Robust Instantaneous Rigid Motion Estimation

Karl Pauwels and Marc M. Van Hulle*
Laboratorium voor Neuro- en Psychofysiologie
K.U.Leuven, Belgium
{*karl.pauwels, marc.vanhulle*}@*med.kuleuven.ac.be*

Abstract

A novel, nonlinear algorithm is introduced for the estimation of rigid camera motion from instantaneous velocity measurements in the calibrated case. It is shown that by minimizing an approximation to the image-reprojection error rather than the actual error, as done by the optimal algorithms, the proposed algorithm achieves largely increased robustness to local minima at the cost of a slightly decreased accuracy.

1. Introduction

Visual motion is one of the more important sensory cues that are used by humans to guide behavior or to navigate a dynamical environment. The instantaneous velocity or optic flow field contains a tremendous amount of information related to the three dimensional (3D) structure of the environment and to independently moving objects that may be present. Knowledge of the egomotion or self-motion of the observer is a necessary prerequisite to obtain this valuable information. Its determination is, however, nontrivial and an active topic of research.

In this paper we propose a novel, nonlinear method for the computation of egomotion from instantaneous velocity measurements, assuming a calibrated camera. The largely reduced variance in the parameter estimates obtained by nonlinear algorithms as compared to linear ones, has been shown on several occasions [2, 10, 12] and is crucial for practical applications where very noisy optic flow serves as input. However, contrary to linear methods, nonlinear algorithms are prone to local minima. These local minima are intrinsic to the egomotion problem itself and depend,

among others, on the feature locations [11]. Additional local minima are induced by the use of, possibly nonlinear, estimation methods from the domain of robust statistics that should be employed in the presence of independently moving objects.

We propose a novel, unbiased egomotion estimation method that provides a balance between accuracy, and robustness to local minima. By means of extensive simulations we demonstrate that it outperforms linear algorithms with respect to variance of the estimates and nonlinear methods with respect to robustness to local minima. The reason for this is the approximation of the error function used by the optimal algorithms, which allows for simplified updating rules in our algorithm.

2. Problem statement

Under a static environment assumption, the motion of all points in space, relative to a coordinate system centered in the nodal point of the observer's eye, is determined by the translational velocity $\mathbf{t} = (t_x, t_y, t_z)^{\mathrm{T}}$ and rotational velocity $\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)^{\mathrm{T}}$ of the moving observer. The 3D velocity $\mathbf{v} = (v_x, v_y, v_z)^{\mathrm{T}}$ of a point $\mathbf{x} = (x, y, z)^{\mathrm{T}}$ is then [6]:

$$\mathbf{v} = -\mathbf{t} - \boldsymbol{\omega} \times \mathbf{x} \, . \tag{1}$$

Under perspective projection and assuming, without loss of generality, a focal length equal to unity, these 3D motion vectors are transformed into a two dimensional (2D) velocity or optic flow field. At feature location $\mathbf{x} = (x, y, 1)^{\mathrm{T}}$, the observed flow $\mathbf{u}(\mathbf{x}) = (u_x, u_y)^{\mathrm{T}}$ equals:

$$\mathbf{u}(\mathbf{x}) = d(\mathbf{x})A(\mathbf{x})\mathbf{t} + B(\mathbf{x})\boldsymbol{\omega} + \mathbf{n}(\mathbf{x}) \, , \tag{2}$$

where

$$A(\mathbf{x}) = \begin{bmatrix} -1 & 0 & x \\ 0 & -1 & y \end{bmatrix} , \tag{3}$$

$$B(\mathbf{x}) = \begin{bmatrix} xy & -1-x^2 & y \\ 1+y^2 & -xy & -x \end{bmatrix} . \tag{4}$$
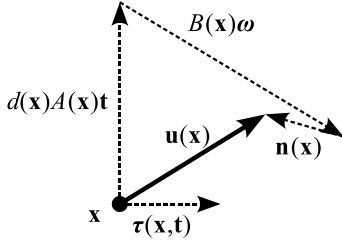
Figure 1. Optic flow components.

The observed flow consists of three parts: a component due to the observer's translation (which also depends on the inverse depth $d(\mathbf{x}) = 1/z$), a component due to the observer's rotation, and the noise $\mathbf{n}(\mathbf{x}) = (n_x, n_y)^T$, which is assumed to be independently (not necessarily identically) distributed. These different components are illustrated in Fig. 1. Also indicated is $\boldsymbol{\tau}(\mathbf{x}, \mathbf{t})$, a unit length vector orthogonal to the translational component of the flow:

$$\boldsymbol{\tau}(\mathbf{x}, \mathbf{t}) = \frac{1}{\|A(\mathbf{x})\mathbf{t}\|} \left(\Big[A(\mathbf{x})\mathbf{t}\Big]_y, -\Big[A(\mathbf{x})\mathbf{t}\Big]_x \right)^T, \quad (5)$$

where $[\mathbf{p}]_x$ and $[\mathbf{p}]_y$ refer to the x- and y-components of the vector $\mathbf{p}$ respectively. When depth is eliminated from Eq. (2), the well-known bilinear constraint [1] on translation and rotation is obtained:

$$\Big(\mathbf{u}(\mathbf{x}) - B(\mathbf{x})\boldsymbol{\omega}\Big)^T \Big(\|A(\mathbf{x})\mathbf{t}\|\, \boldsymbol{\tau}(\mathbf{x}, \mathbf{t})\Big) = 0. \quad (6)$$

This particular notation was chosen since it highlights that the constraint is weighted by $\|A(\mathbf{x})\mathbf{t}\|$. This weight term renders the constraints much simpler algebraically [1] (see also Eq. 17). In the absence of prior knowledge and assuming isotropic, additive noise, it is however unwarranted to weigh the constraints in this arbitrary fashion. Instead, the unweighted constraints should be used [12]:

$$\hat{\mathbf{t}}, \hat{\boldsymbol{\omega}} = \mathrm{argmin}_{\mathbf{t}, \boldsymbol{\omega}} \sum_i \mathbf{r}^2(\mathbf{x}_i), \quad (7)$$

where:

$$\mathbf{r}(\mathbf{x}_i) = \Big(\mathbf{u}(\mathbf{x}_i) - B(\mathbf{x}_i)\boldsymbol{\omega}\Big)^T \boldsymbol{\tau}(\mathbf{x}_i, \mathbf{t}), \quad (8)$$

the normalized, orthogonal, deviations from the epipolar lines. The estimates obtained from Eq. (7) minimize the least-squares image-reprojection error [9]. Algorithms that operate on this error function are commonly referred to as 'optimal' [2, 12].

3. Previous algorithms

A wide variety of egomotion-estimation methods have been proposed in the past. One of the first algorithms was introduced by Bruss and Horn [1] and consists of a straight-forward minimization of the bilinear constraints (Eq. 6) using nonlinear optimization techniques. Heeger and Jepson (H&J) [4] proposed a method to compute the heading (normalized translation) without iterative numerical optimization. Their linear subspace method is based on the construction of a set of constraint vectors that are independent of camera rotation. Another linear algorithm was recently proposed by Ma *et al.* [7] and is conceptually similar to methods that operate on the discrete epipolar constraint. The heading estimates computed with this algorithm have been shown to be identical to those obtained with H&J but the rotation estimates are better.

The heading estimates obtained with the aforementioned algorithms are all systematically biased. Different bias correction procedures can be found in the literature. Kanatani (KAN) [5] introduced a method that subtracts an estimate of the bias from the solution. A second correction procedure was introduced more recently by Maclean (MAC) [8] as an adaptation to H&J. Contrary to KAN, the latter method does not require an estimate of the noise variance. The method we propose uses a bias correction procedure similar to MAC (see Section 4).

An optimal, nonlinear algorithm was introduced by Chiuso *et al.* (CHI) [2]. This algorithm involves a sequence of fixed-point iterations where each part of the sequence involves the solution of a linear least-squares problem. Chiuso *et al.* proposed iterating between estimates of $\mathbf{t}$ and $\{d(\mathbf{x}), \boldsymbol{\omega}\}$. Since a spherical projection model was used in their formulation and the other algorithms assume a traditional pin-hole model, we modified the formulation and implemented the algorithm as follows. Starting from an initial heading estimate $\hat{\mathbf{t}}$, a rotation estimate $\hat{\boldsymbol{\omega}}$ is obtained as the linear least-squares solution to the system of Eq. (8). Using both estimates, the least-squares relative inverse depth estimates $\hat{d}(\mathbf{x}_i)$ are obtained as follows:

$$\hat{d}(\mathbf{x}_i) = \frac{\Big(\mathbf{u}(\mathbf{x}_i) - B(\mathbf{x}_i)\hat{\boldsymbol{\omega}}\Big)^T A(\mathbf{x}_i)\hat{\mathbf{t}}}{\big\|A(\mathbf{x}_i)\hat{\mathbf{t}}\big\|^2}. \quad (9)$$

Next, the estimates $\{\hat{d}(\mathbf{x}), \hat{\boldsymbol{\omega}}\}$ are used to compute a new translation estimate as the linear least-squares solution to the system of Eq. (2). After normalization of this translation estimate, the sequence is repeated until the estimates converge.

Zhang and Tomasi [12] (ZHA) introduced a second optimal algorithm which is very similar to CHI. The rotation $\hat{\boldsymbol{\omega}}$ and relative inverse depth estimates $\hat{d}(\mathbf{x})$ are obtained in the same way but the heading estimate $\hat{\mathbf{t}}$ is updated using a Gauss-Newton optimization scheme. The update $\Delta\mathbf{t}$ is obtained as the linear least-squares solution to the following

set of equations:

$$\left(\mathbf{u}(\mathbf{x}_i) - B(\mathbf{x}_i)\hat{\boldsymbol{\omega}}\right)^\mathrm{T} \boldsymbol{\tau}(\mathbf{x}_i, \hat{\mathbf{t}}) =$$
$$\frac{\delta}{\delta \mathbf{t}}\left[\left(\hat{d}(\mathbf{x}_i)A(\mathbf{x}_i)\mathbf{t}\right)^\mathrm{T} \boldsymbol{\tau}(\mathbf{x}_i, \hat{\mathbf{t}})\right]\Delta \mathbf{t} , \qquad (10)$$

combined with the additional constraint $\hat{\mathbf{t}}^\mathrm{T}\Delta\mathbf{t} = 0$, which ensures a unit length heading estimate.

4. Proposed method

A straightforward fixed-point minimization of the bilinear constraints (Eq. 6), conceptually similar to CHI, can be achieved by alternately estimating the heading assuming the rotation is correct, and estimating the rotation assuming the heading is correct, using linear least-squares for both estimates. There are two problems associated with this approach.

First of all, as explained in Section 1, the constraints are weighted by $\|A(\mathbf{x}_i)\mathbf{t}\|$ and consequently not every flow vector is considered equal. This can be resolved by using weighted least squares, with weights equal to $1/\|A(\mathbf{x}_i)\hat{\mathbf{t}}\|$. Here, $\hat{\mathbf{t}}$ is the heading estimate obtained in the previous iteration. Note the difference with a direct minimization of Eq. (7) which is normalized by the heading estimate itself and thus highly nonlinear.

A second, more difficult problem, concerns the systematic bias to which heading estimates obtained in this way are prone. The correction procedure introduced by Maclean [8] for H&J, can be adapted to work directly on the bilinear constraints. Without loss of generality, in the absence of rotation (or after removal of the estimated rotation), these constraints can equivalently be written as:

$$\mathbf{t}^\mathrm{T}\left(\mathbf{x}_i \times \mathbf{u}(\mathbf{x}_i)\right) = 0 , \qquad (11)$$

with $\mathbf{u}(\mathbf{x}) = (u_x, u_y, 0)^\mathrm{T}$ (note the change in notation as compared to Eq. 2). The least-squares heading estimate $\hat{\mathbf{t}}$ is equal to the eigenvector corresponding to the smallest eigenvalue of the constraint vectors' scatter matrix S. In practice, the flow is noisy and the constraints are:

$$\left(\mathbf{x}_i \times \mathbf{u}(\mathbf{x}_i)\right) + \left(\mathbf{x}_i \times \mathbf{n}(\mathbf{x}_i)\right) = \mathbf{s}_{\mathbf{u}_i} + \mathbf{s}_{\mathbf{n}_i} , \qquad (12)$$

where $\mathbf{n}(\mathbf{x}_i) = (n_{x_i}, n_{y_i}, 0)^\mathrm{T}$ independently distributed, isotropic, noise with zero mean and standard deviation equal to σ_i. The scatter matrix $\tilde{S}$ of these noisy constraint vectors is now equal to:

$$\tilde{S} = \sum_i (\mathbf{s}_{\mathbf{u}_i} + \mathbf{s}_{\mathbf{n}_i})(\mathbf{s}_{\mathbf{u}_i} + \mathbf{s}_{\mathbf{n}_i})^\mathrm{T} ,$$
$$= \sum_i (\mathbf{s}_{\mathbf{u}_i}\mathbf{s}_{\mathbf{u}_i}^\mathrm{T} + \mathbf{s}_{\mathbf{u}_i}\mathbf{s}_{\mathbf{n}_i}^\mathrm{T} + \mathbf{s}_{\mathbf{n}_i}\mathbf{s}_{\mathbf{u}_i}^\mathrm{T} + \mathbf{s}_{\mathbf{n}_i}\mathbf{s}_{\mathbf{n}_i}^\mathrm{T}) ,$$

$$= S + \sum_i (\mathbf{s}_{\mathbf{u}_i}\mathbf{s}_{\mathbf{n}_i}^\mathrm{T} + \mathbf{s}_{\mathbf{n}_i}\mathbf{s}_{\mathbf{u}_i}^\mathrm{T}) + N . \qquad (13)$$

To show that the constraints are biased, we next investigate the expected value of the scatter matrix $\tilde{S}$. The expected value of the second component of $\tilde{S}$ equals the null matrix. The first component, S, corresponds to the scatter matrix of the noise-less constraint vectors, and as a result:

$$\mathrm{E}\{\tilde{S}\} = S + \mathrm{E}\{N\} . \qquad (14)$$

We next investigate the scatter matrix of the noise component. For simplicity the noise variance is assumed constant and equal to σ^2 in the remainder. Note that this assumption is not required and that the analysis remains valid when the noise variance is different at each feature location. The variance σ^2 should then be replaced by the mean variance $\langle \sigma_i^2 \rangle$ in Eqs. (15) and (16). The expected value of the noise scatter matrix can easily be shown to be:

$$\mathrm{E}\{N\} = m\sigma^2 \begin{bmatrix} 1 & 0 & -\langle x \rangle \\ 0 & 1 & -\langle y \rangle \\ -\langle x \rangle & -\langle y \rangle & \langle x^2 \rangle + \langle y^2 \rangle \end{bmatrix} ,$$
$$= m\sigma^2 W , \qquad (15)$$

with m the number of flow vectors. We now have:

$$\mathrm{E}\{\tilde{S}\} = S + m\sigma^2 W . \qquad (16)$$

It can be observed from Eq. (16) that, in the presence of noise, the expected value of the constraints vectors' scatter matrix does not correspond to the scatter matrix of the noise-less constraint vectors. Consequently, the eigenvectors of this matrix and thus the heading estimate will be biased. However, since this bias depends on the feature locations only, a technique called *whitening* can be used to correct for it [8]. Instead of computing the eigenvectors of $\tilde{S}$, the eigenvectors of $W^{-1/2}\tilde{S}W^{-1/2} = W^{-1/2}SW^{-1/2} + m\sigma^2 I_3$ are determined. Since adding a scaled version of the identity matrix I_3 does not change the eigenvectors, the latter will be unbiased. Choosing $\mathbf{e}$, the eigenvector corresponding to the smallest eigenvalue λ of $(W^{-1/2}SW^{-1/2})$, gives: $W^{-1/2}SW^{-1/2}\mathbf{e} = \lambda\mathbf{e}$. The heading estimate is then $\hat{\mathbf{t}} = W^{-1/2}\mathbf{e}$. Since $W^{-1}S(W^{-1/2}\mathbf{e}) = \lambda(W^{-1/2}\mathbf{e})$, this is an eigenvector of $W^{-1}S$. However, since S is composed of noise-free constraint vectors, pre-multiplying S with W^{-1} does not change the eigenvector corresponding to the smallest eigenvalue (which is zero).

Once the unbiased translation estimate $\hat{\mathbf{t}}$ is obtained, it is used to compute the weights, and the rotation is computed as the weighted least-squares solution to the bilinear constraints, alternatively written as:

$$\left((\hat{\mathbf{t}} \times \mathbf{x}_i) \times \mathbf{x}_i\right)^\mathrm{T} \boldsymbol{\omega} - \hat{\mathbf{t}}^\mathrm{T}\left(\mathbf{x}_i \times \mathbf{u}(\mathbf{x}_i)\right) = 0 . \qquad (17)$$

3

In summary, the proposed Fixed-Point algorithm with bias Correction (FPC) minimizes the bilinear constraints (Eq. 6) by alternately solving the reduced problems using weighted linear least-squares. The heading estimate at the previous iteration is used to update the weights. Since the heading constraints are whitened, unbiased estimates are obtained. Note that the least-squares weights also need to be used when averaging over feature locations in the determination of the bias correction matrix W (see Eq. 15).

5. Simulations

A first series of analyses compares the bias and variance of the estimates obtained with the egomotion estimation algorithms explained in Sections 3 and 4 and a second series investigates the nonlinear algorithms only, in terms of their robustness to local minima.

5.1. Bias/variance

We compare H&J, KAN, MAC, Z&T and CHI to the proposed method FPC[1] in terms of the bias and variance of the heading estimates. We did not include the algorithms by Ma *et al.* [7] (the heading estimates of which are identical to H&J's) and by Bruss and Horn [1] (which is a biased, nonlinear algorithm). The rotation estimates are not analyzed since the bias is entirely due to heading estimation and the heading estimates can be visualized and interpreted more easily.

We examined the same configuration of translation and rotation as Zhang and Tomasi [12], namely a translation and rotation direction equal to $[4, -3, 5]$ and $[-1, 2, 0.5]$ respectively. The rotation rate was fixed to $0.23°/$frame and the translational magnitude was chosen so that the speeds of the translational and rotational flow components are identical in the center of the random depth cloud [10]. In each experiment, 500 feature locations were randomly chosen and uniformly distributed over the image. The focal length was set to unity. The 3D positions of the features were uniformly distributed between 1 and 4 units of focal length. Independently distributed Gaussian noise was added to the flow vectors. The signal-to-noise ratio (SNR), defined as: $\left(\mathrm{E}\|\mathbf{u}\|^2/\mathrm{E}\|\mathbf{n}\|^2\right)^{1/2}$, was varied between 10 and 30. For each algorithm, 100 trials were performed, in which the feature locations, depth and noise were randomized. For the nonlinear algorithms, 15 heading initializations, evenly spread on the unit sphere, were used and the solution with the smallest error was retained. Figure 2 contains the heading estimates obtained with all algorithms, for a SNR equal

---

[1] We used implementations provided by Tian *et al.* [10] for H&J and KAN, our own implementations for MAC, CHI and FPC and an implementation provided by Dr. T. Zhang for Z&T.

to 10. The field of view (FOV) was equal to $50°$ and $150°$ in the top, respectively the bottom row. The estimates were mapped to the upper hemisphere and projected onto a circle. The dashed cross marks the true heading. Example flow fields are shown in the rightmost column. For each algorithm and noise level, the bias $e_{\mathbf{t}}$, defined as the angular difference between the mean heading estimate and the actual heading, and a 95% confidence cone θ_α ($\alpha = 0.05$), closely related to the variance of the estimates, were computed using techniques from the domain of spherical statistics [3]. Contrary to the bias/variance measures used in previous studies [10, 12], this more sophisticated analysis clearly brings out the bias in the estimates obtained with some of the algorithms. Table 1 contains both measures for all algorithms, SNRs and FOVs. The bias is underlined in the table if the mean heading estimate is contained within the confidence cone (unbiased). With FOV equal to $50°$, this is the case for all algorithms and noise levels except, as expected, for H&J. We also see that the variance in the estimates is much smaller for the nonlinear algorithms than for the linear ones, as observed in other studies [10, 12]. Note that bias and variance for Z&T and CHI are identical for all configurations. With FOV equal to $150°$, the heading estimates for KAN are clearly biased, which was also observed by Zhang and Tomasi [12]. A performance reduction due to an increased FOV was also reported in [10] for the case of decreasing image resolution (the number of features is fixed) and straight-ahead translation (the translation used here has a relatively large z-component).

In summary, we observe a significant improvement in performance of the nonlinear algorithms with respect to the linear. The algorithm by Maclean outperforms the other linear algorithms, illustrating the quality of this particular bias correction procedure which is also used by our algorithm. Finally, the results obtained with the proposed algorithm are very similar to those obtained with the other nonlinear algorithms, except for a slightly increased variance of the estimates.

5.2. Local minima

Although the use of nonlinear over linear algorithms significantly reduces the variance of the heading estimates, the former may occasionally get stuck in a local minimum corresponding to a grossly erroneous heading estimate. The error surface associated with the egomotion problem becomes flatter in a situation of lateral translation and the number of local minima increases when the feature locations are clustered together, even in the noiseless case [11]. Using this information we constructed a particularly difficult scenario that enabled us to investigate the robustness to local minima of the nonlinear algorithms: Z&T, CHI and FPC. The egomotion consisted of a translation and rotation direc-
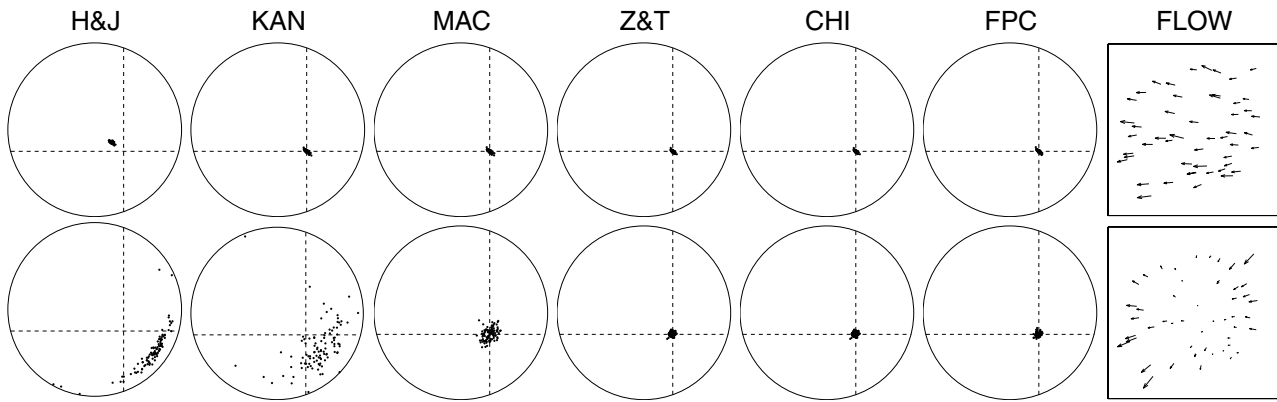
Figure 2. Heading estimates obtained with six different algorithms on 100 random trials. The FOV equals $50°$ and $150°$ in the top, respectively the bottom row (the SNR was equal to 10 for both). An example flow field (subsampled and magnified 10 times) is shown in the rightmost column.

Table 1. Bias and variance of the heading estimates obtained with all six algorithms tested for different FOVs and SNRs. The bias is underlined if the mean heading estimate falls within the 95% confidence cone.

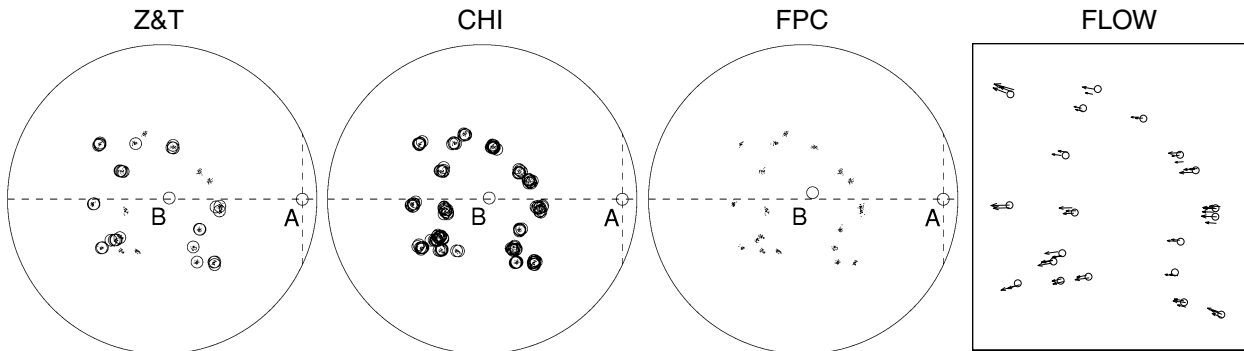|		LINEAR						NONLINEAR					
		H&J		KAN		MAC		Z&T		CHI		FPC	
FOV	SNR	$e_\mathbf{t}$	θ_α	$e_\mathbf{t}$	θ_α	$e_\mathbf{t}$	θ_α	$e_\mathbf{t}$	θ_α	$e_\mathbf{t}$	θ_α	$e_\mathbf{t}$	θ_α
50	30	2.21	0.12	<u>0.03</u>	0.13	<u>0.03</u>	0.12	<u>0.04</u>	0.10	<u>0.04</u>	0.10	<u>0.02</u>	0.10
	20	4.94	0.19	<u>0.05</u>	0.19	<u>0.04</u>	0.19	<u>0.06</u>	0.15	<u>0.06</u>	0.15	<u>0.03</u>	0.16
	10	17.51	0.34	<u>0.19</u>	0.42	<u>0.12</u>	0.42	<u>0.11</u>	0.32	<u>0.11</u>	0.32	<u>0.06</u>	0.35
150	30	9.38	0.51	2.89	0.66	<u>0.18</u>	0.43	<u>0.11</u>	0.16	<u>0.11</u>	0.16	<u>0.11</u>	0.17
	20	18.58	0.91	6.56	1.19	<u>0.27</u>	0.67	<u>0.17</u>	0.25	<u>0.17</u>	0.25	<u>0.18</u>	0.25
	10	35.73	3.68	22.22	4.47	<u>0.51</u>	1.55	<u>0.51</u>	0.55	<u>0.51</u>	0.55	<u>0.46</u>	0.55

Figure 3. Small circles in the leftmost figures correspond to heading estimates obtained with the nonlinear algorithms when initializing with 50,000 distinct random headings. The global minimum is labeled A and the local minimum due to the bas-relief ambiguity is labeled B. Feature locations are indicated with small black dots. The rightmost figure contains the noiseless flow field used (subsampled and magnified 10 times). In this figure, the small circles indicate the feature cluster centers.

5

tion equal to $[1, 0, 0.1]$ and $[0, 1, 0]$ respectively. The depth, translation and rotation magnitudes were chosen as in Section 5.1 and the FOV was set equal to $100°$. A total of 500 features were used but, contrary to Section 5.1, they were not uniformly distributed in the image. Instead, their locations were drawn from 20 spatially distinct clusters, the centers of which were uniformly distributed. The cluster centers are indicated with circles in the rightmost figure of Fig. 3. Also shown in this figure is the (subsampled and scaled) flow field used. No noise was added to the computed flow vectors. Each algorithm was run with 50,000 heading initializations, randomly sampled from the unit sphere, and was allowed 1,000 iterations to reach convergence. This large number of initializations was used to give a detailed account of the algorithms' behaviors over the entire heading space. The first three figures of Fig. 3 contain the estimated headings (black circles) together with the normalized feature locations $\mathbf{x}/\|\mathbf{x}\|$. As before, the dashed cross marks the actual heading. It is apparent from these figures that both Z&T and CHI suffer from a large number of local minima, located near clusters of image pixels, whereas FPC does not suffer from this problem at all and only finds one additional local minimum near the image center (labeled B in Fig. 3). A similar minimum is also found by the other algorithms and is a consequence of the so-called bas-relief ambiguity (for details see [2, 11]). Techniques have been proposed to discriminate between these two strong minima [2]. Note that the second minimum (B) found by FPC differs slightly from that found by Z&T and CHI. This is due to the bias correction which slightly modifies the error function that is used by the algorithm. The results were similar when varying degrees of noise were added to the optic flow.

Figure 4 contains the heading initializations (gray dots) for Z&T and CHI that result in local minima different from A or B. As expected, the locations of these initializations are related to the feature locations (black dots). It is notable that the feature clusters have a rather large spatial extent over which they exert their influence and interactions between clusters are clearly visible.

6. Conclusion

We introduced a novel, unbiased nonlinear algorithm for the computation of egomotion from optic flow in the calibrated case. Contrary to previously introduced, optimal algorithms, the proposed method minimizes an *approximation* to the least-squares image-reprojection error, which allows for much simpler updating rules. By means of extensive simulations, it was shown that this greatly increases the robustness of the algorithm to local minima at the cost of only a slight increase in the variance of the heading estimates. Furthermore, it was shown that this variance is still much smaller than that obtained with linear algorithms.
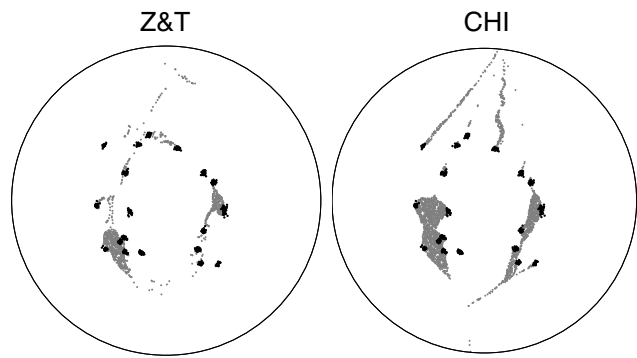

Figure 4. Heading initializations that result in undesired local minima (gray dots), in relation to feature locations (black dots).

References

[1] A. Bruss and B. Horn. Passive navigation. *Computer Graphics and Image Processing*, 21:3–20, 1983.

[2] A. Chiuso, R. Brockett, and S. Soatto. Optimal structure from motion: Local ambiguities and global estimates. *International Journal of Computer Vision*, 39(3):195–228, 2000.

[3] N. I. Fisher, T. Lewis, and B. J. J. Embleton. *Statistical analysis of spherical data*. Cambridge university press Cambridge, 1987.

[4] D. Heeger and A. Jepson. Subspace methods for recovering rigid motion I: Algorithm and implementation. *International Journal of Computer Vision*, 7(2):95–117, 1992.

[5] K. Kanatani. 3-D interpretation of optical flow by renormalization. *International Journal of Computer Vision*, 11(3):267–282, 1993.

[6] H. Longuet-Higgins and K. Prazdny. The interpretation of a moving retinal image. *Proc. R. Soc. Lond. Biol.*, 208:385–397, 1980.

[7] Y. Ma, J. Kosecka, and S. Sastry. Linear differential algorithm for motion recovery: A geometric approach. *International Journal of Computer Vision*, 36(1):71–89, 2000.

[8] W. MacLean. Removal of translation bias when using subspace methods. In *Proceedings of the Eight International Conference on Computer Vision*, pages 753–758, Corfu, Greece, September 20-27 1999. IEEE Computer Society Press.

[9] J. Oliensis. The least-squares error for structure from infinitesimal motion. *International Journal of Computer Vision*, 61(3):1–41, 2005.

[10] Y. Tian, C. Tomasi, and D. Heeger. Comparison of approaches to egomotion computation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 315–320, San Francisco, USA, June 1996.

[11] T. Xiang and L. Cheong. Understanding the behavior of SFM algorithms: A geometric approach. *International Journal of Computer Vision*, 51(2):111–137, 2003.

[12] T. Zhang and C. Tomasi. On the consistency of instantaneous rigid motion estimation. *International Journal of Computer Vision*, 46:51–79, 2002.