In-situ unmanned aerial vehicle (UAV) sensor calibration to improve automatic image orthorectification

CONFERENCE PAPER · JULY 2010

DOI: 10.1109/IGARSS.2010.5652989 · Source: DBLP

4 AUTHORS:

Austin M. Jensen
Utah State University
39 PUBLICATIONS **188** CITATIONS
SEE PROFILE

Norman Wildmann
University of Tuebingen
12 PUBLICATIONS **54** CITATIONS
SEE PROFILE

YangQuan Chen
University of California, Merced
664 PUBLICATIONS **11,003** CITATIONS
SEE PROFILE

Holger Voos
University of Luxembourg
108 PUBLICATIONS **282** CITATIONS
SEE PROFILE

IN-SITU UNMANNED AERIAL VEHICLE (UAV) SENSOR CALIBRATION TO IMPROVE AUTOMATIC IMAGE ORTHORECTIFICATION

Austin M. Jensen[a], Norman Wildmann[a,b], YangQuan Chen[a], Holger Voos[b]

[a]Utah State University
8200 Old Main Hill Logan, UT, USA
Austin.Jensen@aggiemail.usu.edu

[b]Hochschule Ravensburg-Weingarten
Weingarten,Germany

ABSTRACT

Small, low-altitude unmanned aerial vehicles (UAV)s can be very useful in many ecological applications as a personal remote sensing platform. However, in many cases it is difficult to produce a single georeferenced mosaic from the many small images taken from the UAV. This is due to the lack of features in the images and the inherent errors from the inexpensive navigation sensors. This paper focuses on improving the orthorectification accuracy by finding these errors and calibrating the navigation sensors. This is done by inverse-orthorectifying a set of images collected during flight using ground targets and General Procrustes Analysis. By comparing the calculated data from the inverse-orthorectification and the measured data from the navigation sensors, different sources of errors can be found and characterized, such as GPS computational delay, logging delay, and biases. With this method, the orthorectification errors are reduced from less than 60m to less than 1.5m.

Index Terms— Unmanned Aerial Vehicle, General Procrustes Analysis, Calibration, Image Orthorecification, Personal remote sensing system

1. INTRODUCTION

Small, low-altitude unmanned aerial vehicles (UAV)s can be helpful as remote sensing platforms; they can be low-cost, easy-to-use, quick to deploy, and have high spatial resolution. AggieAir™, an unmanned aerial system (UAS) developed at Utah State University (USU), has demonstrated the effectiveness of the UAV as a remote sensing platform [1, 2]. Even though the UAV itself is inexpensive, conventional post processing and georeferencing of the imagery can still take significant time and resources due to the small image footprint and the high number of images needed to cover an area. For example, 100 images may be required to cover a two kilometer squared area. One could georeference all of the images using a basemap and common features in each image; however, this process can be very time consuming and costly, making the UAV platform less competitive with conventional remote sensing platforms. In addition, many cases lack the necessary features in each image to georeference them to a basemap. A quicker and less expensive approach to georeferencing the images would be to orthorectify them based on the position and orientation of the camera at the time the image was taken. This process can be completed autonomously; however, the low-cost sensors that measure position and orientation of the aircraft have some inherent errors, which are then projected onto the ground with the image. Depending on the altitude, this can cause up to 60m of error on the ground. In an effort to address this issue, some have improved the location accuracy of a ground target by loitering above it, sampling the location many times, and filtering the data [3]. Others have improved the location accuracy of a ground target by using multiple UAVs [4]. Since they are applied to a single ground target, these methods may be ineffective for creating maps of an area. The best solution to the errors in the aircraft sensors may be to quantify them by in-situ calibration of the system [5]. This method includes setting up ground control points with defined geodetic positions measured by a precise GPS receiver, taking aerial photographs of the control points, and using the control points to inverse orthorectify the images to find the actual position and orientation of the camera when the picture was taken. Even though Jensen et. al. [5] was able to reduce the errors induced by the aircraft sensors with the in-situ calibration, the method used to inverse orthorectify the images was not able to find the roll and pitch of the aircraft. This limitation prevented the accurate characterization of some of the errors. This paper uses a different, more accurate inverse orthorectification method that can also find the roll and pitch of the aircraft. This more accurate pose estimation allows us to characterize not only the position and

This work is supported in part by UWRL MLF grants WR-1057, WR-1011 and WR-1102.

orientation, but also the delays caused by the GPS receiver and the synchronization between the datalog and the cameras [6].

2. GENERAL PROCRUSTES ANALYSIS

General Procrustes Analysis (GPA) uses least-squares fitting to find the transform matrix between two 3D point sets [7, 8]. For this application, the 3D point sets are the locations of the ground control points on the images (camera frame), and the geodetic position of the ground control points on the earth (navigation frame). Given the point sets, GPA calculates the centroid of the set of points in each coordinate system, finds the distance from each point in the coordinate system to the centroid, and solves the least squares problem by minimizing equation 1, where d_{ci} is the distance of each point to the centroid in the camera frame, m_{ci} is the distance of each point to the centroid in the navigation frame and R is the transform matrix. It is assumed that the measured position and orientation of the UAV are the coordinates of the camera frame and are used as the initial conditions of the least-squares fitting.

$$\sum \|d_{ci} - Rm_{ci}\|^2 \tag{1}$$

After a transform matrix R is found, which minimizes equation 1, the translation between a point d in the camera frame and a point m in the navigation frame can be found with equation 2.

$$T = d - Rm \tag{2}$$

At this point, R and T can be used to find the position and orientation of the camera. However, for this application, better results are found by using non-linear least-squares fitting and iterating until the result converges. Figure 1 shows a plot of the sum of the squared errors versus the iteration number from data collected from a test flight. The significant reduction in error shows the benefit of using the non-linear least-squares method.

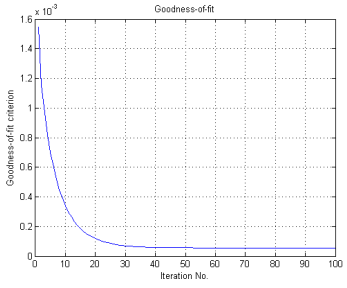

Fig. 1. Sum of squared errors vs iteration number.

3. CAMERA CALIBRATION

For accurate orthorectification, the intrinsic and extrinsic (orientation of the camera with respect to the UAV) parameters of the camera must be known. CalCam, a professional camera calibration program by MosaicMill Ltd. Finland [9], is used to find the intrinsic parameters: principal point, focal length, affinity, skew, radial distortion and tangential distortion.

When the images are inverse-orthorectified, the GPA method will find the orientation of the camera frame with respect to the navigation frame. However, to calibrate the navigation sensors, the orientation of the body frame (aircraft coordinate system) with respect to the navigation frame will be needed. Therefore, the orientation of the camera frame with respect to the body frame must be found. This can be done before the flight by using the UAV to take multiple pictures of a chessboard poster, and using GPA to find the orientation from the inner corners of the chessboard on the image and their true positions on the chessboard poster. Given the orientation of the camera with respect to the chessboard (using GPA) and the orientation of the aircraft body with respect to the chessboard (using orientation data from IMU), the orientation of the camera with respect to the body can be found. Figure 2 shows some test data from this process before a flight. The data was collected in an open area where the magnetic field was not distorted by buildings, power lines, etc. For each of the plots, the y axis is the orientation of the camera and the x axis is the orientation of the UAV. As expected, all of the plots are very linear with a slope of about 1. The bias in each axis represents the orientation of the camera frame with respect to the body frame.

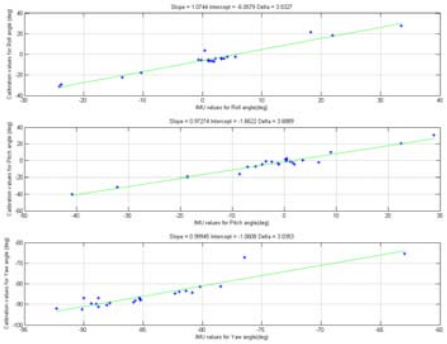

Fig. 2. Results from extrinsic parameter calibration.

4. CALIBRATING AIRCRAFT SENSORS

To calibrate the aircraft sensors, ground control points are set up on the ground in an unsymmetrical pattern to get a unique solution from GPA. The ground control points are also placed so they are well distributed in each of the images even though the UAV is flown at a range of altitudes. After the ground

control points are laid out and geolocated using a survey grade GPS receiver, the UAV is flown over the targets using a flower navigation routine. The flower navigation routine flies the UAV over a given point until it reaches a predetermined distance from that point. Once the distance is reached, the UAV will turn around and fly over the point again. Since the altitude of the UAV can be changed at any point, this routine ensures the images will be taken over the targets with a wide range of altitudes and headings. After the UAV lands, the targets in each of the images are identified and located using automatic target recognition. Using the target locations in the images and the target locations on the ground, the GPA method is used to calculate the position and orientation of the UAV for each image. Comparing the position and orientation found by GPA with the sensor data that was logged in synchronization with the cameras allows us to find, quantify, and correct many features of the sensors: the biases in the altitude, yaw, and position, the misalignment of the cameras with the aircraft coordinate system, and the time delays caused by the GPS receiver and the synchronization between the datalog and the cameras.

5. RESULTS

One large contributer to the orthorectification error is delay. There are two delays in the system: the delay between the camera exposure and when the UAV data for that image is logged, and a computational delay in the GPS. Unlike the roll, pitch and yaw of the aircraft, which are only affected by the logging delay, the GPS data is affected by both the computational and the logging delay. Therefore to find the computational delay, both the GPS delay (computational plus logging) and the logging delay must be known.

One way to find the GPS delay is by using the UAV log, which has all of the UAV data at a higher frequency, and checking the error while stepping forward or backward through the log from the time the picture was assumed to have been exposed. Figure 3 shows the latitude and longitude error with respect to the time shift in the high frequency log. The point at which the error is smallest indicates the point in the log with the correct position of the UAV when the picture was actually exposed. The shift is approximately 0.75s.

The GPS delay can also be found using the kinematics of the UAV. Since the UAV is moving, the speed of the aircraft is related to the position error and the time delay (eq. 3).

$$TimeDelay = \frac{PositionError}{Speed} \tag{3}$$

Figure 4 shows the x and y components of speed plotted with the x and y components of the position error. The plots have a slope of 0.89 and 0.85, which gives some support to the delay found in the previous method.

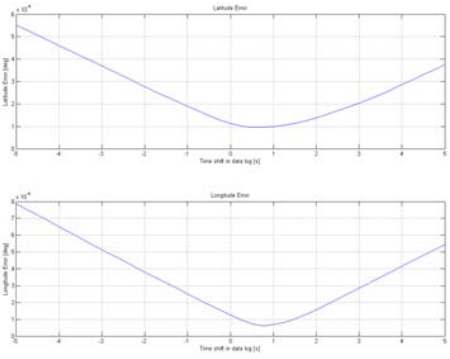

Fig. 3. Delay using GPS position.

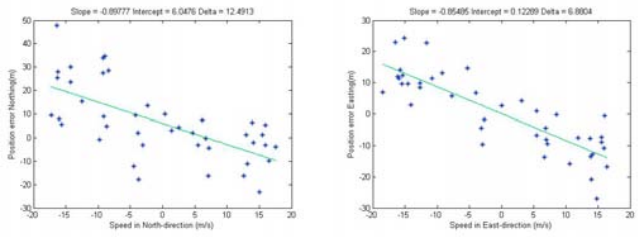

Fig. 4. Delay using speed.

After applying the GPS delay to the data from the GPS (position and altitude), the errors in position improved by 15m and the errors in altitude improved by 2m.

Once the GPS delay is known, it can be separated into the logging delay and the GPS computational delay by using the roll and pitch data. Similar to the method used to find the GPS delay, the UAV log can also be stepped through to find the logging delay. However, instead of using the roll and pitch data directly, the slope of the linear fit between the measured data and calculated data is used. This is due to the noisy IMU data; the statistical approach of a linear fit gives a much clearer result. Figure 5 shows the slope of the roll and pitch data in the high frequency log versus the time shift from the point at which the UAV data for the image was logged. Since the accuracy improves as the slope goes to 1, the plot shows that the UAV data for the image was logged 0.25 seconds after the actual exposure. Subtracting the GPS delay (0.75s) gives a computational delay of 1s (fig. 6). This is confirmed by comparing the improvement in the GPS data after applying the GPS delay (15m) and the average speed of the UAV during the experiment (15 m/s).

After correcting for the delay, figures 7 and 8 show the relationship between the calculated values and the measured values for altitude, yaw, roll and pitch. The altitude and yaw plots look very linear with a slope of about 1 and biases of 3.6m and 11 deg respectively. The results from the roll and pitch, however, are not as conclusive. This is probably due to the limited range of values that were sampled (10 degrees for

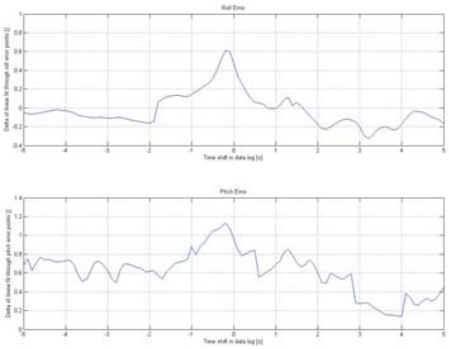

Fig. 5. Delay using roll and pitch.

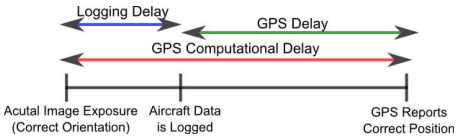

Fig. 6. Delay Explanation.

pitch, 20 degrees for roll, 360 degrees for yaw).

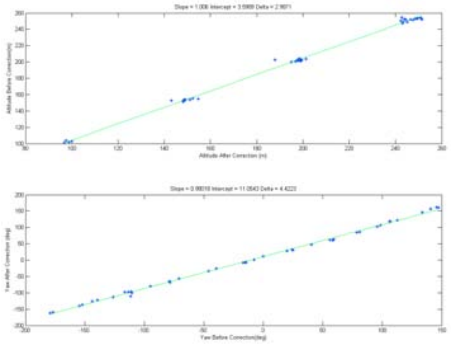

Fig. 7. Altitude and yaw plot.

After correcting the aircraft data for each of the images, the position error for each of the ground control points (through orthorectificiation) is reduced to less than 1.5m.

6. CONCLUSIONS AND FUTURE WORK

Using General Procrustes Analysis to inverse orthorectify the images made it possible to characterize many sources of error from the navigation sensors on the UAV: GPS computational delay, logging delay, and biases in the altitude, yaw, roll, and pitch. Future improvements include finding a way to increase the range of values sampled from the roll and pitch to improve their characterization and finding a real-time method that takes less time to set up in order to make the process more practical for regular use.

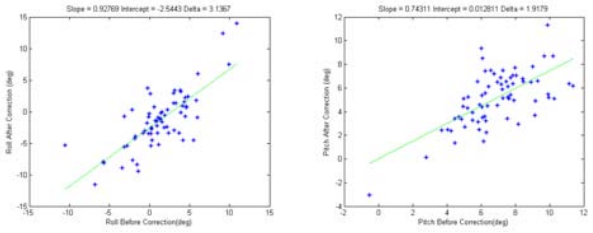

Fig. 8. Roll and pitch after delay correction.

7. REFERENCES

[1] Austin M. Jensen, YangQuan Chen, Thom Hardy, and Mac McKee, "AggieAir - a low-cost autonomous multispectral remote sensing platform: New developments and applications," in *Proc. IEEE International Conference on Geoscience and Remote Sensing Symposium IGARSS 2009*, July 2009.

[2] Haiyang Chao, Austin M. Jensen, Yiding Han, YangQuan Chen, and Mac McKee, *Advances in Geoscience and Remote Sensing*, chapter AggieAir: Towards Low-cost Cooperative Multispectral Remote Sensing Using Small Unmanned Aircraft Systems, pp. 463 – 489, INTECH, 2009.

[3] J.D. Redding, T.W. McLain, R.W. Beard, and C.N. Taylor, "Vision-based target localization from a fixed-wing miniature air vehicle," in *Proc. American Control Conference*, June 2006.

[4] J. Tisdale, A. Ryan, Zu Kim, D. Tornqvist, and J.K. Hedrick, "A multiple UAV system for vision-based search and localization," in *American Control Conference, 2008*, June 2008, pp. 1985–1990.

[5] Austin M. Jensen, Yiding Han, and YangQuan Chen, "Using aerial images to calibrate inertial sensors of a low-cost multispectral autonomous remote sensing platform (AggieAir)," in *Proc. IEEE International Conference on Geoscience and Remote Sensing Symposium IGARSS 2009*, July 2009.

[6] Norman Wildmann, "Techniques towards increased precision in direct-georeferencing for AggieAir, a low-cost personal remote sensing system," M.S. thesis, Hochschule Ravensburg-Weingarten, February 2010.

[7] K. S. Arun, T. S. Huang, and S. D. Blostein, "Least-squares fitting of two 3-D point sets," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, no. 5, pp. 698–700, Sept. 1987.

[8] M. Devrim Akca, "Generalized procrustes analysis and its applications in photogrammetry," 2003.

[9] MosaicMill, "http://www.mosaicmill.com/," June 2010.