

On Automatic Absorption Detection for Imaging Spectroscopy: A Comparative Study

Zhouyu Fu Antonio Robles-Kelly Terry Caelli Robby T. Tan
NICTA*, RSISE Bldg. 115, ANU, Canberra, ACT 0200, Australia

Abstract

In this paper, we aim at presenting a survey on automatic absorption recovery methods for imaging spectroscopy. We commence by viewing the algorithms in the literature from a technical perspective and presenting an overview of the derivative analysis, Fingerprint and Maximum Modulus Wavelet Transform (MMWT) techniques. In addition to these methods, we also present a novel absorption recovery approach based upon unimodal regression and continuum removal. With this technical review of the methods under study at hand, we perform a complexity analysis and examine the implementation issues pertaining each of the alternatives. We show how detected absorption bands can be used for purposes of material identification. We conclude the paper by providing a performance study and providing identification results on hyperspectral imagery. To this end, we make use of a number of distance measures to evaluate the quality of the recovered absorptions as compared to continuum-removed spectra.

Keywords: Absorption features, Fingerprint, Maximum Modulus Wavelet Transform, unimodal segmentation, hyperspectral image processing.

1 Introduction

With the advent and development of new sensor technologies, it is now possible to capture image data in tens or hundreds of wavelength-resolved bands covering a broad spectral range.

*National ICT Australia is funded by the Australian Governments Backing Australia's Ability initiative, in part through the Australian Research Council.

Compared to traditional monochrome and trichromatic cameras, hyperspectral image sensors provide an information-rich representation of the spectral response for the material under study over a number of wavelengths. This has opened-up great opportunities and posed important challenges on material identification and classification. Due to the high dimensional nature of the spectral data, many classical algorithms in pattern recognition and machine learning have been naturally borrowed and adapted so as to perform feature extraction and classification [22]. Techniques such as Principle Component Analysis (PCA) [19], Linear Discriminant Analysis (LDA)[14], Projection Pursuit [18] and their kernel versions [12] treat raw pixel spectra as input vectors in a higher-dimensional space, where the dimensionality is given by the number of bands. The idea is to recover statistically optimal solutions to the classification problem in hand by reducing the dimensionality via a projection of the feature space.

From an alternative viewpoint, hyperspectral imaging and material chemistry can be used for purposes of material identification. This hinges on the notion that different materials have a characteristic response as a function of wavelength. A proof of concept is provided by Clark *et al.* [8], who showed that an unknown spectrum can be identified efficiently by simply matching the dominant absorption feature in its reflectance spectrum with those present in a spectral library so as to recover the best fit. The reason is that reflections and absorptions are two complimentary concepts of light behavior. When light is incident on the surface of a material, some of it is reflected, some is absorbed, and, in certain cases, some is transmitted. The problem with absorptions is that they are less straightforward to recover compared to reflection or reflectance which can be directly measured by image sensors. Nevertheless, absorptions are closely related to the chemistry of the material and, therefore, the presence of an absorption band at a certain spectral range is often a "signature" for chemicals and their concentrations. Thus, compared to the dominant statistical approaches, the use of absorption features for material identification has a clear physical meaning. Furthermore, this principle has been successfully applied in systems such as the Tetracorder [8]. The Tetracorder is a software tool for imaging spectroscopy that can be regarded as a collection of algorithms that permits the analysis of conditions in the spectra.

Despite the success of the Tetracorder system, it is still a semi-automatic one. A few typical spectra are selected for reference and their absorption bands are manually labelled in order to perform matching tasks against test spectra. This process involves non-trivial

input of expert knowledge. The Tetracorder is based upon spectral feature identification algorithms such as that in [9], where a least-squares fit is used to match the spectrum under study to that in a reference library. As a result, the system performance is dependent on the selected reference spectral bands and purpose-labelled absorptions.

Therefore, it is not only theoretically important, but practically useful to recover absorption features automatically. This is even more relevant since a number of absorption features, such as that of water, chlorophyll and some gases find application in a wide variety of areas in geosciences. For instance, chlorophyll absorption (between approximately 0.4 and 0.7 μm), and leaf-cell structure, which reflects near-infrared light, are the basis for the computation of the normalised difference vegetation index (NDVI) [23]. The NDVI is defined as the difference between the chlorophyll absorption and near-infrared (NIR) bands over their sum and has been employed to estimate biomass [2, 28]. Moreover, water absorption features and the NDVI are known to be related to the leaf area index (LAI) [30, 31]. This is understandable since vegetation is expected to have a high water content, whose absorption dominates the NIR band in plant spectral reflectance. In a related application to that of measuring and quantifying biomass, plant stress has been assessed using spectral absorption features. Adams *et al.* [1] have used an approximation of the spectra based upon three bands to measure chlorosis of leaves in stressed plants. Smith *et al.* [32] have tested hyperspectral features in the red-edge of the spectrum so as to detect plant stress caused by soil-oxygen depletion, a common occurrence in natural gas leaks.

A number of methods have been proposed for the automatic detection of absorption bands. These range from the spectra derivative analysis [24] to scale space algorithms, such as the Fingerprint [29] and the MMWT [17]. All these algorithms are based upon the local analysis of the differential structure of the spectra and the subsequent recovery of absorption-boundary points. The modelling of absorption features has attracted attention in the geoscience community. Sunshine *et al.* [34] have proposed a e Modified Gaussian Modeling (MGM) method so as to resolve spectra into their physical components. Brown [5] has proposed a method to perform curve fitting of absorption bands in hyperspectral data.

In this paper, we propose a method based on continuum extraction and unimodal segmentation for the recovery of absorption bands. This method, in contrast with competing algorithms, is based upon analysis techniques which are global in nature. Since there has been a lack of comparative studies for the methods elsewhere in the literature, we also present

a performance comparison of automatic absorption band detection algorithms. Moreover, the absorption modelling methods in [34] and [5] model the shape of absorption bands represented in the energy-reflectance domain with parametric models whose parameters can be recovered using nonlinear least squares methods. Despite providing much insight into the position and relative strength of the absorption bands, they are more suited for chemistry-level analysis of material absorptions. This contrasts with the methods we are studying in this paper, which aim at detecting both, the position and the spread of the absorption band without the need for a parametric representation. Thus, we introduce a Bayesian winner-take-all mapping method which employs absorptions represented in a non-parametric fashion for purposes of material classification. This renders the absorption methods reviewed here applicable in a real-world setting for purposes of processing hyperspectral imagery.

The paper is organised as follows. In Section 2, we presented the theoretical basis of this paper, including a brief overview of conventional absorption recovery methods followed by our unimodal segmentation algorithm and absorption feature representation. In Section 3 we show how detected absorption bands can be used for purposes of material identification. In Section 4, we provide a comparison of our unimodal segmentation method with Fingerprint [29] and MMWT [17] for automatic absorption band detection and the application of detection results to material mapping in hyperspectral imaging. We conclude on the work in Section 5.

2 Absorption Band Detection

2.1 Automatic Absorption Band Recovery

2.1.1 Derivative Analysis

We begin by reviewing derivative analysis [6, 10, 13, 24], which is based on the extraction of interest points from the spectra making use of its derivatives with respect to wavelength index. Since an absorption band can be treated as a “dip” in the spectrum, we can locate these “dips” by recovering the spectral bands at which its second derivative changes sign or the first derivative is at an extremum, i.e. the inflection points of the spectrum. This can be achieved making use of derivative analysis in two ways. The first of these is by taking the first order derivative of the spectrum and finding the local maxima. The second route

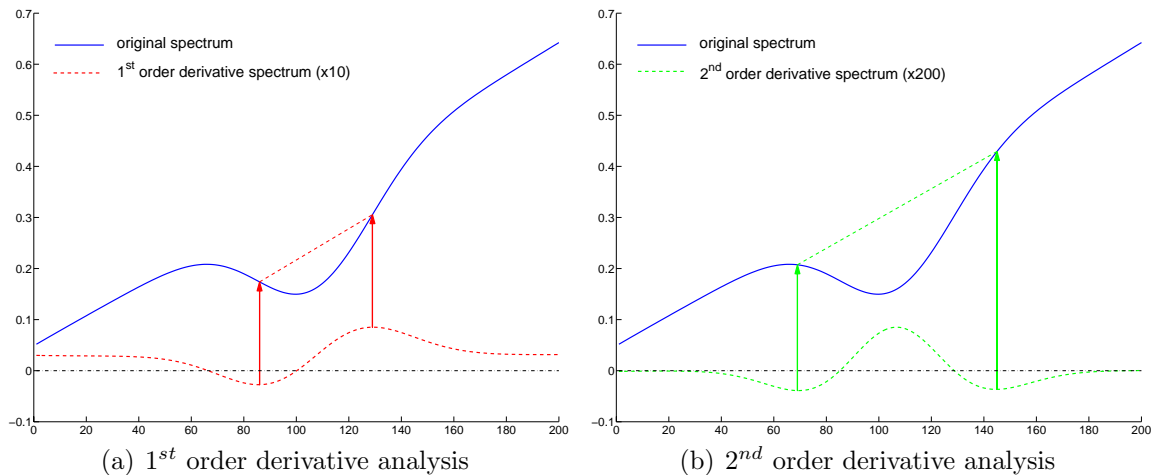


Figure 1: Derivative analysis

is by computing the second order derivative and finding its zero-crossing points, i.e. where the trace of the second derivative changes sign.

An example of this is shown in Figure 1, where the sample spectrum has been plotted using a solid curve. The first-order and second-order derivatives are those traces plotted using dashed lines. In Figure 1, the sample spectrum has a single large absorption whose end-points are given by the local minimum and maximum of the first order derivative, as indicated by the dashed lines and upward-pointing arrows in the panel. This is equivalent to the recovery of the zero-crossing points in the second-order derivative of the example spectrum. For the sake of clarity, we have scaled the first and second-order derivative traces by a factor of 10 and 200, respectively. In our plots, we have also shown the zero-reference level making use of a dotted horizontal line. As shown in Figure 1(a), the absorption band recovered by the first-order derivative analysis is quite narrow and does not cover the full range of the absorption feature. This can be resolved, in part, by computing the second-order derivative of the spectrum and, instead of making use of the zero-crossings for the absorption recovery process by employing its local maxima as an alternative, as shown in Figure 1(b). Here, the end-points of the absorption band correspond to the local maxima of the second-order derivative of the spectrum, with a local minima in the extreme point of the recovered absorption. It is clear that the absorption band recovered by the second order derivative analysis covers a wider range than that corresponding to the first order derivative analysis. This observation is an important one since it constitutes the foundation for the MMWT method [17], which will be introduced later in the paper.

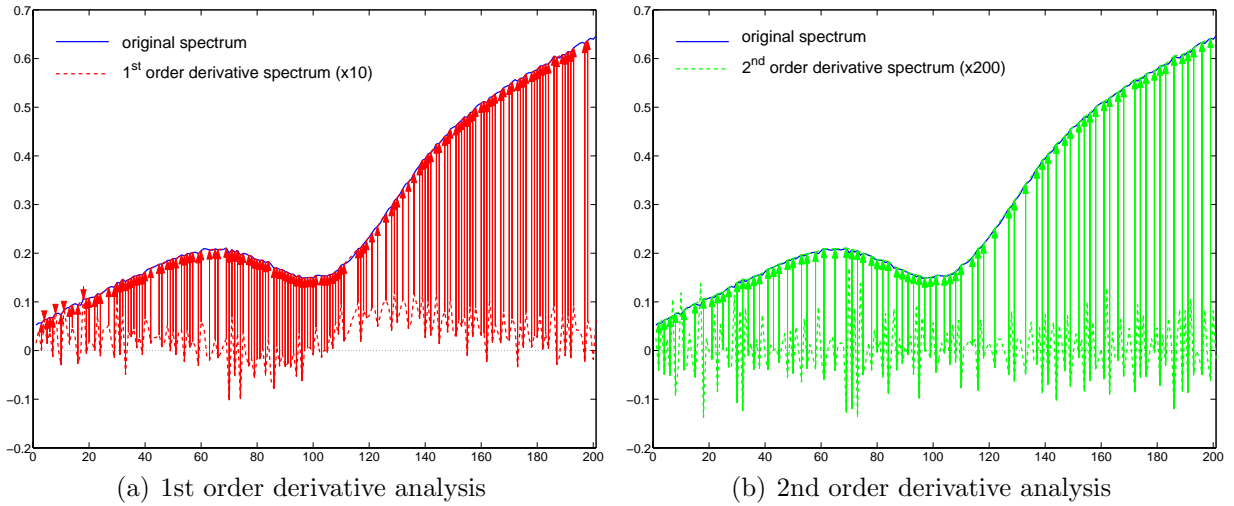


Figure 2: Derivative analysis applied to a sample noise-corrupted spectrum.

The derivative analysis method, however, is extremely sensitive to noise, even after applying filtering and signal processing techniques as a preprocessing step. To illustrate this, we have added Gaussian noise to the example spectrum in Figure 1. We apply derivative analysis to the noise-corrupted spectrum shown in Figure 2. In Figures 2(a) and (b), we show the recovered first and second order derivative traces, plotted in a similar fashion to that in Figure 1.

From the panels, it's clear that derivative analysis is prone to error due to noise. In Figure 2, the effect of the added Gaussian noise is significant, despite a small variance and a large signal-to-noise ratio (SNR). Moreover, from the traces, we can conclude that the second-order derivative is even more prone to noise corruption than the first-order one.

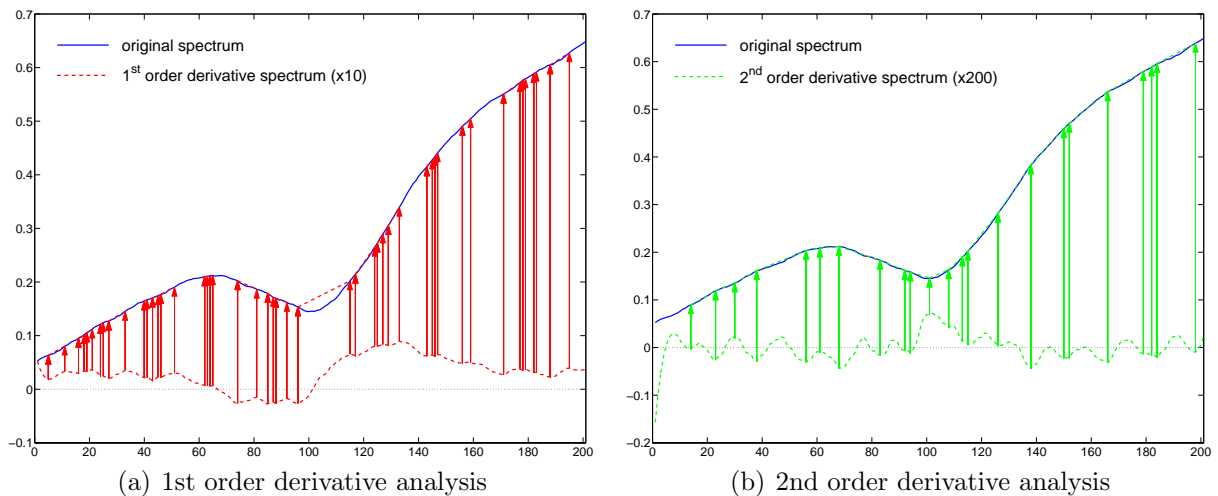


Figure 3: Derivative analysis applied to a sample noise-filtered spectrum.

To overcome noise effects, Lillesand *et al.* [24] have proposed the use of Savitsky-Golay filters as a preprocessing step in hyperspectral derivative analysis. Despite reducing the influence of noise, this approach aims at removing the noise before the spectrum is processed. Thus, the performance of the algorithm is dependent on the parameters of the filter bank. An example of this is shown in Figures 3(c) and (d). Here, we have used a third-order Savitsky-Golay filter and exemplified the effects of tuning the filter in a non-optimal fashion. Note the presence of numerous small absorption features whose end-points, indicated by the upward-pointing arrows, have been recovered making use of the local extrema of both, the first and the second-order spectrum derivatives.

2.1.2 Scale Space Methods

Both the fingerprint [29] and the MMWT method [17] are scale space methods. The purpose of these methods is to analyse the signal under study at different scales so as to recover suitable operators for the task in hand from its scale-space representation [36]. In the case of spectral analysis, this scale-space representation can be achieved via a convolution with gaussian kernels of continuously increasing variances [20]. For purposes of absorption detection, the problems that concern us are twofold. Firstly, the identification of absorption features in the spectra. Secondly, the localisation of the relevant absorption bands. For robust identification, we aim at removing noise at a coarse scale. However, for accurate localisation, we make use of detailed information at finer scales. This leads to the general framework of scale-space analysis: make use of the coarse scale to identify the absorption feature and locate the relevant bands by down-tracking them from coarse-to-fine scales.

Fingerprint

Fingerprint [29] is a method for recovering the absorption bands of the spectrum from its scale-space image. The scale-space image is a set of progressively smoothed traces computed from the original spectrum. As mentioned earlier, the smoothing of the spectrum is effected by convolving it with a Gaussian kernel with variance σ and mean μ as follows

$$F(R(\lambda), \sigma) = R(\lambda) * G(\mu, \sigma) \quad (1)$$

where $G(\mu, \sigma) = \exp\left(-\frac{(\lambda-\mu)^2}{2\sigma^2}\right)$ is a Gaussian kernel and $R(\lambda)$ is the raw spectrum as a function of the wavelength λ . The scale-space image is then comprised by all the functions $F(R(\lambda), \sigma)$ for a set of zero-mean Gaussian kernels with increasing values of σ . An example

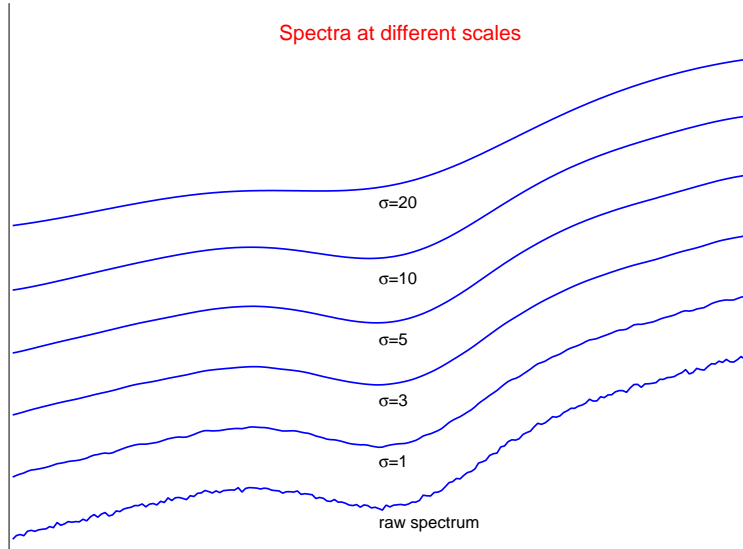


Figure 4: A sample spectrum plotted using different scale-values.

of the scale-space for a sample spectrum is shown in Figure 4. In the figure, we show the coarse-to-fine representation of the spectrum in a top-to-down fashion, i.e. the coarsest scale is represented by the top-most trace. As we can see, as the value of σ increases, i.e. the scale is coarser, the spectrum is smoother.

The Fingerprint method hinges on the rationale that inflexion points which are stable across the scale-space representation of the spectrum should capture absorption features in a robust manner. The method recovers these inflexion points, for each of the spectrum scale-images, using first order derivative analysis. There are many false inflexion points at the fine scales due to signal noise. These false positives will be eliminated making use of coarser scales. If an extremum is detected in nearby positions at adjacent scales, it is then likely to be either the beginning or the end of an absorption band. However, the inflexion points recovered from the coarser scales may be inaccurate, as the spectrum at these scales is often oversmoothed. Thus, the absorption end-points are recovered by backtracking from coarse-to-fine scale images.

At this point, it is worth noting that the smoothing and differentiation steps can be computed efficiently by convolving the spectrum with the derivative of the Gaussian kernel. This is due to the fact that the convolution and the differentiation operations are commutative. This procedure can then be viewed as that of smoothing the spectrum with the Gaussian kernel $G(\mu, \sigma)$ and then taking its derivatives. We can formulate this by using the shorthand

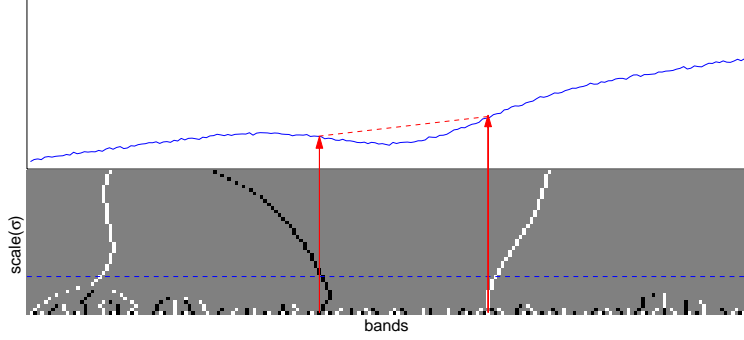


Figure 5: Sample spectrum in scale-space, as recovered by the Fingerprint method.

$x = \lambda - \mu$ as follows

$$F(R(\lambda), \sigma)|_{\sigma_i} = R(\lambda) * G_x(\mu, \sigma) \quad (2)$$

where we have written $F(R(\lambda), \sigma)|_{\sigma_i}$ to imply that the scale-image is taken at the i^{th} value of σ and

$$G_x(\mu, \sigma) = \frac{\partial G(\mu, \sigma)}{\partial x} = \frac{x}{\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) \quad (3)$$

is the derivative of the Gaussian kernel with respect to the difference between the mean μ and the spectrum $R(\lambda)$.

By finding the maxima and minima of $F(R(\lambda), \sigma)$, we are then recovering the inflexion points of the spectrum at the i^{th} scale. This is the Fingerprint representation of the spectra. A Fingerprint example is shown in Figure 5. The white traces in the panel indicate the local maxima of the first order derivatives at different scales. The black lines indicate the local minima of the first-order derivative. An absorption band is recovered by a pair of minima-maxima lines. The minima traces indicate the beginning of the absorption, whereas the maxima evidences the end of the feature. Detection is accomplished at a large value of σ , i.e. a coarse scale. However, the end-points of the absorption feature are given by the extrema at finer scales, which have been backtracked from the scale-images corresponding to larger values of σ . In the figure, we have indicated the end-points of the absorption feature using red vertical arrows.

As we can see in Figure 5, Fingerprint is quite robust to noise-induced artifacts in the spectrum. Nonetheless, the recovered absorption appears narrower as compared to the true feature. Moreover, the scale at which detection is made plays an important role in the identification of the absorption features. The variance σ at which the detection task takes place can be viewed as a cut-off value τ in the scale-space (shown in the figure as a dashed

horizontal line), which is often chosen empirically.

Modulus Maximum Wavelet Transform (MMWT)

The Fingerprint method can be generalised to MMWT [17] making use of second-order derivatives. The idea is to apply a wavelet transform to the spectrum using the second-order Derivatives of Gaussian (DoG) wavelets with continuously varying scale parameters. As in the Fingerprint approach, the absorption is then recovered using the extrema in the scale-space spanned by the wavelet transforms. For the i^{th} scale, the wavelet transform of the spectrum $R(\lambda)$ is given by

$$F(R(\lambda), \sigma)|_{\sigma_i} = \int R(\lambda)\psi_{\mu,\sigma}(\lambda)d\lambda \quad (4)$$

where $\psi_{\mu,\sigma}(\lambda) = \frac{1}{\sqrt{\sigma}}\psi\left(\frac{\lambda-\mu}{\sigma}\right)$ is a scaled and translated Gaussian wavelet function.

Note that for the second-order DoG wavelet, Equation 4 is equivalent to the second-order derivative analysis for fixed values of σ . This is due to the fact that the wavelet transform can be viewed, for the second-order DoG wavelet, as the smoothing and computation of the second-order spectrum derivative.

We are interested in the the local maxima and minima of the scale space representation $F(R(\lambda), \sigma)$. By relating the local maxima and minima at different scales to one another, we can obtain the positive and the negative minimal traces across the scale-space. Each absorption is then characterised by endpoints which correspond to negative minimal traces with a positive maximal line in the middle point of the feature. This follows directly from the result of second-order derivative analysis in Figure 1. As in the Fingerprint method, the detection is made at large values of σ and is dependent of a cut-off value. The localisation is achieved by back-tracking the modulus maximal lines down to the finest scale.

An example of MMWT for absorption band detection is shown in Figure 6. The white traces indicate the positive maximal lines and the black traces show the negative minimal lines. The absorption feature is indicated using upward pointing arrows and the cut-off value is given by the horizontal dashed line. From the figure, we see that the absorption recovered by this method covers a larger range than that recovered by the Fingerprint algorithm.

2.1.3 Unimodal Segmentation

All the methods for automatic absorption band recovery presented so far are based on the local analysis of the spectra derivatives, either in a single scale or a multiple-scale fashion.




Figure 6: Sample spectrum in scale-space, as recovered by the MMWT.

In this paper, we introduce a novel approach to the detection of absorptions which is based upon a global analysis of the spectrum under study.

Our method is based on the assumption that an absorption can be defined in terms of the global shape of the spectrum. If the global shape of the spectrum is convex, and a small region in it is concave, then this region will be regarded as an absorption. Thus, we can recover absorption features in a two-step fashion. First, we can model the shape of the spectrum as a convex hull with the continuum of the spectrum as an upper-bound. Second, the absorption bands can be then viewed as local spectral segments which break the convexity at their corresponding local neighbourhoods. Viewed in this manner, the spectrum, as a function of the wavelength λ , can be treated as points distributed in a two dimensional space. The convex hull, which is the minimum convex polygon enclosing all these points, can be found using readily available algorithms. To illustrate this treatment, in the upper panels of Figures 7(a) and 7(b) we show, in dashed lines, the continuum curve recovered using the convex hull for two sample spectra.

As shown in Figure 7, the continuum of a spectrum is given by its convex hull so as to connect the spectrum local maxima. As a result, it can be viewed as an envelope that isolates the absorption bands in the reflectance spectra. In other words, the use of the convex hull splits the original spectrum into several spectral segments defined over local maxima. These segments assume a concave shape against the continuum.

After recovering the continuum, we perform continuum removal on the spectrum. This is a standard procedure in remote sensing to isolate the local absorption feature from other photometric effects [7]. The removal operation is a simple normalisation of the spectrum $R(\lambda)$ with respect to the continuum curve $C(\lambda)$, such that the continuum removed spectrum




Figure 7: Continuum removal results. Left-hand panel: Continuum removal for a unimodal distribution; Right-hand panel: Continuum removal for a multimodal distribution.

$r(\lambda)$ is given by

$$r(\lambda) = \frac{R(\lambda)}{C(\lambda)} \quad (5)$$

where, as before, λ is the wavelength variable.

Examples of continuum removed spectra are shown in the bottom panels of Figure 7. The spectra shown here have been normalised to unity and inverted in order to turn absorptions into peaks. As we can see, the spectral segments in Figure 7(a) are unimodal, i.e. they only describe a single peak in the spectrum. In contrast, the segment in Figure 7(b) is multimodal with two peaks. These two peaks represent two absorption bands.

To take our analysis further and recover the absorption bands, we make use of unimodal regression. The idea is to further split any number of segments into absorption bands, each of which satisfy unimodal constraints via unimodal regression [33].

Unimodal regression can be implemented through isotonic regressions. The purpose of an increasing/decreasing isotonic regression is to fit a monotonically increasing/decreasing piecewise smooth line to the original sequence so as to minimise the distance between them. Mathematically, denote by $\{y_i | i = 1, \dots, n\}$ a sequence of spectral measurements of order n .

The isotonic regression of the data is the set $\{\hat{y}_i | i = 1, \dots, n\}$ that minimises

$$\begin{aligned} & \sum_{i=1}^n |y_i - \hat{y}_i|^2 && (6) \\ \text{s.t. } & \hat{y}_1 \leq \dots \leq \hat{y}_n && \text{for increasing isotonic regression} \\ & \hat{y}_1 \geq \dots \geq \hat{y}_n && \text{for decreasing isotonic regression} \end{aligned}$$

In the equations above, we have used the L_2 Euclidean distance. Note that other distance measures can also be used. An advantage of using L_2 distance is that isotonic regression can be determined in $O(n)$ time by using Pool Adjacent Violators [3]. Pool Adjacent Violators (PAV) for increasing isotonic regression are defined as follows. For the spectral measurements $\{y_i | i = 1, \dots, N\}$ ordered in decreasing rank, i.e. $y_i \geq y_{i+1} \geq \dots \geq y_j$, the PAV is given by

$$\hat{y}_k = \begin{cases} \frac{\sum_{k=i}^j y_k}{j - i + 1} & k \in [i, j] \\ y_k & \text{otherwise} \end{cases} \quad (7)$$

That is, each decreasing rank sequence is replaced by the mean value on the interval $[i, j]$ while non-decreasing values are preserved. PAV for decreasing isotonic regression can be defined likewise. In Figure 8(a), we show an example of increasing isotonic regression. In the figure, the original signal has been plotted using a solid blue trace. The corresponding increasing isotonic regression result is given by the dash-dotted plot.

Unimodal regression can be effected by locating a splitting point y_m in the sequence. Once y_m is at hand, we apply increasing isotonic regression to the left-hand segment of the sequence and decreasing isotonic regression to the right-hand segment. This is done so as to minimise the error function in Equation 6 subject to the unimodal constraint $\hat{y}_1 \leq \dots \leq \hat{y}_m \geq \hat{y}_{m+1} \dots \geq \hat{y}_N$. Examples of isotonic and unimodal regression are shown in Figure 8.

A straightforward manner to implement unimodal regression is to treat each point in the sequence as a candidate splitting point and apply increasing/decreasing isotonic regression to the left/right segments. The optimal splitting point is that whose sum of squared errors is minimum for the pair of isotonic regressions. Note that the approach above has a computational complexity of $O(n^2)$. An improved unimodal regression algorithm is proposed in [33] by making use of the prefix isotonic regression, which is a method to determine the errors of isotonic regression for the interval $\{y_1, \dots, y_i\}$ with $i = 1, \dots, N$ in a single scan of the sequence. This is achieved via dynamic programming. A detailed review on prefix

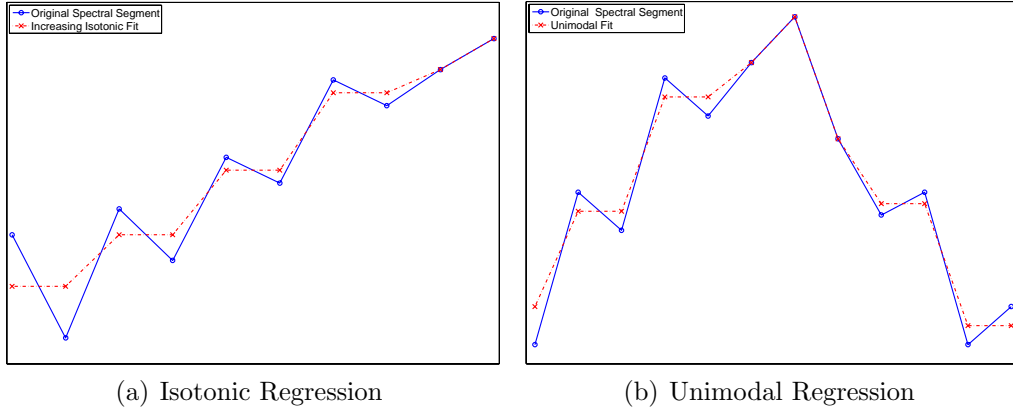


Figure 8: Illustration of isotonic and unimodal regression.

isotonic regression is beyond the scope of this paper and interested readers are referred to [33]. By applying prefix isotonic regression to both, the original and the inverted sequences, the optimal splitting point can be obtained with a complexity of $O(n)$.

With these ingredients, the step sequence for our absorption band detection algorithm is as follows

1. Recover the convex hull over the spectrum using the Graham scan algorithm [16].
2. Apply continuum removal and invert the continuum-removed spectrum.
3. Perform unimodal regression. The stopping criterion for the fitting procedure is given by the threshold ϵ . If the segment is not unimodal, i.e. the fitting error is above the threshold ϵ , split the segment at the local minima.
4. Merge any two adjacent segments if the unimodal constraint holds for their union. This is done, again, based upon the threshold ϵ .
5. Interleave steps 3 and 4 until no further splitting and merging operations can be performed.

The main drawback of unimodal regression is its tendency to generate false positives. In our aim of computation, this implies that the algorithm may yield many spurious, small absorptions when processing noisy spectra. To overcome this problem, a postprocessing step can be applied by eliminating the absorptions whose magnitude is negligible or those that cover a very small number of bands. Despite this, notice that the main absorption features will still be well preserved by unimodal segmentation. Moreover, our experience

suggests that unimodal regression results are less sensitive to the variation of threshold and parameter values than the alternatives.

Examples of absorption band detection using our method for two example spectra are shown in the top row of Figure 9. The second and third rows show the results yield by the Fingerprint and the MMWT, respectively. In contrast with the Fingerprint method, the absorption bands recovered by our approach are in better accordance with the actual absorptions in the spectra. Furthermore, the MMWT cannot recover all the absorptions in either of the spectra under study. The MMWT is less robust than the other two methods due to the use of second order derivatives.

2.2 Complexity Analysis

In this section, we examine the computational and algorithmic complexity of the methods mentioned above. To commence, recall that the Fourier transform played an important role in the the automatic absorption algorithms presented as alternatives to our unimodal segmentation method. Notice that, by making use of the FFT, the Fourier transform can be computed in $\mathcal{O}(n \log n)$ floating-point operations, where n is the number of bands in the spectrum under study. Similarly, for a convolution mask of size k , multiplication in the frequency domain only takes $\mathcal{O}(kn)$ floating-point operations. Thus, the overall complexity is $\mathcal{O}(kn) + 2\mathcal{O}(n \log n)$, which corresponds to the combined costs of the application of a FFT, its inverse and the multiplication operation in the frequency domain.

In case of using signal processing techniques in the frequency domain, the complexity of the derivative analysis method will be given by $\mathcal{O}(kn) + 2\mathcal{O}(n \log n)$ plus the overhead introduced by the filtering algorithm or preprocessing step of choice. In our implementation, we have employed Savitsky-Golay filters as a preprocessing step. These filters can be viewed as a Finite Impulse Response filter [26] governed by a least squares fit to the spectrum, where the window size and the order of the fitting polynomial are parameters set by the user.

Scale-space methods, i.e. Fingerprint and MMWT, are more computationally demanding due to the fact that first and second-order derivatives have to be computed for every scale. Here, the number of scales and their increments are parameters chosen by the user. In terms of complexity, for m scales, this is given by $m(\mathcal{O}(kn) + 2\mathcal{O}(n \log n))$ plus the cost of the backtracking detection algorithm used to find the maxima and minima in the scale-space.

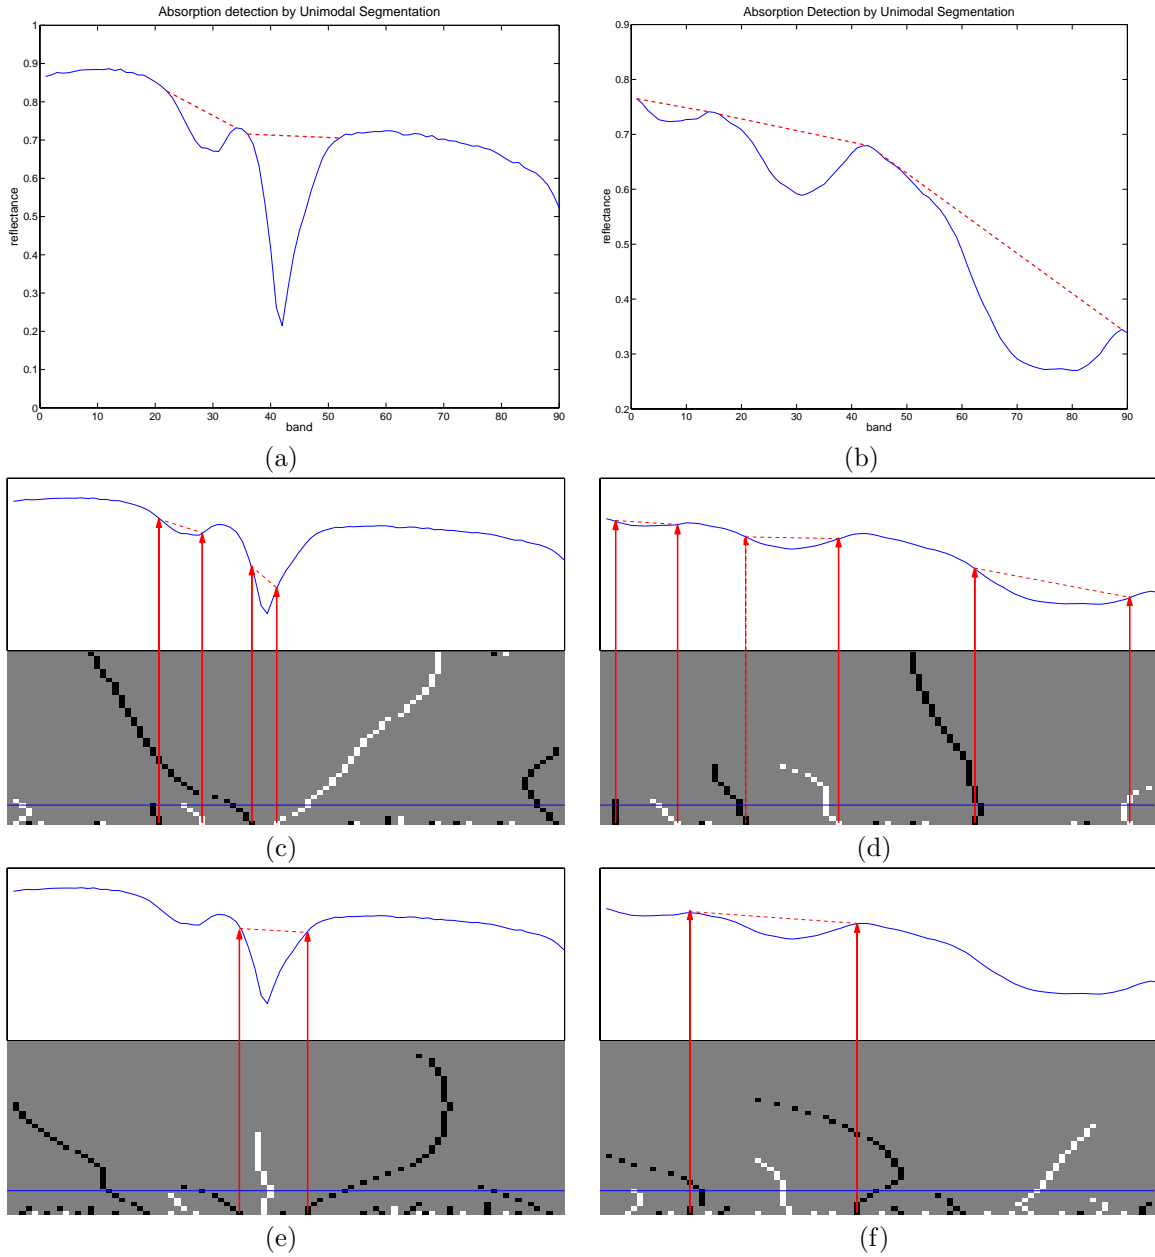


Figure 9: Absorption detection results. Top row: results obtained using continuum removal and unimodal segmentation. Middle row: results yield by the Fingerprint method; Bottom row: results yield by the MMWT.

In our case, we have used a simple backtracking approach based upon a nearest-neighbour search [35] over the minima and maxima at every set of successive scales.

For unimodal segmentation, the bottle neck is in the extraction of the convex hull and the unimodal regression. Along these lines, the complexity for convex hull extraction is $O(n \log n)$, where n is the length of the spectrum. By using the implementation in [33], unimodal regression can be achieved with a cost of $O(n \log n)$ floating-point operations. In

fact, the cost of unimodal regression is much reduced by the nature of our application since we seldom have to apply it to the whole spectrum. This is due to the continuum removal operation where those bands which comprise the convex hull can be removed from further consideration at the unimodal regression step.

In addition to the computational cost, we also examine the algorithmic complexity of these methods. This is important due to the fact that algorithmic complexity often determines the scalability of the method. Hence the simpler the algorithm, the more general situations it applies. In the other hand, the more complex the algorithm, the more accurate of the underlying model. This is in agreement of the well known Occam’s razor theory [25] which states that models and, hence the algorithms they embody, should not be complicated beyond necessity. This complexity can, to some extent, be characterised by the number of parameters used in the algorithm.

In our implementation of derivative analysis, as mentioned earlier, we use two parameters, i.e. the Savitsky-Golay polynomial order and window size. In the case of scale-space methods, both Fingerprint and MMWT require four parameters to be set by the user. These are the range defined by the initial and final values of the scale, i.e. σ_o and σ_m , the step of the scale increment, and the the cut-off value τ . Despite being inefficient, the window size can be chosen as the same length of signal. Nonetheless, the choice of other parameters need some scrutiny and will influence the performance. For example, in Figures 9(c) and 9(d), the value of τ has been set to 5, if it had been chosen to be 10, then some absorption bands would not had been detected. However, $\tau = 10$ might be the right choice for other spectra. Similarly, the scale increment value is also important. In theory, it should be chosen as small as possible, but it is worth noting that the computational cost increases as does the number of scales m in the range $[\sigma_o, \sigma_m]$.

The unimodal segmentation method, on the other hand, has only one parameter, i.e. the tolerance ϵ of the fitting error for the regression step. Hence, the performance of the unimodal segmentation algorithm is easier to control with only one parameter to tune. Furthermore, by normalising the spectrum to unity, the choice of the tolerance parameter is better constrained and the algorithm results are more robust to variations of ϵ .




Figure 10: Plots of parametric functions used to represent absorption features.

Functional	MGM [34]	Voight [5]	Lorentzian [5]
$g(R(\lambda))$	$s \cdot \exp - \frac{(R(\lambda)^\vartheta - \mu^\vartheta)^2}{2\sigma^2}$	$\frac{v}{[1 + 0.5\beta^2(R(\lambda) - \mu)^2/\sigma^2]^{1/\beta^2}}$	$\left[\frac{2v\sigma^2}{2\sigma^2 + (R(\lambda) - \mu)^2} \right]$

Table 1: Summary of the functionals used elsewhere in the literature. In the table above, β and ϑ are user-defined parameters.

3 Absorption Representation and Classification

As mentioned earlier, absorption bands have been modelled via parametric representations [34, 5]. Here, an absorption band is represented by a parametric function after applying a predefined transformation to the original reflectance spectrum as a pre-processing step. Typical transformations comprise converting wavelength to wavenumber in the x-axis and reflectance to negative log reflectance [5] or applying continuum removal to the log-reflectance spectrum plotted in the wavenumber space [34]. The main purpose of these pre-processing steps is to convert reflectance to apparent absorbance represented as a function of energy, as related to the wavenumber.

To represent the spectrum $R(\lambda)$, a standard choice of the parametric function in the transformed spectrum space is the Gaussian function

$$g(R(\lambda)) = v \cdot \exp \left(- \frac{(R(\lambda) - \mu)^2}{2\sigma^2} \right) \quad (8)$$

where v is a proportionality constant and μ and σ are the center and the bandwidth of the absorption band, respectively. Other choices include the modified Gaussian model (MGM) [34], the Lorentzian function [5] and the Voight functional [5]. These functions are plotted in Figure 10 and summarised in Table 1.

Absorption band modelling techniques are specially useful in resolving overlapping absorptions. This is exemplified in Figure 11(a), where three absorption bands are modelled as single Gaussian functions. Another advantage is their inherent relation to material chemistry. However, absorption modelling does not provide a direct means for matching different absorption features for purposes of material mapping and classification. A straightforward choice here is to use the parameters themselves as the input vectors to the mapping process. For instance, in the case of Gaussian functions, we can compare different absorptions by employing their center positions μ and bandwidths σ . Yet, this representation inevitably loses discriminant information due to the lower order statistics which it comprises.

By applying the proposed unimodal segmentation method to the spectrum, we can recover two absorption bands, which have been plotted using red dotted traces in Figure 11(b). Here, we adopt a non-parametric representation of the absorptions by using the spectral values within the absorption bands as feature vectors. This enables us to discriminate between both, separable and overlapping absorption bands without making use of lower order statistics. As a result, robust material mapping can be effected by discriminating between spectra with different absorption bands based on their shapes. To illustrate this, consider the spectrum shown in Figure 11 has been pre-process using continuum removal [7]. If one of the overlapping bands in Figures 11(b) vanishes, the resulting spectrum becomes that shown in Figure 11(c). Making use of the spectral band-values of the continuum removed spectrum segment within the two remaining absorptions, its still possible to discriminate between the two different spectra in the figures. Their non-parametric representations are plotted together in Figure 11(c) using cyan and red dotted lines, respectively.

Another advantage of non-parametric absorption feature representation for material mapping is low computational cost. For purposes of classification, we only have to recover the




Figure 11: Parametric and non-parametric representation of absorption feature shapes.

absorption bands of reference spectra and store their spectral signatures for purposes of testing. Absorption modelling methods rely on curve fitting, which must be applied to each reference spectrum as well as unknown testing spectra. Therefore, parametric modelling of absorption bands for purposes of material mapping on large image-sizes is potentially computational expensive. Moreover, for our non-parametric representation, we can apply local analysis to the spectral range of interest as determined by the absorption bands detected from the reference spectra. This further reduces the computational complexity of the classification process.

In this paper, we study two classification methods which have been widely used in the geoscience community, the Spectral Angle Mapper (SAM) [21] and Spectral Feature Fitting (SFF) [9]. SAM is an automatic method for comparing image and library spectra. By taking apparent reflectance at input, i.e. the true reflectance up to an unknown scaling factor, the correlation coefficient between two spectra can be used as a similarity measure due to scale invariance. The Spectral Angle (SA) between two absorption band segments is then defined as the arc cosine of their correlation coefficient as follows,

$$\alpha(\mathbf{x}, \hat{\mathbf{x}}) = \arccos \left(\frac{\mathbf{x}^T \hat{\mathbf{x}}}{\|\mathbf{x}\| \|\hat{\mathbf{x}}\|} \right) \quad (11)$$

where \mathbf{x} and $\hat{\mathbf{x}}$ are the vector forms of the absorption band segments under comparison.

SAM performs material mapping based on SA. The smaller the SA, the more similar the two spectra under comparison are and the more likely the two materials correspond to one another. As a result, we can use SAM for comparison of local spectral segments and assign a likelihood value based upon their similarity. The likelihood for the spectral segments \mathbf{x} and $\hat{\mathbf{x}}$ in the interval $[\lambda_i, \lambda_j]$ becomes

$$f(\mathbf{x}, \hat{\mathbf{x}} | (\lambda_i, \lambda_j)) = \exp \left(\frac{-\alpha(\mathbf{x}, \hat{\mathbf{x}})^2}{2\sigma^2} \right) \quad (12)$$

where α is a constant and σ is a predefined variance parameter.

The other method examined here, SFF, is based on the local comparison of continuum removed spectra. Let us denote the continuum removed spectra of \mathbf{x} and $\hat{\mathbf{x}}$ as \mathbf{y} and $\hat{\mathbf{y}}$, respectively. SFF aims at matching the reference spectrum $\mathbf{y} = [y(\lambda_i), y(\lambda_{i+1}), \dots, y(\lambda_j)]^T$ to the unknown candidate spectrum $\hat{\mathbf{y}} = [\hat{y}(\lambda_i), \hat{y}(\lambda_{i+1}), \dots, \hat{y}(\lambda_j)]^T$ by recovering a scaling factor κ , which is determined by solving the following Least Squares (LS) problem

$$\min \sum_{\lambda_k \in [\lambda_i, \lambda_j]} (a + by(\lambda_k) - \hat{y}(\lambda_k))^2 \text{ and } \kappa = \frac{1 - b}{b} \quad (13)$$

Input: the set of reference spectra $\mathbf{x}_1, \dots, \mathbf{x}_N$ and the hyperspectral image

- Training:

Apply unimodal segmentation to each reference spectrum and recover the absorption bands $\mathcal{C}^l = \{(\lambda_i^{(k)}, \lambda_j^{(k)}) | k = 1, \dots, m(l)\}$, where \mathcal{C}^l denotes to the set of absorption bands for the i th spectrum, $\lambda_i^{(k)}$ and $\lambda_j^{(k)}$ refers to the first and the last channel for the k^{th} absorption band of the l^{th} reference spectrum and $m(l)$ denotes the the total number of absorptions for the reference spectrum indexed l .

- Testing:

For each candidate spectrum $\hat{\mathbf{x}}$ in the image

- For each reference spectrum \mathbf{x}_l

- 1 Extract the absorption features for the candidate spectrum by using unimodal segmentation.
- 2 Compute the likelihood of the candidate spectrum $\hat{\mathbf{x}}$ being a match to the reference spectrum \mathbf{x}_l . The likelihood is given by

$$L(\mathbf{x}_l; \hat{\mathbf{x}}) = \sum_k w_l^{(k)} f(\mathbf{x}_l, \hat{\mathbf{x}} | (\lambda_i^{(k)}, \lambda_j^{(k)})) \quad (9)$$

where $w_l^{(k)}$ is the weight of the k th absorption band for reference spectrum \mathbf{x}_l . Here $w_l^{(k)}$ is proportional to the area of absorption band such that $\sum_k w_l^{(k)} = 1$. Note that $f(\mathbf{x}_l, \hat{\mathbf{x}} | (\lambda_i^{(k)}, \lambda_j^{(k)}))$ is the likelihood for the absorption band indexed k , which can be calculated via Equations 12 or 15 for either SAM or SFF.

- Endfor

- Compute the posterior probability for the candidate spectrum $\hat{\mathbf{x}}$ given the reference spectrum \mathbf{x}_l using Bayes rule. To do this, we assume equal prior probability for each reference spectrum

$$P(\mathbf{x}_l | \hat{\mathbf{x}}) = \begin{cases} \frac{L(\mathbf{x}_l; \hat{\mathbf{x}})}{\sum_l L(\mathbf{x}_l; \hat{\mathbf{x}})} & \max_l L(\mathbf{x}_l; \hat{\mathbf{x}}) > \varepsilon \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

- Endfor
-

Figure 12: Algorithm for Material Mapping based on Absorption Bands

The solution to the above equation is given by

$$\begin{aligned}
 a &= \frac{1}{n} \left(\sum_{\lambda_k \in [\lambda_i, \lambda_j]} \hat{y}(\lambda_k) - b \sum_{\lambda_k \in [\lambda_i, \lambda_j]} y(\lambda_k) \right) \\
 b &= \frac{\sum_{\lambda_k \in [\lambda_i, \lambda_j]} \hat{y}(\lambda_k) y(\lambda_k) - \frac{1}{n} \sum_{\lambda_k \in [\lambda_i, \lambda_j]} \hat{y}(\lambda_k) \sum_{\lambda_k \in [\lambda_i, \lambda_j]} y(\lambda_k)}{\sum_{\lambda_k \in [\lambda_i, \lambda_j]} y(\lambda_k)^2 - \frac{1}{n} \left(\sum_{\lambda_k \in [\lambda_i, \lambda_j]} y(\lambda_k) \right)^2}
 \end{aligned} \tag{14}$$

where n is the number of bands under comparison, i.e. $n = i - j + 1$, and k can be calculated from b making use of Equation 13.

The likelihood or affinity of the two spectra compared making use of SFF is determined by two factors, the scale factor κ and the goodness of fit (GoF), i.e. the minimum value achieved by the LS problem in Equation 13. The two spectra under comparison are deemed to be similar for large scale factors κ and small GoFs. With these ingredients, we can define the likelihood or fitting score for two local spectra segments \mathbf{y} and $\hat{\mathbf{y}}$ as follows being similar

$$f(\mathbf{y}, \hat{\mathbf{y}} | (\lambda_i, \lambda_j)) = \exp \left(\frac{-GoF(\mathbf{y}, \hat{\mathbf{y}})^2}{2\sigma^2\kappa^2} \right) \tag{15}$$

To perform absorption feature based material classification, we use the values of $f(\mathbf{y}, \hat{\mathbf{y}} | (\lambda_i, \lambda_j))$ yield by either SAM or SFF and compute the posterior probability of a reference spectrum \mathbf{x} given a candidate spectrum $\hat{\mathbf{x}}$. This is, we use the likelihood for the local spectra segments to recover a posterior probability that can be used for purposes of material classification. The posterior indicates whether a candidate spectrum is likely to correspond to a reference one. Therefore, classification becomes a Bayesian winner-take-all procedure [11]. The algorithm for the computation of the posterior probability is shown in Figure 12. It is worth noting that, in calculating the posterior probabilities with respect to different reference spectra, we do not enforce a solution if the likelihood values are all very small. This is due to the fact that, if the posterior is close to zero, this means that the image spectra correspond to a material which has not been represented by any of the reference spectra. Therefore, we just set the posterior to 0 for all classes.

4 Experiments

In this section, we show two sets of experiments. The first of these is a comparative study on the absorption detection results yield by our unimodal segmentation algorithm and two alternatives, namely fingerprint and MMWT. We then turn our attention to material mapping on a hyperspectral image based on spectral feature matching.

4.1 Absorption Recovery

To compare the performance of the absorption detection algorithms, we have used two datasets for our experiments. The first of these is a subset of the USGS spectra library, which contains 481 spectra of different minerals. The spectral samples were selected from the AVIRIS sensor image data, whose spectral range is between 395.1nm and 2560.0nm with non-uniform samplings over a total number of 420 bands. The second dataset contains spectra collected in-house using an ASD spectrometer. This data comprises 146 spectra of different natural materials, such as leaves, wheat, apples, nuts and human skin. This spectral data is uniformly sampled over the range between 400nm and 1075nm with 1nm steps, i.e. 676 bands, under an artificial sunlight.

For purposes of validation, we have performed continuum removal and denoising on the spectral data. Each spectrum is smoothed by a 2^{nd} order Savitsky-Golay filter with window size of 11 bands before recovering the continuum-removed spectra using the algorithm in [7]. This reduces the effects of noise and provides us with the filtered spectral data that will then be used in the following section.

In our experiments, we have focused our attention in the Fingerprint, MMWT and the unimodal segmentation methods. The reasons for this are twofold. Firstly, as mentioned earlier, the derivative analysis is very sensitive to noise corruption. As a result, even making use of a Savitsky-Golay filter as a preprocessing step, it is greatly outperformed by the alternatives. Secondly, the scale-space methods, i.e. Fingerprint and MMWT, are both based upon first and second-order derivative analysis on the scale images of the spectra. Thus, performing experiments on these methods can be viewed as an evaluation of derivative analysis in a scale-space setting.

4.1.1 Noise-filtered Spectra

We commence by testing the absorption detection methods on the noise-filtered spectral samples and compared their results the continuum removed spectra, which we view as our ground truth. The parameters of these methods have been chosen by cross-validation so as to ensure optimal performance on the ground-truth data. For the Fingerprint and the MMWT, we have set $\sigma_o = 3$, $\sigma_m = 52$ and $\tau = 25$. In the other hand, for the unimodal segmentation method, we have used $\epsilon = 0.002$.

Dataset	Absorption Recovery Method		
	Fingerprint	MMWT	Our method
USGS dataset (185 spectra, 420 bands)	0.06227	0.05924	0.03816
Our dataset (92 spectra, 676 bands)	0.13859	0.08098	0.06348

Table 2: Average processing time, in seconds, for the datasets used in our experiments as a function of absorption recovery method.

In Table 4.1.1, we show the average processing time spent per processed spectrum for the methods under study. All our experiments were performed on a Pentium 4, 3 GHz PC with 1 GB of RAM. All algorithms have been implemented making use of optimised convolution (Fingerprint and MMWT) and unimodal regression (unimodal segmentation) routines. As the choice of the number of scales being processed will have a direct effect on the efficiency of the Fingerprint and MMWT, the results listed here should be viewed as a reference rather than an exhaustive computational cost study. However, from Table 2, we can conclude that the unimodal segmentation algorithm provides a margin of computational advantage over the scale space methods. This is particularly true, in general, if a large number of scales are used.

We now proceed to assess the performance of the methods above. We do this by comparing the detection results yield by these methods to the continuum-removed spectra. In Figure 13, we show example absorption detection results for some selected sample spectra from the USGS database. Example results for our data are shown in Figure 14. From the figures, we can conclude that, especially for the mineral spectra in the USGS database, the unimodal segmentation algorithm achieves better results in terms of the area and separation of the recovered absorptions. Fingerprint and MMWT can also locate most of the absorption bands but they are unable to recover the full absorption area, with the MMWT generally recovering a larger absorption bandwidth than the Fingerprint method. Scale space methods may not be able to recover the absorption feature in cases where the local spectral curve is not differentiable, which may be the case with two adjacent absorption bands. This is a common occurrence in many mineral spectra, as shown by the results on the Alunite and the Neodymium Oxide spectra in Figure 13.

The unimodal segmentation algorithm can recover the full extent of the absorption feature. Nonetheless, it may fail to perform well in cases where the continuum is not convex-shaped.

Examples of this are the results on Eucalypt leaf powder shown in Figure 14. Furthermore, the unimodal assumption on the absorption features may hold for mixed absorptions, as in the case of the Navel leaf spectra shown in Figure 14. Here, the two absorption features, one shallow segment around 600nm and the deep band around 670nm, were mistakenly detected as a single absorption by the unimodal segmentation algorithm. The reason is that the absorption around 670nm is too strong and it makes the shallow band less prominent. As a result, their mixture forms only one peak after continuum removal and cannot be separated simply based on the unimodal constraint. From our comparison results, we can assert that there is no single method which consistently give optimal results for all the spectra in the datasets. In our experience, the unimodal segmentation algorithm performs better for the mineral spectra in the USGS database. This is due to the fact that, for mineral spectral data, the cusps and valleys in the spectra are more obvious. For organic material, like plants and leaves, whose absorption features are “shallow”, the MMWT is, in general, a good option.

Finally, we have performed a quantitative analysis on the results. To this end, we compare the detection results with the ground truth making use of a number of error measures. Here, we compare the absorption detection results with the continuum-removed spectra based on the rationale that different detection results will yield different continua. As a result, absorption detection results which are close to one another should produce similar continuum-removed spectral bands, whereas results that are far removed from each other are penalised to a greater extent. To do this, we use the detection results to perform continuum removal by using the line over connecting the absorption end-points as an analogue of the continuum yield by the method in [7]. Accordingly, for the examples shown in Figures 14 and 13, the continua are given by the broken line in each panel, which has been computed making use of the absorption detection results yield by the algorithms under study.

Thus, the step sequence for the procedure above is as follows

- 1.- After recovering the absorption features, perform continuum removal making use of the continuum yield by the detection results.
- 2.- Normalise and “flip” both spectra.
- 3.- Compute the similarity between the two spectra making use of the distance measure of choice.

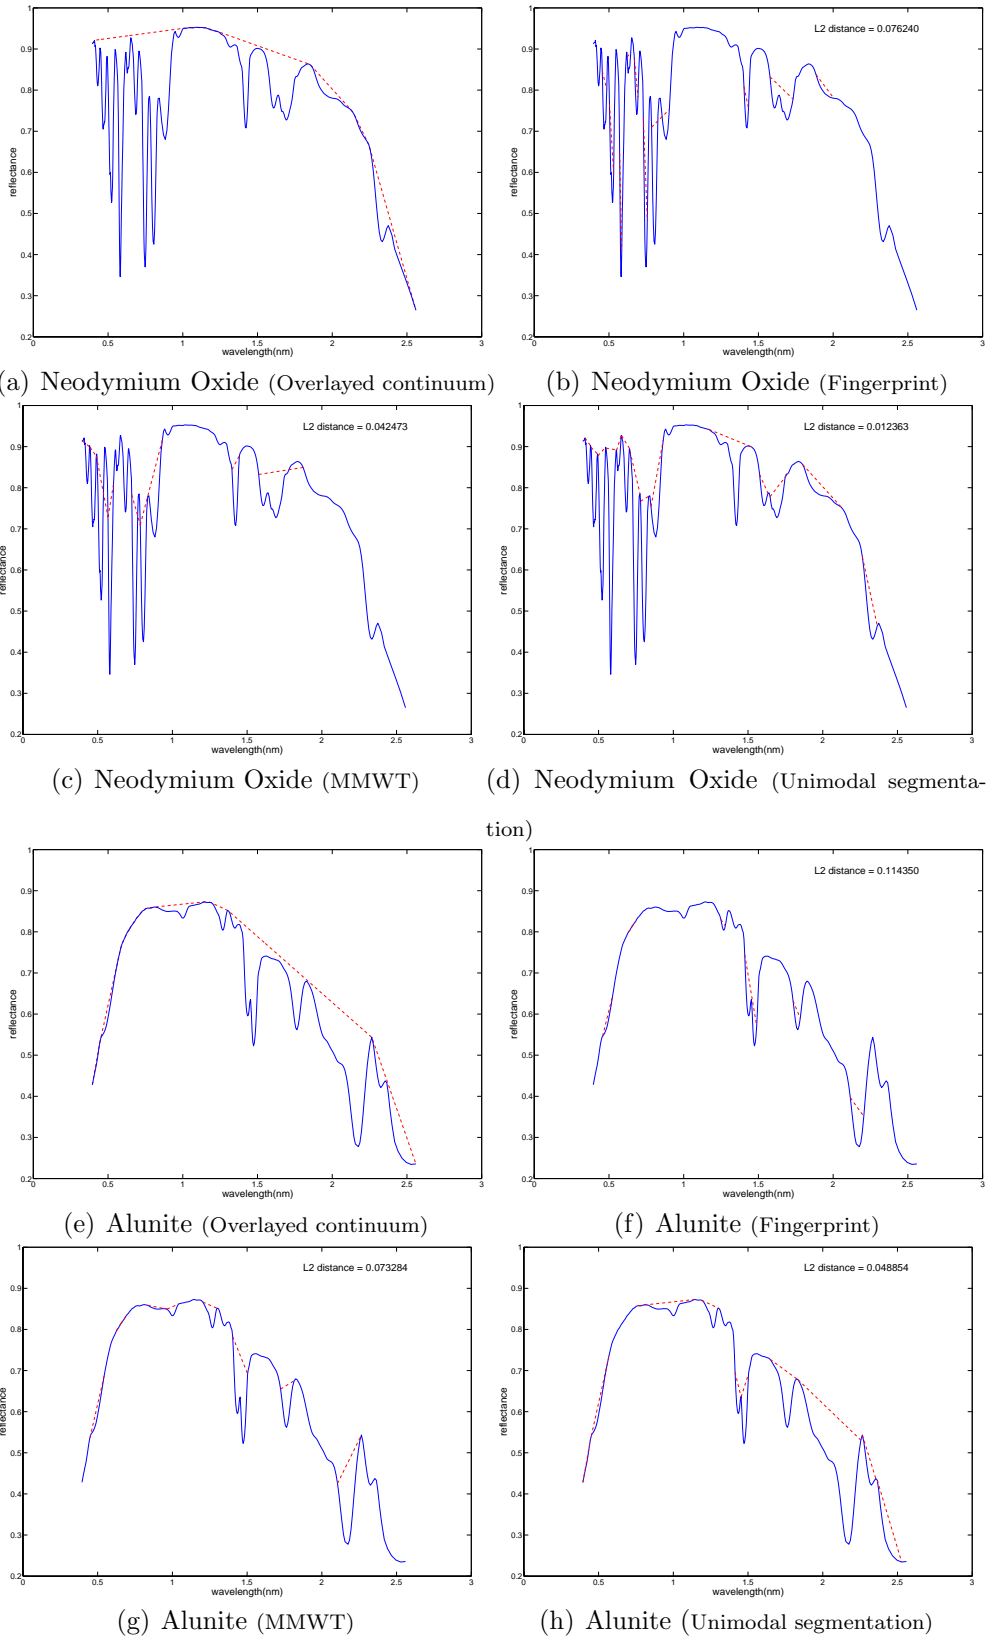


Figure 13: Examples of absorption detection results for USGS spectra

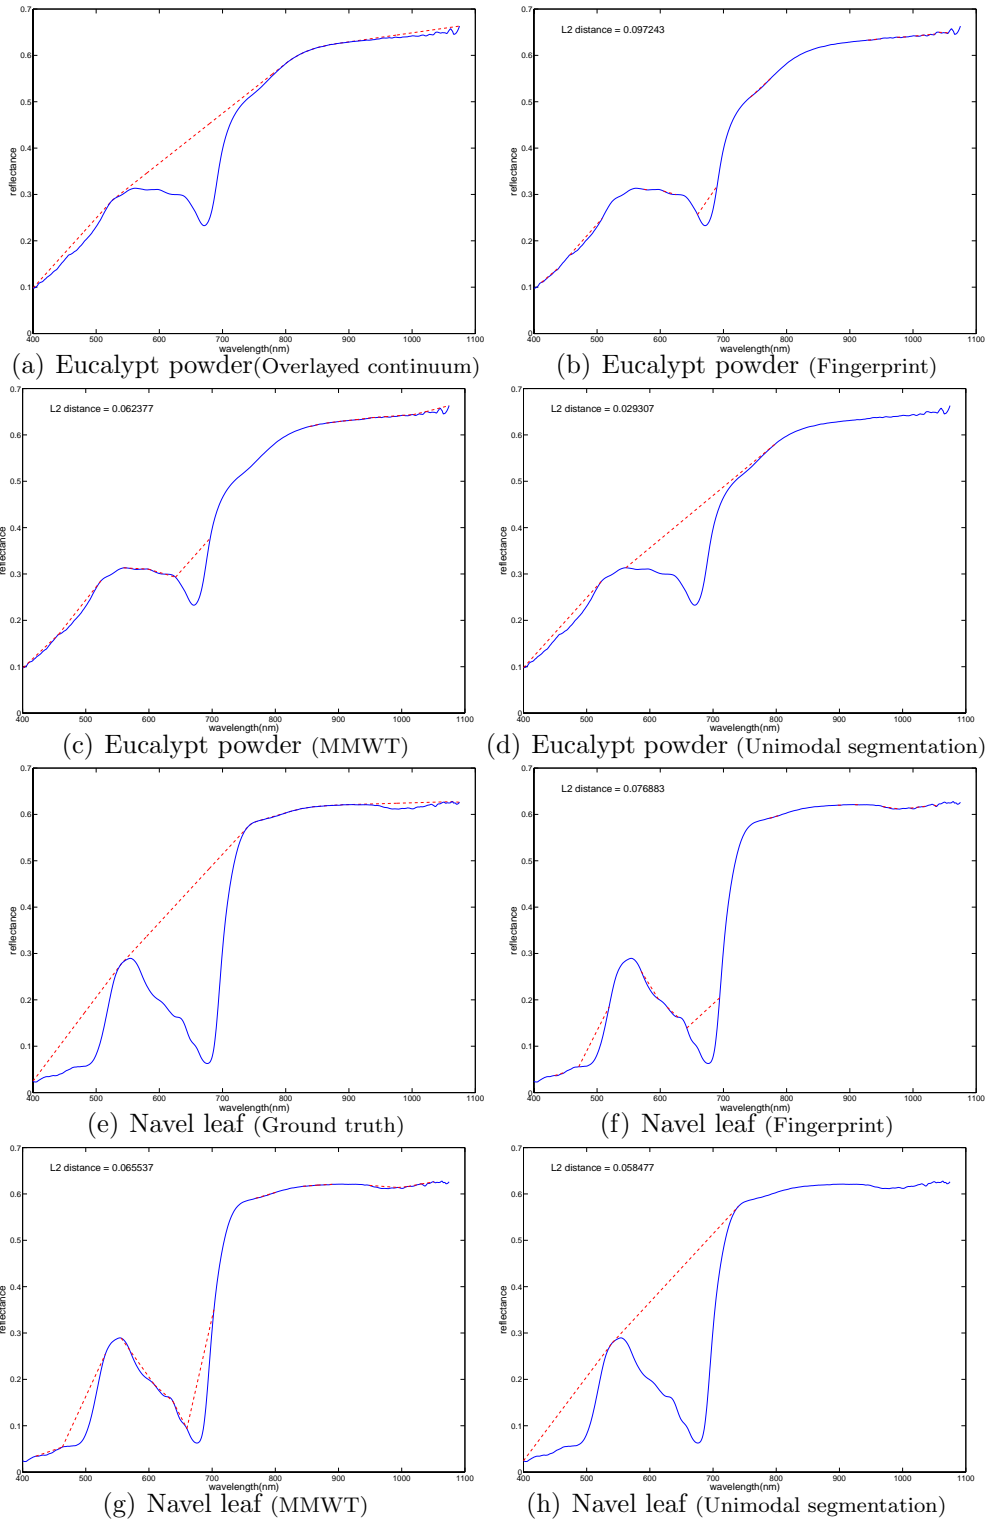


Figure 14: Examples of absorption detection results for our spectroscopy data

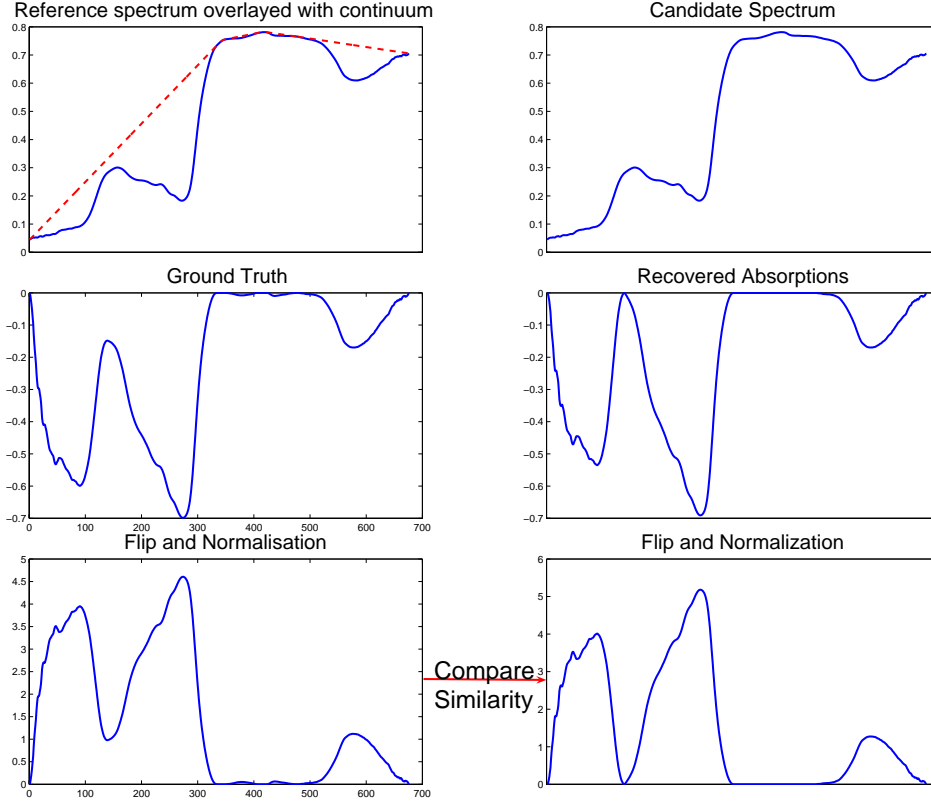


Figure 15: Steps for comparing the absorption detection results with the ground truth

In this manner, our quantitative study accounts for the quality, as well as the number, of the recovered absorption features. This process is illustrated in Figure 15. Similarity measures as taken above will penalise large deviations from the ground truth, i.e. the continuum-removed spectra yield by [7]. Further, the similarity between the ground truth and the recovered absorptions is a measure of the average matching result within each spectrum. For purposes of performance assessment, it makes more sense than enforcing a hard error measure by means of matching the spectra since it takes into account the global nature of the results rather than the local structure of the absorption recovery results.

Here, we have used four different distance measures [27]. These are the Kullback-Leibler divergence, the L2 norm, the Bhattacharya distance, and the χ^2 divergence. The smaller the distance, the higher the similarity between the ground truth and the detection results. Therefore, as the accuracy of the detection method with respect to the ground truth increases, the distance between the spectra decreases. In Table 3, we list the results of our approach and the alternatives for the spectral data from the USGS database. In the table, we provide the mean and standard deviation for the distance over all the spectra in the dataset for each of the

Distance measure	Absorption Recovery Method		
	Fingerprint	MMWT	Unimodal segmentation
KL distance	11.4630 ± 3.8084	5.9812 ± 3.3735	1.0411 ± 1.4125
Euclidean distance	0.1024 ± 0.0468	0.0621 ± 0.0295	0.0129 ± 0.0131
Bhattacharya distance	0.5038 ± 0.0858	0.3751 ± 0.0901	0.1363 ± 0.0832
χ^2 distance	0.3178 ± 0.1062	0.1856 ± 0.0895	0.0307 ± 0.0337

Table 3: Mean and standard deviation between the ground truth and the results yield by different absorption recovery methods for the USGS spectral data.

four alternatives. The results for our spectral database are listed in Table 4. From the tables, we can conclude that, as far as the different distance measures are concerned, all of them give rise to consistent evaluation results favoring unimodal segmentation followed by MMWT and Fingerprint. In contrast with the Fingerprint method, the unimodal segmentation and MMWT achieve similar results, the former being slightly better than the latter.

Distance measure	Absorption Recovery Method		
	Fingerprint	MMWT	Unimodal segmentation
KL distance	12.6738 ± 2.8432	6.1432 ± 2.3884	0.4751 ± 0.7653
Euclidean distance	0.0795 ± 0.0183	0.0608 ± 0.0123	0.0082 ± 0.0064
Bhattacharya distance	0.5604 ± 0.0625	0.4499 ± 0.0583	0.1137 ± 0.0596
χ^2 distance	0.3970 ± 0.0871	0.2781 ± 0.0692	0.0223 ± 0.0190

Table 4: Mean and standard deviation for the distances between the ground truth and the results yield by different absorption recovery methods for our spectral data

4.2 Absorption Feature based Material Mapping

Having presented comparative results for the absorption band recovery methods under study, we now turn our attention to the application of absorption band matching for material mapping in spectral images. To this end, we have used a hyperspectral image captured by the AVIRIS sensor flown over Cuprite, Nevada in 1995. The region of interest has 400 samples, 350 lines and 220 bands covering the visible, near (NIR) and short wave infrared (SWIR) spectral ranges. For our material mapping experiment, we used the last 50 bands,

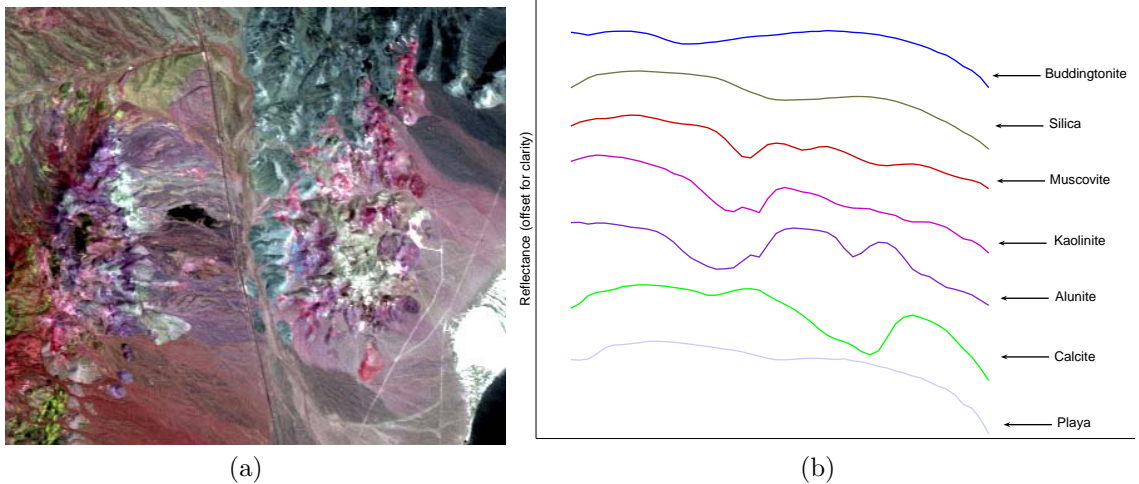


Figure 16: (a) Pseudo-color image of the scene. (b) End-members of different minerals.

from 1990.8nm to 2479nm in the SWIR spectral range. This is due to the higher signal-to-noise ratio (SNR) in the SWIR range. Radiometric calibration was performed using the ATmospheric REMoval program (ATREM) [15] so as to convert radiance spectra to apparent reflectance spectra, i.e. the true reflectance up to a constant yet unknown scaling factor. The image was then processed using the Empirical Flat Field Optimized Reflectance Transformation (EFFORT) [4] to remove residual calibration errors and atmospheric effects.

The pseudo color image of 3 different bands $2.1010\mu m$ (red), $2.2008\mu m$ (green) and $2.3402\mu m$ (blue) is shown in Figure 16(a). We used 7 reference spectra of different minerals selected from the USGS spectral library as end-members for material mapping and classification. The mineral spectra employed were the Playa, Calcite, Alunite, Kaolinite, Muscovite, Silica and Buddingtonite. In Figure 16(b), we have plotted, in a bottom-up order, the end-member spectra. In the figure, for the sake of clarity, we have offset the plots.

For purposes of pixel-level mineral mapping, we have performed experiments making use of the method in [21], the SFF method in [9] and the method proposed in Section 3 using the likelihood values yield by Equation 12. We have omitted results on our method using Equation 15 in the sake of brevity and clarity. Also, in our experience, the results yield employing Equation 12 provide a margin of improvement over those obtained using Equation 15. To ensure optimal performance, the parameters for each of the methods have been separately tuned. The corresponding results are displayed in Figure 17.

In our results, SAM yields a high number of false positives for the Playa end-member. Furthermore, it misses the detail on mineral distributions across the scene. SFF and our

method yield much better results, with the latter gaining a margin of improvement in performance. Nonetheless it exhibits better performance than the SAM, SFF maps a large area on the northeastern corner of the scene to calcite, while most calcite is distributed in the southwestern corner of the scene, as can be observed by a small green-coloured patch in Figure 16(a). The dark green color region on the northeast of the scene is actually formed by Jarosite and various other minerals, which are not included in the reference spectra. In contrast with the SAM and SFF, our absorption-based method can infer the Calcite distribution more accurately by effectively rejecting false matching. Moreover, since our method is based on the local comparison of absorption bands, whose positions are known a priori from the reference spectra, we only have to extract the spectral segment from the bands of interest for each image spectrum and remove the continuum via Equation 5, where $C(\lambda)$ here is simply a straight line joining the first and the last band in the absorption interval.

It is also worth noting that SFF is based on a global shape comparison between the continuum-removed spectra and the candidate spectrum. In order to compare the image spectra with the end-members, we have to apply continuum removal to both, the candidate and the reference spectra. This implies an additional cost of $O(n \log n)$, where n is the number of bands, for each spectrum in the image before spectral matching. This computational cost of the matching operation is of $O(n)$ for both, SAM and SFF. For our method, the computational complexity is $O(n')$, where n' is the total number of spectral samples in the absorption band of the end-member spectrum. Since n' is always less or equal than n , our method is not only more accurate, but also much more efficient than traditional mapping methods for hyperspectral image analysis.

5 Conclusion

In this paper, we have presented a survey on automatic absorption feature recovery algorithms. To this end, we have described a number of absorption recovery algorithms elsewhere in the literature, i.e. the derivative analysis, Fingerprint and MMWT methods. We have also introduced a novel method for absorption feature recovery based upon unimodal segmentation. Departing from a technical overview of the existing methods, we have presented a complexity analysis and elaborated on the implementation issues for each of the alternatives. By making use of four distance measures and ground-truth spectral data, we have provided

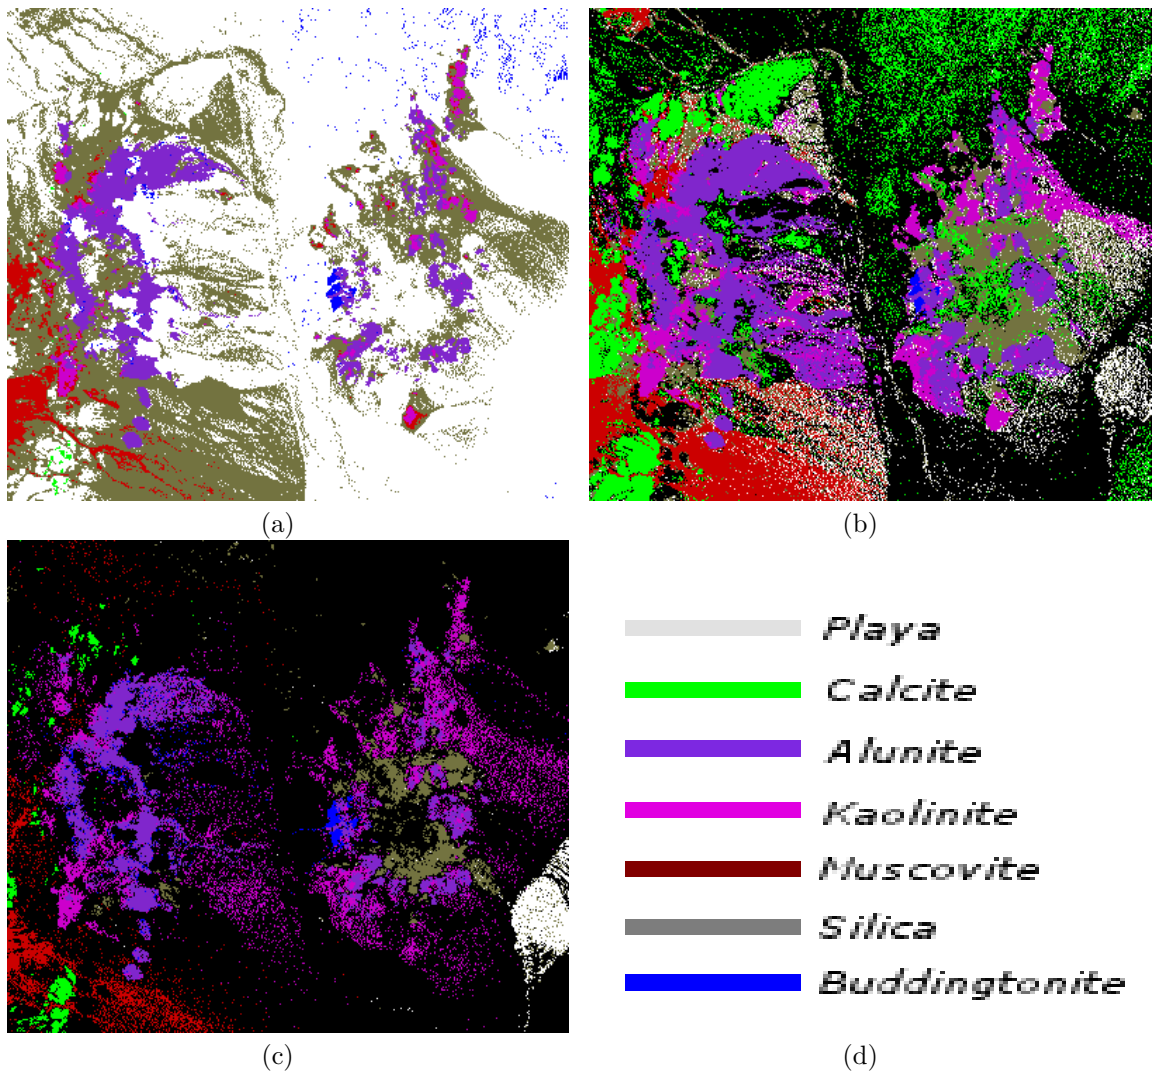


Figure 17: Material mapping results. (a) Spectral Angle Mapper [21]; (b) Results yield by SFF; (c) Our method using the likelihood yield by Equation 15. (d) Legend

a performance study which takes into account the quality of the absorption recovery results. We have also illustrated the use of absorption features for purposes of material mapping in real-world hyperspectral imagery.

The work presented here can be further extended and improved in a number of ways. Firstly, the nature of the unimodal segmentation algorithm is such that allows, in a straightforward manner, the use of Gaussian mixtures to model spectra. This opens-up the possibility of using information theoretic measures to assess and further improve the absorption band recovery process. Secondly, the Bayesian nature of our classification method permits the use of statistical pattern recognition techniques to further refine the classification results. This is not only theoretically important, but practically useful since any improvement in material mapping results has widespread applications in geoscience.

Acknowledgment

The authors are indebted with Mr. Guillaume B. Mainbourg, who kindly provided them with the implementation of the MMWT and the Fingerprint method used in the experimental section of this paper.

References

- [1] M. L. Adams, W. D. Philpot, and W. A. Norvell. Yellowness index: an application of spectral second derivatives to estimate chlorosis of leaves in stressed vegetation. *International Journal of Remote Sensing*, 20(18):3663–3675, 1999.
- [2] P. Atkinson, R. Webster, and P. Curran. Cokriging with airborne mss imagery. *Remote Sensing of Environment*, 50:335–345, 1994.
- [3] M. Ayer, H. Brunk, G. Ewing, W. Reid, and E. Silverman. An empirical distribution function for sampling with incomplete information. *Annals of Mathematical Statistics*, 26:641–647, 1955.
- [4] J. Boardman and J. Huntington. Mineral mapping with 1995 aviris data. In *Summaries of the 6th Annual JPL Airborne Research Science Workshop*, volume 1, pages 9–11, 1996.
- [5] A. Brown. Spectral curve fitting and for automatic hyperspectral data analysis. *IEEE Transaction on Geoscience and Remote Sensing*, 44(6):1601–1608, 2006.
- [6] W. L. Butler and D. W. Hopkins. Higher derivative analysis of complex absorption spectra. *Photochemistry and Photobiology*, 12:439–450, 1970.

- [7] R. Clark and T. Roush. Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research*, 89:6329–6340, 1984.
- [8] R. Clark, G. Swayze, K. Livo, R. Kokaly, S. Sutley, J. Dalton, R. McDougal, and C. Gent. Imaging spectroscopy: Earth and planetary remote sensing with the usgs tetracorder and expert system. *Journal of Geophysical Research*, 108(5):1–44, 2003.
- [9] R. N. Clark, A. J. Gallagher, and G.A. Swayze. Material absorption band depth mapping of imaging spectrometer data using a complete band shape least-squares fit with library reference spectra. In *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer Workshop*, pages 176–186, 1990.
- [10] T. H. Demetriades-Shah, M. D. Steven, and J. A. Clark. High resolution derivatives spectra in remote sensing. *Remote Sensing of Environment*, 33:55–64, 1990.
- [11] R. O. Duda and P. E. Hart. *Pattern Classification*. Wiley, 2000.
- [12] M. Dundar and D. Landgrebe. Toward an optimal supervised classifier for the analysis of hyperspectral data. *IEEE Transaction on Geoscience and Remote Sensing*, 42(1):271–277, 2004.
- [13] A. F. Fell and G. Smith. *Higher Derivative Methods in Ultraviolet, Visible and Infrared Spectrophotometry*. Anal. Proc., 1982.
- [14] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, NY, 2nd edition, 1990.
- [15] B. Gao and A. Goetz. Column atmospheric water vapor and vegetation liquid water retrievals from airborne imaging spectrometer data. *Journal of Geophysical Research*, 95(D4):3549–3564, 1990.
- [16] R. L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Information Processing Letters*, 1:132–133, 1972.
- [17] P. Hsu. *Spectral Feature Extraction of Hyperspectral Images using Wavelet Transform*. PhD thesis, Dept. of Survey Engineering, National Cheng Kung University, 2003.
- [18] L. Jimenez and D. Landgrebe. Hyperspectral data analysis and feature reduction via projection pursuit. *IEEE Transaction on Geoscience and Remote Sensing*, 37(6):2653–2667, 1999.
- [19] I.T. Jolliffe. *Principal Component Analysis*. Springer, 2002.
- [20] J Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
- [21] F. Kruse, A. Lefkoff, J. Boardman, K. Heidebrecht, A. Shapiro, J. Barloon, and A. Goetz. The spectral image processing system (sips) - interactive visualization and analysis of imaging spectrometer data. *Remote Sensing of Environment*, 44:145–163, 1993.
- [22] D. Landgrebe. Hyperspectral image data analysis. *IEEE Signal Processing Magazine*, 19, 2002.
- [23] S. Liang. *Quantitative Remote Sensing of Land Surfaces*. Wiley-Interscience, 2004.

- [24] T. Lillesand and W. Ralph. *Remote Sensing and Image Interpretation*. John Wiley and Sons, 4th edition, 2000.
- [25] D. J. C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [26] A. V. Oppenheim, R. W. Schaffer, and J. R. Buck. *Discrete-Time Signal Processing*. Prentice Hall, 1999.
- [27] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. Mc Graw Hill, 1984.
- [28] S. R. Phinn, J. Franklin, D. Stow, A. Hope, and L. Huenneke. Biomass distribution mapping in a semi-arid environment using airborne video imagery and spatial statistics. *Journal of Environmental Management*, 46(2):139–165, 1996.
- [29] M. Piech and K. Piech. Symbolic representation of hyperspectral data. *Applied Optics*, 26:4018–4026, 1987.
- [30] D. A. Roberts, K. Brown, R. O. Green, S. L. Ustin, and T. Hinckley. Investigating the relationships between liquid water and leaf area in clonal populus. In *Summaries of the 7th Annual JPL Earth Science Workshop*, volume 1, pages 335–344, 1998.
- [31] D. A. Sims and J. A. Gamon. Estimation of vegetation water content and photosynthetic tissue area from spectral reflectance: a comparison of indices based on liquid water and chlorophyll absorption features. *Remote Sensing of Environment*, 84:526–537, 2003.
- [32] W. Smith, A. Robles-Kelly, and E. R. Hancock. Skin reflectance modelling for face recognition. In *Proc. of the Int. Conf. on Pattern Recognition*, pages III:210–213, 2004.
- [33] Q. Stout. Optimal algorithms for unimodal regression. *Computing Science and Statistics*, 32, 2002.
- [34] J. Sunshine, C. Pieters, and S. Pratt. Deconvolution of mineral absorption bands: An improved approach. *Journal of Geophysical Research*, 95:6955–6966, 1990.
- [35] A. Webb. *Statistical Pattern Recognition*. Wiley, 2002.
- [36] A. Witkin. Scale-space filtering. In *Proc. Intl. Joint Conf. on Artificial Intelligence*, pages 1019–1022, 1983.