

Vision-Based Road-Following Using a Small Autonomous Aircraft¹²

Eric Frew, Tim McGee, ZuWhan Kim, Xiao Xiao, Stephen Jackson, Michael Morimoto, Sivakumar Rathinam, Jose Padiar, and Raja Sengupta

AINS Center for Collaborative Control of Unmanned Vehicles
University of California – Berkeley
2105 Bancroft Way, Berkeley, CA

{efrew@path, tmcgee@newton, zuwhan@cs, xiaoxiao@uclink, action@uclink, mmorimot@uclink, rsiva@uclink, jopadiar@uclink, raja@path}.berkeley.edu

Abstract—This paper describes the vision-based control of a small autonomous aircraft following a road. The computer vision system detects natural features of the scene and tracks the roadway in order to determine relative yaw and lateral displacement between the aircraft and the road. Using only the vision measurements and onboard inertial sensors, a control strategy stabilizes the aircraft and follows the road. The road detection and aircraft control strategies have been verified by hardware in the loop (HIL) simulations over long stretches (several kilometers) of straight roads and in conditions of up to 5 m/s of prevailing wind. Hardware experiments have also been conducted using a modified radio-controlled aircraft. Successful road following was demonstrated over an airfield runway under variable lighting and wind conditions. The development of vision-based control strategies for unmanned aerial vehicles (UAVs), such as the ones presented here, enables complex autonomous missions in environments where typical navigation sensor like GPS are unavailable.



Figure 1 – The experimental autonomous aircraft is a modified Sig Rascal radio-controlled airplane with 110 in. wingspan.

TABLE OF CONTENTS

I. INTRODUCTION	1
II. PROBLEM STATEMENT.....	2
III. COMPUTER VISION.....	4
IV. LATERAL CONTROL	5
V. EXPERIMENTAL SYSTEM	7
VI. RESULTS	7
VII. CONCLUSIONS	8
VIII. ACKNOWLEDGEMENTS	9
IX. REFERENCES	9
X. BIOGRAPHY	9

I. INTRODUCTION

The development of vision-based control strategies for unmanned aerial vehicles (UAVs) is critical to the success of future autonomous missions. Applications such as traffic monitoring, border patrol and homeland defense, search and rescue, aerial surveillance, crop inspection, and law enforcement, require high levels of system robustness and

performance. Standard techniques which rely heavily on GPS for navigation and control will have difficulty in urban or hostile environments in which satellite signals can easily be blocked or jammed.

Computer vision is an important sensor for UAVs operating in natural environments. A sequence of video images contains large amounts of information that can be used for vehicle navigation and control, object detection and identification, obstacle avoidance, and many other tasks. Unlike radar- or laser-based systems, computer vision is passive and emits no external signals. As a result, vision systems can be small in size, causing little payload burden to the UAV platform, and can operate undetected in hostile environments.

The goal of the work presented here is to enable vision-based following of a roadway -- using only the natural features of the scene and no additional navigation sensors -- by a small, autonomous aircraft (Figure 1). Previous work

¹ 0-7803-8155-6/04/\$17.00 © 2004 IEEE

² IEEEAC paper #1479, Version 3, Updated January 5, 2004

on vision-based control of UAVs has focused on autonomous helicopters. In particular, vision-based landing has been studied extensively and solutions have used structured landing areas to aid the vision processing [1, 2]. Likewise, artificial markings have been used to perform visual target following [3]. Several techniques have also been proposed for vision-based navigation, using depth maps to locally update occupancy grids [4] and using known databases to perform visual landmark navigation [5]. Another vision-based system is being developed to operate over diverse geographical terrain containing road and traffic networks, with a focus on knowledge representation and high level reasoning [6].

Compared to helicopters, visual control of small autonomous aircraft is complicated by the speeds with which the vehicles fly. Although basic stability is easier to achieve with an aircraft, precise lateral control (needed for road following) is challenging [7, 8]. Stability for a micro-air vehicle has been attempted using visual horizon detection [9]. Stereo-based vision for pose estimation and monocular vision for navigation have been proposed in [10], although no results have been reported. Simultaneous localization and map building (SLAM) for an unmanned aircraft has been performed using artificial markings [11]. A biomimetic approach to vision-based control has been developed using simple optical flow sensors for terrain following and obstacle avoidance [12].

While vision-based road following has never been addressed for a small autonomous plane, the similar problem of lane tracking and following for autonomous ground vehicles traveling at freeway speeds has been studied [13, 14]. There are two main differences between the two applications. First, the dynamics of a small aircraft are significantly different from those of an automobile. In particular, the coupling between aircraft bank angle and turning rate complicate the sensing of the aircraft position relative to the road. Second, the perspective of cameras on the aircraft and automobile vary greatly from another. The automobile effectively sits on the ground and looks along the road, causing the lane or roadway to diverge in the image. By comparison, the aircraft flies above the road looking downward, causing the road to appear as several parallel lines.

Vision-based road following can be split into two important sub-problems: i.) detecting the road within an image and calculating the distance of the aircraft away from it and ii.) controlling the flight of the aircraft in order to follow the road and bring this cross-track error to zero. The solutions to both sub-problems presented here build on previous approaches developed for other applications.

Using only the natural characteristics of the roadway, the computer vision system segments the image by color into background and road pixels, performs a connected-component analysis to identify the road, applies the Hough

transform to calculate an initial estimate of the road direction, and performs a robust fitting to locate the roadway within the image. Several different approaches are developed in order to calculate the lateral distance of the aircraft away from the road. The first assumes the aircraft is at trim and that the pitch and roll of the aircraft is zero. The second uses the onboard inertial sensors to calculate the aircraft orientation and correct for errors in the lateral distance measurement. The final approach uses a stereo vision system to calculate the orientation of the aircraft relative to the ground plane.

Once the distance between the aircraft and roadway is determined, a lateral control strategy drives this distance to zero, thus achieving autonomous road following. Three strategies are presented that can perform the vision-based control task by commanding the aircraft turning rate. The first uses simple proportional-integral-derivative (PID) control on the lateral distance. The second and third each use nonlinear controllers to aim the aircraft toward a point on the road. The second strategy explicitly calculates both the desired and estimated headings in order to generate the control signal while the third strategy attempts to enforce a geometric relationship between the aircraft velocity and lateral error.

This paper presents the development of the road detection algorithms and the new lateral control strategies. A hardware-in-the-loop (HIL) simulation environment is used to verify the performance of these algorithms. Additionally, flight tests are conducted on a small autonomous aircraft (Figure 1) outfitted with a commercial avionics package modified for these experiments. The new road-following system stably tracks a simulated road for over 2 kilometers under constant background winds of up to 5 m/s. Using the experimental hardware, visual tracking of an airfield runway is achieved under various lighting conditions and light, gusting winds with control errors of approximately 10-20 meters from an altitude of 100 meters.

II. PROBLEM STATEMENT

Vision Geometry

The goal of the computer vision system is to detect the road in the camera images and calculate the relative distance and heading between the aircraft and the road. In the basic pinhole camera model, the location of a point in the camera coordinate system (x_c, y_c, z_c) is related to a point in the image (u, v) by the projection:

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{C} \cdot \begin{bmatrix} x_c \\ z_c \\ y_c \\ z_c \\ 1 \end{bmatrix}, \quad (1)$$

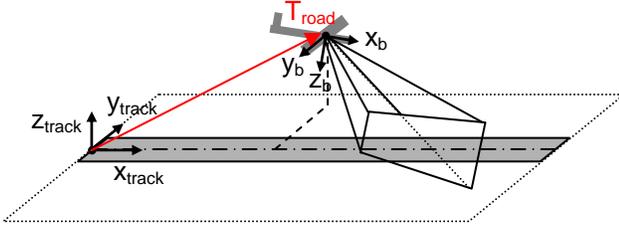


Figure 2 - Aircraft in road coordinate system

$$\mathbf{C} = \begin{bmatrix} f_u & 0 & u_0 \\ 0 & f_v & v_0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (2)$$

where \mathbf{C} is the matrix of intrinsic parameters [15], (f_u, f_v) are the focal lengths of the camera in pixel dimensions, and (u_0, v_0) is the pixel coordinate of the image center. Letting $w = z_c$, $\mathbf{u} = (u \cdot w, v \cdot w, w)^T$, and $\mathbf{X}_c = (x_c, y_c, z_c)^T$,

$$\mathbf{u} = \mathbf{C} \cdot \mathbf{X}_c. \quad (3)$$

Typically the camera coordinate system and the aircraft coordinate system will differ by a rotation \mathbf{R}_e and translation \mathbf{T}_e , known as the camera's extrinsic parameters and determined by camera calibration, such that

$$\mathbf{X}_c = \mathbf{R}_e \cdot \mathbf{X}_b + \mathbf{T}_e, \quad (4)$$

where \mathbf{X}_b is the location of a point in the aircraft coordinate system, so

$$\mathbf{u} = \mathbf{C} \cdot (\mathbf{R}_e \cdot \mathbf{X}_b + \mathbf{T}_e). \quad (5)$$

The aircraft moves in a world reference frame fixed to the road (Figure 2). The location of the aircraft in this frame is $\mathbf{T}_{road} = (x_{track}, y_{track}, z_{track})^T$ and the orientation between the road and aircraft frames is represented by the rotation \mathbf{R}_{road} , which is a function of the roll, pitch, and relative yaw of the aircraft. A point in the road frame (\mathbf{X}_{road}) is related to a point in the aircraft frame (\mathbf{X}_b) such that

$$\mathbf{X}_{road} = \mathbf{R}_{road} \cdot \mathbf{X}_b + \mathbf{T}_{road}. \quad (6)$$

Combining (5) and (6) yields

$$\mathbf{u} = \mathbf{C} \cdot (\mathbf{R}_e \cdot (\mathbf{R}_{road}^{-1} \cdot (\mathbf{X}_{road} - \mathbf{T}_{road})) + \mathbf{T}_e) \quad (7)$$

Road detection is the problem of finding the set of image points \mathbf{u}_i that correspond to points on the center line of the

road. Given this image set, the intrinsic and extrinsic camera parameters, and the fact that $\mathbf{X}_{road_i} = (\alpha_i, 0, 0)^T$, calculating the lateral distance and orientation between the roadway and the aircraft is equivalent to determining \mathbf{T}_{road} and \mathbf{R}_{road} .

Control Problem

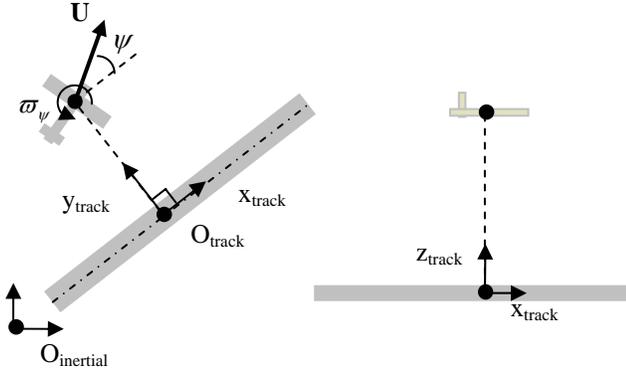
Vision-based road following is essentially an outer (navigation) loop control problem that can be decoupled into two sub-problems: lateral control relative to the road and altitude hold above it (Figure 3). One major assumption in the work presented here is that GPS is not available, so control and navigation must be performed using computer vision and other onboard sensors.

Aircraft stability can be achieved with four simple PID controllers that feed back airspeed, pitch rate, yaw rate, and roll rate to the aircraft throttle, elevator, rudder, and aileron servomotors [16]. Onboard sensors used to close these control loops include static and pitot tubes to measure air speed and rate gyros to measure the angular rates of the aircraft.

Once the inner loops are stabilized, outer loop control can be used to navigate the aircraft. Altitude control of small unmanned aircraft is typically achieved by feedback from barometric pressure sensors to aircraft thrust (throttle) [8]. The same approach is used here. Position control is achieved by commanding the aircraft turning rate, which is directly related to the aircraft orientation. Turning rate commands are therefore achieved by mapping the turning rates into desired roll and pitch angles which are in turn used by the inner PID controllers.

A full non-linear model of the autonomous aircraft has been developed and used for the HIL simulations described below. Assuming tight inner loop control and good turning rate command tracking, a planar kinematic model is used to develop the lateral control strategy. Given a straight roadway with any orientation in space, one point on the roadway is chosen as the origin of the control frame. The x-axis of this frame is defined along the road and the y-axis is perpendicular to it (Figure 3). The vehicle kinematic model is:

$$\begin{aligned} \dot{x}_{track} &= U_{TAS} \cos(\psi) + V_{w_x} \\ \dot{y}_{track} &= U_{TAS} \sin(\psi) + V_{w_y}, \\ \dot{\psi} &= \varpi_{cmd} \end{aligned} \quad (8)$$



Lateral control Altitude hold
Figure 3 - Road following control problem

where U_{TAS} is the aircraft's true airspeed, ψ is the heading of the aircraft relative to the road, ψ_{cmd} is the commanded turning rate, and (V_{w_x}, V_{w_y}) are wind disturbances.

Lateral control of the aircraft is achieved by commanding the vehicle turning rate ψ_{cmd} such that the cross-track distance y_{track} is brought to zero. An inner control loop based on the inertial sensors converts the turning rate commands into the corresponding aileron and rudder commands [16].

III. COMPUTER VISION

Real-time road detection and localization

Real-time road detection and localization is an essential part of the vision-based control. The detection and localization algorithm should be fast and robust at the same time. To achieve this, we applied various techniques such as Bayesian pixel classification, connected-component analysis, Hough transformation, and robust line fitting.

After first rectifying the image (Figure 4b) to account for lens distortion, we applied a Bayesian pixel classifier to find a rough location of the road. We used RGB values of the pixel to classify road pixels from non-road pixels. Our goal is to find the probability $P(\text{road} | r, g, b)$ that a pixel belongs to the road given its RGB values. According to Baye's rule:

$$P(\text{road} | r, g, b) = \frac{P(r, g, b | \text{road}) \cdot P(\text{road})}{P(r, g, b)} \quad (9)$$

Since we do not know the prior distribution of RGB values, $P(r, g, b)$, we use likelihood $L(\text{road} | r, g, b)$ instead:

$$\begin{aligned} L(\text{road} | r, g, b) &\equiv \frac{P(\text{road} | r, g, b)}{P(\text{not road} | r, g, b)} \\ &= \frac{P(r, g, b | \text{road}) \cdot P(\text{road})}{P(r, g, b | \text{not road}) \cdot P(\text{not road})} \end{aligned} \quad (10)$$

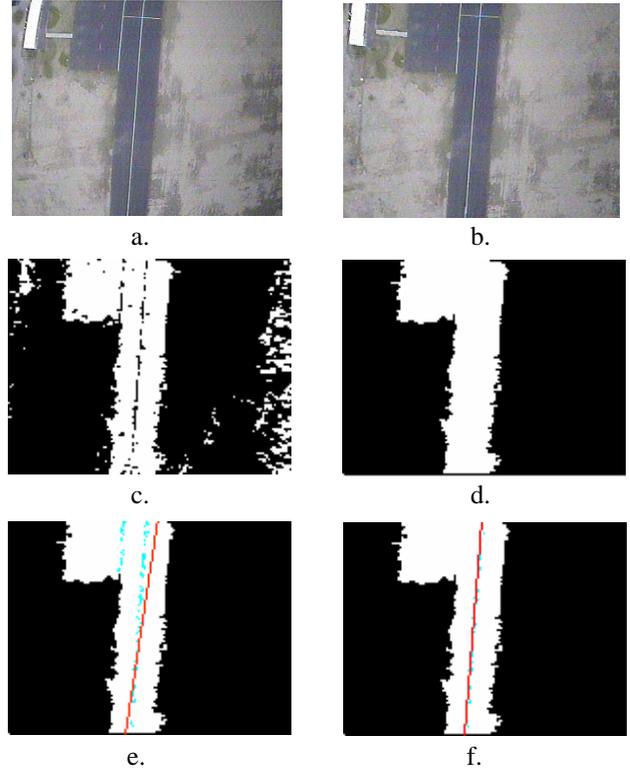


Figure 4 - Road detection algorithm. a.) Original image. b.) Corrected for distortion. c.) Bayesian pixel classification. d.) Connected-component analysis. e.) Hough transform. f.) Least-trimmed square line fitting.

We used multivariate Gaussian distributions to represent $P(r, g, b | \text{road})$ and $P(r, g, b | \text{not road})$. We developed a user interface to gather the RGB values of over 20,000 road- and non-road pixels (labeled). Then, we estimated the parameters of the distributions. Figure 4c shows the classification result.

Once the road pixels are classified, we apply connected-component analysis to remove noise. We applied two-pass connected-component analysis to remove the holes on the center of the road (the lane markings) as well. Figure 4d shows the detected road after the connected-component analysis.

After we detect the road, we search for lane markings on the road. We applied the same Bayesian pixel classification algorithm to find the lane markings. After we detected the lane markings, we apply Hough transformation to test multiple candidate lanes. Figure 4e shows the detected lane markings (shown in cyan) and the fitted lane with the best score (in red). The Hough transformation only gives a rough discretization of the lane. We applied a robust line fitting (least-trimmed square) to finalize the position and the orientation of the center lane marking (Figure 4f).

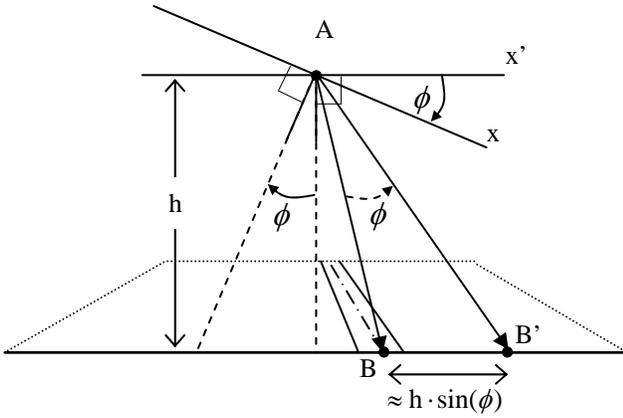


Figure 5 - Lateral distance error due to roll angle error. If the aircraft is banked at a small angle (x -axis) but assumes it is not (x' -axis), it will mistakenly calculate that the object at point B to is at

Estimating lateral distance

Our road detection and localization algorithm gives the position and the orientation of the road in image coordinates. To convert them to the world coordinates, we need to know the position and the orientation of the UAV with respect to the ground. Aircraft altitude is currently measured by a barometric pressure sensor. In order to calculate the height of the aircraft above the road, the altitude of the road must also be known. Future systems will use stereo cameras to determine the aircraft height.

The orientation of the aircraft is obtained in one of several ways. First, we assume the aircraft is in trim flight and the roll and pitch angles are zero. This assumption is typically valid for straight flight; however, significant roll (which can be as small as 5 degrees) is required to turn the aircraft. The resulting errors in roll lead to errors in the cross-track distance measurement (Figure 5) of magnitude

$$e_{y_{track}} \approx h \cdot \sin(\phi) \quad (11)$$

where h is the height of the aircraft above the road and ϕ is the roll angle. For the second method, onboard inertial sensors that measure yaw rate can be used to estimate roll angle. Assuming the aircraft makes a coordinated turn and the pitch angle is small, the roll angle is directly related to the yaw rate, $\dot{\psi}$, (which can be estimated from the rate gyros) by the relationship

$$\tan(\phi) = (U \cdot \dot{\psi}) / g \quad (12)$$

where U is the aircraft speed and g is the acceleration due to gravity. Onboard magnetometers could also be used to improve this estimation process. Finally, a second camera is used to perform stereo triangulation of ground features. By assuming the ground is a flat plane, the orientation of the aircraft relative to this plane can be determined. While a stereo system has been developed for this research, it is not

part of the system presented here.

Once the orientation of the aircraft with respect to the ground is determined, we pick any two points on the road line (on the image) and project them onto the ground plane (with known internal camera parameters) to find the world coordinates of the road. The shortest distance between the road line (in world coordinates) and $p = (0,0,0)^T$ is the lateral distance. We assign the sign of the lateral distance by checking whether p is on the left or right side of the road line. The yaw is calculated in the same manner by projecting vector $(0,0,1)$ on the ground plane and calculating the angle between the projected vector and the road line.

IV. LATERAL CONTROL

Three lateral control strategies are developed to enable autonomous road following. The first uses direct PID feedback between the cross-track (lateral) error and the commanded turning rate. The second and third strategies each use nonlinear controllers to aim the aircraft at a point along the road. The second explicitly calculates both the desired and estimated headings and feeds the heading error to the turning rate command. The third strategy enforces the geometric relationship between the aircraft velocity and the desired intersection point on the road.

PID Controller

A simple PID control law can be used to drive the cross-track errors to zero. Since the controller is run as a discrete system and the only measurement available is the cross-track distance, the turning rate given by

$$\omega_{cmd}(t_i) = k_p \cdot y_{track}(t_i) + k_d \cdot \dot{y}_{est}(t_i) + k_i \cdot y_{est}^{int}(t_i), \quad (13)$$

where

$$\dot{y}_{est}(t_i) = \frac{y_{track}(t_i) - y_{track}(t_{i-1})}{t_i - t_{i-1}} \quad (14)$$

and

$$y_{est}^{int}(t_k) = \sum_{k=0}^i \left(\frac{(y_{track}(t_k) + y_{track}(t_{k-1}))}{2} \cdot (t_k - t_{k-1}) \right). \quad (15)$$

A similar control law was shown to perform poorly for lateral tracking control of waypoint navigation using GPS [7], especially in the face of excessive errors in either cross-track position or velocity. Since the cross-track velocity is estimated from the lateral displacement we expect significant noise. Therefore this control law was never used on the aircraft and the two nonlinear controllers presented below were developed.

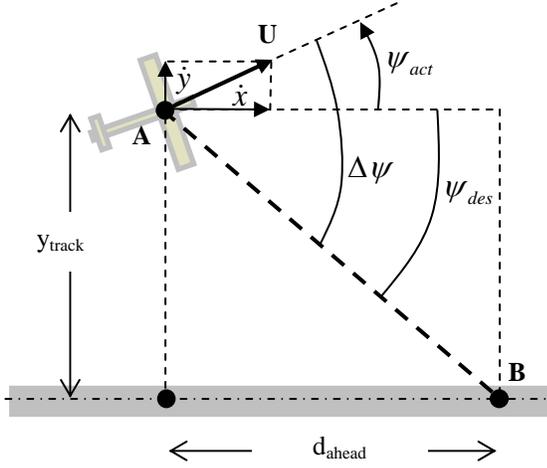


Figure 6 – Arctan nonlinear control strategy

Arctan Controller

Assuming the aircraft speed remains constant, the first nonlinear control strategy aims the aircraft, located at point A in Figure 6, at the point B a specified distance (d_{ahead}) ahead on the roadway. Given this desired intersection point on the road, the heading error is calculated and fed back to the turning rate command through proportional control.

Given point A at $(x_{track}, y_{track})^T$ and point B at $(x_{track} + d_{ahead}, 0)^T$ in the control frame, the desired heading angle is

$$\psi_{des} = \arctan(-y_{track}, d_{ahead}). \quad (16)$$

The actual aircraft heading is calculated from the components of its velocity:

$$\psi_{act} = \arctan(\dot{y}_{track}, \dot{x}_{track}). \quad (17)$$

The error signal is defined as the difference between the desired and actual headings, $e = \Delta\psi = \psi_{act} - \psi_{des}$, and this signal is driven to zero by proportional feedback with a saturation limit of ± 0.2 rad/sec.

In order to perform this control strategy, an estimate of the aircraft velocity in the control frame must be calculated. Given the global velocity of the aircraft and its orientation relative to the road, the velocity in the control frame could be calculated. However, without navigation sensors, the global velocity of the aircraft cannot be determined. Instead, the components of the aircraft velocity in the control frame are estimated from the time history of the cross-track distance $y_{track}(t)$ and the indicated airspeed $U_{IAS} = U_{TAS} + v_{pdyn}$ (true airspeed as measured from an onboard dynamic pressure sensor, where v_{pdyn} is the sensor

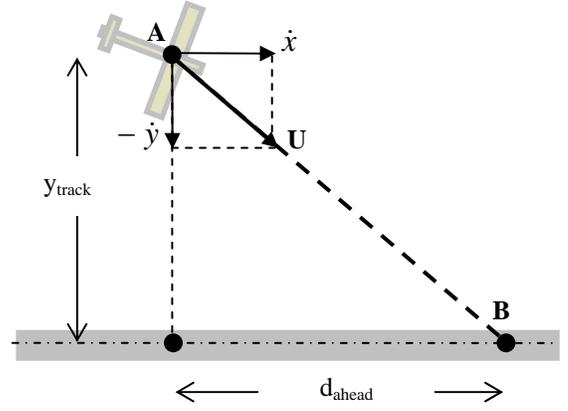


Figure 7 – Velocity ratio control strategy

noise). The cross-track velocity estimate is calculated by (11) while the along-track component is

$$\dot{x}_{est}(t_i) = \sqrt{U_{IAS}^2 - \dot{y}_{est}^2(t_i)}. \quad (18)$$

By directly estimating \dot{y}_{est} from y_{track} the effects of the cross-track wind disturbance v_{w_x} are explicitly incorporated.

However, the wind disturbances effect \dot{x}_{est} by causing a discrepancy between indicated airspeed as measured by the dynamic pressure sensor and the true speed $U = U_{IAS} + v_{pdyn} + \sqrt{V_{w_x}^2 + V_{w_y}^2}$. The vision system is capable of determining the relative orientation (yaw) of the aircraft with respect to the road. However, due to the wind and sideslip, the relative yaw is not necessarily equivalent to the aircraft heading.

Velocity Ratio Controller

A third nonlinear strategy, based on [7], is developed to control the lateral motion of the aircraft relative to the road. Like the second, keeping speed constant, the controller commands the aircraft to aim at the point B a specified distance (d_{ahead}) ahead on the roadway (Figure 7). Given this desired heading, a control signal is derived by establishing the geometric relationship between the desired aircraft position and aircraft velocity

$$\frac{y_{track}}{d_{ahead}} = \frac{-\dot{y}}{\dot{x}}, \quad (19)$$

where the cross-track and along-track velocities are estimated by (11) and (16). In order to enforce this relationship the error signal is defined as

$$e = \dot{x} \cdot y_{track} + \dot{y} \cdot d_{ahead} \quad (20)$$

and is driven to zero by proportional feedback control of the commanded turning rate with a saturation limit of ± 0.2 rad/sec. A value of $k = 0.00001$ was found to achieve



Figure 8 - Cloud Cap Piccolo avionics package

good tracking performance.

V. EXPERIMENTAL SYSTEM

The vision-based control system described here was developed for testing on a Sig Rascal model aircraft (Figure 1). This aircraft has a wingspan of 2.8m, an empty weight of 5.5 kg and a gross weight of 10 kg, allowing 4.5 kg for fuel, avionics, and payload. Low level aircraft control and stabilization is performed by a commercial Cloud Cap Piccolo avionics package (Figure 8), which weighs 212 g [16].

Two downward looking CCD cameras with 320x240 pixel resolution and 62° field-of-view are mounted to the wing struts of the aircraft with a baseline separation of 1.35m. The vision processing is performed on the ground by a PC104 stack.

The cameras send analog video to the ground using two 2.4GHz transmitters. The measurements from the vision system are sent to a Piccolo ground station, via a laptop interface. The ground station sends control inputs to the onboard avionics. This experimental setup is depicted in Figure 9. For future tests, the vision processing will be placed onboard the aircraft and the PC104 will communicate directly with the Piccolo avionics box through a serial port.

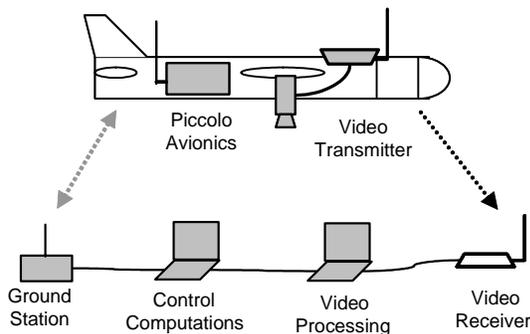


Figure 9 - Flight hardware setup

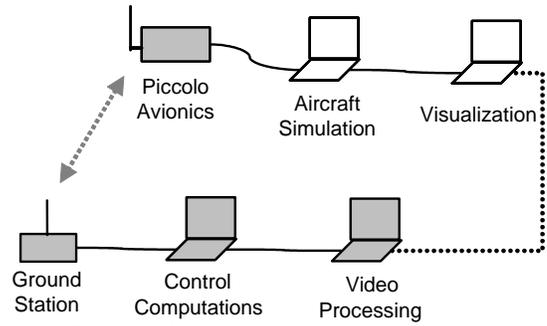


Figure 10 - HIL simulation setup

In addition to the actual aircraft test bed, a hardware-in-the-loop (HIL) simulation, illustrated in Figure 10, was also developed in order to perform initial tests of the vision system and the control algorithms. The flight avionics, ground station, and video processing hardware are used in the HIL simulation. The aircraft dynamics are replaced by a high-dimensional nonlinear simulation provided with the Piccolo avionics package. The camera outputs are simulated using the Vega software package. This software package provides real time visualization of three dimensional models of the testing environment. Figure 11 depicts two Vega generated views. The first image is the simulated camera view, after being processed by the road detection algorithm. The second image depicts a view from a virtual observer behind the aircraft.

VI. RESULTS

The road detection algorithm was first tested independently of aircraft control. The road detection algorithm was able to correctly identify the road in 90% of the images of real roads it was given.

The road following demo was tested using the HIL simulation to follow a simulated 2.5 km stretch of straight road. During the initial tests, the roll and pitch of the aircraft were assumed to be zero. The vision based control allowed the aircraft to successfully keep the road in view during the HIL testing.

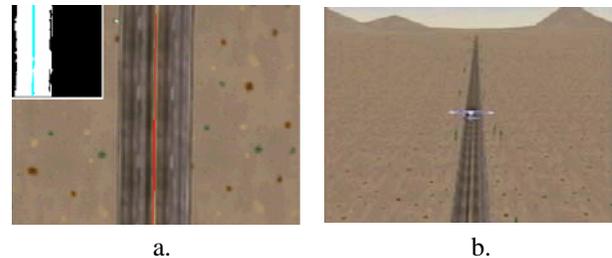
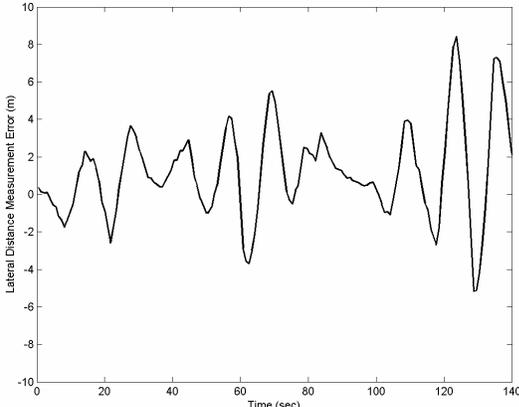
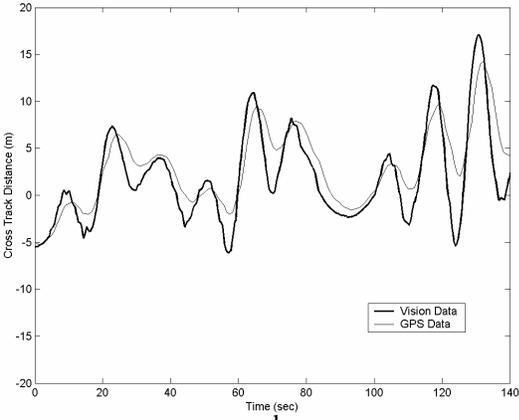


Figure 11 - Vega generated visualization. a.) Camera view. b.) View from behind aircraft.



a.



b.

Figure 12 – Vision-based lateral distant measurements. a.) Vision measurement error vs. time. b.) Vision and GPS lateral distance measurements vs. time

The performance of the vision-based cross-track distance measurements is evaluated by comparing the vision measurements to the simulated GPS flight measurements. Over the course of the 2.5 km experiment, the vision measurements had a mean error of 1.19 meters and a standard deviation of 2.51 meters (Figure 12a). The disparity between the vision and GPS measurements can be attributed largely to the assumption of zero roll, which results in larger vision-based measurements as the aircraft banks to turn toward the road. By overlaying the GPS and vision measurements on the same plot (Figure 12b) and comparing to the vision error plot, this correlation can be seen. Future tests will take into account roll measurements either from stereo vision or onboard inertial sensors.

Figure 13 shows the tracking performance of the aircraft along the straight road during the HIL simulation using the velocity-ratio control law. The aircraft successfully tracks the road for 2.5 km while oscillating with a magnitude of approximately 5 meters. The background wind for this run had a constant velocity of (2.0, 2.0) meters per second, accounting for the slow drift of the aircraft oscillation away from the road.

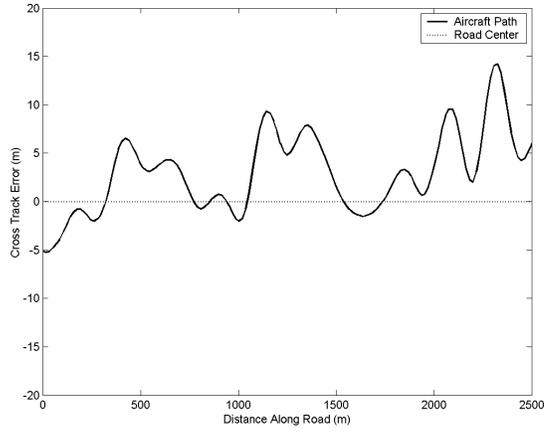


Figure 13 – Tracking performance of vision-based lateral control of HIL simulation.

Single camera road following, with the zero pitch and roll assumption and the velocity-ratio controller, was also tested on the Sig Rascal aircraft at the TIMPA airfield in Tucson, Arizona. Figure 14 shows the results of the aircraft tracking a 230 m long, 15 m wide model aircraft runway. The actual aircraft position was measured using GPS.



Figure 14 - Results of road tracking on actual aircraft.

Although the road following algorithm showed some encouraging results when tested with a single camera, the performance was limited by the assumption of zero roll. In order for this assumption to be valid, only small gains could be used in the control algorithm to prevent excessive banking of the aircraft. Integration of a stereo algorithm or inertial attitude sensors with the road following should allow more aggressive gains and improved tracking performance.

VII. CONCLUSIONS

A vision-based system for tracking and following a road using a small autonomous aircraft is presented. Hardware-in-the-loop (HIL) demonstrations and actual flight tests verify the performance of the system and offer encouraging results for more aggressive control in the future. The performance of the control strategy is directly related to the

accuracy of the lateral distance measurements, which are in turn directly related to errors in the aircraft attitude. Future work will address the optimal fusion of all sensors onboard the aircraft.

VIII. ACKNOWLEDGEMENTS

The authors would like to acknowledge Advanced Ceramics Research Inc. for their help maintaining our autonomous aircraft and providing access to the TIMPA airfield in Tucson, Arizona. This research was supported by the Office of Naval Research AINS program led by Dr. Allen Moshfegh, under contract number #N00014-03-C-0187, SPO #016671-004.

IX. REFERENCES

- [1] S. Saripalli, J.F. Montgomery and G.S. Sukhatme, "Vision-based autonomous landing of an unmanned aerial vehicle," in Proceedings of the 2002 International Conference on Robotics and Automation, 2002.
- [2] O. Shakernia, R. Vidal, C.S. Sharp, Y. Ma and S. Sastry, "Multiple View Motion Estimation and Control for Landing an Unmanned Aerial Vehicle," in Proceedings of the 2002 International Conference on Robotics and Automation, 2002.
- [3] S.M. Rock, E.W. Frew, H.L. Jones, E. LeMaster and B. Woodley, "Combined CDGPS and Vision-Based Control of a Small Autonomous Helicopter," in Proceedings 1998 American Control Conference, pp. 694-698, 1998.
- [4] B. Sinopoli, M. Micheli, G. Donato and T.J. Koo, "Vision Based Navigation for an Unmanned Air Vehicle," in Proceedings of the IEEE International Conference on Robotics and Automation (ICRA2001), pp. 1757-1765, 2001.
- [5] S. Furst and E.D. Dickmanns, "A vision based navigation system for autonomous aircraft," in Intelligent Autonomous Systems. IAS-5, pp. 765-774, 1998.
- [6] P. Doherty, G. Granlund, K. Kuchcinski, E. Sandewall, K. Nordberg, E. Skarman and J. Wiklund, "The WITAS unmanned aerial vehicle project," in Proceedings of 14th European Conference on Artificial Intelligence, pp. 747-755, 2000.
- [7] M. Niculescu, "Lateral track control law for Aerosonde UAV," in Proceedings of the 39th AIAA Aerospace Sciences Meeting and Exhibit, 2001.
- [8] J.S. Jang and C.J. Tomlin, "Autopilot Design for the Stanford DragonFly UAV: Validation through Hardware-in-the-Loop Simulation," in Proceedings of the AIAA Guidance, Navigation, and Control Conference, 2001.
- [9] S.M. Ettinger, M.C. Nechyba, P.G. Ifju and M. Waszak, "Vision-guided flight stability and control for micro air vehicles," in 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 2134-2140, 2002.
- [10] F. Pipitone, B. Kamgar-Parsi and R.L. Hartley, "Three dimensional computer vision for micro air vehicles," in Enhanced and Synthetic Vision 2001.
- [11] J. Kim and S. Sukkarieh, "Airborne Simultaneous Localisation and Map Building," in Proceedings of IEEE International Conference on Robotics and Automation, 2003.
- [12] G.L. Barrows, J.S. Chahl and M.V. Srinivasan, "Biologically inspired visual sensing and flight control," Aeronautical Journal, vol. 107, pp. 159-168 2766, March, 2003.
- [13] C.J. Taylor, J. Malik and J. Weber, "Real-time approach to stereopsis and lane-finding," in Proceedings of the 1996 IEEE Intelligent Vehicles Symposium, pp. 207-212, 1996.
- [14] C.J. Taylor, J. Kosecka, R. Blasi and J. Malik, "Comparative study of vision-based lateral control strategies for autonomous highway driving," International Journal of Robotics Research, vol. 18, pp. 442-453, 1999.
- [15] B.J.P. Horn, "Tsai's camera calibration method revisited," 2000.
- [16] Cloud Cap Technology, Inc, "Piccolo User's Guide."

X. BIOGRAPHY



Dr. Eric Frew is a postdoctoral research scientist at the UC Berkeley Center for Collaborative Control of Unmanned Vehicles. He has over eight years of experience working with unmanned systems, including the HUMMINGBIRD autonomous helicopter and MARS micro rover systems developed at the Stanford University Aerospace Robotics Laboratory. He received his Ph.D. and M.S. from the Aeronautics and Astronautics Department at Stanford University and his BSME from Cornell University. His research interests include unmanned systems, computer vision, self-directed collaborative navigation, and optimal distributed sensing.



Tim McGee is a graduate student at the University of California, Berkeley, working in the Center for Collaborative Control of Unmanned Vehicles. His interests include control systems, computer vision, and robotics. He received his BSME from the University of

Illinois, Urbana-Champaign and his MSME from the University of California, Berkeley.



Dr. ZuWhan Kim is currently a Research Engineer in computer science division and California PATH at the University of California at Berkeley. His research interests include computer vision, artificial intelligence, and their application on various areas including vision-based unmanned air vehicle control, intelligent transportation systems, and geographic information systems. He received his Ph.D. in computer science from the University of Southern California in 2001. He is a member of IEEE Computer Society.

Xiao Xiao is a first year graduate student at the University of California, Berkeley studying controls engineering in the mechanical engineering department, from where he had just earned his B.S. degree. He has enjoyed flying model airplanes since high school and is glad to participate in a research project that is so closely related to the beloved hobby. His research is on the improvement of the current autopilot system.

Stephen Jackson is an undergraduate at UC Berkeley focusing on controls and dynamics. He began working at the Center for Collaborative Control of Unmanned Vehicles in 2003. He will graduate in 2005 as a member of the mechanical engineering's honors society (Pi Tau Sigma) and plans to further his education in controls as a graduate student.



Michael Morimoto is an undergraduate junior in the mechanical engineering department at UC Berkeley who is just getting introduced to the world of controls. He plans to pursue controls after graduating because it is an interesting and exciting field of study.



Sivakumar Rathinam is currently a graduate student in the Systems program at the University of California, Berkeley. His research interests include optimization, control and mechanics. He obtained his B.Tech. degree in mechanical engineering from Indian Institute of Technology-Madras, India in 1999 and M.S. degree in mechanical engineering from Texas A & M University, College Station in 2001.

Jose Padial is undergraduate students in the Mechanical Engineering Department at the University of California, Berkeley.



Dr. Raja Sengupta is currently an Assistant Professor engaged in research on Information Technology for Complex Systems in the department of Civil and Environmental Engineering at the University of California at Berkeley. His research interests are in unmanned air vehicle systems, wireless networking for transportation, and sensor networks for forest fire management. He is Deputy Director for the Center for Collaborative Control of Unmanned Vehicles and Director of the PATH wireless laboratory. Prior to his current appointment he worked as a research engineer in the California PATH program where he developed the automotive wireless networking research program. He received his Ph.D. in Systems Engineering from the University of Michigan in 1995. He led the system safety research for the National Highway System Consortium till 1997 and worked briefly for Mitretek Systems. He is Associate Editor of the IEEE Control Systems Magazine and of the Journal of Intelligent Transportation Systems.