Invariant Hyperspectral Un-mixing and Material Identification

K Smart, J Nothard and W Hirsch
QinetiQ, Cody Technology Park,
Iveley Road, Farnborough, GU14 0LX

Abstract

A new algorithm based on a support vector machine (SVM) has been used to un-mix airborne hyperspectral imagery. The algorithm, called Platypus, is a learning machine classifier, and has been trained using field measurements from hand-held devices of spectra of typical military targets and backgrounds. These high quality field measurements are degraded to match the airborne sensor's spectral resolution, and were only grossly matched to the scene contents. Multiple training examples are used to capture the natural variation for each material type, which cannot be readily done in existing conventional un-mixing methods. Excellent results have been obtained on data from a range of airborne sensors, compared to standard un-mixing techniques. The new algorithm's ability to naturally perform non-linear un-mixing is especially useful in the visible to short wave infrared waveband. The technique might be extended further to un-mix the atmospheric effects, and provide invariant material identification and "difficult" target detection in one step.

Keywords: un-mixing, Hyperspectral, material identification, invariant, atmosphere

Introduction

Many algorithms have been developed and much analysis has been performed on hyperspectral data sets over the past ten years. However, fundamental problems still exist in the robustness of material identification techniques, which often struggle to cope with the variation in material and atmospheric spectra, and rely on high quality spectral library data. Existing spectral un-mixing algorithms typically rely on finding spectral endmembers "in scene" from the airborne hyperspectral imagery, and have been assessed by previous work in the DTC [1] and in the US [2] as the weak link in the hyperspectral processing chain.

The aim of this short pre-cursor study was to investigate a new algorithm, called Platypus, for un-mixing of airborne hyperspectral imagery. The algorithm, based on a support vector machine (SVM) classifier, also approximately matches the material type, so effectively performs spectral un-mixing, target detection and broad material identification in a single step.

The idea behind the algorithm is very simple, and works by associating classifier confidence values with mixing ratios. Platypus is an SVM capable of calculating confidences associated to its classifications [3]. The confidences are calculated by fitting a sigmoid through binary pairs of the material spectra.

The main military benefit of the proposed technique is that better spectral un-mixing should deliver enhanced capability to detect long range or "difficult" targets such as those employing CC&D. In addition, the use of prior knowledge of likely scene content to perform un-mixing, target detection and material identification in one step cuts the processing chain down greatly, making a real-time on-platform solution easily practical.

Background

The idea resulted from research by QinetiQ into a very different application of hyperspectral sensing. Platypus was trained to recognise several different material characteristics, e.g. wool, cotton, silk, and human skin. Analysis of spectral measurements taken from a human subject indicated Platypus was un-mixing the hyperspectral pixels by assigning a confidence to that class proportional to the mixing contribution from it, illustrated in Figure 1.

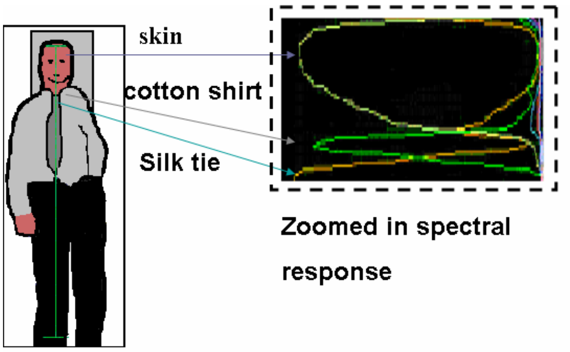

Figure 1 - Example of spectral un-mixing using Platypus, where the right hand plot shows the SVM confidence values swapping from one class to another with mixed pixels assigned intermediate confidences for all the contributing classes

Figure 1 shows analysis of spectra taken along the vertical green line (the field of view of the "line-scan" spectral sensor) on the subject. From the shirt collar where the tie begins, is probably a mix of skin, cotton and silk. Similarly, the centre of the face has a clear peak in the confidence of skin, and this is the purest skin region. There appeared to be a correlation between classifier confidence value and mixing proportions. This work set out to confirm and extend that apparent effect to airborne hyperspectral imagery.

Selection of training data

Platypus is a learning machine and needs to be trained with spectra of typical military targets and backgrounds, measured at short range using hand-held devices in the field or laboratory. However, these measurements are not required to be high quality, since they will be degraded to match the airborne sensor, and need only approximately match the expected scene contents. For example, any examples of leaves and grasses can be used to define "green vegetation", and there is no requirement to have the right kind of tree at the right time of year, as needed by many conventional spectral library matching approaches.

Several different atmospherically corrected image sets were chosen for our analysis of the following airborne sensors:

- AHI (long wave infrared)
- HYMAP (visible to short wave infrared)
- ASI (visible to near infrared)

The training classes were chosen to roughly match the scene contents, which were primarily rural. The following broad classes were found to be appropriate for the visible to short wave infrared (V/SWIR) waveband:

- Healthy vegetation (leaves and grass)
- Other vegetation (dead leaves and grass, bark and acorns)
- Dirt and rock
- Green painted targets
- Other targets

It was also found possible to split the healthy vegetation into two sub groups: illuminated and shaded. The broad classes selected for the long wave infrared (LWIR) waveband were different, principally because the phenomenology in the LWIR is very bland but never the less provides essential discriminative information:

- Vegetation
- Dirt and rock
- Painted targets

Un-mixing AHI imagery with Platypus

For the three LWIR classes identified above, Platypus was trained with field measurements taken from an independent trial (and in another part of the world, time of year etc.) modelled by integration to match AHI spectral bands. The linear kernel was found to give the best performance for AHI, and an example "class image" is shown in Figure 2.

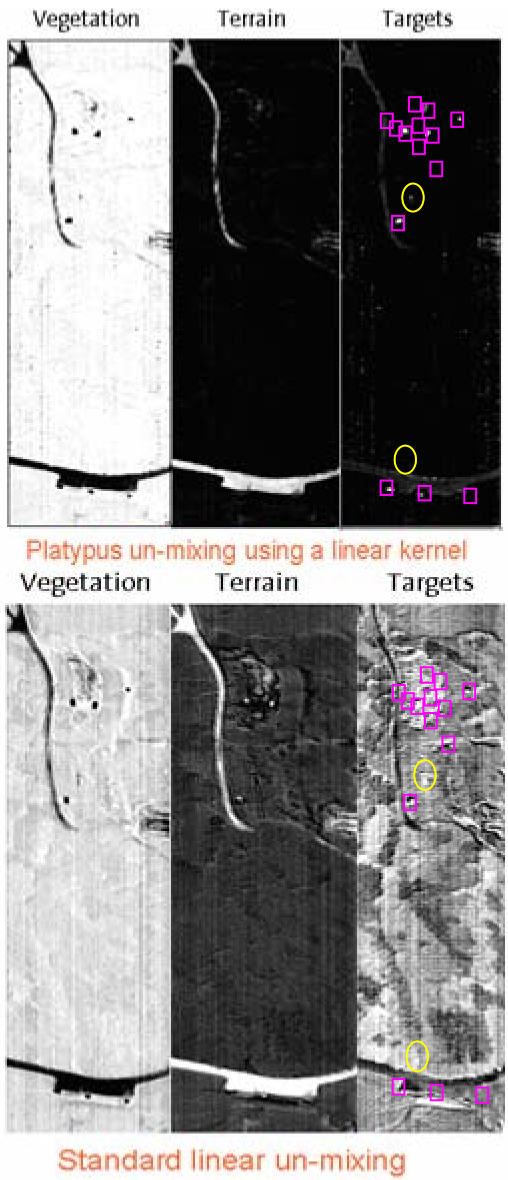

Figure 2 - Comparison of un-mixing using Platypus (top) with standard linear un-mixing (bottom). Each technique generates a "class image" for each of the three trained classes, in which intensity is proportional to probability of the pixel belonging to that class. Note that the actual positions of man-made objects are indicated by the pink squares, and the military vehicle targets by the yellow circles

The standard linear un-mixing algorithm (taken from the ENVI™ software package) associates a high target probability to most of the pixels in the image, which would result in a very high false alarm rate. By contrast, the Targets class image generated by the new algorithm contains several bright spots associated with targets, and would result in very few false alarms. In general, Platypus performs much better on unseen airborne data than the existing non-linear un-mixing. However, the target in hide (which is indicated by the lower yellow circle in Figure 2) is not found by either of the un-mixing methods, due to atmospheric correction degrading the spectral information, which is illustrated in Figure 3.

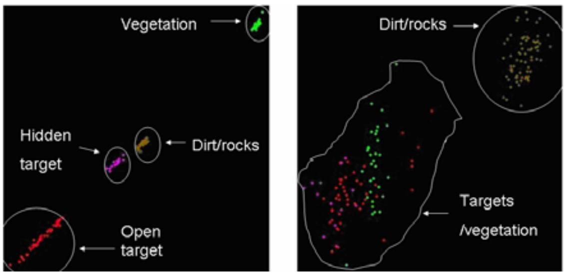

Figure 3 – Scatter plots of AHI radiance (left) and atmospherically corrected data (right) projected into 2D space

Figure 3 shows high dimensionality AHI data projected onto a randomly selected 2D plane. For the uncorrected radiance data, the plot indicates that all the materials are clearly distinct and separable. However, atmospheric correction by this (state of the art) method clearly degrades the spectra with only the dirt/rock spectra remaining distinct.

Finally, an experiment was designed where selected training spectra were deliberately mixed, using linear mixing. Platypus was then used to separate them and to estimate the mixing ratios. A linear kernel gave the results summarised in Table 1. Note the last example shows two targets at 40%, this represents the fact that two different target spectra have been used to make the known mixture.

Known mixing ratios	Veget-ation	Terrain	Targets
80% veg 20% terrain	78.5	15.4	6.1
80% veg 20% target	78.5	14.6	7.0
30% veg 70% target	18.8	66.2	15.0
90% terrain 10% target	3.8	92.0	4.1
20% terrain 40% target 1 40% target 2	10.9	57.3	31.8

Table 3 - Example un-mixing of LWIR, vegetation modelled on Gaussian varied Planck spectrum, using a linear kernel

The un-mixing ratios are not accurate, but do reflect the general trends. It is recommended that the relationship between the confidence value and the mixing ratio is further investigated in any follow-up work, especially for non-linear kernels. Alternative confidence measures should also be considered.

Un-mixing HYMAP and ASI imagery with Platypus

Using the five V/SWIR classes Platypus was trained with field measurements modelled by integration to match HYMAP spectral bands. (Again the measurements were from an independent trial in part of the world, time of year etc.) The triangle kernel (Euclidean distance) was found to give the best performance for HYMAP and for ASI. Note that this is a non-linear kernel indicating that the mixtures in the V/SWIR region may be non-linear.
The trained algorithm was applied to a range of HYMAP and ASI imagery, and the performance analysed in terms of target detection and false alarm rates. The detections were generated by simply thresholding the Targets class image, associating pixels into connected components and rejecting components below a certain size. Moving the threshold generates a receiver operating characteristic (ROC) curve which can be compared to ROC curves generated in a similar way for other un-mixing and target detection methods.

The results for ASI are slightly better than for Hymap, but in both cases the new algorithm significantly out-performed standard linear un-mixing techniques, both in reduction of false alarm rates and in detection of partially hidden or camouflaged targets. The remaining false alarms are mainly from the tree-line, and targets against the tree-line are also difficult to detect, suggesting that invariance to incidence angle and/or shadowing still needs refining.

In a second experiment, the Hymap data was split into two sets: just the visible to near infrared (V/NIR) and the full spectral range (V/SWIR). Un-mixing and target detection performance was compared in the two wavebands. In both cases a non-linear kernel was required to separate the mixtures, indicating non-linear mixing mechanisms, possibly due to adjacency and multiple scattering effects. The V/SWIR results were significantly better than the V/NIR, indicating that much of the discriminatory information comes from the SWIR part of the waveband.

Analysis of histograms of classifier confidence values for all the materials shows their distributions are multimodal. This suggests that an alternative confidence measure is needed to accurately estimate mixing ratios for non-linear kernels.

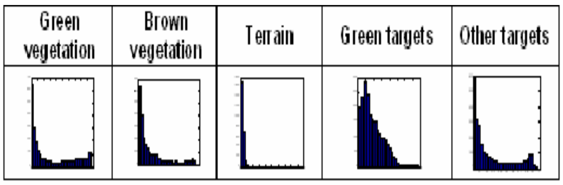

Figure 4 - HYMAP V/NIR bands, target confidences are multi-modal distributed, resulting in inaccurate mixing ratio estimates

Conclusions

A novel un-mixing method has been developed based on the confidence values output by a support vector machine classifier. The classifier was trained using field measurements of typical spectra and tested against airborne imagery from a variety of sensors flown at a number of different trials.

The results indicate that:

- Un-mixing performance is significantly better than standard techniques, leading to much reduced false alarm rates
- Detection of hidden or camouflaged targets should be enhanced by using the new technique
- The training examples need to be only roughly matched to the scene contents
- Non-linear mixing is a significant effect in the V/SWIR region
- Current atmospheric correction methods in the LWIR destroy the subtle spectral information needed to detect difficult targets

The recommendations for a follow-on programme of work are twofold: to investigate the relationship between confidence and mixing ratio, especially for non-linear kernels, and derive new confidence measures if required; and to evaluate non-linear un-mixing of the atmosphere in the same way as the other spectral components, removing the need for conventional atmospheric correction.

References

1. Peter Yuen, Gary Bishop, Adrian Blagg, '*Geometric and statistical spectral unmixing for subpixel target detection*' EMRS DTC annual conference 2005
2. Heesung Kwon, Nasser M. Nasrabadi, '*Kernel-based subpixel target detection in hyperspectral images*' Neural Networks IEEE, 2004
3. Platt JC, '*Probabilistic Outputs for Support Vector Machines and Comparison to Regularized Likelihood Methods*' www.Microsoft.com March 1999

Acknowledgements

The authors gratefully acknowledge the help of Phil Clare, Dstl, UK, with the atmospheric correction of LWIR imagery.
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 SELEX Sensors and Airborne Systems, Thales Defence, Roke Manor Research and Filtronic.