Atmospheric Invariants for Hyperspectral Image Correction

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

Abstract

The degrading effect of the atmosphere on hyperspectral imagery has long been recognised as a major issue in applying techniques such as spectrally-matched filters to hyperspectral data. There are a number of algorithms available in the literature for the correction of hyperspectral data. However, most of these approaches rely either on identifying objects within a scene (e.g. water whose spectral characteristics are known) or by measuring the relative effects of certain absorption features and using this to construct a model of the atmosphere which can then be used to correct the image. In the work presented here,an alternative approach is reported which makes use of the fact that the effective number of degrees of freedom in the atmosphere (transmission, path radiance and down-welling radiance with respect to wavelength) is often substantially less than the number of degrees of freedom in the spectra of interest. This allows the definition of a fixed set of invariant features (which may be linear or non-linear) from which reflectance spectra can be approximately reconstructed, irrespective of the particular atmosphere. The technique is demonstrated on a range of data across the visible to near infra-red, mid-wave and long-wave infra-red regions, where its performance is quantified.

Keywords: Hyperspectral, Atmospheric Correction, Atmospheric Invariant

Introduction

This paper discusses an alternative approach to the conventional atmospheric correction schemes. This approach is predicated upon the observation that although the atmosphere can vary quite a lot, there is always a very high degree of correlation between different atmospheric components, and the functional forms of the relevant quantities are very highly constrained. In other words, the effective number of independent degrees of freedom available over the complete set of all possible transformations implied by atmospheric degradation is rather small. In many cases it can be considerably smaller in terms of dimensionality than the number of bands available in an imaging spectrometer and therefore it may be possible to recover information about the underlying material signature without assuming (or knowing) anything about the *particular atmosphere* through which the radiation propagated. Previous work in this area [1] has concentrated upon linear subspaces and seeking known materials, or the use of sophisticated forward models that include factors such as shadows [2].

In the work described in this paper the aim is to find a single (not necessarily linear) transformation that will map the incoming radiation into a space where it is invariant to a very wide range of atmospheres. This aim is motivated by the need to perform fast atmospheric correction onboard Uninhabited Airborne Vehicles (UAVs) without human intervention, which can be quite extensive in conventional approaches. It is recognised that the transformations used must entail information loss, since not knowing about the particular atmosphere must destroy some information. It is also likely that an automatic method cannot have such good performance as the human-in-the-loop approaches. However, in many applications these more accurate techniques are not applicable due to timelines or availability of skilled operators and processing tools. Another important factor in the approaches described in this paper is the desire to include measures of confidence along with the algorithm's output. If one uses an approximate or lossy technique, its results should come equipped with a reliable confidence measure so that downstream processing, or an operator, can know when to trust its output and when to treat it with caution. In the work described in this a paper some novel statistical approaches for supplying valid confidences or tolerance regions are discussed.

The atmospheric correction technique QUAC [3] (Quick Atmospheric Correction) is also applicable for UAV/tactical imaging spectrometers (in the visible/near infra-red band). It is also worth noting that QUAC is based upon a statistical assumption about the distribution of materials within the scene (specifically the spatial variance measured per band is approximately a constant function of wavelength for a sufficiently diverse scene), but makes no atmospheric assumptions. The approach described here makes a (limited) atmospheric assumption, but no assumptions about the scene content. Thus, the two techniques naturally complement each other.

The focus of the research conducted within this project was on vertical viewing situations, and the only atmospheric assumption in this work is that Modtran is capable of representing the envelope of possible atmospheres. That is, by generating a very large database of Modtran runs, and by creating a random set of all input parameters, one can obtain a set of forward atmospheric transforms (from reflectance to at-sensor radiance) that is reasonably complete. By this it is meant that any specific atmospheric forward transform due to a real atmosphere somewhere will be close to one of these sample points, or 'inside' their envelope. Within the scope of this project, the cases of shadowing and multiple reflections were explicitly excluded.

Monte-Carlo Modtran Runs

For this project a slightly simplified version of the radiation-transport equation was used. In general, the path radiance can depend upon the nature of the background surface in a complex and nonlinear way. However, for much of the wavelength range of interest this is very much a second-order effect and can be represented by linearly interpolating between a background that is a perfect absorber (black body) and a perfect reflector.

In the following equation the wavelength dependence of all the quantities is left as implicit. The radiation transport equation used is given by:

$$s = T(1-r)\tau + dr\tau + p_0 + (p_1 - p_0)r_b$$

where the signal at the imager is given by s, T is the radiance due to the temperature of the target (computed from Planck's law), r is the reflectance of the target, r_b is the reflectance of the background, d is the down-welling radiance, p_0 and p_1 are the path radiance values computed for a black body background and perfectly reflecting background respectively, and τ is the atmospheric transmission.

Each of the quantities in the radiation transport equation is a function of wavelength. It is the particular functional dependence on wavelength that was exploited in this work. For example, the atmospheric transmission cannot be just any function of wavelength - even if the atmosphere being described is not stipulated, the functional forms available for the transmission are quite limited. Similar constraints apply to the other quantities in this equation too, and these should allow the prediction of information about reflectance even though only the signal is observed.

In order to generate a large database of samples for use in the above equation Modtran was run through a custom interface based upon Microsoft Excel. This allowed any chosen parameters to be set via random numbers selected from a range of distributions. In order to compute the terms required above, three separate Modtran runs were required: one to determine the transmission and down-welling radiance, and two using multiple scattering to determine the path radiance for a black body and a perfect reflector. This allowed the reasonably accurate interpolation formula above to be used to compute path radiance as a function of background. The reflectance spectra were randomly sampled from a number of different databases. In addition, the reflectances were linearly mixed and had noise added to them since the primary purpose of the data was to allow the subsequent processing to learn about generic atmospheric properties rather than material properties.

Learning Inverse Mappings

The data generated in the Monte-Carlo runs described above was used as the input for a learning algorithm. The remit for this algorithm was the deduction of a relationship between the input at-sensor radiance and the truth (noisy reflectance mixture). It was decided to make use of the straightforward technique of Ridge Regression because it has a simple analytical solution, is easily extended to nonlinear mappings via the so-called Kernel-trick, and has been applied within the confidence framework described later.

Confidence Techniques

The theory of Conformal Prediction has recently been devised as a means of making predictions that automatically come equipped with a valid confidence measure [5]. Here ‘valid’ means, for example, that when a confidence of 90% is output then the probability that this output is wrong is 10% for an exactly-calibrated confidence measure and less than 10% for conservatively-calibrated confidence measures. In this context, it is the confidence region that is important since the output is a continuous multi-dimensional quantity (reflectance spectra) rather than a discrete set of classes (as it would be for example in classification). A confidence region is a (hyper-)volume (e.g. a hypersphere) centred on the output reflectance spectra with an associated confidence value such as 90%. In this case validity means that the probability of the true reflectance lying outside this sphere is less than 10%.

Most results in Conformal Prediction have been derived within an ‘on-line’ framework in which it is required to make sequential predictions, and where the truth is revealed after each prediction is made. Essentially, any prediction technique can be employed within the framework and the methods allow both for the making of predictions and the computation of confidence values which are guaranteed to be valid under certain conditions. Most notably, that the data about which the predictions are being made are all Independent and Identically-Distributed (IID) samples from some unknown distribution. The point being that it is not necessary to know *what* the underlying distribution is, merely that it exists. The major advantage of this method is that the IID assumption is the only assumption made. In contrast, to correctly apply Bayesian methods the stochastic mechanism generating the data must be known in every detail.

The on-line mode of operation where everything is updated as soon as the truth data for the last prediction is received is sometimes known as Transductive Inference. In Transductive Inference one is only concerned with answering a specific question (about the current sample), so it may not be necessary to construct a model, or if it is, the model is thrown away again afterwards and a different model constructed to make the next prediction.

In the usual atmospheric correction applications we have in mind in this project the sequential/Transductive approach is probably not appropriate since there is no automatic mechanism to tell the algorithm what the truth values are after it has made its predictions. The framework does allow for substantially delayed revelation of the truth, and not all predictions need have the truth revealed. However, there is also a more conventional Inductive version of the theory which allows progress to be made [6].

In order to apply Inductive confidence techniques to the problem at hand, the training data needs to be further partitioned into a 'proper training set' and a 'calibration set'. The main premise is to train the algorithm on the proper training set and then exercise it on the calibration set. To each member of the calibration set a 'strangenes value' α_i is associated which measures (according to any convenient metric) the discrepancy between the predicted value and the true value. Given a new sample to use as input for the next prediction and assuming *k* elements in the calibration set, it is possible to hypothesise potential output prediction values (denoted *y*) and consequently compute their strangeness values α_{k+1} using the same metric as before. Then the p-value associated with the new sample is given by:

$$p(y) = \frac{\#\{i = 1, ..., k+1 : \alpha_i \geq \alpha_{k+1}\}}{k+1}$$

where the '#' denotes a counting operation that returns the number of elements in the set. Given a desired significance level, δ, then the confidence region output at this level of significance is the set of points:

$$\{y : p(y) > \delta\}$$

The 'shape' of this region clearly depends upon the choice of metric, but it is proven in [6] that confidence regions defined in this way are valid in the sense described above.

This approach has been implemented within the algorithm described in this paper and is therefore able to compute valid tolerance regions for the recovered reflectance spectra. The standard Euclidean metric was used to compute discrepancies between truth and prediction, and hence the tolerance regions are nested hyperspheres whose dimensionality is equal to the number of bands. Although such sets can, in principle, be processed by further automated algorithms with suitable modifications, a more interesting question is how this information should be represented and displayed to humans. One straightforward approach is simply to quote the radius of the hypersphere for a given tolerance region, but this does not necessarily provide the necessary intuition to a user. An alternative method which is being explored is to generate a number of different example spectra from the uniform distribution inside the tolerance sphere so that the user can clearly see what magnitude of variability would be considered equivalent at a specific confidence level.

Results

The algorithms as described above, and trained entirely on the Modtran generated Monte-Carlo data, were used to invert some cubes downloaded from the NASA website. AVIRIS is a 224-band, visible to near infra-red sensor. The AVIRIS data also comes in atmospherically-corrected versions which can be used as truth data against which the correction method can be judged.

Five cubes from the Jasper Ridge sequence were corrected using the algorithms described above. In order to assess the benefit of the technique, the correlation of the raw radiance data with the truth reflectance was also computed and is expressed in Table 1. This table also shows the results of the correlation between the truth data and the raw radiance, and the truth data and the atmospherically recovered spectra for the Invariant algorithm. The value shown is the overall mean – averaged over the entire 614 x 512 pixels in the cubes.

The numbers in Table 1 show that the algorithm is able to correct spectra independently of the atmosphere and achieve a 60%-70% correlation with the true spectra (as obtained by more traditional human-expert in-the-loop approaches). The final column shows the standard deviation of this correlation over the cube.

Cube	**Raw Radiance Correlation with Truth**	**Atmospheric Invariants Correlation with Truth**	**σ**
SC01	0.19	0.63	0.23
SC02	0.19	0.58	0.25
SC03	0.28	0.61	0.23
SC04	0.29	0.70	0.19
SC05	0.25	0.66	0.26

Table 1: Invariant algorithm results

Table 2 shows the corresponding results for the QUAC algorithm. This indicates that whereas QUAC performs slightly better on this data it displays much more variability than the Invariants algorithm.

Cube	**QUAC correlation**	**σ**
SC01	0.77	0.39
SC02	0.81	0.33
SC03	0.79	0.31
SC04	0.81	0.28
SC05	0.85	0.26

Table 2: QUAC algorithm results

Table 3 lists the results for the fused combination of the Invariant algorithm with QUAC. This algorithm not only shows the best overall performance, but also displays the smallest variability over the pixels of the cubes.

Cube	**Fused algorithm**	**σ**
SC01	0.84	0.13
SC02	0.88	0.12
SC03	0.87	0.11
SC04	0.87	0.1
SC05	0.91	0.08

Table 3: Fused algorithm results

Conclusions

In this paper a method has been described by which reflectance spectra can be approximately recovered independently of the atmosphere through which they were observed. Recovery is by means of a fixed transform that is atmospherically independent, thereby making the technique applicable for tactical UAV applications, or application where an unskilled operator needs an automated approach to atmospheric correction. In addition, methods have been described by which valid confidence statistics can be associated with each corrected pixel so that an operator, or downstream processing algorithm, can treat the information appropriately.

The results obtained are highly encouraging and are on a par with results that can be obtained on this data using the QUAC algorithm (which is also suitable for automated use in the visible to near infra-red regime). Since these two algorithms are entirely complementary in their motivating assumptions, it was natural to combine them to produce a fused variant. This algorithm has been shown to display the best overall performance and the smallest variability on the data presented here.

There are a range of natural extensions of this work, including:

- Matched filtering
- Combination with conventional atmospheric correction techniques
- Material identification using databases
- Target detection

References

1. G. Healey, D. Slater, "Models and methods for automated material identification in hyperspectral imagery acquired under unknown illumination and atmospheric conditions", IEEE Trans. Geosci. Remote Sensing, vol. 37, No. 5, pp2706-2717, Nov. 1999.
2. Goa,P.E. Skauli,T. Kasen, I. Haavardsholm, T.V. Rodningsby, A. Physical Subspace Models for Invariant Material Identification: subspace composition and detection performance. Proc. SPIE Vol. 5573, pp203-214, 2004
3. Bernstein, L.S. et. al. A new method for atmospheric correction and aerosol optical property retrieval for VUS-SWIR multi- and hyperspectral imaging sensors: QUAC (Quick Atmospheric Correction). Proc. IGARSS 2005, Vol. 5. pp3549-3552, 2005
4. Hastie, T. Tibshirani, R.Friedman, J.H. The Elements of Statistical Learning, Springer 2001
5. Vovk, V. Gammerman, A. Shafer, G. Algorithmic learning in a random world. Springer, New York, 2005.
6. Papadopoulos, H. Proedrou, K. Vovk, V. Gammerman, A. Inductive Confidence Machines for Regression. Machine Learning: ECML 2002 (Lecture Notes in Computer Science), Vol. 2430/2002 pp185-194. 2002
7. http://aviris.jpl.nasa.gov/

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.