# Vision-aided IMU for handheld pedestrian navigation

C. Hide, *IESSG, University of Nottingham*
T. Botterill, M. Andreotti, *GRC, University of Canterbury*

## BIOGRAPHIES

Chris Hide is a Research Fellow at the IESSG, University of Nottingham. He has a degree in Mathematics and Topographic Science from the University of Wales, Swansea and a PhD in Engineering Surveying from the IESSG, University of Nottingham. He worked as a Senior Research Scientist between 2006 and 2009 at the Geospatial Research Centre in New Zealand. He has been working with GNSS and integrated systems for over 10 years.

Tom Botterill obtained his MA degree in Maths and Computer Science at Cambridge University in 2005 and is and is currently studying towards a PhD with the Geospatial Research Centre and Department of Computer Science at the University of Canterbury, Christchurch, New Zealand. His research focuses on navigation using computer vision.

Marcus Andreotti received BSc. and MSc. Degrees in Electrical Engineering in Brazil and a PhD in Engineering Surveying and Electrical Engineering from the IESSG, University of Nottingham. He has over 18 years R&D experience in the area of hardware and software development in signal processing and sensor integration. He is currently holding the position of Senior Research Scientist at the Geospatial Research Centre in New Zealand, where his main area of research is on signal processing methods to enable GNSS positioning in indoor environments, together with general Electronics and Communication development.

## ABSTRACT

Low cost inertial sensors are often promoted as the solution to indoor navigation. However, in reality, the quality of the measurements is poor, and as a result, the sensors can only be used to navigate for a few seconds at a time before the drift becomes too large to be useful. Therefore, it is necessary to regularly update the sensors with measurements from external systems such as GPS or other sensors useful for navigation. One such sensor is provided by the computer vision community where a camera can be used to obtain information about the relative translation and rotation between successive images.

This paper describes the use of a camera attached to a low cost IMU for navigation in areas where GPS is unavailable such as indoors or deep urban canyons. It is assumed that a pedestrian user is walking with the mobile device held out in front of them with the camera pointing approximately towards the ground. Features are matched between successive frames, and the robust RANSAC framework is used to identify which of these lie on the ground plane, while estimating the camera's orientation and 3 dimensional body frame translation relative to its previous position. This information is used to aid the IMU using a Kalman filter to reduce the position drift.

This paper describes the implementation of the combined computer vision and inertial navigation approach. A tactical grade IMU is used for initial testing since it provides more reliable measurements and enables us to provide a reference by which to compare the measurements obtained from the computer vision algorithm. It is demonstrated that even with a good quality IMU, the algorithm is able to significantly improve the performance of INS navigation when GPS measurements are unavailable.

## INTRODUCTION

Accurate pedestrian positioning still provides one of the most difficult research problems for the navigation community. Users not only navigate outside where GNSS signals are available, but they also spend large periods of time indoors where GNSS signals are degraded or not available at all. A recent advance in pedestrian navigation has been the use of low cost foot-mounted IMUs (see, for example, Foxlin 2005 and Hide et al 2009). This has the significant advantage that zero velocity updates (ZUPTs) can be used as a measurement every step the user takes. The time between steps is typically very short, which allows low cost IMUs to become a realistic solution for pedestrian navigation. However, for many applications, it is undesirable to mount a sensor on the user's foot, instead, it is preferred that the navigation sensors could be self contained on a device such as a smart phone or PDA.

The reason that foot mounted inertial sensors work so effectively is because they are constantly being updated

by external measurements, in this case ZUPTs. In fact, even high quality tactical grade inertial sensors still experience drifts of hundreds of metres after periods as short as 5 to 10 minutes. Therefore, it is clear that for any low cost IMU to work effectively, it will need to be constantly aided by an external source. One increasingly feasible area of aiding measurements comes from the computer vision community. Cameras can be used to give navigation information using sensors that are already provided by many mobile devices.

This paper examines the use of velocity aiding from a camera attached to an IMU. It is assumed that a pedestrian user is walking with the mobile device held out in front of them with the camera pointing approximately towards the ground. The camera therefore has a view of the ground beneath, and immediately in front of the user. The sequential images are then used to compute the 3 dimensional body frame translation of the camera as well as 3 dimensional rotation. Images are collected at the frame rate of the camera such as 15 or 30Hz, and features are identified within the image using standard computer vision algorithms. A robust estimator is then used to find corresponding features between subsequent images. This is achieved by assuming the camera views two scenes that form a plane (ie the ground the user is walking over is approximately flat) and that the features in both images lie on the plane. This homography can be used to remove features that do not lie on the plane, such as the user's feet and legs as they are taking a step. The robust estimator also deals with issues such as incorrect feature correspondences by removing measurements that do not conform to the model.

The computer vision algorithm results in a translation and rotation that can be used to continuously aid the IMU. One immediate issue that occurs is that only the translation is known (ie a direction vector) and not the body frame velocity. The body frame velocity can be computed if the height of the camera above the ground is known. An initial approximation of the ground height can be obtained by taking into consideration the average height that users would hold the camera above the ground, but the value needs to be estimated when other measurements are available such as GPS. The ability to estimate this value is discussed in the paper.

This paper describes our proposed implementation of a computer vision-aided IMU. In order to test the algorithm effectively, a good quality tactical grade IMU is used so that the characteristics of the computer vision measurements can be quantified. It is demonstrated that even with a good quality IMU, the computer vision measurements significantly improve the position accuracy with errors reducing from 327 metres after 6 minutes, to 14 metres. Some of the limitations of the technique are discussed including the need to estimate the camera height above ground as well as some of the issues of working in low light areas.

## INERTIAL NAVIGATION

Inertial Navigation provides the foundation of the proposed algorithm. An IMU is used that consists of three gyros and accelerometers that are used to compute the position and orientation of the mobile device. The initial position must be obtained using an estimate from an external sensor such as a GNSS receiver, or potentially other systems such as Wi-Fi. The initial orientation of the IMU is computed using a short period (typically less than 1 second) of stationary accelerometer measurements. The accelerometer measurements are used to resolve the initial roll and pitch of the IMU by comparison of the outputs with the local gravity vector. For low cost IMUs, the initial heading of the IMU is typically computed using a one-off heading measurement derived from a magnetometer. If the IMU can be considered to be non-rotating, an initial estimate of the gyro drifts can also be computed by calculating the zero-offset from the averaged gyro measurements.

After initialisation, the gyro and accelerometer measurements are used to update the 3 dimensional position and orientation using strapdown navigation equations such as those described in Titterton and Weston (1997). The orientation is stored as a quaternion that transforms measurements from the body frame to the navigation frame termed $q_b^n$. The quaternion is updated using the differential equation:

$$\dot{q}_b^n = 0.5 Q \omega_{bn}^b \qquad (1)$$

where $Q$ is a 4 by 3 matrix formed using the elements of $q_b^n$ (see Titterton and Weston (1997) for details), $\omega_{bn}^b$ is a term including the rotation vector measured by the gyros in the body frame, along with terms to correct for the rotation of the Earth and the transport rate. The latter terms can be ignored for pedestrian navigation using low cost sensors since the user is traveling in small areas at low velocity, and the IMU used is not sufficiently sensitive to measure the rotation of the Earth.

The velocity (and hence position after double integration) is computed in the navigation frame by numerically integrating the differential equation:

$$v^n = C_n^b f^b - (2\omega_{ie}^n + \omega_{en}^n) \times v^n + g^n \qquad (2)$$

where $v^n$ is the velocity in the local (North, East, Down) navigation frame; $C_n^b$ is the rotation matrix that transforms measurements from the navigation frame $n$ to body frame $b$; $f^b$ is the accelerometer measurement in the body frame; $\omega_{ie}^n$ is rotation rate of the Earth in the navigation frame; $\omega_{en}^n$ is the transport rate of the navigation frame and $g^n$ is the gravity vector in the

navigation frame. Again the middle term can be dropped depending on the application and quality of IMU used. By numerically integrating the above equations, the position, velocity and orientation of the IMU can be constantly updated using the high rate (typically 100Hz) IMU measurements. This combination of IMU and navigation equations forms an INS. However, for low cost IMUs, inertial sensor characteristics such as biases, scale factor errors, axis non-orthogonality and noise result in significant position, velocity and orientation drift after only short periods of time depending on the quality of the IMU used.

A Kalman filter is used to estimate the navigation and IMU errors. The state vector is defined as:

$$x = \begin{pmatrix} \delta p & \delta v^n & \delta \phi & \delta g^b & \delta a^b \end{pmatrix}^T \qquad (3)$$

where $\delta p$ is vector of latitude, longitude and height errors; $\delta v^n$ is the vector of navigation frame velocity errors; $\delta \phi$ is the navigation frame axis misalignment; $\delta g^b$ is the vector of gyro bias errors; $\delta a^b$ is the vector of accelerometer bias errors. The Kalman filter is used to estimate the errors using a linearised inertial navigation model such as that described in (Titterton and Weston, 1997; Farrell and Barth, 1999). The model describes the interaction between different error states and can be used to estimate the full state vector using position or velocity measurements and sufficient dynamics. The filter is used in feedback form so that when a measurement is available from a sensor, the error is computed using the Kalman filter which is then used to correct the inertial sensor measurements and navigation parameters. This is to ensure the navigation errors remain small and hence keep the linearised model valid. More information on Kalman filters and Kalman filters for inertial navigation can be found in (Hide, 2003; Foxlin, 2005; Farrell and Barth, 1999).

## COMPUTER VISION ALGORITHM

This section describes how computer vision is used to compute the motion of the camera. The camera captures a sequence of images, which show the ground plane; however the images also contain other features, such as the pedestrian's legs, feet and shadow and other objects on the ground. For each image, the relative position and orientation of the camera is estimated, relative to its position when it captured the previous image.

When an image is captured, the first stage is to detect point features in the image. The FAST corner detector (Rosten et al, 2008) is used to detect approximately 300 points in each image which are likely to be detected in other images showing the same scene. The image around each FAST corner is described using a small patch of the image: a patch sized 27x27 pixels centered on each corner is scaled down to 9x9 pixels. The similarity of two of these 9x9 patches is measured by computing by the sum-of-squared differences (SSD) between corresponding pixel values.

Each detected feature location is transformed using the camera's calibration matrix, and shifted to correct radial lens distortion.

Each patch feature from an image is matched to the most similar patch feature in the previous image. These feature matches ('correspondences') are found by computing the SSD between all pairs of patch features, and choosing the closest match to each. This takes around 100ms per pair of frames (a more efficient procedure could be used, such as a kd-tree; Beis and Lowe, 1999). When a patch feature appears similar to several patch features in the other image, all possible matches between pairs of patch features are used as correspondences.

Many of these correspondences will give the location in each image of some feature visible in both images. When these correspondences also lie on the ground plane, they are related by a perspective homography, *H*, which is a 3x3 matrix mapping homogeneous point locations in one image, (x,y,1), to homogeneous point locations in the other image, (x',y',1) (following normalization so that the third component is one). *H* can be computed from four or more correspondences using the Discrete Linear Transform, or DLT (Hartley and Zisserman, 2003, Section 4.1), a least-squares approach.

Some correspondences are not on the ground plane however, and many others will be incorrect matches caused by similar-looking features, and matches between moving features. These outlier correspondences must be removed before a least-squares approach can be used.

To remove outliers while simultaneously fitting a homography to inliers, the BaySAC framework is used (Botterill et al. 2009). BaySAC is based on the RANSAC framework (Fischler and Bolles, 1981), but enables matches between multiple similar-looking points to be used efficiently. To compute *H* using RANSAC, many random subsets (`hypothesis sets') of four points are selected. Each hypothesis set is used to generate a homography. The number of correspondences compatible with each homography is counted. When a homography compatible with many correspondences is found, this model is usually correct, and the correspondences found are inliers.

In BaySAC, hypothesis sets are selected based on the prior inlier probabilities of each match (estimated from the number of potential match candidates) and the history of hypothesis sets which have been tried. This enables many low-quality matches between multiple similar features to be used efficiently, without increasing the computational cost of RANSAC.

Once a homography and inlier set are found, the set is refined by using the DLT to fit a new homography to all of the inliers found and re-computing which points are compatible with the new model. The DLT is then used again to fit a homography to all of these points.

The homography, *H*, has the property that

$$H = R + d^{-1}\boldsymbol{t}\boldsymbol{n}^T \qquad (4)$$

where $R$ is the rotation from the previous camera location, $\boldsymbol{t}$ is the camera motion vector, $\boldsymbol{n}$ is a unit vector normal to the ground plane, and, and $d$ is the distance between the camera and the ground. *H* is decomposed to find *R*, $\boldsymbol{n}$ and $\boldsymbol{t}$/d, using Levenberg-Marquardt's algorithm (Hartley and Zisserman, 2003, Appendix 6), $\boldsymbol{t}$ is calculated from the estimated height of the camera above the ground.

Occasionally, when few matches between frames are correct (for example because there are no distinctive features on the ground, or motion is too fast and either frames contain motion blur, or consecutive frames do not overlap), BaySAC will fail to compute a homography compatible with many correspondences. In this case, no update will be made. By only accepting estimates compatible with many correspondences, incorrect measurements are avoided.

## VISION VELOCITY IMU AIDING

As described in the previous section, the computer vision algorithm provides two types of measurements: a translation and rotation. For this work, the rotation information is not used as initial analysis indicates that the accuracy of the rotation from the Homography is lower than the performance of the gyros, not just in terms of noise, but also containing biases. This will be considered in future work. Instead, we focus on using the translation information to restrict the drift of the IMU. To achieve this, we extend the state vector in the Kalman filter to include the camera height. An approximate estimate for the camera height can be obtained by considering the height above ground an average user would hold the device. For this work, we assume 1.2m.

For the camera measurements, we use the following error model:

$$\tilde{v}^b = (1+\delta s)C_c^b v^c + e_v \qquad (5)$$

where $\tilde{v}^b$ is the estimate of the IMU body frame velocity from the camera; $v^c$ is the true velocity in the camera frame; $\delta s$ is the scale factor error for the camera height; $e_v$ is the measurement noise; and $C_c^b$ is the rotation matrix from the camera frame to the body frame. The rotation matrix can be formed, assuming small rotations, using the approximation:

$$C_c^b \approx \begin{bmatrix} 1 & -\omega_z & \omega_y \\ \omega_z & 1 & -\omega_x \\ -\omega_y & \omega_x & 1 \end{bmatrix} \qquad (6)$$

where $\omega_x$, $\omega_y$ and $\omega_z$ are the small rotations between the IMU and camera. The computation of the frame offsets are not considered in this work, although essentially we have a problem similar to boresight calibration in aerial photogrammetry, and hence automated techniques from this field could be applied to solve the problem. Alternatively, the state vector could be extended to solve for these parameters. However, at this stage we assume the offset to be zero and consider the results with respect to this assumption.

Following the derivation in Shin (2005) to use vehicle frame measurements, and assuming no axes offset, we have the following INS error equation:

$$\begin{aligned} \hat{v}^b &= \hat{C}_n^b \hat{v}^n && (7) \\ &\approx C_n^b\left[I + (\delta\phi\times)\right](v^n + \delta v^n) && (8) \\ &\approx v^b + C_n^b \delta v^n - C_n^b (v^n\times)\delta\phi && (9) \end{aligned}$$

where $\hat{v}^b$ is the predicted body frame velocity from the IMU. Therefore the observation equation can be formed as:

$$\begin{aligned} \delta z &= \hat{v}^b - \tilde{v}^b && (10) \\ &= C_n^b \delta v^n - C_n^b (v^n\times)\delta\phi - \delta s(v^b) - e_v && (11) \end{aligned}$$

This observation equation is used to relate the body frame measurements from the IMU and camera to the states that are being estimated in Equation 3, therefore Equation 3 is extended to include the scale factor error term.

## TRIAL

A trial was conducted at the University of Canterbury, New Zealand in March 2010. The trial was conducted with a NovAtel SPAN integrated GPS/INS system using an iMAR FSAS IMU. A good quality IMU was used so that full analysis of the quality of the computer vision measurements could be conducted; however the intention is to use a lower cost IMU in the future. Furthermore, the trial was conducted in a field with clear view of the sky even though the intention will mainly be to use the system for indoor navigation. This was so that an accurate reference trajectory could be generated for testing,

The iMAR FSAS IMU contains three fibre-optic gyros with a drift of less than 0.75deg/hr, and three accelerometers with a bias uncertainty of 1mg. Non-differential code GPS measurements were used to generate the real-time GPS solution which was used as the reference trajectory for analysis. The equipment was configured as shown in Figure 1. The GPS receiver, data

logger and power supply were fixed inside a backpack with a Trimble Zephyr antenna attached to the top. A Canon IXUS 65 consumer camera generating 640x480 pixel video at 30 frames per second was fixed to the IMU with the axes approximately aligned. The small rotational and translational offset of the IMU was not calibrated and the results are considered with this in mind.


**Figure 1 Field trial configuration**

A trajectory was walked around the field with a 6 minute stationary period at the start followed by 17 minutes of walking and 5 minutes stationary at the end. The alignment was achieved using a stationary coarse alignment followed by a fine alignment using GPS position and velocity measurements in a Kalman filter. The University of Nottingham's POINT software was modified to include the error model described in this paper. Time synchronisation was achieved using cross correlation of the z-axis body frame gyro and computer vision estimates since it was not possible to time stamp the images from the camera.

## RESULTS

Figure 2 shows the computed body frame velocity from the integrated GPS/INS data using the POINT software compared to the estimate generated by the computer vision algorithm. Here the figure shows that there is a clear correlation between the two measurements. The computer vision estimates were approximately scaled using the camera height so that they could be compared to the integrated solution. The figure shows that the majority of the motion took place in the y-axis which was approximately pointing forward. Some velocity was also sensed in the x and z axes due to the direction the camera was pointing. In the figure it is clear that there are some noisier computer vision estimates occurring around 700 seconds. This period is where the user walks in an area with a mixture of strong shadow and bright light and the camera produces poor quality images. This issue is discussed further towards the end of the paper. Figure 3 shows a shorter period of data from the computer vision and integrated GPS/INS. In this figure, the 2Hz motion of the IMU is clearly shown as the user takes steps.

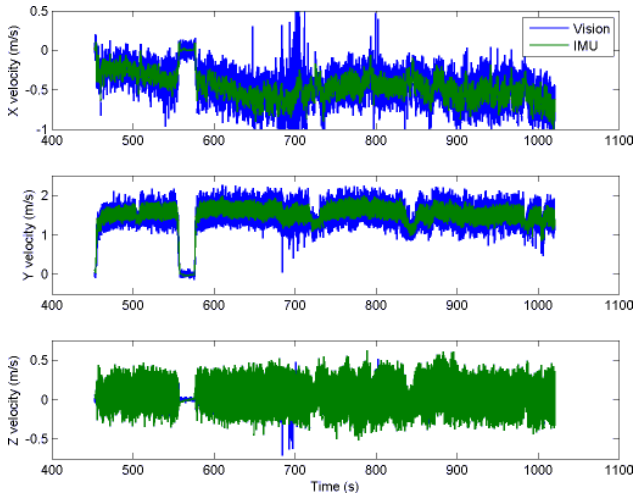


**Figure 2 Comparison of body frame velocity**

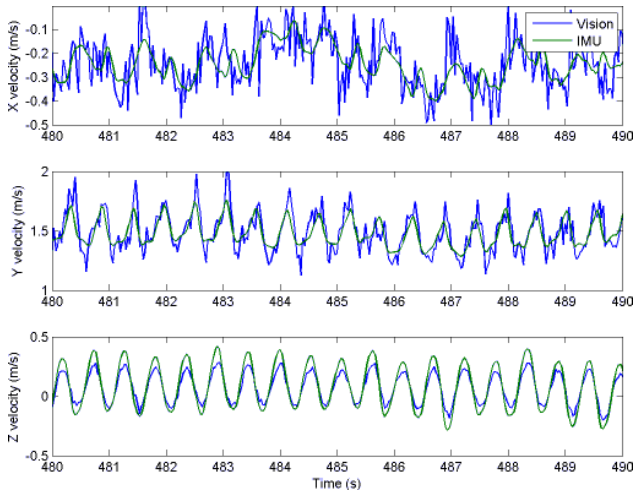


**Figure 3 Comparison of body frame velocity**

Figure 4 shows the integrated GPS/INS solution with a simulated GPS outage lasting 6.5 minutes. After 9.5 minutes, the GPS outage begins and the position solution is generated only from the INS. The position accuracy during the outage is summarised in Table 1. Here it is clear that for the first 60s, the position drift is relatively small with a drift of 4.2m and 2.2m in the North and East axes respectively. This is typical of a good quality IMU although it is also possible that better accuracy would be achieved with a higher dynamic dataset (such as from a vehicle or aircraft) since INS error observability is improved with higher dynamics.

Table 1 also shows that the position error increases rapidly, particular after 240s where the total horizontal error totals 84.5 metres. Again this is a typical characteristic of inertial navigation where the position error increases rapidly over time. This demonstrates that even with a good quality IMU, INS-only navigation is not practical for an indoor application where a user can walk indoors for long periods of time.

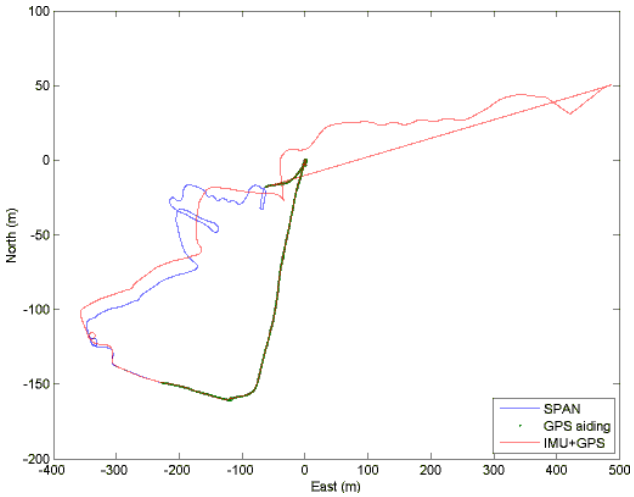


**Figure 4 INS-only trajectory for simulated outage**

| Time (s) | North error (m) | East error (m) |
|---|---|---|
| 60 | -4.2 | 2.2 |
| 120 | -10.8 | -8.5 |
| 180 | -15.2 | 11.8 |
| 240 | -28.3 | 79.6 |
| 300 | -55.5 | 174.4 |
| 360 | -83.5 | 315.7 |

**Table 1 INS-only error during simulated outage**

Figure 5 shows the integrated solution using measurements from the computer vision algorithm along with the GPS position and velocity (with simulated outage) used in Figure 4. The results are summarised in Table 2. Here we demonstrate that the use of computer vision measurements significantly improves the position accuracy when GPS measurements are removed. Table 2 shows that the total position error after 6 minutes with GPS measurements is now only 4.5m and -13.5m in the North and East directions respectively. Although the position error is increasing, the use of computer vision measurements significantly reduces the speed at which the INS drifts. It is also noted that the small misalignments of the camera and IMU have not been modeled which may affect the obtained performance.

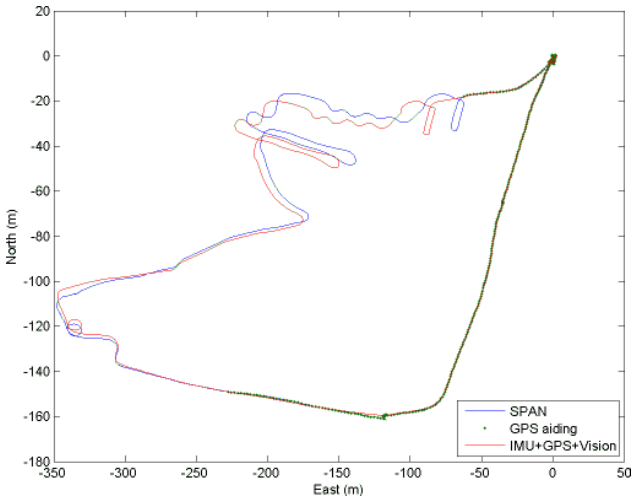


**Figure 5 Trajectory for integrated GPS/INS/Vision solution**

| Time (s) | North error (m) | East error (m) |
|---|---|---|
| 60 | -2.7 | 0.6 |
| 120 | 0.8 | -0.9 |
| 180 | 1.6 | -3.5 |
| 240 | 2.3 | -8.9 |
| 300 | 4.7 | -7.3 |
| 360 | 4.5 | -13.5 |

**Table 2 Integrated INS/Vision error during simulated outage**

Figure 6 shows the estimated scale factor error for the position solution shown in Figure 5. The initial estimate of the scale factor was set to 1.2m to correspond to the approximate height the user was holding the camera. This figure shows that initially an error of 15% is estimated which corresponds to a height error of 18cm. Over time the scale factor estimate slowly reduces to 5% (approximately 6cm). It is thought that this corresponds to the user ending up holding the IMU closer to the ground by the end of the dataset as the IMU is relatively heavy (approximately 2Kg). However, the estimate is very sensitive to the process noise used in the filter. For this test we used a process noise of 0.1mm/s. Detailed analysis of the observability of the scale factor estimate will need to be conducted in future work.

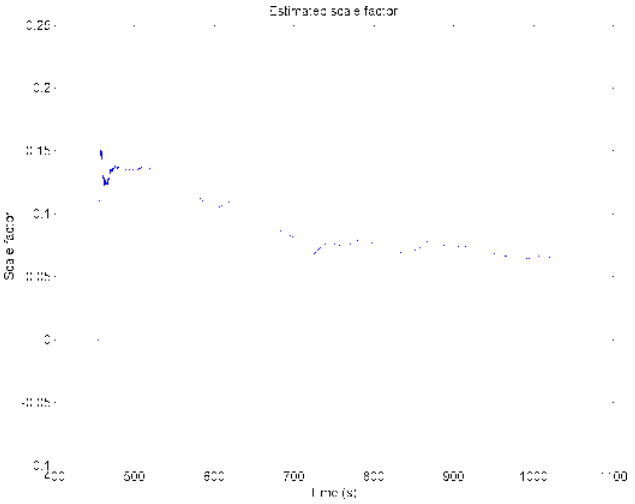


**Figure 6 Estimated camera height error scale factor**

## DISCUSSION

The results from the integrated GPS/INS/Vision system are very promising indicating that the use of computer vision estimates can significantly reduce INS position drift when GPS measurements are unavailable. Using the good quality IMU, the total position drift over the outage was approximately 1 metre per minute. This would be excellent dead reckoning accuracy for integrating measurements from other sensors. In particular a combination of Wi-Fi positioning and map matching approaches seems to offer great potential for an indoor positioning system. However, the IMU used was tactical grade quality and hence further work will be required to investigate how well the vision aiding approach will translate to a much lower grade of IMU.

The main limitations of the approach include computer processor requirements and using the camera in low light conditions. Firstly, for the work shown in this paper, the computer vision algorithm requires approximately 60ms per frame on a modern 3Ghz desktop PC. Although this could potentially work in real-time using 15 frames per second video, it would not be suitable for implementing on a modern smart phone. However, it is thought that with some optimisation of parameters such as the number of features used, it may be possible to significantly reduce this value. Furthermore, the INS position and orientation measurements could be used to immediately remove large outliers if the two systems were to be combined more closely.

Another limitation identified with this approach is that the quality of the images from the camera is significantly reduced when walking in low light conditions such as those typically encountered indoors. In fact, reduced accuracy was demonstrated in this paper in Figure 2 where the camera is located in a high shadow area. Since the camera is typically looking towards the ground, the relative speed of the ground with respect to the camera is fast which results in low image quality. Obvious solutions are to increase the sensitivity of the camera (some consumer cameras offer some control over this); to use a larger lens, or to use artificial lighting. Of course, larger lenses or artificial lighting are not necessarily practical options for mobile devices. Instead, it would be necessary to demand a much more sensitive sensor than is currently available on smart phones, which will result in more noisy images. Although this is a significant current limitation, it is hoped that future technology may enable such an approach.

## CONCLUSIONS

This paper has demonstrated the integration of GPS, IMU and camera measurements as a potential system to enable indoor navigation with a handheld device. It has been shown that using a typical tactical grade IMU, position errors have been reduced from 327m after 6 minutes to only 14m by aiding the INS with estimates of camera translation from a computer vision algorithm. The ground plane homography algorithm has been demonstrated to be robust to data containing a large number of outliers. Furthermore the idea of using the images collected as the user walks with the device held out of in front of them appears to be a realistic proposition.

Future work is required to address a number of issues. These include identifying how well the computer vision aiding measurements will transfer to using a low cost IMU. Secondly, some optimisations will be required in order to make the algorithm more efficient so it can be used on a mobile device. This may include, for example, reducing the number of features used, or taking advantage of closer coupling between the IMU and computer vision algorithm. Finally, the issue of capturing images in low light conditions such as those encountered indoors needs to be examined. However, in summary, the approach shows promise as a future component of improving indoor positioning.

## REFERENCES

Beis, J. S., Lowe, D. G., 1999. Indexing without invariants in 3D object recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence 21-10, p1000-1015.

Botterill, T., Mills, S., Green, R., 2009. New Conditional Sampling Strategies for Speeded-Up RANSAC. In Proceedings of the British Machine Vision Conference.

Foxlin, E., November 2005. Pedestrian tracking with shoemounted inertial sensors. IEEE Computer Graphics and Applications 25, 38–46.

Farrell, J. A., Barth, M., 1999. The Global Positioning System and Inertial Navigation. McGraw-Hill.

Fischler, M., Bolles, R., 1981. Random sample consensus: a paradigm for model fitting with applications to image

analysis and automated cartography. Communications of the ACM 24-6, p381-395.

Hartley, R., Zisserman, A., 2003. Multiple View Geometry in Computer Vision 2nd Edition. Cambridge University Press.

Hide, C., September 2003. Integration of GPS and low cost INS measurements. Ph.D. thesis, University of Nottingham.

Hide, C., Botterill, T., Andreotti, M., 2009. An integrated IMU, GNSS and image recognition sensor for pedestrian navigation. In Proceedings of the Institute of Navigation GNSS Conference. Fort Worth, Texas.

Rosten, E., Porter, R., Drummond, T., 2010. FASTER and better: A machine learning approach to corner detection. IEEE Trans. Pattern Analysis and Machine Intelligence. 32 pp. 105-119.

Shin, E. H., 2005. Estimation Techniques for Low-Cost Inertial Navigation. PhD Thesis, MMSS Research Group, Department of Geomatics Engineering, University of Calgary, Calgary, AB, Canada, UCGE Report No. 20219.

Titterton, D. H., Weston, J. L., 1997. Strapdown Inertial Navigation Technology. Institution of Electrical Engineers.