# Database Verification using an Imaging Sensor

Philippe Simard and Frank P. Ferrie

McGill University, Montreal, Quebec, Canada

## ABSTRACT

In aviation, synthetic vision systems produce artificial views of the world to support navigation and situational awareness in poor visibility conditions. Synthetic images of local terrain are rendered from a database and registered through the aircraft navigation system. Because the database reflects, at best, a nominal state of the environment, it needs to be verified to ensure its consistency with reality. This paper presents a technique for real-time verification of databases using a single imaging device, of any type. It is differential and as such, requires motion of the sensor. The geometric information of the database is used to predict how the sensor image should change. If the measured change is different from the predicted change, the database geometry is assumed to be incorrect. Geometric anomalies are localized and their severity is estimated in absolute terms using a minimization process. The technique is tested against real flight data acquired by an helicopter to verify a database consisting of a digital elevation map. Results show that geometric anomalies can be detected and that their location and importance can be evaluated.

**Keywords:** synthetic vision databases, real-time verification, imaging sensor, geometric anomalies, anomaly detection, anomaly correction

## 1. INTRODUCTION

Synthetic vision systems in aviation render artificial views of the world (from a database and pose information) to support navigation and situational awareness in low visibility conditions. To achieve this goal, such systems require geometrically correct databases of the areas over which the aircraft will be flown. Due to both its static nature and inherent modeling errors, the database introduces anomalies in the synthetic imagery. Consequently, it must be periodically verified in order to, at the very least, detect errors. A sample view of such a database is shown in Figure 1a along with a corresponding picture of the real world (Figure 1b). Note that the overall geometry is quite accurate but that the mountain's shape is slightly incorrect. A good database verification technique should be able to detect, localize and evaluate such an anomaly.

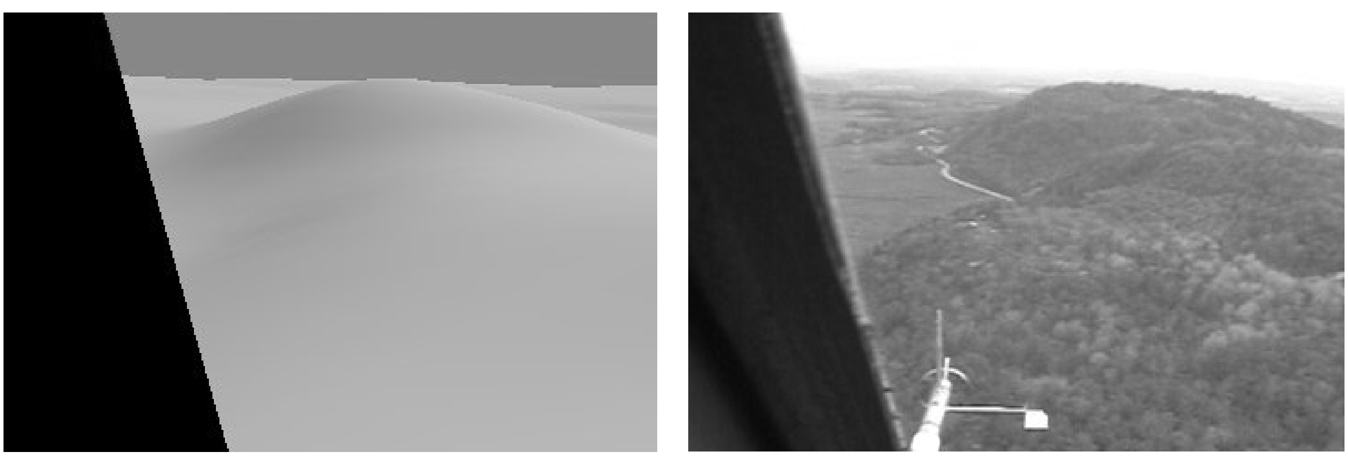

(a) (b)

**Figure 1. (a) Sample view of the database, (b) Corresponding real world view**

This paper presents an approach to estimate both the location and size of geometric anomalies in synthetic vision databases. The approach requires a single imaging device (of any type), as long as its projective geometry is known and is intended to be used real-time. It is also differential and thus, requires motion of the sensor (which is assumed to be known). It uses the geometric information of the database to predict the appearance of the sensor image, after a motion is induced. If the measured image does not match the predicted image, the geometry of the database is assumed to be inconsistent with reality. The algorithm produces a two-dimensional anomaly map that is used to identify database locations where a geometric inconsistency was detected. These locations are corrected using a minimization approach where the number of anomalies in the map is used as a cost function. The basic idea is to modify local geometry so that the predicted image matches the observed image. The correction process results in an updated database that is coherent with sensor observations.

## 2. BACKGROUND

This work was performed as part of a program with CAE Inc., the National Research Council (NRC) of Canada, BAE Systems Canada, the University of Toronto, York University, and the Department of National Defence (Canada) to develop an Enhanced and Synthetic Vision System (ESVS) for low-visibility flight in the Search and Rescue application[12]. The ESVS includes an on-board image generator to render synthetic images of the terrain and an infrared sensor to provide a correlated out-the-window image. Both images are presented on a helmet-mounted-display with the infrared image fused as an inset in the center of the synthetic image field of view[17,18]. Flight symbology is superimposed on the final image to assist the pilot in flying the aircraft. The ESVS display concept is presented in Figure 2.

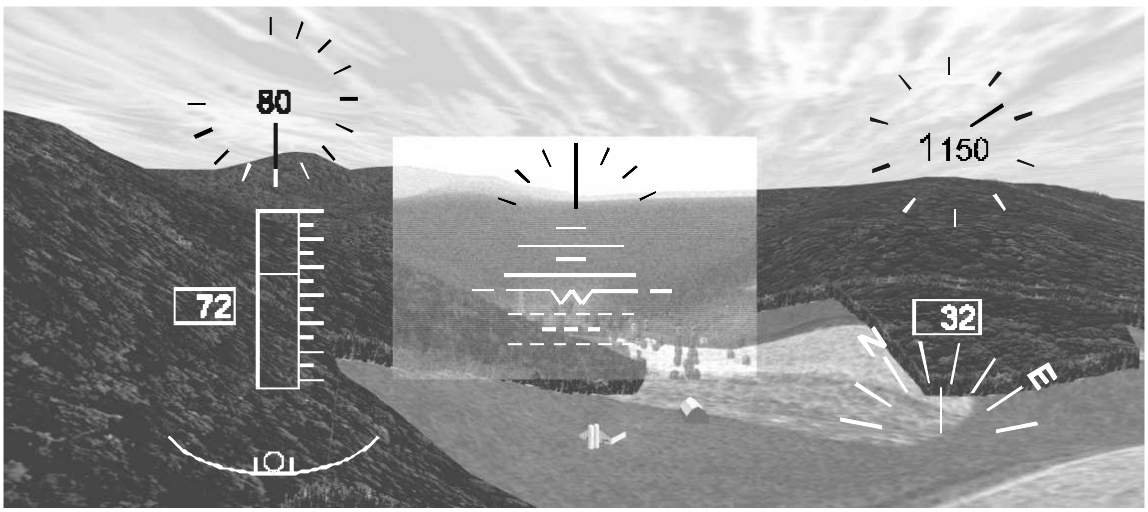


**Figure 2. ESVS concept**

The specific goal of this work was to provide a technique to verify the geometry of an existing database (or 3D model). The verification process basically consists of detecting database errors and of estimating their location and magnitude. The most relevant techniques to this problem are model updating approaches. The main idea of these algorithms is to actually build models in an incremental fashion. Model points are estimated and incrementally refined in a lot of cases using Kalman-based approaches. Examples of such techniques include work by Matthies et al.[16], Kumar et al.[13], Azarbayejani et al.[1] and Beardsley et al.[3].

Although these approaches could probably be used for estimating geometric errors in already existing models, they do not lend themselves well to real-time applications. The main reason is that they tend to update a large number of model points at each iteration. In fact, having a rather accurate model (or database) from the start will require, in most cases, to correct only a small set of points. If incorrect model points could be quickly identified, the verification process could be greatly simplified. The necessary step before the actual correction is therefore to detect geometric errors in the model or in other words, to validate the model.

Model validation algorithms have been proposed for dealing with specific environments or objects. Bejanin et al.[4] and Huertas and Nevatia[11] have developed a technique to confirm the presence of man-made structures (buildings) in aerial images by matching model features to image features. Recent work by Hsu et al.[10] on pose estimation given a scene model (consisting of a set of buildings) and video data has also led to a validation technique. The validated objects are then refined by computing a height map that represents required changes.

In addition, model validation has been used for real-time obstacle detection on a moving vehicle. The basic idea is to assume that the road is a planar surface and to identify anything that violates the assumption as an obstacle. Enkelmann[9] has developed an algorithm based on the idea that the motion field can be predicted, assuming only a translational camera motion parallel to the road plane. To detect obstacles, differences between calculated optical flow and observed flow vectors are evaluated.

Bertozzi et al.[5] have developed a system for obstacle detection in a pair of images acquired by a stereo vision device installed on a moving vehicle. The inverse perspective mapping technique (IPM)[15] is used to remove the perspective effect and to produce a new image representing the texture of the road as it was observed from the top. The IPM is applied on both the left and right images and if the two extracted textures are different, an obstacle is identified. Their method has been generalized in two ways. First, it has been modified to use a single camera[2]. The second extension was the support of non-flat roads[6]. Another obstacle detection technique proposed by Carlsson and Eklundh[7] uses the planar surface assumption to build a predicted sensor image, given known motion of the vehicle. All points in the image that are not on the planar surface will be erroneously predicted and will not match the observed image.

The main contribution of this paper is to provide a database verification framework where geometric anomalies can be quickly detected and corrected. The correction technique is also novel in that it uses anomaly maps as a metric for iterative updating. The proposed validation approach builds on the work by Carlsson and Eklundh[7] but assumes very unconstrained environments and motion. It removes the planar surface assumption, allowing validation of models with arbitrary shapes. It also eliminates the necessity of having model features, a condition that cannot be guaranteed in general.

## 3. THEORY

The anomaly detection technique is divided in two steps. The first one consists of detecting and localizing geometric anomalies in the database. A predicted image is generated and compared against an observed image to produce an anomaly map. The map is then used to identify database locations having incorrect geometry. The second step serves to estimate the severity of anomalies in absolute terms. Geometric changes are iteratively applied to the database to minimize the number of detected anomalies. The final solution provides not only the estimated locations of geometric problems but also an estimate of their importance in absolute terms. The estimated discrepancy between the database and the world can further be used to update the database.

### 3.1 Anomaly Detection

The process of rendering a predicted image mainly consists of computing the projected motion field. Given that the motion of the camera and the scene geometry are both known, the movement of pixels between two images (taken from two different viewpoints) can be exactly predicted. A depth value (computed from the geometric information of the database and current sensor pose) is first associated to each pixel of the sensor image. Given this depth value, the sensor motion vector and projective geometry, the projected motion field can be calculated. This field basically consists of the displacement vectors in the image plane and relates pixels positions in the sensor image before and after the motion. Using the motion field, the predicted image is constructed by warping the sensor image gathered before the motion. The underlying assumption is that the transformation of the image intensity is completely determined by the geometric transformation of the image points (for example, non-Lambertian surfaces and imaging sensor noise violate this assumption). This rendering process is usually referred to as image-based rendering in the literature[8,14].

The predicted image is then compared with the observed sensor image gathered after the induced motion to construct an anomaly map. If the motion field prediction is incorrect and if it generates an intensity difference between the predicted and the observed image, an anomaly is detected. The predicted image is compared with the observed image using a region-based approach. The metric used is the sum-of-absolute differences (SAD) which is computed for each location of the image over a given neighborhood. The anomaly map consists of the set of SAD values and has the same dimensions as the sensor images. The SAD will be large if the two regions are different, meaning that the geometry of the model is incorrect for that particular location. Conversely, if it is close to zero, the geometry is assumed to be correct. To allow for locations where the intensity transformation cannot be exactly predicted by the geometric transformation of image points, the anomaly map is thresholded by keeping only large SAD values. The anomaly detection process is summarized in Figure 3.

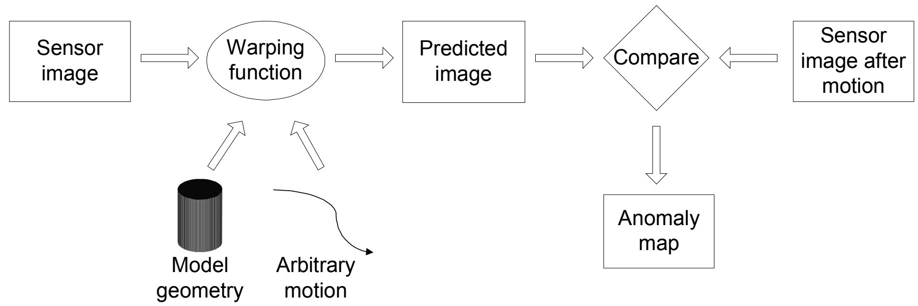


**Figure 3. Anomaly detection process**

Since the anomaly map is expressed in the two-dimensional imaging space, it can be directly used for displaying geometric anomalies. For example, the map can be superimposed onto its corresponding synthetic image to highlight regions where the database geometry is incorrect. However, it might be preferable to use the map to mark anomalous database locations. This can be easily achieved since there is a direct association between detected anomalies in the map and three-dimensional locations in the database (remember that the database geometry was used to compute the motion vectors necessary for the predicted image generation). Information about the correctness of the geometry can thus be directly included in the database. Mathematical details on the anomaly detection algorithm can be found elsewhere[19].

### 3.2 Estimation of the Severity of Anomalies

Unfortunately, a large SAD value does not necessarily imply a large discrepancy in geometry between the database and the world. Instead, it means that the geometric anomaly generated a large difference in intensity between the predicted and the observed image. As a consequence, an intensity value in the anomaly map cannot be directly translated into an estimate of the error in the database. To estimate the severity of anomalies in absolute terms, the database geometry is corrected such that the prediction matches the observation. The correction problem is cast as a minimization process where the anomaly map is used to define a cost function. The idea is to refine incorrect database locations in order to minimize the number of detected anomalies. The final solution consists of the database corrections leading to a minimum number of anomalous points in the anomaly map.

Incorrect database locations are first identified using the detection algorithm presented in the previous section. The minimization itself is performed using a gradient descent approach. Database locations are iteratively adjusted in parallel so as to reduce the number of anomalies in the map. At each iteration, the direction of the gradient (or the update vector consisting of database corrections) is evaluated by systematically varying the geometry of incorrect database locations. The database is then updated using the computed vector. This process is repeated until convergence, which occurs given the following assumptions. Because database errors are usually bounded, the starting solution should be within the same convexity region as the desired minimum. Also, the cost function must be sensitive to small perturbations of the database geometry. This occurs when sensor images have a high intensity gradient.

## 4. METHODOLOGY

The algorithm was used to verify a database consisting of a digital elevation map (DEM) provided by the Canadian Government. The DEM has grid posts every 93m in latitude and 65m in longitude and known elevation errors of up to 25m. A standard digital video camera having a field-of-view of $45^{o}$x$30^{o}$ with a resolution of 720x480 was mounted inside a Bell 205 helicopter (Figure 4). The pilot flew at relatively low altitudes over a forested terrain in the Gatineau Hills region near Ottawa (Canada) according to a predetermined flight path. A differential GPS and inertial sensors (with 0.5m accuracy and less than $0.5^{o}$ rotation errors respectively) were used to log the position/attitude information of the aircraft.


(a)


(b)

**Figure 4. (a) Helicopter used for data collection, (b) Camera attached to window frame**

The depth values necessary for computing the motion field were generated by rendering synthetic images of the model using an OpenGL rendering package and by extracting the depth buffer. In order to achieve this, the correspondence between the camera reference frame and the world reference frame (in which the model is expressed) had to be established. The intrinsic parameters of the camera were determined by a standard calibration procedure. The location and orientation of the camera relative to the GPS receiver and the attitude sensors were also measured. As the position calculated by the GPS was expressed in the same coordinate system as the DEM (i.e. latitude/longitude/altitude), no additional processing was necessary to register them. Initial alignment in attitude was performed manually: the yaw component was aligned by using the positional data given by the GPS (the angle of the flight path was computed and compared with the measured yaw) and the roll/pitch using the horizon (present in both the synthetic and sensor imagery) as a landmark. This led to a transform relating the two reference frames. The precision achieved when generating an image of the database is about one pixel.

A simple interpolation scheme was implemented for the predicted image generation (to cope with, for example, image expansions due to forward motions). Also, the SAD used for the generation of the anomaly map was computed using 3x3 neighborhoods. Although the data was processed offline to validate and correct the model geometry, the required computations could be done in real-time onboard using a standard PC and optimized software.

## 5. RESULTS AND DISCUSSION

The verification algorithm was first tested against a flyby scenario near the Ottawa airport. In that region, the model only consisted of a flat terrain (Figure 5a). The camera images before and after a combination of a forward and right motion of about 100m in each direction are shown in Figure 5b and 5c respectively (note that the left portion of the image is obscured by the window mounting and that the white object in the lower/center region is the air data boom attached to the aircraft). The predicted image (Figure 5d) was constructed by warping the camera image acquired before the motion according to the motion field prediction, which again is governed by the geometric information of the database and the measured motion. Note that in regions where the world was in fact flat, the appearance of the features was correctly predicted. However, the appearance of the building (which is missing from the model) is tilted and taller.

The corresponding anomaly map is shown in Figure 5e. Note that locations where the algorithm has detected geometric anomalies are shown in light tones. The building was correctly identified as missing from the model. The map was used to identify incorrect database locations, which are shown in white tones in Figure 5f. To deal with the mismatch between the image and model resolutions, a nearest neighbor approach was used to mark invalid database locations. These points were then refined such that the prediction matches the observation, thus minimizing the number of detected anomalies. Finally, the refined depth values were used to update the original DEM.

The resulting updated prediction and the associated anomaly map are shown in Figure 5g and 5h respectively. Note that the general appearance of the building is much better than the original prediction (the dark region to the right of the building corresponds to the disocclusion produced by the added structure). Also observe that the number of detected anomalies has been greatly reduced. The corresponding updated model is presented in Figure 5i. Note that a pyramid shaped structure is now present in the database, which roughly corresponds to the missing structure (in position and size). Given the relatively low resolution of the database, the solution represents the best approximation to the missing cylindrical structure.

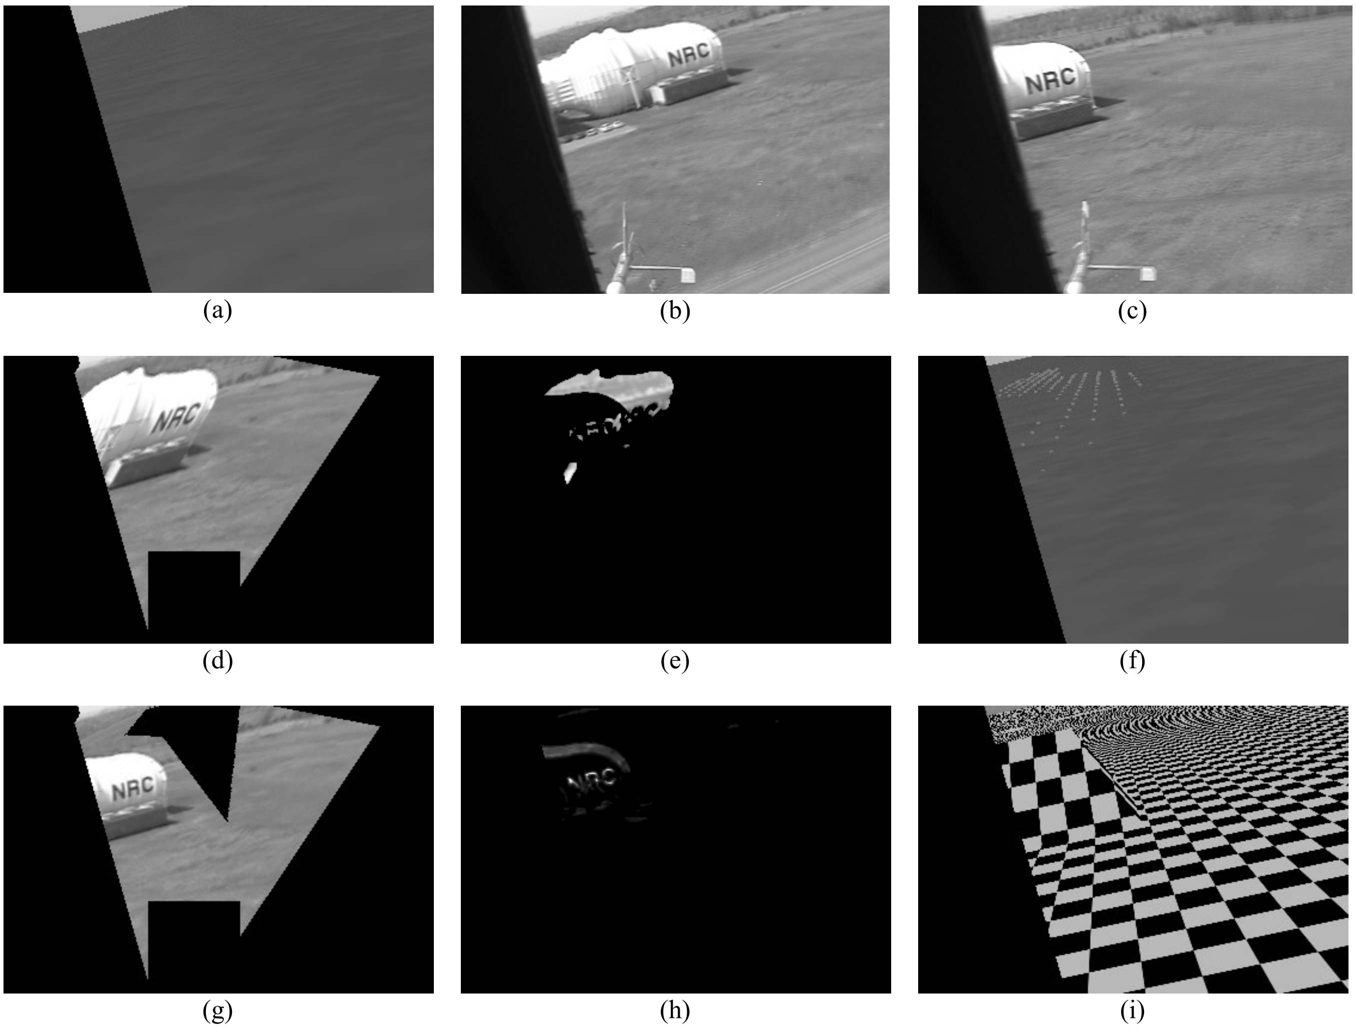


**Figure 5. (a) Database view, (b) Camera image, (c) Camera image after forward/right motion, (d) Predicted image, (e) Anomaly map, (f) Incorrect database locations, (g) Updated prediction, (h) Updated anomaly map, (i) Updated database**

A second test was performed during an approach to a mountain, which in the model (Figure 6a) has slightly incorrect geometry. The camera image before and after a forward motion of 250m are shown in Figure 6b and 6c respectively. The corresponding predicted image is presented in Figure 6d. Observe that the general appearance of the scene is correct. However, the anomaly map (Figure 6e) reveals that the mountain appearance is largely correct but that the right part of the ridge contains some errors (Figure 6e, 6f, 6h and 6i are close-up views of the region indicated by the dashed box in Figure 6c).

Using the map, erroneous database locations were identified and are represented in dark tones in Figure 6f. The model was refined using the minimization algorithm. The resulting updated prediction, updated anomaly map and corrected model are shown in Figure 6g, 6h and 6i. The improvement to the geometry can be observed by comparing Figure 6f and 6i. Indeed, this also shows up in the map since the number of detected anomalies has been reduced.

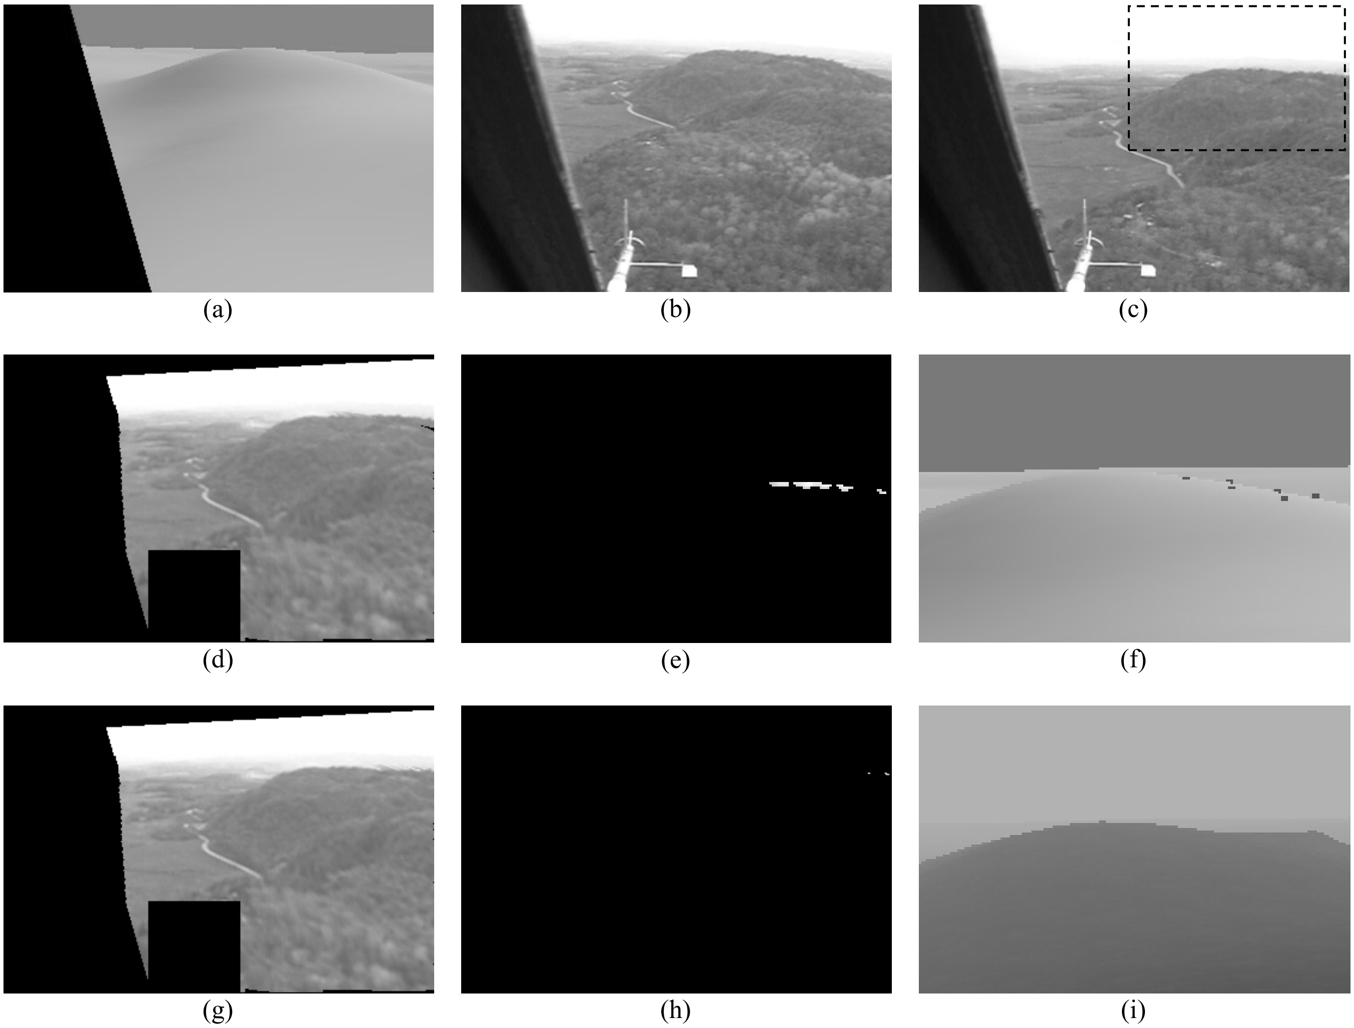


**Figure 6. (a) Database view, (b) Camera image, (c) Camera image after forward motion, (d) Predicted image, (e) Anomaly map, (f) Incorrect database locations, (g) Updated prediction, (h) Updated anomaly map, (i) Updated database**

## 6. CONCLUSIONS

A differential approach was presented to verify the geometry of databases used in synthetic vision systems. The technique requires a single imaging device and differs from traditional approaches by quickly identifying geometric anomalies. This significantly reduces the number of locations that require correction, making the technique applicable to real-time applications. The algorithm is sensor independent and can deal with databases of arbitrary shapes or scale. The verification technique produces an anomaly map which can be used for the identification of incorrect database locations and as a cost function for database correction.

A standard camera mounted on a helicopter was used to verify a database consisting of a terrain. Results demonstrated that geometric anomalies can be detected/estimated and even automatically fixed. The algorithm was able to detect and correct anomalies of different type, size and scale. In particular, missing man-made structures as well as terrain geometric anomalies were correctly handled by the technique.

Future work include the development of mechanisms to better cope with, for example, non-Lambertian surfaces such as lakes. Techniques for adaptively changing the model resolution (in cases where the model cannot account for the observations) will also be investigated.

## 7. ACKNOWLEDGMENTS

The authors would like to thank Sion Jennings from the Flight Research Laboratory of the National Research Council for providing real flight data and Ronald V. Kruk from CAE Inc. for his help in planning the data acquisition.

## 8. REFERENCES

1. Azarbayejani, A., Horowitz, B. and Pentland, A., "Recursive Estimation of Structure and Motion using Relative Orientation Constraints", *Conference on Computer Vision and Pattern Recognition*, pp. 294-299, 1993.
2. Batavia, P.H., Pomerleau, D.A. and Thorpe, C.E., "Overtaking Vehicle Detection Using Implicit Optical Flow", *IEEE Intelligent Transportation Systems Conference*, pp. 729-734, 1997.
3. Beardsley, P.A., Zisserman, A. and Murray, W., "Sequential Updating of Projective and Affine Structure from Motion", *International Journal of Computer Vision*, pp. 235-259, 1997.
4. Bejanin, M., Huertas, A., Medioni G. and Nevatia R., "Model Validation for Change Detection", *IEEE Workshop on Applications of Computer Vision*, pp. 160-167, 1994.
5. Bertozzi, M., Broggi, A. and Fascioli, A., "A Stereo Vision System for Real-Time Automotive Obstacle Detection", *IEEE International Conference on Image Processing*, pp. 681-684, 1996.
6. Bertozzi, M., Broggi, A. and Fascioli, A., "An Extension to the Inverse Perspective Mapping to Handle Non-Flat Roads", *IEEE Intelligent Vehicules Symposium*, pp. 305-310, 1998.
7. Carlsson, S. and Eklundh, J., "Object detection using model based prediction and motion parallax", *European Conference on Computer Vision*, pp. 297-306, 1990.
8. Chen, S.E. and Williams, L., "View Interpolation for Image Synthesis", *SIGGRAPH '93*, 1993, pp.279-288.
9. Enkelmann, W., "Obstacle Detection by Evaluation of Optical Flow Fields from Image Sequences", *European Conference on Computer Vision*, pp. 134-138, 1990.
10. Hsu, S., Samarasekera, S., Kumar, R. and Sawhney, H.S, "Pose estimation, model refinement, and enhanced visualization using video", *Conference on Computer Vision and Pattern Recognition*, pp. 488-495, 2000.
11. Huertas A. and Nevatia R., "Detecting changes in aerial views of man-made structures", *International Conference on Computer Vision*, pp. 73-80, 1998.
12. Kruk, R.V., Link, N.K., MacKay, W.J., Jennings, S., "Enhanced and Synthetic Vision System for Helicopter Search and Rescue Mission Support", *Proc. American Helicopter Society 55th Annual Forum*, Montreal, Quebec (Canada), 1999.
13. Kumar, R., Sawhney, H.S. and Hanson, A.R., "3D Model Acquisition from Monocular Image Sequences", *Conference on Computer Vision and Pattern Recognition*, pp. 209-215, 1992.
14. Laveau, S. and Faugeras, O., "3-D Scene Representation as a Collection of Images", *International Conference on Pattern Recognition*, pp.689-691, 1994.
15. Mallot, H.A., Bülthoff, H.H., Little, J.J. and Bohrer, S., "Inverse Perspective Mapping Simplifies Optical Flow Computation and Obstacle Detection", *Biological Cybernetics*, vol. 64, pp. 177-185, 1991.
16. Matthies, L., Kanade T. and Szeliski R., "Kalman Filter-based Algorithms for Estimating Depth from Image Sequences", *International Journal of Computer Vision*, 3, pp. 209-238, 1989.
17. Simard, P., Link, N.K., Kruk, R.V., "Feature detection performance with fused synthetic and infrared imagery", *Proc. Human Factors and Ergonomics Society 43rd Annual Meeting*, pp. 1108-1112, 1999.
18. Simard, P., Link, N.K., Kruk, R.V., "Evaluation of algorithms for fusing infrared and synthetic imagery", *Proc. of the SPIE Conference on Enhanced and Synthetic Vision*, vol. 4023, pp. 127-138, 2000.
19. Simard, P., Ferrie, F.P., "Online database updating by change detection", *Proceedings of the SPIE Conference on Enhanced and Synthetic Vision*, vol. SPIE-4363, pp. 103-111, 2001.