Knowledge-Based Multi-Spectral Image Classification

Mark J. Carlotto
Victor T. Tom
Paul W. Baim
Richard A. Upton

The Analytic Sciences Corp.
1 Jacob Way
Reading, Mass 01867

Abstract

A new approach to the problem of classifying surface materials in satellite multi-spectral imagery is described and demonstrated in this paper. Surface material classes are defined heuristically using rules which describe the typical appearance of the material under specified conditions in terms of relative image measures. A knowledge-based approach allows expert knowledge of the domain to be used directly to develop classification rules. An expert system is currently being developed in the Zetalisp/Flavors programming environment on the Symbolics 3600 Lisp Machine. An example of its use in classifying Landsat Thematic Mapper imagery is presented.

1. Introduction

Traditionally, statistical pattern recognition techniques have been used in multi-spectral image classification applications in land remote sensing[1]. In using the signal statistics of spectral bands computed in a training region to represent land cover (surface material) classes, signal variability between data sets can become a problem in classification. The goal of the present effort is to develop an alternate approach to image classification. One in which expert knowledge of how materials typically appear in imagery obtained from a particular sensor under specified conditions can be used to develop rules for classifying imagery. The intent being to develop a representation that will allow material types to be recognized over a wider range of scene conditions than is possible with pure signal-based techniques. The approach is based on the use of rules to describe the typical appearance of surface materials in terms of relative image measures.

The organization of this paper is as follows: Section 2 reviews the Landsat Thematic Mapper sensor, with particular attention to its spectral bands. Section 3 points out some of problems encountered with the use of absolute signal representations for surface material classes, and motivates the development of relative image measures. Section 4 discusses how these relative image measures are computed and how they can be used to represent surface materials in the form of rules. The architecture of the expert system is discussed in Section 5, and a demonstration of its use in classifying Landsat Thematic Mapper imagery is contained in Section 6.

2. Landsat Thematic Mapper

The Thematic Mapper (TM) is one of several types of instruments carried aboard the Landsat satellite. The TM collects data from seven spectral bands in the visible and infrared (IR). The spectral bands are summarized in Table 1. The spatial resolution of the TM is 30 m, except 120 m for the thermal IR band. The locations and bandwidths of the TM bands were selected to allow the discrimination of vegetation, land use, and other resources. For example, band 1 is useful for hydrographic studies and the differentiation of coniferous and decidous forests. Band 3 is centered on a chlorophyll absorption band to aid plant species differentiation. Band 4 can detect ferric absorption (an important indicator for some types of mineralization), and is also important for biomass estimation and the delineation of water. Band 5 can be used to estimate the moisture content in vegetation, and also serves to differentiate between snow and clouds. Band 7 detects a hydroxyl absorption band which is an important indicator of cetain clay minerals.

Table 1 Landsat Thematic Mapper Spectral Bands

Band Number	Wavelength (μm)	Region
1	0.45 - 0.52	visible
2	0.52 - 0.60	visible
3	0.63 - 0.69	visible
4	0.76 - 0.90	near IR
5	1.55 - 1.75	middle IR
7	2.08 - 2.35	middle IR
6	10.4 - 12.5	thermal IR

3. Characterization of Surface Materials

For classification purposes, land cover or surface material classes are generally represented in terms of the means and covariances of spectral bands and transformations of spectral bands in prototypical regions in a training data set[2]. This statistical description is then used to define the decision boundaries in feature space which are subsequently used to classify test data set(s). In attempting to use the training statistics computed in one image to classify another, signal variability becomes a problem[3]. For example, one image may be hazier than the other; or images taken at different times may be different due to changes in the illumination, or in the biomass. So while there is generally sufficient information at the signal level for discrimination, the information is not sufficient for classification over a wide range of conditions. In short, a unique and invariant signature for each surface material or land cover class does not exist at the signal level[4].

The image analyst (IA) familiar with a particular type of sensor is generally able to recognize the typical appearance of surface materials in an image over a wide range of conditions. IAs are able to interpret imagery under variable conditions because they know the kind of scene they're looking at (hence the types of objects and materials to expect in the scene), and are able to reason about the appearance of various materials and structures in the image not only in the visible, but in the infrared and microwave regions as well. Since humans tend to describe things in relative terms (e.g., wet fields are darker than dry fields in the visible and infrared) rather than absolutely, it seems appropriate to develop a representation which is based on relative image measures.

4. Relative Image Measures

Currently, two kinds of relative image measures are being investigated for the purpose of characterizing surface materials. They involve an analysis of multi-spectral images pixel-by-pixel across wavelength (spectral signature analysis), and at a particular wavelength or spectral band across intensity (histogram analysis).

The spectral signature at (x,y) in the image consists of the set of individual sensor responses $\{ i_k(x,y) \}$ to the incident radiation $I(x,y,\lambda)$ where

$$i_k(x,y) = \int_0^{\infty} I(x,y,\lambda)\, r_k(\lambda)\, d\lambda$$

and $r_k(\lambda)$ is the detector response in the k-th spectral band. By examining the spectral signature, a rough indication of the appearance of the material as a function of wavelength can be obtained. Thus, one can speak of trends in the signature, e.g., the spectral response peaking at a particular wavelength, and use this kind of description to characterize certain materials. The cell structure of vegetation (crops, trees) causes most of the incident radiation in band 4 to be either reflected or transmitted, and much of the radiation in band 5 to be absorbed due to water in the cell structure (Fig 1a). Soil-like materials on the other hand tend to reflect increasing amounts of radiation as one progresses into the far infrared (Fig 1b). A comparision of Landsat TM bands 4,5, and 7 permits vegetation and soils to be easily discriminated (Fig 2).

The histogram of a spectral band

$$f_k(i) = \#\,(\, i_k(x,y) = i \,)$$

summarizes the relative frequency of intensities in that band over the entire image. If one knows something about the contents of the scene, it is possible to relate various modes in the histogram to instances of particular materials in the image. For example, since the reflectivity of water in the infrared is quite low, if the scene contains water, then the darkest regions in the infrared are likely to be water. Dense foliage (e.g., crops) tends to reflect most of the incident radiation in the near infrared, so if crops are present in the scene, then the brightest regions in the near infrared are likely to be crops.

Histogram analysis[5] involves: smoothing the histogram by convolving it with a Gaussian to remove spurious peaks, locating zero-crossings in the second derivative of the smoothed histogram, estimating the parameters of an underlying probability model using a maximum likelihood estimation technique, and assigning pixels into classes (e.g., relatively dark or bright) using a minimum probability of error criterion. The histogram is approximated by a Gaussian mixture, whose initial parameters (i.e., relative frequency $P(\omega_q)$, mean μ_q, and variance σ_q^2 for each distribution) are obtained from the zero-crossing analysis. These parameters are then refined using an iterative maximum likelihood estimation technique. The estimated mixture is

$$\bar{f}(i) = \sum_{q=0}^{Q-1} P(\omega_q)\, p(i \mid \omega_q)$$

where

$$p(i \mid \omega_q) = \frac{1}{\sqrt{2\pi}\,\sigma_q} e^{\frac{-(i-\mu_q)^2}{2\sigma_q^2}}$$

is the conditional probability density for the *q-th* mode.

Several modes may be required to describe the distribution of image intensities in a particular band for a given surface material. For example, six modes were found in the TM band 4 histogram shown in Fig 3a. The distributions due to water, crops, and other materials are identified in Fig 3b. The darkest two modes corresponded to deep water (reservoir), and shallow water (river and small ponds), respectively.

To determine if a pixel belongs to a particular mode, the *a posteriori* probabilities

$$P(\omega_q \mid i(x,y)=i) = \frac{P(\omega_q)\, p(i(x,y)=i \mid \omega_q)}{p(i(x,y)=i)}$$

for $q=0,1,\ldots,(Q-1)$ are computed. The pixel is assigned to the p-th mode if

$$P(\omega_p \mid i(x,y)) > P(\omega_q \mid i(x,y))$$

for $p \neq q$. In the present system, histogram analysis is performed prior to classification to determine if a pixel is *relatively-dark*, or *relatively-bright* in a given spectral band; i.e., if a pixel belongs to the darkest or bightest mode(s) in the intensity histogram for that band.

5. Expert System Architecture

A multi-spectral image analysis expert system is currently under development at TASC. The system is being implemented in Zetalisp/Flavors on the Symbolics 3600 Lisp Machine. The Symbolics 3600 is a single-user multi-tasking workstation built around a special purpose processor developed by Symbolics. Zetalisp is a rich implementation of Lisp built as a superset of both Maclisp and Common Lisp. The Flavors object-oriented programming language is a subset of Zetalisp.

The multi-spectral image classifier architecture is based on the use of object oriented programming to modularize the software into "expert system modules", and the organization of these modules into a cooperative hierarchy. Object-oriented programming allows the fusing of data structures and procedures for accessing and manipulating them into *objects*. Once an object has been defined, copies of that object may be *instantiated*. In addition to the modularity inherent in object-oriented programming, interactions between an object instance and the outside world takes the form of uniform *messages* which are handled by the object without requiring the sender to be aware of the processing involved. Thus, object-oriented programming offers a powerful tool for creating systems that are highly modular and quite flexible due to this "virtual" character of message passing[6].

Expert system modules are based on object definitions which function as inference engines. Each inference engine has associated with it an interpretable knowledge base containing rules, a factual knowledge base containing values of attributes, and a procedural knowledge base containing methods for computing attribute values. Each inference engine module incorporates a control scheme (e.g., backward chaining) as part of its definition.

Interpretive or expert knowledge for recognizing surface materials is expressed in the form of production rules. A rule consists of an antecedent (if-part) and a consequent (then-part). The antecedent contains one or more selectors[7], all of which must be true for the rule to fire (i.e., for the action specified in the consequent part to be performed). Each selector contains an attribute and one or more predicate-value pairs. For example, the selector

(band-4 > band-3)

defines the part of the image where the intensity in band 4 is greater than the intensity in band 3. The rule

If: (band-4 > band-3) and (band-4 > band-5)
Then: (assert vegetation)

states that wherever the intensity in band 4 is greater than that in bands 3 and 5, one can infer that vegetation is present.

Given these modules, the construction of an expert system is straightforward and linking systems together into a cooperative hierarchy (Fig 4) requires minimal additional programming. The classifier domain and the information provided by the multiple spectral bands partitions easily into a recursive structure in which the initial image is first examined to determine major classes (e.g., water, vegetation, and soils). The parts of the image identified as belonging to each major class are then passed to other expert systems which perform further differentiation within their sub-domains. This process continues recursively until the lowest level classification has been performed.

By using multiple expert systems in a recursive hierarchy, the rules required to reach low-level classifications are greatly simplified. In addition, this architecture allows the system to selectively stop the depth of classification as desired and to provide approximate (i.e. higher-level classification) when low-level classification cannot be performed due to limitations in the data and/or the knowledge base.

6. Landsat-TM Classification Example

To demonstrate the expert system, a Landsat TM data set over Lawrence, Kansas is classified into general surface material (land cover) categories. The thermal band (band 6) is not used in this experiment due to its lower spatial resolution. Sharpening techniques currently under development at TASC[8] will permit its use in subsequent versions of the system. USGS topographic maps and aircraft photos were used to infer ground truth for assessing the accuracy of the classification.

The scene in Fig 5 is first decomposed into major classes: water, vegetation, and soil-like materials (Fig 6). The rules used to extract these major classes are

If: (band-4-relative-intensity is dark)
Then: (assert water)

If: (band-4 > band-3) and (band-4 > band-5)
Then: (assert vegetation)

If: (band-4 < band-5)
Then: (assert soil-like)

Next, crops were separated from vegetation, and plowed fields and concrete-silt were separated from soil-like materials. Plowed fields have the general spectral characteristics of soils (i.e., increasing response in the IR), but tend to be darker than most soils due to their higher moisture and organic content. On the other hand, concrete and silt, which also have the characteristic signature of soils in the IR, are bright relative to other soil-like materials in the visible. This information was used to sub-classify soil-like materials (Fig 7). In this particular scene, crops gave rise to the highest response in band 4. The following rule extracts crops when applied to regions of vegetation (Fig 8):

If: (band-4-relative-intensity is bright)
Then: (assert crops)

The usefulness of this discriminant will depend on the period of the growth phase. This being one case in which collateral information (i.e., the time of year) would be required.

7. Summary

A knowledge-based approach for classifying surface materials in satellite multi-spectral imagery (Landsat TM) was described. Expert knowledge for identifying surface materials is represented in the form of production rules. The rules embody knowledge of how a material appears in imagery obtained from a particular sensor. The current implementation uses a cooperative hierarchy of expert system modules. Each contains rules for decomposing a class into sub-classes, and uses backward-chaining inference applied on a pixel by pixel basis to classify the sub-domain.

References

(1) D.A.Landgrebe, "Analysis Technology for Land Remote Sensing," *Proceedings of the IEEE*, Vol.69, No.5, May, 1981.

(2) P.H.Swain, "Fundamentals of Pattern Recognition in Remote Sensing," pp. 136-187, in *Remote Sensing: The Quantitative Approach*, P.H.Swain and S.M.Davis (Editors), McGraw-Hill, New York, 1978.

(3) R.M.Hoffer, "Biological and Physical Considerations in Applying Computer-Aided Analysis Techniques to Remote Sensor Data," pg. 270, in *Remote Sensing: The Quantitative Approach*, P.H.Swain and S.M.Davis (Editors), McGraw-Hill, New York, 1978.

(4) R.G.Henderson, "Signature Extension using the MASC Algorithm," *IEEE Transaction on Geoscience Electronics*, Vol.GE-14, No.1, January, 1976.

(5) M.J.Carlotto, "Image Segmentation Via Zero-Crossing Analysis," (in preparation)

(6) P.W.Baim, M.J.Carlotto, and R.A.Upton, "An Expert System for Target Identification in Radar Imagery Using Object-Oriented Techniques," submitted to *The First Conference on Artificial Intellegence Applications*, Denver, Colorado, December 5-7, 1984.

(7) R.S. Michalski, "Variable-Valued Logic: System VL1," *1974 International Symposium on Multiple-Valued Logic*, May, 1974, West Virginia University, Morgantown, West Virginia.

(8) V.T. Tom, M.J. Carlotto, D.K. Scholten, "Spatial Resolution Improvement of TM Thermal Band Data," *SPIE's 28-th Annual Internation Technical Symposium*, August, 1984, San Diego, Cal.

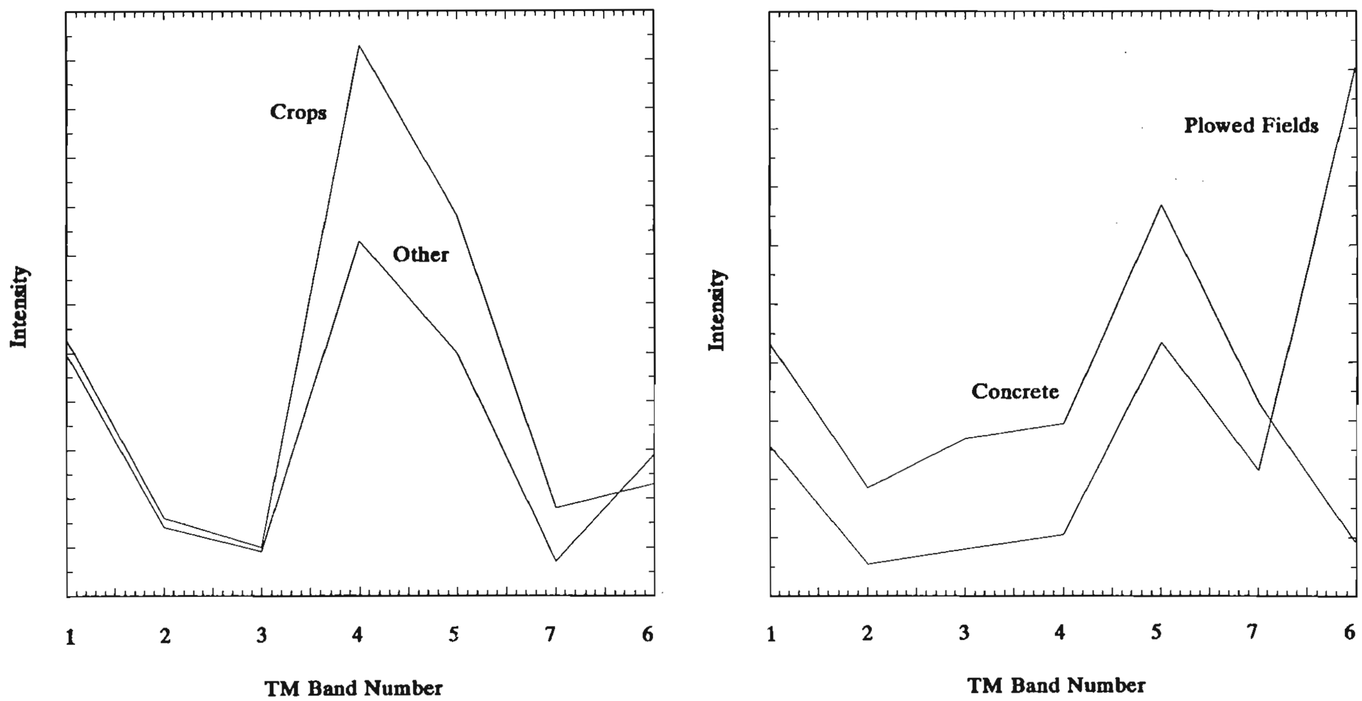

Fig 1. Spectral signatures of vegetation and soil-like materials

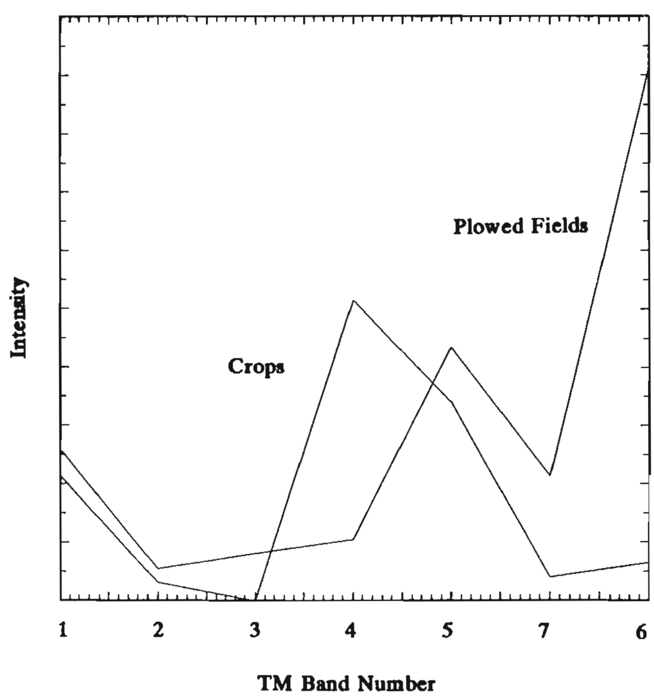

Fig 2. Comparison of spectral signatures of crops and plowed fields

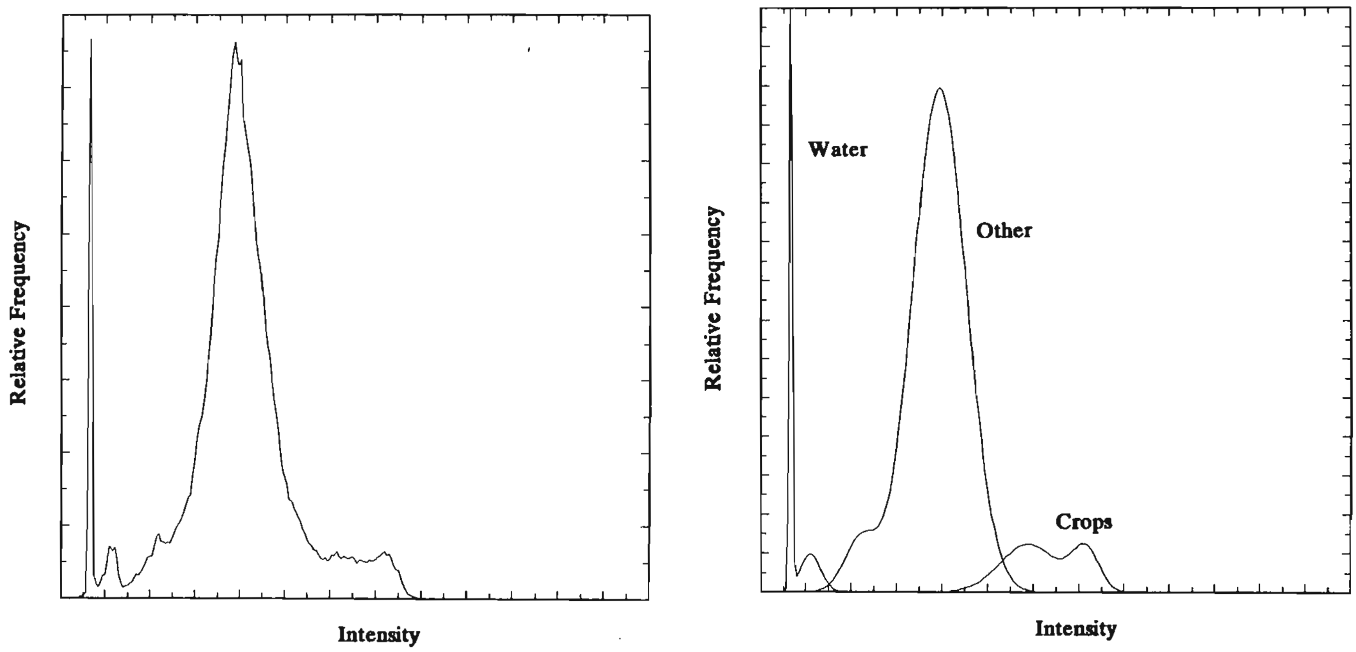

Fig 3. Histogram of TM band 4 with components due to water, crops, and other materials identified

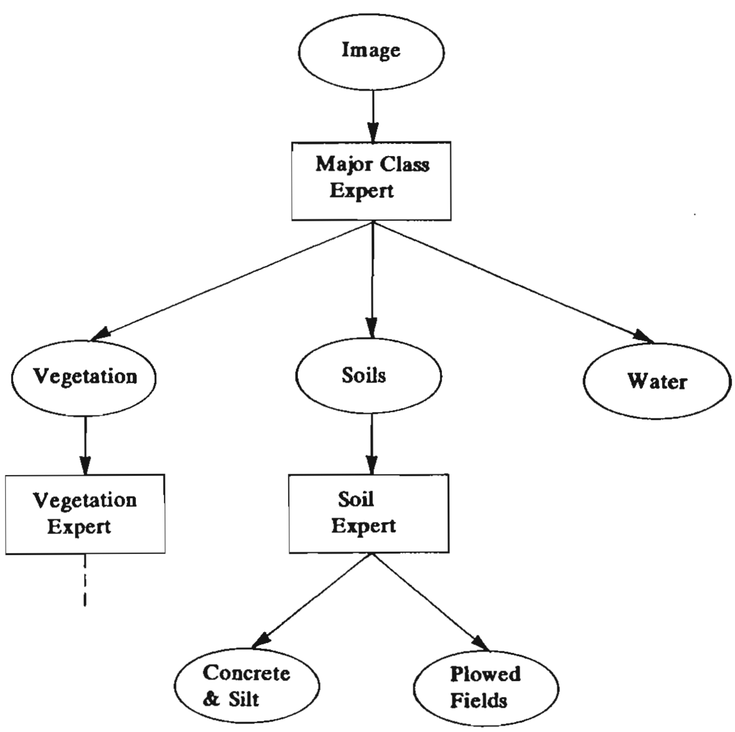

Fig 4. Expert system modules organized into a hierarchy for classification

Fig 5 Thematic mapper image over Lawrence, Kansas

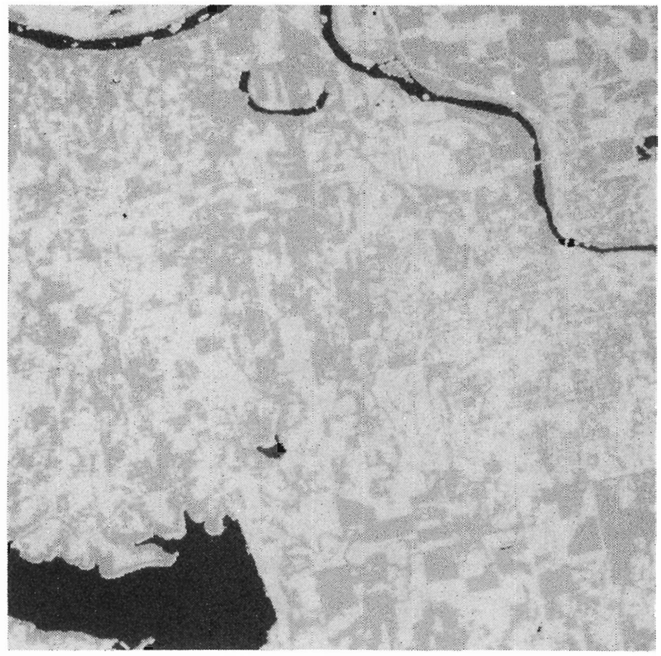

Fig 6 Major classes: water (black), vegetation (grey), soil-like materials (white)

Fig 7 Soil-like materials: concrete and silt (white), plowed fields (dark-grey), other kinds of soil-like materials (dark-grey)

Fig 8 Vegetation: crops (white), trees and other kinds of vegetation (grey)