

Passive GPS-Free Navigation for Small UAVs

Jack Langelaan
Steve Rock

Department of Aeronautics and Astronautics
Stanford University
Stanford, CA 94305, USA
jackl@stanford.edu

Abstract— A method for passive GPS-free navigation of a small Unmanned Aerial Vehicle with a minimal sensor suite (limited to an inertial measurement unit and a monocular camera) is presented. The navigation task is cast as a Simultaneous Localization and Mapping (SLAM) problem. While SLAM has been the subject of a great deal of research, the highly non-linear system dynamics and limited sensor suite available in this application presents a unique set of challenges which have not previously been addressed. In this particular application solutions based on Extended Kalman Filters have been shown to diverge and alternate techniques are required.

In this paper an Unscented Kalman Filter is applied to the navigation problem, which leads to a consistent estimate of vehicle and feature states. This paper presents: (a) simulation results showing mapping and navigation in three dimensions; (b) preliminary hardware test results showing navigation and mapping using an off-the-shelf inertial measurement unit and camera in a laboratory environment.

TABLE OF CONTENTS

- 1 INTRODUCTION
- 2 NAVIGATION
- 3 ESTIMATION PROCESS
- 4 SIMULATION RESULTS
- 5 HARDWARE TEST RESULTS
- 6 CONCLUSION
- 7 ACKNOWLEDGEMENTS

1. INTRODUCTION

Navigation without reference to some global coordinate system (i.e. GPS) is a pivotal problem in the field of mobile robotics including aerospace navigation applications [4]. In many situations GPS is unavailable due to jamming or environmental considerations and other navigation techniques are required.

This is of particular interest for navigation of small UAVs, which may be operated in areas where GPS signals are

jammed or are obscured by natural or man-made features. Small UAVs, however, present a unique set of challenges ultimately stemming from the limited payload capacity of the vehicle. Sensing is greatly limited, often including only an inertial measurement unit (IMU) and a monocular camera. Furthermore, since small UAVs are typically intended to be inexpensive, the use of high-quality IMUs is precluded, making a purely inertial navigation solution infeasible. Hence the problem is to obtain a localization solution which is accurate enough and is updated quickly enough to be useful for navigation using an extremely limited sensor suite (the on-board IMU and monocular camera).

By fusing inertial measurements and bearings to fixed objects on the ground we can obtain a navigation solution. Fusing inertial measurements with bearings to known features has been applied to the problem of autonomous ship-board landing [8], indoor navigation by humans [5] and navigation of robots [13], [11], [12]. When features are not known their locations must be included in the vector of states to be estimated. This is known as Simultaneous Localization and Mapping (SLAM), sometimes referred to as Concurrent Localization and Mapping.

SLAM has been the subject of a great deal of research. It enables navigation by simultaneously estimating the state of the robot and of landmarks in the environment using only on-board sensing, which typically includes odometry (for wheeled robots) or inertial measurements to measure vehicle motion and relative measurements (i.e. range and bearing) to landmarks. SLAM has been applied both in simulation and on hardware in situations as diverse as navigation of wheeled ground vehicles (e.g. [14], range and bearing provided by a scanning laser range finder) and subsea navigation of autonomous underwater vehicles (e.g. [16], range provided by sonar). An algorithm has even been developed for a fixed wing UAV [9]. However these were fairly large aircraft (10kg payload) and measurements included both bearing and range to landmarks using a camera and a laser rangefinder.

Many SLAM implementations use Extended Kalman Filters (EKF) to estimate vehicle and feature states. However when sensing is limited to inertial measurements and bearings to features experience has shown that EKF-based approaches often lead to divergent solutions [10].




Figure 1. Navigation Scenario

In [10] we presented a solution based on an Unscented Kalman Filter (UKF). Simulation results in 2D showed that consistent unbiased estimates of both vehicle state and obstacle positions could be obtained and that these estimates could be used by a generic obstacle avoidance algorithm to enable collision-free navigation through a previously unsurveyed forest. In this paper we extend the previous work to three dimensions, incorporate methods to perform data association and incorporate methods to enable addition of new features to the map. This is demonstrated with simulation results in three dimensions and preliminary hardware test results using an off the shelf IMU and camera showing mapping and navigation in a laboratory environment.

2. NAVIGATION

The specific navigation scenario considered here is shown schematically in Figure 1. A small UAV equipped with a monocular camera and an IMU flies through an unsurveyed forest. The camera obtains bearing measurements to obstacles (tree trunks) and the IMU provides accelerations and angular rates in the body-fixed frame.

Trees are located at \mathbf{x}_i in frame O , an inertial NED frame. A transformation matrix \mathbf{T} resolves a vector in frame O to the body frame B . $\boldsymbol{\omega}$ represents the rotational velocity of the vehicle expressed in the body frame and $[u \ v \ w]^T$ represents the velocity of the aircraft in the body frame. The position of the aircraft is defined by $[x \ y \ z]^T$, expressed in frame O .

The problem is to estimate the locations of the tree trunks and of the aircraft given the limited sensor suite. This information can subsequently be used by a path planning algorithm to generate a collision-free trajectory through the forest.

The estimation problem contains two significant nonlinearities. The first is due to the rotational degrees of freedom of the UAV; the second is due to the vision system's projection of the three-dimensional world onto the two-dimensional image plane. Thus the system dynamics and measurement equations are highly non-linear and depend on

the actual state of the system.

Since only bearing measurements are available, observability of states is highly dependent on motion of the camera (i.e. the trajectory flown by the UAV). During motion directly towards or away from an object (along the bearing) there is no new information which improves range estimate. Transverse motion is required produce a useful estimate of object position. In the case of obstacle avoidance, transverse motion has the added benefit of ensuring that a collision is avoided, but the presence of multiple obstacles places conflicting demands on the trajectory which must be flown. Additional demands on the trajectory may arise if certain areas are to be avoided or if a final goal location must be reached. However, trajectory-related issues are not addressed in this paper.

To summarize, given a sensor suite limited to a monocular camera and an IMU an estimate of aircraft state and obstacle locations must be generated.

3. ESTIMATION PROCESS

The use of low-cost inertial sensors complicates system design. Both bias and random noise errors are present and may be quite large in low cost inertial sensors, which can lead to unbounded errors in the integrated quantities. Hence sensor biases must be included in the UAV state vector.

System Equations

The estimator propagates a state vector (which includes aircraft state and obstacle states) and the associated covariance:

$$\hat{\mathbf{x}} = \begin{bmatrix} \hat{\mathbf{x}}_v \\ \hat{\mathbf{x}}_o \end{bmatrix} \quad (1)$$

$$\mathbf{P} = \begin{bmatrix} \mathbf{P}_{vv} & \mathbf{P}_{vo} \\ \mathbf{P}_{ov} & \mathbf{P}_{oo} \end{bmatrix} \quad (2)$$

The feature state vector is concatenated from the states of each individual feature:

$$\hat{\mathbf{x}}_o = \begin{bmatrix} \hat{\mathbf{x}}_1 \\ \hat{\mathbf{x}}_2 \\ \vdots \\ \hat{\mathbf{x}}_m \end{bmatrix} \quad (3)$$

The vehicle state vector is

$$\hat{\mathbf{x}}_v = [x \ y \ z \ \phi \ \theta \ \psi \ u \ v \ w \ \boldsymbol{\alpha}^T \ \mathbf{b}_a^T \ \mathbf{b}_\omega^T]^T \quad (4)$$

The IMU error states must be included because of the possibility of drift.

Time Update (Vehicle Kinematics)

The time update of the filter is driven by the kinematics of the vehicle and dynamics of the IMU biases. Measurements of acceleration and angular rate supplied by the inertial measurement unit. Features are assumed to be stationary.

In discrete form, the system dynamics are

$$\mathbf{x}_{k+1} = \begin{bmatrix} \mathbf{x}_{v,k+1} \\ \mathbf{x}_{o,k+1} \end{bmatrix} = \begin{bmatrix} f(\mathbf{x}_{v,k}, \mathbf{u}_k) \\ \mathbf{x}_{o,k} \end{bmatrix} + \begin{bmatrix} \mathbf{n}_{ext} \\ 0 \end{bmatrix} \quad (5)$$

where

$$\mathbf{u}_k = \begin{bmatrix} \mathbf{z}_{imu,k} \\ \mathbf{n}_{\alpha,k} \\ \mathbf{n}_{b,k} \end{bmatrix} \quad (6)$$

Both \mathbf{u} and \mathbf{n}_{ext} are assumed to be Gaussian random variables. States related to the vehicle are driven by the inertial measurements $\mathbf{z}_{imu,k}$ and the IMU bias states are driven by a zero-mean random walk. Note that measurements from the IMU are likely to be available at a much higher rate than measurements from the camera, allowing the time update to proceed at a higher rate than the vision correction.

The nonlinear function $f(\mathbf{x}_{v,k}, \mathbf{u}_k)$ captures the kinematics of the vehicle. Since inertial measurements are noisy this must be included as process noise in the time update. We approximate process noise as

$$\mathbf{Q}_k = \begin{bmatrix} \nabla_{\mathbf{u}} f \Sigma_{\mathbf{u}} \nabla_{\mathbf{u}} f^T + \Sigma_{ext} & 0 \\ 0 & 0 \end{bmatrix} \quad (7)$$

This is identical to the EKF process noise approximation, and is adequate for the application being considered. Since obstacles (i.e. trees) are assumed to be perfectly stationary only components corresponding to vehicle states have non-zero process noise.

Measurement Update (Vision Correction)

Measurements from the camera are incorporated in the measurement update step of the estimator. The vision model reflects the projection of a vector in 3D (the vector from the vehicle to the feature) onto the 2D image plane. This has the form

$$\mathbf{z}_{cam,k} = h(\mathbf{x}_{v,k}, \mathbf{x}_{o,k}) + \mathbf{n}_c \quad (8)$$

Here \mathbf{z}_{cam} represents bearings from the vehicle to stationary features. Measurement noise is represented by the zero-mean Gaussian random variable \mathbf{n}_c .

The projection function $h(\mathbf{x}_{v,k}, \mathbf{x}_{o,k})$ is highly non-linear, complicating the problem of accurately calculating the uncertainty associated with a measurement prediction. In general the uncertainty in vehicle state will be large enough that the measurement equations cannot be accurately approximated by linearization over the range of possible vehicle states.

Unscented Kalman Filter Implementation

The vehicle kinematics equations, vision equations and measurements (inertial and vision) are used to implement a filter to generate estimates of the state vector. An extended Kalman filter approximates a non-linear system by linearizing the system equations about the current estimate. However, the states about which the dynamics are linearized are uncertain, leading to potentially significant unmodelled uncertainty in the

linearization. An Unscented Kalman Filter [7] [15] instead approximates the probability distribution of the state which is to be estimated. The probability distribution is represented by a set of *Sigma Points* which capture the mean and covariance of a random vector $\mathbf{x} = \mathcal{N}(\hat{\mathbf{x}}, \mathbf{P})$:

$$\mathbf{X} = [\hat{\mathbf{x}} \quad \hat{\mathbf{x}} + \eta\sqrt{\mathbf{P}} \quad \hat{\mathbf{x}} - \eta\sqrt{\mathbf{P}}] \quad (9)$$

where η is a scale factor and $\sqrt{\mathbf{P}}$ is an orthogonal matrix square root of the covariance of the distribution. Note that this matrix square root is not unique: it can be chosen based on issues such as numerical stability, computational cost or to exploit known structure in the problem (Huster [6] makes use of this property: he chooses sigma points so that the range to the target is always positive). These Sigma Points are propagated through the full non-linear system equations and the estimate mean and covariance are then recovered. This allows the statistics of the time update step and measurement update to be captured more accurately.

This implementation allows generation of a world model for a 6DOF vehicle with a limited sensor suite (IMU and monocular vision only) operating in close proximity to the objects being observed. There remains two issues which must be addressed: data association and feature initialization.

Data Association—A key problem with navigation in an environment with indistinguishable features is data association. This occurs in any SLAM problem where features are not uniquely identifiable and for Kalman-filter based SLAM implementations correct data association is critical for a convergent navigation solution. Indeed, a central assumption of Kalman filters is correct data association.

Most data association algorithms are ultimately based on maximizing an expectation. To obtain a global maximum one must find the association which maximizes the expectation over all possible associations. To maximize the expectation we minimize the error between the actual and predicted measurements:

$$E = (\mathbf{z} - \mathbf{A}\bar{\mathbf{z}})^T \mathbf{P}_{zz}^{-1} (\mathbf{z} - \mathbf{A}\bar{\mathbf{z}}) \quad (10)$$

where \mathbf{A} is the $(m \times n)$ data association matrix, where an element has value 1 if a measurement is associated with a feature and 0 otherwise.

As the number of features and measurements grows this quickly becomes intractable to compute. A second possible approach is to minimize the error of individual data associations, in effect uncoupling the problem.

$$E_{ij} = (\mathbf{z}_i - \bar{\mathbf{z}}_j)^T \mathbf{P}_{zz,ij}^{-1} (\mathbf{z}_i - \bar{\mathbf{z}}_j) \quad (11)$$

It is less likely that a global minimum will be found in this case, but this has the advantage of being quick to compute. Problems arise when there are multiple possible associations with similar errors.

In the current implementation the likelihood that a measurement is associated with an existing feature is computed by first evaluating the expectations E_{ij} . Those which satisfy

$$E_{ij} < \gamma \quad (12)$$

are kept as potential associations. A nearest neighbor approach is then used to assign correlation for each measurement. While this approach is brittle, it has been adequate for this application.

When a particular bearing is not assigned to any current feature the estimate state vector is augmented with a new feature. The choice of γ thus incorporates a probability that a particular measurement will be assumed to refer to a new feature: a large value of γ reflects a belief that a bearing is likely to be associated with a previously seen feature, while a small value of γ increases the likelihood that a new feature will be created.

Feature Initialization—Landmark initialization for bearings-only SLAM has proven to be a difficult issue to resolve in a consistent manner. However solutions have been proposed which defer inclusion of a new feature until an appropriate level of ‘‘Gaussian-ness’’ has been achieved in its position estimate [1] or which define the feature’s position along a line [3]. For the problem of UAV flight in a forest we implement a heuristic: given the estimate of the vehicle’s altitude, a new feature is generated at the intersection of the bearing to the feature and the ground plane. The uncertainty in the new feature position is obtained from the uncertainty in vehicle state and the noise associated with the bearing measurement. While this does not give a solution which is strictly consistent (the correlation between the uncertainty in the new feature position and the vehicle state uncertainty is currently ignored) it does allow good localization of new features.

If γ (in essence a likelihood that a bearing is associated with a previously seen feature) is chosen to be too small, new features may be spuriously created. This can be seen in the map when several features are generated very close together, indicating that they are most likely multiple instances of what is actually a single true feature on the ground. While spurious features can be detected and combined, this is not implemented in the present work.

4. SIMULATION RESULTS

We present results of simulation of low-altitude flight (5m) in an unknown environment to shown navigation and mapping, results of monte-carlo simulation of low-altitude flight to show estimate consistency and results of monte-carlo simulation at both low and medium (100m) altitude flight to show error growth.

Navigation and Mapping

The UAV is flown in a circular trajectory of radius 20m at a velocity of 10m/s and an altitude of 5m through an initially

unknown forest of randomly distributed trees. Initial vehicle localization is assumed to be precise. In effect, this defines a coordinate frame relative to the vehicle’s initial position.

A sequence of images is shown in Figure 2. In the first frame (t=0.1s) only 4 features have been seen and mapped: note the large uncertainty in feature positions. As the vehicle follows its trajectory feature localization becomes more precise. Eventually all but 5 of the features have been mapped (these 5 features never came into view of the camera), and as seen in Figure 3 the vehicle has maintained a consistent estimate of its own position.

Monte Carlo Simulation

The UAV is flown in a circular trajectory of radius 20m at a velocity of 10m/s and an altitude of 5m through a forest of randomly distributed trees. A new random forest was generated at the start of each run.

Consistency of the estimate is checked by comparing the 2-norm of the estimate error with the estimated variance:

$$D_k = \frac{\mathbf{E}\sqrt{(\hat{\mathbf{x}}_k - \mathbf{x}_k)^T(\hat{\mathbf{x}}_k - \mathbf{x}_k)}}{\mathbf{E}\sqrt{\text{Tr}P_{kk}}} \quad (13)$$

where, for a consistent estimator:

$$D_k = 1 \quad (14)$$

This is plotted in Figure 4. Note that the ratio of true error and estimated variance, while greater than unity, is small. Closer inspection of the results revealed that a significant part of the difference was due to generation of multiple features, which the current implementation does not check. In any case, the ratio of true error and estimated variance does not show the unbounded growth associated with a divergent estimator, indicating that the estimates are generally reliable.

Error Growth

Error can be characterized as a percentage of distance traveled or as an absolute position error. For long flights through new terrain (exploration) it is most appropriate to express vehicle position error as a function of distance traveled, while for station keeping flights (i.e. orbiting around a fixed point) position error relative to the desired orbit point is more appropriate.

This paper does not address the question of optimizing trajectories to minimize navigation error: constant altitude, constant speed circular trajectories were chosen arbitrarily for ease of implementation and visualization.

Exploration—In exploratory flight the UAV is likely to traverse long distances. This may occur at low (i.e. nap of the earth) or at higher altitudes (~ 100 m). In this mode new features are continually coming into view of the camera while ‘old’ features leave the field of view: in effect the vehicle navigates using visual odometry. Error growth is affected by




Figure 2. Simulation showing navigation and mapping in an unknown environment. True feature locations are shown as blue dots, estimated feature locations and associated 3σ error ellipsoids are shown in red. True vehicle location is shown in blue, estimated vehicle position is shown in red. The dotted blue and red lines show vehicle position history over the previous 5 seconds for the true and estimated states, respectively.




Figure 3. Vehicle position error and 3σ error bounds. Note that the vehicle position estimate generally remains within the 3σ error bounds, indicating consistency of the vehicle position estimate.






Figure 4. Estimate consistency. The solid line shows the mean ratio of expectation of the true estimate error and the expectation of the computed variance. For an ideal estimator this ratio is unity. The variation from unity in this case is in large part due to generation of multiple features



(a) Absolute error growth with time



(b) Percent error as function of distance

Figure 5. Monte Carlo simulation results showing mean vehicle position error growth as function of time (a) and mean percent error as function of distance traveled. The vehicle flies a circle of radius 20m at a speed of 10m/s at 5m altitude. After initialization error grows linearly with distance traveled (approximately 1.7% of distance traveled) until previously seen features are re-observed (occurring at 10 seconds). The cyclically varying error is caused by initial uncertainty in vehicle heading.

many factors: the number of features in view at any given time, the length of time a particular feature remains in view and the video frame rate are all critical parameters. These in turn are affected by feature density, vehicle speed and altitude and camera field of view.

Figure 5 shows both absolute error growth as a function of time and error as a percentage of distance traveled for a Monte Carlo simulation of low altitude flight. While new terrain is being explored error grows linearly with time (i.e. error as a percentage of distance traveled is constant) until previously seen features are re-acquired, leading to a large reduction in error. For the cases investigated here the error growth was approximately 1.7% of distance traveled.

In addition to these Monte Carlo simulation results, figure 6 shows percentage error for a longer exploration flight. Again the position error as a percentage of distance traveled remains roughly constant.




Figure 6. Vehicle position error growth as a percentage of distance traveled for a single longer duration “exploration” run. The percent error remains approximately constant over the duration of the run.




Figure 7. Mean absolute error for 100 run Monte Carlo simulation of station keeping. The cyclic variation in error is caused by initial uncertainty in vehicle heading.

Station keeping—Station keeping is more likely to occur with the vehicle flying at higher altitudes (i.e. above nap of the earth) orbiting a fixed point. In this case it is more appropriate to express error relative to the desired orbit point. Over several orbits the same features will be seen repeatedly, enabling cyclic reductions in vehicle position error. Alternatively, if the camera is pointed towards the center of the orbit the same features may remain in view continuously, further reducing position error. Again the error characteristics are dependent on the number of features in view, the length of time a particular feature remains in view and the video frame rate.

Simulations consisting of multiple orbits about the origin with the camera pointed inwards were conducted. Orbit radius was 200m, speed was 20m/s, altitude was 100m. A set of five features were randomly placed for each run so that each feature would be in continuous view of the camera. Each run was initialized with random vehicle position and velocity error. A single run consisted of approximately 4.2 orbits. Figure 7 shows mean absolute error. After stabilization the average vehicle position error was observed to vary cyclically over each orbit.

A single simulation for a longer flight (1 hour) was conducted. A set of 20 features was randomly distributed over an area of 250m×250m with the vehicle flying a circle of radius 200m. Approximately 57 orbits were performed and the




Figure 8. Absolute position error for longer duration (one hour) station keeping. Green line shows absolute position error, blue line shows computed estimate error. The cyclic variation (period equal to orbit period) is caused by the initial heading uncertainty.

error varied as a random walk superimposed on a cyclic error. Absolute position error is shown in figure 8, errors in x , y , and z (altitude) are shown in figure 9.

5. HARDWARE TEST RESULTS

Using a Crossbow DMU-6X inertial measurement unit [2] and a Sony XC-999 color camera navigation in a laboratory environment using pylons as features was demonstrated. Color segmentation is used to identify features in the image and a bearing to the centroid of each pylon in view is returned by the vision system.

At the beginning of the trajectory the camera/IMU pair was kept level and stationary for 20 seconds while acceleration and angular rate data was collected. The average and standard deviation was calculated and used as initial values of bias and of measurement noise. Inertial data was collected at a rate of 50 Hz, vision updates occurred at 10 Hz.

Following initialization the camera/IMU pair was hand carried through the lab along a curved path approximately 8 meters in length, taking approximately 10 seconds. Truth data is available for each feature location and for the vehicle at the start and end points of the trajectory.

Figure 10 shows a sequence of camera images and vehicle position and landmark locations for navigation with unknown features in the laboratory environment. The run took approximately 20 seconds to complete and followed a path of approximately 8 meters. Table 1 gives results of vehicle position error at the end of the run and worst-case feature localization error over the entire run. Over the eight meter trajectory the final vehicle position error was 3.25% of the distance traveled.

6. CONCLUSION

This paper has presented simulation and preliminary hardware test results demonstrating navigation and mapping in an unknown environment using only a monocular camera and an




Figure 9. Components of vehicle position error for longer duration (one hour) station keeping. Green line shows true error, dotted blue lines show computed 3σ error bounds.

Table 1. Navigation and mapping results, laboratory test.

	$ \mathbf{x} - \hat{\mathbf{x}} $	σ	$\frac{ \mathbf{x} - \hat{\mathbf{x}} }{\sigma}$
final position	26 cm	8.9 cm	2.93
absolute feature position	57 cm	17 cm	3.35
“least consistent” feature position	25 cm	5.8 cm	4.3

inertial measurement unit. Results of 3D simulation show that consistent, unbiased estimates of vehicle and object states can be obtained, and this has been demonstrated in the laboratory.


The appropriate measure of error growth is dependent on the vehicle’s mode of operation: for exploration error is expressed as a percentage of distance traveled, for station keeping (i.e. orbits around a fixed point) absolute error is used. Error growth is dependent on several factors including vehicle speed, feature density and video frame rate. The laboratory test resulted in a vehicle position error of less than 4% of the total distance traveled.

7. ACKNOWLEDGEMENTS

Portions of this work were carried out using computation facilities provided by Santa Clara University.




(a)t=8 seconds



(b)t=8 seconds



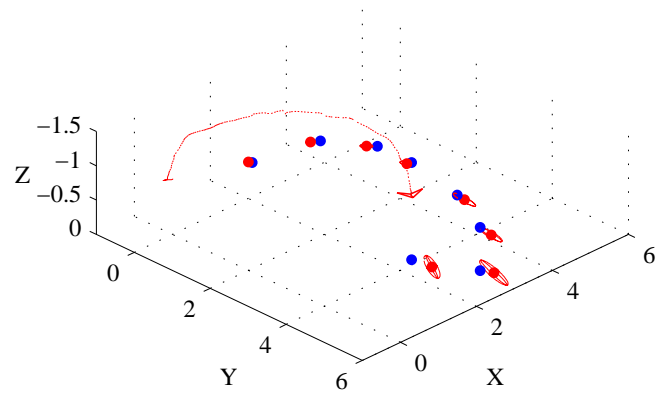
(c)t=13 seconds



(d)t=13 seconds



(e)t=18 seconds



(f)t=18 seconds

Figure 10. Navigation and mapping in a laboratory environment. The left column show a representative sequence of images from the camera. Color segmentation is used to identify pylons (used as navigation features), and the vision system provides bearings to the centroid of each pylon. The right column shows the sequence of estimated vehicle and feature locations. True feature locations are shown as blue dots, estimated feature locations and associated 3σ error ellipsoids are shown in red.

Estimated vehicle position is shown in red. The dotted line shows the time history of the estimated vehicle position.

REFERENCES

- [1] Tim Bailey. Constrained initialisation for bearing-only slam. In *IEEE International Conference on Robotics and Automation (ICRA)*, Taipei, Taiwan, 2003. IEEE.
- [2] Crossbow Technology Inc., 41 East Daggett Drive, San Jose CA 95134. *DMU User's Manual*, 1.3 edition, February 1999.
- [3] Andrew J. Davison. Real-time simultaneous localisation and mapping with a single camera. In *International Conference on Computer Vision*, Nice, France, October 2003.
- [4] Greg Duckworth. Gps-free navigation. <http://www.darpa.mil/DARPAtech2004/pdf/scripts/DuckworthGPSScript.pdf>, 2004.
- [5] Eric Foxlin and Leonid Naimark. Vis-tracker: A wearable vision-inertial self-tracker. In *IEEE VR2003*, Los Angeles, California USA, March 2003. IEEE, IEEE.
- [6] Andreas Huster and Steve Rock. Relative position sensing by fusing monocular vision and inertial rate sensors. In *11th International Conference on Advanced Robotics (ICAR)*, 2003.
- [7] 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*, 45(3):477–482, March 2000.
- [8] Isaac Kaminer, Antonio Pascoal, and Wei Kang. Integrated vision/inertial navigation system design using nonlinear filtering. In *Proceedings of the American Control Conference*, San Diego, California, June 1999.
- [9] Jong Hyuk Kim and Salah Sukkarieh. Recasting slam—towards improving efficiency and platform independency. In *11th International Symposium of Robotics Research (ISRR)*, Siena, Italy, 2003.
- [10] Jack Langelaan and Steve Rock. Navigation of small uavs operating in forests. In *AIAA Guidance, Navigation and Controls Conference*, Providence, RI USA, August 2004. American Institute of Aeronautics and Astronautics.
- [11] Jorge Lobo and Jorge Dias. Integration of inertial information with vision. In *24th Conference of IEEE Industrial Electronics Society (IECON)*, pages 1263–1267, Aachen, Germany, 1998. IEEE.
- [12] Thomas Netter and Nicolas Franceschini. A robotic aircraft that follows terrain using a neuromorphic eye. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Lausanne, Switzerland, 2002.
- [13] Stergios I. Roumeliotis, Andrew E. Johnson, and James F. Montgomery. Augmenting inertial navigation with image-based motion estimation. In *IEEE International Conference on Robotics and Automation (ICRA)*, Washington, DC, 2002. IEEE.
- [14] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. A real-time algorithm for mobile robot mapping with applications to multi-robot and 3d mapping. In *IEEE International Conference on Robotics and Automation (ICRA)*, San Francisco, CA, 2000. IEEE.
- [15] Rudolph van der Merwe, Eric Wan, and Simon Julier. Sigma point kalman filters for nonlinear estimation and sensor fusion: Applications to integrated navigation. In *AIAA Guidance, Navigation and Controls Conference*, Providence, RI USA, August 2004. American Institute of Aeronautics and Astronautics.
- [16] Stefan B. Williams, Gamini Dissanayake, and Hugh Durrant-Whyte. Field deployment of the simultaneous localisation and mapping algorithm. In *15th IFAC World Congress on Automatic Control*, Barcelona, Spain, June 2002.



Jack Langelaan is a Ph.D. candidate in the Aerospace Robotics Laboratory at Stanford University. His research interests include estimation and control for individual robots and teams of cooperating robots. Mr. Langelaan obtained an M.S.A.A. from the University of Washington in 1994. Prior to joining the ARL he worked at Bombardier Aerospace in Toronto. Mr. Langelaan is a member of the American Institute of Aeronautics and Astronautics and the Canadian Aeronautics and Space Institute.



Steve Rock is Professor of Aeronautics and Astronautics and Director of the Aerospace Robotics Laboratory at Stanford University. Prof. Rock's research interests include the development and experimental verification of advanced control techniques for robotic and vehicle systems. Areas of emphasis include underwater remotely operated vehicle control, planning for teams of cooperating robots and aircraft systems. Prior to joining the Stanford faculty, Dr. Rock led the Controls and Instrumentation Department of Systems Control Technology, Inc. Prof. Rock is an Associate Fellow of the American Institute of Aeronautics and Astronautics.