# Evaluation of Algorithms for Fusing Infrared and Synthetic Imagery

Philippe Simard[a], Norah K. Link[b] and Ronald V. Kruk[b]

[a]McGill University, Montreal, Quebec, Canada
[b]CAE Electronics Ltd., St-Laurent, Quebec, Canada

## ABSTRACT

Algorithms for image fusion were evaluated as part of the development of an airborne Enhanced/Synthetic Vision System (ESVS) for helicopter Search and Rescue operations. The ESVS will be displayed on a high-resolution, wide field-of-view helmet-mounted display (HMD). The HMD full field-of-view (FOV) will consist of a synthetic image to support navigation and situational awareness, and an infrared image inset will be fused into the center of the FOV to provide real-world feedback and support flight operations at low altitudes. Three fusion algorithms were selected for evaluation against the ESVS requirements. In particular, algorithms were modified and tested against the unique problem of presenting a useful fusion of information from high quality synthetic images with questionable real-world correlation and highly correlated sensor images of varying quality. A pixel averaging algorithm was selected as the simplest way to fuse two different sources of imagery. Two other algorithms, originally developed for real-time fusion of low-light visible images with infrared images, (one at the TNO Human Factors Institute and the other at the MIT Lincoln Laboratory) were adapted and implemented. To evaluate the algorithms' performance, artificially generated infrared images were fused with synthetic images and viewed in a sequence corresponding to a search and rescue scenario for a descent to hover. Application of all three fusion algorithms improved the raw infrared image, but the MIT-based algorithm generated some undesirable effects such as contrast reversals. This algorithm was also computationally intensive and relatively difficult to tune. The pixel averaging algorithm was simplest in terms of per-pixel operations and provided good results. The TNO-based algorithm was superior in that while it was slightly more complex than pixel averaging, it demonstrated similar results, was more flexible, and had the advantage of predictably preserving certain synthetic features which could be used support obstacle detection.

**Keywords:** image fusion, enhanced and synthetic vision, infrared imagery, synthetic imagery

## 1. INTRODUCTION

The present study is part of a Canadian Forces Search And Rescue (CF SAR) Technology Demonstrator project to provide an Enhanced and Synthetic Vision System (ESVS) to SAR helicopter crews in poor visibility conditions. The ESVS includes an on-board Data Base and Computer Image Generator to generate a synthetic image of local terrain, registered through the aircraft navigation system, and a near visual wavelength (Infra Red) sensor to provide a correlated out-the-window image. Both images are presented on a Helmet-Mounted-Display with the IR image fused as an inset in the center of the synthetic image field of view.

The IR sensor typically has small field of view or low resolution, it is subject to degradation due to weather effects, and it can be quite noisy. Synthetic images can be generated with both large field of view and high resolution. However, they will suffer real-world correlation problems due to the resolution of the polygonal representation and the resolution and accuracy of available source data. Objects may be displaced or not represented at all because of modeling compromises or because they were not present in the source data.

Previous work has been done in fusing low-light visible CCD and IR imagery, which is a priori very similar to our problem. Algorithms developed to address the CCD-IR fusion problem were therefore considered strong candidates for our evaluation. These algorithms, which may involve only operations on corresponding pixels or be based on pyramid representations, share the ESVS goal of trying to preserve the highest information content from the two image sources.

There are, however, unique problems in fusing synthetic and IR imagery for an ESVS. Although low-light visible CCD is somewhat close to synthetic imagery, there are differences in the sense that the ESVS synthetic imagery is not affected by weather conditions and therefore has consistently high contrast and sharp edges. For an ESVS however, sensor imagery should be privileged when it shows good contrast because it is the modality that best represents the outside world.

## 2. ESVS CONCEPT AND ARCHITECTURE

### 2.1 ESVS Concept

The ESVS provides a means to the pilot to perceive the necessary visual information from three different image sources: a synthetic computer generated image, an enhanced image from an electro-optical sensor and aircraft instrument symbology. The ESVS display concept is presented in Figure 1.

The synthetic image provides the pilot with a wide field-of-view terrain display that can be used to maintain a global sense of position and orientation and to navigate en-route by recognising landmarks that would otherwise be hidden by poor visibility. It is generated in real time from aircraft navigation system data and the pilot's viewpoint based on head orientation. The terrain database will be augmented by obstacles detected by a weather-penetrating active range sensor and obstacle detection system.

The enhanced image is a high confidence sensor image that can be used during manoeuvring close to the ground. It occupies a smaller field-of-view to enable a high-resolution output. This is required to produce high-definition terrain and obstacles, necessary to support close manoeuvring. As the pilot's head turns, the head tracker reads the pilot's head position and sends a signal to the sensor platform telling it to follow the pilot's head movement.

The central region of the displayed image comprises the enhanced sensor image fused with the synthetic image. The image fusion process combines the best aspects of each image, so that the sensor image is automatically given more emphasis close to terrain and obstacles, and as the image quality improves with weather penetration. This accommodates problems of accuracy and missing data in the synthetic terrain database during critical phases of Search and Rescue missions.

Flight symbology can be superimposed on the final image to assist the pilot in flying the aircraft. The flight symbology provides the necessary primary flight references such as airspeed, altitude, attitude, power, a turn co-ordinator, and a compass.

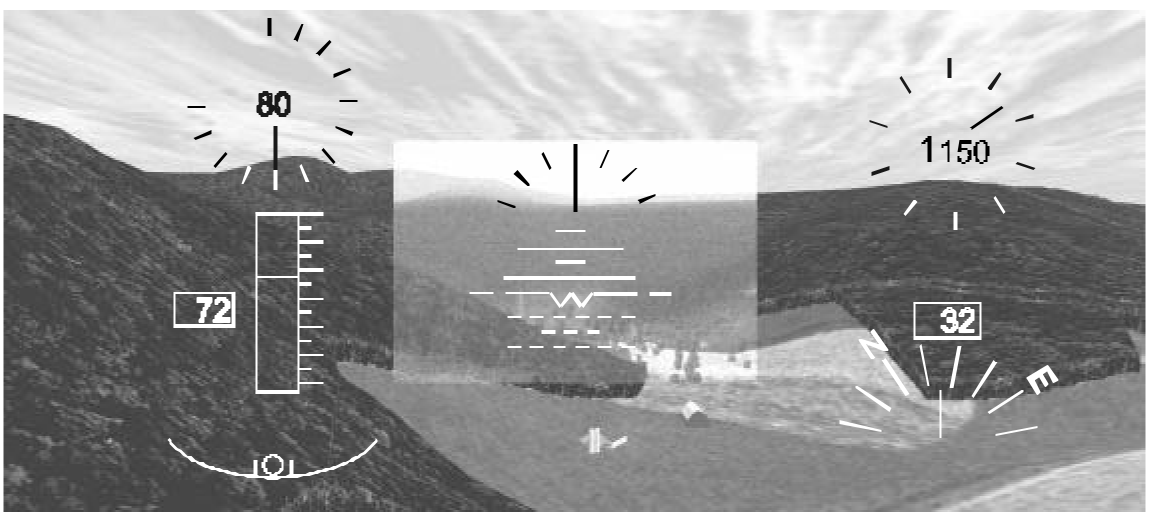


**Figure 1. ESVS concept**

### 2.2 ESVS Architecture

The ESVS architecture (Figure 2) is based to the maximum extent possible on using low-cost Commercial Off-The-Shelf (COTS) hardware. The enhanced vision system component is composed of two subsystems: an infrared camera mounted on the aircraft camera platform and a Digital Image Processing Unit (DIPU).

The synthetic vision component is comprised of two subsystems: a PC-based real-time image generator, and an object detection system based on data received from an active ranging system. The synthetic image viewpoint will be located in the database according to inputs received from the navigation system. The object detection system will incorporate a module to interpret range data, detect significant obstacles, and incorporate them into the synthetic terrain database. It will also be used to correct database inaccuracies due to erroneous or incomplete source data and, where possible, to maintain alignment of the database to the real world despite navigational system drift.

The ESVS will provide image fusion between the electro-optic sensor of the enhanced vision system and the output of the synthetic vision system. The fusion software will be implemented within the DIPU. More details on the system architecture are in a paper by Kruk et al[1].

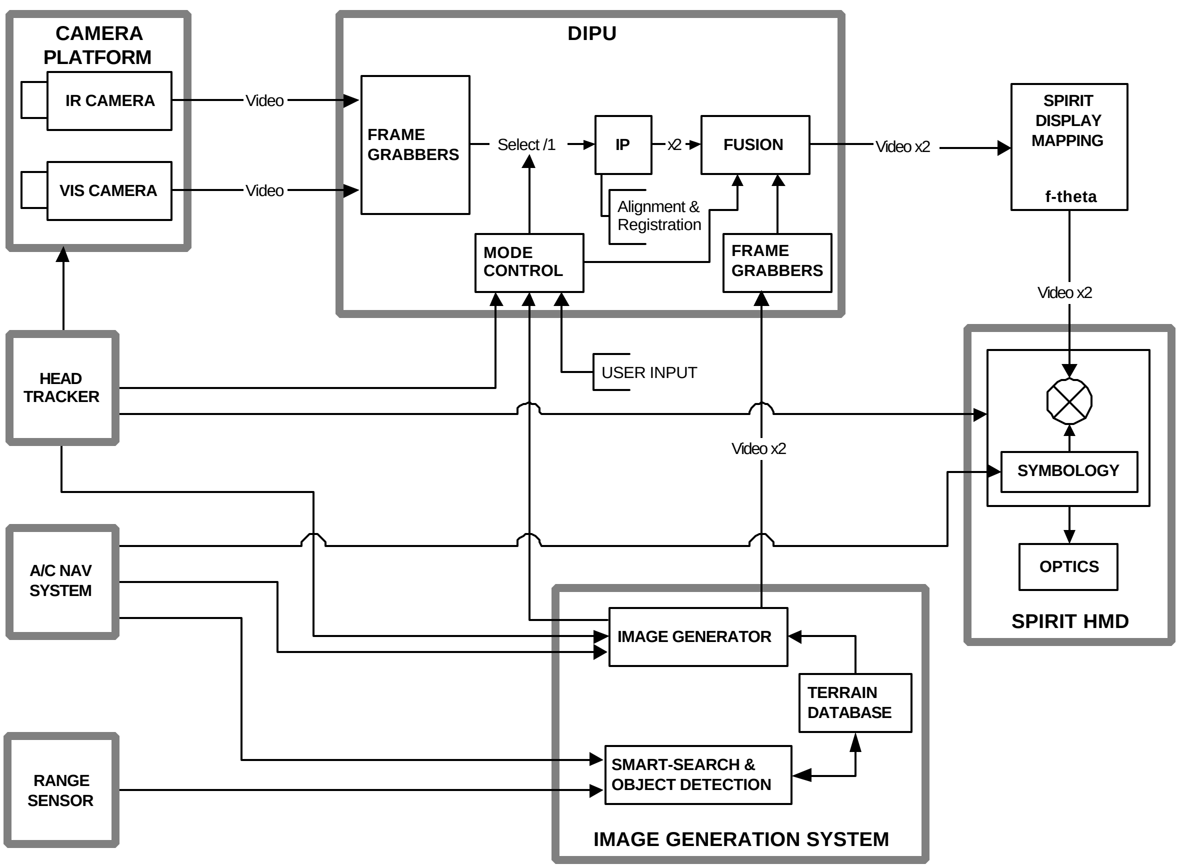


**Figure 2. ESVS system architecture**

### 2.3 Supporting Studies

A comprehensive research program has been developed to understand the limitations of the demonstration ESVS hardware, to promote a fundamental understanding of human-machine cockpit interfaces, and to gain experience in the integration of sophisticated cockpit technology. Without this knowledge, the system designers are limited in their ability to predict the benefits and costs of ESVS use in an operational environment. Designers also run the risk of inadvertently increasing workload and reducing situational awareness if they fail to account for human performance limitations. The development effort is concentrated at four facilities: CAE's Research and Development facilities in St. Laurent, Quebec; the CMC Systems Integration Facility in Kanata, Ontario; the Full Flight Simulator (FFS) at the University of Toronto Institute for Aerospace Studies; and at the National Research Council's Flight Research Laboratory (NRC FRL) in Ottawa, Ontario. Supporting research has also been conducted at Carlton, York and McGill Universities, Defense Research Establishment Valcartier, and the Centre for Research in Earth and Space Technology (CRESTech).

Supporting research for ESVS has examined issues of pilot performance against parameters such as field-of-view, design eye, roll compensation in the head-tracked sensor platform, system (temporal) delays, scene content, and navigation data stability[1]. Mission simulations have also been used to compare results from simulation trials with aircraft trials to validate the use of the simulator as an experimental device. The current study was developed to assess image fusion algorithms for ESVS.

## 3. IMAGE FUSION FOR ESVS: PERFORMANCE EVALUATION

An experiment was designed to evaluate the performance of the three different algorithms within the ESVS for Search and Rescue context. A matrix of test conditions was constructed which reflects issues unique to fusing sensor and synthetic images. The infrared sensor was selected as the most likely candidate to be used in these missions. Image sequences were generated using a generic synthetic database typical of the type of terrain encountered during such missions in Canada. Both synthetic and sensor images were simulated using CAE proprietary technology to provide control over the test conditions. Evaluation of the algorithm performance was based on the observers' ability to detect key features critical to the flight task under the various conditions.

### 3.1 Matrix of Test Conditions

#### 3.1.1 Image Fusion Issues

The main issues that needed to be considered for the fusion of infrared and synthetic imagery were:

1. Performance under different sensor conditions (e.g. black hot / white hot).
2. Performance with varying quality of the sensor image due to weather penetration capability and sensor noise characteristics.
3. The effect of synthetic and enhanced sensor scene content mismatches, arising from typical synthetic image inaccuracies (due to source data errors, missing source data, modelling inaccuracies, and navigational drift), on the usability of the central fused region.

The key issues that were evaluated in this study are further described in the following sections.

#### 3.1.2 Sensor Conditions

Sensor modality (white hot or black hot) and weather (visibility) conditions were varied to produce six conditions of infrared sensor image quality to test different fusion situations. These are listed in Table 1. The three visibility ranges correspond to high humidity percentages of above 90%.

| | Sensor White Hot | Sensor Black Hot |
|---|---|---|
| High Visibility – 3 nm | Sensor Condition 1 | Sensor Condition 2 |
| Medium Visibility – 1.5 nm | Sensor Condition 3 | Sensor Condition 4 |
| Low Visibility – 0.5 nm | Sensor Condition 5 | Sensor Condition 6 |

**Table 1. Sensor conditions**

#### 3.1.3 Registration Conditions

Inconsistencies can be expected between the content of synthetic images generated from a stored database and the IR images in an ESVS. These inconsistencies are referred to as registration problems. By affecting the relative geometry, or registration, of the sensor and synthetic images, these conditions would affect the fusion process. Nine different registration conditions were defined to study separately the different effects that misregistration between the two image sources would have on fusion algorithms. The conditions are listed in Table 2.

| | |
|---|---|
| **Registration Condition 1** | Control: Identical database and viewpoint to sensor image (40 m terrain elevation posts). |
| **Registration Condition 2** | Synthetic viewpoint inaccuracies introduced on flight path (± 15m). |
| **Registration Condition 3** | Missing objects, object position offsets. |
| **Registration Condition 4** | Global database offset (extreme - 1°, 5 m). |
| **Registration Condition 5** | Decreased terrain resolution (mid - 160 m). |
| **Registration Condition 6** | Local terrain elevation errors (extreme - 45m). |
| **Registration Condition 7** | Global database offset (mid - 0.5°, 3 m). |
| **Registration Condition 8** | Decreased terrain resolution (extreme - 320 m). |
| **Registration Condition 9** | Local terrain elevation errors (mid - 15 m). |

**Table 2. Synthetic scene registration conditions**

### 3.2 Experimental Setup

#### 3.2.1 Apparatus / Image Generation

Twenty Pentium II - 266 MHz PC's, running 24 hours per day for 20 days were used to generate the 240 source images (both sensor and synthetic) and 2592 fused images used in the experiment, as well as over 3000 additional source images used to generate dynamic sequences for further evaluation.

#### 3.2.2 IR Sensor and Synthetic Image Configuration

Figure 3 shows the image configuration. This configuration was designed to register the pixels of the sensor and synthetic images automatically. It also provided an opportunity to assess boundary conditions between the inset fused image and the background synthetic image. The IR images were generated as a post-process to image generation using standard sensor controls and responses as developed at CAE Electronics Ltd. This includes random noise generation, misalignment of the sensor elements, image filtering (noise removal) and brightness/contrast and black hot/white hot controls. Objects in the database were color tuned to mimic their thermal signature for a fixed time of day and time of year (around 5 p.m. on an early fall day, ambient temperature approximately 10°C). Finally, scene fading due to atmospheric effects was simulated.

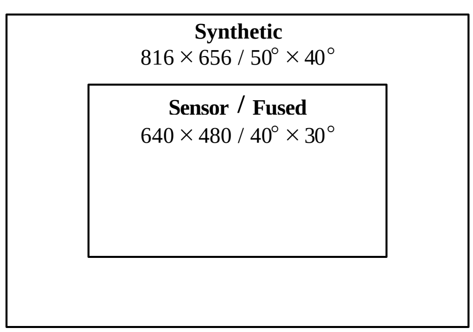


**Figure 3. Image configuration**

3.2.3 Flight Path

A flight path was modeled to simulate a typical SAR approach up a box canyon into a crash site in hilly, forested terrain. It consisted of an approach descending from 500 ft to 30 ft over rising terrain (see Figure 4). The crash site was located on a hillside just below a saddle ridge.

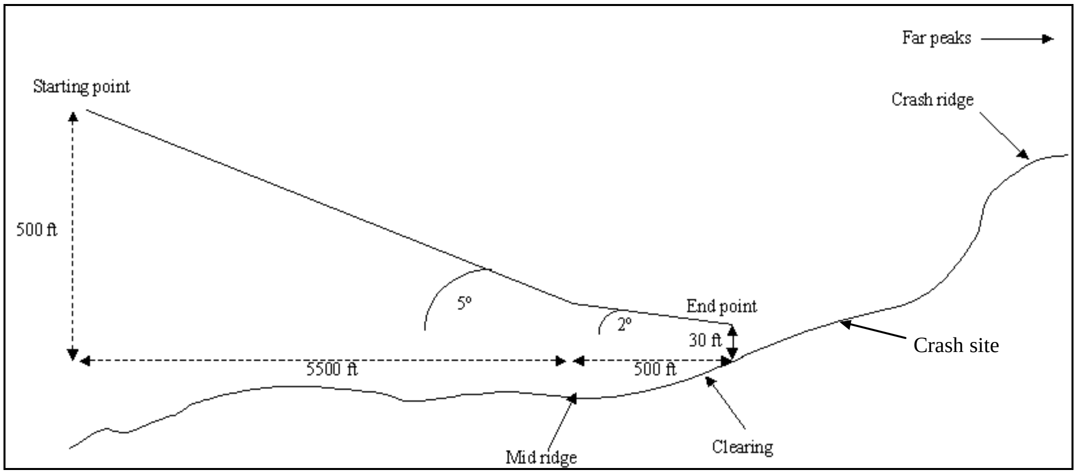


**Figure 4. Flight path and terrain profiles**

3.2.4 Evaluation

Two types of evaluation were used to study image fusion: a static and a dynamic evaluation. Both shared the same goal of determining if the fusion of synthetic and infrared imagery would improve information content. Still images were mainly used to estimate the impact of sensor and registration conditions on fusion. Also, 30 Hz sequences were generated to evaluate the dynamic performance of the algorithms. These two types of evaluation are referred to as static and dynamic respectively.

For the static evaluation, three experienced psychophysical observers were instructed to perform a target detection task which consisted of assessing the visibility of given features and in verifying if there were conflicts between objects or terrain features due to registration problems. The features (located in Figure 4) were: the far peaks (required for general route planning); a mid ridge rising to the left of the flight path (terrain obstacle on approach); the clearing and clearing obstacles; the ridge behind the crash site (also a terrain obstacle and required for route planning); and the crash site itself (visible only in the sensor image). Subjects were required to view 16 images along the flight path and to note at which image the features first became visible and which images contained terrain or object conflicts.

## 4. FUSION ALGORITHMS

Three fusion algorithms of different complexity were investigated and modified to fit the ESVS requirements. In this context, a good fusion algorithm should include as much infrared sensor information as possible because it more closely represents the real world. The synthetic image should therefore be used when the sensor content is very low, and to represent the scene outside of the sensor coverage. Algorithms should respond automatically to the sensor content, and also must be applicable to both white hot and black hot sensor modalities.

The pixel averaging algorithm was selected for study as the simplest way to fuse two images. The two other algorithms implemented for evaluation were based on the TNO algorithm (developed in The Netherlands at the TNO Human Factors Institute by Toet et al.[3]) and the MIT algorithm (developed at MIT Lincoln Labs by Waxman et al.[4-7]). The latter two algorithms were originally developed for real-time (or near real-time) fusion of low-light visible images with IR images. Both were derived from biological models of fusion of visible light and infrared radiation, the TNO algorithm being a simple approximation of the MIT algorithm.

More complex algorithms based on pyramid representations were also considered (but not implemented). Although such approaches may produce good results, it was decided that these algorithms might introduce processing latencies undesireable in a real-time system such as ESVS.

### 4.1 Pixel Averaging Algorithm

The pixel averaging algorithm assigns the average of each corresponding pixel from the infrared sensor image and the synthetic image to the fused image. However, pure pixel averaging presents major problems when fusing sensor and synthetic imagery. An equal contribution of both images usually leads to one of the two following situations: 1) A poor quality (low content) sensor image obscures synthetic data without adding any value to the image; 2) A good quality sensor image is obscured by uncorrelated synthetic data.

For that reason, the algorithm was modified to take into account the quality of the sensor image. The simple average was transformed to a weighted average that provides more weight to high-quality sensor (i.e. high content/contrast) images and less weight to low content images. The problem then becomes one of being able to measure the quality of the sensor image. Different metrics to estimate the quality of the sensor images were considered in both the spatial and frequency domains. However, operations in the frequency domain were not investigated because their complexity could represent an obstacle to real-time performance. The content of the sensor image is therefore evaluated by calculating the intensity standard deviation, under the assumption that greater deviations result from more scene content present in the sensor image while lower deviations result from sensor images obscured by atmospheric effects or noise. As a result, the relative weight applied to a sensor image in the average is directly proportional to its intensity standard deviation.

While this scheme produces satisfactory results, sensor images can have interesting regions, i.e. regions of high scene content, while other regions are obscured. Such a sensor image (taken from the experiment, simulated with 0.5nm visibility) is shown in Figure 5. Note that the top part of the image contains only noise while the bottom part has interesting features such as individual trees and a ridge line. Using the same relative weights throughout the image brings out the interesting features of the sensor image in the foreground but brings out too much of the obscured background and therefore decreases the contrast of the synthetic features in the background where they are most needed. Superior performance was obtained by dividing the image into subparts. The contrast measured for these subparts is used to calculate sensor weights for the center of each region, and the weight at each pixel is calculated as a bilinear interpolation of the region-center weights. This interpolation eliminates boundary problems where adjacent tiles have significantly different weights. In tuning the algorithms prior to the experiment, a $6 \times 6$ grid seemed to offer the best compromise between large regions that could discard important features of the sensor image and small regions for which the standard deviation would not be meaningful.



**Figure 5. Sample sensor image with both high and low content portions**

Finally, in order to further improve the contrast of the sensor-synthetic fused frame, the fused image is remapped to use the full range of available display intensities. The minimum and maximum intensities of the image are used to compute a gain and offset to linearly remap all pixel intensity values. This improves the contrast of the fused image without changing the intensity distribution, resulting in a more natural appearance than that achieved with other contrast enhancement techniques such as histogram equalization.

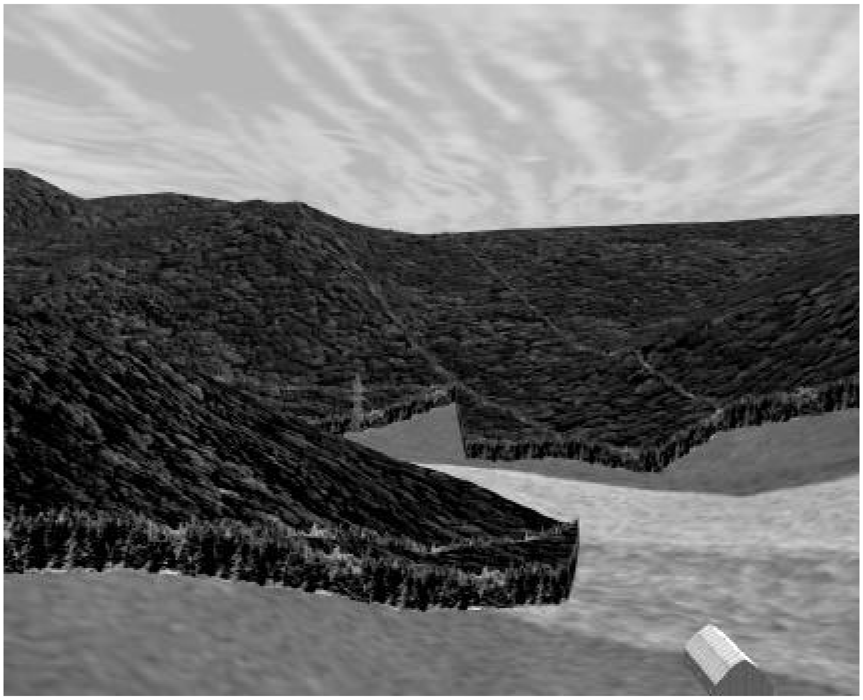

**Figure 6. Synthetic image**

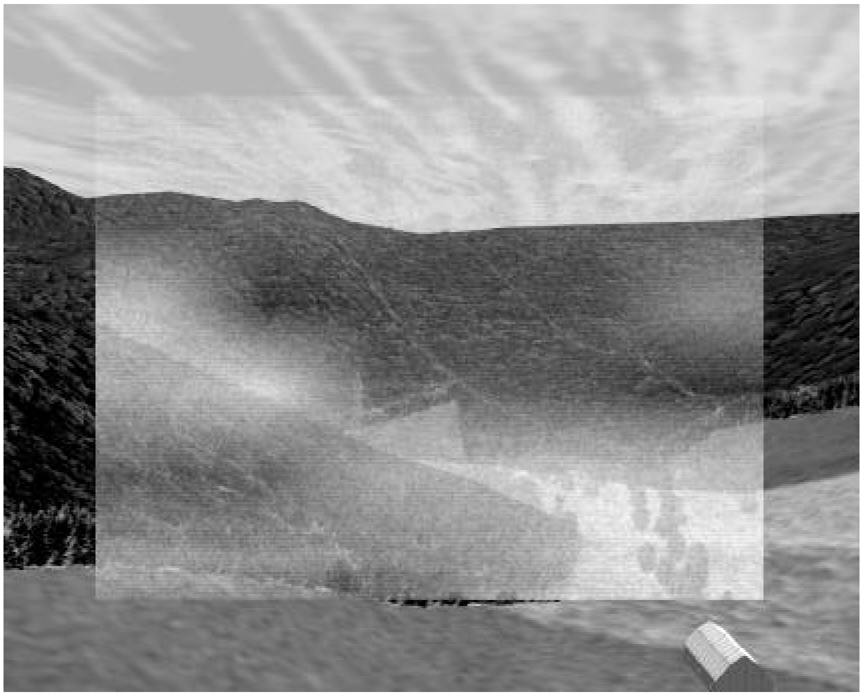

**Figure 7. Fused image using pixel averaging algorithm**

To illustrate the performance of the pixel averaging algorithm, consider a representative synthetic image (Figure 6) which corresponds to the same location in a correlated database (registration condition 3) as that used to simulate the sensor image in Figure 5. The result of fusing these two images is presented in Figure 7. Notice that the fused image contains the visible features of the sensor image while not obscuring the synthetic objects where only sensor noise was present.

### 4.2 TNO-Based Algorithm

The TNO algorithm was adapted and tuned for the purpose of infrared-synthetic fusion to account for sensor image quality and synthetic scene content. The algorithm can be described as follows (note that the algorithm operates strictly on the intensities of individual pixels).

First, the common component of the two original input images is determined. This is simply implemented as a local minimum operator. Next, the common component is subtracted from the original images to obtain the unique component (characteristic) of each image. This unique component represents the details that are unique to each image. The unique component of the synthetic image is then subtracted from the sensor image to enhance the representation of sensor specific details in the final fused result. Finally, a weighted average is computed with the original synthetic image and the enhanced sensor image. The characteristic components of the individual images is further underscored by including into the weighted average the absolute value of the difference between the two characteristic images. As in the pixel averaging algorithm, both the sensor and the difference channel weights applied in the average are determined from the standard deviation of intensities in the original sensor image. More weight is given to the sensor enhanced image and to the difference channel image for higher deviations. The weights are calculated for sub-windows and interpolated in the same manner as in the pixel averaging algorithm.

An example of an image fused with the TNO-based algorithm is illustrated in Figure 8, also using Figure 5 and Figure 6 as the input sensor and synthetic images. This algorithm presents similar results to the averaging algorithm. However, one major advantage of this algorithm is that white objects in the synthetic image can survive the fusion process. This would allow certain important objects (e.g. objects modeled after range sensor returns) to be modeled specifically to remain in the fused image even when the sensor image quality is good.

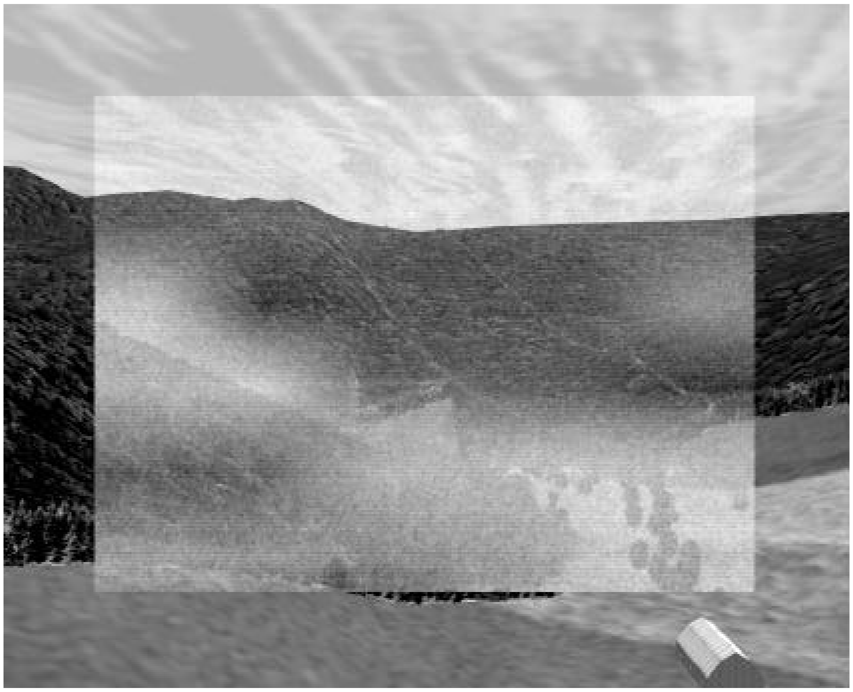

**Figure 8. Fused image with TNO-based algorithm**

### 4.3 MIT-Based Algorithm

Although the MIT algorithm is very effective in fusing low light visible images with IR images, it leads to some anomalous results for infrared-synthetic fusion. The center-surround color opponent cells act as a selection process to preserve the best contrast of the two source images to produce the fused image. As a consequence of the high contrast of the synthetic images compared to that of the sensor images, very little of the sensor image comes through in the fused images even for relatively high quality sensor images. Therefore, the algorithm was modified to take this into account by incorporating parts of the TNO and pixel averaging algorithms. The following is an overview of the modified algorithm.

First, a center-surround cell which represents a spatial differential filter is applied to obtain an initial fusion of the synthetic and sensor original images. The center of the cell is fed by the sensor image and the surround by the inverse synthetic image. This results in a high-contrast image that combines the highest contrast features of both images. Second, the sensor characteristic contribution is calculated by subtracting the synthetic image from the initial fused image. In the same fashion, the synthetic characteristic contribution is calculated by subtracting the sensor image from fused image. The final fused image is then produced using a weighted average of the synthetic characteristic contribution image, the sensor characteristic contribution image, the original sensor image and the original synthetic image. Note that this algorithm uses the global contrast of the entire image to compute the relative weights in the final average. Local contrast is effectively taken into account in the initial MIT-style center-surround fusion.

Figure 9 illustrates the result of fusing the same sensor and synthetic images using the MIT-based algorithm. Observe that the forest canopy on the far hills has reversed contrast. This is due to the particular response of the center-surround cell. Notice also that the mountain ridges at the top of the image are highlighted. This effect can also be explained by the particular response of the center-surround cell. Because it is a neighborhood operation and the surround of the cell is fed with the synthetic image, the original high contrast fused image contains the features of the synthetic frames, but they are blurred. When the synthetic frame is subtracted from this fused image, an edge highlighting is performed on the synthetic image. This can be seen as an advantage for highlighting coarse features such as ridgelines and forest canopy when the synthetic image begins to be obscured by sensor noise.

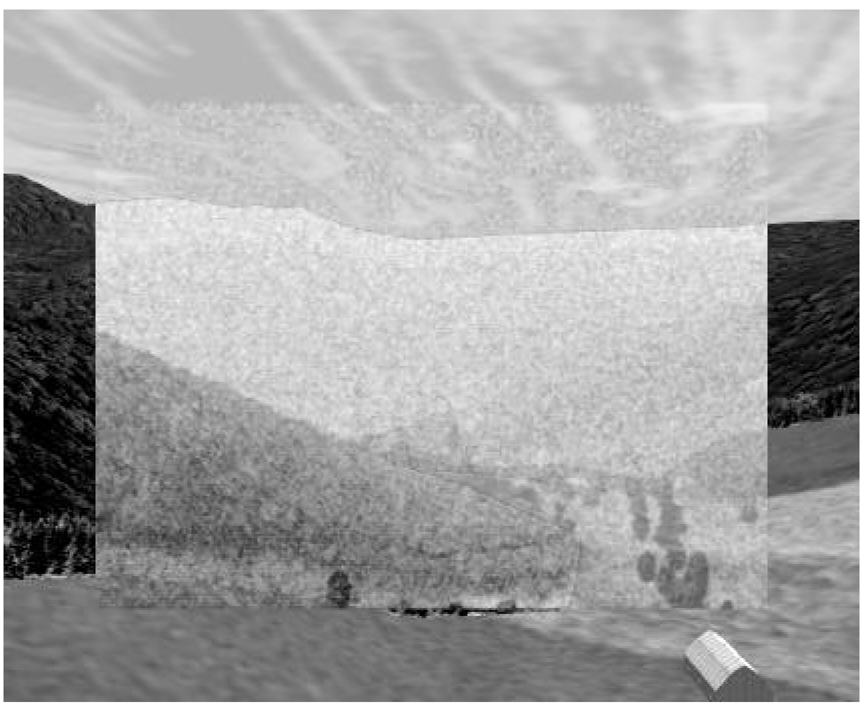

**Figure 9. Fused image with MIT-based algorithm**

# 5. RESULTS AND DISCUSSION

## 5.1 Static Evaluation

The static evaluation verified that fusion improves the useful content of independent synthetic and sensor images. All three algorithms examined provided observers the capability to detect important features from greater distances than sensor alone. The MIT-based algorithm displayed a good overall performance but had some undesirable effects such as contrast reversals. It was also very difficult to tune and is computationally intensive. The pixel averaging algorithm was found to be the simplest in terms of required operations per pixel. The TNO-based algorithm, although slightly more complex, demonstrated results similar to the pixel averaging and has the advantage of preserving certain synthetic features in a predictable way that would facilitate modeling and display of range sensor features. Although white hot and black hot sensor modes seemed to be equivalent, the white hot polarity demonstrated better results in the fused images. The three subjects also found the white hot polarity to have a more natural appearance.

Experimental results[2] indicate that viewpoint noise (due to inaccuracies in the aircraft navigation system) and global database offset would not have a severe impact on the system. Rather, such errors appear to be tolerable (and compensated for) by the observer. Similarly, lower terrain resolutions did not have a great impact on interpretation of the fused images.

However, terrain elevation errors from source data had more severe effects on registration and fusion. Subjects experienced confusion and had difficulty distinguishing real terrain from synthetic terrain. This effect may be less severe when the synthetic image has a lower resolution content (e.g. polygonal forest canopies). More significantly, it was discovered that the viewpoint may drop below synthetic terrain that is misaligned due to large elevation errors in source data. In this situation, the synthetic system would provide unpredictable results, generating confusing images with surreal scene content that should be obscured. Further development of techniques to deal with these problems is required.

## 5.2 Dynamic Evaluation

Overall performance with all three algorithms was markedly improved with the dynamic sequences, with an increased visibility of significant features. Many of the inconsistencies in the static image sequences with the MIT-based algorithm were resolved, although anomalies such as contrast reversal of the forest canopy remained. We attribute this not to algorithm characteristics, but to superior processing capability in the human visual system when dealing with dynamic imagery. On being presented with dynamic information, the human visual system integrates over several frames to permit correlation of scene elements from one frame to the next, reducing noise and cleaning up persistence effects in the display of noisy images. This permitted objects to be separated from noisy backgrounds and to be detected at lower thresholds. It also allowed some separation of synthetic and sensor content where conflicts had previously been reported, especially for objects which at some point crossed the border between the fused inset and the purely synthetic background.

In initial tests, dynamic sequences were subject to severe flickering, which was traced to two problems:

1. Variation of the relationship between the computed contrast in adjacent sub-windows (used to calculate the weights in the pixel averaging and TNO-based algorithms) due to small shifts in the direction of regard.
2. Small changes from one frame to the next in the intensity distribution of the fused images.

The first flickering effect that was observed for the pixel averaging and TNO-based algorithms was manifested as a flashing at the borders between the regions of the weight grid. The cause was traced to changes in the noise from frame to frame, and more particularly to non-uniform contrast changes in the sub-windows, which led to disproportionate changes in the averaging weights. The function used to translate contrast measures into weights has a relatively high gain, which results in a large variation of weights even for small variations in contrast.

The second problem, apparent in image sequences obtained using any of the three algorithms, was a brightness flicker of the whole image. The cause was traced to the final remapping of the fused images to the full intensity range. It uses the minimum and maximum pixel intensities found in the image and because of the noise, these intensities vary between two successive frames. As a consequence, the remapping function varies too much between two images.

The solution that was implemented was to stabilize both the weights and the remapping values so that they would vary more slowly. This was achieved by maintaining a history buffer of the values computed in previous frames. At each frame, the weights and parameters were stored in a first-in first-out (FIFO) buffer and an average of the history buffer computed.

## 6. CONCLUSIONS

The results of the present study showed that fusion improves the useful content of independent synthetic and infrared images. Hence, image fusion could significantly enhance pilot performance in terrain and obstacle avoidance in poor visibility conditions. From the three algorithms studied, the simpler ones, i.e. the pixel averaging and TNO-based algorithms, seem to perform the best. The TNO-based method had the additional advantage of preserving certain synthetic features that could be used to incorporate range data into the synthetic terrain database. The dynamic sequences showed a further improvement in the performance resulting in an increased visibility of key features.

An ESVS Technology Demonstrator is currently under development at CAE, CMC, and NRC FRL. It is scheduled for flight tests in the NRC FRL Bell 205 Airborne Research Simulator in August 2000

## 7. ACKNOWLEDGEMENTS

The ESVS program is primarily supported financially by the Canadian Department of Defense - Chief, Research and Development, and the Search and Rescue Directorate. LCol R. Thompson, leader of the Advanced Cockpit Technologies initiative, has been instrumental in assuring that the program remains on track and is properly supported. Additional funding has come from the Centre for Research in Earth and Space Technology (CRESTech), an Ontario (Canada) government center of excellence.

## 8. REFERENCES

1. 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), May 1999.
2. 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.
3. Toet, A. & Walraven, J, "New false colour mapping for image fusion", *Optical Engineering, 35(3)*, pp. 650-658, 1996.
4. Waxman, A.M. et al, "Color night vision: fusion of intensified visible and thermal IR imagery", *Proc. SPIE Conference on Synthetic Vision for Vehicle Guidance and Control, vol. SPIE-2463* (pp. 58-68), 1995.
5. Waxman, A.M. et al, "Electronic imaging aids for night driving: low-light CCD, thermal IR, and color fused visible/IR", *Proc. SPIE Conference on Transportation Sensors and Controls, SPIE-2902*, 1996a.
6. Waxman, A.M. et al, "Progress on color night vision: visible/IR fusion, perception & search, and low-light CCD imaging", *Proc. SPIE Conference on Enhanced and Synthetic Vision, vol. SPIE-2736* (pp. 96-107), 1996b.
7. Waxman, A.M. et al, "Color night vision: opponent processing in the fusion of visible and IR imagery", *Neural Networks, 10(1)*, 1-6, 1997.