© EARSeL and Warsaw University, Warsaw 2005. *Proceedings of 4th EARSeL Workshop on Imaging Spectroscopy. New quality in environmental studies.* Zagajewski B., Sobczak M., Wrzesień M., (eds)

A MATHEMATICAL MORPHOLOGY BASED APPROACH TO LOCATING SPECTRAL ENDMEMBERS

Conrad Bielski and Pierre Soille

Joint Research Centre of the European Commission, Land Management Unit, Ispra (VA), Italy; Conrad.Bielski@jrc.it , Pierre.Soille@jrc.it

ABSTRACT

Spectral endmembers are needed in order to apply Spectral Mixture Analysis (SMA) based on imaging spectrometer data. Currently there are several techniques able to find image endmembers in N spectral dimensions and are wholly based on the spectral feature space. However, modelling only the outside 'hull' of the spectral scatter provides only a basic understanding of the spectral variability. Our ultimate goal is to model the n-dimensional spectral space from the inside, out thereby providing a map of image spectral objects. The research presented in this paper are initial findings applying mathematical morphology to imaging spectrometer data. Mathematical morphology was applied in two ways: i) by measuring the morphological hyperspectral scalar gradient at each pixel position using a structuring element of predefined size and shape and ii) by determining whether this pixel falls within the smallest enclosing hyperbox containing the pixels belonging to the structuring element centred at this position. While the results are interesting in themselves, they were unable to automatically provide the image endmembers. However, the results provide a wealth of information on the local variations in spectral profiles which will be used to model their behaviour within the spectral feature cloud.

INTRODUCTION

The ability of imaging spectroscopy to acquire spectral profiles is both a blessing and a curse for analysts. Such fine spectral resolution in the VNIR and SWIR region of the electromagnetic spectrum provides important details on the spectral characteristics of materials. However, such fine spectral resolution also generates large amounts of data that must be processed in order to extract pertinent information required by the user. Information that is relevant to many remote sensing analysts is the discovery of the proportion of materials present within a single spatially mixed pixel. Spectral Mixture Analysis (SMA) (i) as well as Multiple Endmember Spectral Mixture Analysis (MESMA) (ii) are techniques that provide the analyst with such information. Both techniques require spectral endmembers to be defined. An image endmember is a pixel found within the acquired image that is considered to be 'pure' in the sense that spectral mixing is minimised. Such pixels will generally be found at the furthest reaches of the spectral scatter (Figure 1) and current endmember searching algorithms such as the Pixel Purity Index (PPI) (iii,iv) and NFINDR (v) try to locate these extreme regions. Once the endmembers are identified, the image pixels can be unmixed using the above mentioned techniques. These techniques assume linear mixing and are able to derive fractional abundances of materials in all the image pixels based on the chosen endmembers.

In this study, we propose a spatial/spectral method of automatically deriving endmember spectral profiles based on the available imagery. This method is based on a mathematical morphology framework (vi,vii,viii) and takes the spatial neightbourhood relationships between the image pixels into account. The research presented in this paper is a stepping stone to reaching our ultimate goal of mapping the spectral scatter from the inside out thereby providing in depth knowledge of the real behaviour of spectral mixtures within a given image.

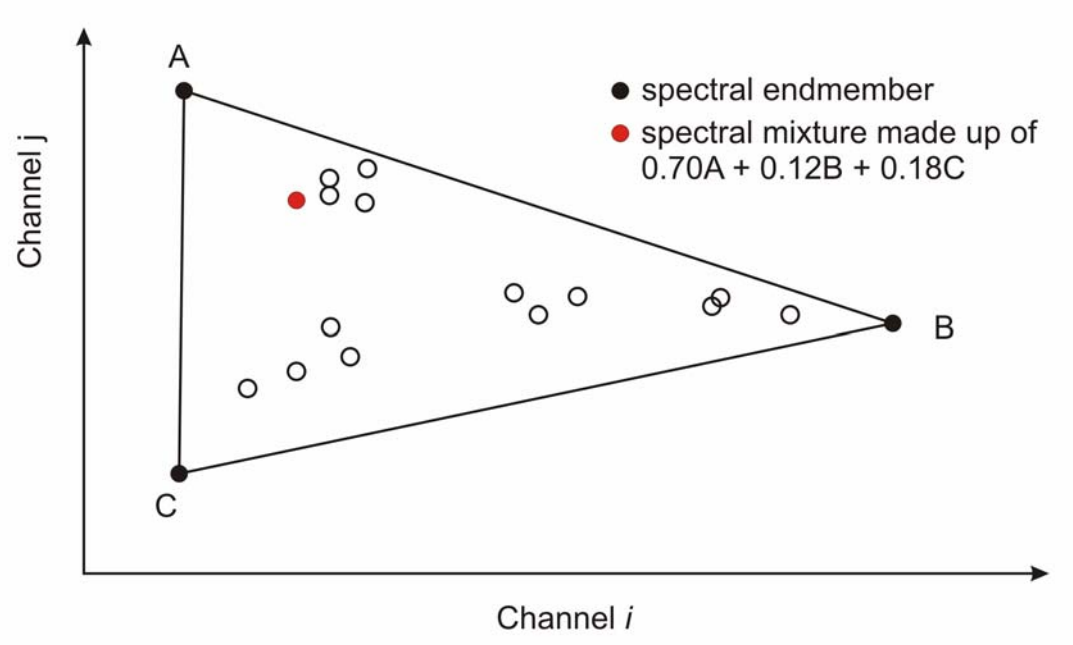

Figure 1: Idealised spectral scatter in two dimensions with image endmembers. The spectral mixture of the red point is computed based on the chosen endmembers using a linear mixture model. Within the scatter, there are regions where spectra will group together and are related to real objects with a specific spectral characteristic.

METHODS

Two techniques were tested based on mathematical morphology in order to see whether they could be used for the automatic extraction of image endmembers. Mathematical morphology is the study of spatial structures and is applied as an image analysis technique that examines the shapes and forms of objects by analysing spatial relationships between a set of predefined neighbouring pixels called structuring elements (6). In this study, mathematical morphology is applied in two ways: i) by measuring the morphological hyperspectral scalar gradient at each pixel position using a structuring element (SE) of predefined size and shape and ii) by determining whether this pixel falls within the smallest enclosing hyperbox containing the pixels belonging to the SE centred at this position. Both approaches are detailed hereafter. It was reasoned that these two techniques could be used to describe local spatial variability within the spectral cloud; as such they would also provide information on the locations of the image endmembers. We denote by f_i a given mono-channel (single valued) image $\mathbf{f}$ a multichannel image, i.e. $\mathbf{f} = (f_1, ..., f_n)$ where n refers to the number of channels.

The hyperspectral scalar gradient

Various techniques have been developed for combining the components of the vector gradient in order to create a scalar gradient image. Usually, the sum or the maximum of the magnitude of the gradient components is considered (ix). Rather than combining the different gradient components, we define the hyperspectral scalar gradient ρ_s at a pixel $\mathbf{x}$ of a multispectral image $\mathbf{f}$ with a SE B as the largest Euclidean distance d_ε separating $\mathbf{f}(\mathbf{x})$ from its neighbours $\mathbf{f}(\mathbf{x}+\mathbf{b})$ where $\mathbf{b} \in B$:

$$\left[\rho_B(\mathbf{f})\right](\mathbf{x}) = \max_{\mathbf{b}\in B} d_\varepsilon\left[\mathbf{f}(\mathbf{x}), \mathbf{f}(\mathbf{x}+\mathbf{b})\right]. \qquad (1)$$

This definition is illustrated in Figure 2. In the case of a panchromatic image, the proposed gradient gives the maximum intensity variation between the pixel and its neighbours.

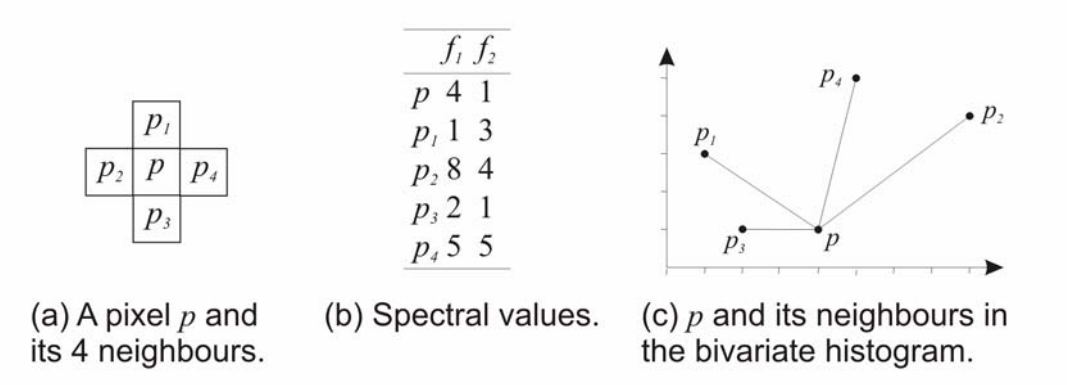

Figure 2: Multispectral scalar gradient ρ_s at pixel p of a multispectral image $\mathbf{f}$ having two components f_1 and f_2. According to Equation 1, the gradient at p equals d_ε [$\mathbf{f}(p)$, $\mathbf{f}(p_2)$] in this example. (Adapted from (6)).

The hyperbox

The main assumption behind current endmember detection algorithms such as the PPI is that the endmembers lie at the 'corners' of the convex hull of the image data values projected into a Euclidean space whose dimensionality equals the number of available spectral channels. We also use the notion of the convex hull but we compute it locally for a set of pixels falling within a given SE centred at the pixel being analysed.

The convex hull of a set can be defined in terms of morphological closing using half-planes as SEs. The convex hull of a set corresponds to the intersection of all half-spaces containing this set because the intersection of all half-spaces of a given orientation and containing a set corresponds to the morphological closing of the set using the underlying half-plane as SE. It follows that convex hulls of a given order can be defined by using a restricted set of orientations (x). For example, in the 2-D Euclidean space, the smallest rectangle enclosing a set defines the convex hull of order 0 and is obtained by performing the intersection of all vertical and horizontal half-planes containing this set, or, equivalently, the intersection of the morphological closing with all four possible vertical and horizontal half-planes. That is, in a n-dimensional space, the convex hull of order 0 of a set corresponds to the smallest hyperbox whose sides are parallel to coordinate axes and which contain this set. From herein the convex hull transformation of order 0 is denote by CH_0. Higher orders are not considered in this paper given the difficulty of computing convex hulls in spaces of more than 2 dimensions.

The methodology proceeds as follows. Given a pixel $\mathbf{x}$ of a multichannel image $\mathbf{f}$ and a set of neighbouring pixels defined by a SE B, we propose to determine whether $\mathbf{x}$ falls within or outside the convex hull of order 0 of the neighbours of $\mathbf{x}$ mapped in the n-dimensional space. This result is stored in an indicator function χ which is obtained as follows:

$$\chi(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{f}(\mathbf{x}) \in \mathrm{CH}_0\left[\bigcup_{\mathbf{b}\in B} \mathbf{f}(\mathbf{x}+\mathbf{b})\right], \\ 1, & \text{otherwise.} \end{cases} \qquad (2)$$

Note that the SE B should not contain its origin because the indicator function χ would always return 0.

Owing to the orthogonality of the boundaries of the convex hulls of order 0, we may perform calculations one dimension at a time by checking, for each individual channel i, whether the value of a given pixel $\mathbf{x}$ falls within or outside the range defined by the SE centred at this pixel:

$$\chi_i(\mathbf{x}) = \begin{cases} 0, & \text{if } \min_{\mathbf{b}\in B} f_i(\mathbf{x}+\mathbf{b}) < f_i(\mathbf{x}) < \max_{\mathbf{b}\in B} f_i(\mathbf{x}+\mathbf{b}), \\ 1, & \text{otherwise.} \end{cases} \qquad (3)$$

In practice, we are also interested in counting the number of times a given pixel is outside the range defined by its neighbours in each successive channel. We denote the resulting image by $\mathbf{X}$:

$$X(\mathbf{x}) = \sum_{i=1}^{n} \chi_i(\mathbf{x}). \qquad (4)$$

Note that $\mathbf{x}$ falls outside the hyperbox (convex hull of order 0) defined by the values falling within the SE B centred at $\mathbf{x}$ if and only if $\mathbf{X}(\mathbf{x}) = 0$:

$$\chi(\mathbf{x}) = \begin{cases} 0, & \text{if } X(\mathbf{x}) = 0 \\ 1, & \text{otherwise.} \end{cases} \qquad (5)$$

The trials were performed on a subset of imaging spectrometer data acquired by the Airborne Visible and Infra-Red Imaging Spectrometer (AVIRIS) over Cuprite, NV, USA on June 19, 1997 (xi). The imagery was corrected to ground reflectance and bad bands were excluded prior to processing (bad bands = 369.85 – 419.06, 1343.95 – 1443.53, 1801.92 – 1958.59, 2477.08 – 2506.81 nm).

RESULTS

The hyperspectral scalar gradient was computed for the Cuprite image and is shown in Figure 3a. The SE used to compute these results looked at the four closest pixels to point p (as shown in figure 2a).

Figure 3b presents the results of computing the hyperbox. The hyperbox results are presented as the number of times that a pixel was found outside the CH0. Therefore, the higher the sum the more times the pixel p was considered as 'extreme' with respect to the SE. Along with computing the CH_0, derivative results were also computed in order to understand the local variability of the spectral scatter. The lengths of each of the sides of the hyperbox were summed together and used as a proxy for the volume of the hyperbox (Figure 3c). These results summarise and try to encapsulate local spectral behaviour in a single image. However, each spectral band that was processed also contains the interim results which are also of interest. For example the minimum, maximum and range of CH_0 for each spectral dimension was also computed and shows the variability in n-dimensional spectral space which may or may not be possible to describe in a single band.

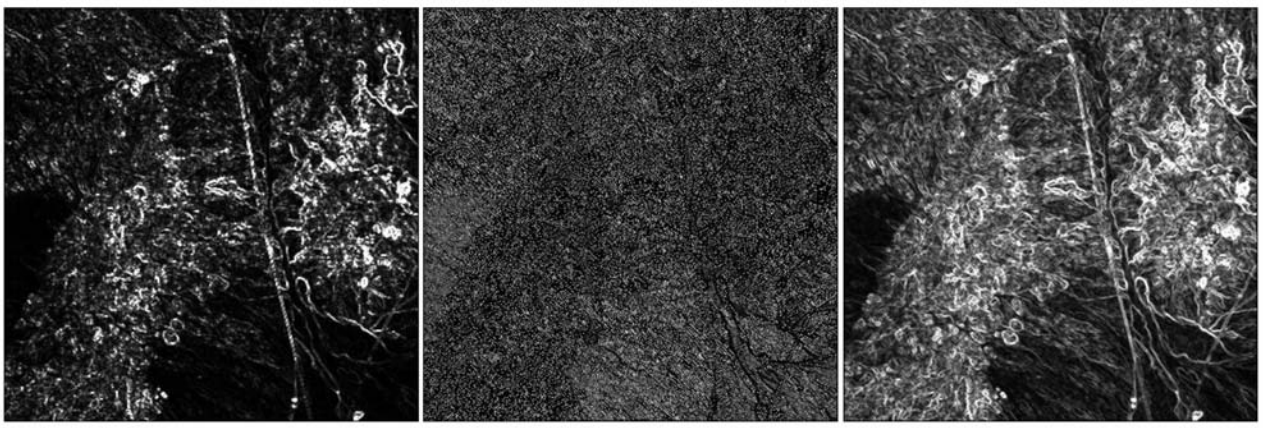

Figure 3: Results of computing the hyperspectral scalar gradient and hyperbox.

DISCUSSION OF RESULTS

The results presented above are interesting however in their present form were unable to locate the spectral endmembers automatically. The hyperspectral scalar gradient image (Figure 3a) exposes areas in the imagery where there are significant spectral differences between local pixels. The hyperspectral scalar gradient values varied from 13.31 to 79575.44. While the majority of the points (95%) where below the 4381.43 level, at this time there is no way of knowing based on these results what the cut-off should be between deciding whether a pixel is similar or not to the neighbour. The spatial distribution of the values does however provide an overview of the expected locations of spectrally similar pixels. For example, structures in the imagery that are most likely objects of interest such as the road are apparent. Based on this information, our next focus based on the boundaries created by the hyperspectral scalar gradient will be to see whether the objects found within the bounds can be distinguished as a similar material. The fact that the majority of image pixels have a small hyperspectral scalar gradient is most likely an indication that most pixels are spectrally mixed pixels because you would expect low values of the hyperspectral scalar gradient to be found where little spectral change occurs. However, based solely on local variability it would be impossible to tell without doubt that those regions where high values of the hyperspectral scalar gradient are found are regions where the image endmembers would most likely be found. In order to find the image endmembers the hyperspectral scalar gradient must be scaled to the entire image thus requiring a more global approach.

The results based on the hyperbox computations also presented us with a better understanding of the ability of this technique as well as the complexity of the inner spectral variability of an image's spectral scatter. However, such a locally based approach was again unable to find the image endmembers in the present state of the methodology.

The image in Figure 3b shows the sum of the number of times the pixel p was outside the hyperbox. A total sum of 185 (i.e. the number of spectral bands in the image) would indicate that the pixel p was always outside the bounds. In a local sense this is an extreme pixel however the results at the moment cannot be scaled so one could not say that those same pixels are 'globally extreme'. On the other hand, when the sum is equal to zero one would expect that locally the pixel p would be a mixture. This is not the case because there are instances where an 'extreme' pixel could be considered locally as never falling outside the bounding hyperbox.

To illustrate this problem, Figure 4 presents a schematic diagram of a hyperbox in two dimensions. Based on a four element SE as was used in this study, the hyperbox was generated based on the SE. The example presents p_4 and p_1 as the extreme values of the bounding box in the first dimension and p_3, p_4 in the second dimension. The point p therefore has three different possibilities: i) in both dimensions point p is found within the bounds and graphically is within the solid square where x=0, ii) in one dimension the point p is found outside the bounding square while in the other dimension it is found inside where x=1 and iii) where point p is found outside the bounds in both dimensions where x=2. Intuitively, one would expect that if a local mixture is found, i.e. a pixel p that is equal to x=0, one would also have a global mixture. However, the difficulty in finding the image endmembers based on the hyperbox technique occurs (in the two-dimensional case) when the SE pixels are greater in the lower left and upper right directions compared to pixel p which could be somewhere in the upper left or lower right position of the hyperbox. In such a case, pixel p will never be outside the hyperbox even though it may lie on the hull of the spectral scatter. A two band image is shown in Figure 5 to clarify this problem.

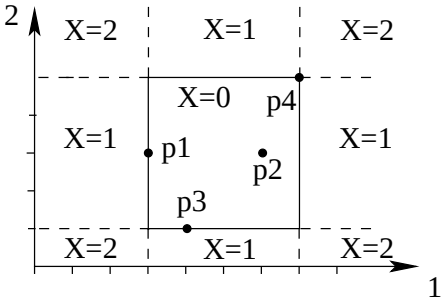

Figure 4: In this example, the SE B is defined by the 4-neighbours of a pixel. These neighbours are mapped in 2-D feature space (assuming the number of channels n equals 2). The analysed pixel is then mapped in this space and depending on its position, the value of X can range from 0 to 2.

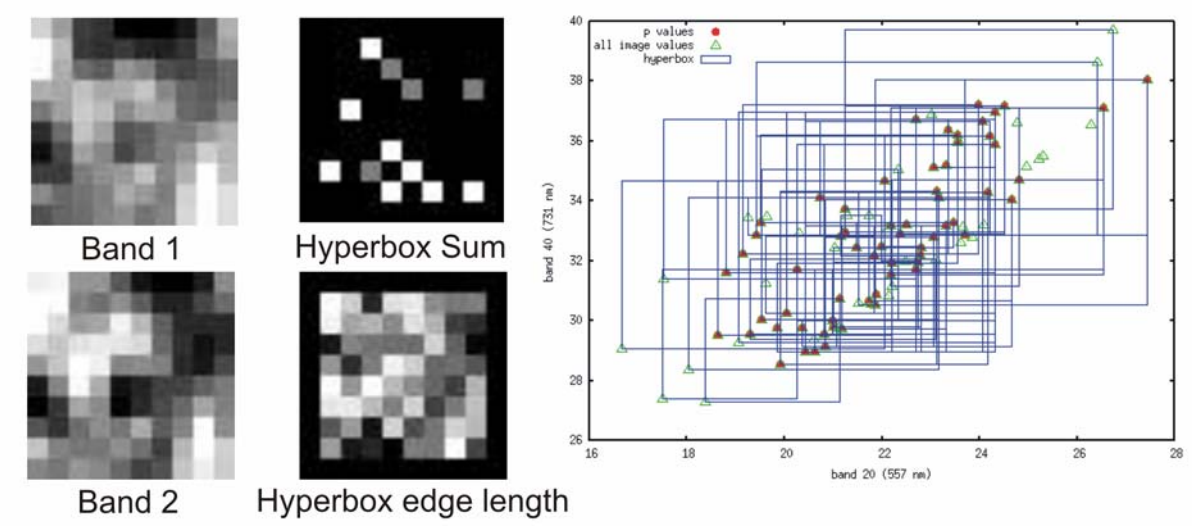

Figure 5: Two band test image. Leftmost pair of images are reflectance images. Above centre presents the results of computing the hyperbox with X ranging from 0 (pixel falls in the defined hyperbox) to 2 (pixel lies in the corner domains of the hyperbox). Lower centre image shows the hyperbox edge length for each pixel. Right graph shows the hyperboxes computed for the entire image.

The sample image is made up of two bands of the Cuprite scene with dimensions of 10 by 10 pixels. The plot on the right side of figure 5 shows the results of computing hyperboxes for this two band example. The hyperboxes do a fine job of delineating the extreme pixels located in the bottom left and upper right corners of the spectral scatter. However, the pixels that expand the scatter in the other directions (i.e. towards the upper left and lower right corners) are not delineated properly because the bounding shape is rectangular. Furthermore, the bounding boxes are based on local pixel variability which cannot be extrapolated to the global variability of the images spectral scatter. Therefore, the search for the image endmembers must include scaling of the hyperbox technique the entire image.

CONCLUSIONS

The current automatic and semi-automated endmember search techniques provide the user with global image spectral endmembers. While this is the required information for applying SMA, it does not provide the analyst with information about the spectral variability within the N-dimensional spectral scatter. For our research, we require knowledge of the image spectral endmembers as well as the ability to map spectral variability within the spectral scatter. The research presented showed that describing n-dimensional spectral variability is difficult and that relating such variability in spectral and spatial domain space is not trivial. However, the use of local SEs, the spectral scalar gradient and the computation of hyperboxes does provide initial information that can be used to

reach our goal. Computing hyperboxes and hyperspectral scalar gradients are fast and therefore useful when dealing with large data sets such as imaging spectroscopy data.

Future research will concentrate on scaling the current method to the entire image. This will most likely involve subsequent levels of processing that will be based on the hyperbox sum and edge length variables.

References

i Adams, John B., Smith, Milton, O. And Gillespie, Alan, R., 1993. Imaging Spectroscopy: Interpretation Based on Spectral Mixture Analysis, chapter 7 (Cambridge University Press) Pieters and Englert eds.

ii Roberts, D. A., Gardner, M., Church, R., Ustin, S., Scheer, G., and Green, R. O., 1998. Mapping chaparral in the Santa Monica Mountains using multiple endmember spectral mixture models, Remote Sensing of Environment, 65: 267-279.

iii Boardman, J. W., 1993. Automated spectral unmixing of AVIRIS data using convex geometry concepts. In Fourth JPL Airborne Earth Science Workshop, (JPL, Pasadena) 11-14.

iv Boardman, J. W., Kruse, F. A. and Green, R. O., 1995. Mapping target signatures via partial unmixing of AVIRIS data. In Fifth JPL Airborne Earth Science Workshop, (JPL, Pasadena) 23-26.

v Winter, M. E., 1999. N-FINDR: An algorithm for fast autonomous spectral end-member determination in hyperspectral data. In Proceedings SPIE Imaging Spectrometry V, (SPIE) 266-275.

vi Soille, P., 2004. Morphological Image Analysis: Principles and Applications, (Springer-Verlag), corrected 2nd printing of the 2nd edition.

vii Soille, P. and Pesaresi, M., 2002. Advances in mathematical morphology applied to geoscience and remote sensing, IEEE Transactions of Geoscience and Remote Sensing, 40(9): 2042-2055.

viii Plaza, A., Marinez, P., Perez P., and Plaza, J., 2002. Spatial/spectral endmember extraction by multidimensional morphological operations, IEEE Transactions on Geoscience and Remote Sensing, 40(9):2025-2041.

ix Di Zenzo, S., 1986. A note on the gradient of a multi-image, Computer Vision, Graphics, and Image Processing, 33:116-125.

x Soille, P., 2000. From binary to grey scale convex hulls, Fundamenta Informaticae, 41(1-2):131-146.

xi Imagery downloaded from the AVIRIS website http://aviris.jpl.nasa.gov