

A VISUAL NAVIGATION SYSTEM FOR UAS BASED ON GEO-REFERENCED IMAGERY

Gianpaolo Conte and Patrick Doherty

Artificial Intelligence and Integrated Computer System Division
Linköping University
SE-581 83 Linköping, SWEDEN
gianpaolo.conte@liu.se

KEY WORDS: UAS, Terrain-aided Navigation, Sensor Fusion, Visual Navigation, Inertial Sensors

1 ABSTRACT

This article presents an approach to the terrain-aided navigation problem suitable for unmanned aerial vehicles flying at low altitude. The problem of estimating the state parameters of a flying vehicle is addressed in the particular situation where the GPS information is unavailable or unreliable due to jamming situations for instance. The proposed state estimation approach fuses information from inertial and image sensors. Absolute localization is obtained through image-to-map registration. For this purpose, 2D satellite images are used. The algorithms presented are implemented and tested on-board an industrial unmanned helicopter. Flight-test results will be presented.

2 INTRODUCTION

Operating Unmanned Aerial Systems (UAS) in non-segregated or civil airspace represents a technical challenge. Several problems have to be solved in order to safely operate such systems in populated areas. As an example, an important problem is represented by the capability of an UAS to autonomously perceive and avoid obstacles. Such characteristic is often addressed in the literature as the "See and Avoid" capability. A general robust solution to this problem is non trivial.

Another important problem, which is the focus of this article, is related to GPS vulnerability. Modern GPS receivers are capable of providing global position and velocity information with high accuracy. Such information is of vital importance for autonomous navigation for two reasons: it gives an absolute localization means and it can be used to estimate and compensate for the inertial sensors errors. Normally, a UAS navigation system relies on inertial sensors and GPS for estimating the platform states. The main states of a UAS are the position, velocity, acceleration, attitude angles and attitude rates. A precise estimation of the states is essential for accurate autonomous control of the platform. When the GPS is unreliable or unavailable the estimation of the states becomes inaccurate and the UAS autonomous capabilities are not fulfilled.

The inertial sensors, or inertial navigation system (INS), are composed of 3 gyroscopes which measure the 3D angular rates, and 3 accelerometers which measure the 3D acceleration. Time integration of the inertial sensors provides platform state information at high rate. The drawback is that, during the integration, the inaccuracies of the inertial sensors cause a drift of the computed states from the real platform states. The drift problem can be solved with an appropriate sensor fusion algorithm combining the INS solution with an external position information. The Kalman filter represents the standard fusing algorithm which optimally solves this kind of state estimation problems.

The major concern when using a GPS system is its vulnerability to jamming (Volpe, 2001). The availability of GPS jamming



Figure 1: The Yamaha Rmax UAV helicopter.

technology on the market makes a UAS unsafe because susceptible to jamming actions. The scope of this article is to propose a solution to the state estimation problem in situations of GPS unavailability. To achieve this goal, the proposed solution employs an image sensor (monocular video camera) to supply for the missing GPS information. It is evident that it is very difficult, for any alternative system, to compute position and velocity information with accuracy comparable to the one provided by state of the art GPS receivers. The solution proposed provides a graceful degradation of the state estimation preserving the autonomous flight capabilities of the platform.

The vision system is composed by an image sensor and an image processing software. The main components of the image processing software are the *visual odometry* and *geo-referenced image registration*. Visual odometry and geo-referenced imagery provide a complementary source of information. Visual odometry provides information on the UAS displacement (relative position) while geo-referenced imagery provides absolute position information. Both information are required in order to obtain a robust solution. The growing availability of high resolution satellite images (for example provided by Google Earth) makes this topic quite attractive. Approaches which use terrain information for navigation are also referred to as terrain-aided navigation (TAN).

TAN problems are usually non-linear and have a multi-modal probability distribution. For such problems Kalman filter techniques cannot be applied. One possible approach is to decompose the state estimation problem in two parts. One part is related with the position estimation, which is usually the non-linear and multi-modal part, while the other part addresses the velocity and attitude estimation where Kalman filter applies. This approach is followed in the work presented in this article and extended details can be found in (Conte, 2009).

TAN approaches can be divided in two categories, approaches which require an a priori terrain map and approaches which tries to build the map of the environment on-line.

Approaches which tries to build a map on-line are better known under the acronym SLAM (Simultaneous Localization and Mapping). The goal of SLAM is to localize a robot in the environment while mapping it at the same time. Prior knowledge of the environment is not required. In SLAM approaches, an internal representation of the world is built on-line in the form of a landmarks database. Such a representation is then used for localization purposes. For indoor robotic applications SLAM is already a standard localization technique. More challenging is the use of such technique in large outdoor environments. Some examples of SLAM applied to aerial vehicles can be found in (Karlsson et al., 2008, Kim and Sukkarieh, 2007).

Terrain databases used for TAN can be of divers nature. Some approaches make use of terrain elevation database (Bergman, 1999). A direct measurement of the flight altitude relative to the ground is required. Matching the ground elevation profile, measured with a radar altimeter for instance, to an elevation database allows for aircraft localization. A commercial navigation system called TERNAV (TERrain NAVigation) is based on such a method. Navigation systems based on this principle have been implemented successfully on some military jet fighters and cruise missiles. In case of small UAVs and more specifically for unmanned helicopters, this method does not appear to be appropriate. Compared to jet fighters, UAV helicopters cover short distances at very low speed so the altitude variation is quite poor in terms of allowing ground profile matching.

The terrain database used in this work is represented by 2D orthorectified reference images. The use of reference images for aircraft localization purposes is also investigated in (Sim et al., 2002). The work focuses mainly on the image processing issues and gives less emphasis to architectural and fusion schemes which are of focus in this article.

TAN problem is of great interest not only for terrestrial applications but also for future space missions. One of the requirements for the next lunar mission in which NASA/JPL is involved, is to autonomously land within 100 meters of a predetermined location on the lunar surface. Traditional lunar landing approaches based on inertial sensing do not have the navigational precision to meet this requirement. A survey of the different terrain navigation approaches can be found in (Johnson and Montgomery, 2008) where methods based on passive imaging and active range sensing are described.

The scope of this article is to present experimental results of a TAN approach where a single video camera is used to extract two different type of information: relative displacement (odometry) and absolute position (image registration). This is a very practical and innovative concept. The approach presented is implemented on a Yamaha Rmax unmanned helicopter (Figure 1) and tested on-board during autonomous flight-test experiments.

3 SENSOR FUSION ALGORITHMS

The sensor fusion architecture presented in this article is depicted in Figure 2.

The state estimation problem is divided in two parts. A Kalman filter (KF) is used to estimate the INS errors $\delta\hat{x}^k$ while a point-mass filter (PMF) is used to estimate the horizontal (north-east) position \hat{x}^P . The reason why the state estimation problem has been decomposed in this way is that the estimation of $\delta\hat{x}^k$ can be modeled as unimodal linear Gaussian process while for \hat{x}^P a full probability distribution has to be maintained over the 2D image registration area.

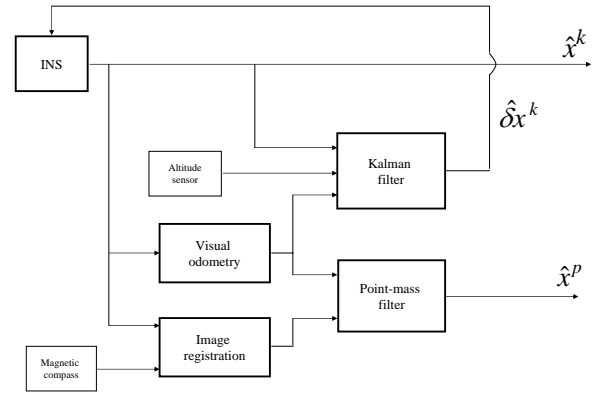


Figure 2: Sensor fusion architecture.

3.1 Kalman filter

The KF implemented estimates the INS errors instead of the platform states directly (error dynamics formulation). The advantage of such formulation lies in the fact that the dynamics of the INS errors are much slower than the platform dynamics itself, hence, the error dynamics KF can be implemented with a slower rate. The estimated errors are used to correct the INS navigation solution. This is realized in practice by feeding back the estimated errors to the INS as it is shown in Figure 2. For the KF implementation a linearized INS error dynamic model is used. More details can be found in (Britting, 1971). The state-space model used in the KF has the following standard linear structure:

$$\delta\dot{x}^k = \mathbf{F}\delta x^k + \mathbf{G}u \quad (1)$$

$$\delta y^k = \mathbf{H}\delta x^k + e \quad (2)$$

where Equation 1 represents the dynamics model of the state evolution and Equation 2 is the measurement equation.

$$\delta x^k = [\delta h \ \delta v_n \ \delta v_e \ \delta v_d \ \epsilon_\phi \ \epsilon_\theta \ \epsilon_\psi \ \delta a_x \ \delta a_y \ \delta a_z]^T$$

represents the KF error-state vector (altitude, 3D velocity, 3 attitude angles and 3 accelerometer biases) and

$$\delta y^k = [h^{ins} - h^{baro}, \ v_n^{ins} - v_n^{odom}, \ v_e^{ins} - v_e^{odom}]^T \quad (3)$$

represents the measurement vector (difference between the INS solution and sensor measurement). For the altitude measurement a barometric sensor is used (h^{baro}). The barometer measures the absolute atmospheric pressure. This sensor is usually used to compute the differential pressure between a point of known altitude (usually the take-off position) and the current flight position. The differential pressure is then converted in altitude variation (meters) relative to the take-off position. The velocity north (v_n^{odom}) and velocity east (v_e^{odom}) measurements are derived from the odometry. Finally, u and e are the process and measurement, Gaussian distributed, white noises. The optimal estimation of the error-state vector $\delta\hat{x}^k$ is computed from the well known Kalman filter recursion.

3.2 Point-mass filter

The point-mass filter (PMF) is used to fuse measurements coming from the visual odometry and image registration modules. The PMF computes the solution to the Bayesian filtering problem on a discretized grid. In (Bergman, 1999) such a technique was applied to a TAN problem where a digital terrain elevation database was used instead of digital 2D images. The PMF is particularly suitable for this kind of problems since it handles general probability distributions and non-linear models.

The discretized dynamic model employed to solve the 2D localization problem is represented as follows:

$$x_t^p = x_{t-1}^p + u_{t-1}^{odom} + v \quad (4)$$

$$p(y_t^p | x_t^p) \quad (5)$$

where Equation 4 represents the dynamic equation for the state evolution and 5 the measurement likelihood.

$$x^p = [p_n \ p_e]^T$$

is the 2D state vector representing the north and east position, u^{odom} is the position displacement coming from the visual odometry and v is the process noise. The process noise is assumed white and Gaussian distributed.

It is worth to notice that, while the state evolution Equation 4 is described as a linear Gaussian process, the measurement likelihood (5) is expressed in the most general form, which means that no assumptions are made on its shape (in the KF the likelihood is assumed to be Gaussian). The measurement likelihood $p(y^p | x^p)$ is provided by the image registration step as it will be explained later in the article.

The solution \hat{x}^p of the state estimation problem (4)-(5) is obtained by solving the Bayesian filtering recursion on a discretized 2D grid mesh. The grid used in this application is uniform with N number of points and resolution δ . The discretized recursion is:

$$p(x_t^p(k) | Y_{t-1}) = \sum_{n=1}^N p_v(x_t^p(k) - u_{t-1}^{odom} - x_{t-1}^p(n)) p(x_{t-1}^p(n) | Y_{t-1}) \delta^2 \quad (6)$$

$$\alpha_t = \sum_{n=1}^N p(y_t(n) | x_t^p(n)) p(x_t^p(n) | Y_{t-1}) \delta^2 \quad (7)$$

$$p(x_t^p(k) | Y_t) = \alpha_t^{-1} p(y_t(k) | x_t^p(k)) p(x_t^p(k) | Y_{t-1}) \quad (8)$$

$$\hat{x}_t^p = \sum_{n=1}^N x_t^p(n) p(x_t^p(n) | Y_t) \delta^2 \quad (9)$$

$$P_t = \sum_{n=1}^N (x_t^p(n) - \hat{x}_t^p)(x_t^p(n) - \hat{x}_t^p)^T p(x_t^p(n) | Y_t) \delta^2 \quad (10)$$

with $Y_{t-1} = y_{1:t-1}$ and $k = 1 \dots N$. The terminology $p(x_t^p(k) | Y_t)$ indicates the value of the probability density function corresponding to the point k on the grid. Equations 6 and 8 represent the time and measurement updates respectively while the estimation of the

minimum variance state vector and its uncertainty are computed from Equations 9 and 10 respectively.

The convolution operation in the time update (6) requires N^2 operations and computationally is the most expensive operation of the recursion. For this reason the PMF approach becomes computationally unfeasible for real-time applications when the state dimension is greater than 3 or 4. In this application we have a 2D state vector and the PMF works excellently even in real-time. In this application a uniform static grid mesh gives a sufficient accuracy but adaptive grid algorithms also can be used. The idea with the adaptive grid is to dynamically increase the grid resolution in regions with high probability while decreasing the resolution in other regions. This should produce a more accurate estimation solution. More details on the PMF implementation can be found in (Bergman, 1999).

A different algorithmic approach to the Bayesian filtering problem is represented by the particle filter. In the particle filter approximation, the Bayesian recursion is solved on a stochastic grid where the grid is represented by particles. The particles become more dense in regions with high probability and vice versa. The particle filter is computationally more efficient than the PMF. On the other hand other kind of problems arise when using the particle filter approach. In any case the particle filter represent a valid alternative to the PMF.

4 VISUAL ODOMETRY

Visual odometry for aerial navigation has been object of great interest during the last decade. In the early work (Amidi, 1996) it was shown how a visual odometry is capable of stabilizing an autonomous helicopter in hovering conditions. Recent works (Fritsch et al., 2008, Michaelsen et al., 2004, Caballero et al., 2005, Ready and Taylor, 2007) on visual odometry for airborne applications are based on homography estimation under a planar scene assumption. In this case the relation between points of two images can be expressed as $x_2 \propto Hx_1$ where x_1 and x_2 are the corresponding points of two images 1 and 2 expressed in homogeneous coordinates and H is the 3x3 homography matrix. The relation is valid up to a scale factor. The visual odometry described in this article is based on robust homography estimation.

The camera rotation and displacement between two camera positions $c1$ and $c2$, can be computed from the homography matrix decomposition (Hartley and Zisserman, 2003):

$$H = K \left(R_{c1}^{c2} + \frac{1}{d} \bar{t}^{t2} \bar{n}^{c1T} \right) K^{-1} \quad (11)$$

where K is the camera calibration matrix determined with a camera calibration procedure, \bar{t}^{t2} is the camera translation vector expressed in the camera 2 reference system, R_{c1}^{c2} is the rotation from camera 1 to camera 2, \bar{n}^{c1} is the unit normal vector to the plane being observed and expressed in the camera 1 reference system and d is the distance of the principal point of the camera 1 to the plane. For this work a pin-hole camera model is used.

The homography matrix H is estimated from a set of corresponding corner features being tracked from frame to frame. H can be estimated using direct linear transformation (Hartley and Zisserman, 2003) with a minimum number of four feature points (the features must not be collinear). In practice the homography is estimated from a higher number of corresponding features (50 or more). The feature tracker used in this work is based on the

well-known KLT algorithm. The algorithm selects a number of features (corners) in an image according to a "goodness" criteria described in (Shi and Tomasi, 1994). Then it tries to re-associate the same features in the next image frame. In any case, incorrect feature association can occur. A method to identify incorrect feature correspondences is therefore desirable. One popular method is the random sample consensus (RANSAC) algorithm (Fischler and Bolles, 1981). In Figure 3 the RANSAC algorithm has been applied on a set of features tracked in two consecutive frames. On the left are represented the feature displacements computed by the KLT algorithm while on the right the set of outlier features has been detected and removed using RANSAC.




Figure 3: On the left is displayed the set of features tracked with the KLT algorithm. On the right the outlier feature set has been identified and removed using the RANSAC algorithm.

More advanced features like SURF (Bay, 2006) could be also used in this context. SURF features are invariant to image translation, scaling and rotation enhancing the odometry robustness.

Once the homography matrix has been estimated it can be decomposed into its rotation and translation components in the form of Equation 11 using singular value decomposition as described in (Hartley and Zisserman, 2003). However, homography decomposition is a poorly-conditioned problem especially when using cameras with a large focal length (Michaelsen et al., 2004). The problem arises also when the ratio between the flight altitude and the camera displacement is high. For this reason it is recommendable to use inter-frame rotation information from other sensor sources if available. The odometry presented in this work uses rotation information from the INS. Equation 11 can be rearranged as follows:

$$\vec{t}^n \vec{n}^{nT} = d \left(R_n^{c2T} K^{-1} H K R_{c1}^n T - I \right) \quad (12)$$

where \vec{t}^n is the camera displacement vector expressed in the navigation reference frame (north-east-down), R_n^{c2} is the rotation matrix from the navigation frame to the camera 2 frame and R_{c1}^n is the rotation matrix from camera 1 to the navigation frame. In (Conte, 2009) can be found more details on the derivation of Equation 12.

The rotation matrices R_n^{c2} and R_{c1}^n are taken from the compensated INS solution (Figure 2). This fact creates an undesirable correlation between the measurement noise and the process noise in the KF since the odometry information is used to update the KF. This problem was analyzed in (Conte, 2009) through Monte Carlo simulations where a dependency between the altitude d and the estimation accuracy was found. It was shown that below 200 meters altitude the state estimation accuracy was marginally affected by the correlation problem. On the other hand, the accuracy degrades at higher altitude.

In order to compute \vec{t}^n from Equation 12, the parameters d and \vec{n}^n must be found. The distance to the ground d can be measured from an on-board laser or radar altimeter. Alternatively,

a terrain altitude database could be used, if available, to derive both d and \vec{n}^n . The flight-tests presented in this article were performed over a flat area, therefore the differential barometric altitude, measured with an on-board barometric sensor, was directly used as a measurement of the ground altitude. The flat terrain condition imposes also $\vec{n}^n = [0, 0, 1]$.

Finally, indicating with RHS the right hand side of Equation 12, the north and east helicopter displacement computed by the odometry are:

$$\begin{aligned} t_{north}^n &= RHS_{1,3} \\ t_{east}^n &= RHS_{2,3} \end{aligned} \quad (13)$$

The computed (north-east) odometry displacement are used as input \vec{t}^{odom} in Equation 4. For the Kalman filter update a north and east velocity (v_n^{odom} , v_e^{odom}) were required (Equation 3). This was simply obtained dividing the odometer displacements by the image processing rate.

5 IMAGE REGISTRATION

Image registration is the process of overlaying two images of the same scene taken at different times, from different viewpoints and by different sensors. The registration geometrically aligns two images (the *reference* and *sensed* images). Image registration has been an active research field for many years and it has a wide range of applications. A literature review on image registration can be found in (Zitova and Flusser, 2003). The assumption of small attitude angles has been adopted here. This means that perspective deformation has not been taken into account during the registration process. The assumption is somehow reasonable in this specific context since it can be assumed that, in case of emergency (like the loss of GPS signal), the UAV flight envelope it is restricted.

For the application described in this article the registration position of the two images is not required but only the degree of similarity of the sensed image for every location of the reference image. The similarity criterion used is the normalized cross-correlation of intensity images (Pratt, 1991). After the correlation map between the two images is found, the likelihood $p(y^p|x^p)$, which is used as observation in the PMF, is computed by a normalization step.

Before computing the correlation map, the sensed image is scaled and aligned to the reference image using the altitude information from the pressure sensor (as for the odometry) and heading information from a magnetic compass. Pitch and roll angles are small for this application therefore the associated image deformation is neglected. After the alignment and scaling steps, the cross-correlation is computed. If S is the sensed image and I is the reference image, the expression for the two-dimensional normalized cross-correlation is:

$$C(u, v) = \frac{\sum_{x,y} [S(x, y) - \mu_S][I(x-u, y-v) - \mu_I]}{\sqrt{\sum_{x,y} [S(x, y) - \mu_S]^2 \sum_{x,y} [I(x-u, y-v) - \mu_I]^2}} \quad (14)$$

where μ_S and μ_I are the average intensity values of the sensed and the reference image respectively. Figure 4 depicts a typical




Figure 4: Typical normalized cross-correlation result between the sensed image and reference image. The distribution represents the likelihood of the measurement observation used to update the PMF.

cross-correlation result between a sensed image taken from the Rmax helicopter and the reference image of the flight-test site.

For computational reasons, the cross-correlation is performed on a restricted window of the reference image. The window is centered around the current platform position estimates. It must be pointed out that the correlation between the reference and sensed images does not correspond to the helicopter absolute position likelihood directly. A horizontal offset correction must be applied to the cross-correlation map in order to compensate for the camera view angle. The offset correction is easily found using the flight altitude and attitude angles information.

6 EXPERIMENTAL RESULTS

The proposed filter architecture has been implemented and tested on-board the industrial Yamaha Rmax UAV helicopter (Figure 1). The total helicopter length is about 3.6 m. It is powered by a 21 hp two-stroke engine and it has a maximum take-off weight of 95 kg. The avionic system was developed at the Department of Computer and Information Science at Linköping University and has been integrated in the Rmax helicopter. The platform developed is capable of fully autonomous flight from take-off to landing.

The inertial measurement unit used is composed of 3 accelerometers and 3 gyros and it is built-in the Yamaha Rmax helicopter. The unit outputs accelerometer data at 66Hz and gyro data at 200Hz. The data were pre-processed with an anti-aliasing filter (20Hz) before sampling. The data were sampled at 50Hz rate. The barometer and compass output data at 40Hz rate. The video camera sensor is a standard CCD analog camera with approximately 45 degrees horizontal angle of view. The camera frame rate is 25Hz and the images are reduced to half resolution (384x288 pixels) at the beginning of the image processing pipeline to reduce the computational burden.

The filter architecture is implemented on 2 on-board pc104 PentiumIII 700MHz computers. Network communication between computers is physically realized with serial line RS232C and Ethernet. Ethernet is mainly used for remote login and file transfer, while serial lines are used for hard real-time networking. The implementation of the filtering architecture has been split among the two computers. In the first pc104 is implemented the Kalman filter running at 50Hz (and the control system for autonomous flight). On the second pc104 are implemented the visual odometry, image registration and point-mass filter all running at 7Hz rate. Data were exchanged between the two PCs through serial line. For the PMF a 80x80 meters grid of one meter resolution was used. Consequentially, the image registration was computed

on a moving 80x80 meters restricted portion of the reference image. The flight-tests were performed in an emergency services training area in the south of Sweden. The reference image of the area used for this experiment is an ortho-rectified aerial image of 1 meter/pixel resolution. The results of the terrain-aided navigation algorithm are compared to a navigation solution given by an INS/GPS Kalman filter assumed as reference. The RTK GPS used has a sub-meter position accuracy.

Figure 5 depicts the evolution of the filter probability density function (PDF). The initial density for x_0^p is assumed Gaussian distributed. It can be observed how, after a few iterations, the PDF evolves developing several peaks confirming the multi-modal nature of the problem.




Figure 5: Evolution of the PMF probability density function. The capability of the filter to maintain the full probability distribution over the grid space can be observed.

Figure 6 shows the horizontal position error. The position converges in about 20 sec to a position error below 5 meters.




Figure 6: Convergence of the position estimate.

Figure 7 shows the pitch angle error. As for the position, the pitch error converges in about 20 sec. The pitch accuracy is below 1 deg after convergence.




Figure 7: Convergence of the pitch estimate.

Figure 8 shows a UAV flight path flown at the flight-test site. The helicopter flew a closed loop path of about 1 km length at 60 m constant altitude above the ground and with 3 m/s velocity. The picture presents a comparison between the position estimated from the PMF, visual odometer and GPS reference. The PMF position is always close to the GPS position indicating a satisfactory localization performance of the approach. In addition, it can be observed how the position of the visual odometry

alone is affected by a drift which is correctly compensated using the information extracted from the geo-referenced image of the flight-test site. The results in this plot were obtained running the algorithm off-line using real flight-test data




Figure 8: The figure displays the comparison between the flight path computed with the point-mass filter (red), the odometry alone (dashed white) and the GPS reference (blue).

In the first phase of the development, the off-line tests have been helpful for tuning and refining the algorithms. In the second phase, the algorithms have been implemented on-board the Rmax helicopter and tested in flight. Figure 9 shows the flight-test results. The helicopter was partially flown in autonomous mode with the control system using the state estimated from the sensor fusion approach described in this article. The PMF path in the figure was computed on-board the Rmax during the flight-test and the flight data were downloaded after the flight. The path was flown at an altitude of 55 m above the ground with a maximum speed of 4 m/s.



Figure 9: Real-time on-board position estimation results. Comparison between the flight path computed with the PMF (red) and the GPS reference (blue).

7 CONCLUSIONS

The terrain-aided navigation approach proposed in this article is suitable for implementation on-board a UAV platform and can serve as a back-up state estimation approach in case the GPS system fails. Global localization capabilities have been demonstrated over an urban-like area with houses and a road network. Global localization will be more problematic over a rural area in absence of structured features. In any case, visual odometry performs quite well even in rural areas and it can still be used to greatly reduce the INS drift.

REFERENCES

- Amidi, O., 1996. An Autonomous Vision-Guided Helicopter. PhD thesis, Carnegie Mellon University, Pittsburgh, PA.
- Bay, H., T. T. G. L., 2006. Surf: Speeded up robust features. In: Proceedings of the ninth European Conference on Computer Vision.
- Bergman, N., 1999. Recursive Bayesian Estimation, Navigation and Tracking Applications. PhD thesis, Department of Electrical Engineering, Linköping University, Sweden.
- Britting, K. R., 1971. Inertial Navigation Systems Analysis. John Wiley & Sons, Inc.
- Caballero, F., Merino, L., Ferruz, J. and Ollero, A., 2005. A visual odometer without 3d reconstruction for aerial vehicles. applications to building inspection. In: IEEE International Conference on Robotics and Automation, Barcelona, Spain.
- Conte, G., 2009. Vision-Based Localization and Guidance for Unmanned Aerial Vehicles. PhD thesis, Dept. of Computer and Information Science, Linköping University, Sweden.
- Fischler, M. A. and Bolles, R. C., 1981. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Communications of the Association for Computing Machinery 24(6), pp. 381–395.
- Frietsch, N., Meister, O., Schlaile, C., Seibold, J. and Trommer, G., 2008. Vision based hovering and landing system for a vtol-mav with geo-localization capabilities. In: AIAA Guidance, Navigation and Control Conference and Exhibit, Honolulu, Hawaii.
- Hartley, R. and Zisserman, A., 2003. Multiple View Geometry, Second Edition. Cambridge University Press, Cambridge.
- Johnson, A. E. and Montgomery, J. F., 2008. Overview of terrain relative navigation approaches for precise lunar landing. In: Proc. of the IEEE Aerospace Conference, Big Sky, Montana.
- Karlsson, R., Schön, T. B., Törnqvist, D., Conte, G. and Gustafsson, F., 2008. Utilizing model structure for efficient simultaneous localization and mapping for a uav application. In: Proc. of the IEEE Aerospace Conference, Big Sky, Montana.
- Kim, J. and Sukkarieh, S., 2007. Real-time implementation of airborne inertial-slam. Robot. Auton. Syst. 55(1), pp. 62–71.
- Michaelsen, E., Kirchhoff, M. and Stilla, U., 2004. Sensor pose inference from airborne videos by decomposing homography estimates. Altan MO (ed) International Archives of Photogrammetry and Remote Sensing 35, pp. 1–6.
- Pratt, W. K., 1991. Digital Image Processing, 2nd ed. Wiley, New York.
- Ready, B. and Taylor, C., 2007. Improving accuracy of mav pose estimation using visual odometry. In: American Control Conference, 2007. ACC '07.
- Shi, J. and Tomasi, C., 1994. Good features to track. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR'94), Seattle.
- Sim, D. G., Park, R. H., Kim, R. C., Lee, S. U. and Kim, I. C., 2002. Integrated position estimation using aerial image sequences. IEEE Trans. Pattern Anal. Mach. Intell. 24(1), pp. 1–18.
- Volpe, 2001. Vulnerability assessment of the transportation infrastructure relying on the global positioning system. Technical report, John A. Volpe, US Department of Transportation.
- Zitova, B. and Flusser, J., 2003. Image registration methods: a survey. Image and Vision Computing (21), pp. 977–1000.