

SPATIAL REASONING IN REMOTELY SENSED DATA*

J. Campbell
R. W. Ehrich
D. Elliott
R. M. Haralick
S. Wang

Virginia Polytechnic Institute and State University
Blacksburg, VA 24061

ABSTRACT

Photointerpreters employ a variety of implicit spatial models to provide interpretations from remotely sensed aerial or satellite imagery. The process of making the implicit models explicit and the subsequent use of the explicit models in compute processing is difficult.

In this paper we illustrate one application: how ridges and valleys can be automatically interpreted from LANDSAT imagery of a mountainous area and how a relative elevation terrain model can be constructed from this interpretation. We show how to examine valleys for the possible presence of streams or rivers and explore how a spatial relational model can be set up to make a final interpretation of the river drainage network.

1. INTRODUCTION

Remote sensing practitioners are highly skilled in deriving information from complicated geographical structures represented on aerial photography despite the fact that they may have great difficulty explaining their reasoning processes. Because humans have great innate facility for extracting information from visual shapes, forms, and textures, photointerpreters often do not devote much conscious thought to their analyses of the detailed relationships between light and dark that convey information about the content of a visual scene. The reasoning process may be implemented through a series of implicit steps that are not immediately obvious even to those who conduct the interpretation. As a result, to mechanize this process it is necessary to devote some effort to constructing explicit definitions of many of the subtle (possibly vague) processes applied by a skilled photointerpreter as he analyses a scene.

This paper describes our initial work to formulate a relational model for reasoning about the contents of a visual scene. This work is being done in cooperation with a skilled photointerpreter, and we are attempting to determine the primitives and the relationships for one task domain -- deducing the

* The work described in this report was done at the Spatial Data Analysis Laboratory at VPI & SU with the support of the National Aeronautics and Space Administration through contract No. NAS5-26276

network of streams and rivers represented by Landsat data. There are two reasons for our choice of this application area. First, data are abundant, and they are available for geographic regions familiar to us. Second, despite the apparent triviality of the problem, the reasoning required to reconstruct the drainage network is in fact very complicated, in part because it requires simultaneous analysis both at the detailed local level and at very global levels. Furthermore, reasoning about rivers requires a great deal of knowledge of geography and physics that is not directly part of the relational model. This knowledge, for example, includes such facts as "water runs downhill" and "in the morning, the sun is in the east". Thus far we have examined spatial relationships within a single spectral channel, and have not yet attempted to apply the additional information that may be conveyed in other spectral records.

We began our work on the spatial reasoning problem by assuming that the radiometric information we would need could be derived from edge information obtained by applying gradient-type operators, provided we used the proper edge model. As work progressed, it was determined that edges derived from region boundaries were more reliable, and the work has progressed to the point where we are extracting mountain ridges and valleys with a reliability that is sufficient for us to concentrate our efforts on the formulation of a relational model.

2. STUDY AREA AND DATA

This research examines an area in southeastern West Virginia, shown in Figure 1.




Figure 1 - Original Landsat Scene

This region is a portion of the Appalachian Plateaus physiographic province, within the "unglaciated Allegheny plateau" defined by Thornbury (1965). In general, this region is a thoroughly dissected plateau-like surface. It receives 40 inches of precipitation each year and, as depicted on topographic maps, has a moderate drainage density. Drainage is through tributaries of the New (Kanawha) River, which flows west into the Ohio drainage system.

The overall drainage pattern within our region is that of a relatively large sinuous channel (the Gauley River) superimposed over the finer texture of a dendritic pattern formed by first, second, and third order streams. A result of this entrenched channel is that a number of the small first or second order streams flow directly into the large channel. Thus the overall pattern is composed of a mixture of many very small stream segments, many with very steep gradients, a prominent major channel with a relatively low gradient, and relatively few stream segments of intermediate length and gradient.

Throughout the area, flood plains (when present) are narrow and tend to closely follow the course of the stream channel. Valleys are narrow, with steep sides; the Gauley River, for example, follows a valley that is typically 150 meters deep but only 100 meters wide. Uplands often consist only of ridge crests; although plateau-like upland regions are present, they are not continuous or extensive. The area is forested with a dense cover of deciduous trees (Kuchler's "mixed mesophytic forest"). Cleared areas for agriculture (chiefly pasture) tend to follow the valleys of intermediate-sized streams. Settlements are small and dispersed, usually positioned in valleys.

This region appears on the Charleston, West Virginia/Ohio USGS 1:250,000 quadrangle (NJ 17-5). Our investigations include areas in Nicholas County, W. VA and neighboring counties; many of our examples are within the Gauley Bridge 7.5 minute USGS topographic quadrangle. This area was imaged by the LANDSAT-I MSS on April 13 1975 (scene id: 5360-14502; path 18, row 34). This date reflects important qualities of the scene. First, at this date the atmosphere was unusually clear--there is no evidence of scattering or atmospheric degradation of the data. Also, at this date in the spring most of the forested areas are without leaves, especially at higher elevations. Lower elevations have a cover of newly-emerged leaves and grasses. Within a few weeks leaves will have emerged in vegetation throughout the entire regions, but at this time in April, there is a sharp spectral contrast between the vegetation cover of the higher elevations and that of some of the valleys.

For the most part, it is not possible to directly observe the drainage network of the LANDSAT data--instead the image must be examined to interpret the indirect manifestation of the drainage. Interpretation of this indirect evidence permits us to approximate the major features of the drainage pattern.

The data in band 7 provides a singular good example. Energy in this spectral region (.8-1.1 μ m) is largely free of atmospheric scattering, so shadows tend to be dark and clearly defined. Water bodies absorb infrared radiation, so they appear as clearly delineated dark bodies on the band 7 image. Living vegetation reflects strongly in this portion of the infrared, so areas of living, green, vegetation will appear as bright regions on band 7.

In this scene the solar beam illuminates the landscape from the southeast (azimuth = 119 degrees) at an elevation of 45 degrees above the horizon. Five separate cases can be defined:

1. Slopes that face the illumination are highlighted, and are recorded as bright areas on the image. Those that face the northwest are in shadow and appear as dark or black regions. Thus the overall pattern, greatly simplified, is one of alternating bright and dark strips in regions where the topography is oriented parallel to the solar beam. Contacts ("edges") between these light and dark strips on the image correspond to ridges and valleys on the ground.
2. In this region the topography is very complex, so there are numerous valleys and ridges that are oriented parallel to the solar illumination. In these instances shadows are minimized so the drainage course is visible on

the image only if the channel is wide enough to fill one or more MSS pixels. Thus, the course of a sinuous stream is visible on the image as the edge between a light and dark area that is broken into segments at irregular intervals. The breaks correspond to changes in direction that bring the stream course parallel to the solar beam.

3. Only the widest streams are wide enough to be directly imaged. The Gauley River channel--some 100 meters-wide is the only instance in our area. On band 7 this river is represented as a continuous chain of dark pixels of varying width. Where the river flows at right angles to the solar beam, the river (and its valley) is clearly depicted as a wide dark strip, often several pixels wide, caused by the combined effects of shadow and open channel. The brightly illuminated slope facing the sun is also visible. However, where the channel is parallel to the illumination the shadows are absent and the river is visible only because the channel is wide enough to appear in its own right. If the channel is not at least one pixel wide, there is a gap at this point in the image representation of the channel.
4. A few streams on the image are portrayed as bright areas on band 7. In these areas there is no shadowing due to the width or orientation of the valley. The newly emerged leaves of shrubs and grasses in pastures and meadows cause a bright response in the infrared radiation of band 7. The strong reflection contrasts with the darker response of the upland forests where leaves have not yet emerged.

This examination of the landscape and its image representation allows us to define some of the key characteristics of the patterns we intended to interpret:

1. Some streams can be defined by the edge between bright and dark patches on the image. However some edges are also ridges, so not all such edges correspond to streams. It will be necessary therefore, to be able to decide which edges correspond to ridges and those that represent valleys.
2. Some of the larger rivers appear as continuous chains of dark pixels.
3. Some of the smaller streams appear simply as short dark strips of pixels.
4. Some of the valleys will be defined not by shadow, or by the open water of a river channel, but by the bright infrared response of leaves on shrubs and grasses in non-forested valleys.
5. For 1, 2, and 3, we must expect to encounter gaps in the linear features. We must be prepared to extrapolate between the ends of segments to fill in these gaps.

3. A RELATIONAL MODEL

Before giving the algorithms by means of which we have been extracting information from the Landsat test scene, it is appropriate to present the details of the relational model and to show how it is applied to this particular problem.

There are numerous ways in which a relational model might be formulated. One that we have considered involves partitioning each connected valley and ridge system into a set of segments that contain no junctions or termination points except at their ends. Each such segment is called a unit, and associated with each one is a list of measured properties. The goal of the reasoning process is to assign to each unit one of the following labels in such a way

that no logical conflicts occur with the spatial knowledge database. The labels are:

- 1) Valley containing a stream
- 2) Valley not containing a stream
- 3) Ridge

The spatial knowledge database would consist of a set of relations, R, of the form

$$((s_1, p_1, l_1), \dots, (s_n, p_n, l_n), r, N).$$

Each relation is a generic configuration, r , among n units with particular properties and labels which are given the name, N . More specifically, s_1, \dots, s_n are segments with property lists p_1, \dots, p_n and labels l_1, \dots, l_n . To see how these relations are used, consider a small example of three units u_1, u_2 , and u_3 on the image with respective property lists q_1, q_2 , and q_3 that actually correspond to three segments that form a stream junction. The reasoning system scans the relations R until it finds one that satisfies the information that has been observed. A relation is satisfied if there exists a correspondence of units in the scene with segments in the relation such that the corresponding units have the properties specified in the relation and such that the units are in the spatial relationship, r .

In the case of the example, suppose that

$$((s_1, p_1, l_1), (s_2, p_2, l_2), (s_3, p_3, l_3), r, N) \in R$$

is found such that u_1 corresponds with s_2 , u_2 corresponds with s_1 , and u_3 corresponds with s_3 , properties p_2 are among the properties q_1 of s_1 , p_1 is contained in q_2 , and p_3 is contained in q_3 . Then if u_1, u_2 , and u_3 have spatial relationship r , label l_2 is a possible label of u_1 , l_1 is a possible label of u_2 , and l_3 is a possible label of u_3 . Also, the name N junction is given to that particular configuration of u_1, u_2, u_3 . The variable r may simply be considered the name of the procedure that determines whether or not u_1, u_2 , and u_3 are indeed the N junction relationship.

The goal of the reasoning process is to assign labels to units in such a way that there is global consistency. Suppose, for example, that a unit u enters into one relationship that supports the assignment of label l to u . If there are other valid relationships involving u , all of them must support l as a label. Otherwise l is removed as a candidate label for u . One of the features of this process is the way which names are assigned to relationships. Units in such relationships may themselves be regarded as units that participate in higher level relationships (Haralick, 1978).

There are numerous properties that may be measured and associated with each segment. These include:

- 1) Local texture
- 2) Global texture
- 3) Global curvature
- 4) Length
- 5) Number of segments incident at each end
- 6) Endpoint separation
- 7) Global orientation
- 8) Local orientation at ends
- 9) Contrast (or confidence)
- 10) Ridge or valley
- 11) Width
- 12) Location of endpoints

We are experimenting with the above properties but do not know yet whether we have determined a suitable menu.

4. SEGMENTATION

In this section the segmentation algorithm is described by means of which valleys and ridges are extracted from the original scene for inclusion in a database. This database stores line segments, and for each line segment a list of properties is stored by means of which the spatial reasoning algorithm marks those valleys that are candidates for "streamhood" according to the stored relations.

The picture elements of the test image in Figure 1 have been classified into the four classes shown in Figure 2 by grey level thresholding.




Figure 2 - Thresholded Image

The resulting regions are termed very dark, dark, bright, or very bright according to the intensities of their picture elements. The regions with very dark pixels may belong to a large body of water or to the dark side of the mountain. The regions with dark pixels usually belong to the dark side of the mountain. Regions with very bright pixels belong to vegetation areas. Since in the work reported here we do not wish to differentiate open water (very dark pixels) from other regions, the dark and very dark pixels are placed together into the same category.

Thresholding to separate bright from very bright pixels is relatively successful because of the distribution of gray levels in the test image. However, the distinction between bright and dark picture elements is not clear and the regions that result after thresholding require further processing. First the dark regions are obtained by setting a relatively low threshold. Then, each 8-connected component (i.e. pixels that are transitively connected through 8-neighbor adjacencies) is assigned a unique region number. Next, each connected component is enlarged slightly by analyzing each pixel directly adjacent to the region boundary. Each such pixel is added to the region if its intensity is closer to the region's threshold value than to the next higher

(brighter) threshold value. This processing step tends to increase the contrast of the region boundaries. Even though regions may touch one another after the growing process, they still have different region numbers and are, therefore, distinct regions. However, since they do touch each other, the distinct regions are hard to display in a black and white photo. Figure 3 is an attempt to show the different regions by encoding them with different intensities.



Figure 3 - Intensity Coded Regions

The boundaries of each of the closed regions just determined are easy to extract by using a boundary tracker that follows the perimeter of a labeled region. Since most of the trees in this image are unfoliated, the strongest region boundaries are shadow boundaries rather than tonal, and the strongest boundaries are those at the extremes of steep slopes oriented normal to the sun direction. Since the sun illumination is predominantly east-west, a boundary that is dark on the left and bright on the right will correspond to a ridge, and the reverse will correspond to a valley. For east-west region boundaries, the illumination model fails. Where east-west boundaries exist, some are ridges and some are valleys. We have labeled all such segments as valleys unless they are adjacent to ridges on both ends. The valleys, of course, are of interest because the streams and rivers will flow through some of them. Figure 4 shows the ridge and valley maps for the test image, and Figure 5 shows the stream map obtained by analysis of corresponding topographic maps.



Figure 4 - Ridge and Valley Maps



Figure 5 - Stream Map

5. FROM RIDGES AND VALLEYS TO A TERRAIN MODEL

Once the ridges and valleys have been identified, it is possible to construct a digital terrain model by solving a two-dimensional boundary value problem. To do this we associate boundary values with each valley and ridge pixel. The valley pixels get a relative elevation of 0 and the ridge pixels

get a relative elevation of 100. We then seek a surface having the given boundary values and whose Laplacian is everywhere 0.

A digital Laplacian mask is shown in Figure 6.

	-1		
-1	4	-1	
	-1		

Figure 6 - A Digital Laplacian Mask

The terrain model we wish to determine has elevation values which satisfy that the application of the digital Laplacian mask to each 3 X 3 neighborhood of the image yields 0.

An iterative procedure accomplishes the required solution to the boundary value problem. For the initial step, those pixels which have not been marked ridge or valley are given the intermediate value 50. In each iteration, all the pixels of the current terrain model are scanned in a left to right, top to bottom manner. If the pixel is a ridge or valley pixel nothing is done. If the pixel is not a ridge and not a valley, then its value is replaced with minus the sum of its north, south, east, and west neighbors. The replacement is done recursively using the most recent neighbor values.

Figure 7 illustrates the relative terrain model constructed from the ridge and valley images of Figure 4.



Figure 7 - Relative Terrain Model

The realism is suprisingly good. The fact that the terrain model is only relative because true elevations of the ridges and valleys were not used limits the application of the terrain model. However, the fact that the terrain model is precisely registered to the Landsat scene is an important advantage.

6. SEMANTICS OF STREAM NETWORKS

In this section we attempt to enumerate the properties of well-formed stream networks by means of which we can deduce whether a valley segment in the valley database is a valid stream segment and whether a gap between several valley segments ought to be filled on the basis of contextual information and labeled stream segments. The following list enumerates what we regard as the most important criteria for forming valid stream networks.

- (a) completed pattern is dendritic in form.
- (b) completed network has no enclosed areas (i.e., stream segments are connected only at the downstream ends).
- (c) completed network is without gaps (i.e., all segments are connected; no isolated segments).
- (d) unconnected ends of the smallest tributaries point towards ridgelines.
- (e) tributaries form acute angles with main streams, right angles possible; obtuse angles unlikely.
- (f) streams and valleys can't cross ridgelines.
- (g) ridgelines can't cross streams.
- (h) junctions of 2,3, or 4 segments possible.

- (i) small tributaries to large streams typically occur parallel or nearly parallel to each other, and often join the larger stream at a right angle.
- (j) "larger" streams may be recognized by a larger radius of curvature.
- (k) "smaller" streams may be recognizable by a smaller radius of curvature.
- (l) each stream segment can be labeled with a flow direction such that each stream junction has exactly one outward flowing stream.

Gaps in the stream data may not be filled if

- (a) the new link will cross a ridge
- (b) the new link will create a network with a closed loop
- (c) the new link will create a sharp angle in a large stream
- (d) the new link connects the unconnected ends of first order streams (see above)
- (e) the new link forms a tributary entering a larger stream at an obtuse angle

It is not hard to see that the translation of these stream semantics into relational form may require a substantial amount of work. It is also possible that relational form may turn out to be unnatural for people to use. One reason for this may be that each relation in R admits possibilities other than excludes them. The specification of exclusion relations may considerably simplify the task of specifying problem knowledge.

7. PERFORMANCE

Errors can be categorized as errors of omission (streams or ridges present on the ground but not identified on the interpretation) or errors of commission (identification of a stream or a ridge where in fact none are present on the ground). For this interpretation we detect no instances of identification of streams as ridges or vice versa, though the possibility of that happening must always be considered.

The most noticeable errors of commission are small segments, often spatially isolated, that apparently do not correspond to real streams or ridges. Some of these segments may in fact be "real", but they cannot be recognized as such from the topographic maps. These segments tend to occur as short segments of stream or ridge not connected to the larger more continuous networks. Those segments identified as streams often are located in the far upstream region of the drainage basin, and may possibly be the result of the fine-textured, complex topography and shadowing present in this part of the area. To the human interpreter the lengths, positions, and orientations of these segments permit a fairly clear separation of these segments from the major drainage features, although some of these errors could easily be interpreted as real features.

Another error that occurs less frequently but is more serious are the errors of commission that bridge the gaps between the unconnected tips of the smallest tributaries in separate drainage basins. This kind of error is small in extent but is very important as it can completely confuse the interpretation of the drainage network. Errors 1-3 in Figure 8 demonstrate these kinds of errors.




Figure 8 - Valley Map with Overlaid Stream Map Sketch

If this kind of error is present, streams join at both upstream and downstream ends. If only a small local area is considered, it may be very difficult to detect such errors using only the information present in the interpretation. If a much larger area is examined, it may be possible to detect and resolve these errors if they are few in number and the remainder of the network has been interpreted correctly.

The most important errors of omission are missing segments in the stream, often caused by changes in the orientation of a stream or valley that brings its course parallel to the solar beam, (5-7 in Figure 8) thereby eliminating or diminishing the shadows used for the interpretations. In other instances these gaps have no immediately obvious cause (4 in Figure 8). Usually these errors are readily recognizable (upon manual inspection of the product of the machine interpretation) because the segments on both sides of the gap are often fairly long and continuous, and it seems clear from examination of the global pattern that these gaps should be filled in to form a continuous network.

When the ground truth stream map is superimposed on the valley map as shown in Figure 8, the results look very encouraging. Indeed, it would be considerable effort for a skilled photointerpreter to do this well. However, there are a number of problems. One of most serious is that there are a large number of dense clusters of marked picture elements that make the valley connections ambiguous. Since there are also numerous short gaps, the number of hypothesized valley segments will tend to be large.

The lack of reliability of the data is another serious problem. Not only may the valley-ridge labels be wrong, but any of the properties may be incorrect. For example, it is very difficult to measure the incidence angle of a

tributary with a main stream, and the reasoning process will have to be modified to be tolerant of measurements that are totally wrong. There is also a possibility that it may not be possible to determine exactly where an error has been made. Error 1 in Figure 8 results because an east-west ridge that is completely invisible in the imagery blocks an apparently continuous valley. While this error is easy to detect, its cause may be elusive because of the missing data.

Because of the subtlety of some of the deductions that are being made, image imperfections may also cause problems. In this particular image, 6th band striping is easily visible, and in several instances the ridge-valley maps have been influenced by false intensity shifts between adjacent scan lines. Nevertheless, we feel that it will be possible to achieve a significant degree of automation of the spatial reasoning task in this application area, and we are continuing our work to refine the relational reasoning model.

8. REFERENCES

- [1] Kuchler, A. W., 1964, Potential Natural Vegetation of the Conterminous United States. American Geographical Society Special Publication No. 36.
- [2] Thornbury, William D., 1967, Regional Geomorphology of the United States, Wiley, New York, pp. 609
- [3] Haralick, R. M., "Scene Analysis, Arrangements, and Homomorphisms," Computer Vision, (Hanson and Riseman, Eds.) Academic Press, 1978, pp. 199-212.