on the ground plane with translation towards its

[2] see §4

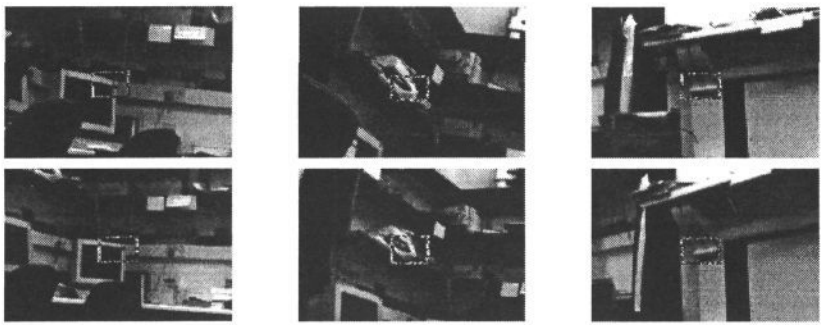

Figure 6: Feature patches matched by correlation after vehicle motion.

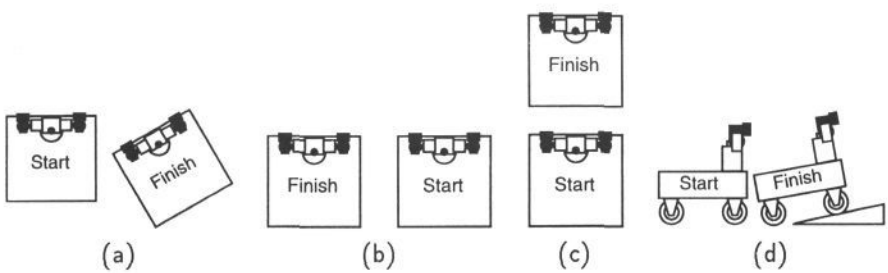

Figure 7: The vehicle motions used in experiments 1 (a), 2 (b,c) and 3 (d).

right and backwards, and a rotation to the left measured as 30°. The raw angular results obtained for fixation on ten feature points from each view and the frontal plane coordinates derived from them are displayed in Table 1. The veridical and measured direction of this motion in the $xz-$plane were

$$\hat{\mathbf{t}}_v = \begin{pmatrix} -0.855 \\ -0.142 \\ -0.498 \end{pmatrix} \quad \hat{\mathbf{t}}_m = \begin{pmatrix} -0.855 \\ -0.142 \\ -0.498 \end{pmatrix} ,$$

and the veridical and measured rotation matrices were

$$[\mathbf{R}]_v = \begin{bmatrix} 0.867 & 0.000 & 0.500 \\ 0.000 & 1.000 & 0.000 \\ -0.500 & 0.000 & 0.867 \end{bmatrix} \quad [\mathbf{R}]_m = \begin{bmatrix} 0.871 & -0.011 & 0.490 \\ 0.014 & 1.000 & -0.003 \\ -0.490 & 0.010 & 0.872 \end{bmatrix} .$$

Experiment 2: translation only. In the second experiment, the vehicle was first moved in a simple sideways translation to the left as in Figure 7(b) and then forwards as in Figure 7(c). The veridical and measured motions in the first motion, where 11 feature correspondences were obtained, were

$$\hat{\mathbf{t}}_{v,m} = \begin{pmatrix} 1.000 \\ 0.000 \\ 0.000 \end{pmatrix}, \begin{pmatrix} 0.999 \\ -0.003 \\ 0.051 \end{pmatrix} \quad [\mathbf{R}]_{v,m} = \begin{bmatrix} 1.000 & 0.000 & 0.000 \\ 0.000 & 1.000 & 0.000 \\ 0.000 & 0.000 & 1.000 \end{bmatrix}, \begin{bmatrix} 1.000 & -0.004 & -0.009 \\ 0.004 & 1.000 & 0.006 \\ 0.009 & -0.006 & 1.000 \end{bmatrix} .$$

First Position				Second Position			
Joint Angles		Frontal Plane		Joint Angles		Frontal Plane	
$\theta_e(^\circ)$	$\theta_v(^\circ)$	x	y	$\theta'_e(^\circ)$	$\theta'_v(^\circ)$	x'	y'
2.35	-13.11	-0.233	0.041	2.69	-37.42	-0.766	0.047
13.33	-14.20	-0.260	0.237	16.38	-37.01	-0.786	0.294
23.94	-20.31	-0.405	0.444	31.59	-41.64	-1.044	0.615
33.30	-35.52	-0.854	0.657	52.09	-53.14	-2.171	1.284
24.18	-41.24	-0.961	0.449	43.53	-61.48	-2.538	0.950
5.71	-44.10	-0.974	0.100	11.91	-68.24	-2.561	0.211
40.20	30.10	0.759	0.845	32.33	11.00	0.230	0.633
36.50	21.78	0.497	0.740	32.82	1.30	0.027	0.645
5.99	3.41	0.060	0.105	6.22	-21.04	-0.387	0.109
18.26	29.92	0.606	0.330	15.32	6.49	0.118	0.274

Table 1: Joint angles and frontal plane coordinates to fixate on features in successive positions of the vehicle in Experiment 1.

For the second translation, the veridical and measured motion estimated for 12 correspondences were

$$\hat{\mathbf{t}}_{v,m} = \begin{pmatrix} 0.000 \\ 0.000 \\ 1.000 \end{pmatrix}, \begin{pmatrix} -0.016 \\ -0.003 \\ 1.000 \end{pmatrix} [\mathbf{R}]_{v,m} = \begin{bmatrix} 1.000 & 0.000 & 0.000 \\ 0.000 & 1.000 & 0.000 \\ 0.000 & 0.000 & 1.000 \end{bmatrix}, \begin{bmatrix} 1.000 & 0.020 & -0.042 \\ -0.020 & 1.000 & -0.009 \\ 0.041 & 0.010 & 1.000 \end{bmatrix}.$$

Experiment 3: rotation and translation out of plane. To demonstrate that recoverable motions are not confined to the ground plane, the front of the vehicle was raised up such that the vehicle's base made an angle of 10° with the ground as in Figure 7(d). Using 19 matched points, the recovered motions were

$$\hat{\mathbf{t}}_{v,m} = \begin{pmatrix} 0.000 \\ 0.175 \\ 0.985 \end{pmatrix}, \begin{pmatrix} 0.109 \\ 0.209 \\ 0.972 \end{pmatrix} [\mathbf{R}]_{v,m} = \begin{bmatrix} 1.000 & 0.000 & 0.000 \\ 0.000 & 0.985 & 0.174 \\ 0.000 & -0.174 & 0.985 \end{bmatrix}, \begin{bmatrix} 1.000 & 0.005 & 0.001 \\ -0.005 & 0.983 & 0.183 \\ -0.000 & -0.183 & 0.983 \end{bmatrix}.$$

3.1.1 The effects of viewing angle.

One of the significant advantages of an active camera is that it effectively provides a wide angle view without the distortion associated with wide angle optical devices. We have briefly explored the how the quality of motion recovery is affected by the angular range of points used in the 8-point algorithm. A large number (23) of features was located in the vehicle's starting position, roughly half of which were within the field of view of a static camera (about 40°), and half of which were at wide angles only accessible to an active camera. The vehicle was then turned through a measured angle of 20° in the $xz-$plane, and the point correspondences established in the new frame. The Weng, Huang and Ahuja algorithm [7] was then applied three times: first to entire set of point matches, then to ten of those found in the narrow field of view, and finally to ten from the wide angles. The calculation based on the complete set was taken to be our best estimate of the rotation. It gave a rotation matrix:

$$[\mathbf{R}]_v = \begin{bmatrix} 0.939 & 0.004 & -0.345 \\ -0.002 & 1.000 & 0.008 \\ 0.345 & -0.007 & 0.939 \end{bmatrix},$$

meaning a rotation angle of 20.189°.

Using just the ten points from the narrow field an angle of 20.733° was calculated, wheras the ten wide points gave the value 20.367°. While by no means conclusive, this gives some indication of the advantage of a broader range of points.

4 Discussion

In this paper we have given two examples of the modelling of an active camera as a projective pointing mechanism.

Whilst the first example provides a practical method of trajectory recovery which has been used elsewhere [8], we do not suggest that the second method is a satisfactory method of frame by frame visually guided navigation. Clearly the serial search for points would be too time consuming. The method has more immediate promise as a method of re-orientation, perhaps after the vehicle becomes interrupted during some purposive navigation task by a more urgent reactive task.

More important however is that the idea makes much of the research into motion understanding using a static projective camera immediately accessible to the active camera. As a simple example, one might consider the following: "Suppose the active camera on a vehicle views an static object and then the vehicle makes some known motion. How should the camera move to search for the object?" Using the frontal plane the answer is obviously that the fixation point should be moved along the epipolar line corresponding to the first viewed position $\mathbf{x}$. The equivalent search in joint space can then be derived from the inverse kinematics.

There appear several avenues in which the idea might be advanced. Can some way be found of breaking the need to fixate in several directions at one position by using and improving an estimate of vehicle motion? A second is to incorporate the ability to ignore the visibility constraint built into the 8-point algorithm, a constraint not required as the active head could look behind. Here work on cheirality in projective geometry should be of use.

References

[1] H.C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. *Nature*, 293:133–135, 1981.

[2] R Y Tsai and T S Huang. Uniqueness and estimation of three-dimensional motion parameters of rigid objects with curved surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6:13–27, 1984.

[3] J.L. Mundy and A. Zisserman. Projective geometry for machine vision. In J. L. Mundy and A. Zisserman, editors, *Geometric Invariance in Computer Vision*, pages 463–519. MIT Press, Cambridge MA, 1992.

[4] I. D. Reid, D. W. Murray, and K. J. Bradshaw. Towards active exploration of static and dynamic scene geometry. In *IEEE International Conference on Robotics and Automation*, pages 718–723, San Diego, May 1994.

[5] K. J. Bradshaw, P. F. McLauchlan, I. D. Reid, and D. W. Murray. Saccade and pursuit on an active head/eye platform. *Image and Vision Computing*, 12(3):155–163, 1994.

[6] D. W. Murray, P. F. McLauchlan, I. D. Reid, and P. M. Sharkey. Reactions to peripheral image motion using a head/eye platform. In *Proc. 4th Int'l Conf. on Computer Vision, Berlin*, pages 403–411, Los Alamitos, CA, 1993. IEEE Computer Society Press.

[7] J. Weng, T. S. Huang, and N. Ahuja. Motion and structure from two perspective views: algorithms, error analysis and error estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(5):451–476, 1989.

[8] K. J. Bradshaw, I. D. Reid, and D. W. Murray. Active recovery of motion trajectories and their use in prediction. Preprint. Submitted to IEEE Trans on Robotics and Automation, 1995.