Recursive Structure and Motion Estimation Based on Hybrid Matching Constraints*

Anders Heyden, Fredrik Nyberg, and Ola Dahl

Applied Mathematics Group
School of Technology and Society, Malmö University, Sweden
{heyden,fredrik.nyberg,ola.dahl}@ts.mah.se

Abstract. Motion estimation has traditionally been approached either from a pure discrete point of view, using multi-view tensors, or from a pure continuous point of view, using optical flow. This paper builds upon a novel framework of hybrid matching constraints for motion estimation, combining the advantages of both discrete and continuous methods. We will derive both bifocal and trifocal hybrid constraints and use them together with a structure estimate based on filtering techniques. A feedback from the structure estimate will be used to further refine the motion estimate. This gives a complete iterative structure and motion estimation scheme. Its performance will be demonstrated in simulated experiments.

1 Introduction

Structure from motion is one of the central problems in computer vision and has been extensively studied during the last decades. The objective is to compute the motion of the camera and the structure of the scene from a number of its two-dimensional images. The standard method is to first estimate the motion of the camera, based on matching tensors, obtained from corresponding points in an image sequence. Then, given the motion of the camera, the structure of the scene is obtained as a sparse set of 3D-points, which can be used as a starting point for surface estimation or texture mapping, cf. [1].

The most common method for estimation of the matching constraints is based on a discrete setting, where e.g. the fundamental (or essential) matrix is estimated between an initial view and another view obtained later in the sequence, c.f. [2]. In order to deal with long image sequences several matching constraints are then pasted together, giving a consistent set of matching constraints from which the motion of the camera can be estimated, cf. [3]. Another approach, closely related to optical flow, is to use a continuous setting and estimate the motion parameters from continuous time matching constraints based on image point positions and velocities, c.f. [4,5,6].

Attempts has been made to combine the discrete and the continuous methods. In [7], a number of differential matching constraints were derived and an algorithm for updating the fundamental matrix along an image sequence was

* This work was partially supported by the SRC project 621-2002-4831.

indicated. However, no experimental evidence or details about the algorithm were given. These ideas have been taken up and extended to a form suitable for on-line structure and motion estimation, cf. [8], and preliminary results on an on-line structure and motion system have been reported in [9].

The main purpose of the present paper is to develop suitable methods for *on-line recursive structure and motion estimation* for long image sequences. By this, we mean methods that can update a current estimate of the position of the camera and the structure of the scene, when a new image in the sequence becomes available. Such methods have been presented in [10] and [11], where in both cases complex non-linear procedures are used to update the structure. We will propose a novel method where the motion estimation is separated from the structure estimation, enabling simpler and more stable update schemes.

In this work we derive and utilize two types of matching constraint, called *hybrid matching constraints* (HMC), for the estimation and update of motion parameters. The first one is an extension of the epipolar constraint to a *hybrid epipolar constraint*, where both corresponding points in two images as well as their motion in the second image are used. The second one is an extension of the trifocal constraint to a *hybrid trifocal constraint*, where both corresponding points in three images as well as their motion in the third image are used. These hybrid constraints will enable us to update the current motion estimate *linearly* based on at least *three corresponding points*. This will be shown theoretically, by proving the exact number of linearly independent constraints obtained from each corresponding point. The HMC are here fused with a continuous-discrete extended Kalman filter for the state estimation, in order to construct an algorithm for recursive estimation of both structure and motion. We also introduce a linear reprojection error constraint, where feedback of the structure estimates is used to recursively obtain corrections to the motion estimates. This constraint connects the structure and motion estimation processes in a consistent way, and is significantly improving the performance of the method.

2 Problem Description

2.1 Camera Model and Notation

We assume the standard pinhole camera model,

$$\lambda \mathbf{x} = P\mathbf{X} \quad , \quad (1)$$

where $\mathbf{x}$ denotes homogeneous image coordinates, P the camera matrix, $\mathbf{X}$ homogeneous object coordinates and λ a scale factor. The camera matrix P is usually written as $P = K[R \mid -b]$, where K denotes the intrinsic parameters and (R, b) the extrinsic parameters (R being an orthogonal matrix). We will from now on assume that the camera is calibrated, i.e. K is known and that the image coordinates have been transformed such that P can be written as $P = [R \mid -b]$. When several images of the same point are available, (1) can be written as

$$\begin{cases} \lambda(t)\mathbf{x}(t) = P(t)\mathbf{X}, t \in [0, T] & \text{or} \\ \lambda_i\mathbf{x}_i = P_i\mathbf{X}, i = 1, \dots, M \end{cases} \quad (2)$$

in the continuous time case and the discrete time case respectively. The camera matrix P is assumed to have the form

$$P(t) = [R(t) \mid -b(t)] \quad \text{or} \quad P_i = [R_i \mid -b_i] , \quad (3)$$

in the continuous case and in the discrete case respectively. We furthermore assume that the object coordinate system has been chosen such that $R(0) = R_1 = I$ and $b(0) = b_1 = 0$, implying that $P(0) = P_1 = [I \mid 0]$.

2.2 Problem Formulation

A structure and motion estimation problem can now be formulated as the task of estimating both the structure $\mathbf{X}$ in (2) and the motion parameters $R(t)$ and $b(t)$ in (3) at the time t , given the set of perspective measurements $\mathfrak{M}_t = \{\mathbf{x}(t_i) \mid \forall i : t_i \leq t\}$. A *recursive* structure and motion problem can be formulated as given an estimate of the structure and motion parameters up to time t , update this estimate based on measurements obtained up to time $t + \Delta t$.

2.3 Discrete Matching Constraints

The *discrete matching constraints* are obtained by using the discrete version of (2), for several different i and eliminating the object coordinates $\mathbf{X}$ from the resulting equations. In the case of two views we obtain the well-known *epipolar constraint*

$$\mathbf{x}_1^T E \mathbf{x}_2 = 0, \quad \text{with} \quad E = R^T \hat{b} , \quad (4)$$

where we for simplicity have used the notation $R_2 = R$ and $b_2 = b$ and $\hat{b}$ denotes the skew-symmetric matrix corresponding to the vector b . The matrix E in (4) denotes the well-known *essential matrix*. This constraint can be used to estimate the motion parameters linearly from at least eight corresponding points. The three- and four-view constraints are obtained similarly, by starting with three (or four) images of the same point and eliminating the object coordinates $\mathbf{X}$, from the resulting system of equations, see [12,13,14]. In the three-view case the trifocal constraints and the trifocal tensor are obtained and in the four-view case, the quadrifocal constraints and the quadrifocal tensor are obtained. These can be used to estimate the camera motion linearly, from at least 7 and 6 corresponding points respectively.

2.4 Continuous Matching Constraints

The *continuous time matching constraints* are obtained from the camera matrix equation (2) in continuous form and its time derivative (where for simplicity the time dependency is expressed using an index):

$$\begin{aligned} \lambda_t \mathbf{x}_t &= [R_t \mid -b_t] \mathbf{X} = R_t \tilde{\mathbf{X}} - b_t, \\ \lambda'_t \mathbf{x}_t + \lambda_t \mathbf{x}'_t &= [R'_t \mid -b'_t] \mathbf{X} = \hat{w}_t R_t \tilde{\mathbf{X}} - b'_t = \hat{w}_t (\lambda_t \mathbf{x}_t + b_t) - b'_t , \end{aligned} \quad (5)$$

where $\tilde{\mathbf{X}}$ denotes the first 3 components of the vector $\mathbf{X}$ and we have assumed that the $\mathbf{X}$ is normalized such that the fourth component is equal to 1. Furthermore, we have used the fact that the derivative of a rotation matrix can be written as $R'_t = \hat{w}_t R_t$, where w_t represents the momentary rotational velocity of the camera at time t . Similarly b'_t denotes the momentary translational velocity of the camera at time t . Define

$$\nu_t = b'_t - \hat{w}_t b_t \quad (6)$$

(representing the momentary translational velocity in a local coordinate system) and multiply the last equation in (5) with $\nu_t \times \mathbf{x}_t$ giving

$$\mathbf{x}'^T \hat{\nu} \mathbf{x} - \mathbf{x}^T \hat{\nu} \mathbf{x} = 0 \quad , \quad (7)$$

which is the well known *continuous epipolar constraint*. This constraint can be used to estimate the motion parameters from at least eight corresponding points. Higher order continuous multi-view constraints are obtained by taking higher order derivatives of (2) and then eliminating the structure, cf. [4].

3 Hybrid Matching Constraints

In this section we will introduce a novel concept, called *the hybrid matching constraints*, that will be used to update the motion parameters in two different ways.

3.1 Epipolar Hybrid Matching Constraints

Write down the camera matrix equations for time 0, time t and time $t + \Delta t$:

$$\begin{cases} \lambda_0 \mathbf{x}_0 = [I \mid 0] \mathbf{X} \\ \lambda_t \mathbf{x}_t = [R_t \mid -b_t] \mathbf{X} \\ \lambda_{t+\Delta t} \mathbf{x}_{t+\Delta t} = [R_{t+\Delta t} \mid -b_{t+\Delta t}] \mathbf{X} \end{cases} \quad (8)$$

Using $R_t = e^{t\hat{w}_t}$ as a first order approximation of R , valid between t and $t + \Delta t$, implying that

$$R_{t+\Delta t} = e^{\hat{w}_t \Delta t} R_t \approx R_t + \hat{w}_t R_t \Delta t \quad ,$$

and

$$\begin{aligned} b_{t+\Delta t} &= b_t + d_t \Delta t, \text{ corresponding to } d_t = b'_t \quad , \\ \lambda_{t+\Delta t} &= \lambda_t + \mu_t \Delta t, \text{ corresponding to } \mu_t = \lambda'_t \quad , \\ \mathbf{x}_{t+\Delta t} &= \mathbf{x}_t + \mathbf{u}_t \Delta t, \text{ corresponding to } \mathbf{u}_t = \mathbf{x}'_t \quad , \end{aligned} \quad (9)$$

and eliminating $\mathbf{X}$ using the first equation in (8) and expanding until the first order in Δt gives

$$\underbrace{\begin{bmatrix} R_t \mathbf{x}_0 & \mathbf{x}_t & 0 & b_t \\ \hat{w}_t R_t \mathbf{x}_0 & \mathbf{u}_t & \mathbf{x}_t & d_t \end{bmatrix}}_{M_d} \begin{bmatrix} -\lambda_0 \\ \lambda_t \\ \mu_t \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \quad , \quad (10)$$

implying that $\text{rank } M_d < 4$, which will be called *the hybrid epipolar constraints*. Assuming normalization such that $\mathbf{x} = (x, y, 1)$, $\mathbf{u} = (u_x, u_y, 0)$, and expanding the minors of M_d give the following different constraints in the motion parameters:

1. Minors containing the first three rows give the epipolar constraint.
2. Minors containing two rows out of the first three give linear constraints in d_t and w_t , in total nine such linear constraints.
3. Minors containing the three last rows give non-linear constraints on the motion parameters.

For our purposes, only the second type of constraints are useful. In fact there are exactly two linearly independent constraints on the motion parameters from the nine constraints of the second type above. This implies that the essential matrix can be updated from at least three corresponding points, which is a huge improvement compared to the standard discrete approaches, where five corresponding points give highly non-linear constraints, and at least eight corresponding points are needed to obtain reasonable simple linear constraints.

In order to prove that there exist exactly two linearly independent constraints among the nine ones of the second type above, we have to prove that there are exactly seven independent linear dependencies between the constraints. Start by extending M_d with the second column, giving the following 6×5 -matrix:

$$\begin{bmatrix} R_t \mathbf{x}_0 & \mathbf{x}_t & 0 & b_t & \mathbf{x}_t \\ \hat{w}_t R_t \mathbf{x}_0 & \mathbf{u}_t & \mathbf{x}_t & d_t & \mathbf{u}_t \end{bmatrix}.$$

The 5×5 minors of this matrix are obviously identically zero. Expanding the three minors obtained by removing each one of the last three columns give

$$\begin{aligned} x_t L_{11} - y_t L_{21} + L_{31} - u_x y_t E + u_y x_t E &= 0 \\ x_t L_{12} - y_t L_{22} + L_{32} - u_x E &= 0 \\ x_t L_{13} - y_t L_{23} + L_{33} + u_y E &= 0, \end{aligned} \tag{11}$$

where L_{ij} denotes the hybrid constraint obtained by removing row number i from the first three rows of M_d and row number j from the last three rows of M_d and E denotes the epipolar constraint between time 0 and t . Thus (11) contains exactly three linear dependencies among the epipolar hybrid constraints, assuming that the epipolar constraint between views 0 and t is fulfilled. Repeating the same procedure by instead adding the third column and expanding the minors obtained by removing each one of the first three columns give three further linear dependencies. Finally, adding both the second and the third column and expanding the determinant of the resulting 6×6 matrix, give the last linear dependency. Thus only two linearly independent hybrid constraints remain (assuming the epipolar constraint is fulfilled). For practical purposes, there might be more than two linearly independent constraints when the epipolar constraint is not exactly fulfilled. However, these are numerically unconditioned in the sense that they are close to spanning a two-dimensional space.

Observe that when d_t and w_t has been recovered, the new motion parameters and the new essential matrix can easily be obtained from

$$\begin{cases} R_{t+\Delta t} = e^{\hat{w}_t \Delta t} R_t \\ b_{t+\Delta t} = b_t + d_t \Delta t \end{cases} \Rightarrow E_{t+\Delta t} = R_{t+\Delta t}^T \hat{b}_{t+\Delta t} . \quad (12)$$

The update in (12) guarantees that the new essential matrix fulfils the nonlinear constraints. Observe also that although the update is nonlinear in the motion parameters it can be performed efficiently using e.g. Rodriguez formula. Finally, observe that either ν_t or d_t may be used as a parameter for the translational velocity, since they are related according to (6).

3.2 Trifocal Hybrid Matching Constraints

Write down the camera matrix equations for time 0, s , t and $t + \Delta t$:

$$\begin{cases} \lambda_0 \mathbf{x}_0 = [I \mid 0] \mathbf{X} \\ \lambda_s \mathbf{x}_s = [R_s \mid -b_s] \mathbf{X} \\ \lambda_t \mathbf{x}_t = [R_t \mid -b_t] \mathbf{X} \\ \lambda_{t+\Delta t} \mathbf{x}_{t+\Delta t} = [R_{t+\Delta t} \mid -b_{t+\Delta t}] \mathbf{X} . \end{cases} \quad (13)$$

Eliminating $\mathbf{X}$ using the first equation and expanding until the first order in Δt give

$$\underbrace{\begin{bmatrix} R_s \mathbf{x}_0 & \mathbf{x}_s & 0 & 0 & b_s \\ R_t \mathbf{x}_0 & 0 & \mathbf{x}_t & 0 & b_t \\ \hat{w}_t R_t \mathbf{x}_0 & 0 & \mathbf{u}_t & \mathbf{x}_t & d_t \end{bmatrix}}_{N_d} \begin{bmatrix} -\lambda_0 \\ \lambda_s \\ \lambda_t \\ \mu_t \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} , \quad (14)$$

implying that $\text{rank } N_d < 5$, which will be called *the hybrid trifocal constraints*. The minors of N_d gives the following different constraints in the motion parameters:

1. Minors containing only one row out of the first three give the previously derived differential-algebraic epipolar constraint.
2. Minors containing only one row out of the last three give the standard discrete trifocal constraints.
3. Minors containing two rows out of the first three, one row out of the middle three rows and two rows out of the last three rows give linear constraints in d_t and w_t , in total 27 such linear constraints.

For our purposes, only the last type of constraints are useful. It turns out that there only exist two linearly independent constraints on the motion parameters from the nine constraints of the second type above, which can be proven in the same way as above. This implies that the trifocal tensor can be updated from at least three corresponding points.

3.3 Motion Estimation Using HMC

Given a current estimate of the motion parameters and a new image with at least three corresponding points, the motion parameters can be updated using a linear system of equations of the type

$$M \begin{bmatrix} w_t \\ d_t \end{bmatrix} = m, \quad (15)$$

where $M = M(\mathbf{x}_0^k, \mathbf{x}_t^k, \mathbf{u}_t^k, R_t, b_t)$ for $k=1, \dots, n$, $n \geq 3$ and $m = m(\mathbf{x}_0^k, \mathbf{x}_t^k, \mathbf{u}_t^k, R_t, b_t)$ in the epipolar case and similarly for the trifocal case, where also $\mathbf{x}_s^k$ appears as a parameter.

3.4 State Estimation Using the Continuous-Discrete EKF

Given the motion parameters it is possible to employ a number of algorithms for recursive structure recovery. Here we optionally select a continuous-discrete extended Kalman (EKF) filter for the state estimation process [15]. For further details, see [9] or [16].

3.5 Motion Estimation Refinement by Reprojection Constraints

Given motion estimates R_t and b_t obtained using the HMC through (15), the measurement $\mathbf{x}_t$, and the 3-D estimate $\mathbf{X}$ from the EKF, we seek correction vectors $\alpha, \beta \in \mathbb{R}^{3 \times 1}$ of small magnitude, such that improved motion estimates R_t^+ and b_t^+ are given by the reprojection constraint

$$\lambda_t^+ \mathbf{x}_t = [R_t^+ \mid -b_t^+] \mathbf{X}, \quad R_t^+ = e^{\hat{\alpha}} R_t, \quad b_t^+ = b_t + \beta. \quad (16)$$

Expanding the first equation in (16) to a first order approximation gives

$$\lambda_t \mathbf{x}_t \approx R_t \tilde{\mathbf{X}} + \hat{\alpha} R_t \tilde{\mathbf{X}} + b_t + \beta \Rightarrow \widehat{R_t \tilde{\mathbf{X}}} \alpha + \beta = \epsilon := R_t \tilde{\mathbf{X}} + b_t - \lambda_t \mathbf{x}_t, \quad (17)$$

where ϵ can be interpreted as the reprojection error. Observe that (17) is a linear constraint in the correction vectors α and β . Since (17) contains two linear constraints on these 6 parameters (λ_t is also a free parameter) in the correction vectors, a linear update on the motion parameters can be made from at least three corresponding points.

The inclusion of the reprojection constraint correction step significantly enhances the performance of the estimation procedure, leading to more accurate and robust estimates of both structure and motion.

3.6 Structure and Motion Algorithm

Using the results of the previous sections, the following algorithm can now be employed for recursive structure and motion recovery:

1. Preparations

- Assume that images are obtained sequentially at time instants $t_i, i=0, 1, 2 \dots$, equally spaced by Δt . Also assume some initial values for the state vector and the error covariance matrix in the EKF.
- Given the images at times $t_0 = 0$ and $t_1 = \Delta t$ with at least eight point correspondences, get initial parameter estimates w_0 and ν_0 using e.g. the continuous eight-point algorithm.
- Compute $R_{t_1} = e^{\hat{w}_0 \Delta t}$ and $b_{t_1} = d_0 \Delta t$.

2. Estimation loop - for $i = 1, 2 \dots$ do

- Using at least three point correspondences, set up the hybrid epipolar or trifocal matching constraints in (15).
- Solve the linear system (15) for the new parameter estimates w_{t_i} and d_{t_i} .
- Update the rotation matrix and the translation vector according to (12).
- Use w_{t_i} and ν_{t_i} in the EKF to get structure estimates over the time interval $[t_i, t_i + \Delta t]$.
- Refine the motion estimate according to (17).

Note that since we are estimating both structure *and* motion, the estimates are inherently subjected to a scale ambiguity. In the above algorithm the scale issue is resolved by assuming the translational velocity vector ν to be of unit length in the initialization procedure. This together with the assumption of normalized image coordinates fixes the scale for the subsequent parameter estimates through (15).

4 Experiments

Since the initial parameter values obtained by the initialization process generally can be assumed quite accurate, the truly interesting case will be when one or both of the parameter vectors w and ν are time varying. The hybrid-based method can then be evaluated by its ability to follow the time-variations in the parameters, as well as by its ability to correctly recover the 3-D structure.

For purpose of illustration we simulate images of an object consisting of eight points in a general configuration on a grid of stepsize 10^{-4} , and with the parameter vectors

$$w(t) = \frac{2}{3} (1, -1, 1) \quad , \quad \nu(t) = \frac{1}{\sqrt{1.32}} (-1, -0.4, 0.4)^T + \frac{1}{2} (t, t, -0.5t)^T \quad .$$

Perspective measurements were computed at time instants separated by $\Delta t = 0.01$. The estimates of the components of the rotational velocity w and the translational velocity ν together with the true values based on the hybrid epipolar constraints, are shown in Fig. 1(a) and Fig. 1(b) respectively. The resulting 3-D estimation error for one of the observed object points is shown in Fig. 1(c).

We also conducted a similar experiment on the same data based on the hybrid trifocal constraints. The same procedure as before, based on the initialization using the continuous epipolar constraint, and the recursive estimation using the epipolar hybrid constraints, was used until time $t = 0.4s$. After that, the trifocal hybrid constraints was used, with $s = t/2$, see Fig 2.

Fig. 1. Estimation results obtained using the hybrid epipolar constraints: (a) True (solid) and estimated (dashed) translational velocity ν , (b) True (solid) and estimated (dashed) rotational velocity w , (c) 3-D estimation errors for one of the observed object points

Fig. 2. Estimation results obtained using the hybrid trifocal constraints: (a) True (solid) and estimated (dashed) translational velocity ν , (b) True (solid) and estimated (dashed) rotational velocity w , (c) 3-D estimation errors for one of the observed object points

5 Conclusion

We have proposed an algorithm for recursive estimation of structure and motion from perspective measurements in a continuous-discrete setting, utilizing the novel concept of the hybrid epipolar and trifocal matching constraint for the estimation of the velocity parameters, combined with a state estimator, here optionally selected as the continuous-discrete EKF. The structure and motion estimation processes are connected by recursive feedback of the structure estimates, resulting in reprojection error constraints used to obtain refined motion estimates. Simulated experiments are included to illustrate the applicability of the concept. The main advantages of the presented method is that only three corresponding points are needed for the sequential update and correction of the velocity parameter estimates. Further, both these update schemes are linear.

Note that it is *not* necessary that the same three points are tracked throughout the whole image sequence. It is easy to change to any other triplet of point correspondences when needed.

References

1. Hartley, R., Zisserman, A.: *Multiple View Geometry*. Cambridge (2003)
2. Longuet-Higgins, H.C.: A computer algorithm for reconstructing a scene from two projections. *Nature* 293, 133–135 (1981)
3. Torr, P., Fitzgibbon, A.W., Zisserman, A.: Maintaining multiple motion model hypotheses over many views to recover matching and structure. In: *Proc. International Conference on Computer Vision*. pp. 485–491 (1998)
4. Åström, K., Heyden, A.: Continuous time matching constraints for image streams. *International Journal of Computer Vision* 28(1), 85–96 (1998)
5. Viéville, R., Faugeras, O.: The first order expansion of motion equations in the uncalibrated case. *Computer Vision and Image Understanding* 64(1), 128–146 (1996)
6. Ma, Y., Košecká, J., Sastry, S.S.: Linear differential algorithm for motion recovery: A geometric approach. *International Journal of Computer Vision* 36(1), 71–89 (2000)
7. Triggs, B.: Differential matching constraints. In: *Proc. of the International Conference on Computer Vision and Pattern Recognition*. vol. 1, 370–376 (1999)
8. Heyden, A.: Differential-algebraic multiview constraints. In: *Proc. International Conference on Computer Vision*, vol. 1, pp. 159–162. IEEE Computer Society Press, Los Alamitos (2006)
9. Nyberg, F., Heyden, A.: Iterative estimation of structure and motion using hybrid matching constraints. In: *2nd Workshop on Dynamical Models for Computer Vision, DV2006 (at ECCV2006)* (2006)
10. Soatto, S.: 3-D structure from visual motion: Modeling, representation and observability. *Automatica* 33(7), 1287–1312 (1997)
11. Azarbayejani, A., Pentland, A.P.: Recursive estimation of motion, structure, and focal length. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 17(6), 562–575 (1995)
12. Triggs, B.: Matching constraints and the joint image. In: *Proc. International Conference on Computer Vision*. pp. 338–343 (1995)
13. Heyden, A.: Tensorial properties of multilinear constraints. *Mathematical Methods in the Applied Sciences* 23, 169–202 (2000)
14. Faugeras, O., Long, Q.T.: *Geometry of Multiple Images*. The MIT Press, Cambridge (2001)
15. Maybeck, P.S.: *Stochastic Models, Estimation, and Control*, vol. 2. Academic Press, San Diego (1982)
16. Huster, A.: *Relative Position Sensing by Fusing Monocular Vision and Inertial Rate Sensors*. PhD dissertation, Stanford University, Department of Electrical Engineering (2003)