

Decoupling the 3D Motion Space by Fixation

Konstantinos Daniilidis^{1*} and Inigo Thomas²

¹ Computer Science Institute, University of Kiel

² GRASP Laboratory, University of Pennsylvania

Abstract. Fixation is defined as the ability of an active visual system to keep the projection of an environmental point stationary in the image. We show in this paper that fixation enables the decoupling of the 3D-motion parameters by projecting appropriately the spherical motion field in two latitudinal directions with respect to two different poles of the image sphere. Both computational steps are based on one-dimensional searches along meridians of the image sphere. We do not use the efference copy of the fixational rotation of the camera. Performance of the algorithm is tested on real world sequences with fixation accomplished either off-line or during the recording using an active camera.

1 Introduction

The ability to perceive the three-dimensional motion relative to the environment is crucial for every robot acting in a dynamically changing world. The estimation of 3D motion parameters has been addressed in the past as a reconstruction problem: Given a monocular image sequence the goal was to obtain the relative 3D motion to every scene component as well as a relative depth map of the environment. Solutions given suffer under instability problems and require an immense computational effort which excludes a real time reactive behavior. In this study, we will show the computational advantages with respect to 3D motion estimation of the fixation on an environmental point. There is a large amount of work in biological and computer vision research on how fixation is achieved [1, 2]. The evident advantages of overcoming the field of view, foveal sensing, and reducing the motion blur have been considered sufficiently justifying the fixational movements so that only sporadic approaches delved into the computational advantages of fixation.

We show in this paper that the ability to fixate on a stationary point combined with the appropriate representation of the motion field enables the decoupling of the 3D-motion parameters. We use a spherical image surface which can be mapped 1:1 to the image plane. We do not use any information from the motor encoders or from the input in the fixation feedback loop (called the efference copy in biology). Fixation is formulated only as a constraint on the

* Preusserstr. 1-9 24105 Kiel, Germany, kd@informatik.uni-kiel.de. We heavily acknowledge the constructive discussions with Gerald Sommer, Ruzena Bajcsy, Yiannis Aloimonos, and Hans-Peter Mallot.

motion field. This constraint reduces the number of unknowns from five to four. The translation direction remains unknown (two parameters) but instead of the angular velocity (three unknowns) we obtain only the torsion - rotation about the target direction- and the time to collision to the fixated scene point. The new representation for the fixated motion field is based on two projections. Assuming that the fixated target point is the pole of the sphere we show that the latitudinal projection of the motion field has the property of being constant along a meridian. The constant value is equal to the torsion and the meridian contains the heading direction. Taking as a new pole the normal to this meridian we again project the flow field in the latitudinal direction and obtain a similar pattern: A meridian with respect to the new pole where the new latitudinal projection is constant and equal to the time to impact to the target. This new meridian fully constrains the heading direction. We are, thus, able to compute the heading direction by applying only two onedimensional searches. In case of a heading direction outside the field of view we replace the second projection with the solution of an equation in the two remaining unknowns.

2 Problem Statement

We assume that the imaging surface is a sphere with unit radius. We denote by \hat{p} the points on this sphere resulting from the projection $\hat{p} = \mathbf{P}/\|\mathbf{P}\|$. The mapping of the planar imaging surface to a spherical surface is one to one. Let $\mathbf{x} = \mathbf{P}/\hat{z}^T \mathbf{P}$ be a point on the image plane $Z = 1$ with the optical axis parallel to the Z -axis with unit vector \hat{z} . If $\dot{\mathbf{x}}$ is the motion field on that plane then it can be easily proved that the spherical motion field reads

$$\dot{\mathbf{p}} = \frac{1}{\|\mathbf{x}\|}(\hat{\mathbf{p}} \times (\dot{\mathbf{x}} \times \hat{\mathbf{p}})). \quad (1)$$

Most of the authors assume that for a small field of view the two fields are approximately equal. However, for a large field of view the above equation should be used. Special care should be taken in the mapping of the planar discretization noise onto the sphere. We assume that the observer is moving with instantaneous linear velocity \mathbf{v} and angular velocity $\boldsymbol{\omega}$ relative to the environment so that the velocity of a scene point \mathbf{P} can be written as $\dot{\mathbf{P}} = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{P}$. In case of pure ego motion all equations are valid with the opposite sign for the velocities \mathbf{v} and $\boldsymbol{\omega}$. The spherical motion field reads

$$\dot{\mathbf{p}} = \frac{1}{\|\mathbf{P}\|}(\hat{\mathbf{p}} \times (\mathbf{v} \times \hat{\mathbf{p}})) + \boldsymbol{\omega} \times \hat{\mathbf{p}} \quad (2)$$

where we can observe the classical decomposition into a translational component depending on the environment ($\|\mathbf{P}\|$) and the rotational term depending only on the image position. The spherical motion field vector lies on the tangential plane at point $\hat{\mathbf{p}}$ so that $\dot{\mathbf{p}}^T \hat{\mathbf{p}} = 0$. As we mentioned at the beginning we suppose that a control algorithm exists that makes a target point $\hat{\mathbf{t}}$ on the sphere be fixated which means $\dot{\hat{\mathbf{t}}} = 0$. From (2) follows that

$$-\frac{\mathbf{v} \times \hat{\mathbf{t}}}{\|\mathbf{P}\|} + \boldsymbol{\omega} \quad \text{is parallel to} \quad \hat{\mathbf{t}}$$

where $\|\mathbf{T}\|$ is the distance to the target scene point. Hence, the angular velocity in case of fixation reads

$$\boldsymbol{\omega} = \gamma \hat{\mathbf{t}} + \frac{\mathbf{v} \times \hat{\mathbf{t}}}{\|\mathbf{T}\|} \quad (3)$$

It is constrained to be a function of the linear velocity and possesses only one degree of freedom γ : the torsion around the target point $\hat{\mathbf{t}}$. Thus, after fixation the flow field contains three components (Fig. 1): A translational one due to \mathbf{v} , a fixational equal to the second term $\mathbf{v} \times \hat{\mathbf{t}}/\|\mathbf{T}\|$ of (3), and a torsional component $\gamma \hat{\mathbf{t}}$.




Fig. 1. The three components of a fixated motion field.

After inserting the fixation angular velocity (3) into (2) the spherical motion field of a point $\hat{\mathbf{p}}$ different from the target reads

$$\dot{\hat{\mathbf{p}}} = \hat{\mathbf{p}} \times (\mathbf{v} \times (\frac{\hat{\mathbf{p}}}{\|\hat{\mathbf{p}}\|} - \frac{\hat{\mathbf{t}}}{\|\mathbf{T}\|})) + \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}}). \quad (4)$$

After eliminating the structure information $\|\mathbf{P}\|$ by taking the scalar product with $\mathbf{v} \times \hat{\mathbf{p}}$ we obtain the “epipolar” equation for the fixated motion field

$$(\mathbf{v} \times \hat{\mathbf{p}})^T (\dot{\hat{\mathbf{p}}} - \frac{\hat{\mathbf{t}}}{\|\mathbf{T}\|} \hat{\mathbf{p}}^T \mathbf{v} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}})) = 0 \quad (5)$$

which corresponds to the instantaneous version of epipolar equation for general motion

$$(\mathbf{v} \times \hat{\mathbf{p}})^T (\dot{\hat{\mathbf{p}}} - \boldsymbol{\omega} \times \hat{\mathbf{p}}) = 0. \quad (6)$$

We see that the depth-free equation (5) contains three unknowns for the scaled linear velocity $\mathbf{v}/\|\mathbf{T}\|$ plus one unknown for the torsion γ around the target. Furthermore, the equation (5) is quadratic in the components of \mathbf{v} and bilinear in (\mathbf{v}, γ) .

The first and most important result obtained by Bandopadhyay and Ballard [3] and by Aloimonos et al. [4] was that fixation reduces the number of unknowns from five to four. Their fixation constraint was that $\boldsymbol{\omega} = (v_y, -v_x, \gamma)$ which is direct implication of (3) if we set the target parallel to the optical axis: $\hat{\mathbf{t}} = \hat{\mathbf{z}}$. In the work of Fermüller and Aloimonos [5] fixation is exploited to compute the line on the image which passes through the FOE. Using only normal flow the location of the FOE on this line is found by matching patterns to the repeatedly detranslated flow. Then the FOE is localized using the rotational control signals

of the camera movements in two subsequent time points and making the assumption that translation direction is almost constant despite fixation. The equation of fixational motion field (4) is used by Taalebinezhad [6]. The flow field in the Brightness Change Constraint Equation (BCCE) is substituted by the fixational motion field. As the BCCE at every pixel introduces a new unknown (depth) an additional assumption of minimal variation of depth near the fixation point is added. To convert the resulting minimization into an eigenvalue problem it is further assumed that the torsion γ is already computed in a preceding step. This step is solved assuming local frontoparallel patches. However, this assumption enables a local and linear computation of rotation and translation without fixation. Raviv and Herman [7] study the surfaces in the world that produce constant flow in the image. They show that the level sets of equal latitudinal flow are cylinders and that the longitudinal flow is zero along two planes. The intersection of these planes with the cylinder corresponds to the points in the world that produce zero flow. The first part in [8] is identical to the work by Raviv and Herman [7]. They derive the equal flow cylinders and planes. However, Thomas et. al. [8] apply their findings of zero longitudinal flow to determine the angle between the target and the velocity v . This plane always appears in the image as a line, provided that the FOV is 180 degrees. These results are tested using a novel 180 degrees field of view camera. In [9] fixation is combined with the log-polar transformation. Using the second order spatial derivatives of the fixated log-polar field it is shown that the time to collision can be computed using only the radial component of the velocity. Advantages of the polar transformation in case of fixation are also shown in [10] where the heading direction is computed using two specific lines through the center of the image. The work of Barth and Tsuji [11] addresses the issue of how to fixate in the direction of the translation. Their technique is based on the following heuristic. They group the flow vectors near the point of fixation into two groups: positive and negative flows. The difference in the average of the flow values at these groups indicates the direction of translation with respect to the current fixation direction. Based on this value the robot is controlled to turn towards the direction of translation. The same issue is addressed in [12] using an affine model for the optical flow field. Servoing towards the heading direction is achieved by minimizing the lateral translational components by means of a task function.

3 Projections of the fixated motion field

We proceed by projecting the fixated spherical motion field (4) into two different orthogonal basis systems of the tangential plane at an arbitrary point on the sphere. The first projection assumes that the target direction \hat{t} is the pole of the sphere defining thus a latitudinal and a longitudinal unit vector

$$\hat{\phi}_1 = \frac{\hat{t} \times \hat{p}}{\|\hat{t} \times \hat{p}\|} \quad \text{and} \quad \hat{\theta}_1 = \hat{\phi}_1 \times \hat{p},$$

respectively, lying in the tangential plane of point \hat{p} .

The second projection assumes as a pole the unit vector in the direction of $\mathbf{v} \times \hat{\mathbf{t}}$ yielding a latitudinal and a longitudinal unit vector

$$\hat{\phi}_2 = \frac{(\mathbf{v} \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}}}{\|(\mathbf{v} \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}}\|} \quad \text{and} \quad \hat{\theta}_2 = \hat{\phi}_2 \times \hat{\mathbf{p}},$$

respectively. Through the course of exposition the reader may consult Figure 2 where the projections are illustrated.




Fig. 2. The meridians with respect to the target pole $\hat{\mathbf{t}}$ are drawn on the left sphere. The spherical flow $\dot{\mathbf{p}}$ is projected on the latitudinal direction. The first step of the algorithm is a 1D search for the meridian with constant $\dot{\mathbf{p}}^T \hat{\phi}_1 / \|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|$. In the second step (see the right sphere) the pole is $\hat{\mathbf{n}}_1$ perpendicular to the meridian found in the first step. The flow without torsion $\dot{\mathbf{p}}' = \dot{\mathbf{p}} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}})$ is projected on the new latitudinal directions. A 1D search among the meridians with respect to pole $\hat{\mathbf{n}}_1$ for the meridian with constant $\dot{\mathbf{p}}'^T \hat{\phi}_2 / \|\hat{\mathbf{t}} \times \hat{\mathbf{n}}_1\|$ yields a second big circle. The intersection of the big circles found in the two steps gives the solution for the desired translation direction \mathbf{v} .

The latitudinal projection using the target direction $\hat{\mathbf{t}}$ as a pole reads

$$\dot{\mathbf{p}}^T \hat{\phi}_1 = \frac{1}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|} \mathbf{v}^T (\hat{\mathbf{t}} \times \hat{\mathbf{p}}) \left(\frac{1}{\|\mathbf{P}\|} - \frac{\dot{\mathbf{p}}^T \hat{\mathbf{t}}}{\|\mathbf{T}\|} \right) + \gamma \|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|.$$

Because the angle between the target $\hat{\mathbf{t}}$ and the considered point is known we divide by its sine which is equal to $\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|$:

$$\frac{\dot{\mathbf{p}}^T \hat{\phi}_1}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|} - \gamma = \frac{1}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|^2} \mathbf{v}^T (\hat{\mathbf{t}} \times \hat{\mathbf{p}}) \left(\frac{1}{\|\mathbf{P}\|} - \frac{\dot{\mathbf{p}}^T \hat{\mathbf{t}}}{\|\mathbf{T}\|} \right). \quad (7)$$

We see that the latitudinal component minus the torsion vanishes if the considered point lies on the plane spanned by the target and the translation direction.

Thus, we are able to constrain the translation direction if we find the meridian with longitude η where the term $\frac{\hat{\mathbf{p}}^T \hat{\boldsymbol{\phi}}}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|}$ is constant independent of the latitude $\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|$. Unfortunately this is not the only case where this term becomes constant. Suppose that a part of the environment is planar. Let the equation of the plane be $\hat{\mathbf{N}}^T \mathbf{X} = d$ and assume that the target is on the optical axis. If the plane normal reads $\hat{\mathbf{N}} = (\cos \alpha \sin \beta, \sin \alpha \sin \beta, \cos \beta)$ then it can be easily proved that

$$\frac{1}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|^2} \mathbf{v}^T (\hat{\mathbf{t}} \times \hat{\mathbf{p}}) \left(\frac{1}{\|\mathbf{P}\|} - \frac{\hat{\mathbf{p}}^T \hat{\mathbf{t}}}{\|\mathbf{T}\|} \right) = \frac{1}{d} \mathbf{v}^T \hat{\boldsymbol{\phi}}_1 \sin \beta \cos(\alpha - \eta), \quad (8)$$

which is independent of the latitude $\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|$. Hence, all meridians that are projections of lines on planes in the scene will have a constant latitudinal projection independent of the colatitude angle. Furthermore, the right hand side of (7) will vanish on the meridians that are projections of infinite depths ($1/\|\mathbf{P}\| = 0$) and on the entire field of view if the translation is parallel to the target direction: $\mathbf{v} \times \hat{\mathbf{t}} = 0$.

To summarize the defeating configurations:

1. There may exist meridians with constant latitudinal projection if these meridians are projections of planar parts of the environment or of scene points at infinity.
2. The latitudinal projection is everywhere constant if we fixate on the translation direction or if translation does not exist at all.

Suppose now that the unit vector $\hat{\mathbf{n}}_1$ in the direction $\mathbf{v} \times \hat{\mathbf{t}}$ is given and let it be the new pole. The new pole introduces new meridians and latitudes. Since torsion can be computed in the first projection above we consider the latitudinal projection of the torsion-free flow

$$(\hat{\mathbf{p}} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}}))^T \hat{\boldsymbol{\phi}}_2 = \frac{1}{\|\mathbf{P}\| \|(\mathbf{v} \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}}\|} (\hat{\mathbf{p}} \times (\mathbf{v} \times \hat{\mathbf{p}}))^T ((\mathbf{v} \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}}) + \frac{\|\mathbf{v} \times \hat{\mathbf{t}}\|}{\|\mathbf{T}\|} \|\hat{\mathbf{p}} \times \hat{\mathbf{n}}_1\|,$$

where $\hat{\mathbf{n}}_1$ is the unit vector $\frac{\mathbf{v} \times \hat{\mathbf{t}}}{\|\mathbf{v} \times \hat{\mathbf{t}}\|}$ known from the first projection. Hence, we can divide the left hand side and rewrite the right hand side as following:

$$\frac{(\hat{\mathbf{p}} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}}))^T \hat{\boldsymbol{\phi}}_2}{\|\hat{\mathbf{p}} \times \hat{\mathbf{n}}_1\|} = \frac{1}{\|\mathbf{P}\| \|(\mathbf{v} \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}}\|} \hat{\mathbf{p}}^T (\mathbf{v} \times (\mathbf{v} \times \hat{\mathbf{t}})) + \frac{\|\mathbf{v}\| \|\hat{\mathbf{v}} \times \hat{\mathbf{t}}\|}{\|\mathbf{T}\|}.$$

Considering now meridians through the pole $\mathbf{v} \times \hat{\mathbf{t}}$ we obtain following cases where the torsion-free latitudinal component will be constant.

1. On the meridian with normal $\mathbf{v} \times (\mathbf{v} \times \hat{\mathbf{t}})$.
2. On the meridians containing points with infinite depth.

The detection of the meridian with normal

$$\hat{\mathbf{n}}_2 = \frac{\mathbf{v} \times (\hat{\mathbf{v}} \times \hat{\mathbf{t}})}{\|\hat{\mathbf{v}} \times \hat{\mathbf{t}}\|}$$

allows the full computation of the translation direction

$$\hat{\mathbf{v}} = \hat{\mathbf{n}}_1 \times \hat{\mathbf{n}}_2.$$

Having obtained the heading direction we know the sine of the angle between the heading direction and the target $\|\hat{\mathbf{v}} \times \hat{\mathbf{t}}\|$. The remaining constant after vanishing of the first term in (3) yields $\lambda = \|\mathbf{v}\|/\|\mathbf{T}\|$ which is the fourth and last unknown of the motion problem in case of fixation. The inverse of it can be interpreted as the time to collision to an object at the same distance as the target in the motion direction.

To find meridians of constant value in the first and the second latitudinal projections we compute for every meridian the mean and the variance over the latitude. Then, we search for the meridians on which this variance is minimized. The means on these meridians yield the torsion and the inverse of the time to collision, in the first and second projection respectively.

Although in the first projection all meridians - or sectors of them - were contained in the field of view this is not the case in the second projection where the meridians are with respect to the new pole $\hat{\mathbf{n}}_1$. It is very easy to imagine this case if for example $\hat{\mathbf{n}}_1 = (0, 1, 0)$. We will see in the experiments that in such a case the variance of the second latitudinal projection gets its minimum at the border of the field of view. A corrective saccade can then shift the focus of expansion inside the field of view and the process can be continued with a refixation on a new point. If we want to avoid a corrective saccade we must replace the second search with a procedure as follows. The first step constrains the translational velocity to the plane with normal $\hat{\mathbf{n}}_1$. Thus, we can write

$$\hat{\mathbf{v}} = \cos \chi \hat{\mathbf{t}} + \sin \chi (\hat{\mathbf{n}}_1 \times \hat{\mathbf{t}}), \quad (9)$$

where χ is the remaining degree of freedom of the translation direction or, in the terms of the formulation above, the longitude of the searched meridian in the second step. Let rewrite (5) as

$$(\hat{\mathbf{v}} \times \hat{\mathbf{p}})^T (\hat{\mathbf{p}}' - \lambda \hat{\mathbf{t}} \hat{\mathbf{p}}^T \hat{\mathbf{v}}) = 0, \quad (10)$$

where $\hat{\mathbf{p}}' = \hat{\mathbf{p}} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}})$ is known from the second step and $\lambda = \|\mathbf{v}\|/\|\mathbf{T}\|$ is the inverse of the time to collision. If we insert $\hat{\mathbf{v}}$ from (9) in (10) in the above equation we obtain

$$\cos \chi (\hat{\mathbf{t}} \times \hat{\mathbf{p}})^T \hat{\mathbf{p}}' + \sin \chi ((\hat{\mathbf{n}}_1 \times \hat{\mathbf{t}}) \times \hat{\mathbf{p}})^T \hat{\mathbf{p}}' = \lambda \sin \chi \hat{\mathbf{p}}^T \hat{\mathbf{n}}_1 (\cos \chi \hat{\mathbf{p}}^T \hat{\mathbf{t}} + \sin \chi (\hat{\mathbf{n}}_1 \times \hat{\mathbf{t}})^T \hat{\mathbf{p}}). \quad (11)$$

This is a nonlinear equation in the two unknowns χ and λ which can be solved numerically with nonlinear minimization.

To summarize, we present the algorithmic steps of our method:

1. Choose a sampling step for the longitude angle η with respect to pole $\hat{\mathbf{t}}$ - in reality being always the optical axis if we fixate on the center. Divide the optical flow field in groups with the same longitude η corresponding to meridians. Compute for every group the mean and the variance of

$$\frac{\hat{\mathbf{p}}^T \hat{\phi}_1}{\|\hat{\mathbf{t}} \times \hat{\mathbf{p}}\|}.$$

Carry out an 1D-search for the minimum η_{min} of the variance. The new pole $\hat{\mathbf{n}}_1$ reads $(\sin \eta_{min}, -\cos \eta_{min}, 0)$ if $\hat{\mathbf{t}}$ is the optical axis.

2. Compute for all points the longitude angle χ with respect to the new pole $\hat{\mathbf{n}}_1$ and group the vectors with the same χ . Compute for every group the mean and the variance of

$$\frac{(\hat{\mathbf{p}} - \gamma(\hat{\mathbf{t}} \times \hat{\mathbf{p}}))^T \hat{\phi}_2}{\|\hat{\mathbf{p}} \times \hat{\mathbf{n}}_1\|}$$

and search for the minimum χ_{min} of the variance. Divide the mean by $\|\hat{\mathbf{v}} \times \hat{\mathbf{t}}\|$ in order to obtain the inverse of the time to collision $\|\mathbf{v}\|/\|\mathbf{T}\|$. If χ_{min} is near the border of the field of view then either carry out a saccade towards χ_{min} or apply the nonlinear minimization described above.

4 Experimental Results

We tested the proposed algorithms with synthetic as well as real data. Real data experiments were carried out using sequences recorded by passive as well as active cameras. In the non-fixated sequence we emulated the fixation by appropriately rotating the optical flow field. In all the experiments, the 1D-search of the first step runs over 45 samples of the 180 degrees η -range. The sampling interval for χ in the 1D search of the second step is one degree. If the focus of expansion lies outside the field of view we replace the second step with the alternative nonlinear minimization method applying the Levenberg-Marquardt method.

We produce synthetic motion fields assuming a scene looking like a corridor. In the first experiment we assume a wide field of view of 90 degrees and we apply translations $\mathbf{v} = (\sin \chi_{gt}, 0, \cos \chi_{gt})$ where χ_{gt} is the ground truth angle between translation and target direction. The latter is assumed to coincide with the optical axis. In this as well as all subsequent simulations it turns out that the error in the azimuthal angle η of the translation direction was under 2 deg and the relative error in the torsion γ under 3%. Therefore we will plot only the error in the χ -angle and the inverse of the time to collision λ . In Fig. 3 we show the error in the angle χ for translation directions deviating from 5 to 40 degrees from the target direction. The motion field is corrupted by gaussian noise with relative standard deviation of 10% and 20%. We tested for two torsion values 0 and 0.005, shown in the left and right of Fig. 3 respectively. We observe that the error increases with the deviation of the translation from the target direction and its behavior is not smooth in presence of torsion. The second synthetic experiment concerns a smaller field of view (45 deg) in presence of torsion and relative optical flow error of 10%. Since the second step can be applied only for $\chi < 20$ deg we applied in all steps the nonlinear minimization with respect to χ and λ . The results (Fig. 4) are significantly better than the 1D-search for even a larger field of view (see above) but with the additional cost of an iterative method. The same initial values were used in the nonlinear minimization for all translation directions. In the following image sequences we computed the optical flow with

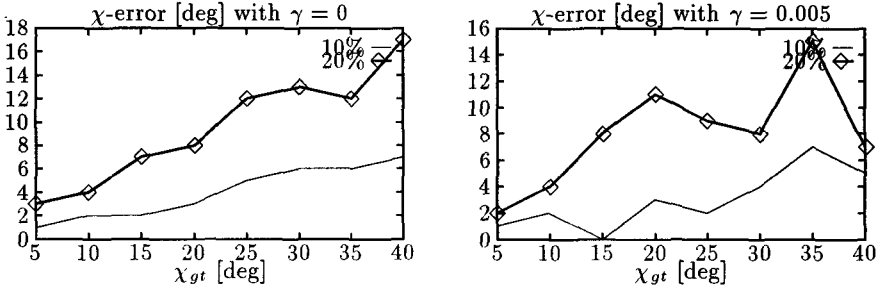


Fig. 3. The error in the χ -angle as a function of the translation direction for a field of view of 90 degrees and two values of torsion: 0 (left) and 0.005 (right). The motion field is corrupted by gaussian relative error with standard deviation of 10% and 20%.




Fig. 4. The error in the χ -angle (left) and the inverse of the time to collision $\lambda = \|\mathbf{v}\|/\|\mathbf{T}\|$ as computed by the alternative to the second step. The field of view is 45 deg and the relative error in the optical flow is 10%.

a standard differential method which assumes a constant flow field in the local neighborhood of every pixel. The spatiotemporal derivatives are computed with binomial kernels which are approximations of the first derivative of a Gaussian. The computed flow field is first mapped to the plane $Z = 1$ using the intrinsic parameters and then transformed to a spherical flow field using (1). The first sequence is the well known synthetic Yosemite sequence (Courtesy of Lynn Quam at SRI) which contains both translation with ground truth ($\eta = 90$ deg, $\chi = -9.84$ deg) and rotation with ground truth $\omega = (0.00023, 0.00162, 0.00028)$. The original and the fixated flow fields (Fig. 5 left) are computed only for the part of the image that contains ego motion (the clouds area is excluded). The minimum of the variance of the first latitudinal projection (Fig. 5, top right) gives an η estimate of 97.37 deg and a torsion estimate of -0.00063 (the opposite sign is due to our formulation of scene motion). Since the minimum of the variance of the second latitudinal projection (Fig. 5, bottom right) is at the limit of the field of view we again apply the nonlinear minimization for the second step and obtain $\chi = -5.96$ and $\lambda = 0.00145$. The second sequence is already fixated during its recording with an active camera (Fig. 6). Up to the fixational movement the motion of the observer is pure translational with ground truth measured manually ($\eta_{gt} = 0$ deg and $\chi_{gt} = 9.2$ deg). Because the focus of expansion is

inside the field of view solutions are obtained by applying both steps of our algorithm yielding the estimates $\eta = -2$ deg and $\chi_{gt} = 5$ deg (Fig. 7). As already observed in the simulations the main error is in the deviation of the translation from the target direction in the second step. The observed robustness of the first step is consistent with the theoretical results by Maybank and Jepson [13, 14] in case of general motion who proved that if the observed surface is irregular the line through the center and the focus of expansion can be robustly estimated.

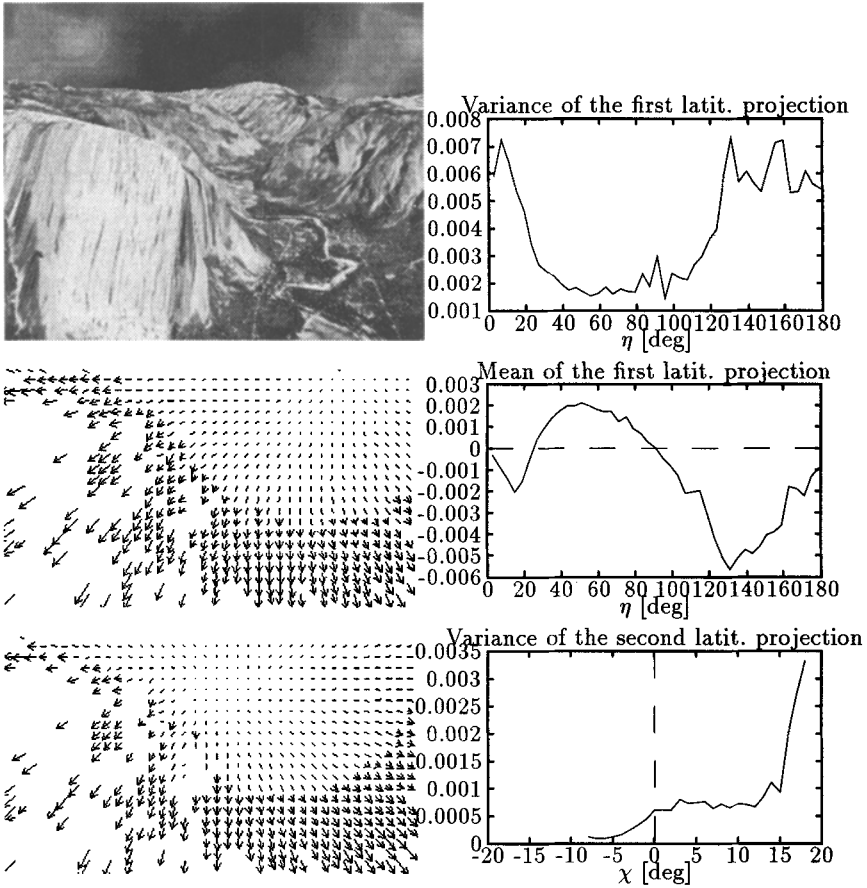


Fig. 5. The 1st and the 14th image of the Yosemite sequence (above), the original flow field (bottom left), and the fixated flow field (bottom right). The variance (top left) and the mean (top right) of the first latitudinal projection for the Yosemite sequence. The minimum of the variance gives the angle η and the mean for this η gives the torsion. The variance of the second latitudinal projection has its minimum at the right bound of χ indicating a focus of expansion outside the field of view.



Fig. 6. The 1st (left) and the 10th (middle) of the real fixated sequence and the computed optical flow field (right)

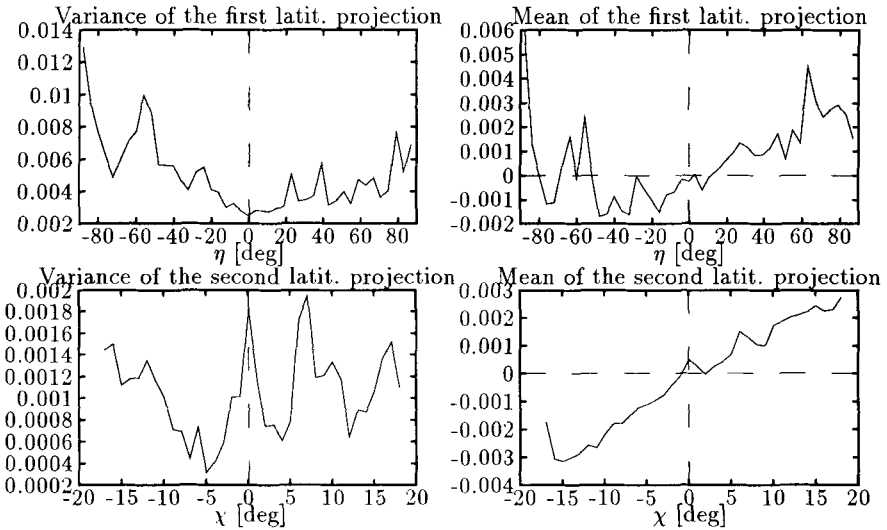


Fig. 7. The variance (top left) and the mean (top right) of the first latitudinal projection for the real fixated sequence. The minimum of the variance gives the angle η and the mean for this η gives the torsion. The variance (bottom left) and the mean (top right) of the second latitudinal projection for the real fixated sequence giving the angle χ at the minimum of the variance and the inverse of the time to collision, respectively.

5 Conclusion

It was proven in the past that fixation reduces the number of unknowns in the structure from motion problem from five to four. We showed in this paper that fixation can further simplify the computation of 3D-motion parameters from a monocular sequence. Appropriate projections of the spherical flow field enable the decoupling of the motion parameters in two groups: The first contains the azimuthal angle of the translation and the torsion. The second parameter group contains the polar angle of the translation direction and the time to collision to the fixated target. Two 1D searches in each of the two projections yield all

four unknowns. In contrast to other algorithms, we do not make any use of the measurements of camera movements necessary for fixation. We assume for the second search that the focus of expansion is inside the field of view. If this is not the case we can apply a two-unknowns nonlinear minimization or even better carry out a correcting saccade that will bring the focus of expansion inside the field of view. The algorithm was tested in three real world sequences fixated off-line or actively. Without applying any special method for accurate computation of the flow we obtained very promising results.

References

1. R.H.S. Carpenter. *Movements of the Eyes*. Pion Press, London, 1988.
2. 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*, 12:155–163., 1994.
3. A. Bandopadhyay and D.H. Ballard. Egomotion perception using visual tracking. *Computational Intelligence*, 7:39–47, 1990.
4. Y. Aloimonos, I. Weiss, and A. Bandyopadhyay. Active Vision. In *Proc. Int. Conf. on Computer Vision*, pages 35–54, London, UK, June 8–11, 1987.
5. C. Fermüller and Y. Aloimonos. The role of fixation in visual motion analysis. *International Journal of Computer Vision*, 11:165–186, 1993.
6. M.A. Taalebinezhaad. Direct recovery of motion and shape in the general case by fixation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14:847–853, 1992.
7. D. Raviv and M. Herman. A unified approach to camera fixation and vision-based road following. *IEEE Trans. Systems, Man, and Cybernetics*, 24:1125–1141, 1994.
8. I. Thomas, E. Simoncelli, and R. Bajcsy. Spherical retinal flow for a fixating observer. In *Proc. IEEE Workshop on Visual Behaviors*, pages 37–41, 1994.
9. M. Tistarelli and G. Sandini. Dynamic aspects in active vision. *CVGIP: Image Understanding*, 56:108–129, 1992.
10. K. Daniilidis. Computation of 3D-motion parameters using the log-polar transform. In V. Hlavac et al. (Ed.), *Proc. Int. Conf. Computer Analysis of Images and Patterns CAIP, Prag*, pages 82–89, 1995.
11. M.J. Barth and S. Tsuji. Egomotion determination through an intelligent gaze control strategy. *IEEE Trans. Systems, Man, and Cybernetics*, 23:1424–1432, 1993.
12. V. Sundareswaran, P. Bouthemy, and F. Chaumette. Active camera self-orientation using dynamic image parameters. In *Proc. Third European Conference on Computer Vision*, pages 111–115. Stockholm, Sweden, May 2–6, J.O. Eklundh (Ed.), Springer LNCS 800, 1994.
13. S. Maybank. *Theory of Reconstruction from Image Motion*. Springer-Verlag, Berlin et al., 1993.
14. A.D. Jepson and D.J. Heeger. Subspace methods for recovering rigid motion II: Theory. Technical Report RBCV-TR-90-36, University of Toronto, 1990.