

CDS

TECHNICAL MEMORANDUM NO. CIT-CDS 94-005

March 7, 1994

**“Robust and Efficient Recovery of Rigid Motion
from Subspace Constraints Solved using Recursive
Identification of Nonlinear Implicit Systems”**

Stefano Soatto and Pietro Perona

**Control and Dynamical Systems
California Institute of Technology
Pasadena, CA 91125**

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 07 MAR 1994		2. REPORT TYPE		3. DATES COVERED 00-00-1994 to 00-00-1994	
4. TITLE AND SUBTITLE Robust and Efficient Recovery of Rigid Motion from Subspace Constraints Solved using Recursive Identification of Nonlinear Implicit Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) California Institute of Technology, Control and Dynamical Systems, Pasadena, CA, 91125				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Robust and Efficient Recovery of Rigid Motion from Subspace Constraints Solved using Recursive Identification of Nonlinear Implicit Systems*

Stefano Soatto and Pietro Perona
California Institute of Technology 116-81
Pasadena-CA 91125
soatto@systems.caltech.edu

Keywords: Dynamic Vision, Rigid Motion Estimation, Nonlinear Identification, Implicit Extended Kalman Filter, Exterior Differential Systems

Abstract

The problem of estimating rigid motion from projections may be characterized using a nonlinear dynamical system, composed of the rigid motion transformation and the perspective map. The time derivative of the output of such a system, which is also called the “motion field”, is bilinear in the motion parameters, and may be used to specify a subspace constraint on either the direction of translation or the inverse depth of the observed points. Estimating motion may then be formulated as an optimization task constrained on such a subspace. Heeger and Jepson [5], who first introduced this constraint, solve the optimization task using an extensive search over the possible directions of translation.

We reformulate the optimization problem in a systems theoretic framework as the identification of a dynamic system in exterior differential form with parameters on a differentiable manifold, and use techniques which pertain to nonlinear estimation and identification theory to perform the optimization task in a principled manner. The general technique for addressing such identification problems [14] has been used successfully in addressing other problems in computational vision [13, 12].

The application of the general method [14] results in a recursive and pseudo-optimal solution of the motion problem, which has robustness properties far superior to other existing techniques we have implemented.

By releasing the constraint that the visible points lie in front of the observer, we may explain some psychophysical effects on the nonrigid percept of rigidly moving shapes.

Experiments on real and synthetic image sequences show very promising results in terms of robustness, accuracy and computational efficiency.

1 Visual motion estimation from a dynamic model

Let a scene be represented by a set of N feature points in 3D space moving rigidly with respect to the viewer; the visual motion problem is *defined* by the rigidity constraint and the perspective

*Research funded by the California Institute of Technology, ONR grant N00014-93-1-0990 and an AT&T Foundation Special Purpose grant. This work is registered as CDS Technical Report CIT-CDS 94-005, California Institute of Technology, January 1994 – revised February 1994

projection equations. If $\mathbf{X}_i \doteq [X_i \ Y_i \ Z_i]^T$ are the coordinates of the i^{th} point and $\mathbf{x}_i \doteq [x_i \ y_i]^T$ the corresponding projection, we may write

$$\begin{cases} \dot{\mathbf{X}}_i = \Omega \wedge \mathbf{X}_i + V & \mathbf{X}(0) = \mathbf{X}_0 \\ \mathbf{x}_i = \pi(\mathbf{X}_i) + n_i & \forall i = 1 : N \end{cases} \quad (1)$$

where n_i represents an error in measuring the position of the projection of the point i and π represents an ideal perspective projection¹. Solving the visual motion problem consists in estimating V, Ω from all the visible points, i.e. reconstructing the input of the above system from its noisy output. We show that it is possible to invert the above system using a technique which has been recently introduced in [14] for identifying systems in Exterior Differential form [2] with parameters on a topological manifold.

The scheme is motivated by the work of Heeger and Jepson [6, 5] and may be considered as a recursive solution of their task using methods which pertain to the field of nonlinear estimation and identification theory. As a result, the minimization task which is the core of the subspace methods for recovering rigid motion, needs not to be performed by brute force search, as it is done in [5]. Instead, an Implicit Extended Kalman Filter (IEKF) [9, 3, 8, 14] is in charge of estimating the motion parameters recursively according to nonlinear prediction error criteria (for an introductory treatment of Prediction Error Methods (PEM) in a linear context, see for example [16]). As an effect, our method exploits in a pseudo-optimal manner the information coming from a long stream of images, making the scheme robust and computationally efficient.

2 Motion reconstruction via (least squares) inversion constrained on subspaces

Consider the following expression of the “motion field”, i.e. the first derivative of the output of the model (1).

$$\dot{\mathbf{x}}(t) = \left[\frac{1}{Z} \mathcal{A}(\mathbf{x}) \mid \mathcal{B}(\mathbf{x}) \right] \begin{bmatrix} V(t) \\ \Omega(t) \end{bmatrix} \quad (2)$$

where

$$\mathcal{A} \doteq \begin{bmatrix} 1 & 0 & -x \\ 0 & 1 & -y \end{bmatrix} \quad \mathcal{B} \doteq \begin{bmatrix} -xy & 1+x^2 & -y \\ -1-y^2 & xy & x \end{bmatrix}. \quad (3)$$

If we observe enough points, we have an overdetermined system which we may solve for the motion parameters in a least-squares sense. Call $\mathcal{C}_i \doteq \left[\frac{1}{Z_i} \mathcal{A}_i \mid \mathcal{B}_i \right]$, we have

$$\begin{bmatrix} \hat{V}(t) \\ \hat{\Omega}(t) \end{bmatrix} = \begin{bmatrix} \vdots \\ \mathcal{C}_i \\ \vdots \end{bmatrix}^\dagger \begin{bmatrix} \dot{\mathbf{x}}_1 \\ \dot{\mathbf{x}}_2 \\ \vdots \\ \dot{\mathbf{x}}_N \end{bmatrix} \doteq \mathcal{C}^\dagger \dot{\mathbf{x}}$$

where the symbol \dagger denotes the pseudo-inverse. Note that \mathcal{C}_i depends on the depth of the point Z_i , which we do not know. By substituting the above expression into eq. (2), we have an *implicit constraint* on Z_i [5], which consists of imposing that $\dot{\mathbf{x}}$ is the null space of the orthogonal complement

¹More articulated camera models may be employed. However, we do not address them here.

of the range of \mathcal{C} . We may try to approximate this constraint in a least squares sense by solving w.r.t Z_i the following nonlinear optimization problem:

$$\hat{Z}_i = \arg \min_{Z_i} \|\dot{\mathbf{x}} - \mathcal{C}(Z) \begin{bmatrix} \hat{V} \\ \hat{\Omega} \end{bmatrix}\| = \|(I - \mathcal{C}\mathcal{C}^\dagger)\dot{\mathbf{x}}\|. \quad (4)$$

i.e. we are looking for $Z_i \forall i = 1 \dots n$ such that $\dot{\mathbf{x}}$ is the null space of the orthogonal complement of the range of \mathcal{C} . If \mathcal{C} was invertible, the above constraint would be satisfied trivially for all motions. However, when $2N > 3$, $\mathcal{C}\mathcal{C}^\dagger$ has rank at most three, and hence $(I - \mathcal{C}\mathcal{C}^\dagger) \neq 0$.

2.1 Recovery of direction of translation

Note that the minimization described above is performed with respect to the depth of each point in space. However, the role of depth and translation may be interchanged, as it is evident from the structure of the matrix \mathcal{C} . We may therefore “pseudo-invert” the system above with respect to depth and rotation and then perform the minimization with respect to the direction of translation in \mathbf{S}^2

(the two-sphere of radius one). For each point i we have $\dot{\mathbf{x}}_i(t) = [\mathcal{A}_i(\mathbf{x}_i)V(\theta, \phi) \mid \mathcal{B}_i(\mathbf{x})] \begin{bmatrix} \frac{1}{Z(t)_i} \\ \Omega(t) \end{bmatrix}$,

where $V \in \mathbf{S}^2$ is represented in local coordinates as $V(\theta, \phi)$. If we observe N points we may write $\dot{\mathbf{x}} = \tilde{\mathcal{C}}(\theta, \phi) [\frac{1}{Z_1}, \dots, \frac{1}{Z_N}, \Omega]^T$, where

$$\tilde{\mathcal{C}}(\theta, \phi) \doteq \begin{bmatrix} \mathcal{A}_1 V & & & \mathcal{B}_1 \\ & \ddots & & \vdots \\ & & \mathcal{A}_N V & \mathcal{B}_N \end{bmatrix}.$$

Now, proceeding in a similar way as before, we could seek for θ, ϕ which satisfy the following subspace algebraic constraint:

$$\left[I - \tilde{\mathcal{C}} \left(\tilde{\mathcal{C}}^T \tilde{\mathcal{C}} \right)^{-1}_{|\theta, \phi, \mathbf{x}} \tilde{\mathcal{C}}^T \right] \dot{\mathbf{x}} = 0. \quad (5)$$

Note that we are trying to “adapt” the orthogonal complement of $\tilde{\mathcal{C}}$, which is highly structured as a function of θ, ϕ , until a given vector $\dot{\mathbf{x}}$ is its null space. Heeger and Jepson [5] solve this problem by minimizing the two-norm of the above constraint using an extensive search over θ, ϕ , or a sampling of the sphere. This procedure does not exploit any of the geometric structure of the problem. Furthermore it does not take into account the measurement noise, which enters into the minimization in a highly structured fashion, and is computationally expensive. Temporal coherence of motion is also not taken into account: at each step we want to exploit all the processing performed at the previous time instant and update recursively the motion estimates.

The method for performing the minimization task described above in a principled way is presented in section 3: we rephrase the problem as the identification of an exterior differential system with parameters on the two-sphere. The method outputs motion estimates together with their reliability in the form of the second order statistics of the estimation error. Such an error may be used in subsequent modules for estimating structure.

2.2 Recovery of the mean distance

In many applications it is of interest to estimate the average distance of an object from the camera (position of the centroid). For this case, it is sufficient to consider the minimization in eq. (4) when

$Z_i = Z_c \forall i$; Z_c is the distance of the centroid. The solution to such a problem is analogous to the recovery of translation, and will be presented in section 3.

2.3 Recovery of rotation and depth

Once the direction of translation has been recovered, we may derive the rotational velocity and inverse depth in a least-squares fashion from their definition:

$$\begin{bmatrix} \frac{1}{Z_1} \\ \vdots \\ \frac{1}{Z_N} \\ \Omega \end{bmatrix} = \tilde{C}^\dagger(\theta, \phi)\dot{\mathbf{x}}.$$

The motion estimates may be fed, together with the variance of their estimation error, into a recursive structure from motion module which processes motion error, such as for example [11, 15].

3 Solving the subspace optimization via identifying an exterior differential system with parameters on a differentiable manifold

Let us define $\alpha \doteq [\theta, \phi]^T$; \mathbf{x}_i are measured up to some error, $\mathbf{y}_i \doteq \mathbf{x}_i + n_i$ $n_i \in \mathcal{N}(0, R_{n_i})$, which induces an error in the derivative: $\mathbf{y}_i' = \dot{\mathbf{x}}_i + n_i'$. Call \mathbf{x} the column vector obtained by stacking the components i , similarly with $\dot{\mathbf{x}}$. Now define $\tilde{C}^\perp(\mathbf{x}, \alpha) \doteq \left[I - \tilde{C} \left(\tilde{C}^T \tilde{C} \right)^{-1}_{|\alpha, \mathbf{x}} \tilde{C}^T \right]$. Finally the subspace constraint (5) may be written as $\tilde{C}^\perp(\mathbf{x}, \alpha)\dot{\mathbf{x}} = 0$. Now

$$\begin{cases} \tilde{C}^\perp(\mathbf{x}, \alpha)\dot{\mathbf{x}} = 0 & V(\alpha) \in \mathbf{S}^2 \\ \mathbf{y}_i \doteq \mathbf{x}_i + n_i & \forall i \end{cases}$$

represents a system in Exterior Differential Form. *Solving for translation is equivalent to identifying the above exterior differential system with parameters on a differentiable manifold* (the sphere in this case) from the noisy data \mathbf{y} .

We have addressed this problem using the general methods presented in [14]. The solution is given by the simple iteration

Prediction step

$$\begin{cases} \hat{\alpha}(t+1|t) = \hat{\alpha}(t|t) & \hat{\alpha}(0|0) = \alpha_0 \\ P(t+1|t) = P(t|t) + R_\alpha(t) & P(0|0) = P_0 \end{cases}$$

Update step

$$\begin{cases} \hat{\alpha}(t+1|t+1) = \hat{\alpha}(t+1|t) + \\ \quad L(t+1)\tilde{C}^\perp(\mathbf{y}(t), \alpha(t+1|t))\mathbf{y}' \\ P(t+1|t+1) = \\ \quad \Gamma(t+1)P(t+1|t)\Gamma^T(t+1) + \\ \quad \underline{L(t+1)D_+(t)R_n(t+1)D_+^T(t)L^T(t+1)} \end{cases}$$

where

$$\begin{cases} L(t+1) = P(t+1|t)\tilde{C}^{\perp T}(t+1)\Lambda^{-1}(t+1) \\ \Lambda(t+1) = \tilde{C}^{\perp}(t+1)P(t+1|t)\tilde{C}^{\perp T}(t+1) + D_+(t)R_n(t+1)D_+^T(t) \\ \Gamma(t+1) = I - L(t+1)\tilde{C}^{\perp}(t+1) \\ D_+(t) \doteq \left(\frac{\partial \tilde{C}^{\perp} \hat{\mathbf{x}}}{\partial \mathbf{x}(t)} \right)_{|\mathbf{y}(t)} \end{cases}$$

as from the appendix A of [14].

3.1 Observability/identifiability of the method

In order to be able to assess the convergence of the above scheme, we must prove its observability/identifiability. It can be shown, using the analysis of Heeger and Jepson [7], that the scheme described above is identifiable under general position conditions. Note that the analysis in [7] is carried out with different motivations; however, when the results are cast into the above estimation/identification framework, they allow inferring the identifiability of the method.

3.2 Enforcing rigid motion: the positive depth constraint

When estimating motion from visible points, we must enforce the fact that the measured points are *in front of the observer*. This may be easily done in the prediction step by computing the mean distance of the centroid, as indicated above, and checking whether it is positive. If it is not, the prediction is reflected on the sphere (the diametral point of the state space sphere is chosen as the prediction).

When we do not impose such a constraint, the filter may converge to a rigid motion which corresponds to points moving behind the observer, and are therefore not physically realizable. However, if we allow such condition to happen by releasing the positive depth constraint, and then feed the estimate to a structure estimation step, such as for example a simple Extended Kalman Filter [10, 11, 15] initialized with points at positive depth, the result is a *rubbery interpretation of structure* which has been observed also in psychophysical experiments.

The geometric interpretation of the rubbery percept is illustrated in figure 1. Note that both affine 3D motion and similarity transformations viewed under projection admit a geometric invariant, which is the absolute conic [4]. On the contrary the orientation (determinant of the transformation) is not invariant under projection.

4 Experimental assessment

We have experimented the scheme on real and noisy synthetic image sequences. For the same data set used in [15], the scheme proves far more robust to the effect of measurement noise. Convergence is reached from *arbitrary initial condition* and noise in the image plane coordinates up to 10 pixel std. The scheme works also with higher noise levels when properly initialized.

The estimate of the two components of the direction of translation with 8 pixel std noise is shown in figure 4, together with the estimation error. An estimate for more usual error levels (one pixel std) is reported in figure 3. In both cases the positive depth constraint has been enforced. The estimates of rotational velocity are plotted in figure 4.

A typical plot of the residual function, which is the value of the constraint (5) as a function of θ, ϕ , is shown in figure 5. The bright area indicate a small residual value. The black asterisk indicates the position of the motion (in the local coordinates of the sphere of directions of translation) which generated the residual.

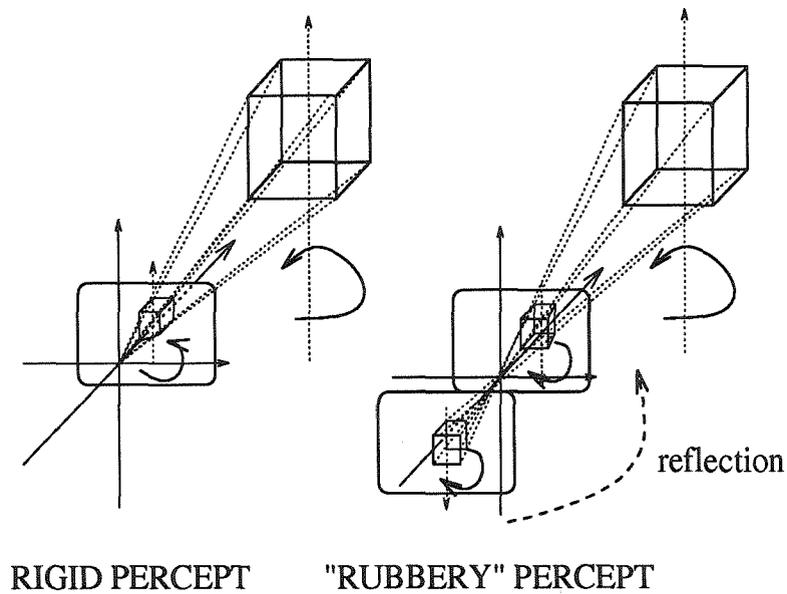


Figure 1: Geometric interpretation of the “rubbery” perception: motion is estimated without imposing the positive depth constraint; this may result in a motion estimate which is compatible with a rigid structure interpretation behind the observer. Once such a structure is reflected in front of the observer, it gives rise to the perception of a rubbery structure rotating in the opposite direction of the true one.

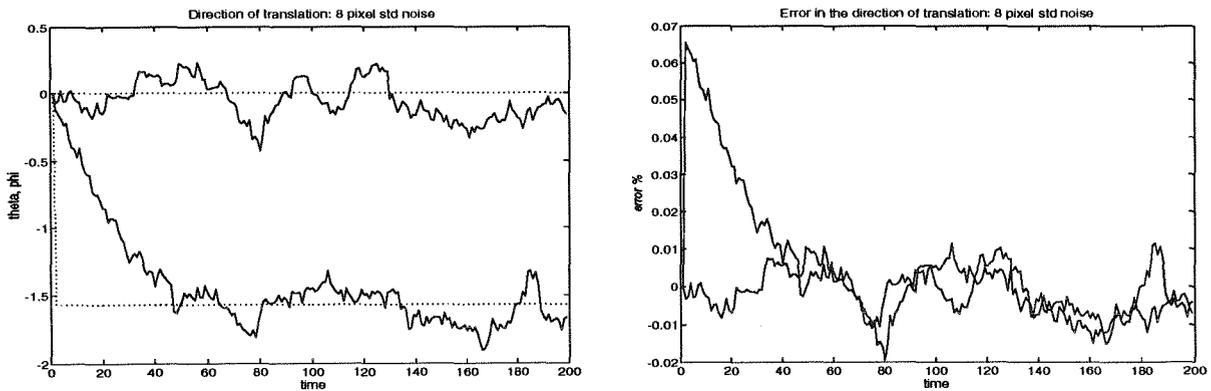


Figure 2: (Left) Estimates of the two components of the direction of translation. The error in the image plane measurements had 8 pixel standard deviation. The initial conditions were zero for both components. The ground truth is in dotted lines (Right) Estimation error for the direction of translation. With noise of 8 pixel std in the data, the estimates are still within 10 % of the true value. The positive depth constraint has been enforced.

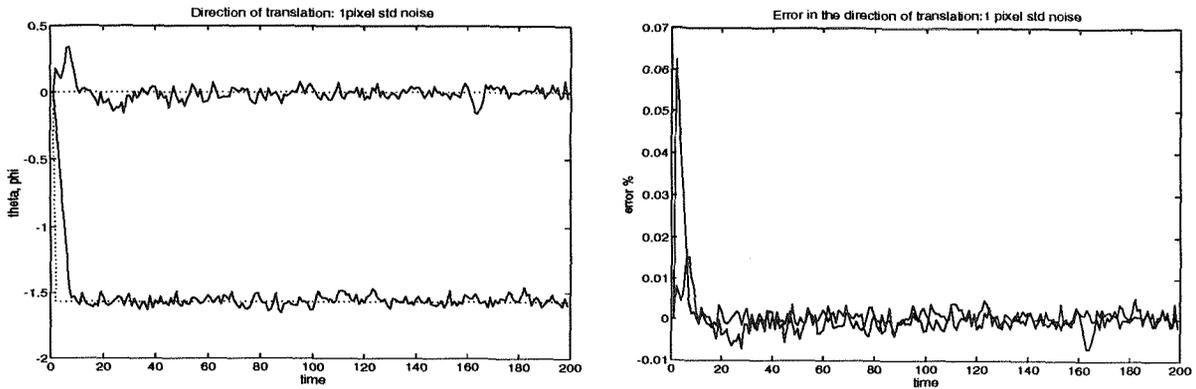


Figure 3: Estimates and errors for the direction of translation when the noise in the image plane has a standard deviation of 1 pixel (according to the performance of common optical flow/feature tracking schemes). Note that convergence is reached from zero initial condition in about 10 steps.

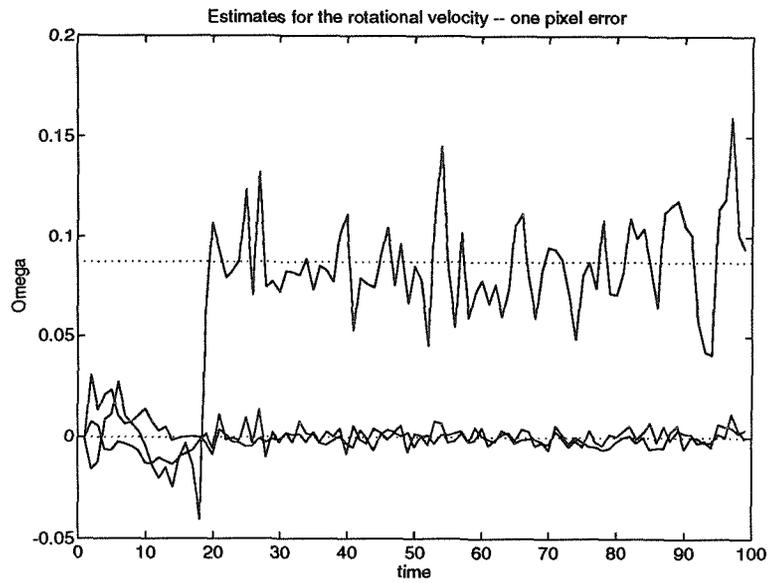


Figure 4: Estimates for the components of rotational velocity

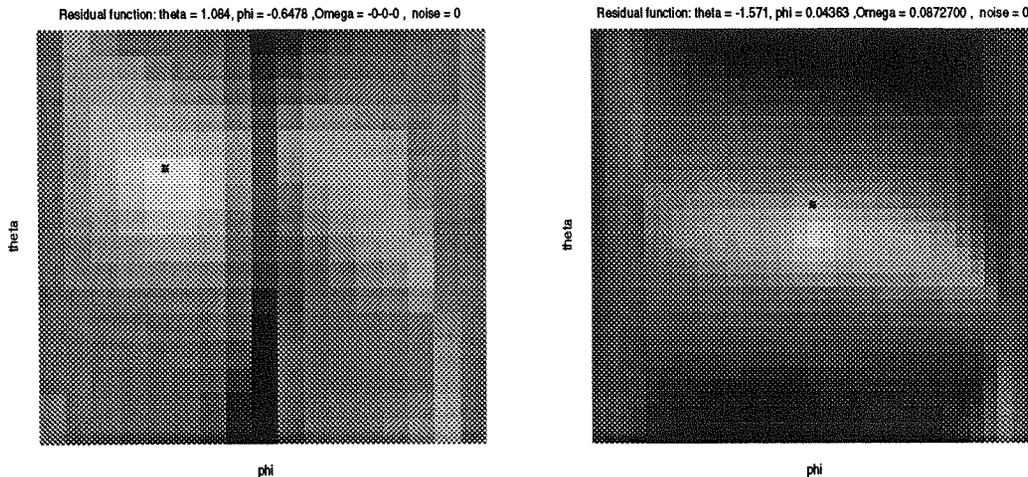


Figure 5: Plots of the residual function. The local coordinates of the sphere of directions of translation is plotted. Bright regions denote small residuals. The black asterisk is the “true” motion which generated the residual. Note that for small rotation (left) the minimum of the residual coincides with the true motion. When translation is large (right) the Euler step approximation is no longer valid, and the minimum moves from the true location.

It is noted that the minimum of the residual is displaced from the true motion when the norm of the rotational velocity is large. This is due to the fact that we approximate the velocity of the projected points (motion field) with first differences; the approximation is good as long as $R = e^{\Omega \Delta} = I + \Omega \Delta$, i.e. as long as the norm of translation is small.

Note the presence of local minima, as it may be seen from the mesh plot of the cost function (figure 6). The filter may temporarily enter into one of these. In figure 4 (right) we show the temporary convergence of the filter to a local minimum. In figure 4 (left) we show the convergence to the “rubbery motion interpretation” when the positive depth constraint is released.

In the following figure 8 we show the convergence of the filter to the rubbery interpretation (left) and rigid motion (right) plotted on the image of the cost function. When the positive depth constraint is not enforced the filter may converge either to the rigid or to the rubbery interpretation (figure 9 left). However, when imposing the positive depth constraint, the estimate is reflected onto the correct rigid interpretation (figure 9 right, see also figure 10 right for the state estimates).

When we feed the motion estimation to a structure from motion module estimating motion error, and initialized with points at positive depth, we may observe either a rigid set of points which move according to the correct motion (a top view of the points is shown in figure 12 left) or to a “rubbery” percept (figure 12 right). This is in accordance with the experience in psychophysical experiments.

4.1 Comparison with the essential filter

The filter proposed in this paper proves far less sensitive to noise in the measurements and to the initial conditions when compared to the essential filter [13]. In particular, for 20 observed points, the essential filter converges for initial conditions within 30 %, while the subspace filter converges

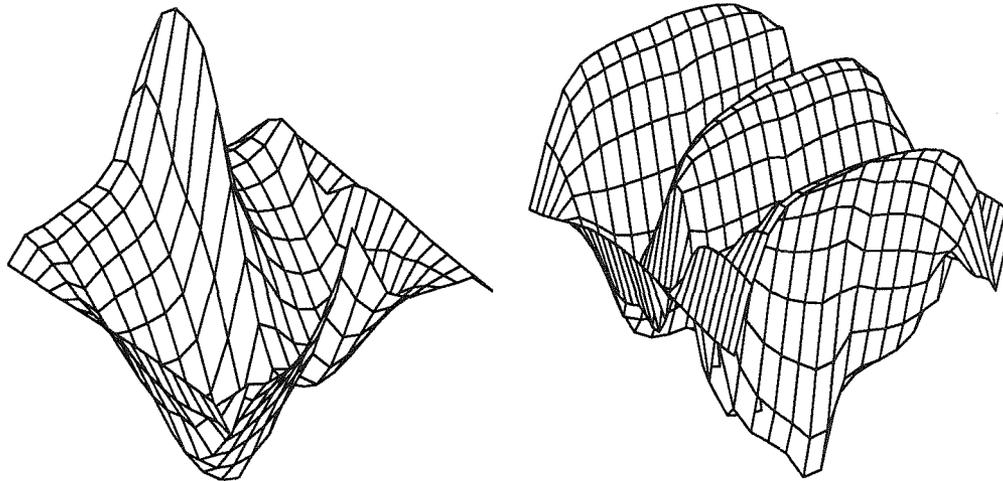


Figure 6: (Left) Mesh plots of some typical residual functions: note the presence of local minima.

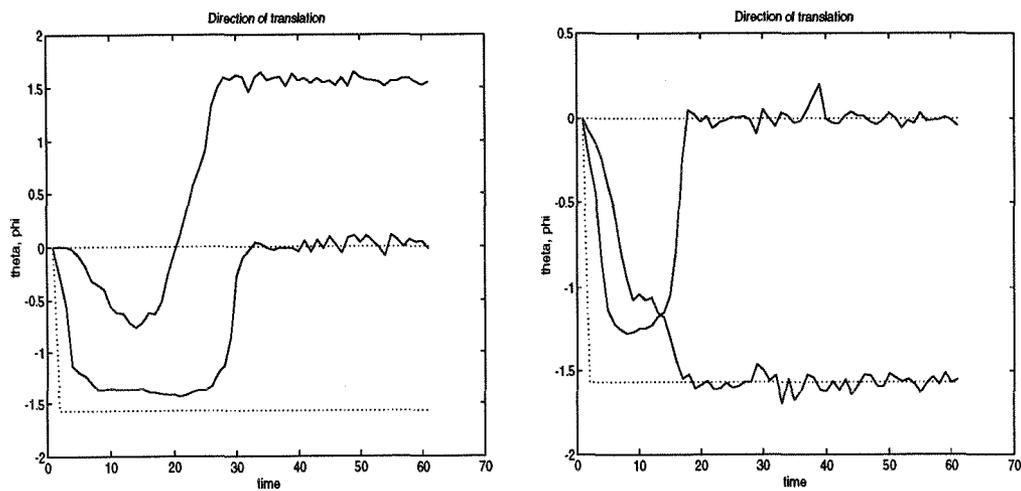


Figure 7: (Left) convergence to the local minimum corresponding to the rubbery interpretation when the positive depth constraint is not enforced. (Right) convergence to a local minimum and then to the correct rigid motion when the positive depth constraint is enforced.

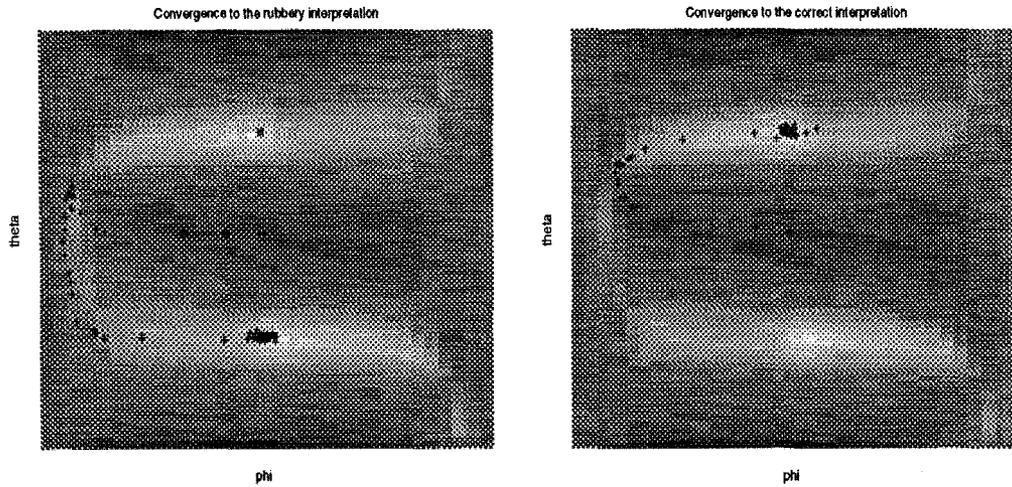


Figure 8: Convergence to the “rubbery interpretation” (left) versus convergence to the rigid motion interpretation

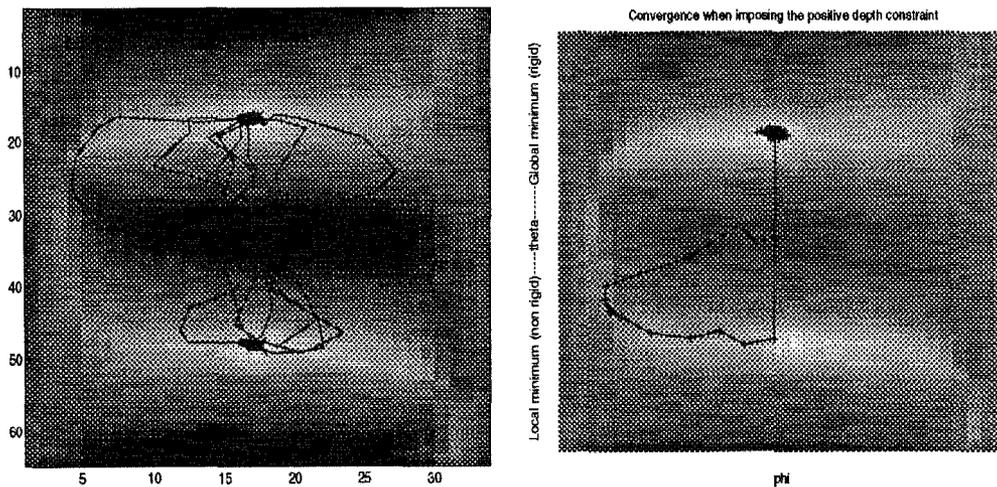


Figure 9: (Left) Convergence when not imposing the positive depth constraint: the filter may converge to either the correct rigid interpretation (top) or to the local minimum corresponding the “rubbery” interpretation (bottom). However, when imposing the positive depth constraint (right), the filter only converges to the correct rigid motion interpretation.

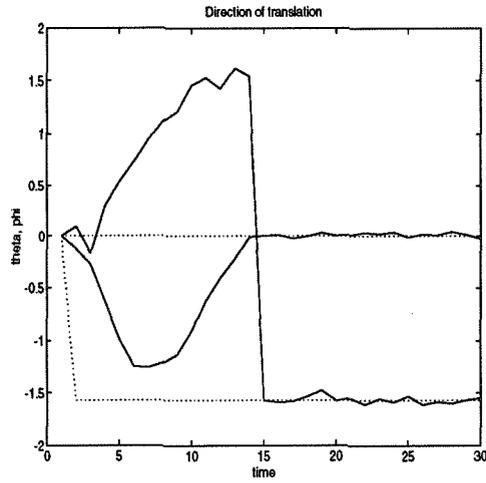


Figure 10: Convergence of the filter when imposing the positive depth constraint

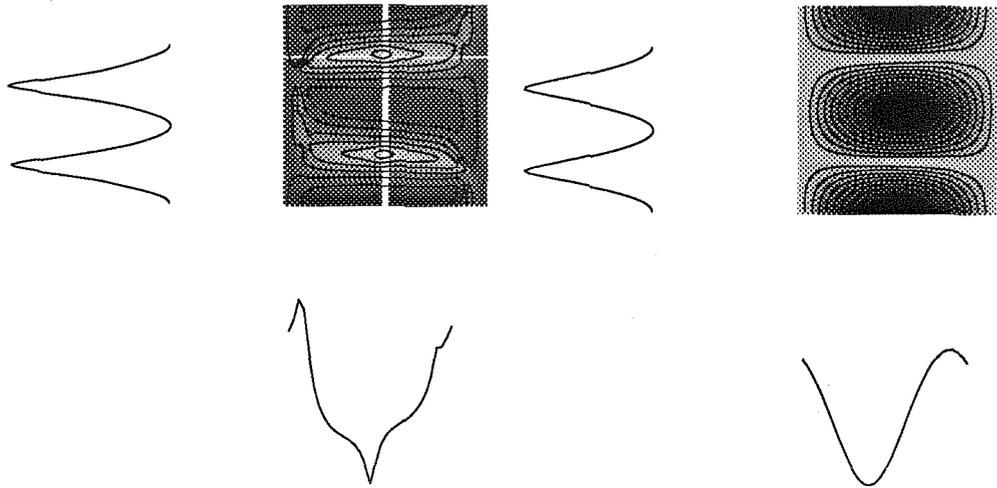


Figure 11: Comparison of the contour plots of the residual function for the subspace filter (left) and the essential filter (right). The slice of the residual surface is plotted along the two dimensions in the coordinates of the minimum. Note that the essential filter seems to have a simpler residual, with no local minima except for the one corresponding to the rubbery structure. However, for the essential filter this is only a two-dimensional slice of the more complicated residual which is on a five-dimensional space.

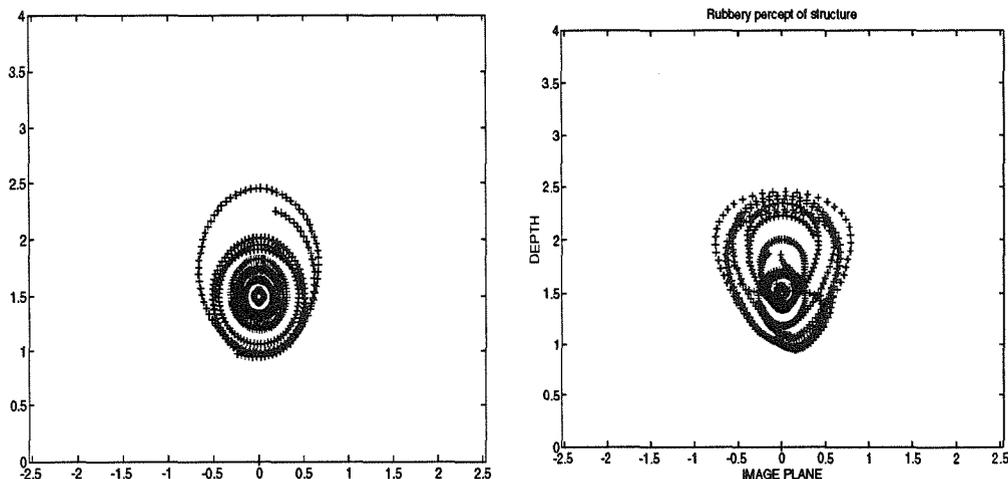


Figure 12: Convergence of a structure from motion module to a rigid interpretation of structure (left) or to a rubbery object rotating in the opposite direction (right). The plots show a top view of the points, with the image plane on the low end.

from any initial condition. The essential filter is faster in converging, reaching regime in 5 to 15 frames, while the subspace filter takes on average 10-20 frames. However, the essential filter is more sensitive to noise, and the subspace filter may tolerate up to 5 times more error on the measured image plane coordinates, i.e. up to more than 10 pixel std. The contour plots of the residual function for the subspace filter and the essential filter may be compared in figure 11. Note that for the essential filter only a two-dimensional slice of the five-dimensional residual is plotted.

In the essential filter the positive depth constraint is encoded directly in the definition of the state space manifold (the essential manifold). The convergence of the essential filter is illustrated in fig. 13: on the left the convergence is shown when starting from the rubbery motion interpretation and imposing positive depth. On the right the positive depth constraint has been released (equivalently reflections are allowed in the essential manifold), and therefore we may observe occasionally convergence to the local minimum corresponding to the rubbery interpretation.

4.2 Experiments with real image sequences

We have tested the scheme on real image sequences: the noise level achieved by the most common feature tracking/optical flow techniques is easily handled by the filter. As an example we report here the filter estimates for the rocket scene, for comparison with [13]. Due to the fact that the filter takes about 20 frames to converge, we have doubled the sequence and used one run as initial condition for the second run, which is displayed in image 14. We are in the process of testing the scheme on longer image sequences.

5 Conclusions

We have formulated a new recursive scheme for estimating rigid motion under perspective via identifying a dynamic model in exterior differential form. The motivation comes from Heeger and

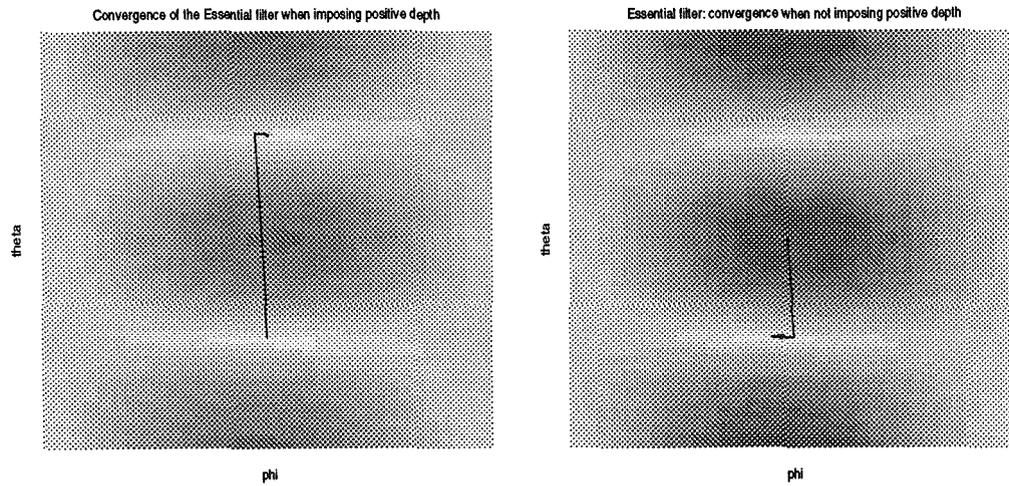


Figure 13: Convergence of the essential filter: (left) the filter imposes automatically the positive depth constraint; even starting from the rubbery state, the filter switches to the correct estimate. (Left) Releasing the positive depth constraint, it is possible for the filter to converge to the rubbery interpretation.

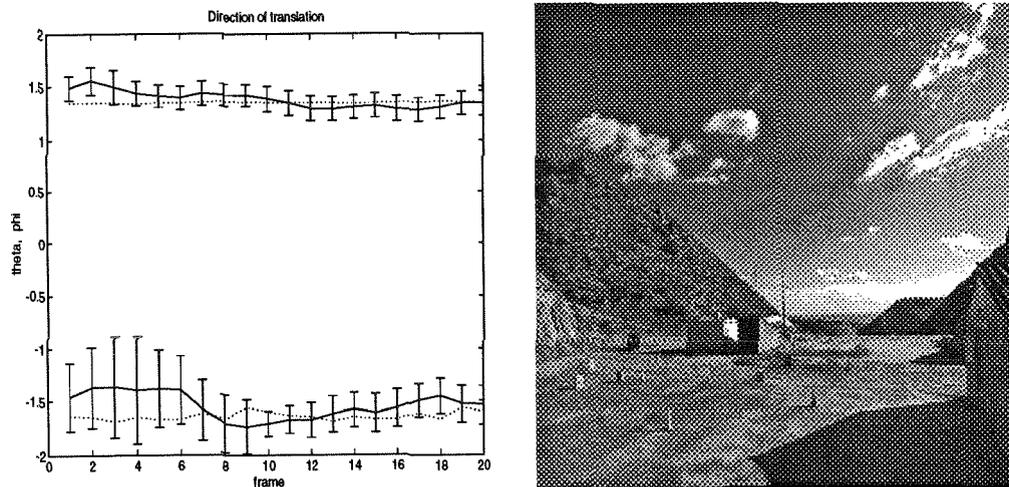


Figure 14: (Left) Estimate of the direction of translation for the rocket scene. (Right) One image of the rocket scene.

Jepson [5], who propose to view motion estimation as an optimization problem constrained to subspace. They solve the minimization by extensive search.

Using results from nonlinear estimation and identification theory, we formulate a motion estimator which is fast, computationally efficient, accurate and more robust than any recursive motion estimation scheme we have implemented. Extensive experiments have been performed that highlight such features.

Acknowledgements

We wish to thank Doug Shy for his help in implementing the scheme using the tensors toolbox of Matlab.

References

- [1] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow techniques. RPL-TR 9107, Queen's University Kingston, Ontario, Robotics and perception laboratory, 1992. Also in Proc. CVPR 1992, pp 236-242.
- [2] Bryant, Chern, Goldberg, and Goldsmith. *Exterior Differential Systems*. Mathematical Research Institute. Springer Verlag, 1992.
- [3] R.S. Bucy. Non-linear filtering theory. *IEEE Trans. A.C. AC-10*, 198, 1965.
- [4] O. Faugeras. *Three dimensional vision, a geometric viewpoint*. MIT Press, 1993.
- [5] D. Heeger and A. Jepson. Subspace methods for recovering rigid motion i: algorithm and implementation. *Int. J. Comp. Vision vol. 7 (2)*, 1992.
- [6] D. Heeger and A. Jepson. Subspace methods for recovering rigid motion i: algorithm and implementation. RBCV TR-90-35, University of Toronto – CS dept., November 1990. Revised July 1991.
- [7] D. Heeger and A. Jepson. Subspace methods for recovering rigid motion ii: theory. RBCV TR-90-35, University of Toronto – CS dept., November 1990. Revised July 1991.
- [8] A.H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, 1970.
- [9] R.E. Kalman. A new approach to linear filtering and prediction problems. *Trans. of the ASME-Journal of basic engineering.*, 35-45, 1960.
- [10] L. Matthies, R. Szeliski, and T. Kanade. Kalman filter-based algorithms for estimating depth from image sequences. *Int. J. of computer vision*, 1989.
- [11] J. Oliensis and J. Inigo-Thomas. Recursive multi-frame structure from motion incorporating motion error. *Proc. DARPA Image Understanding Workshop*, 1992.
- [12] S. Soatto, R. Frezza, and P. Perona. Recursive estimation of camera motion from uncalibrated image sequences. *Technical Report CIT-CDS 94-003, California Institute of Technology*, 1994.
- [13] S. Soatto, R. Frezza, and P. Perona. Motion estimation on the essential manifold. *Proc. of the ECCV 94 – To appear in “Lecture Notes in Computer Sciences”, Springer Verlag*, May 1994.

- [14] S. Soatto, R. Frezza, P. Perona, and G. Picci. Dynamic estimation of rigid motion from perspective views via recursive identification of exterior differential systems with parameters on a topological manifold. *Technical Report CIT-CDS 94-004*, California Institute of Technology, Feb. 1994.
- [15] S. Soatto, P. Perona, R. Frezza, and G. Picci. Recursive motion and structure estimation with complete error characterization. In *Proc. IEEE Comput. Soc. Conf. Comput. Vision and Pattern Recogn.*, pages 428–433, New York, June 1993.
- [16] T. Soderstorm and P. Stoica. *System Identification*. Prentice Hall, 1989.