Proceedings of ICAR 2003
The 11th International Conference on Advanced Robotics
Coimbra, Portugal, June 30 - July 3, 2003

Relative Position Sensing by Fusing Monocular Vision and Inertial Rate Sensors

Andreas Huster and Stephen M. Rock

Aerospace Robotics Lab, Stanford University
Stanford CA 94305, USA
{huster,rock}@arl.stanford.edu

Abstract

*Presented is a system that fuses monocular vision with inertial rate sensor measurements to generate an estimate of relative position between a moving observer and a stationary object. Although vision-only solutions are possible, fusing inertial rate sensors generates a sensing strategy that is robust to vision drop-outs and is able to determine relative position with minimal requirements on the vision system. However, the combination of limited observability and significant nonlinearities, which are inherent to this sensing strategy, creates an estimation problem which cannot be solved with a standard Extended Kalman Filter (*EKF*). This paper describes an estimation algorithm that is uniquely adapted to this sensor fusion problem and presents results from laboratory experiments with a manipulator system. For these experiments, the estimator was implemented as part of a closed-loop control system that can perform an object pick-up task.*

1 Introduction

Relative position sensing and control are core requirements for a wide range of autonomous, intervention-capable robots in terrestrial, space and underwater environments. This paper focuses on a new relative position sensing system that fuses bearing information from monocular vision with inertial rate sensor measurements to estimate relative position, velocity and orientation. The estimate is computed in real time and is suitable for closed-loop control. A feature of this system is that it relies on very simple vision requirements: tracking a single feature in a single camera image. Real environments exacerbate the typical challenges of identifying good visual features, establishing feature correspondences and robust tracking. Reducing the vision requirements by integrating inertial rate sensors has the potential to produce a more robust sensing system than typical vision-only techniques.

The sensing strategy takes advantage of the complementary nature of monocular vision measurements, inertial rate sensor measurements, and observer[1] motion. The motion of the camera between successive images generates a baseline for range computations by triangulation. Inertial rate sensors, whose acceleration and angular rate measurements can be integrated to obtain velocity, position and orientation, can account for the 6-DOF motion of the camera along this baseline. When these measurements are fused, the relative position between the observer and the object can be computed. A key benefit of this system is that the inertial rate sensors continue to maintain a useful estimate of relative position during vision drop-outs (e.g., occlusions, detection errors, lack of correspondence). Furthermore, both inertial rate sensors (for navigation) and monocular vision systems (for science purposes) are already common sensors on many field robots.

However, inertial rate sensors suffer from bias and random noise errors. These lead to unbounded drift during a simple integration of the measurements to account for observer motion. The fusion algorithm estimates these inertial rate measurement errors. This is especially important for low-cost inertial rate sensors, which are subject to significant drift errors.

Inherent to this sensing strategy based on a single bearing measurement is a combination of limited observability and significant nonlinearities. The Extended Kalman Filter (EKF), a standard nonlinear estimation technique, performs poorly in this context, because it relies on linearizations of the nonlinear sensor and process models in order to apply the Kalman Filter equations. The EKF is unable to account for the uncertainty that results from these linearizations when the states about which the models are linearized are uncertain. This is a critical issue for this estimation problem. Consequently, the EKF underpredicts the estimate uncertainty and generates estimates with significant biases.

A new estimator design, which handles nonlinearities without linearizing them, has been created to solve this problem. However, the ability of this estimator to converge depends on the trajectory of the observer. For example, during camera motion directly toward the feature, the estimator has no new information with which to improve its range estimate. Only camera motions transverse to the feature direction provide observability for the range estimate. However, motion directly toward the object is typically required to complete a manipulation task. Furthermore, moving toward the object improves the effectiveness

[1] In this paper, observer refers to the moving sensor platform, and not the estimation algorithm.

Figure 1: OTTER*: a Small Underwater Vehicle Operated in* MBARI*'s Test Tank*

of future transverse motion.

In previous work, we addressed and demonstrated techniques to develop estimators that will be effective for these applications. In [1], we demonstrated that the EKF fails to provide adequate solutions for this sensor fusion problem and explored an alternative approach for a simplified (2D) problem. In [2], we presented a laboratory testbed and demonstration task to evaluate the new sensing system. In [3], we outlined an algorithm to solve the complete estimation problem and described the first application of this sensing strategy to perform a useful manipulation task.

In this paper, we extend these results and explore the performance of the algorithm in the presence of significant process noise. In particular, our goal is to predict the expected accuracy of the estimator when it is applied to an AUV (our target AUV is the OTTER vehicle shown in Figure 1 and described in [4]). Our approach is to use a laboratory test bed manipulator configured to obey equations of motion characteristic of an AUV operating in an ocean environment.

Section 2 defines the sensor fusion problem and presents models for the vision and inertial rate sensor measurements, the dynamics, and disturbances. Section 3 describes the estimator design. Initial experiments, conducted in the laboratory with a fixed-base 7-DOF manipulator arm, are described in Section 4. This platform provides a truth measurement and can be used to investigate competing approaches, to simulate different disturbance environments, and to quantify performance. Section 5 demonstrates the effectiveness of the estimator design by combining the estimator with a trajectory and a controller to perform a simple manipulation task.

1.1 Related Work

Other relative position sensing systems that fuse measurements from a bearing sensor and inertial rate sensors have been discussed in Kaminer et. al. [5] (airplane tracking a ship) and Gurfil and Kasdin [6] (missile intercepting a target). These authors are also concerned with finding alternatives to the EKF. Both papers present simulations of systems with known inertial rate sensor biases, known observer orientation, and perfect gravity compensation. These assumptions do not hold for our work with low-cost inertial rate sensors.

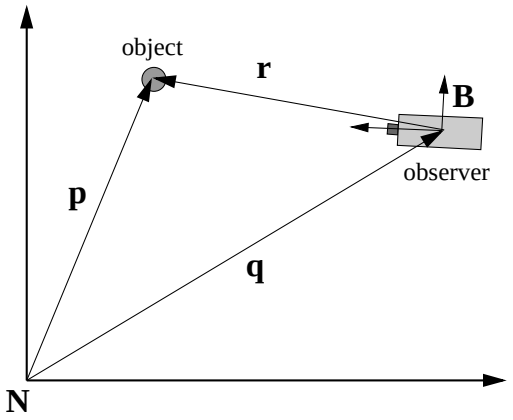

Figure 2: *Geometry of the Estimation Problem*

2 Sensor Fusion Problem

2.1 Estimation Scenario

Figure 2 shows a stationary object and a moving observer, composed of a camera and inertial rate sensors. The camera is tracking the object and the inertial rate sensors are reporting the observer's acceleration and angular velocity. The purpose of the estimator is to determine the relative position and velocity between the object and the observer.

Frame $\mathbf{N}$ is the inertial frame and $\mathbf{p}$ indicates the position of a feature on the stationary object tracked by the camera. $\mathbf{q}$ is the position of the observer body frame $\mathbf{B}$. To simplify the discussion, we assume that the camera and the inertial rate sensors are all co-located at $\mathbf{q}$. In practice, any known position offset between the sensors and the origin of frame $\mathbf{B}$ can be incorporated into the algorithm. The z-axis of the body frame is aligned with the optical axis of the camera. R is the rotation matrix that transforms a vector resolved in inertial coordinates to the body frame. $\boldsymbol{\omega}$ is the associated rotational velocity resolved in the body frame. A leading superscript (e.g., ${}^{B}\mathbf{x}$) indicates that the vector is resolved in a specific frame.

The position of the feature as seen by the observer is $\mathbf{r} = \mathbf{p} - \mathbf{q}$. We assume that the feature is stationary in the inertial frame, so $\dot{\mathbf{p}} = \ddot{\mathbf{p}} = 0$. Therefore, $\dot{\mathbf{r}} = -\dot{\mathbf{q}}$ and $\ddot{\mathbf{r}} = -\ddot{\mathbf{q}}$. Because of this assumption, a measurement of the observer acceleration $\ddot{\mathbf{q}}$ is useful for estimating the relative feature position $\mathbf{r}$.

2.2 System Block Diagram

The block diagram in Figure 3 shows a system that uses the sensing strategy to perform a control task. The dashed box represents the environment and includes the moving observer, the stationary object, the camera, the inertial rate sensors and disturbance forces. This environment can be described by the state vector $\mathbf{x}$. The sensor measurements $\mathbf{z}$ together with the observer command $\mathbf{u}$ are inputs to the estimator, which computes an estimate $\hat{\mathbf{x}}$ of the state. The difference between this estimate and the desired state $\mathbf{x}_{des}$,

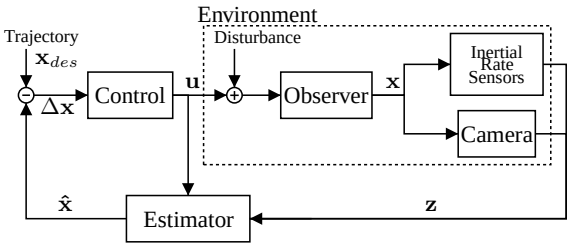

Figure 3: *Block Diagram of a System Based on the Sensing Strategy*

specified by the trajectory, is used in the controller to compute the observer commands.

2.3 Sensor Models

The vision measurement $\mathbf{z}_s$ is the projection of $\mathbf{r}$ onto the image plane, and is modeled as follows:

$$\mathbf{S} = \begin{bmatrix} S_x \\ S_y \\ S_z \end{bmatrix} = R\,{}^N\mathbf{r} = R\,(\mathbf{p} - \mathbf{q}) \quad (1)$$

$$\mathbf{z}_s = \begin{bmatrix} s_x \\ s_y \end{bmatrix} + \mathbf{n}_s = \frac{1}{S_z}\begin{bmatrix} S_x \\ S_y \end{bmatrix} + \mathbf{n}_s \quad (2)$$

$\mathbf{n}_s$ is zero-mean white Gaussian measurement noise. The camera measurements are normalized so the effective focal length is 1.

The accelerometer measures specific force, which includes the acceleration $\ddot{\mathbf{q}}$ of the observer and a component due to gravity. The measurement $\mathbf{z}_a$ also includes sensor biases $\mathbf{b}_a$ and zero-mean white Gaussian sensor noise $\mathbf{n}_a$.

$$\mathbf{z}_a = \boldsymbol{\alpha} R\,(-\ddot{\mathbf{q}} + \mathbf{g}) + \mathbf{b}_a + \mathbf{n}_a \quad (3)$$

$\mathbf{g} = \begin{bmatrix} 0 & 0 & -g \end{bmatrix}^T$ is the acceleration due to gravity in inertial coordinates. $\boldsymbol{\alpha} \approx I_{3\times 3}$ are scale factors induced by the sensor.

The rate gyro measurement includes the rotational velocity $\boldsymbol{\omega}$ of the observer, sensor biases $\mathbf{b}_\omega$, and zero-mean white Gaussian sensor noise $\mathbf{n}_\omega$.

$$\mathbf{z}_\omega = \boldsymbol{\omega} + \mathbf{b}_\omega + \mathbf{n}_\omega \quad (4)$$

Random walk models are used to capture the dynamics of the inertial rate sensor parameters. $\mathbf{n}_{ba}$, $\mathbf{n}_{b\omega}$ and n_α are zero-mean white Gaussian driving terms.

$$\frac{d}{dt}\mathbf{b}_a = \mathbf{n}_{ba} \quad (5)$$

$$\frac{d}{dt}\mathbf{b}_\omega = \mathbf{n}_{b\omega} \quad (6)$$

$$\frac{d}{dt}\alpha = n_\alpha \quad (7)$$

2.4 Process Model

A linear drag model relates the control input $\mathbf{u}$ to observer velocity.

$$\frac{d}{dt}\dot{\mathbf{q}} = \ddot{\mathbf{q}} = R^T\mathbf{u} + \gamma\,(\mathbf{w} - \dot{\mathbf{q}}) \quad (8)$$

$\mathbf{u}$ is known and represents the actuator commands in observer body coordinates. $\mathbf{w}$ represents the water velocity and can be interpreted as the source for a disturbance $\mathbf{d} = \gamma\mathbf{w}$. $\mathbf{d}$ represents unknown external forces on the observer as well as errors in the actuator model. It is modeled by a first-order Gauss-Markov process.

$$\frac{d}{dt}\mathbf{d} = -\frac{1}{\tau}\mathbf{d} + \mathbf{n}_d \quad (9)$$

where $\mathbf{n}_d$ is zero-mean white Gaussian noise.

3 Estimator Design

The sensor fusion algorithm requires a nonlinear estimator, like the Extended Kalman Filter (EKF). However, while the EKF is a useful solution for many nonlinear estimation problems, we have shown in [1] that the direct application of the EKF to this problem fails to generate an adequate solution.

The EKF fails because it relies on linearizations of the sensor and process models in order to apply the Kalman Filter equations. When the states about which these models are linearized are uncertain, the linearization can represent significant uncertainty, which is not accounted for in the EKF. In this problem, the object range is a state that can be very uncertain and also one that is involved in the linearizations. As a result, the EKF underpredicts the estimate uncertainty, which leads to biased estimates.

The estimator design is based on the Kalman Filter framework, but it uses a specific state representation, $\mathbf{x}$, that leads to a linear sensor model. Insisting on a linear sensor model transfers all of the nonlinearities into the state dynamics and avoids linearizations in the measurement update. The time update, which captures these dynamics, is implemented with the unscented transform [7, 8], which does not require linearization. The motivation for this design and additional detail are presented in [3, 9].

Implementing a measurement update without linearization requires a state representation that leads to a linear sensor model $\mathbf{z} = H\mathbf{x}$, where H is constant. Therefore, terms that appear in the measurement models, such as s_x, s_y, $R\mathbf{d}$, $R\dot{\mathbf{q}}$, and $\mathbf{b}_a$, have to appear explicitly in the state vector.

The representation of feature range S_z presents an additional design choice. Because it does not appear in the sensor model, it is not constrained by the requirement of a linear measurement update and we are free to choose a convenient representation. We represent feature range with $\zeta = 1/S_z$, which leads to low-order polynomials as the dominant nonlinearities in the state dynamics. Polynomials tend to result in more accurate estimator time-updates

than, for example, ratios, which are induced by representing range with S_z.

The design includes a simplifying assumption on the accelerometer model. We assume that $\boldsymbol{\alpha} = \alpha I_{3\times 3}$ and that the scale factor variation can be ignored for the contribution due to real accelerations, which is much smaller than the contribution from gravity. This assumption has a small impact on overall accuracy and greatly simplifies the estimator design. The modified accelerometer model is

$$\mathbf{z}_a = R\left(-\ddot{\mathbf{q}} + \alpha \mathbf{g}\right) + \mathbf{b}_a + \mathbf{n}_a. \tag{10}$$

One important remaining source of estimate error that is not handled by this design is the assumption that the states can be modeled as Gaussian random variables, which is not necessarily true for nonlinear problems. If the estimate uncertainties are very large, this assumption can cause problems. In that case, a more elaborate solution, like a Particle Filter design [10], might be necessary. However, for moderate uncertainties encountered by this application, this design has produced good results.

The estimator state vector is given by:

$$\mathbf{x} = \begin{bmatrix} s_x \\ s_y \\ \zeta \\ \mathbf{v} \\ \mathbf{a} \\ \mathbf{b}_a \\ \mathbf{Z} \\ \psi \\ \mathbf{b}_\omega \end{bmatrix} \quad \begin{array}{l} \mathbf{v} = \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = R\dot{\mathbf{q}} \\ \\ \mathbf{a} = R\mathbf{d} \\ \\ \mathbf{Z} = \begin{bmatrix} Z_x \\ Z_y \\ Z_z \end{bmatrix} = R \begin{bmatrix} 0 \\ 0 \\ \alpha \end{bmatrix} \end{array} \tag{11}$$

$\mathbf{Z}$ represents the direction of gravity in the body frame modified by the accelerometer scale factor α. ψ is the heading, or rotation about $\mathbf{Z}$. Together, $\mathbf{Z}$ and ψ define the observer orientation. A similar representation for attitude is described in [11].

The corresponding state dynamics are given by (5), (6), and

$$\begin{aligned}
\frac{d}{dt}s_x &= -v_x\zeta + s_x v_z \zeta \\
&\quad + s_x s_y \bar{\omega}_x - (1+s_x^2)\bar{\omega}_y + s_y \bar{\omega}_z && (12) \\
\frac{d}{dt}s_y &= -v_y\zeta + s_y v_z \zeta \\
&\quad + (1+s_y^2)\bar{\omega}_x - s_x s_y \bar{\omega}_y - s_x \bar{\omega}_z && (13) \\
\frac{d}{dt}\zeta &= v_z\zeta^2 + \zeta s_y \bar{\omega}_x - \zeta s_x \bar{\omega}_y && (14) \\
\frac{d}{dt}\mathbf{v} &= \mathbf{u} + \mathbf{a} - \gamma\mathbf{v} - \bar{\boldsymbol{\omega}} \times \mathbf{v} && (15) \\
\frac{d}{dt}\mathbf{a} &= -\frac{1}{\tau}\mathbf{a} - \bar{\boldsymbol{\omega}} \times \mathbf{a} + \mathbf{n}_d && (16) \\
\frac{d}{dt}\mathbf{Z} &= -\bar{\boldsymbol{\omega}} \times \mathbf{Z} + \frac{1}{\alpha}\mathbf{Z}n_\alpha && (17) \\
\frac{d}{dt}\psi &= \frac{1}{\alpha\left(Z_y^2 + Z_z^2\right)} \begin{bmatrix} 0 & Z_y & Z_z \end{bmatrix} \bar{\boldsymbol{\omega}} && (18)
\end{aligned}$$

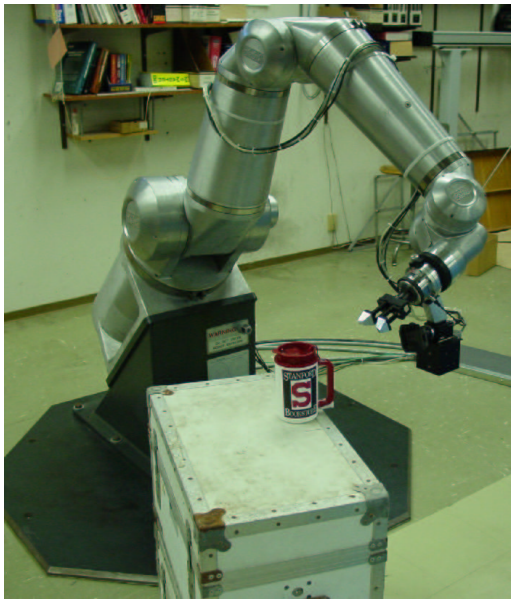

Figure 4: *RRC K-1607 7-DOF Manipulator*

$$\bar{\boldsymbol{\omega}} = \begin{bmatrix} \bar{\omega}_x \\ \bar{\omega}_y \\ \bar{\omega}_z \end{bmatrix} = \mathbf{z}_\omega - \mathbf{b}_\omega - \mathbf{n}_\omega \tag{19}$$

Note that the rate gyro measurement $\mathbf{z}_\omega$ is used directly in the dynamics. The noise sources $\mathbf{n}_\omega$, $\mathbf{n}_d$, $\mathbf{n}_a$, $\mathbf{n}_\omega$ and n_α represent process noise. In practice, ψ is updated with a difference equation that avoids the singularity when $Z_y = Z_z = 0$.

We use the square-root version of the unscented transform [12] to propagate the estimate and covariance forward in time. We chose this algorithm based on accuracy, ease of implementation, and computational efficiency.

The camera measurements $\mathbf{z}_s$ and the accelerometer measurements $\mathbf{z}_a$ are incorporated with a linear measurement update, which is implemented with an array algorithm [13]. The combination of an array algorithm for the measurement update and the square-root version of the unscented transform for the time-update leads to an algorithm that operates on and stores only the square-root of the covariance, and not the actual covariance, which results in better numerical properties and reduced computational cost.

4 Laboratory Testbed

Figure 4 shows the experimental hardware used in this research. It is a K-1607 manipulator built by Robotics Research Corporation[2]. It is a 7-DOF, kinematically redundant manipulator whose endpoint can be moved to any position and orientation in its workspace. We have at-

[2] http://www.robotics-research.com

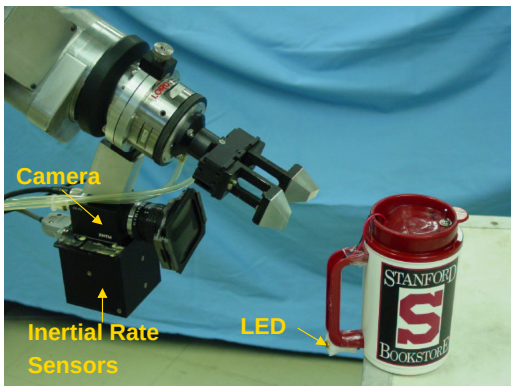

Figure 5: *Manipulator Endpoint with the Camera, Inertial Rate Sensors, and Gripper; and the Cup with Infrared* LED

tached a camera and inertial rate sensors (DMU-6X Inertial Measurement Unit by Crossbow[3]) on the endpoint of the manipulator in order to demonstrate the estimator performance in the context of real sensor measurements. All of the manipulator joints are instrumented with encoders, so that the exact position of the endpoint can also be computed.

We have developed a simple robotic task—picking up an object—to demonstrate the relative position estimator. The object is a large plastic cup which the robot can pick up using a pneumatic gripper. The manipulator endpoint, with the gripper, camera, and inertial rate sensors, as well as the cup are shown in Figure 5. An LED on the handle of the cup is the only feature that the camera can see. The relative position estimate obtained by fusing the bearing to this LED with inertial rate sensor measurements is used to control the motion of the robot.

This experiment is based on the system in Figure 3. A more detailed description of this experiment as well as a discussion on trajectories for this task appears in [3].

5 Results

This section presents the performance of the estimator in the presence of disturbances scaled to be typical of underwater environments. Specifically, the time constant was estimated to be $\tau = 60\,s$ and the standard deviation of $\mathbf{d}$ was estimated to be $\sigma(\mathbf{d}) = 0.012\,m/s^2$. This leads to $\sigma(\mathbf{n}_d) = 0.0022$. Note that the manipulator has been programmed to emulate the motion of an underwater vehicle subject to disturbances.

Before each run of the experiment, the observer (manipulator endpoint) was moved to a random initial position $\mathbf{q}_0$ and the estimator was reset to the initial estimate $\hat{\mathbf{x}}_0$, which assumes that the initial target range is $0.65\,m$ and that all other states, including inertial rate sensor biases, are unknown. The system uses the current state estimate to control the observer along a precomputed trajectory, which defines the desired relative position and observer orientation. Throughout the experiment, the actual feature position $\mathbf{p}$ remains constant.

[3] http://www.xbow.com

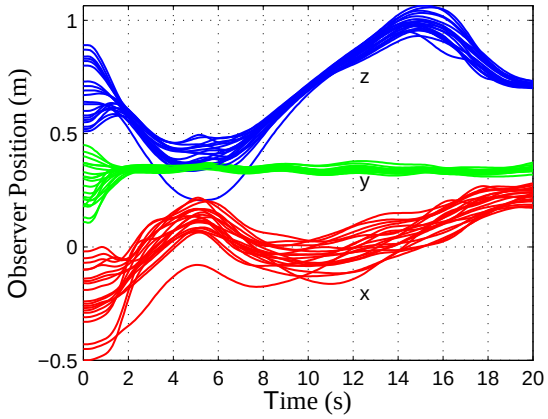

Figure 6: *Observer Position for Twenty Runs with Different Initial Conditions*

Table 1: *Mean and Standard Deviation of Final Observer Position With Artificial Disturbance*

	mean (m)	standard deviation (m)
q_x	0.230	0.032
q_y	0.343	0.011
q_z	0.715	0.009

Figure 6 overlays the observer position $\mathbf{q}$ (a truth value computed from the manipulator joint encoders) for twenty runs, each with different initial observer positions $\mathbf{q}_0$. These range from -0.50 to $-0.02\,m$ in x, 0.11 to $0.45\,m$ in y, and 0.51 to $0.89\,m$ in z. This plot shows that the controller, which has access only to the relative position estimate and not the truth measurement $\mathbf{q}$, is successful in moving the observer from various initial positions toward a desired position near the object. Table 1 shows the mean and standard deviation of the final observer position for these runs. Both the plot and the table show that the x-coordinate of position, which corresponds to the optical axis of the camera, is the most difficult to estimate.

This experiment shows that the estimator, implemented as part of a real-time control system, can determine relative position by fusing monocular vision measurements of a single feature and inertial rate sensor measurements.

Table 2: *Standard Deviation of Final Observer Position Without Artificial Disturbance (previously reported in [3])*

	standard deviation (m)
q_x	0.006
q_y	0.003
q_z	0.002

Further, the estimator performs well in the context of real sensors and realistic disturbance environments. In particular, these results predict an accuracy of 3.2 cm when the estimator is applied to the OTTER AUV experiments.

For comparison, Table 2 reproduces the standard deviations of the final observer position that were reported in [3] for the case when no artificial disturbances were used.

6 Conclusions

This paper discusses a sensing strategy that fuses monocular vision measurements of a stationary object with inertial rate sensor measurements to estimate relative position. This capability satisfies a core requirement for many autonomous, intervention-capable robots: a robust, real-time estimate of the relative position between a moving observer and a stationary object. The sensing strategy inherits the advantages of vision in unstructured environments, but provides greater robustness than vision-only solutions and can operate with only a single trackable feature.

We have defined this sensor strategy, outlined an estimation algorithm to solve the sensor fusion problem, and performed laboratory experiments to evaluate the approach. The experimental results demonstrate the effectiveness of this sensing strategy in the context of real sensors and typical disturbances. We have shown that the estimator can be used as part of a closed-loop control system for a manipulator arm performing an object pick-up task. Our future work will focus on demonstrating this sensor strategy on underwater vehicles.

Acknowledgments

This research was supported in part by the Packard Foundation under Grants 98-3816 and 98-6228.

References

[1] Andreas Huster and Stephen M. Rock, "Relative position estimation for intervention-capable AUVs by fusing vision and inertial measurements," in *Proceedings of the 12th International Symposium on Unmanned Untethered Submersible Technology*, Durham, NH, August 2001, Autonomous Undersea Systems Institute.

[2] Andreas Huster and Stephen M. Rock, "Relative position estimation for manipulation tasks by fusing vision and inertial measurements," in *Proceedings of the Oceans 2001 Conference*, Honolulu, November 2001, MTS/IEEE, vol. 2, pp. 1025–1031.

[3] Andreas Huster, Eric W. Frew, and Stephen M. Rock, "Relative position estimation for AUVs by fusing bearing and inertial rate sensor measurements," in *Proceedings of the Oceans 2002 Conference*, Biloxi, MS, October 2002, MTS/IEEE, vol. 3, pp. 1857–1864.

[4] H. H. Wang, S. M. Rock, and M. J. Lee, "OTTER: The design and development of an intelligent underwater robot," *Autonomous Robots*, vol. 3, no. 2-3, pp. 297–320, June-July 1996.

[5] Isaac Kaminer, Wei Kang, Oleg Yakimenko, and Antonio Pascoal, "Application of nonlinear filtering to navigation system design using passive sensors," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, no. 1, pp. 158–172, January 2001.

[6] Pini Gurfil and N. Jeremy Kasdin, "Two-step optimal estimator for three dimensional target tracking," in *Proceedings of the 2002 American Control Conference*, Anchorage, AK, May 2002, vol. 1, pp. 209–214.

[7] Simon Julier and Jeffrey K. Uhlmann, "Data fusion in nonlinear systems," in *Handbook of Multisensor Data Fusion*, David L. Hall and James Llinas, Eds., chapter 13. CRC Press, 2001.

[8] Simon Julier, Jeffrey Uhlmann, and Hugh F. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 477–482, March 2000.

[9] Andreas Huster, *Relative Position Sensing by Fusing Monocular Vision and Inertial Rate Sensors*, Ph.D. thesis, Stanford University, Stanford, California, 2003, also available at http://arl.stanford.edu/∼huster/dissertation/.

[10] A. Doucet, J. F. G. de Freitas, and N. J. Gordon, Eds., *Sequential Monte Carlo Methods in Practice*, Springer Verlag, New York, 2001.

[11] Henrik Rehbinder and Xiaoming Hu, "Drift-free attitude estimation for accelerated rigid bodies," in *IEEE International Conference on Robotics and Automation*, Seoul, South Korea, May 2001, vol. 4, pp. 4244–4249.

[12] Rudolph van der Merwe and Eric A. Wan, "The square-root unscented Kalman filter for state and parameter-estimation," in *International Conference on Acoustics, Speech, and Signal Processing*, Salt Lake City, Utah, May 2001, IEEE, pp. 3461–3464.

[13] Thomas Kailath, Ali H. Sayed, and Babak Hassibi, *Linear Estimation*, Prentice Hall, 2000.