Spectral Super Resolution

J.D. Rogers and M. Bernhardt
Waterfall Solutions, Parklands,
Guildford, Surrey GU2 9JX.

Abstract

The utility of reconnaissance imagery can be enhanced by appropriate interpretation of its spectral information. Surveillance tasks such as recognition and identification can benefit greatly from this information, but typically employ sensors which only possess three or four spectral bands. This work explores novel processing approaches to exploit variations in atmospheric conditions and sensor characteristics to increase the spectral resolution of standard reconnaissance imagery. An analysis of the properties of spectra and of the atmospheric conditions is presented which provides an understanding of the underlying mechanism. The variation in spectral response within a sensor is shown to provide great potential for this type of technique, which is demonstrated here on real and synthetic data.

Keywords: Spectral, Super-Resolution, Surveillance, Reconnaissance

Introduction

Spectral information represents a particularly useful quantity for reconnaissance image analysis tasks such as detection, recognition and identification. Currently, in-service sensors usually employ three or four spectral bands and this situation is expected to remain for some time to come. Such limited spectral resolution inevitably constrains the information content of the imagery, and hence its utility. If imagery of higher spectral resolution could be obtained, a better interpretation of the scene being viewed could result, as this would provide more detailed spectral information. The easiest way to achieve such an outcome is obviously to add more sensors, or a hyper-spectral device, but the premise of the idea reported here is that the end result can be achieved through smart processing rather than extra sensors. It is acknowledged that enhancing the spectral information from a typical sensor is a non-trivial task, but it represents a low-cost way of obtaining additional image information and could be provided as an upgrade to in-service equipment. This paper reviews a three month feasibility study into spectral super-resolution techniques for the Electromagnetic Remote Sensing (EMRS) Defence Technology Centre (DTC).

The research reported here explored the super-resolution mechanism in two ways. The first method examined the utility of exploiting the variation in atmospheric corruption on surveillance imagery. This variation in atmospheric conditions can alter the spectral content of imagery in different ways because the atmosphere acts as a spectral filter which corrupts the original spectra. When performing strategic reconnaissance it is often the case that multiple views of the same scene will be gathered at various points in time and through different atmospheres. Therefore, this study examined the utility of extracting further spectral information from these multiple views and analysed the role of spectral smoothness properties and atmospheric variation.

The second approach analysed the variation in spectral response within a sensor array. Similar to the first part of this research, this activity exploited the variation in spectral information from different views. However, instead of employing the variation in atmospheric conditions, the variation in spectral response within a single sensor was examined, along with how this variation could be combined. This concept was explored by simulating the role of a sensor and reducing the spectral information of the original imagery.

The first part of this paper discusses the case of exploiting variations in atmospheric conditions to perform super-resolution, and overviews the investigation into the underlying mechanics of these methods. The second part of the paper summarises research that was conducted into performing super-resolution using the variation in spectral response within an imaging array and presents the results of this technique.

Exploiting Variation in Atmospheric Conditions

Using a library of real spectra for different materials and a set of real atmospheres from the Modtran database, a data set was created to which a transform could be applied. The objective of this transform was to re-create an original set of high-resolution spectra from a collection of lower-resolution spectra that had been corrupted by different atmospheres. A regression technique was employed to learn the transform and a kernel method was used to create a non-linear mapping and identify more complex relationships between the atmospherically corrupted views and their higher-resolution counterparts.

The Role of Spectral Smoothness

When learning a mapping to increase spectral resolution, it is important to consider the underlying distribution of the higher-resolution data. If the distribution is relatively smooth, it can also be viewed as being predictable and this aids the task of constructing a super-resolution transform because it will be implicitly learned by the regression mechanism.

The role of the smoothness properties within the original spectra was examined in two ways. Firstly, in addition to the non-linear regression methods that were used to learn the transform, a simpler method was also applied which performed a linear interpolation of a single lower resolution view to create the higher resolution result. This formed a very simple model that was not expected to perform well if the underlying spectra are complicated to predict. The second method of examination replaced the target spectra with random values to create a distribution that was very hard to learn. In this case, the transform was not aided by the distribution and was restricted to learning the inverse of the atmospheric corruption.

The data set was randomly partitioned into a training set for constructing the transform and a test set for its evaluation. This was repeated over 100 trials and the averaged results are shown in Table 1. The performance metric used here was the average correlation between the predicted spectra and the truth super-resolution values. It can be seen that applying linear interpolation to the normal material spectra performed surprisingly well considering the simplicity of the method, which suggests that the underlying spectra are relatively smooth and predictable. Applying non-linear regression to the data gave significantly better results than the interpolation, although the benefit of employing multiple views was marginal. Supplying both views to the regression model only provided an increase in average correlation of approximately 0.02 over the single view case, and the benefit of averaging the two predictions was even less significant.

From examination of the average correlation values obtained when random spectra were used, it can be seen that the regression performance was dramatically reduced for all of the methods. The benefit that was gained through combining the multiple views when random spectra were applied is noticeable, but the smoothness properties of real spectra are clearly a much more important factor for determining the success of this type of super-resolution transform.

	Normal Spectra	Random Spectra
Linear Interpolation	0.7186	-0.1470
Single View	0.8617	0.2313
Averaged Single View	0.8667	0.2444
Combined Views	0.8853	0.3034

Table 1 – Average correlation between predicted spectra and the truth for different transform methods on normal and random spectra

The Role of Atmospheric Variability

The variation between different atmospheric conditions is another factor which can affect this technique. If atmospheres are very similar, the way in which the material spectra are filtered will also be similar and the learning of a super-resolution transform will become easier. However, the benefit that is gained from having multiple views of a scene is related to the diversity of these views.

To explore this idea, the atmospheric parameters were generated in a manner which allowed their variation to be controlled. The covariance matrix of the atmospheric parameters can be decomposed through singular value decomposition to yield an alternative representation, as shown in Equation 1 .

$$\Sigma_\alpha = USU^T \qquad [1]$$

Here, Σ_α represents the covariance matrix of the atmospheric parameters, U is a matrix consisting of orthogonal vectors, and S is a diagonal matrix containing the singular values. The number of non-zero singular values dictates the dimensionality of the subspace within which the values of the atmospheric parameters can exist. The size and number of these non-zero singular values then directly affects the variation of the resultant atmospheres. Using this representation, the diagonal elements of S were replaced by the expression in Equation 2, where λ was a parameter that was used to control the atmospheric variation.

$$S_{ii} = e^{-\frac{i}{\lambda}} \qquad [2]$$

Sets of atmospheric parameters were then generated for different values of λ using Equation 3, where r is a random unit vector.

$$\alpha = rU\sqrt{S}U^T \qquad [3]$$

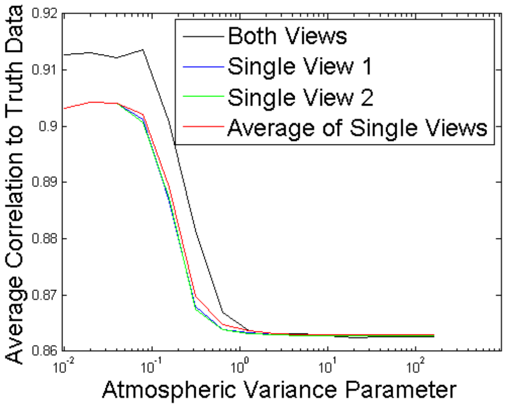

Figure 1 - Super-resolution performance for varying levels of atmospheric variation

Figure 1 illustrates the performance of the different techniques for various values of the variance parameter λ. From the figure, it can be seen that increasing the value of λ reduces the performance of the super-resolution methods. A greater variation in atmospheric corruption results in a more difficult task for the super-resolution techniques, which causes this reduction in performance. Again, the benefit of combining the multiple views is present but is only marginal when compared to the single view case.

When Atmospheres are Known

The experiments reported thus far all constructed a super-resolution transform using spectra that had been corrupted through a variety of different atmospheres, for which no knowledge was available. The next experiment constructed a transform in the same manner as before, but maintained the same two atmospheres throughout. In this case, the transform was able to learn the particular corruption that was produced by these two atmospheres. In doing so, lessons were learned regarding the ability of the super-resolution transforms when knowledge of atmospheric conditions is available. Table 2 illustrates the correlation performance of this technique on normal and the synthetic random material spectra.

	Normal Spectra	Random Spectra
Single View	0.9507	0.6963
Multiple Views	0.9984	0.8020

Table 2 - Correlation performance for normal and random spectra when atmospheric corruption is known

The benefit of combining the two views can be seen to be very significant for both normal and random spectra. This clearly demonstrates that good super-resolution is possible if the atmospheres are known, but it is acknowledged that this is less likely to be the case in practise. In the typical case of the atmospheres being unknown, the benefit is not significant.

Exploiting Variation in Spectral Response of Imaging Sensor

Another application for a super-resolution transform is the exploitation of the variation in spectral response within a sensor array. It was envisaged that if the variation is sufficient and can be learned, this information could lead to the extraction of further spectral information. If multiple images of the same scene can be obtained using the same sensor but with variations in sensor position and orientation, parts of the scene will have been viewed by several different parts of the sensor array. Therefore, the potential exists for combining these different views of the same scene to produce a single view with a greater amount of spectral information.

To test the efficacy of this type of method, standard-resolution imagery was converted into lower-resolution imagery in a manner which reflects a variation in spectral response for the different pixels. A super-resolution transform was then applied to this data in an attempt to reconstruct the original imagery, and the results were then compared to the original image for evaluation. Here, standard RGB imagery was used and converted to multiple grey-scale images. In standard colour to grey-scale conversion, the three colour bands are averaged using standard weighting coefficients. Here, these coefficients were replaced with random values to synthesise the variation in spectral response and this was repeated twice to create two different views of the scene. Colour images were then converted by these two synthetic sensors into two grey-scale images. Each pixel of the original image was then represented by two grey-scale versions and the objective of the super-resolution transform was to recover the original colour pixels from these two grey-scale views.

The left panel of Figure 2 displays a test image of a colour grid and one of the grey-scale versions into which it was converted.

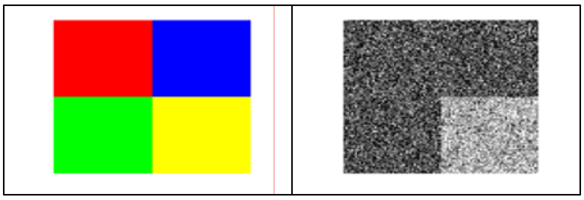

Figure 2 - Colour grid example image (left) and random grey-scale conversion (right)

Cross-validation was applied to this data in order to construct a super-resolution transform and predict the original colour values. Cross-validation ensured that the data being predicted was not the same as that used to construct the transform. As before, non-linear regression techniques were used to train the transform and the input to the transform consisted of the grey-scale value of each pixel along with its corresponding conversion coefficients.

Figure 3 depicts the reconstruction that resulted from using only one of the grey-scale images, and of applying both views. The benefit of using multiple views is demonstrated with a better colour depiction in the right panel, which is particularly noticeable in the red quadrant. The incorporation of a greater number of views is anticipated to improve this further.

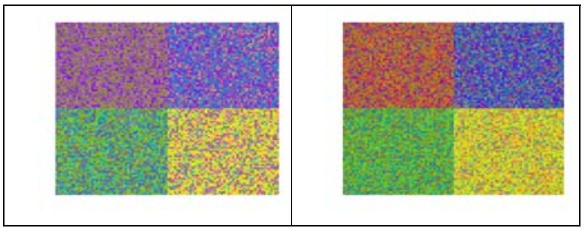

Figure 3 - Reconstructed colour image from single grey-scale image (left) and multiple grey-scale images (right)

The technique was also applied to 3-band imagery from Google-Earth and 4-band imagery from the IKONOS Earth Observation Satellite. The average correlation of the predicted spectral values to the known truth is presented in Table 3.

	Single View	Multiple Views
Colour Grid	0.5897	0.7670
Google Image	0.7602	0.8226
IKONOS Imagery	0.6083	0.7354

Table 3 - Correlation performance of reconstructed colour imagery from grey-scale views

For all of the images, the benefit of applying multiple views is significant. However, the performance of a single-view model can be seen to be reasonably good. This is because although the exact pixels being predicted were never used for training the transform, the training pixels came from the same image, and hence had a similar distribution. Therefore, the transform was able to implicitly learn the underlying distributions of the images.

The previous experiments used random values instead of the standard grey-scale conversion coefficients. The actual variation in spectral response within a sensor array is unknown, and so the objective here was to examine the role of this variation. This method employed the standard conversion weights, but with extra values added. These extra values were selected randomly from a uniform distribution in the range $[0,\delta]$. This caused a variation from the standard coefficients and the value of δ determined the level of this variation. As mentioned previously, the super-resolution transform can implicitly learn the distribution of the images which aids the transform construction. The experimental method here used randomly generated 3-band imagery to ensure that the image distribution does not help the transform. The same process as before was applied to the random image and the level of spectral variation was varied by altering the value of δ. The correlation performance was then measured and shown in Figure 4.

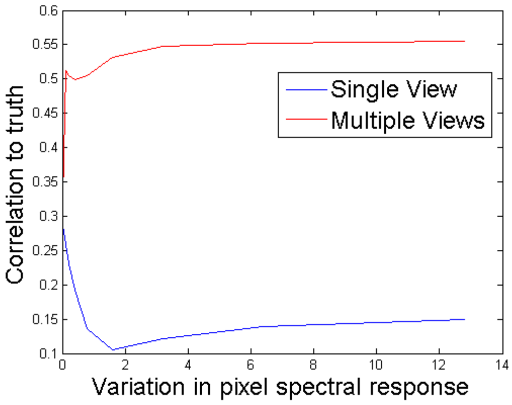

Figure 4 - Plots showing correlation performance of 3-band image reconstruction using different levels of spectral variation for single and multiple views

The figure clearly shows that the transform performed relatively well for multiple views, even though it was not aided by the image distribution. The advantage of combining multiple views is evident for even very small values of δ. This is supported by the fact that the average correlation of the reconstructed image from two views is 0.513 when the value of δ is fixed at 2.6 x 10^{-6}.

Discussion

The first part of this study explored the concept of using multiple views of the same scene taken through different atmospheres to generate a single view with increased spectral content. The application of non-linear regression techniques to this type of data yielded promising results. However, after careful examination of the super-resolution mechanism, it was discovered that the reasons for this good performance were related to the smoothness properties of the material spectra and of the level of atmospheric variation. The smooth and predictable qualities of material spectra enable a good level of regression performance to be achieved from a single view through a single atmosphere. The benefit of employing a second view was found to be only marginal.

The role of the atmospheric variation was also investigated through the generation of synthetic atmospheres from a manipulated covariance matrix. Raising the level of atmospheric variation was shown to reduce the regression performance and could be seen to have a larger impact than combining multiple views.

The reason for achieving only a marginal benefit in combining multiple views through several atmospheres is due to the fact that the atmospheres are unknown. It was shown that if the data is constructed with known atmospheres, the transform can learn the corruption more effectively. The benefit in this case is therefore much greater, although in practise this atmospheric information is less likely to be available.

The second part of the study concerned exploiting the variation of spectral response within an imaging array to perform spectral super-resolution. Synthetic data was created by converting RGB imagery into grey-scale imagery, and then this was used to learn a super-resolution transform capable of recovering the original image. This was found to be successful, and a significant benefit was noted when multiple views were combined. It is believed that if the variation within an imaging array is sufficient and can be learned, this type of technique can increase the spectral information of standard imagery.

The simulated variation in spectral response was shown to enable the extraction of further spectral information. Further work is required to establish the degree of variation within a real imaging array and whether it is sufficient to enable a practical application of this type of technique. If so, it is envisaged that a super-resolution transform, that is specific to the sensor, can be learned under laboratory conditions. Once the transform has been learned, it can be applied to any imagery that is viewed by the sensor and used to increase the spectral information. The imagery would first be registered to establish which components of the scene have been viewed by which parts of the sensor. The transform could then be applied to the observed responses, and the appropriate learned sensor characteristics, to yield a super-resolution output.

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.