Hyperspectral Thermal Infrared Remote Sensing of the Land Surface and Target Identification using Airborne Interferometry

Dr Jonathan P Taylor, Dr Jean-Claude Thelen and Dr Stephan Havemann
Met Office, Fitzroy Road, Exeter, Devon, UK. EX1 3PB

INTRODUCTION

The main applications of this research project relate to the real-time target acquisition and target identification, for military purposes, using various hyperspectral sensors. The advantages for military personnel are twofold. Firstly, they would be able to identify, in real-time, man-made structures in a natural environment and secondly they would be able to ascertain to a certain extent whether these structures were civilian or military and in the latter case whether they were friendly or not. The technique for target acquisition/identification presented here relies on the fact that artificial surfaces and natural surfaces have, in general, widely different surface emissivity spectra. Thus, in theory, measuring the surface emissivity of various objects should allow us to distinguish between them.

Retrieval of surface emissivity is a non-trivial problem as the radiance signature seen at altitude by a sensor on an aircraft or satellite is a complex combination of the surface properties and the intervening atmospheric structure including the temperature and water vapour profile and presence of clouds or aerosols. These atmospheric properties affect the radiance in different ways dependent on the wavelength and this requires modelling of the effects using a radiative transfer code. Operational weather forecasting has to address these same problems as the initial structure of the atmosphere needs to be deduced from many data sources which are dominated by satellite observations. Within operational numerical weather prediction (NWP) the data are assimilated in to the forecast model using a technique called variational data assimilation. This technique aims to solve a cost function that identifies the optimum description of the atmosphere that minimises the difference between the observed and simulated spectra whilst fully in to account the known error covariance characteristics of the observations and the model.

The final problem that needs to be overcome both within NWP and in military applications is the speed of the radiative transfer calculations. Meteorological satellite instruments like the Infrared Atmospheric Sounding Interferometer (IASI) measures the radiance spectrum between 3 and 15 microns at a spectral resolution of 0.25cm-1 resulting in some 8461 radiance channels. Computing the radiative transfer of these hyperspectral data is historically very time consuming. In this report we present early results using a new fast radiative transfer scheme – the Havemann Taylor Fast Radiative Transfer Code (HT-FRTC) - which conducts the radiative transfer in principal component space which is coupled to a 1d-variational data assimilation technique to retrieve the full atmospheric and surface state from hyperspectral sounder data.

PRINCIPAL COMPONENT RADIATIVE TRANSFER AND VARIATIONAL DATA ASSIMILATION

If one constructs a large representative set of radiance spectra which has a spectral resolution significantly finer than the observation of interest then conducting Singular Value Decomposition of this array of spectra delivers three matrices a set of Empirical Orthogonal Functions (EOF), a diagonal matrix of singular values and a matrix of principal component (PC) scores. These EOF represent the spectral

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
OCT 2009	**N/A**	**-**

4. TITLE AND SUBTITLE
Hyperspectral Thermal Infrared Remote Sensing of the Land Surface and Target Identification using Airborne Interferometry

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Met Office, Fitzroy Road, Exeter, Devon, UK. EX1 3PB

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADB381583. RTO-MP-SET-151 Thermal Hyperspectral Imagery (Imagerie hyperspectrale thermique). Meeting Proceedings of Sensors and Electronics Panel (SET) Specialists Meeting held at the Belgian Royal Military Academy, Brussels, Belgium on 26-27 October 2009., The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**SAR**	**12**	

physical characteristics of the atmospheric gases, temperature, clouds, aerosols, surface properties and the instrument. Each element of the PC score matrix represents the actual atmospheric state. Furthermore analysis of the singular values shows that the Eigenvalue decreases rapidly with the latter singular values adding less and less information. Analysis shows that only the leading singular values (of order 100) are required to reconstruct the spectrum to a sufficient accuracy. This means that if one can compute the leading 100 PC scores for a new atmospheric state then by simple matrix multiplication one can recombine the training set EOF and singular value array with the new PC score matrix to generate a hyperspectral radiance spectrum.

The HT-FRTC code allows the rapid calculation of the PC scores and hence the prediction of hyperspectral radiance spectra. To do this a cluster analysis has been performed that identifies the key monochromatic radiance calculations that contain the information. The HT-FRTC takes a new atmospheric state, computes the mono-chromatic radiances and then through a series of polynomial relationships computes the corresponding PC scores. If required these PC scores can be used to generate the radiance spectrum through a simple matrix multiplication.

Given a knowledge of the background error covariance of the forecast model that provides the initial guess and a knowledge of the observation and forward model error covariance a variational data assimilation technique can be used to minimise the cost function and find the best atmospheric and surface state that matches the observations. For speed there is no need to deal in radiance space at all, the observation radiance spectra is projected on to the EOF set to deliver a set of PC scores which are truncated and compared to the PC scores computed by the HT-FRTC code for the first guess profile. The HT-FRTC simultaneously computes the Jacobians of the PC scores which identify the sensitivity (expressed as a derivative) of the PC scores to changing atmospheric state. The variational data assimilation technique adjusts the first guess iteratively to find the solution that best matches the observed PC scores. The outputs from the variational data assimilation technique are profiles of temperature, water vapour and ozone, surface temperature and spectrally varying emissivity.

HOW TO DEAL WITH SPECTRALLY VARYING EMISSIVITY

The infrared surface emissivity of natural and anthropogenic materials is highly variable showing considerable variation in magnitude and spectral structure. The advantage of this is that materials may be identified by their IR emissivity. However in understanding the transport of radiation through the atmosphere to the sensor it is necessary to understand the fine spectral detail of the emitted radiance. For a hyperspectral sensor, which may have 1000's of channels it is exhaustive to explicitly deal with independent emissivity values for each wavelength of interest. Conducting an eigenvalue analysis of a representative set of emissivity spectra has shown that the full detail of the spectral variation of emissivity can be accurately represented using only 15 Principal Components which significantly simplifies the problem. In the data assimilation technique described previously the retrieval scheme returns values for these 15 PC scores of emissivity and these can then be reconstructed to deliver an emissivity spectrum.

AIRBORNE RETRIEVAL RESULTS

Data gathered over Oklahoma in the USA have been utilised in this study. The radiance spectra were gathered with the Airborne Research Interferometer Evaluation System (ARIES) that is operated on the Facility for Airborne Atmospheric Measurements (FAAM) BAe 46-301 aircraft. The FAAM BAe146 aircraft is jointly run by the Met Office and the

Natural Environment Research Council. In addition to the naturally occurring spectra some

synthetic radiance measurements were computed for testing purposes employing the forward component of the HTFRTC code. All of these synthetic radiances were calculated using the U.S. standard atmosphere with a surface temperature of 288K. Furthermore, it was assumed that the measurements were taken at a height of 321 hPa (≈ 9km), which corresponds to the height at which the ARIES radiances were measured. Moreover, all the retrievals in this report were done using a Met Office Unified Model profile with a surface temperature of 289.2 K as an a priori guess. Our emissivity database contains 321 surface emissivity profiles which have been taken from the JPL ASTER and the MODIS UCSB spectral libraries. The profiles have been interpolated onto a common wavenumber grid spanning the range from 810-2710 cm^{-1} with a resolution of 1 cm^{-1}. In a first test we investigated whether our scheme is capable of distinguishing between the natural background and artificial surfaces consisting of a single surface type. By artificial, we basically mean man-made surface types such as paint or metallic surfaces which have very distinct spectra from natural surfaces. This experiment was performed by inserting synthetic radiance measurements, based on the surface emissivities of aluminium and aluminium paint, into the first 210 ARIES measurements obtained during the flight over Oklahoma (19/04/07). It was assumed that the ARIES measurements represent the natural background surface while the synthetic measurements represent some man-made structure. Once the simulated radiances were introduced into the ARIES dataset, the first 210 surface emissivity profiles were retrieved with an a priori surface emissivity of bare soil. Only a selected subset of the ARIES channels, namely those between 810 cm^{-1} -2710 cm^{-1} has been used. The channels between 2710 cm^{-1} and 3000 cm^{-1} have been neglected because of too small a signal to noise ratio while the channels between 500 cm^{-1} and 810 cm^{-1} have been neglected since the surface emissivity spectra in our database do not cover this range. The results are displayed in figures 1 and 2. Figure 1 shows the first 210 retrieved spectra as a function of wavenumber, the red line denoting the averaged retrieved background surface emissivity. Figure 2 shows the same information in form of a contour plot which displays the emissivity spectra as a function of the distance travelled. The dark lines in the contour plot clearly indicate where the surface type changes from the background to aluminium/aluminium paint and back to the background. By plotting these ’dark’ spectra as a function of wavenumber the two lower emissivity spectra (with an emissivity of about 0.4 and 0.1) in figure 1 were obtained. By comparing these three retrieved surface emissivities, i.e. the two lower curves and the red curve in figure 1 to those in our database we were able to identify them as aluminium (at 0.1), as aluminium paint (at 0.4) and as some form of soil. The retrieved soil emissivity spectrum is different from the emissivity spectrum that was used as an a-priori guess.

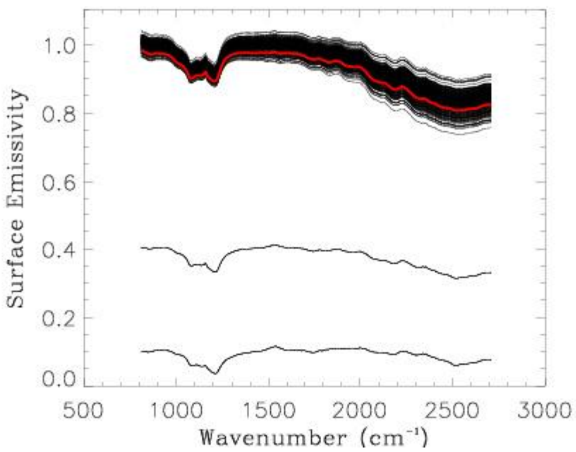

Figure 1 210 retrieved surface emissivity spectra as function of wavenumber. The red line shows the a averaged surface emissivity spectrum.

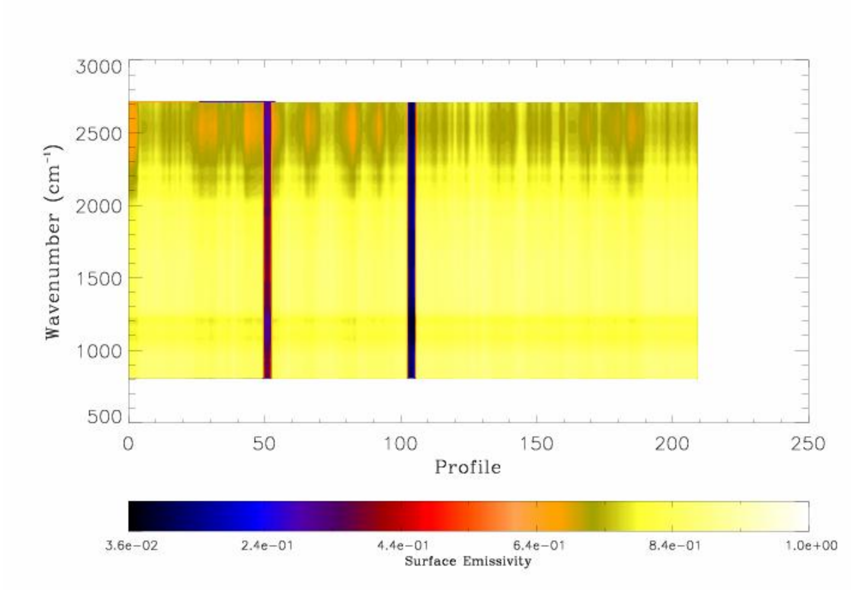

Figure 2 210 retrieved surface emissivity spectra shown in the form of a contour plot, the two synthetic anthropogenic spectra are clearly visible.

MIXED SCENE RETRIEVALS

In a second experiment we tested whether it is possible to retrieve mixed surface emissivities and to identify their constituent components. To this end synthetic radiance measurements based on the surface emissivities of sand, a mixture of sand (50%) and granite (50%) and a mixture of sand (80%) and olive paint (20%) were generated. We could, for example, imagine the scenario depicted in figure 3 where some camouflaged man-made structure is hidden in an area surrounded by sand and granite.

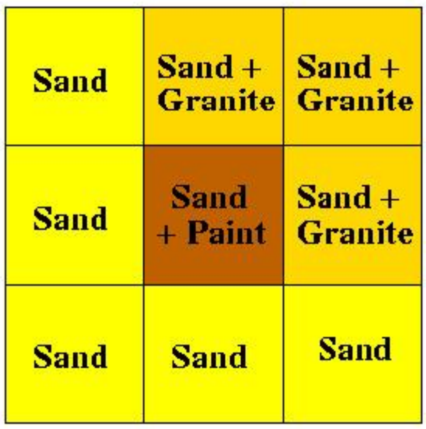

Figure 3 The artificial scene used to demonstrate the possibility of retrieving and identifying mixtures of two different surfaces.

Furthermore, let us imagine that some sensor has provided us with 9 radiance spectra, each spectra having a given footprint as shown in figure 3. The question we would like to answer is whether our scheme is capable of identifying the camouflaged object from the radiances alone? Theoretically this should be possible as long as the spectral signature of the sand/paint mixture is sufficiently different from the spectral signature of the surrounding background. Practically, however, this is not as straightforward, as the various surface types can vary slightly in their emissivity profile. For example, our database of surface emissivities contains nearly 50 types of soil, all of which have slightly different emissivity spectra. Moreover, it is highly unlikely that the retrieval scheme provides us with a spectrum that matches the spectrum of the sand/paint mix exactly, making the identification much more difficult. These problems will be discussed again at a later stage. Several assumptions have been made in this experiment. Firstly, it was assumed that nothing is known about the underlying background surface. Secondly, it was assumed that each square in figure 3 only contains two different surface types. This is done mainly for convenience as the computational cost of running the ’identification’ algorithm increases rapidly with each additional surface type. However, as this report constitutes a proof of concept it was felt that this was not a huge restriction. The ’identification’ algorithm will be discussed again at a later stage. Finally, the surface temperature within each square is assumed to be constant although different squares may have different surface temperatures. This assumption is not very realistic since natural surfaces, such as grass, and artificial surfaces, as for example runways or trucks, will have very different surface temperatures for most of the day. Unfortunately, this limitation is hardwired into our scheme since our radiative transfer code is one-dimensional and therefore only a single value can be specified for the surface temperature in each box. Figure 4 displays the surface emissivities which were retrieved from the above three synthetic radiance measurements.

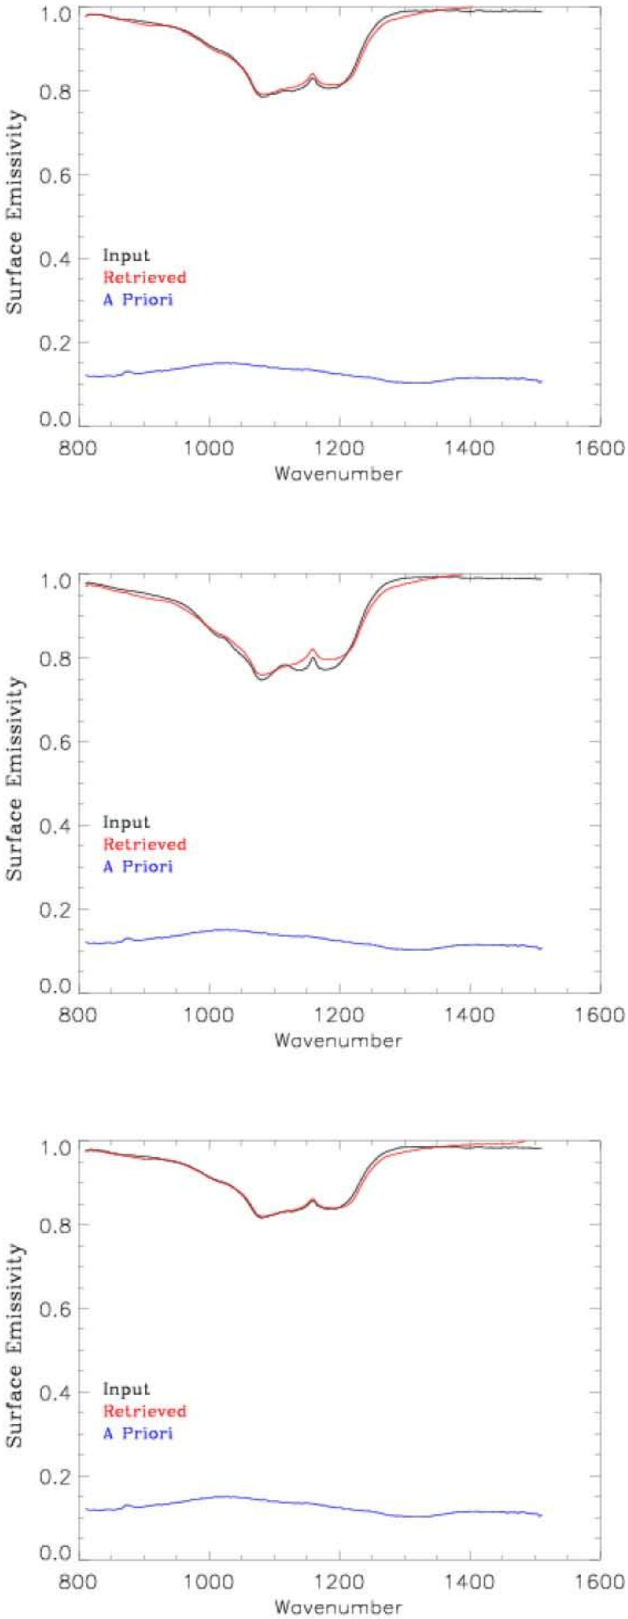

Figure 4 **Retrieved surface emissivities for sand (top), mix of 50% sand and 50% granite (middle) and a mix of 80% sand and 20% olive paint (bottom).**

The top graph shows the retrieved emissivity for sand, while the middle and the bottom graphs show the retrieved emissivities for the sand/granite and sand/paint mixtures respectively. In these figures, the red line denotes the retrieved surface emissivity, the blue line denotes the a priori guess and the black line denotes the 'truth', i.e. the actual surface emissivity that was used in computing the synthetic radiances. Only those channels lying inside the atmospheric window, i.e. between 8 µm and 13 µm (770 cm^{-1}-1250 cm^{-1}), are shown since outside this region the radiances contain little information about the surface. In general, the agreement between the retrieved and the actual surface emissivities is very good, with the largest discrepancy occurring for the sand/granite mixture. In this case the retrieval fails to capture the small scale features within the atmospheric window which suggest that our scheme fails to adjust the higher order principal components. This in turn implies that our error covariance matrix (B-matrix) constrains the higher order principal components too tightly. We now try to identify the components of the retrieved surface emissivity. We compute the distance, **d**, between the retrieved surface emissivity and all the combinations in our database.

$$d = \sum_{i}^{N} \sqrt{\left(x_i - \hat{x}_i\right)^2}$$

Here **N** denotes the number of frequencies at which d is calculated, **x** denotes the retrieved surface emissivity and $\hat{x}$ the surface emissivities in the database. The mixture with the smallest distance is assumed to correspond to the retrieved surface emissivity and hence to characterise the underlying surface. Some care needs to be taken here. As it was mentioned earlier it is highly unlikely that our scheme retrieves the 'true' surface emissivity. Thus if the distance between an emissivity spectrum in the database and the retrieved one were to be zero it would not necessarily imply that this corresponds the true underlying surface type. A better approach might be to compute the error bars for the retrieved emissivity and to check which surface emissivities within our database fall within the error bars. This approach is being tested. The components for these ten closest matches are listed in table 1. Although, we are not able to identify the retrieved emissivity spectra correctly as 20% olive paint and 80% sand, two out of the ten closest matches suggest that the surface is made up out of 20% black paint and 80% sand and five out of the ten matches suggest that the surface consist of a mixture of artificial and natural surfaces. Moreover, in all the cases 'sand' has been identified as one of the components (assuming that terra cotta tiles contain sand). For the sand/granite (see table 2), on the other hand, only one of the ten best matches suggest that the surface consists of a mixture of artificial/natural surfaces. A similar result is obtained in the case of pure sand. In this case our scheme fails to identify the surface as pure sand. However, the only artificial/natural mixture that comes up in the ten best matches is construction concrete and sand, which is maybe not too surprising as concrete is composed of cement as well as other cementitious materials such as fly ash and slag cement and various aggregates which generally include sand (see Wikipedia - http://en.wikipedia.org/wiki/Concret). Thus, our scheme has some skill in distinguishing between mixtures of natural surfaces and mixtures of natural and certain artificial surfaces, such as paint or metallic surfaces, even though we are not able to identify the mixture exactly. However, from this example it is also evident that problems would arise if we were trying to identify concrete structures in a sandy environment.

Distance	Surface Type
0.0607622	Construction Asphalt(40%) and Sand(60%)
0.0647122	Salty Soil(30%) and Sand(70%)
0.0745404	Copper Metal(30%) and Sand(70%)
0.0751578	Powder Sample(40%) and Sand(60%)
0.0775418	Sand(40%) and Soil (60%)
0.0775681	Flat Black Paint(20%) and Sand(80%)
0.0775681	Flat Black Paint(20%) and Sand(80%)
0.0805284	Sandstone(40%) and Sand(60%)
0.0837398	Construction Concrete(40%) and Sand(60%)
0.0837398	Granite(30%) and Terra cotta Tiles(70%)

Table 1: 10 mixtures that provide the closest match to the retrieved surface emissivity in the case of 20% olive paint and 80% sand.

Distance	Surface Type
0.115353	Soil(50%) and Alkalic Granite(50%)
0.119375	Flagstone(30%) Alkalic Granite(70%)
0.121467	Construction Concrete(50%) and Granite(50%)
0.121779	Monzonite(40%) and Sand(60%)
0.122726	Soil(50%) and Monzonite(50%)
0.124413	Soil(50%) and Monzonite(50%)
0.125505	Soil(30%) and Alkalic Granite(70%)
0.126394	Soil(30%) and Alkalic Granite(70%)
0.128253	Sterling ST-620 Sand Tile(20%) and Sand(80%)
0.128345	Common California Tile(20%) and Sand(80%)

Table 2: 10 mixtures that provide the closest match to the retrieved surface emissivity in the case of 50% sand and 50% granite.

ASSESSING THE IMPACT OF USING FEWER CHANNELS

In the above test cases roughly 5000 channels, spaced equally over the spectral interval 500 cm^{-1} - 3000cm^{-1}, were used to retrieve and identify the underlying surface. However, for most military and commercial sensors currently in use the number of channels does in general not surpass a few hundreds. In this section we will briefly investigate how the degradation of the spectral resolution affects our retrieval scheme. Two numerical experiments were performed to this effect. In the first experiment the number of channels was reduced from 5000 channels to 250 channels and then further to 50 channels. In the second example all channels outside the atmospheric window (800 cm^{-1} − 1400cm^{-1}) were switched off and within the window the number of channels was decreased to 170 channels respectively to 17 channels. In both experiments the channels were equally spaced over the respective spectral interval. The results from these two numerical tests are displayed in figure 5, the dashed red line indicating the relative error corresponding to the 5000 channel retrieval. Here it was assumed that the ARIES instrument error covariance matrix, R, remains correct even if the number of channels is reduced. This implies that the R-matrix is the correct instrument covariance matrix for our 'imaginary' sensors with 250,170, 50 and 17 channels respectively. This is obviously not the case as the instrument noise characteristics for ARIES are probably not the same as those of a 17 channel sensor.Thus by doing these experiments we assume that changing the R matrix has only a small effect on the retrieved solution. This is not a unreasonable assumption since the retrieved solution should ideally be independent of the the instrument noise characteristics. The most important thing to notice in figure 5 is that, as long as some channels are located within the atmospheric window, the effects of degrading the number of channels on the retrieved solution is very small. In all the cases considered here the retrieved solutions lie within the error bars of the full 5000 channel solution. Moreover, the magnitude and the spectral features of the surface emissivity are retained, even for as little as 17 channels. This is a fortunate consequence of using Empirical Orthogonal Functions (EOF) in order to retrieve the surface emissivity and would not be possible if the surface emissivity was retrieved directly. In our case the surface emissivity is obtained from a linear combination of the EOFs which contain all the relevant spectral information. Thus as long as we are able to determine the principal component scores, i.e. the linear coefficients, correctly the linear combination of EOFs will result in a retrieved emissivity spectrum with the correct spectral features. The above 50 channel retrieval suggest that even 10 or 12 channels located within the atmospheric window will be enough to achieve this. Reducing the number of channels will, not surprisingly, have an impact on the 'identification' algorithm, as the algorithm is very sensitive to the shape and the magnitude of the retrieved solution. In the case considered here, only the full 5000 channel retrieval, the 250 channel and, paradoxically, the 17 channel retrieval are able to identify the underlying surface as a mixture of artificial and natural surfaces. These three cases also suggest that the surface consists of a mixture of paint and sand. The retrieval obtained from the other three cases mainly suggest natural-natural mixtures. The fact the identification fails in the latter three cases is not particularly surprising as these are also the solutions that have moved furthest away from our synthetic 'true' solution. This suggests the need to develop an 'identification' algorithm that is less sensitive to changes in the retrieval scheme. It is important to note however that as one moves to fewer spectral channels then the atmospheric correction elements of the retrieval become more reliant on the background profile – in other words if you have fewer channels you have less information about the temperature, water vapour and aerosol profile of the atmosphere in your measurements and hence you are more reliant on the accuracy of the NWP forecast which provides the first guess to the retrieval scheme.

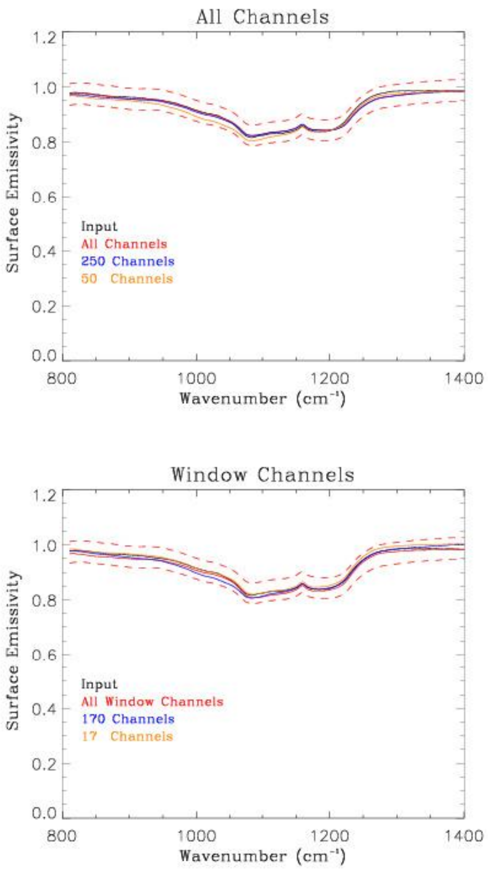

Figure 5 Retrieved emissivities for mix of 80% sand and 20% olive paint. Channels spaced evenly between 500cm-1 and 3000cm-1 (top), channels spread equally over the atmospheric window region 800cm-1 to 1400cm-1only.

EMISSIVITY RETRIEVALS FROM SPACE

The European Metop satellite was launched in November 2006 as part of a global network of satellites used for meteorological observation. The Infared Atmospheric Sounding Interferometer (IASI) is one of the suite of instruments on Metop. IASI measures infrared radiances between 3 and 15 microns with 8461 spectral channels. The data is widely used by NWP centres around the world to aid in weather prediction. Currently use of the data is very conservative and is restricted over land to the use of channels that are insensitive to the land surface because of the complexity of the land surface emissivity. We have utilised the techniques described here for airborne observations to try and make greater use of the IASI data by allowing the data assimilation technique to explicitly retrieve the surface emissivity. Figure 6 shows initial results for the surface emissivity over Oklahoma retrieved from space based IASI observations. The emissivity retrievals compare well with independent observations of emissivity made by University of Wisconsin who have built up a land atlas of emissivity and from low level measurements made by the FAAM BAe146 aircraft (not shown).

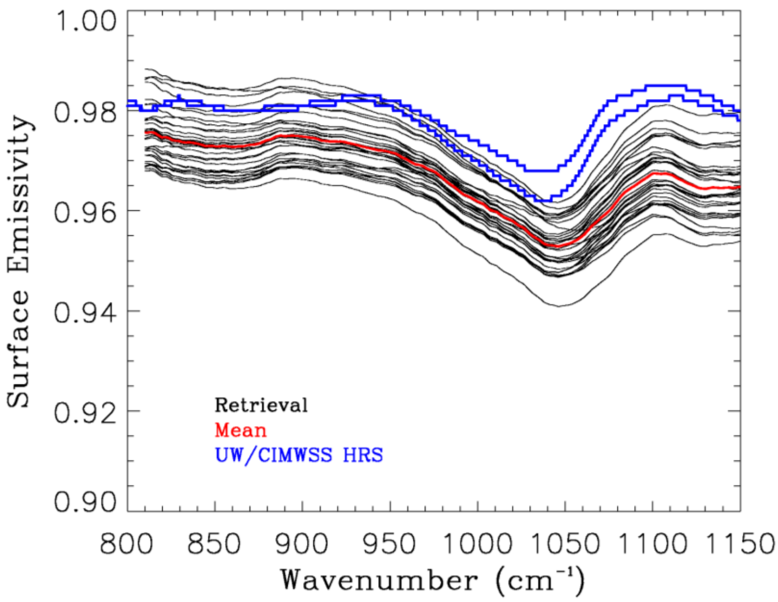

Figure 6 Emissivity retrieval from IASI (black lines) compared with independent land atlas (blue lines).

SUMMARY

This report has shown that hyperspectral radiance spectra contain significant details on the atmospheric state as well as surface properties. Furthermore with by utilising a variational assimilation technique and a state of the art Numerical Weather Prediction model it is possible to retrieve spectrally varying surface emissivity spectra as well as conduct atmospheric correction. This technique relies on a good understanding of the error covariance of the observations and the Numerical Weather Prediction model used to deliver the first guess.

In this report it has been shown, firstly, that our scheme is capable of accurately retrieving the surface emissivity of both pure and mixed surface types and, secondly, that the retrieval is not very sensitive to the number of channels used. Strictly speaking, about ten channels, spread equally over the atmospheric window, will be sufficient to retrieve the surface emissivity with a good degree of accuracy. Moreover, our retrieval/identification scheme also demonstrated some skill in determining whether the underlying surface is made up of a mixture of natural surfaces or whether is consists of a mixture of natural and man-made surfaces. It has also been shown that it is very difficult to identify the underlying surface exactly since the retrieved surface emissivity will, in general, not correspond exactly to the underlying surface. The best possible outcome is that our ’identification’ algorithm comes up with a mixture of surface types similar to that of the true surface. At the moment, this algorithm works best if the spectral features of the man-made surfaces are very distinct from their natural counterparts. Moreover, the identification algorithm is not very robust in the sense that it is very sensitive to changes made to the retrieval algorithm such as the reduction in the number of channels. These results are very encouraging since the technique proposed here allows us to accurately retrieve surface emissivities and also shows some skill in identifying the constituent parts of mixtures of surface types. Further work with the Defence Science and Technology Laboratory aims to apply more sophisticated scene identification techniques.