Cooperative navigation for UAV and ground units

Jonas Nygårds and Joakim Rydell
Department for C4ISR
FOI, Swedish Defence Research Agency
Box 1165
581 11 Linköping, Sweden
SWEDEN

jonas.nygards@foi.se, joakim.rydell@foi.se

ABSTRACT

This paper presents results from experiments on cooperative navigation performed during the NATO SET-229 RTG. Image-based navigation is used to estimate the position of a UAV, while dismounted ground units are positioned using foot-mounted inertial sensors. When the ground units are observed by the UAV, their position estimates are refined based on these observations. The positioning systems utilized on the ground units and UAV are described, along with a cooperative navigation algorithm for fusing data from the two systems. Results from a trial performed by the RTG are reported, and show that this type of cooperative navigation has the potential to reduce the ground units' position error significantly.

1. INTRODUCTION

This paper presents experiments on cooperative navigation that were performed in Sennybridge, Wales during the NATO SET-229 RTG. In this work, cooperative navigation refers to methods for exchanging navigation information between different units, with the purpose of improving the accuracy of their respective navigation solutions. In this paper, we present how cooperative navigation can be applied to a UAV and one or more ground units, operating in GNSS denied conditions.

In the examined approach, the UAV estimates its position by comparing images from an onboard camera to a georeferenced orthographic photo of the environment, while the ground units use foot-mounted IMUs. When moving ground units are detected by the UAV, it tracks and estimates their positions. These position estimates and the corresponding uncertainties are thereafter transferred to the ground units, which use these estimates to improve their own navigation accuracy.

2. FOOT-MOUNTED NAVIGATION FOR GROUND UNITS

The ground units have a stand-alone system based on foot-mounted sensors using dead reckoning. At the foot, an IMU is mounted and using the fact that the foot is stationary regularly for a short period during each step, a zero velocity update allows for compensation of IMU drift where each step can be integrated separately [1]. The system still exhibit a drift especially in the heading direction but this approach can provide a decent position accuracy as a stand-alone system, even using low-cost MEMS-based IMU's, for shorter periods. The foot-mounted system used herein is based on the OpenShoe project [2].

3. IMAGE-BASED UAV NAVIGATION

The UAV estimates its position by matching images from an onboard camera to a georeferenced orthographic photo of the operation environment. The position of the UAV is determined by finding the part of the georeferenced photograph, which is most similar to the current image. To avoid ambiguities, a particle filter combines the image matching results with velocity estimates from the stream of camera images.

The georeferenced image can be taken from online map services such as Google Earth, or be created from own images using photogrammetry software. It is important that the area has not changed too much since the image was acquired, and that the image resolution (pixels per meter) is sufficient. What constitutes a significant change and a sufficient resolution varies with the flight altitude. When flying higher, small details in the reference image have less effect on the matching result, which reduces the resolution requirement and allows more changes to the environment. In our experiments, including the Sennybridge trials, the altitude has generally been approximately 35 meters. At that altitude, a resolution of 1 – 2 pixels/meter is sufficient. Vehicles and smaller objects do not significantly affect the matching, but buildings, roads and other large objects need to be accurately represented in the reference image.

The onboard camera is mounted in a gimbal, which stabilizes it in the nadir direction (downwards). Acquired images are therefore unaffected by the UAV's pitch and roll angles, which simplifies the image matching process. Altitude and heading are estimated using data from the barometer and magnetometer in the UAV's flight controller, leaving only the latitude and longitude to be determined from image matching. This is considerably easier than solving the full pose estimation problem, where altitude, roll, pitch and yaw also need to be estimated. After scaling and rotating the camera image to match the scale and orientation of the reference image, matching scores for each position in the reference are computed using image correlation.

The velocity is estimated by computing the optical flow (the apparent movement between consecutive images). This provides a measurement in pixels per second in image coordinates, which is converted to an actual velocity in world coordinates by rotating and scaling according to the heading and altitude measurements from the flight controller.

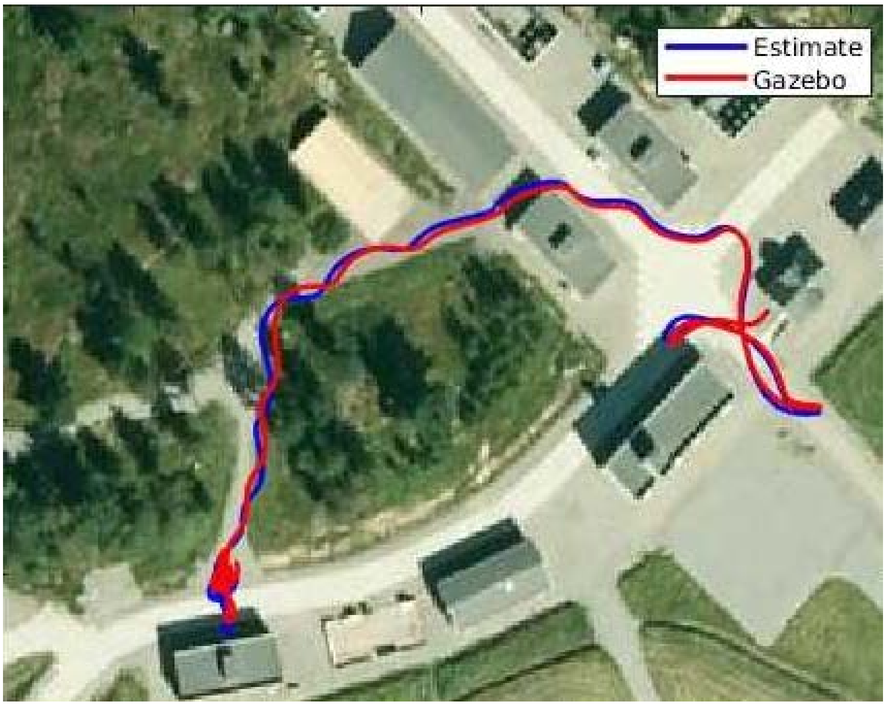

Figure 1: Actual (red) and estimated (blue) UAV trajectories during a simulated flight, overlaid on the reference image used for navigation.

Finally, the velocity is used to propagate particles in a particle filter. Each particle represents one hypothesis about the latitude and longitude of the UAV. The particles' weights are updated using the correlation values at their respective positions in the reference image, and the weighted mean of all particles is used as the final estimate of the UAV's position.

The positioning algorithm is implemented in the Robot operating system (ROS, [3]), and runs in realtime on an Nvidia Jetson TX2 computer which is mounted on the experimental UAV. It also runs in a simulated environment, which facilitates development and evaluation. Figure 1 shows actual and estimated flight trajectories from an experiment in the simulated environment. The estimate generally agrees with the ground truth data reported by the simulator, but errors exist. These errors are correlated over time, which may cause problems when using these position estimates for collaborative navigation.

The algorithm is inspired by [4], and is presented in more detail in [5].

4. DETECTION OF GROUND UNITS IN AERIAL IMAGES

Moving ground units are detected in the images by finding image regions where the optical flow does not match the estimated UAV velocity. For each such observation, its position in world coordinates is computed using information about the position, orientation, and lens properties of the UAV camera. Again, the horizontal position is obtained from image matching; altitude and heading come from onboard barometer and magnetometer sensors; the pitch and roll are always zero since the camera is stabilized in the downward direction by a gimbal; and the lens properties are known from camera calibration.

False detections are removed by tracking the potential ground units in world coordinates. A potential ground unit is considered valid when it has been tracked for a few seconds, if it has moved a sufficient distance from its initial position. The latter condition eliminates most false detections caused by parallax, where for example a roof appears to move relative to the ground because of its shorter distance from the camera. Such false detections generally move only a small distance, which depends on the relation between the height of the detection above the ground, and the altitude of the camera above the ground.

The algorithm does not distinguish between different types of ground units, and can be used to detect dismounted ground units as well as manned or unmanned vehicles.

Figure 2 shows a reference image (in grayscale), with an image from the UAV camera overlaid (in color). The large blue marker shows the estimated UAV position, while the squares show tracked ground units (walking persons in this example). The circles shown in the color image show moving objects detected in this image. Note that the two images have different orientations.

Figure 2: UAV positioning and tracking of ground units. (Images ©FOI, with thanks to SENTA Range Commandant)

5. FUSION FOR IMPROVED GROUND UNIT NAVIGATION ACCURACY

As tracks from the UAV are received by a ground unit, they can be associated with its own motion. The cooperative navigation can be seen as a special case of distributed data fusion where [6] is a good introduction. As with other tracking tasks the association problem is the first obstacle to handle. In the current experiments the amount of clutter was low and the initial uncertainty of the dismounted unit low enough that a simple nearest neighbor combined with a chi2-test was sufficient. For more complex association tasks we have experimented on delayed association by matching parts of the trajectory to achieve sufficient certainty. The error in the tracks from the UAV will in general be correlated over time and to handle this the fusion with the foot-mounted system is achieved using both a filtering method in combination with one smoothing method. The filter used provides consistent estimates even for correlated errors by using the inverse covariance intersection method (ICI, [7]). The inverse covariance is seen as the information in a Fisher-information sense. In the ICI method a worst-case estimate of the common information in two information sources is calculated giving consistent results for the fusion of two sources of information even under an unknown amount of common information. By using the worst-case assumption the ICI provides a conservative estimate. In situations when the correlation can be modeled/known an estimator based on such models will have a smaller variance. However, in the present case the communicated information between units does not allow this more precise modeling.

The ICI results are used to provide an approximation of the common information to be used in an update to a smoothing filter based on the Georgia Tech Smoothing and mapping package (GTSAM, [8]). The approximation is based on subtracting the common information from the information provided by the UAV to keep the covariance mostly consistent. However in the ICI the common information is compensated for in the total sum, but when subtracting from only the UAV-provided information the remaining information may become negative. Under the approximation used, only zero information was used in such cases otherwise the GTSAM solution became unstable.

Figure 3: Cooperative navigation results. The red dots represent unaided navigation of the foot-mounted system. Blue circles are the result of aiding by the UAV tracking results. (ortophoto background © Google & Bluesky)

The results for the experiments during SET 229 are illustrated in Figure 3. In the scenario, a supporting ground unit is waiting by the end of the woods and initially has a good GNSS signal coverage, but after loosing the GNSS signals errors accumulate. This was emulated by the uncertainty accumulated when walking to the intial point. Note that the unit then is almost stationary at the staging point for a long while, accumulating orientation errors while stepping around at almost the same point. Due to mainly orientation errors the path illustrated by red dots deviates so that it’s unclear at what side of the house the dismounted unit arrives at. By supporting with measurements from the UAV the error is reduced and the path follows the right hand side of the road and ends in front of the house just as in real life.

5. DISCUSSION & FUTURE WORK

The SET-229 trials showed, in early experiments, the value of supporting the navigation systems on dismounted and mounted platforms with UAV camera measurements. The experiments where conducted in relevant scenarios and complemented by other interesting experiments. We are grateful for DSTL providing access to the training area and providing the necessary operation of the UAV. For future work we intend to examine ways to better account for correlated information within a smoothing and mapping framework like GTSAM [8] and the early smoothing and mapping results in [9] by for instance also using shared landmarks, marginalisation [10] and better approximation for the worst case common information as provided by ICI.

The image-based navigation currently requires the UAV to observe the scene from above. This requirement stems from the nadir-facing camera as well as the image matching process, where the camera image is assumed to be similar to some region in the orthographic reference image. In order to relax these constraints so that the UAV can operate freely, even at very low altitude (eye level), a new matching algorithm is being developed. The new algorithm matches images from the onboard camera to rendered images from a 3D model representing the area of operation. It therefore allows any position and orientation of the UAV and its camera (and can also be used on UGVs, dismounted soldiers, and other platforms). This comes at a higher computational cost, but by combining absolute position measurements from the matching algorithms with positioning based on visual odometry, relatively sparse position fixes from the matching algorithm are expected to be sufficient.

REFERENCES

[1] J. Elwell, "Inertial navigation for the urban warrior," in Proc. SPIE 3709, Digitization of the Battlespace IV, 1999.

[2] J.-O. Nilsson, I. Skog, P. Händel and K. V. S. Hari, "Foot-mounted INS for everybody-an open-source embedded implementation," in Proceedings of IEEE/ION PLANS 2012, 2012.

[3] Stanford Artificial Intelligence Laboratory et al., 2018. Robotic Operating System, Available at: https://www.ros.org.

[4] G. Conte and P. Doherty, “Vision-based unmanned aerial vehicle navigation using geo-referenced navigation,” EURASIP Journal of Advances in Signal Processing, 2009.

[5] J. Rydell, E. Bilock and M. Tulldahl, “Computationally Efficient Vision-based UAV Positioning”, ION International Technical Meeting (ITM), Reston, Virginia, USA, January 2019.

[6] M.E. Campbell, and N.A Ahmed, “Distributed data fusion: Neighbors, rumors, and the art of collective knowledge,” IEEE Control Systems , Vol. 36, No. 4, pp. 83-109, 2016.

[7] B. Noack, J. Sijs, M. Reinhardt and U. D. Hanebeck, ”Decentralized data fusion with inverse covariance intersection”, Automatica, vol. 79, pp. 35-41, 2017.

[8] F. Dellaert, Factor graphs and GTSAM: A hands-on introduction, 2012. https://gtsam.org/tutorials/intro.html

[9] L. Andersson, and J. Nygards. "C-SAM: Multi-robot SLAM using square root information smoothing." 2008 IEEE International Conference on Robotics and Automation. IEEE, 2008.

[10] P. Thompson, and S. Sukkarieh, “Tracking multiple features including cross-feature correlations, with observation parameter uncertainties,” 9th International Conference on Information Fusion, 2006, pp. 1-8, 2006