Hyperspectral Spatial Enhancement

W. Oxford
Waterfall Solutions, Parklands,
Guildford, Surrey GU2 9JX.

Abstract

Hyperspectral (HS) imaging systems have a proven military capability for tasks such as the defeat of targets employing camouflage concealment and deception (CC&D) techniques. Broadband, colour and MS imaging systems will dominate the battlespace for the foreseeable future, whereas HS imaging systems are likely to remain rare assets. The research reported here explored techniques for the fusion of HS and MS imagery in order to retain their complementary properties of spectral and spatial resolution. Techniques for the automatic registration and fusion of the imagery are presented and example results are given based on real HS imagery. The results show that a fused product can be generated with only a small amount of spectral distortion, and is therefore suitable for further spectral exploitation.

Keywords: Spectral, Spatial Enhancement, Fusion, Surveillance, Reconnaissance

Introduction

Hyperspectral (HS) imaging systems have a proven military capability over and above that of multispectral (MS) and broadband systems for tasks including the detection of vehicles under canopy and employing camouflage, and the discrimination of real targets and decoys.

However, HS imaging systems are likely to be rare assets in any future battlespace, whilst broadband and colour imaging systems will dominate into the foreseeable future. These sensors possess high spatial and temporal resolution, whereas a typical HS sensor possesses a high spectral (and hence discriminatory) resolution. It would be desirable to fuse the benefits of these two modalities to enable improved exploitation (i.e. the discrimination of smaller targets via their spectral signature).

This paper reviews a three month feasibility study into the fusion of HS and MS imagery for the Electromagnetic Remote Sensing (EMRS) Defence Technology Centre (DTC).

The fusion of any two imaging modalities begins with spatial image registration (assuming that they are temporally synchronised). In this work the registration problem was limited to only considering near-nadir imagery collected by high altitude platforms, and so the spatial registration problem was not significant. However, the fusion of MS and HS imagery does present an additional challenge in that there is the potential for very large differences in spatial resolution to be dealt with.

This study investigated the integration of spectral information into image-based and feature-based registration algorithms with the intention of improving the speed, accuracy and robustness of the registration process.

The process of image fusion was addressed firstly in spectral feature space and secondly in spatial image space. The relationship between the data in the MS and HS feature spaces was captured by an optimal linear mapping which allowed the MS data to be mapped into the HS feature space at the full spatial resolution of the MS imagery (which is generally greater than the resolution of HS imagery).

The resulting estimated HS image cube occupies a sub-space of the original HS data and the role of spatial fusion was to inject HS information at the low spatial scales of the HS image. A number of spatial image fusion techniques were implemented, including multi-resolution and frequency-based algorithms.

A software framework was developed which can synthesise MS and HS imagery from real HS data, apply image registration and image fusion algorithms, and assess the output at each stage of the processing chain.

This paper now introduces elements of the framework followed by an analysis of some typical results.

Image Synthesis

Throughout the study, elements of the framework were tested using HS and MS images synthesised using imagery from the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) sensor.

MS imagery was synthesised by spectrally resampling the AVIRIS imagery using the spectral response of IKONOS, a commercial space-based sensor (see Figure 1), yielding red, green, blue (RGB) and near infra-red (NIR) wavebands.

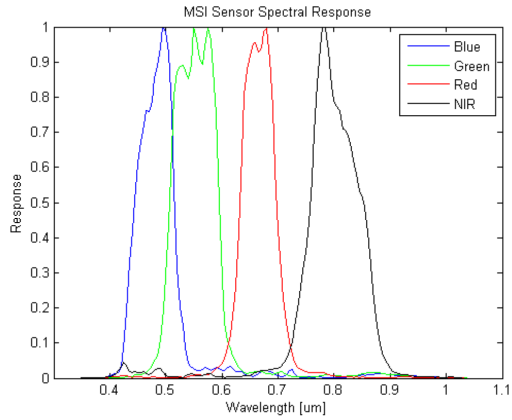

Figure 1 – MS spectral response curves

HS imagery was synthesised by spatially aggregating the AVIRIS imagery using linear spatial scales of 2, 4 and 8 (or 4, 16 and 64 in area). Only spectral bands covering the range of the MS imagery were retained (~0.4-1µm).

Spatial Registration

Image registration is an often overlooked or assumed pre-requisite of any image fusion process. This study has investigated the utility of incorporating spectral information into typical registration algorithms, with both image-based and feature-based techniques being addressed.

The spatial relationship between the MS and HS imagery was assumed to be described by a global projective transform. This transform is appropriate to high altitude, nadir-viewing platforms for which parallax and obscuration can be ignored.

One of a number of high-performance image-based registration techniques developed by WS under SEAS DTC funding was selected for assessment. The algorithm chosen uses a multi-scale gradient descent algorithm to find the global minima of a given cost function. The baseline cost function selected was the global mean-square-error (MSE) between the registered and unregistered images. The minimisation routine of the baseline algorithm required the gradient of the cost function with respect to each parameter of the projective transform as well as the overall measure between the registered and unregistered images

The MSE cost function was applied to the first principal component (PC) of the MS and HS imagery. This was done to maximise the variance, and hence contrast, in the images provided to the registration algorithm.

Two additional cost functions were implemented into the technique: spectral angle and Euclidean distances. These were selected due to their widespread use as spectral metrics.

A feature-based registration algorithm was developed that found the optimal transform between matching features identified in the MS and HS images. Candidate matches were generated using a corner detector and each was described using surrounding image values. These features were then pair-wise matched based on a similarity metric accounting for a potential relative rotation between the prospective matches. The similarity metrics used were based on correlation (again using the first PC bands), spectral angle and Euclidean distance.

Matching pairs were filtered based on their similarity score and relative rotation. The remaining matches were then randomly partitioned a large number of times into two sets – one set was used to compute the optimal projective transform, and the spatial error associated with this transform was evaluated on the second set. The projective transform yielding the lowest error was assumed to describe the global projective transform between the MS and HS images.

The performance of the image and feature-based registration algorithms was tested using synthesised MS and HS data at differing spatial scales and for increasing levels of initial misregistration.

The results showed that for the image-based registration the spectral cost functions were capable of registering the imagery at low levels of initial misregistration but failed to register the images at high levels of initial misregistration. The MSE cost function was capable of registering the images to within ½ of a HS image pixel in all test cases.

The performance of the spectral cost functions was lower than expected. One reason for this may have been an inappropriate pairing of the cost function and the minimisation routine (essentially a local search). The spectral cost functions may exhibit multiple local minima which would be addressed more appropriately by a guided stochastic search approach such as genetic algorithms or simulated annealing.

Tests using the available data highlighted the fact that the feature-based registration technique was not capable of registering the images at the highest spatial scales when there was a lack of distinct corner features in the HS image. When feature-based registration did perform well, the error was slightly inferior to that given by the image-based registration technique due to the relatively poor localisation of the feature matches (i.e. 1 pixel at best).

The robustness to noise of the spectral cost functions in the image-based registration technique was investigated and the results are shown in Figure 2.

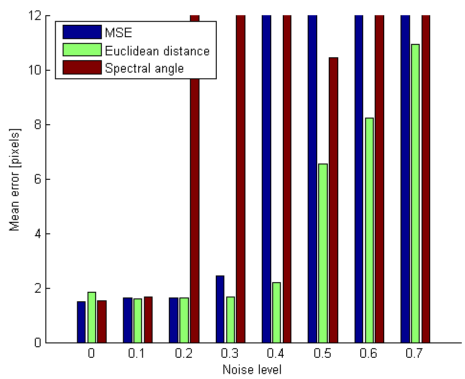

Figure 2 – Robustness to noise

The results show that the Euclidean distance cost function was significantly more robust to noise than the MSE cost function. This implies that the metric would find utility where robustness to noise is a requirement.

Spectral Fusion

Spectral fusion describes the mapping, G, from the low dimensional space occupied by the MS image, D_M, to the high dimensional space occupied by the HS image, D_H, with error, ε, as shown by Equation 1.

$$GD_M = D_H + \varepsilon \qquad [1]$$

The optimal linear mapping (in the least squares sense) is given by Equation 2:

$$G = \left(D_M D_M^T\right)^{-1} D_M^T D_H \qquad [2]$$

The result of applying this mapping to the MS image was an estimated HS image at the spatial resolution of the MS image. Equation 2 was augmented to account for offsets between the HS and HS radiance images that might be present due to path radiance effects.

The accuracy of the mapping was evaluated by comparing the estimated HS image with the truth HS image at the resolution of the MS image. The mean fractional Euclidean distance (or spectral) error was used to quantify the accuracy of the fusion process.

Figure 3 shows the spectral error for the original (i.e. down sampled) HS image and estimated HS image after spectral fusion.

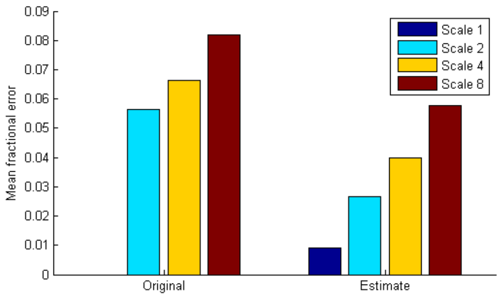

Figure 3 – Spectral fusion spectral error

Figure 3 shows that, with the exception of a spatial scale factor of one (which does not involve down-sampling), the estimated HS image, although based upon a linear mapping of only four MS bands, describes the truth HS image better than the reduced spatial resolution HS image.

The linear mapping was extended to non-linear mappings in an attempt to reduce spectral error. However, no significant effect was observed for a second-order mapping. This could be explained by the fact that the test scene was comprised of few ground-cover types (mainly vegetation) and that the process of down-sampling simplifies the HS feature space.

The concept of the linear mapping can be extended to RGB and broadband sensors. Figure 4 shows the spectral error for MS, RGB and NIR sensors.

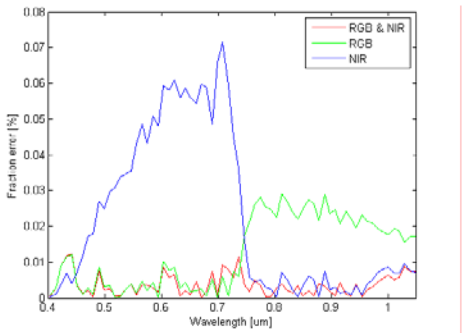

Figure 4 – Spectral error for MS, RGB and NIR sensors

Figure 4 shows that the spectral error increases in spectral regions not covered by the MS sensor. Figure 4 also implies that the NIR region contains a significant amount of information - the latter can be explained by variations in the chlorophyll edge of the largely vegetative land cover.

Spatial Fusion

Spatial fusion is used to fuse the estimated HS image, as a result of spectral fusion, and the down-sampled HS image. Spatial fusion, on a per-band basis, attempts to increase the dimensionality of the estimated

HS image by introducing HS data at the spatial resolution of the HS image.

Three classes of spatial fusion algorithm were implemented: weighted average, multi-resolution and frequency based.

The weighted average techniques were included as baseline methods because of their simplicity and prevalence, but they do not acknowledge the saliency of the information in the sources images at different spatial scales and generally give poor-quality fused images. The schemes implemented were the simple weighted average and PC average (in which the weights are chosen to maximise the variance of the fused image).

The multi-resolution algorithms attempt to describe the detail within the source images at a number of length scales. Detail is preserved at each length scale and represented in the final fused output. The distinction between the different techniques is given by the variety of means employed to derive detail information between resolution levels. The multi-resolution techniques implemented were the Laplacian, Contrast, Gradient, Ratio, and Gaussian.

At each resolution level the detail from each source image was fused according to a weighted average rule. The weights chosen acknowledged the mismatch in spatial resolution between the source images and were found to provide better results than taking the largest value at each resolution level.

The frequency-based approach decomposed each source image into its discrete cosine transform. The transform of the HS image was filtered by a two-pole Butterworth filter with a cut-off frequency appropriate to its spatial resolution, and the MS image was filtered by the complement of the filter. The output image was found by the inverse transform of the sum of these filtered transforms.

The spatial fusion techniques were assessed for different spatial scales of the HS image. The results for are several techniques are shown in Figure 5 which shows that, as the spatial scale increases, the performance of all the techniques becomes more similar and the benefit afforded by the best technique over the HS estimate diminishes. The figure also shows that the best performing technique was the Laplacian algorithm.

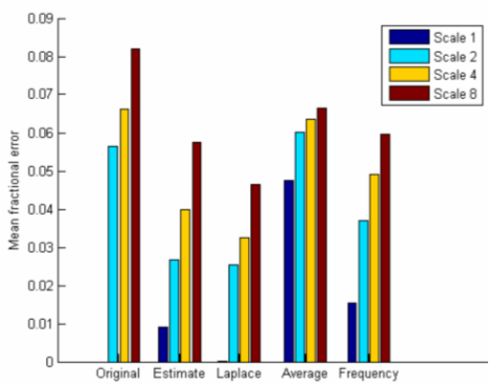

Figure 5 – Spectral error for spatial fusion techniques

Fused Product

The assessment of the spectral and spatial techniques presented thus far used a global measure of spectral error. It is informative to examine other measures to illuminate the behaviour of the fusion processes.

Figure 6 shows a typical eigenvalue spectrum of the HS images before fusion, after spectral fusion, and after spatial fusion.

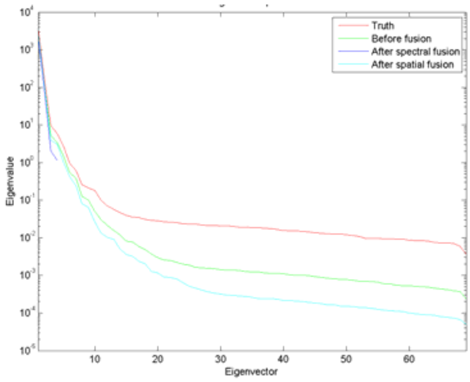

Figure 6 – Eigenvalue spectrum before & after fusion

Figure 6 shows the reduction in the eigenvalues caused by spatial aggregation and the truncated dimensionality of the image after spectral fusion. The effect of spatial fusion is to increase the dimensionality of the image, although the dimensionality before fusion is not fully recovered. This can be regarded as the cost of trading spectral information for spatial resolution.

Figure 7 shows a typical cumulative histogram of spectral error when the images were compared to the truth HS image.

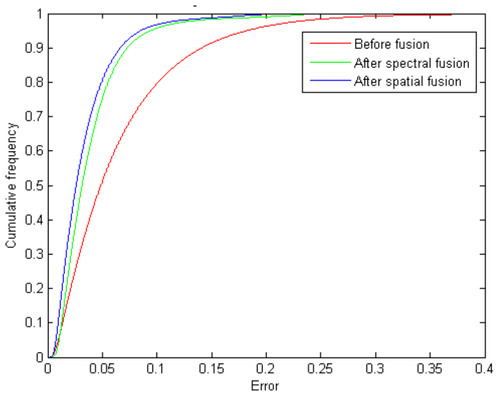

Figure 7 – Cumulative spectral error histogram

Figure 7 shows the overall benefit of first spectral and then spectral fusion. In terms of the impact on pixel spectra, Figure 8 shows spectra from the fused and truth HS images for the 5%, 50% and 95% percentile pixels when ranked in order of spectral error.

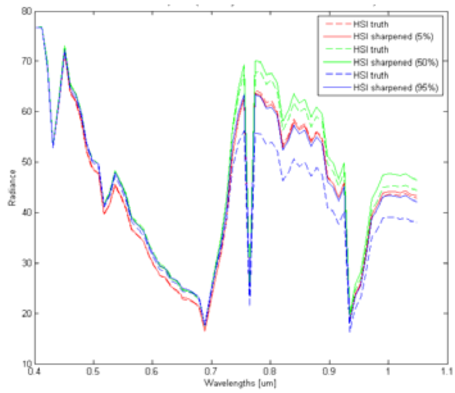

Figure 8 – Pixel spectra by spectral error

Figure 8 indicates that the majority of the error occurs towards longer wavelengths, but that, even at the 95% pixel error, the shape of the spectra are very well preserved.

Discussion

The techniques developed covered image registration as well as spectral and spatial fusion. New image registration algorithms were developed to incorporate spectral information. Overall, image-based registration was shown to be preferable to feature-based registration due to the potentially low spatial resolution of the HS imagery. Within the scope of the image-based registration assessment, it was shown that the utility of spectral information is limited to cases for which robustness to noise is a strong requirement. However, additional research is required into the use of minimisation routines capable of coping with the more complex spectral cost functions employed.

A spectral fusion technique was developed that does not require the MS and HS images to be converted into reflectance since it can cope with offsets potentially caused by path radiance effects. Although a second-order non-linear mapping was shown to have negligible effect, further research into non-linear mappings may provide benefit here, as they have done in many other areas of spectral exploitation.

A number of popular spatial fusion techniques were assessed and the multi-resolution Laplacian algorithm was found to be the best performing. The combination of spectral and spatial fusion has been shown to yield a fused product with low spectral distortion.

These results are encouraging and indicate that the fusion techniques developed in this project will have utility in many spectral exploitation disciplines. The tasks most likely to benefit are those in which the objects of interest are unresolved in the HS image, yet are unable to be spectrally identified in the MS image (e.g. target detection and mapping).

Acknowledgements

The work reported in this paper was funded by the Electro-Magnetic Remote Sensing (EMRS) Defence Technology Centre, established by the UK Ministry of Defence and run by a consortium of SELEX Galileo, Thales UK, Roke Manor Research and Filtronic.