EXTRACTING SURFACE FEATURES IN MULTISPECTRAL IMAGERY

M J Carlotto

The Analytic Sciences Corporation
1 Jacob Way
Reading, Massachusetts 01867 USA

ABSTRACT

A method for recognizing surface features in multispectral imagery is described. The problem is formulated in two steps. First, the surface material composition (SMC) is inferred using a combination of spectral shape and histogram analysis techniques. A symbolic description is then computed in which each connected area in the SMC map is represented by a node in a graph structure. Surface features are represented by a set of constraints over groupings of connected areas (nodes in the graph structure) in the form of rules. Examples illustrating the extraction of selected surface feature classes from Landsat Thematic Mapper imagery are included.

Key words:
Multispectral classification, feature extraction, image understanding, image processing, artificial intelligence, structural pattern recognition

1. INTRODUCTION

The identification of surface features in aerial imagery is an important activity in applications such as cartography, land use monitoring/planning, and agriculture. A variety of techniques for multispectral image classification have been developed by the remote sensing community since the launch of Landsat in 1972 (Ref. 1). On the other hand, the development of scene analysis techniques has taken place almost exclusively in the image understanding community (Refs. 2 and 3). The purpose of this paper is to draw together methodologies from the two communities for the purpose of recognizing surface features such as built-up areas, forests, agriculture, bodies of water, and roads from multispectral imagery acquired from satellites such as Landsat and SPOT.

The organization of the paper is as follows. Section 2 formulates the feature extraction process as a transformation from the multispectral image to an intermediate physical representation (surface material map), followed by a transformation from the physical representation to a semantic interpretation. Section 3 discusses methods for inferring surface material composition using spectral shape and histogram analysis techniques. Section 4 discusses how a symbolic description of the imaged scene is built and how surface feature classes can be extracted using rules. In Section 5 preliminary results for Landsat TM imagery are presented.

2. MULTISPECTRAL IMAGE UNDERSTANDING

The problem of recognizing semantically-significant features in aerial imagery such as built-up areas, forests, agriculture, and bodies of water may be viewed as a two-step process which transforms the sensed imagery (signal level) into an intermediate (physical level) representation, which is then transformed into a semantic-level interpretation of the imaged scene. An image understanding (IU) approach is characterized as one in which representations of a scene are built at a variety of levels (i.e., signal, physical, and semantic). Thus an IU approach does not attempt to directly infer meaning from signal information, rather it builds an intermediate level representation that provides a physical (or perceptual) basis to drive subsequent (semantic) processing.

IU systems can build intermediate level representations in a variety of ways. Early systems which used b&w imagery segmented the image into regions having similar grey-level or texture, or into edges. Since image intensity depends on illumination, shadowing, albedo, and the reflectance properties of surface materials, it is difficult to extract regions or edges that correspond to physically distinct phenomena in the imaged scene. The single band of information provided by a b&w image precludes the use of statistical inferencing techniques developed in the remote sensing community for surface material classification. This is one reason why attention in the IU community tends to focus on methods for representing objects in terms of their structure rather than their composition.

Here we seek to build an intermediate representation which describes the surface material composition (SMC) of the imaged scene. Computing a surface material map is possible when multispectral (or color) information is available using statistical inferencing techniques (e.g., supervised statistical classifiers). The surface material map provides a physical basis for building a symbolic representation of the scene in terms of its composition. Once built, such a representation provides the means for describing surface feature classes as sets of constraints over groupings of connected areas in the SMC map. The method discussed in this paper is thus similar in concept to the 2d image understanding systems developed at the University of Massachusetts (Ref. 3) and Kyoto University (Ref. 2).

3. SURFACE MATERIAL CLASSIFICATION

The first step in the feature extraction process is to infer surface material classes such as water, vegetation, bare soil, crops, and concrete from the multispectral imagery. The objective is to partition the image into general surface material categories which serve as the primitives for representing and describing surface features in terms of their constituent parts. This section describes the use of rule-based surface material classification techniques which are based on the analysis of multispectral data across bands (spectral shape analysis), and across intensity within a band (histogram analysis).

3.1 Spectral shape classification

In traditional multispectral classifiers, surface material classes are represented in terms of statistics such as the mean and covariance. Due to the confounding effects of illumination, albedo, and surface orientation, the spectral signature may vary considerably within a region that is composed of the same surface material. As a result classifiers which use the original multispectral imagery or linear combinations of multispectral bands (e.g., as in a principal components transformation) will introduce error in areas that are shadowed or have significant topographic relief.

These observations motivated the development of an alternative approach for recognizing surface materials by the shape of their spectral signature (Ref. 4). Rules are used to describe each surface material class in terms of the relative intensity between pairs of bands. For example, the rule for clear water

IF: (band 1 > band 2) & (band 2 > band 3) ....
THEN: (SMC = water)

reflects the qualitative observation that the reflectivity of clear water decreases from the visible towards the infrared.

The classification process involves segmenting the multispectral image into areas that have the same spectral shape. An inference engine then applies a set of rules to classify each area based on its spectral shape. The inference engine evaluates the left hand side (lhs) or "if-part" of each rule in the ruleset. If the lhs is true, the assertion in the right hand side (rhs) or "then-part" of the rule is made. Multiple rules may fire if a spectral shape matches more than one lhs (i.e., if there is overlap between classes), or no rule may fire if a spectral shape does not match any lhs. In the current implementation, each spectral shape is classified by selecting the rule with the highest match score (defined as the ratio of the number of band relations that are true in the lhs to the total number of band relations in the lhs of the rule).

Preliminary results indicate that the error rate (number of misclassifications within a training set) is comparable to a minimum distance classifier. However, an important difference in performance is that the spectral shape classifier tends to confuse materials that are similar in spectral shape (e.g., sparse and dense vegetation) while the minimum distance classifier tends to confuse materials that are quite different such as water and asphalt because they are dark. Thus the performance of the spectral shape classifier appears to degrade gracefully in the presence of confusing information.

3.2 Histogram analysis

It may not always be possible to separate SMCs based on just the spectral signature. For example, SPOT has only three multispectral bands and so provides only six unique spectral shape possibilities. Even for sensors with a greater number of bands such as the Landsat Thematic Mapper, certain SMCs may not always be separable in terms of spectral shape, e.g., crop fields from other kinds of vegetation at certain points in the growing season.

Histogram analysis complements the analysis of spectral shape for surface material classification by allowing prior knowledge concerning the relative intensity between materials in a particular band to be factored into the classification process. For example, to subclassify vegetation into crops the following procedure may be used:

- Histogram vegetation in band 4
- Find at least two modes in the histogram
- Assume that if crops are present, the brightest mode corresponds to crops
- Label the brightest mode as crops if the following constraints are satisfied:
  100 < mean < 200
  20 < standard deviation < 50
  relative frequency of occurence > 10%.

The identification of bright and dark modes is accomplished with a histogram analysis technique that decomposes histograms into sums of normal distributions (Ref. 5). Initial SMC hypotheses are generated by the spectral shape classifier. The histogram analyzer applies an independent ruleset to test these hypotheses by examining the relative intensities between classes within a particular band.

The above techniques have been embedded into an interactive computer program that allows a user to build hierarchical classification trees using combinations of spectral shape and histogram analysis techniques (Ref 6).

4. SURFACE FEATURE RECOGNITION

Once the surface material map has been derived, a physical basis exists for representing and extracting surface features. For example if trees can be extracted, then a simple rule for extracting forests might be to look for connected areas made up of trees that are greater than a certain number of pixels in size. The following subsections describe methods for extracting areal and linear features in Landsat TM imagery.

4.1 Areal Feature Extraction

For areal features (e.g., bodies of water, forests, urban areas), the surface material map is preprocessed using a local mode filter in order to eliminate isolated pixel "noise" and thin areas (e.g., roads). Connected areas in the preprocessed surface material map are labeled and geometrical and relational attributes for each connected region are computed. (The preprocessing step significantly reduces the number of connected areas in order to speed up subsequent symbolic processing.) For each connected area in the image, an instance of a symbolic data structure called a token is created for storing the attributes of the corresponding area in the image. For example,

```
#<TOKEN 32652117> :
label                     6561
xbar                      66
ybar                      412
area                      283
compactness               0.74
smc                       plowed-field
surface-feature-class     nil
adjacent-tokens           ( #<TOKEN 32655351> )
```

describes a connected area in the surface material map whose composition is plowed-field that is adjacent to another area that happens to be a crop-field. A list of such tokens constitutes the symbolic description of the image which may be depicted schematically as a graph (Fig. 1). As shown in the figure, surface features which consist of groupings of connected areas in the image correspond to groupings of tokens (subgraphs) in the symbolic description.

Surface feature classes, like surface material classes, are represented in the form of rules which contain a left hand side (if part) and a right hand side (then part). If the left hand side is true, the conclusion(s) in the right hand side are made. A simple rule for recognizing forested areas may be stated as follows:

> "If a connected area is composed of trees and is greater than 100 pixels in area, then it is a forest".

A more complicated rule for recognizing composite features (i.e., features consisting of a grouping of connected areas) such as an agricultural regions my be stated as follows:

> "If a connected area is crop-field (plowed-field) and is adjacent to at least one area that is plowed-field (crop-field), then it is an agricultural area."

Since connected areas may be assigned to more than one surface feature class, the class with the greatest number of constraints (i.e., the greatest number of conditions in the left hand side) is favored in the current implementation.

4.2 Linear Feature Extraction

For linear features such as roads, a multisource knowledge integration technique is used which computes minimum cost road networks. Since the resolution of sensors such as Landsat and SPOT is about the same as most roads, it is difficult to extract connected road networks. An extension of the F* algorithm (Ref. 7) has been developed for extracting connected road networks. The F* algorithm computes minimum cost paths through cost arrays where the cost at a point in the image is inversely proportional to the likelihood that a road is present at that point. In the F* algorithm, evidence about roads is provided by two basic types of operators:

- Type 1 - low false alarm rate but a high miss rate for roads, and
- Type 2 - low miss rate but a high false alarm rate for roads.

Type 2 information provides the framework which allows partial road segments produced by Type 1 operators to be filled in. Minimum cost paths are computed from starting and stopping points (provided by Type 1 operators) in a cost array derived from Type 2 information. Since in this application roads are typically not much wider than a pixel and tend to be lighter than the background in the visible, Type 2 information was provided by convolving a visible band with a Laplacian edge detector. The cost array is then modified using the surface material map as an independent source of knowledge. If, at that point in the surface material map, the SMC is road-like (e.g., concrete or asphalt), the cost is set to a low value; if the SMC is water or any other kind of material that cannot support a road or be confused with a road, the cost is set to a very high value; if the SMC is one that could obscure a road or be construed as a road-like material (e.g., bare earth), the value of the Type 2 operator at that point is used as the cost.

5. LANDSAT THEMATIC MAPPER EXAMPLE

The Landsat TM image in Fig. 2 was acquired over Leavenworth, Kansas. The image (512 x 512 pixels) was initially segmented into 83 areas based on spectral shape. The segmentation step reduces the computational complexity three orders of magnitude from 512^2 to 83. A rulebase was then applied to initially classify these areas into one of ten classes based on spectral shape: still water (ponds), turbid water (river), trees, vegetation, plowed field, sparse vegetation, bare ground, concrete, asphalt, and crop field. Crop-fields could not be separated from vegetation at this point in the classification process due to the similarity in the shape of the their spectral signatures (Fig. 3).

Histogram analysis was then performed and an independent ruleset was applied to refine the initial classification map. At this point, crops were separated from vegetation as was discussed in Section 3.2. The refined surface material map is shown in Fig. 4 and clearly shows the separation between crop fields and vegetation.

Areal and linear features were then separated by preprocessing the surface material map. A symbolic representation of areal features in the surface material map was derived and rules for recognizing built-up areas, bodies of water, forests, and agricultural areas were then applied. The F* algorithm was then used to derive connected road networks as outined in Section 4.2. The results in Fig. 5 show areal features overlayed on the road network extracted by the F* algorithm.

6. SUMMARY

A two step method for recognizing surface features in multispectral imagery was described and tested on Landsat TM imagery. First, surface material composition was inferred from multispectral imagery using a combination of spectral shape and histogram analysis techniques. Second, a symbolic description of connected areas in the surface material map was computed and used as the basis for inferring areal surface features using a rule based approach. This result was overlayed on a connected road network derived by the F* algorithm.

On-going work is involved with testing the above system on a wider range of imagery, and on integrating collateral sources (e.g., terrain elevation data) for identifying features such as marshes and flood plains.

References

1) M. J. Carlotto, "Techniques for multispectral image classification," Society of Photo-optical Instrumentation Engineers, Vol. 528, Jan. 1985.
2) M. Nagao and T. Matsuyama, A Structural Analysis of Complex Aerial Photographs, Plenum Press, 1980.
3) A. R. Hanson and E. M. Riseman, "VISIONS: A computer system for interpreting scenes," in Computer Vision Systems, Academic Press, 1977.
4) M. J. Carlotto and V. T. Tom, "Surface material classification based on spectral shape," IGARSS, Amherst, MA, Oct. 1985.
5) M. J. Carlotto, "Histogram analysis using a scale-space approach," Computer Vision and Pattern Recognition Conference, San Francisco, CA., June 1985.
6) R. D. Ferrante, M. J. Carlotto, J. L. Pomarede, and R. A. Upton, "Multispectral image analysis system," First Conference on Artificial Intelligence Applications, Denver, CO., Dec. 1984.
7) M. A. Fischler, J. M. Tenenbaum, and H. C. Wolf, "Detection of roads and linear structures in low-resolution aerial imagery using a multisource knowledge integration technique," Computer Graphics and Image Processing, Vol. 15, pp 201-223, 1981.

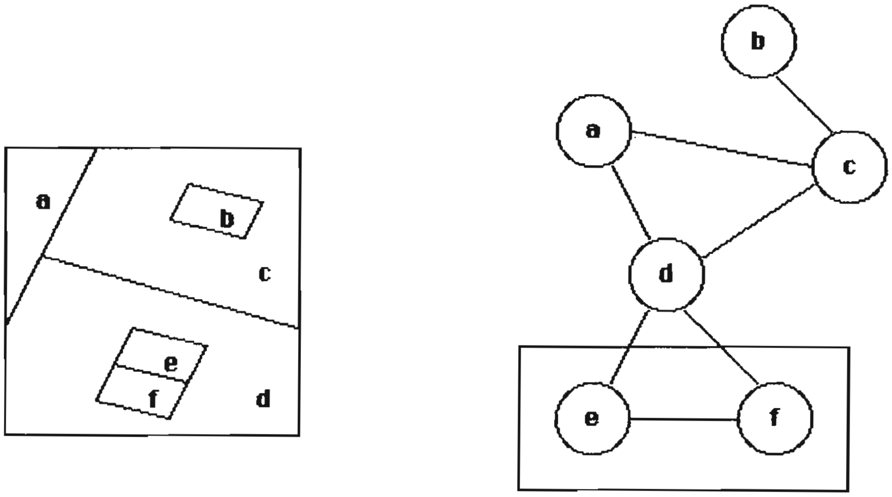

Labeled connected areas in image **Corresponding graph representation**

Fig.1 Symbolic representation of connected areas in surface material map

Fig. 2 Landsat TM image with bands 2, 3, and 5 in blue, green, and red.

Fig. 3 Initial surface material map after spectral shape classification.

Fig. 4 Refined surface material map after histogram analysis with turbid and still water (blue), trees (green), vegetation (dark green), crop fields (bright green), plowed fields (red), bare soil and sparse vegetation (brown), concrete (yellow), and asphalt (white).

Fig. 5 Results of surface feature classification with roads (white), built-up areas (red), agriculture (green), forests (dark green), and open water (blue).