# Vision-Based Estimation for Guidance, Navigation, and Control of an Aerial Vehicle

**M. K. KAISER**
Air Force Research Laboratory

**N. R. GANS,** Member, IEEE
University of Texas, Dallas

**W. E. DIXON,** Senior Member, IEEE
University of Florida

**While a Global Positioning System (GPS) is the most widely used sensor modality for aircraft navigation, researchers have been motivated to investigate other navigational sensor modalities because of the desire to operate in GPS denied environments. Due to advances in computer vision and control theory, monocular camera systems have received growing interest as an alternative/collaborative sensor to GPS systems. Cameras can act as navigational sensors by detecting and tracking feature points in an image. Current methods have a limited ability to relate feature points as they enter and leave the camera field of view (FOV).**

**A vision-based position and orientation estimation method for aircraft navigation and control is described. This estimation method accounts for a limited camera FOV by releasing tracked features that are about to leave the FOV and tracking new features. At each time instant that new features are selected for tracking, the previous pose estimate is updated. The vision-based estimation scheme can provide input directly to the vehicle guidance system and autopilot. Simulations are performed wherein the vision-based pose estimation is integrated with a nonlinear flight model of an aircraft. Experimental verification of the pose estimation is performed using the modelled aircraft.**

Manuscript received August 8, 2008; revised February 12, 2009; released for publication March 2, 2009.

IEEE Log No. T-AES/46/3/937960.

Refereeing of this contribution was handled by M. Braasch.

This research is supported in part by the NSF CAREER award CMS-0547448, AFOSR Contracts F49620-03-1-0381 and F49620-03-1-0170, AFRL Contract FA4819-05-D-0011, and by Research Grant US-3715-05 from BARD, the United States–Israel Binational Agricultural Research and Development Fund at the University of Florida.

Authors' addresses: M. K. Kaiser, Air Force Research Laboratory, Munitions Directorate, Eglin AFB, FL 32542-6810; N. R. Gans and W. E. Dixon, Department of Mechanical and Aerospace Engineering, University of Florida, Gale-Lembrand Dr., Gainesville, FL 32611, E-mail: (wdixon@ufl.edu).



## I. INTRODUCTION

Global Positioning System (GPS) is the primary navigational sensor modality used for vehicle guidance, navigation, and control. However, a comprehensive study, referred to as the Volpe report [1], indicates several vulnerabilities of GPS associated with signal disruptions. The Volpe Report delineates the sources of interference with the GPS signal into two categories, unintentional and deliberate disruptions. Some of the unintentional disruptions include ionosphere interference (also known as ionospheric scintillation) and RF interference (broadcast television, VHF, cell phones, and two-way pagers); whereas, some of the intentional disruptions involve jamming, spoofing, and meaconing. Some of the ultimate recommendations of this report were to, "create awareness among members of the domestic and global transportation community of the need for GPS backup systems…" and to "conduct a comprehensive analysis of GPS backup navigation…" which included instrument landing systems (ILS), long range navigation (LORAN), and inertial navigation systems (INS) [1].

The Volpe report inspired a search for strategies to mitigate the vulnerabilities of the current GPS navigation protocol. Nearly all resulting strategies followed the suggested GPS backup methods that revert to archaic/legacy methods. Unfortunately, these navigational modalities are limited by the range of their land-based transmitters, which are expensive and may not be feasible for remote, hazardous, or adversarial environments. Based on these restrictions, researchers have investigated local methods of estimating position when GPS is denied.

Given the advancements in computer vision and estimation and control theory, monocular camera systems have received growing interest as a local alternative/collaborative sensor to GPS systems. One issue that has inhibited the use of a vision system as a navigational aid is the difficulty in reconstructing inertial measurements from the projected image. Current approaches to estimating the aircraft state through a camera system utilize the motion of feature points in an image. A geometric approach is proposed (and our preliminary results [2, 3]) that uses a series of homography relationships to estimate position and orientation with respect to an inertial pose. This approach creates a series of "daisy-chained" pose estimates (see [4] and [5]), in which the current feature points can be related to previously viewed feature points to determine the current coordinates between each successive image. Through these relationships previously recorded GPS data can be linked with the image data to provide position measurements in navigational regions where GPS is denied. The method also delivers an accurate estimation of vehicle attitude, which is an

open problem in aerial vehicle control. The position and attitude (i.e., pose) estimation method can be executed in real time, making it amenable for use in closed-loop guidance control of an aircraft.

The concept of vision-based control for a flight vehicle has been an active area of research over the last decade. Recent literature focused on vision-based state estimation for use in control of a flight vehicle can be categorized by several distinctions. One distinction is that some methods require simultaneous sensor fusion [6, 7], while other methods rely solely on camera feedback [8]. Research can further be categorized into methods that require a priori knowledge of landmarks (such as pattern or shape [9, 10], light intensity variations [11], and runway edges or lights [12, 13]) versus techniques that do not require any prior knowledge of landmarks [14–16].

Another category of research includes methods that require the image features to remain in the field of view (FOV) [14] versus methods that are capable of acquiring new features [17]. Finally, methods can be categorized according to the vision-based technique for information extraction, such as: optic flow [18], simultaneous localization and mapping (SLAM) [19, 20], stereo vision [21], or epipolar geometry [14, 22–24]. This last category might also be delineated between methods that are more computationally intensive and therefore indicative of the level of real-time, on-board, computational feasibility.

Methods using homography relationships between images to estimate the pose of an aircraft are presented by Caballero et al. [24] and Shakernia et al. [14] (where it is referred to as the "planar essential matrix"). The method presented by Caballero et al. is limited to flying above a planar environment, and this method creates an image mosaic, which can be costly in terms of memory. Shakernia's approach does not account for feature points entering and exiting the camera FOV. The method developed here is designed for use with a fixed wing aircraft, thus the method explicitly acquires new feature points when the current features risk leaving the image, and no target model is needed, as compared to other methods [9–13]. The requirement of flying over a constant planar surface is also relaxed to allow flight over piecewise planar patches, which is more characteristic of real-world scenarios.

The efforts used here share some concepts associated with visual SLAM (VSLAM). There is no strict definition of VSLAM, and there are many different approaches. Some authors (e.g., [25–27]) make a distinction between "local VSLAM" and "global VSLAM." In this categorization local VSLAM is concerned with estimating the current state of the robot and world map through matching visual features from frame to frame, and global VSLAM is concerned with recognizing when features have been previously encountered and with updating estimates of the robot and map (sometimes referred to as "closing loops"). To address both issues many researchers use invariant features, such as scale-invarient feature transform (SIFT) [28], which can be accurately matched from frame to frame or from multiple camera viewpoints. Many VSLAM approaches use probabilistic filters (e.g., an extended Kalman filter or a particle filter) [26, 27, 29–31], that typically estimate a state vector composed of the camera/robot position, orientation and velocity, and the 3D coordinates of visual features in the world frame. The use of epipolar geometry [25, 32] is an option to filter-based approaches. A final possible category is methods that build a true 3D map (i.e., a map that is easily interpreted by a human being, such as walls or topography) [25–27, 30, 31] and those that build a more abstract map that is designed to allow the camera/robot to accurately navigate and recognize its location but that is not designed for human interpretation.

The method presented can provide an estimate of the position and attitude of unmanned air vehicle (UAV), and the method can be extended to map the location of static landmarks in the world frame. Hence this approach can be used in VSLAM of the UAV, with applications toward path planning, real time trajectory generation, obstacle avoidance, multi-vehicle coordination control and task assignment, etc. By using the daisy-chaining strategy, the coordinates of static features that have moved out of the FOV can also be estimated. The estimates of static features can be maintained as a map, or they can be used as measurements in existing VSLAM methods. The daisy-chaining method in this paper is best suited to a local VSLAM approach, where the global problem of closing loops can be addressed though existing methods. Unlike typical VSLAM approaches camera/vehicle odometry is not required as an input for the approach in this paper.

To investigate the performance of the developed method, a numerical simulation is provided for a nonlinear, six degrees-of-freedom model of an Osprey UAV. The simulation illustrates the ability of the estimation method to reconstruct the UAV pose in the presence of disturbances, such as errors in the initial altitude measure, image quantization, and noise. To illustrate the potential use of the estimates in a feedback loop, an autopilot was also included in the simulation, with inputs complimentary with the outputs from the estimation method. A specific maneuver is created to perform a simultaneous rolling, pitching, and yawing motion of the aircraft, combined with a fixed-mounted camera. The aircraft/autopilot modeling effort and maneuver are intended to test the robustness of the homography-based estimation method as well as to provide proof-of-concept in

using the camera as the primary sensor for achieving closed-loop autonomous flight.

Based on the outcomes from the simulation, the performance of the developed estimation method is also experimentally tested through two flight tests with an Osprey UAV. One experiment compares the estimated inertial coordinates of the UAV to the GPS position reported by two on-board GPS units. The GPS units provide different outputs, but the homography-based estimation is shown to provide approximately equivalent performance. A second experiment is also performed, where precision laser-measured ground markers were viewed by the flight camera. In this experiment, the images of the known maker locations are used to generate a precise ground truth. The estimation method (which did not use any information about marker locations) matches the reconstructed ground truth data.

Section II provides details on pose reconstruction using the homography-based daisy-chaining method. Simulations are presented in Section III to demonstrate the performance of the method, and Section IV includes the experimental results. Section V describes potential future efforts to improve the developed method for broader application.

## II. POSE RECONSTRUCTION FROM TWO VIEWS

### A. Euclidean Relationships

Consider a body-fixed coordinate frame $\mathcal{F}_c$ that defines the position and attitude of a camera, with respect to a constant world frame $\mathcal{F}_w$. The world frame could represent a departure point, destination, or some other point of interest. The rotation and translation of $\mathcal{F}_c$, with respect to $\mathcal{F}_w$, is defined as $R(t) \in \mathbb{R}^{3\times 3}$ and $x(t) \in \mathbb{R}^3$, respectively. The camera rotation and translation of $\mathcal{F}_c$ between two time instances, $t_0$ and $t_1$, is denoted by $R_{01}(t_1)$ and $x_{01}(t_1)$. During the camera motion a collection of $I$ (where $I \geq 4$) coplanar and noncolinear static feature points are assumed to be visible in a plane $\pi$. The assumption of four coplanar and noncolinear feature points is only required to simplify the subsequent analysis and is made without loss of generality. Image processing techniques can be used to select coplanar and noncolinear feature points within an image such as in [33] and [34]. However, if four coplanar target points are not available, then the subsequent development can also exploit a variety of linear solutions for eight or more noncoplanar points (e.g., the classic eight points algorithm [35, 36] or virtual parallax [37, 38]) or nonlinear solutions for five or more points [39].

A feature point $p_i(t)$ has coordinates $\bar{m}_i(t) = [x_i(t), y_i(t), z_i(t)]^T \in \mathbb{R}^3 \ \forall i \in \{1 \ldots I\}$ in $\mathcal{F}_c$. Standard geometric relationships can be applied to the coordinate systems depicted in Fig. 1 to develop the following relationships:

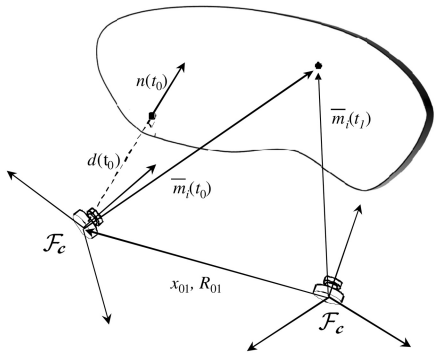


Fig. 1. Euclidean relationships between two camera poses.

$$\bar{m}_i(t_1) = R_{01}(t_1)\bar{m}_i(t_0) + x_{01}(t_1)$$
$$\bar{m}_i(t_1) = \underbrace{\left(R_{01}(t_1) + \frac{x_{01}(t_1)}{d(t_0)} n(t_0)^T\right)}_{H(t_1)} \bar{m}_i(t_0) \tag{1}$$

where $H(t_1)$ is the Euclidean homography matrix at time $t_1$, $n(t_0)$ is the constant unit vector normal to the plane $\pi$ from $\mathcal{F}_c$ at time $t_0$, and $d(t_0)$ is the constant distance between the plane $\pi$ and $\mathcal{F}_c$ along $n(t_0)$. After normalizing the Euclidean coordinates as

$$m_i(t) = \frac{\bar{m}_i(t)}{z_i(t)} \tag{2}$$

the relationship in (1) can be rewritten as

$$m_i(t_1) = \underbrace{\frac{z_i(t_0)}{z_i(t_1)}}_{\alpha_i} H(t_1) m_i(t_0) \tag{3}$$

where $\alpha_i \in \mathbb{R} \ \forall i \in \{1 \ldots I\}$ is a scaling factor. Further details on the Euclidean homography can be found in [40] and [41].

### B. Projective Relationships

Using standard projective geometry the Euclidean coordinate $\bar{m}_i(t)$ can be expressed in image-space pixel coordinates as $p_i(t) = [u_i(t), v_i(t), 1]^T$. The projected pixel coordinates are related to the normalized Euclidean coordinates, $m_i(t)$ by the pin-hole camera model as [40]

$$p_i = A m_i \tag{4}$$

where $A$ is an invertible, upper triangular camera calibration matrix defined as

$$A \triangleq \begin{bmatrix} a & -a\cos\phi & u_0 \\ 0 & \dfrac{b}{\sin\phi} & v_0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{5}$$

In (5), $u_0$ and $v_0 \in \mathbb{R}$ denote the pixel coordinates of the principal point (the image center as defined by the intersection of the optical axis with the image

plane), $a, b \in \mathbb{R}$ represent scaling factors of the pixel dimensions, and $\phi \in \mathbb{R}$ is the skew angle between camera axes.

By using (4), the Euclidean relationship in (3) can be expressed as

$$p_i(t_1) = \alpha_i A H(t_1) A^{-1} p_i(t_0) = \alpha_i G(t_1) p_i(t_0). \quad (6)$$

Sets of linear equations can be developed from (6) to determine the projective and Euclidean homography matrices $G(t_1)$ and $H(t_1)$ up to a scalar multiple. Various techniques [41, 42] can be used to decompose the Euclidean homography to obtain $\alpha_i(t_1)$, $n(t_0)$, $x_{01}(t_1)/d(t_0)$, and $R_{01}(t_1)$. The decomposition methods in [41] and [42] generally return two physically valid solutions. The correct solution can be chosen by knowledge of the correct value for the normal vector or by using an image taken at a third pose. The distance $d(t_0)$ must be separately measured (e.g., through an altimeter or radar range finder) or estimated using a priori knowledge of the relative feature point locations or as an estimator signal in a feedback control.

### C. Chained Pose Reconstruction for Aerial Vehicles

Consider an aerial vehicle equipped with a GPS and a camera capable of viewing a landscape. A technique is developed in this section to estimate the position and attitude using camera data when the GPS signal is denied. A camera has a limited FOV, and motion of a vehicle can cause observed feature points to leave the image. The method presented here chains together pose estimations from sequential sets of tracked of points. This approach allows the system to halt tracking a set of image features if the set is likely to leave the image and allows the system to begin tracking a new set of features while maintaining the pose estimate. Thus the estimation can continue indefinitely and is not limited by the camera's FOV.

The subsequent development assumes that the aerial vehicle begins operating at time $t_0$, where the translation and rotation (i.e., $x_0(t_0)$ and $R_0(t_0)$ in Fig. 2) between $\mathcal{F}_c$ and $\mathcal{F}_w$ at time $t_0$ are known. The rotation between $\mathcal{F}_c$ and $\mathcal{F}_w$ at time $t_0$ can be determined through the bearing information of the GPS, along with other sensors such as a gyroscope and/or compass. Without loss of generality the GPS unit is assumed to be fixed to the origin of the aerial vehicle's coordinate frame, and the constant position and attitude of the camera frame are known with respect to the position and attitude of the aerial vehicle coordinate frame. The subsequent development further assumes that the GPS is capable of delivering altitude, perhaps in conjunction with an altimeter, so that the altitude $a(t_0)$ is known.

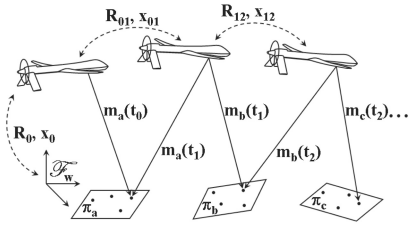


Fig. 2. Illustration of pose estimation chaining.

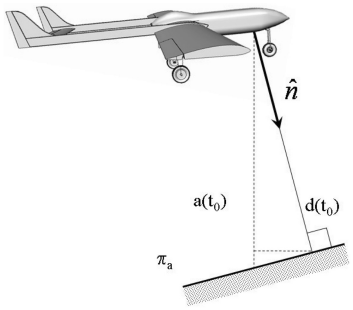


Fig. 3. Depth estimation from altitude.

As illustrated in Fig. 2, when the camera is initially at a pose $\{x_0(t_0), R_0(t_0)\}$, the initial set of tracked coplanar and noncolinear feature points are contained in the plane $\pi_a$. The plane $\pi_a$ is perpendicular to the unit vector $n_a(t_0)$ in the camera frame, and the plane lies at a distance $d_a(t_0)$ from the camera frame origin. The feature points lying in $\pi_a$ have Euclidean coordinates $\bar{m}_{ai}(t_0) \in \mathbb{R}^3 \ \forall i \in \{1 \ldots I\}$ in $\mathcal{F}_c$. An image is captured at this time, resulting in image points with pixel coordinates $p_a(t_0)$. At time $t_1$ the vehicle has undergone some rotation $R_{01}(t_1)$ and translation $x_{01}(t_1)$. At this time the points in $\pi_a$ have Euclidean coordinates $\bar{m}_{ai}(t_1) \in \mathbb{R}^3 \ \forall i \in \{1 \ldots I\}$ in $\mathcal{F}_c$. Another image is captured, resulting in image points with pixel coordinates $p_a(t_1)$.

As described in Section IIB, $R_{01}(t_1)$ and $x_{01}(t_1)/d_a(t_0)$ can be solved from two corresponding images of the feature points $p_a(t_0)$ and $p_a(t_1)$. A measurement or estimate for $d_a(t_0)$ is required to recover $x_{01}(t_1)$. This estimation is possible with distance sensors or with a priori knowledge of the geometric distances between the points in $\pi_a$. However, with an additional assumption, it is possible to estimate $d_a(t_0)$ geometrically using altitude information from the last GPS reading and/or an altimeter. From the illustration in Fig. 3, if $a(t_0)$ is the height above $\pi_a$ (e.g. the slope of the ground is constant between the feature points and projection of the plane's location to the ground), then the distance $d_a(t_0)$ can be determined by the projection:

$$d_a(t_0) = n_a(t_0)^T a(t_0) \quad (7)$$

where $n_a(t_0)$ is known from the homography decomposition.

Once $R_{01}(t_1)$, $d_a(t_0)$, and $x_{01}(t_1)$ have been determined, the rotation $R_1(t_1)$ and translation $x_1(t_1)$

can be determined with respect to $\mathcal{F}_w$ as

$$R_1 = R_0 R_{01}$$
$$x_1 = R_{01} x_{01} + x_0.$$

As illustrated in Fig. 2, at time $t_1$, a new collection of feature points $m_{bi}(t_1)$ is visible in the planar patch denoted by $\pi_b$. Capturing an image of these points gives a set of image points with pixel coordinates $p_b(t_1)$ that can be tracked over time. At time $t_2$ the sets of points $p_b(t_1)$ and $p_b(t_2)$ can be used to determine $R_{12}(t_2)$ and $x_{12}(t_2)/d_b(t_1)$, which provides the rotation and scaled translation of $\mathcal{F}_c$ with respect to $\mathcal{F}_w$. If $\pi_b$ and $\pi_a$ are the same plane, then $d_b(t_1)$ can be determined as

$$d_b(t_1) = d_a(t_1) = d_a(t_0) + x_{01}(t_1) \cdot n(t_0). \tag{8}$$

When $\pi_b$ and $\pi_a$ are the same plane, $x_{12}(t_2)$ can be correctly scaled, and $R_2(t_2)$ and $x_2(t_2)$ can be computed in a similar manner as described for $R_1(t_1)$ and $x_1(t_1)$. These estimates can be propagated by chaining them together at each time instance without further use of the GPS.

In the general case $p_b$ and $p_a$ may not be coplanar, and (8) cannot be used to determine $d_b(t_1)$. If $p_b$ and $p_a$ are both visible for two or more frames, it is still possible to calculate $d_b(t)$ through geometric means. Let $t_{1^-}$ denote some time before the daisy-chain operation is performed, when both $p_b$ and $p_a$ are visible in the image. At time $t_{1^-}$ an additional set of homography equations can be solved for the points $p_b$ and $p_a$ at times $t_1$ and $t_{1^-}$

$$m_{ai}(t_1) = \alpha_a \left( \bar{R}_1(t_1) + \frac{\bar{x}_1(t_1) n_a(t_{1^-})^T}{d_a(t_{1^-})} \right) m_{ai}(t_{1^-}) \tag{9}$$

$$m_{bi}(t_1) = \alpha_b \left( \bar{R}_1(t_1) + \frac{\bar{x}_1(t_1) n_b(t_{1^-})^T}{d_b(t_{1^-})} \right) m_{bi}(t_{1^-}) \tag{10}$$

where $\bar{R}_1(t_1)$ and $\bar{x}_1(t_1)$ are the rotation and translation, respectively, that the camera undergoes from time $t_{1^-}$ to $t_1$, and

$$\alpha_a = \frac{z_{ai}(t_{1^-})}{z_{ai}(t_1)} \quad \text{and} \quad \alpha_b = \frac{z_{ai}(t_{1^-})}{z_{ai}(t_1)}.$$

Note that $\bar{R}_1(t_1)$ and $\bar{x}_1(t_1)$ have the same values in equations (9) and (10), but the distance and normal to the plane are different for the two sets of points. The distance $d_a(t_{1^-})$ is known from (8), and the scaled translations

$$x_a(t_1) = \frac{\bar{x}_1(t_1)}{d_a(t_{1^-})} \quad \text{and} \quad x_b(t_1) = \frac{\bar{x}_1(t_1)}{d_b(t_{1^-})}$$

can be recovered from the homography decomposition in (9) and (10). Given $d_a(t_{1^-})$, $x_a(t_1)$, and $x_b(t_1)$, the translation $\bar{x}_1(t_1)$ can be determined as

$$\bar{x}_1(t_1) = d_a(t_{1^-}) x_a(t_1)$$


Fig. 4. Air and Sea Inc. Osprey UAV used in experiments and modeled for simulations.

and $d_b(t_{1^-})$ can then be determined as

$$d_b(t_{1^-}) = \frac{x_b^T(t_1) \bar{x}_1(t_1)}{\|x_b(t_1)\|}.$$

The distance $d_b(t_1)$ can then be found by using (8), with $d_b(t_{1^-})$ in place of $d_a(t_0)$. Alternatively, additional sensors, such as an altimeter, can provide an additional estimate in the change in altitude. These estimates can be used in conjunction with (8) to update depth estimates. Note that an error in estimating the distance signals $d$ does not affect the orientation estimate, but it does cause a scaling error in the translation estimate. The effects of an error in estimate of $d(0)$ is investigated in Section III.

## III. SIMULATION RESULTS

To facilitate the subsequent flight tests, a high fidelity vehicle simulation was developed to investigate the feasibility of the proposed vision-based state estimation and guidance method. Prior to the simulation the Osprey fixed-wing UAV by Air and Sea Composites, Inc. (See Fig. 4) was selected as the experimental testbed because of cost and payload capacity factors and because of the fact that the pusher prop configuration is amenable to forward looking camera placement. Given that the Osprey UAV was going to be used as the experimental testbed, a full, nonlinear model of the equations of motion and aerodynamics of the Osprey UAV were developed (see [43] for details regarding the model development) along with an autopilot design. The autopilot design allowed for the vehicle to perform simple commanded maneuvers that an autonomous aircraft would typically be expected to receive from an on-board guidance system. The autopilot was constructed to use state estimates that came from the homography-based algorithm. Preliminary modal analysis of the Osprey vehicle flying at a 60 m altitude at 25 m/s indicated a short-period frequency

$\omega_{sp}$ = 10.1 rad/s and damping $\zeta_{sp}$ = 0.85; a phugoid mode frequency $\omega_{ph}$ = 0.34 rad/s and damping $\zeta_{ph}$ = 0.24; a dutch-roll frequency $\omega_{dr}$ = 4.39 rad/s and damping $\zeta_{dr}$ = 0.19; a roll subsidence time constant of $\tau_r$ = 0.08 s; and a spiral mode time-to-double *ttd* = 39.8 s. These rigid-body frequencies are crucial for the autopilot design as well for determining what, if any, of the state estimation values coming from the camera and proposed technique are favorable to be used in a closed-loop sense, due to the imminent phase lag associated with associated camera frame rate (sampling frequency) and phase lag resulting from the signal filtering.

In the simulations five patches of 4 feature points are programmed to lie along a 500 m ground track. For simplicity all planar patches lie in the same plane. The task is to perform the state estimation during a maneuver. The commanded maneuver is to simultaneously perform a 10 m lateral shift to the right and a 10 m longitudinal increase in altitude. This particular maneuver results in the vehicle simultaneously pitching, rolling, and yawing, while translating. For simulation purposes the camera is mounted underneath the fuselage looking downwards. The camera model is intended to be representative of a typical $640 \times 480$ lines of resolution charge-coupled device equipped with a 10 mm lens. To more accurately capture true system performance, pixel coordinates are rounded to the nearest integer to model errors due to camera pixilation effects (i.e., quantization noise), and furthermore a 5% error is added to the estimated vehicle altitude to test robustness. While some modeling errors are present in this analysis in order to address some practical issues, a full error analysis is intentionally avoided since statistical error analysis for camera-based systems is still somewhat ambiguous. Some of the error sources which are present in real-world systems include, but are not limited to: feature point selection/extraction/tracking algorithms, spatial separation of feature points in image plane, lens aberrations, camera calibration error, coplanarity of 4 feature points, distance to feature point plane, exposure time (effectively "shutter speed"), rolling shutter effects (if not a global shutter camera), interlacing effects (if not a progressive scan camera), and numerical errors associated with the construction and decomposition of the homography relationships.

The first simulation is designed to test the accuracy of the vision-based estimation. Vision is not used in the feedback for this maneuver, and the estimated pose is compared to the true pose. The results of this preliminary analysis are given in Figs. 5 and 6. The effects of noise are visible, but the estimated pose is accurate.

The second simulation is intended to examine the effects of using the vision-based estimate as a sensor in closed-loop control. This simulation involves replacing the perfect position and attitude measurements, used in the guidance system and autopilot, with position and attitude estimations determined from the vision-based method. The resulting control architecture and sensor suite for this UAV is given in Fig. 7. The noise content of the estimated position and attitude requires filtering prior to being used by the autopilot to prevent the high-frequency noise from being passed to the aircraft actuators. As expected the noise occurs at 30 Hz and corresponds to the frame rate of the camera. First-order, low pass-filters (cutoff frequency as low as 4 rad/s) are used to filter the noise. The noise also prevents effective differentiation of the position and attitude and necessitates the use of rate gyros for yaw and roll damping, as depicted in Fig. 7. The air data system is also included, as shown in Fig. 7, for the initial altitude measurement since it is more accurate for altitude than current GPS solutions. The results of the camera-in-the-loop system performing the same guidance-commanded autonomous maneuver are given in Figs. 8 and 9.

The simulation results indicate that a camera supplemented with minimal sensors, such as rate gyros and barometric altitude, can be used for completely autonomous flight of a fixed wing vehicle, however some residual oscillation effects due to noise are present in the vehicle attitude response. A majority of the noise source can directly be attributed to camera pixilation effects and the corresponding phase lag introduced by the first order filtering.

To test the performance of the system over long time periods, a Monte Carlo analysis is performed. This analysis includes a sequence of constant altitude "s-curves" for 1, 5, 10, 15, 30, 45, and 60 sec. Each s-curve takes 1 min to perform, and the UAV's forward velocity is constant at 10 m/s. Quantization noise is included. Each simulation is performed 1000 times, and mean and standard deviation of the pose and rotation error are plotted verses time and total distance traveled and total rotation performed. Total distance refers to the length of the curve followed, not the straight distance from start to finish. Similarly, as the rotation evaluates to sinusoidal values to achieve the s-curve, the total rotation refers to the rms of the angular values over time.

The results are seen in Figs. 10 and 11. The expected error for both simulations is roughly linear with the time/distance travelled. The expected pose error after an hour of flight and a travelled distance of 30 km is 0.66 km. The expected rotation error over this same time period is 0.21 rad. This amounts to an error of 2.2% of distance travelled and 0.39% of total rotation performed.

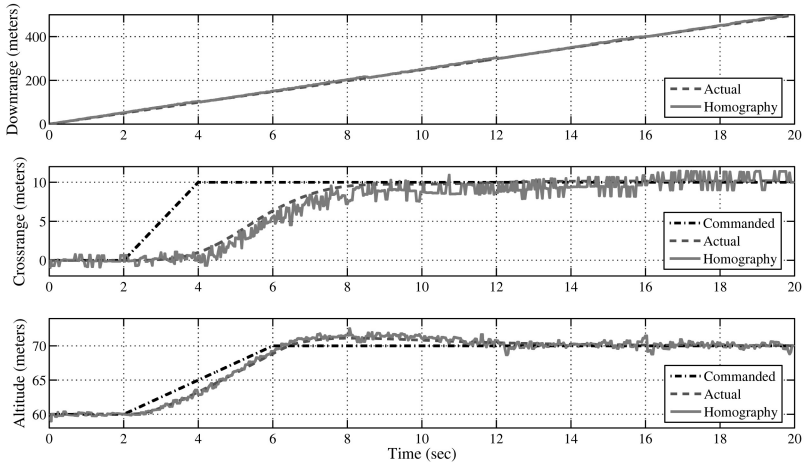


Fig. 5. Actual aircraft position versus estimated position. State estimation only.

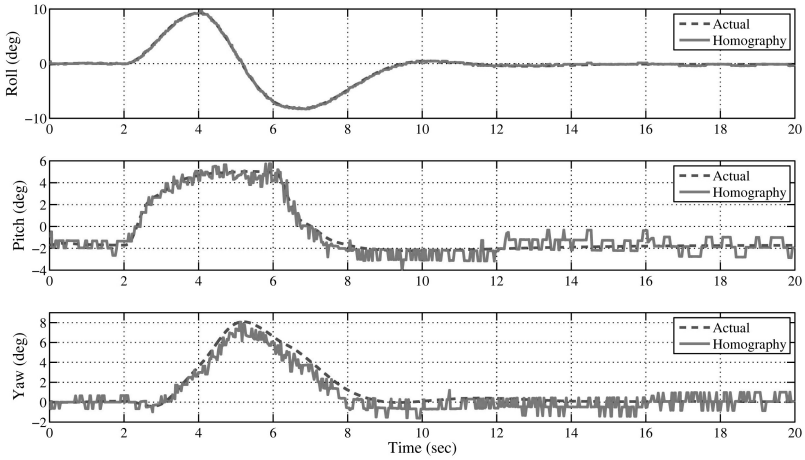


Fig. 6. Actual aircraft attitude versus estimated attitude. State estimation only.

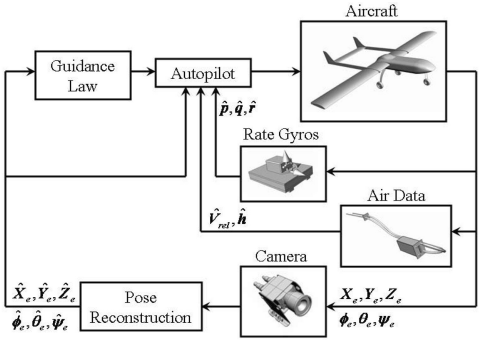


Fig. 7. Autonomous control architecture.

## IV. EXPERIMENTAL RESULTS

Based on the results of the simulation, two sets of experimental results are conducted to establish the feasibility of the proposed vision-based state estimation method. To give broad results each experiment features different equipment and different methods to generate a ground truth comparison. In both flight tests artificial, red markers (i.e., fiducial features) are placed along a stretch of the runway. The runway is fairly flat, thus each patch of points are roughly coplanar, although this was not assumed to be true in the experiment. Due to the orientation changes of the aircraft during the flight, the normal vector $n$ is estimated at each handoff. A camera mounted on the radio controlled UAV (Fig. 4) records the view of the runway; video data is captured using a digital, video tape recorder and analyzed offline. Sample video frames from the flights are provided in Figs. 14 and 16. The features are extracted and tracked using a color-based thresholding method. Fig. 12 represents a 3D surface constructed from a nonlinear combination of the red, green, and blue color space values of a particular image frame. Note that the large spikes correspond to the location of the red landmarks. An example of the output of the tracking algorithm is also given in Fig. 13. The four trajectories in this particular case represent the path of the landmarks from the first patch as it enters and exits the FOV (from top to bottom). As a point of interest, note that the center of the image plane does not correspond to the location of the optical axis. This apparent oddity is in

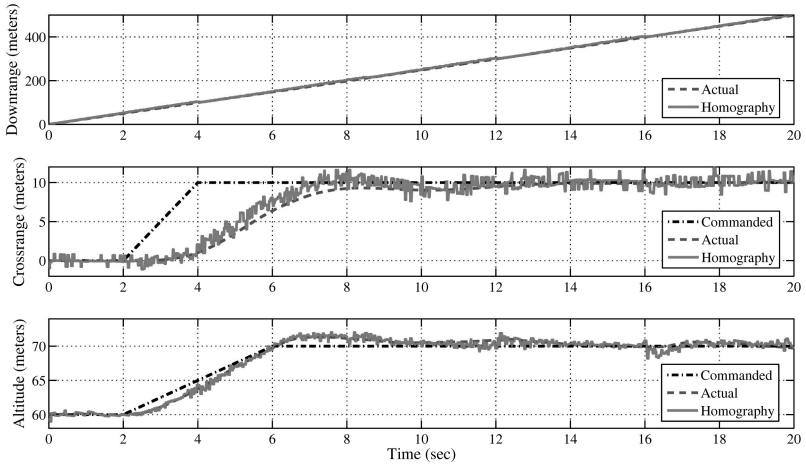


Fig. 8. Actual aircraft position versus estimated position in closed-loop control.

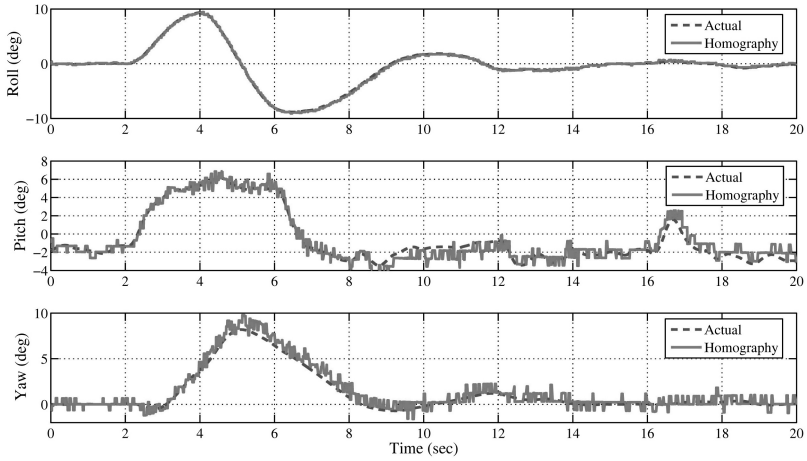


Fig. 9. Actual aircraft attitude versus estimated attitude in closed-loop control.

keeping with the camera calibration results, and it is most likely congruous with a low end imager and lens.

## A. Daisy Chained Homography versus GPS

The first flight experiment compares position estimates from the daisy-chained pose estimation to position, which is estimated using a Garmin GPS 35 receiver. A copy of the video is overlaid with GPS, as seen in Fig. 14, which shows the effects of video interlacing. A second GPS unit (manufactured by Eagle Tree Systems, LLC) is also onboard to test inter-GPS accuracy. The use of two GPS units provides a comparison for the vision-based method, which is intended to compute GPS-like information. Poor image quality, including focus, motion blur, and interlacing of the digital video, necessitates a downsampling of the feature extraction, by a factor of five, to reduce noise. The downsampling results in a 5 Hz input signal.

Results of the experiment are given in Fig. 15. In the legend for Fig. 15, GPS2 represents the overlaid GPS data, and GPS1 represents the onboard data

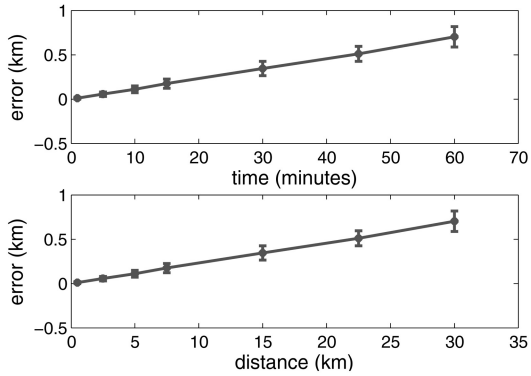


Fig. 10. Error plots of Monte Carlo analysis, showing mean position error and standard deviation of position error over long time period and long distance. Position error refers to norm of difference between known 3D position vector and estimated 3D position vector.

logger GPS values. A "*" on the time axis indicates a time when a daisy-chaining hand-off is performed, and pose reconstruction is performed using a new set of feature points. The results from this test appear to be very promising. Significant disagreement exists between the two GPS measurements (low cost GPS

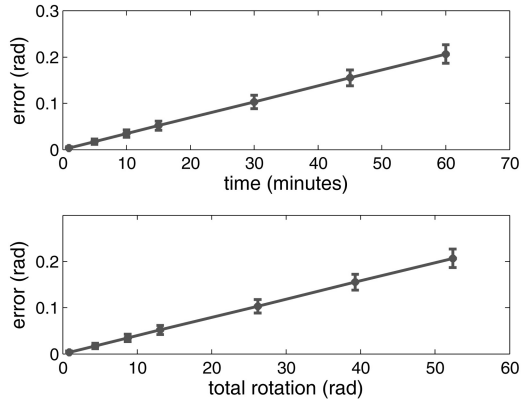


Fig. 11. Error plots of Monte Carlo analysis, showing mean orientation error and standard deviation of orientation error over long time period and long distance. Orientation error refers to norm of difference between known 3D orientation and estimated 3D orientation, where orientation is represented as an angle/axis vector.

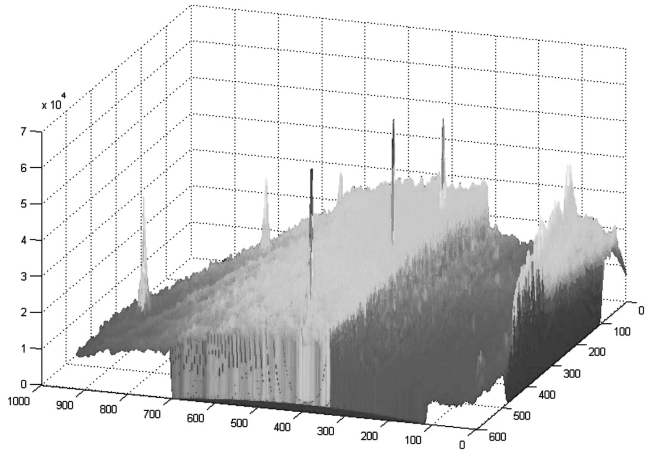

Fig. 12. Example of 3D contour plot generated from matrix designed as nonlinear combination of red and green color space matrices.

receivers are notorious for having a poor altitude measurement [44]), and the vision-based estimation remains proportionate to the two GPS measurements.

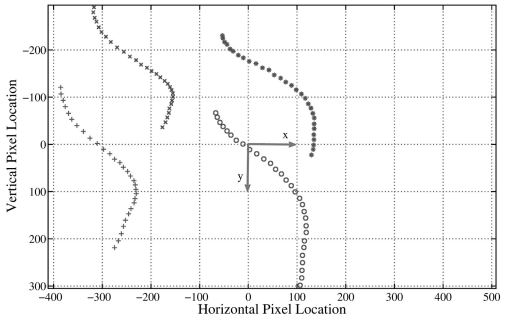


Fig. 13. Image plane trajectories made by landmarks from patch1entering and exiting FOV.


Fig. 14. Single video frame with GPS overlay illustrating landmarks placed along inside edge of runway.

Furthermore, the estimates agree closely with GPS2, for downrange and crossrange translation, and with GPS1, for altitude translation. There is no discernible discontinuity or increased error at the daisy-chain hand-off times. Note that the resolution of the vision-based estimation (5 Hz) is also higher than that of both GPS units (1 Hz). The image processing and pose estimation code can be executed in real time ($> 30$ Hz) on a Pentium M laptop (single core) with 1 Gb RAM. The disagreement between the two GPS

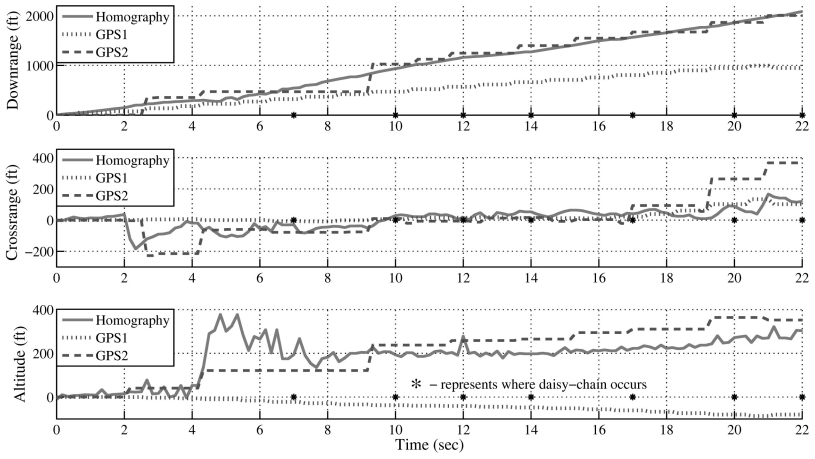


Fig. 15. Results from first experimental flight test. Estimated position is compared with two GPS signals.

units raises concerns about their accuracy. This, in part, inspired the experiment in Section IVB.

### B. Daisy Chained Homography versus Precise Geometric Reconstruction

The accuracy of the recorded GPS data proved to be questionable, as the two GPS records had significant disagreement. A second flight test was designed that was intended to enable the vehicle position to be determined with greater accuracy than what could be achieved with the inexpensive GPS units used in the first round of testing. An additional drawback with the first flight test was that, due to cost constraints, an inertial measurement unit (IMU) was not employed to capture attitude information. In addition to the issues with the GPS output being considered as truth data, it was also of interest to investigate whether using a wider angle FOV lens on a more forward pointing camera could improve the results. The fundamental assumption of this additional testing was that if the locations of the ground targets were known precisely, then vehicle position and attitude could be determined through geometric reconstruction. With the trajectories of the four landmarks of each patch recorded, vehicle localization and orientation could then be performed offline for analysis purposes. While this geometric reconstruction was also based on image data, the precise knowledge of the fiducial position in 3D world coordinates allowed for a very accurate estimate. If the geolocation of the points were known (e.g. through previous physical mapping or template matching with an accurate satellite images), this would represent an accurate means to geolocate the UAV. In essence this experiment compared the best vision-based estimate, given all available knowledge of the precise Euclidean position of the fiducial markers with the daisy-chained homography estimate, which required no knowledge other than initial position and attitude. Details on the geometric reconstruction of attitude and position information for the ground truth in this experiment is described in [43]. A similar procedure using three points that remained in the FOV is described in [15].

By using a wider FOV lens, the resulting video image exhibited severe lens distortion effects that required correcting in order for the homography to operate correctly. The distortion can be seen in Fig. 16, where the four landmarks in the lower portion of the frame appear to be located on a spherical surface as opposed to a plane. This same image, with the radial distortion removed, is given in Fig. 17. The interlacing effect, prevalent in Fig. 14, has also been removed using freeware video editing software VirtualDub.[1]

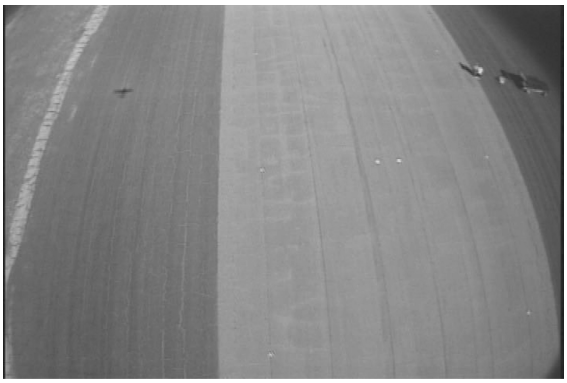

Fig. 16. Single video frame from second flight test experiment illustrating effect of forward-looking, wider FOV camera.

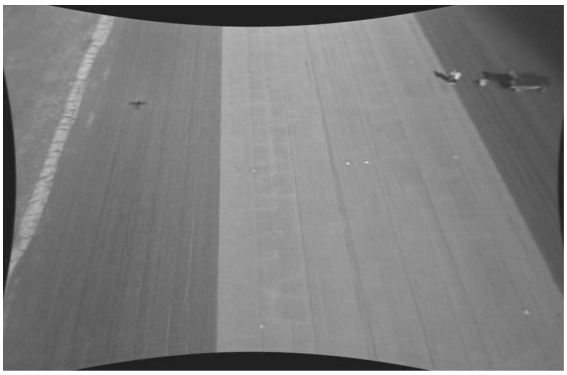

Fig. 17. Single video frame from second flight test experiment with lens distortion removed.

Results of this flight experiment are given in Figs. 18–21. The initial value of $d(0)$ is estimated from the detailed knowledge of the first four points, but subsequent estimates of $d(t)$ are generated using the methods in Section IIC. A "*" in the plot indicates a time when a daisy-chaining hand-off was performed and when pose reconstruction is performed using a new set of feature points. The position estimate of the daisy-chained homography estimate agrees closely with the position estimate of the geometric reconstruction. The attitude estimate of both methods shows much more variation than the translation estimates, and there is less agreement between the two estimates. Attitude is expressed in roll/pitch/yaw angles. Despite the disagreement, the relative shape of both signals is similar, and the homography tracks the geometric reconstruction well, with the exception of the first 3 s of the roll angle and a possible drift in the heading angle towards the end of the flight. Since the developed estimation is based on chaining successive approximations, some error propagation is inevitable. (See Figs. 11 and 10.) The focus of this

[1]VirtualDub—Copyright (c) by Avery Lee http://www.virtualdub.org/index.html.

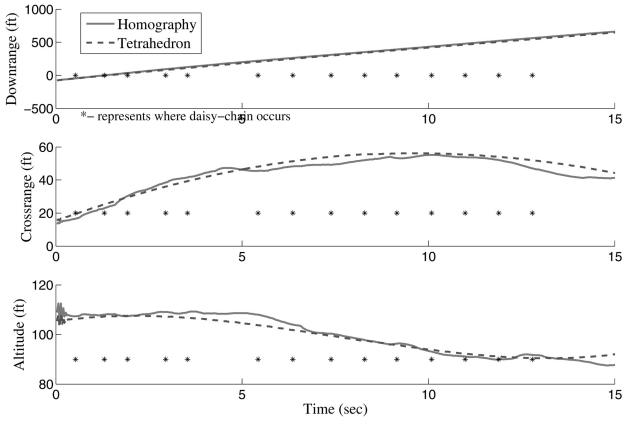


Fig. 18. Results from second experimental flight test. Estimated position is compared to construction using complete geometric knowledge of ground targets.

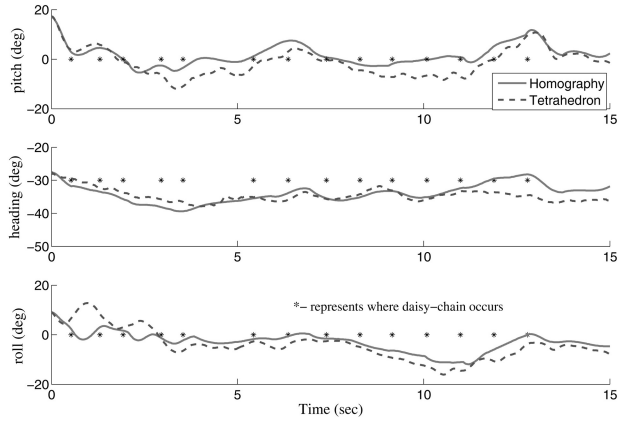


Fig. 19. Results from second experimental flight test. Estimated attitude and bearing is compared to construction using complete geometric knowledge of ground targets.

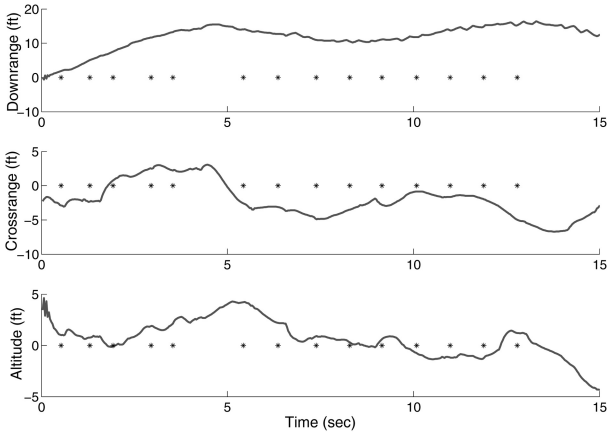


Fig. 20. Results from second experimental flight test shows disagreement between daisy-chained position estimate and position estimation using complete geometric knowledge of ground targets.

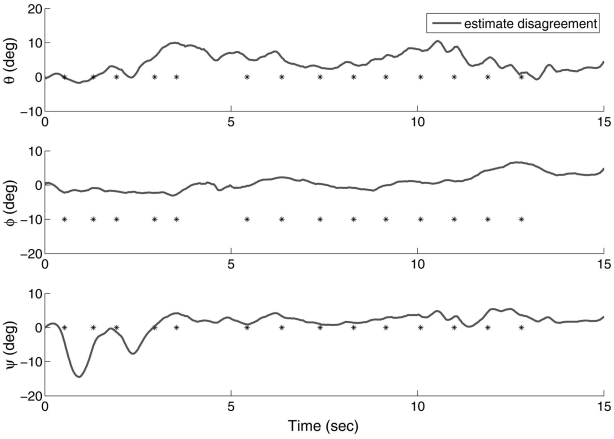


Fig. 21. Results from second experimental flight test shows disagreement between daisy-chained orientation estimate and orientation estimation using complete geometric knowledge of ground targets.

study was to investigate the feasibility of using the camera feedback for pose estimation. Future efforts will focus on quantifying the error propagation, and these efforts will investigate methods that may be used to mitigate the error propagation, such as collaborative sensing between multiple UAVs, sensor integration with other sensors, error correction by sporadic GPS feedback, etc.

## V. CONCLUSIONS

The efforts in this paper integrated new vision-based, pose estimation methods for guidance, navigation, and control of an aerial vehicle. This method is based on Epipolar geometry, with a novel "daisy-chaining" approach that allows image features to enter and leave the FOV while maintaining pose estimation. Furthermore, no known target is required. A nonlinear aircraft flight model for an Osprey UAV was developed to investigate the performance of the developed method in simulations. Simulations also include testing the pose estimation method in closed-loop control of the aircraft. Based on the successful simulation results, two flight tests were also performed using a radio controlled Osprey UAV. The first flight test indicated that the estimation method provided approximately equal pose estimates as the two GPS units. However, conclusive results were difficult to obtain due to discrepancies in the GPS units and due to lack of attitude information. The second flight test used precise information about the Euclidean position of fiducial markers as a ground truth for the estimation method. The daisy-chaining estimation method was successfully demonstrated to reconstruct both position and attitude measurements in comparison to the reconstructed ground truth using detailed geometric knowledge.

In addition to providing pose estimates for control, the daisy-chaining method in this paper is suitable for use as a localization algorithm in local VSLAM applications. It is feasible to keep track of the location of previously tracked feature points, resulting in a map useful for global VSLAM applications. Future efforts will work to further develop such VSLAM

applications. Future research will also investigate methods to improve inevitable error propagation through the inclusion of additional sensors, such as IMUs and intermittent GPS. Additional future work will focus on determining what, if any, bounds can be placed on estimate errors over time and under various common error sources. Future research will investigate methods to improve inevitable error propagation through the inclusion of additional sensors, such as IMUs and intermittent GPS.

REFERENCES


[1] J. A. V. N. T. S. Center
Vulnerability assessment of the transport infrastructure relying on the global positioning system.
Office of the Assistant Secretary for Transportation Policy, U.S. Department of Transportation, Report, Aug. 2001.

[2] Kaiser, K., Gans, N., and Dixon, W.
Position and orientation an aerial vehicle through chained, vision-based pose reconstruction.
In *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, 2005.

[3] Kaiser, K., Gans, N., and Dixon, W.
Localization and control an aerial vehicle through chained, vision-based pose reconstruction.
In *Proceedings of the American Control Conference*, 2007, 5934–5939.

[4] Hu, G., Mehta, S., Gans, N., and Dixon, W. E.
Daisy chaining based visual servo control. Part I: Adaptive quaternion-based tracking control.
In *Proceedings of the IEEE Multi-conference on Systems and Control*, 2007, 1474–1479.

[5] Hu, G., Gans, N., Mehta, S., and Dixon, W. E.
Daisy chaining based visual servo control. Part II: Extensions, applications and open problems.
In *Proceedings of the IEEE Multi-conference on Systems and Control*, 2007, 729–734.

[6] Chatterji, G., Menon, P., and Sridhar, B.
GPS/machine vision navigation system for aircraft.
*IEEE Transactions on Aerospace and Electronic Systems*, **33**, 3 (July 1997), 1012–1025.

[7] Roberts, J. M., Corke, P. I., and Buskey, G.
Low-cost flight control system for a small autonomous helicopter.
In *Proceedings of the Australasian Conference on Robotics and Automation*, 2002.

[8] Zhang, H. and Ostrowski, J.
Visual servoing with dynamics: Control of an unmanned blimp.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 1, May 1999, 618–623.

[9] Sharp, C. S., Shakernia, O., and Sastry, S. S.
A vision system for landing an unmanned aerial vehicle.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2001, 1720–1727.

[10] Jones, C. G., Heyder-Bruckner, J. F., Richardson, T. S., and Jones, D. C.
Vision-based control for unmanned rotorcraft.
In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2006.

[11] Saripalli, S., Montogomery, J. F., and Sukhatme, G. S.
Visually guided landing of an unmanned aerial vehicle.
*IEEE Transactions on Robotics and Automation*, **19**, 3 (2003), 371–380.

[12] Chatterji, G. B., Menon, P. K., and Sridhar, B.
Vision-based position and attitude determination for aircraft night landing.
*Journal of Guidance, Control, and Dynamics*, **21**, 1 (1998), 84–92.

[13] Liu, T. and Fleming, G.
Videogrammetric determination of aircraft position and attitude for vision-based autonomous landing.
Presented at the AIAA Aerospace Sciences Meeting and Exhibit, 2006.

[14] Shakernia, O., Ma, Y., Koo, T. J., and Sastry, S.
Landing an unmanned air vehicle: Vision based motion estimation and non-linear control.
*Asian Journal of Control*, **1** (1999), 128–145.

[15] Yakimenko, O., Kaminer, I., Lentz, W., and Ghyzel, P.
Unmanned aircraft navigation for shipboard landing using infrared vision.
*IEEE Transactions on Aerospace and Electronic Systems*, **38**, 4 (Oct. 2002) 1181–1200.

[16] Koch, A., Wittich, H., and Thielecke, F.
A vision-based navigation algorithm for a VTOL-UAV.
Presented at the AIAA Guidance, Navigation, and Control Conference, 2006.

[17] Jianchao, Y.
A new scheme of vision based navigation for flying vehicles—concept study and experiment evaluation.
In *Proceedings of the IEEE International Conference on Control, Automation, Robotics And Vision*, 2002, 643–648.

[18] Barber, D. B., Griffiths, S. R., McLain, T. W., and Beard, R. W.
Autonomous landing of miniature aerial vehicles.
*Journal of Aerospace Computing, Information, and Communication*, **4** (2007), 770.

[19] Bryson, M. and Sukkarieh, S.
Observability analysis and active control for airborne SLAM.
*IEEE Transactions on Aerospace and Electronic Systems*, **44**, 1 (Jan. 2008), 261–280.

[20] Kim, J. and Sukkarieh, S.
Autonomous airborne navigation in unknown terrain environments.
*IEEE Transactions on Aerospace and Electronic Systems*, **40**, 3 (July 2004), 1031–1045.

[21] Trisiripisal, P., Parks, M. R., Abbot, A. L., Liu, T., and Fleming, G. A.
Stereo analysis for vision-based guidance and control of aircraft landing.
Presented at the AIAA Aerospace Sciences Meeting and Exhibit, 2006.

[22] Wu, A. D., Johnson, E. N., and Proctor, A. A.
Vision-aided inertial navigation for flight control.
*Journal of Aerospace Computing, Information, and Communication*, **2** (2005), 348–360.

[23] Prazenica, R. J., Watkins, A., Kurdila, A. J., Ke, Q. F., and Kanade, T.
Vision-based Kalman filtering for aircraft state estimation and structure from motion.
Presented at the AIAA Guidance, Navigation, and Control Conference, 2005.

[24] Caballero, F., Merino, L., Ferruz, J., and Ollero, A.
Improving vision-based planar motion estimation for unmanned aerial vehicles through online mosaicing.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2006, 2860–2865.

[25] Se, S., Lowe, D., and Little, J.
Global localization using distinctive visual features.
In *Proceedings of the IEEE/RAS International Symposium on Robotics and Automation*, 2002, 226–231.

[26] Eustice, R., Singh, H., Leonard, J., Walter, M., and Ballard, R.
Visually navigating the RMS titanic with SLAM information filters.
In *Proceedings of Robotics: Science and Systems*, June 2005.

[27] Jensfelt, P., Kragic, D., Folkesson, J., and Bjorkman, M.
A framework for vision based bearing only 3D SLAM.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2006, 1944–1950.

[28] Lowe, D.
Object recognition from local scale-invariant features.
In *Proceedings of the IEEE International Conference on Computer Vision*, 1999, 1150–1157.

[29] Davison, A. J., Reid, I. D., Molton, N. D., and Stasse, O.
MonoSLAM: Real-time single camera SLAM.
*IEEE Transactions on Pattern Analysis and Machine Intelligence*, **29**, 6 (2007), 1052–1067.

[30] Kim, J-H. and Sukkarieh, S.
Airborne simultaneous localisation and map building.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2003.

[31] Jung, I-K. and Lacroix, S.
High resolution terrain mapping using low attitude aerial stereo imagery.
In *Proceedings of the IEEE International Conference on Computer Vision*, 2003, 946–951.

[32] Goncalves, L., di Bernardo, E., Benson, D., Svedman, M., Ostrowski, J., Karlsson, N., and Pirjanian, P.
A visual front-end for simultaneous localization and mapping.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2005, 44–49.

[33] Baillard, C. and Zisserman, A.
Automatic reconstruction planar models from multiple views.
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1999, 559–565.

[34] Okada, K., Kagami, S., Inaba, M., and Inoue, H.
Plane segment finder: Algorithm, implementation and applications.
In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2001, 2120–2125.

[35] Longuet-Higgins, H.
A computer algorithm for reconstructing a scene from two projections.
*Nature*, **293** (Sept. 1981), 133–135.

[36] Hartley, R.
*Computer Vision—Proceedings of the 1992 European Conference on Computer Vision, Lecture Notes in Computer Sciences*.

[37] Boufama, B. and Mohr, R.
Epipole and fundamental matrix estimation using virtual parallax.
In *Proceedings of the IEEE International Conference on Computer Vision*, 1995, 1030–1036.

[38] Malis, E. and Chaumette, F.
2 1/2 D visual servoing with respect to unknown objects through a new estimation scheme of camera displacement.
*International Journal of Computer Vision*, **37**, 1 (2000), 79–97.

[39] Nister, D.
An efficient solution to the five-point relative pose problem.
*IEEE Transactions on Pattern Analysis and Machine Intelligence*, **26** (2004), 79–97.

[40] Y. Ma, S. Soatto, J. Koseck, and S. Sastry
*An Invitation to 3-D Vision*.
New York: Springer, 2004.

[41] Faugeras, O. D. and Lustman, F.
Motion and structure from motion in a piecewise planarenvironment.
*International Journal of Pattern Recognition and Artificial Intelligence*, **2**, 3 (1988), 485–508.

[42] Zhang, Z. and Hanson, A.
3D reconstruction based on homography mapping.
Presented at the ARPA Image Understanding Workshop, Palm Springs, CA, 1996.

[43] Kaiser, M. K.
Vision-based estimation, localization, and control of an unmanned aerial vehicle.
Ph.D. dissertation, Department of Mechanical and Aerospace Engineering, University of Florida, Gainesville, FL, May 2008.

[44] Kim, J. and Sukkarieh, S.
A baro-altimeter augmented INS/GPS navigation system for an uninhabited aerial vehicle.
In *Proceedings of the International Symposium on Satellite Navigation Technology Including Mobile Positioning and Location Services*, Melbourne, Australia, 2003.

**Michael Kent Kaiser** received his B.S. degree in aerospace engineering from Auburn University, Auburn, AL, in 1990, an M.S. degree in mechanical engineering from the University of Houston, Houston, TX, in 1995, and a Ph.D. in aerospace engineering from the University of Florida, Gainesville, in 2008.

He has nearly 20 years of work experience within the aerospace industry, which includes his current employment with Mosaic, ATM, as well as previous employment with the Air Force Research Labs (AFRL) and Lockheed Martin Corporation in Houston, TX, Palmdale, CA, and Marietta, GA. Some of his research interests include system identification, aerodynamic performance based vehicle design, flight dynamics and control, and vision-based navigation, guidance, and control.

**Nicholas Gans** (M'05) received his B.S. degree in electrical engineering from Case Western Reserve University, Cleveland, OH, in 1999. He earned his M.S. in electrical and computer engineering and his Ph.D. in systems and entrepreneurial engineering from the University of Illinois at Urbana-Champaign in December 2005.

After completing his doctoral studies he worked as a postdoctoral researcher with the Mechanical and Aerospace Engineering Department at the University of Florida. He is currently a postdoctoral associate with the National Research Council, where he conducts research for the Air Force Research Laboratory at Eglin AFB in Florida. His research interests include nonlinear and adaptive control, with focus on vision-based control and estimation, robotics, and autonomous vehicles.

Dr. Gans is a member of the IEEE Robotics and Automation Society (RAS), and the IEEE Control System Society (CSS), and he has served on the program committees for several conferences.

**Warren Dixon** (M'00—SM'05) received his Ph.D. degree in 2000 from the Department of Electrical and Computer Engineering from Clemson University, Clemson, SC.

After completing his doctoral studies, he was selected as a Eugene P. Wigner Fellow at Oak Ridge National Laboratory (ORNL), where he worked in the Robotics and Energetic Systems Group. In 2004 he joined the faculty of the University of Florida in the Mechanical and Aerospace Engineering Department. Dr. Dixon's main research interest has been the development and application of Lyapunov-based control techniques for uncertain nonlinear systems.

Dr. Dixon has published 2 books, an edited collection, 4 chapters, and over 180 refereed journal and conference papers. He was awarded the 2001 ORNL Early Career Award for Engineering Achievement for his contributions to Lyapunov-based control methods. He was awarded the 2004 DOE Outstanding Mentor Award for his student advising at ORNL. He was awarded an NSF CAREER award, in 2006, for new development and application of Lyapunov-based control methods. He was also awarded the 2006 IEEE Robotics and Automation Society (RAS) Early Academic Career Award. He serves on the IEEE CSS Technical Committee on Intelligent Control, is a primary member of the ASME DSC Division Mechatronics Technical Committee, is a member of numerous conference program committees, and serves on the conference editorial board for the IEEE CSS and RAS and the ASME DSC. He served as an appointed member to the IEEE CSS Board of Governors for 2008. He is currently an associate editor for *IEEE Transactions on Systems Man and Cybernetics: Part B Cybernetics*, *Automatica*, *International Journal of Robust and Nonlinear Control*, and *Journal of Robotics*.