the loss from the reduced scale.

A fair comparison is afforded between the results from the A-9 Autograph and those normally expected from an A-7. Both are built by the same company, both operate on the same mechanical principle, and the grid tests indicate that the precision of the A-9 tested is about equivalent to that of the A-7 Autograph. Judging from results over several years from the A-7 at GIMRADA and others owned by private companies under contract, and from the results with the A-9 in these tests, the vertical accuracy of ultra-wide angle models is about $\frac{5}{6}$ that of standard mapping photography.

Since the accuracy of a given system is considered to be essentially a straight-line function of altitude, it follows that accuracy equal to that of standard photography could be achieved by flying ultra-wide-angle photography at $\frac{5}{6}$ the altitude. This, then, would reduce the coverage advantage so that the area per photo would be only a little more than double that of regular photography. Still, cutting in half the number of pictures required for a given area seems to be a worthwhile achievement.

Considerable discussion has centered around the fact that in rough terrain there may be hidden areas caused by the low slope of rays from those areas. This is true, of course, but the problem is not limited to ultra-wide-angle photography. There will be hidden areas wherever the slope of the ground is away from the camera station and is steeper than the rays to the camera. This condition is currently plaguing those trying to map rough areas such as the Rocky Mountains, and would be worse if ultra-wide-angle photography were being used. The conclusions are that there are vast areas in the world that are not too rough to be mapped by ultra-wide-angle photography, that it is not to be expected that it will completely supplant regular mapping, but that there will be sufficient advantage in its coverage aspects to insure its widespread acceptance as an important new tool of photogrammetry.

*An Approach to Automatic Photographic Interpretation**

DR. AZRIEL ROSENFELD,
Budd Electronics,
Long Island City 1, N. Y.

ABSTRACT: *This paper describes an approach to the automatic identification of the images of targets on aerial photographs. The approach described involves the extraction from the photographic image of two basic types of information, one of them relating to the presence in the image of figures having given shapes and sizes, and the other to the "textural" nature of the image.*

TARGET RECOGNITION CRITERIA

A MAJOR aim of photographic interpretation is the recognition of the presence or absence of given types of features ("targets," in the military case) on given aerial photographs. To perform photographic interpretation automatically, the photographs must be accepted as inputs by the data handling system which will make the recognition decisions. A major obstacle to effective automatic photographic interpretation is the wealth of information, measurable in the billions of bits, which a high quality photograph contains. The human interpreter certainly does not consciously process all of this information in making flash recognition decisions; for recognition purposes, most of the information content of the photograph is redundant. It seems reasonable to conclude that a practical method of automatic target recognition will have to markedly reduce the amount of redundant information contained in the photograph. The result of such a reduction in information may be thought of as a simplified description

* This paper was presented at the 1962 National Convention on Military Electronics.

of the photograph which tells just enough about it to make target recognition possible.

Of the many types of extractable descriptive information, the choice of those which are actually extracted must be based on apparent or proved usefulness as aids to target recognition. An analysis of methods of strategic photographic interpretation suggests that there are at least two important kinds of descriptive information which can contribute to target recognition, namely,

a) *Shape and size information*—that is, the presence of images having characteristic ranges of shapes and sizes. Examples of targets for which shape/size information contributes to recognition are airfields, storage tanks, ships and planes, as well as regions of terrain which contain cultural features (these last almost invariably contain straight lines and right angles). Targets of these types are illustrated in Figures 1–2.
b) *"Textural" information*—that is, parameters which describe local distributions of densities and density contrasts in the photographic image. The simplest examples of such parameters are the mean density and the mean "level of detail" (=number of contrasts per unit area). Targets for which textural information contributes to recognition include urban and industrial areas (high level of detail) and targets such as harbors, ships and bridges which are associated with hydrographic features (negligible level of detail in the absence of sun glitter). Such targets are illustrated in Figures 3–4.

Most target types require for their identification the presence of particular *combinations* of shape/size and textural information in *particular relative locations* on the photograph. A useful automatic photographic interpretation scheme should thus extract from a given photograph at least the following types of information:

(1) *The locations of images having various shapes* (in particular: straight lines, circles, right angles, etc.) *and sizes* appropriate to the scale of the photograph.
(2) *The locations of regions having various textural natures* corresponding to terrain and cultural feature types which are relevant to target identification.

In the following sections of this paper, automatic shape recognition techniques as they apply to photographic interpretation are briefly reviewed; approaches to statistical image analysis are then discussed in some detail.

Shape Recognition

Two basic methods of detecting the presence of an image of a given shape on a photograph may be referred to as *direct matching* and *property matching*. In direct matching, the photographic image is correlated with a suitable physical or mathematical template in such a way as to produce signal maxima at points corresponding to the locations of images of the given shape. An important advantage of the direct matching approach is that it selects and recognizes the given shaped images in a single operation; it does not require that portions of the photograph first be selected ("figure extraction") and then tested to determine whether or not they have the given shape ("figure recognition"). This approach has the disadvantage of being sensitive to the relative orientation of the photograph and template; however, it can be performed for all possible orientations at fairly high speeds.

There exists a wide variety of property matching techniques for shape recognition. These techniques generally presuppose that figure extraction has been performed. They then proceed to measure certain descriptive properties of the extracted figures and to compare these measurements with those of the shape which is sought. Many of the most useful of such measurements relate to the boundary of the figure in question. For example, an arbitrary shape is completely determined by specifying the curvature of its boundary (as a function of arc length measured along the boundary). If the shape which is to be recognized is of a relatively simple type, special-purpose recognition techniques can be devised for it; for example, straight lines can be recognized by applying "tracking" techniques. The examples just given indicate that the property matching approach to shape recognition can be made independent of the position and orientation of the shape being sought. However, this approach still requires that the shapes in question be somehow picked out from the photographic image as a whole; this will in general be a highly nontrivial task.

Analysis of Image Texture

A variety of parameters can be used to measure the "texture" of a given portion of a photograph. Among the operations which can be performed on the image density function to yield useful texture information are harmonic

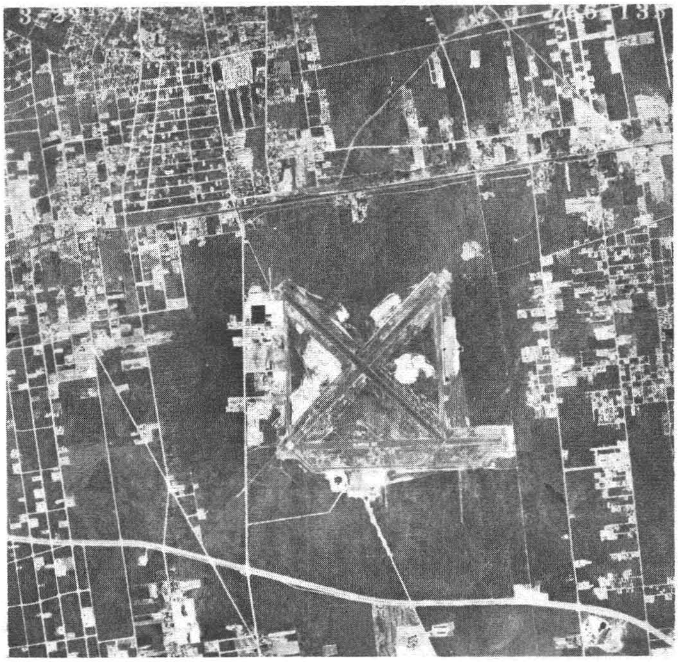

FIG. 1

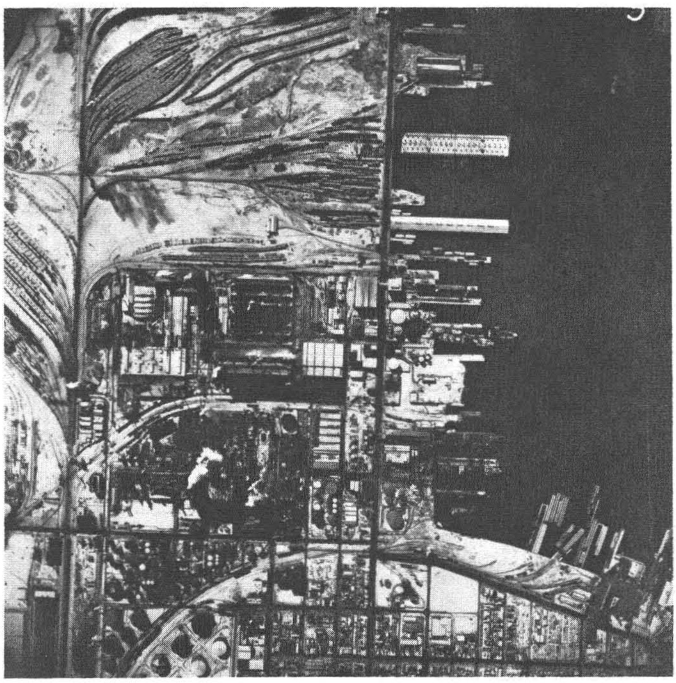

FIG. 2

FIG. 3

FIG. 4

analysis; autocorrelation; differentiation and level slicing; and measurement of statistical moments (mean, variance, etc.). The relative merit of these parameters as texture descriptors has yet to be completely determined. However, the limited studies which have been performed* indicate that such parameters can provide useful texture type classifications which are suitable for automatic photographic interpretation purposes.

A shortcoming of most of the past work on image texture analysis is that it has related to the measurement of texture parameters for a preselected portion of the image. To perform testure analysis for the image as a whole, two approaches are possible:

a) The photograph can first be divided into fixed portions (small squares, for example) and the texture parameters in question measured for each of these portions. As a means for extracting textural information for purposes of photographic interpretation, this approach has the disadvantage that the portions are not preselected with any regard to the content of the photograph. If one of the portions overlaps two differently textured types of regions on the photograph, its parameters will correspond to neither of the types, and may misleadingly correspond to a third type which is not actually present. As a result, important information relating to the boundaries between different types of regions is lost.

b) Alternatively, the texture parameters can be measured for an "aperture" which is systematically scanned through every possible position over the photograph. This approach gives complete textural information about the photographic image. However, this information is actually greater in quantity than the information contained in the image density function, since it includes several texture parameter measurements made over an aperture centered at every point of the photograph. This approach thus increases, rather than decreases the information content of the image, and as such is hardly practical for use in an automatic photographic interpretation system.

The disadvantages of the two approaches just described can be avoided by taking a *global*, rather than a *local*, approach to textural image analysis. In this approach, the photograph is automatically divided into regions of appreciable size, in each of which the texture parameters of interest are substantially constant, but which are such that these parameters differ significantly from region to region. If this can be done, the resulting regions will be essentially the parts into which a human observer might divide the photograph if asked to report on its appearance after very briefly viewing it; they should in general correspond to different types of terrain. The texture parameters of interest can then be measured for these regions. This will result in a radical quantitative reduction in the information content of the photograph; at the same time, very little *useful* information will be lost, since the parameter measurements taken over such uniform regions will not differ widely from the local measured values of the parameters within these regions. The important information relating to the boundaries between different types of regions will also be preserved.

Methods of automatically dividing a photograph into texture (or terrain) types will now be discussed.

Texture Type Separation

The subdivision of a photograph into terrain types is not always easy even for a human interpreter. In some cases the boundaries between terrain types will be sharply defined and the terrain on opposite sides of these boundaries will be significantly different. In other cases the transitions between adjoining terrain types may be quite gradual. For illustrations of both situations, see Figures 3–4.

In automating the process of terrain type subdivision, it is reasonable to begin by treating situations in which the desired subdivision is relatively obvious. Further, it is convenient to consider a one- rather than a two-dimensional case. Suppose, specifically, that an "image density" function is given which varies along a line rather than over a plane. It is desired to subdivide the line into segments within each of which certain "textural" parameters of the given density function remain approximately constant while differing significantly from segment to segment. This corresponds to a subdivision into clearly defined, significantly differing terrain types. Such a subdivision can be achieved if it is possible to determine those points along the line (if any) on opposite sides of which one or more of the parameters in question are sig-

* See, for example, the author's *Automatic Recognition of Basic Terrain Types from Aerial Photographs*, Photogrammetric Engineering, Vol. XXVIII, No. 1, March, 1962.

nificantly different.

Mathematical models for the detection of subdivision points can be formulated using any of the textural parameters defined above. A simple model which used mean density and mean contrast frequency as parameters has been constructed and successfully tested as described elsewhere.* This model may be regarded as a first approximation in the sense that it employed only first statistical moments (means).

In devising more sophisticated subdivision models, it is convenient to make the further simplifying assumption that the input photograph can be adequately approximated by a quantization into elements which are either black or white. This assumption, though dubious at first glance, is actually not too unreasonable; very good black-or-white approximations to aerial photographs can be made using high-resolution "screening" techniques which are standard practice in the graphic arts. A one-dimensional black-or-white density function may be thought of as taking on the values zero and one only. The graph of such a function consists of a succession of plateaus (ones) and valleys (zeros).

Higher order approximations to the definition of texture can be very conveniently formulated in terms of black-or-white density functions. Specifically, one can use as texture parameters the mean, variance and higher moments of the plateau widths and of the valley widths in the function. It appears quite plausible that these parameters do indeed relate closely to the textural appearance of the region which has the given (cross-sectional) density function.

It may be noted that certain other useful texture parameters are closely related to the "width moment" parameters just defined. Thus the mean density and mean contrast frequency, are simple functions of the mean plateau and valley widths. The autocorrelation spectrum is not completely derivable from the first few width moments; however, periodicities in the density function, which are perhaps the most useful properties which can be detected by autocorrelation, can also be detected by measuring the variances of the plateau and valley widths. The relationship of harmonic components to statistical moments is even less immediate. It seems likely, however, that statistical moments are more closely related to human perception of and differentiation among textures than are harmonic coefficients.

The one-dimensional case discussed above can be generalized to two dimensions by measuring one-dimensional texture parameters in every direction through each point of the photograph. This straightforward generalization is probably not completely adequate in the sense that it does not give special weight to inherently two-dimensional aspects of the subdivision problem, such as the shapes of elements and of regions, which are "disproportionately" important to human observers. A more complete generalization requires an analysis of informational properties of perceived shapes and is beyond the scope of the present paper.

Texture-type separation schemes such as described above can be simulated quite inexpensively. The given photograph is first converted to the black-or-white format by high-resolution screening. Plateau and valley widths, and densities if desired, can be measured on this quantized photograph using a medium-power microscope. The resulting numerical data can be processed either manually or on a small general-purpose digital computer such as the IBM 1620. Real-time processing is also possible using an ultra high-resolution flying spot scanner or a motor-driven mechanical scan of the microscope stage; the resulting density signals can be quantized and processed using relatively simple special purpose circuitry.

The techniques discussed in this paper constitute an approach to automatic photographic interpretation along several parallel fronts. While much work has yet to be done before fully automatic photographic interpretation even approaches realization, it is felt that present lines of effort will play important roles in bringing the ultimate goals closer to practicality.

* See the author's *Automatic Recognition Techniques Applicable to High-Information Pictorial Inputs*, to appear in the Proceedings of the 1962 IRE International Convention.