Fly spy: lightweight localization and target tracking for cooperating air and ground robots

Richard T. Vaughan, Gaurav S. Sukhatme, Francisco J. Mesa-Martinez, and James F. Montgomery

Robotics Research Laboratories, University of Southern California,
Los Angeles, CA 90089-0781

vaughan|gaurav|mesamart|monty@robotics.usc.edu

Abstract

Motivated by the requirements of micro air vehicles, we present a simple method for estimating the position, heading and altitude of an aerial robot by tracking the image of a communicating GPS-localized ground robot. The image-to-GPS mapping thus generated can be used to localize other objects on the ground. Results from experiments with real robots are described.

Keywords: robot helicopter localization cooperation MAV

1 Introduction

Payload volume, mass and power consumption are critical factors in the performance of aerial robots. Very small "micro" air vehicles (MAVs) are particularly desirable for reasons of cost, portability and, for military applications, stealth. Building on several years work with small robot helicopters, we are interested in minimalist approaches to localization which reduce the number of sensors that such robots carry with them. This paper presents a simple method for estimating the position, heading and altitude of an aerial robot using a single on-board camera. Many applications of aerial robots such as reconnaissance and target tracking already require an on-board camera. We exploit this sensor to simultaneously perform localization. Since ground platforms are more suited to carry larger sensor payloads, and are typically localized, collaboration and data sharing between ground and aerial robots is employed to provide the extra information needed to localize the air vehicle.

2 Related work

The work that most closely parallels the system discussed in this paper is the "visual odometry" system at CMU [1] which can visually lock-on to ground objects and sense relative helicopter position in real time. The system tracks image pair templates using custom vision hardware. Coupled with angular rate sensing, the system is used to stabilize small helicopters over reasonable speeds.

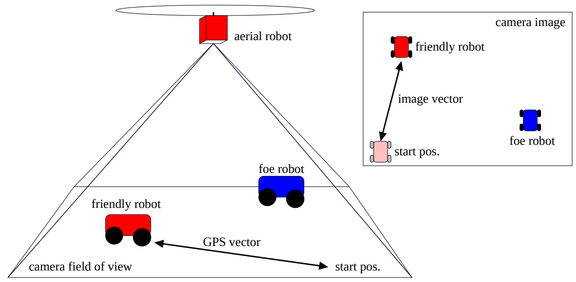

Fig. 1. Cooperating MAV/AGV scenario

Approaches to automate helicopter control include both model-based and model-free techniques. Pure model-based RC helicopter control systems have mostly used linear control. Feedback linearization, a design technique for developing invariant controllers capable of operating over all flight modes, was used in [2] and [3]. In [2], the nonlinear controller robustness was increased using sliding mode and Lyapunov-based control. An approach using output tracking based on approximate linearization is discussed in [4]. Neural networks were combined with feedback linearization in [3] to create an adaptive controller with the same goal in mind.

Various techniques to *learn* a control system for helicopters have been proposed. These include a genetic algorithm-based approach [5] to discover fuzzy rules for helicopter control; a combination of expert knowledge and training data to generate and adjust the fuzzy rule base [6].

We have successfully implemented [7] a behavior-based control approach to the helicopter control problem, described in Section 4.1. Our past work on the subject has included a model-free method for generating a helicopter control system which also has an online tuning capability using teaching by showing.

3 Approach

If an MAV can observe the relative positions of itself and two objects with known location on the ground below it, it can localize itself completely by triangulation. If the objects on the ground were friendly robots ('Unmanned Ground Vehicles' - UGVs) with on-board localization, they could inform the MAV of their positions and even send updates as they moved. This requires only a camera, modest computation and communications. All these resources are likely to be required on the MAV for other tasks; this localization technique requires no dedicated resources.

Alternatively, a single cooperating robot can move over time to establish a baseline on the ground and on the image plane of the MAV's camera, as shown in Figure 1. This is subject to some assumptions about the behaviour of the MAV, but it is the minimal configuration for such a system and is the scenario investigated in the rest of this paper.

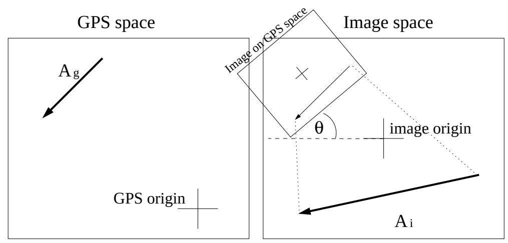

Fig. 2. Mapping a vector in image space $\boldsymbol{A_i}$ to its correlate in GPS space $\boldsymbol{A_g}$

Additional advantages of such a scheme are (i) scalability: a single ground robot can cooperatively localize many MAVs, effectively sharing its expensive and bulky GPS sensor among a swarm of aerial vehicles; (ii) the localization estimate is not subject to drift, and as such could be used to correct parallel high-resolution but drifting inertial sensors.

3.1 Assumptions

We assume these simplifying assumptions and constraints:

1. the camera is mounted perpendicular to the plane of the ground, pointing straight down;
2. the ground in the field of view is essentially flat;
3. the camera is mounted directly under the center of rotation of the MAV (though an off-center mounting can easily be compensated for by a rigid body transform);
4. the MAV is in the same place and pose when the two vector-end samples are taken.
5. over short distances (<100m) GPS space is approximately Cartesian; GPS latitude and longitude coordinates can be converted to meters by a scaling factor that is a function of latitude). The error introduced by this approximation is small compared to errors elsewhere in the system.

3.2 Heading

As the GPS coordinate system is aligned with north and the camera coordinate system is aligned with the helicopter, the vehicle's heading is the rotation angle θ of the coordinate systems, shown in Figure 2.

Given that (x_{i0}, y_{i0}) are the image coordinates of the UGV at $t = 0$ and $x_{g0}, y_{g0})$ are the GPS coordinates of the UGV at $t = 0$:

$$\boldsymbol{A_i} = \begin{pmatrix} x_i - x_{i0} \\ y_i - y_{i0} \end{pmatrix} \quad \boldsymbol{A_g} = \begin{pmatrix} x_g - x_{g0} \\ y_g - y_{g0} \end{pmatrix} \qquad \theta = \angle \boldsymbol{A_g} - \angle \boldsymbol{A_i}$$

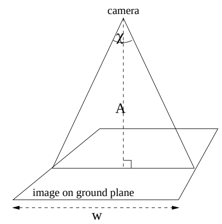

Fig. 3. The altitude calculation

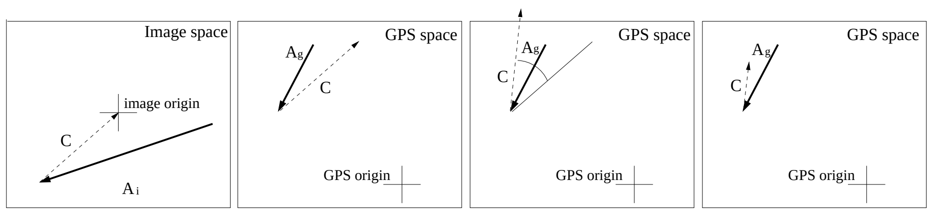

Fig. 4. Mapping the origin of image space into GPS space by translating (left) rotating (middle) and scaling (right) vector C.

3.3 Altitude

The scaling between image and GPS systems is the ratio of the lengths of the vectors in each coordinate frame $S = \frac{|\boldsymbol{A_i}|}{|\boldsymbol{A_g}|}$.

Once the scaling is found, given the camera aperture angle α of the camera and the width W of the image in pixels, the altitude estimate a can be obtained: $a = S \frac{\frac{W}{2}}{\tan\frac{\alpha}{2}}$ (Figure 3).

Now a is the height above the ground in meters. If instead we want to know the absolute altitude above sea level, we can add a to the GPS altitude measured by the UGV.

3.4 Position

The MAV's (x,y) position is assumed to be at the center of the image i.e. at the origin. To find the GPS location of the MAV we find the GPS location of the image origin. To do this we construct the vector $\boldsymbol{C}$ from the position of the current image sample to the image origin (Figure 4 top left) and transform it into GPS space. First we translate C to the position of the current GPS sample (x_g, y_g) (Figure 4 bottom left), then rotate it by θ (Figure 4 bottom middle), then scale by S to complete the coordinate transformation (Figure 4 bottom right). The transformed vector C now points to the corresponding GPS location of the image origin, giving us the MAV's position (x_u, y_u).

$$\begin{pmatrix} x_u \\ y_u \end{pmatrix} = S \left[\begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x_g - x_i \\ y_g - y_i \end{pmatrix} \right]$$

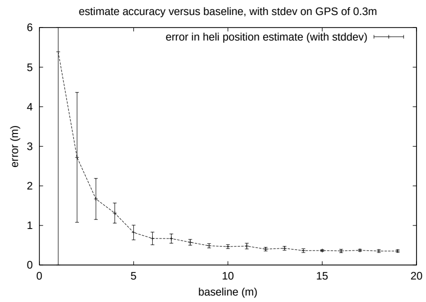

Fig. 5. Empirically determined MAV localization accuracy versus GPS baseline for a standard deviation on GPS measurements of 0.3m.

3.5 Target localization

With the parameters of the image-to-GPS coordinate transform completely known, any pixel in the image can be mapped to GPS space. Thus we can estimate the position of any object tracked in the image. This has application for aerial mapping and reconnaissance, friend/foe identification and target interception.

We demonstrate this in the next Section by estimating the GPS location of a 'foe' robot placed arbitrarily in the MAV's field of view.

3.6 Accuracy constraints

The accuracy of the MAV and target localization estimates are limited by four main factors. Of these, the first two are directly related to the assumptions described above; the last two are due to limitations in the accuracy of the sensors.

1. **Camera movement:** the method requires that the two coordinate frames to not move relative to one another in the time interval between the two samples. In particular, the method is sensitive to rotation of the coordinate frames. However, camera movement can be compensated for if an estimate of the movement is available. This estimate in turn is subject to inaccuracy, but very accurate, very small inertial sensors are available that might afford movement compensation.
2. **Camera misalignment:** any displacement of the camera from vertical will induce error. This also can be compensated for if measurements of the offset angles are available.
3. **Image sampling resolution:** the resolution of the images puts a lower bound on the attainable accuracy of the system, even assuming a perfect object tracking system. For a MAV at high altitudes a ground vehicle may have to move a long way to produce a good image vector.
4. **GPS measurement error:** over small distances, the effect of GPS measurement error is significant. The maximum accuracy attainable is a function of the measurement error and the baseline, the length of the vector between

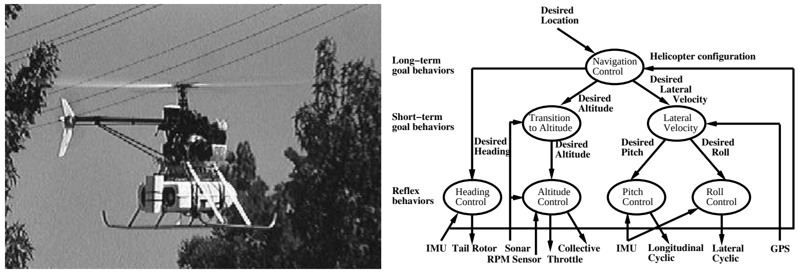

Fig. 6. The AVATAR Unmanned Aerial Vehicle (left) and its control system (right)

the two samples. To achieve good accuracy, the length of a GPS vector must be significantly longer than the variance in the error of its measured end points.

For example our robot helicopter's onboard GPS reports an average standard deviation of its measurements of 0.3m. By generating simulated GPS data with a normally distributed error with this deviation, and applying the localization algorithm above, we obtain the baseline/accuracy curve shown in Figure 5. We must expect the performance of the the real system to be somewhere below this curve.

4 Demonstration

The localization method described above is attractively simple, but has some fairly strong constraints which could make it impractical. In order to assess its utility for real world applications we implemented the system on our real MAV and UGV systems. This section describes the implementation and compares the localization estimates generated by the method with those obtained from an independent combined GPS/IMU sensor.

4.1 MAV platform

The USC AVATAR (Autonomous Vehicle Aerial Tracking And Reconnaissance) MAV is a gas-powered model helicopter fitted with a PC104 computer and custom control electronics (Figure 6 left) [7,8]. It has been developed through three generations over the last 9 years in our lab. It carries a high-quality Inertial Measurement Unit (IMU), a Novatel RT20 GPS receiver/decoder, an engine RPM sensor and a color video camera. A laser altimeter is currently being integrated. Communication with ground workstations is via 2.4GHz wireless Ethernet and 2.3GHz wireless video.

AVATAR can autonomously servo to GPS locations and orient itself to GPS headings. The AVATAR control system is implemented using a hierarchical behavior-based control system architecture [9]. Briefly, the behavior-based control approach partitions the control problem into a set of loosely coupled computing modules called 'behaviors'. Each behavior is responsible for a specific task

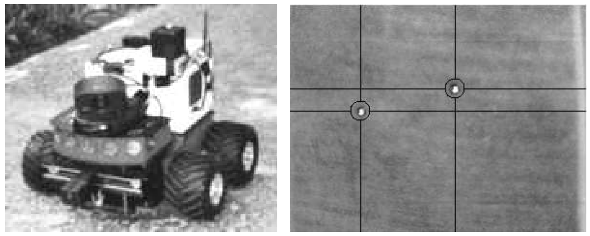

Fig. 7. The Pioneer AT robot UGV (left); an example image from the helicopter camera, showing two Pioneer AT UGVs correctly identified (right)

and they act in parallel to achieve overall robot goals. For example a low-level behaviour controls the vehicle's pitch angle by changing the tilt of the rotor head; a higher-level behaviour controls the forward speed by generating appropriate desired pitch angles; a top-level behavior controls the vehicles position by generating desired forward and lateral speeds. A schematic of the controller architecture is shown in Figure 6 (right).

For the experiments in this paper the AVATAR was required to hover and maintain the same 3D location and heading as accurately as possible. At the time of the experiments the altitude control loop had not been closed, so our pilot controlled altitude manually to maintain as constant an altitude as possible.

4.2 UGV platform

Our UGV platform is an ActivMedia Pioneer AT robot, augmented with a PC104 stack as shown in Figure 7 (left). It has five forward and two side facing sonars, a compass and wireless Ethernet. It carries a Novatel GPS receiver/decoder which is older and less accurate than that carried by AVATAR.

The Pioneer's controller allows us to servo to a specified GPS location and orient to a compass heading while avoiding obstacles using sonar. In this experiment the UGV is required simply to move over the field at a constant speed and heading to establish the GPS and image vectors. The Pioneer broadcasts its GPS location on the local network's broadcast channel at 5Hz. The AVATAR monitors these messages to obtain the GPS vector.

4.3 Object tracking

A simple blob tracking algorithm is used to track the position of high contrast regions in the image. The experiments were carried out over an open grassy field so the black and white UGVs stand out clearly from the background. The images are sampled from the MAV's wireless video stream by a low-cost framegrabber on a Pentium II workstation. An update rate of 25Hz was achieved while reliably tracking two objects. Figure 7 (right) shows an example frame from the MAV video camera with the locations of the two UGVs correctly tracked by the software.

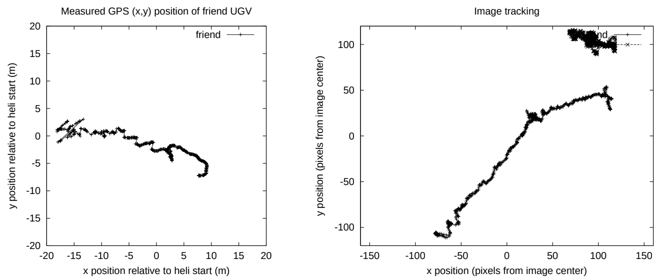

Fig. 8. Inputs to the estimator: Friend UGV GPS path (left) and tracked image positions of friend and foe (right).

4.4 Procedure

The AVATAR was used to represent a modestly capable MAV. It can hold its position for a few seconds, capture video images of the ground below it, receive data from a friendly ground vehicle, and perform some computation. The localization geometry requires relatively little computation; the object tracker is more greedy, but several systems exist that perform this task in hardware.

AVATAR's onboard GPS/IMU systems were used to maintain the hover and their position and pose measurements logged. These logs were used as the ground truth to compare with the estimates produced by the geometric method.

AVATAR was commanded to hover above the center of our test field. Two Pioneer UGVs were placed in arbitrary start positions below it. Once we determined that the UGVs were being tracked by the vision system, the initial image positions were recorded and the trial began. The friend UGV then moves approximately 10m to the west and stops for a few seconds, then moves west another 15m, by which time it has disappeared from the camera's field of view. Each time AVATAR receives a GPS location from the friend it generates an (x, y, z, θ) localization estimate for itself and an (x, y) estimate for the foe UGV based on the vectors from the initial to current positions in GPS and image space. Data were recorded for 80 seconds.

4.5 Results

Figure 8 shows the inputs to the estimator. The GPS plot (left) shows the locations reported by the friend's GPS as it traces a path from $(8, -7)$ to $(-17, 0)$. The scale is in meters with the origin at the initial helicopter position. The image plot (right) shows the friend's path through the helicopter camera image as detected by the tracker.

Figure 9 (left) shows the measured GPS positions of the helicopter and foe UGV. The helicopter drifts slightly around its initial position and the foe UGV is stationary at approximately $(17, -5)$ but the apparent variation in position shows the varying error in its GPS measurement.

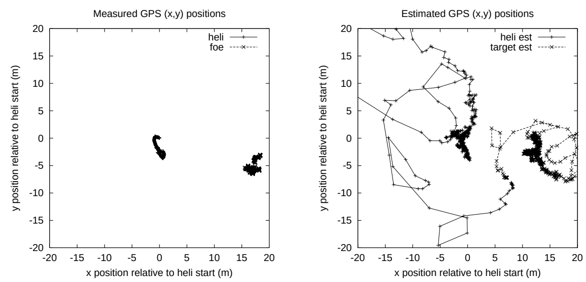

Fig. 9. Helicopter and foe locations recorded by on-board GPS (left) and estimated (right).

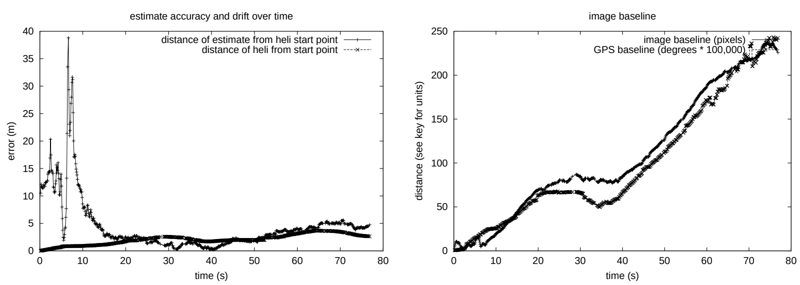

Fig. 10. MAV localization error and MAV drift over time.

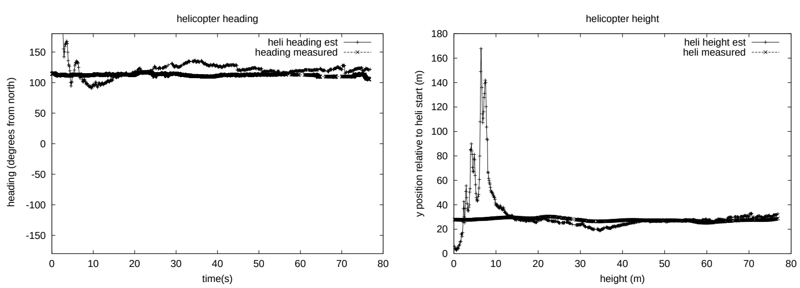

Fig. 11. Measured and estimated MAV heading (left) and altitude (right).

The right hand side of Figure 9 shows the estimated positions of the helicopter and foe. Comparing this plot with the left hand side, both the traces appear to converge reasonably close to the corresponding measured points. Figure 10 (left) shows the distance error in helicopter (x, y) localization over time. The estimate starts out poor at $t = 0s$, but quickly begins to approximate the measured value. We would predict the initially poor but improving estimate from consideration of Figure 5.

The length of the GPS and image baselines are plotted in Figure 10 (right). The slope of these curves is approximately constant, indicating a constant translation speed for the UGV, except between $t = 20s$ and $t = 40s$, when the UGV was stationary. We would expect the estimate to improve as a function of the GPS baseline. However in the error plot Figure 10 (left) we see that the error reaches a minimum at around $t = 40s$, then increases slowly over time. This is due to the lateral drift in the helicopter's position, violating assumption 3 in Section 3.1. The amount of drift is also plotted in Figure 10 and it can be seen that the drift is of a similar order to the error it induces. We conclude that a small lateral drift in MAV position may be tolerable in some applications, though we suspect that this may not be true of a change in heading.

Figure 11 (left) shows the heading of the helicopter as measured by its onboard GPS/IMU hardware and that estimated by our method. Again the error is large for the first few seconds but decreases rapidly; from t=10s until the end of the trial the estimate remains within 10% of the measured value. Note that the measured helicopter heading varied very little, allowing generation of a good estimate.

Figure 11 (right) shows the helicopter's measured and estimated height above the ground. Again the estimate is initially poor, but begins to approximate the measured value after a few seconds. We omit error plots for heading and altitude for lack of space, but as can be seen from the estimate plots, the errors do not increase over time. We conclude that these estimates are also rather insensitive to the lateral drift in helicopter position.

Over the length of the trial, AVATAR's GPS reported average latitude stdev of 0.20m, longitude 0.34 m, an average of 0.27m which is comparable to the simulated plot of stdev 0.3m in Figure 5.

5 Conclusions and further work

We have demonstrated a simple geometric method that approximately localizes an aerial robot using minimal sensing. We suggest it may be suitable for use in micro air vehicles, as it exploits general-purpose sensors that are likely to be carried already and requires only modest computation. It also allows multiple MAVs to share a single GPS receiver carried by a cooperating ground robot. We aim to further evaluate this method and to increase the interaction between air and ground vehicles. In particular we aim to direct the friendly ground robot to intercept the foe robot using the estimate calculated by the MAV. The localization accuracy observed in these experiments suggests that this is feasible.

Acknowledgements

The authors gratefully acknowledge the talent and hard work of Kale Harbick, Gerry Bohne and the rest of the USC AVATAR team. Thanks also to Brian Gerkey and Goksel Dedeoglu for their help with the Pioneer robots. This work is supported by DARPA grant DABT63-99-1-0015, contract F04701-97-C-0021 and contract DAAE07-98-C-L028.

References

1. O. Amidi, T. Kanade, and K. Fujita. A visual odometer for autonomous helicopter flight. In *Proceedings of IAS-5*, 1998.
2. D. Y. Maharaj. *The Application of Nonlinear Control Theory to Robust Helicopter Flight Control*. PhD thesis, Department of Aeronautics, Imperial College of Science, Technology, and Medicine, 1994.
3. B. S. Kim. *Nonlinear Flight Control Using Neural Networks*. PhD thesis, Department of Aerospace Engineering, Georgia Institute of Technology, 1993.
4. T. J. Koo and S. Sastry. Output tracking control design of a helicopter model based on approximate linearization. In *Proceedings of IEEE Conference on Decision and Control*, Orlando, Florida, 1998.
5. C. Phillips, C. L. Karr, and G. Walker. Helicopter flight control with fuzzy-logic and genetic algorithms. *Engineering Applications of Artificial Intelligence*, 9(2):175–184, April 1996.
6. M. Sugeno, M. F. Griffin, and A. Bastian. Fuzzy hierarchical control of an unmanned helicopter. In *17th IFSA World Congress*, pages 179–182, 1993.
7. J. F. Montgomery, A. H. Fagg, and G. A. Bekey. The USC AFV-I: A behavior-based entry in the 1994 International Aerial Robotics Competition. *IEEE Expert*, 10(2):16–22, April 1995.
8. J. F. Montgomery and G. A. Bekey. Learning helicopter control through 'teaching by showing'. In *1998 IEEE Conference on Decision and Control*, December 1998.
9. R. A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal Robotics and Automation*, 2(1):14–23, March 1986.