Using Radar Imagery for Crop Discrimination: A Statistical and Conditional Probability Study

R. M. HARALICK, F. CASPALL, and D. S. SIMONETT Department of Electrical Engineering, and Department of Geography, Center for Research in Engineering Science, The University of Kansas, Lawrence, Kansas.

Abstract

A number of the constraints with which remote sensing must contend in crop studies are outlined. They include sensor, identification accuracy, and congruencing constraints; the nature of the answers demanded of the sensor system; and the complex temporal variances of crops in large areas. Attention is then focused on several methods which may be used in the statistical analysis of multidimensional remote sensing data.

Crop discrimination for radar K-band imagery is investigated by three methods. The first one uses a Bayes decision rule, the second a nearest-neighbor spatial conditional probability approach, and the third the standard statistical techniques of cluster analysis and principal axes representation.

Results indicate that crop type and percent of cover significantly affect the strength of the radar return signal. Sugar beets, corn, and very bare ground are easily distinguishable, sorghum, alfalfa, and young wheat are harder to distinguish. Distinguishability will be improved if the imagery is examined in time sequence so that changes between times of planning, maturation, and harvest provide additional discriminant tools. A comparison between radar and photography indicates that radar performed surprisingly well in crop discrimination in western Kansas and warrants further study.

Introduction

Agriculturists, geographers, and others are constrained in the study of crops with remote sensors by the kinds of agricultural characteristics that the sensors can in fact detect. The sensors used in the various portions of the electromagnetic spectrum are sensitive to dissimilar energy-matter interactions. They detect different characteristics and therefore convey different kinds of information. In the visible region molecular absorptions produce the color effects which convey information about crop condition; in the infrared region, the cycle of thermal response under an insolation load may give information about moisture stresses within crops; in the radar region, the backscattered return is primarily related to surface roughness and dielectric, which may be related to crop type, percent of cover, moisture content, and similar variables.

Although, at first sight, radar cannot sense "color" for its information on crop condition, it may detect related changes in crop geometry and moisture and thus convey information about crop type and state which is not immediately obvious. Investigators may wish radar and other remote sensors to perform the following jobs.

1. Detect the presence of different crops: Is the region entirely homogeneous or are there different crops?
2. Determine the crops actually present.
3. Determine the boundaries and acreages of different crops.
4. Determine the vigor (state) of crops.
5. Determine the agent responsible for any loss of crop vigor.
6. Predict yields.

This paper is concerned with the second item: determining the crops actually present. Studies of crop discrimination are now underway at Purdue University, the universities of Michigan, California, and Kansas, the U.S. Department of Agriculture Experiment Station at Weslaco, Texas, and elsewhere with a view to the ultimate use of remote sensor systems in automated or partially automated data collection and analysis of crop type, condition, and yields. Such data would be valuable in the United States, especially in underdeveloped regions where adequate data on crop type and similar information are not available.

Because radar exhibits the combined advantages of fine resolution, virtual all-weather capability, and independence of the sun for illumination and, moreover, is without serious degradation by the atmosphere, its abilities and shortcomings as a sensor of crops, pastures, and natural vegetation are matters of no small moment to the agricultural community. A great deal of research will be needed to define clearly these abilities and shortcomings, since little work has been carried out on radar as a sensor for agricultural purposes. The study reported here is one of a handful which explores these topics (Meyer, 1967; Schwarz and Caspall, 1967; and Simonett *et al.*, 1967). All are concerned with a single frequency and one (less frequently two or four) polarization(s) only, and thus fall well short of an ultimate capability for multifrequency, polypolarization radars that are feasible for agricultural surveillance. The underlying concept of synchronous data acquisition with multifrequency polarization radars is that it represents the microwave equivalent of multiband spectral reconnaissance in the photographic and thermal IR region, which, as Hoffer (1967) has shown, has produced some encouraging results in tests at Purdue University.

The Need To Use Both Spatial and Temporal Data to Effect Crop Discrimination

In order ideally to separate crop types with a single sensor, we hope that there will be a narrow spread in the probability distribution of the energy recorded by the sensor for each crop; we also hope that the probability distributions of each crop will not overlap. This is of course not the case with radar or *any* sensor because, in the discrimination between crops, the attribute we are seeking—genetic difference between crops—must be approached through the surrogate of crop reflectance and/or emission as modulated by variations in the source, the intervening atmosphere, and so on. The surrogate may be quite imperfectly correlated in a discriminatory sense, with the defining genetic attributes. Furthermore, there is considerable natural variability in the time of planting, varieties planted, and conditions of crops as a function of time of planting, soil type, and so on. Consequently no single sensor is able to effect clear-cut discrimination, at a single time, of all crops in all fields. It becomes necessary therefore to increase the probability of correct discrimination of crops through the use of a number of wavelength bands (and polarizations) in concert, at the same and at different times.

Only through empirical testing of numerous sites in the United States (and elsewhere) at repeated intervals throughout a number of growing seasons (probably at least five) will we be able to define adequately the time-sequential probability distributions of radar and other sensor returns for different crops. The development of "ground truth" sample data on training sets for each flight mission in regions which are homogeneous as to crop assemblages and geography needs to be experimented with to determine, for different areas and conditions, the "appropriate" sampling techniques and rates. It is also apparent, because crops have high spatial and temporal variance from the simple circumstances of crop geography, that the distribution of the recorded energies within a single growing season and between growing seasons will be such that both multiple looks in time and a number of information channels will be required to ensure identification. However, the various trade-offs between number and kind of sensors and number and timing of data collection by the sensors has yet to be examined systematically in relation to regional crop time tables but will be an essential ingredient of continuing studies in the United States.

Consider, for example, the condition of crops and their times of planting and harvesting in the United States. There are many differences in kinds of crops, growing season, and growth patterns between west, center, and east; between humid, subhumid, and irrigated; and between south, center, and north—to name only the more obvious factors involved in crop geography. This means that remote sensors can be usefully employed from March through October because critical changes in the status of crops are occurring somewhere in the United States. In summary, multiple looks through time coupled with judicious ground truth samples for development of training sets for pattern recognition may turn out to be very important in agricultural remote sensing.

Crop Discrimination Accuracy: Identification and Congruencing Problems

Identification Accuracy

We may want various degrees of crop identification accuracy from remote sensing systems. For example, we may wish (a) to identify without error the crops in every field in an area and to determine the acreages of each field with an error of 1% or less; or, less difficult, (b) to identify, with some acceptable error, say (p=0.05), the kind of crops in each field of a region, and to estimate total acreages of each crop in the region accurately to $\pm 5\%$; or, least difficult, (c) to identify a single crop with, say, 5% error at a time of the year when it is most sharply differentiated from other crops.

If we want to determine the exact acreages of a crop being grown in an area, either for statistical reporting purposes or for checking on wheat and feed grain program compliance, we have a situation described by (a) above. In such a circumstance, if ironclad identification cannot be made in a single pass, then multiple passes must be made, with the attendant problems of congruencing: storing spatial-reference coordinates for all points on the ground; and rendering time-separated, different-flight-path data in congruent planimetric geometry. This is no mean task! At present it is a very severe burden to place on a remote sensor system.

Alternatively, as in (b), we may be able to avoid the necessity for many passes by allowing modest errors in identification and by accepting regional average dominant crop acreages with as much as a 5% error. If we select this alternative, a less stringent requirement for data storage and spatial accuracy needs to be met.

Case (c) is notably less demanding than the two preceding ones. A single pass with a few channels may be adequate, depending on the crop to be discriminated and those against which it is to be viewed.

Congruencing

As we have already mentioned, one of the major problems in obtaining high identification accuracy with any automated multisensor system is that of bringing the images into mutually congruent geometry. Each resolution cell of each image corresponds to a small area on the ground. The system must be able to determine the correspondence, in planimetric coordinates, between each resolution cell of each image and each resolution cell of every other image, and between each resolution cell and each small area on the ground. This correspondence must be established automatically if the system is to work quickly and effectively.

The problem of image congruence has been attacked by researchers investigating the visible and infrared region of the spectrum. Multiple wavelength systems using mechanical rotating mirror-scanners have been developed in which the signal received is split into different wavelength bands, each with its own detector. The signals are then recorded similtaneously on different channels of a magnetic tape. A similar technique is feasible with radar systems operating with multiple polarizations and wavelengths, although direct recording of the signal requires a fairly wide-bandwidth recorder. The use of some form of intermediate storage so that only the actual significant bandwidth of the signal need be recorded can reduce this problem (sometimes hundreds of pulses are averaged to make one resolution cell).

A more severe problem of image congruence comes from multiple flights. If an aircraft or spacecraft carrying the radar could fly exactly the same path within a wavelength or so each time, the problem would not be difficult. Unfortunately, aircraft cannot repeat either horizontal or vertical location of the path to anything like wavelength precision; furthermore, operational conditions may require significant differences in altitude which result in distortions between images. Even the best-programmed flight lines sometimes deviate, and then there are horizontal as well as vertical differences.

If we did not need to know exactly where each crop occurred in a region, a relatively crude job of congruencing could be performed. Such unsophisticated congruencing would retain just enough statistical structure to tell us the acreage of each crop, independent of its spatial position, or the percentage of occurrence for each crop in the region. Little, if any, work has been done in this area, and we believe it warrants further thought.

The need for automatic congruencing equipment is clearly urgent. Multiple recording during a single pass solves only the problem for that pass. No system known produces images on successive flights that are initially congruent. Digital correlation and subsequent linear or non-linear correction of scale would be feasible if time were not a factor. However, with large masses of information collected in an extensive flight program, digital techniques for congruencing would be ineffectively slow. Some faster means must be found to automate solutions for the urgent problems in agriculture.

There is another approach to the image-congruencing problem which relies on the fact that there are geographic boundaries which are relatively permanent. Homogeneous regions (determined by "natural" statistical classes), independent of other images, may be found on each image. Boundaries must exist between homogeneous regions. Congruencing may be achieved by distorting the geometry of the images so that there is maximum correlation between the boundaries on each image with those on every other image. The distorted ("congruenced") images may now be used for classification purposes. The end product of the classification may be brought to planimetric correspondence with the ground by examining the boundaries around the categories and distorting the geometry of the categorized image so that it matches the natural boundaries of the area.

Data Analysis

The Threefold Approach of Data Representation, Classification, and Clustering

Sensor systems are designed to obtain information about an interesting statistical structure of an environment, where "interesting" is defined by the investigator. The sensor system must reflect or transmit this interesting environmental structure onto the data structure if anything is to be gained by using the sensor system. Thus, once a sensor system has been conceived and designed, the question that must be answered when it is tested is: "What statistical structure of the environment is actually preserved through the sensor's 'eyes'—in particular, is the interesting structure preserved?"

In order to answer this question, some way of examining the data, some way of determining the structure within the data, must be devised. Then the structure of the data may be compared with the known structure of the environment, and the question can be answered.

The most usual type of environmental structure which is defined as "interesting" is that of a classificatory nature. The investigator arranges the environment into a set of mutually exclusive categories. It is his hope that the sensor system will preserve enough of the environment's structure that data obtained from the sensor system may also be categorized into a set of mutually exclusive classes, where each class corresponds to a unique category in the environment.

The following techniques of data representation, data classification, and data clustering are valuable in determining the structure of the data.

1. The data set may be represented in some geometric way such as a scattergram, in which the data structure is immediately apparent. This representation must be perceivable directly as a "picture."

2. A statistically optimum classification of the data may be made; the optimality criterion is given by the investigator.

3. The natural classes of the data structure may be found and compared with the categories in the environment.

The information received in a "picture" of the data structure gives the investigator an immediate "feel" for what is actually happening. It gives him an intuitive base which he can expect to guide him to more sophisticated techniques. It may also tell him quickly whether there have been faults in the sensor system while the system was gathering the present set of data. If so, then the data would not have to be processed further and the experiment could be rerun. The information gained by the statistically optimum classification of a data set is, of course, most fundamental in the sensor system test. From it the investigator finds out how good the system really is in preserving the classificatory structure of the environment. Finally, by finding out the natural classes in the data structure, the investigator can learn "what language the sensor system speaks." If a natural class in the data structure actually corresponds to two environmental categories, then we would expect that, regardless of the optimality of any statistical classifying rule, these two environmental categories would be confused, and the investigator can do nothing but search for a new sensor system. This confounding of categories must then be taken into account in any interpretation of data from the sensor system.

When the sensors take data in a sequential way and it is known that observations which are close to each other in the sequence are likely to be in the same category, some effort should be made to preserve the similarity of neighboring data points in their classification. For imaging sensors in particular, this means that two neighboring resolution cells should not be classified independently. An exception to this occurs when the image data are like those of Fig. 1, which has the kind of ideal structure illustrated in the scattergram in Fig. 2. Here there is complete isolation between the categories: bare ground occupies the lower third of densities, grain sorghum occupies the

.11 BA	.15 BA	.09 BA	.80 CR
.08 BA	.10 BA	.71 CR	.75 CR
.41 GS	.47 GS	.77 CR	.74 CR
.51 GS	.50 GS	.53 GS	.45 GS

Key:

BA Bare ground
GS Grain sorghum
CR Corn

Fig. 1. Simple image, where the density for each resolution cell is indicated by the number appearing in the top of each box and the category of surface cover which corresponds to the cell is indicated by the two-letter code appearing at the bottom of each box.

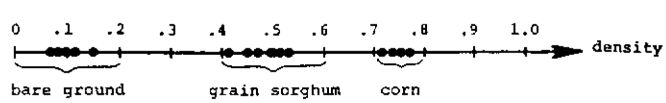

Fig. 2. Scattergrams of simple image data: The categories are all isolated.

middle third of densities, and corn occupies the upper third of densities.

Principal Components Representation

Representing the data structure in a two- or three-dimensional scattergram is a problem when the measurements the sensor system takes are n-dimensional (n being more than 3). Some criterion for choosing the representation must be determined. If we decide that the representation is to be attained by a linear transform, because linear transforms are easy to find and quick to work with, then in effect we are projecting the data on the two-dimensional subspace which optimally retains the basic data structure, losing only the fine features. If we translate the data by its mean, this linear transform is defined by the two eigenvectors of the covariance matrix with largest eigenvalues (see Kendall and Stuart, 1966). We call this representation the principal axes representation.

It can be shown that this representation corresponds to finding a plane in n-dimensional space where the data projected onto this plane have maximum variance. Mathematically the situation can be represented as finding the projection operator P which has rank 2 and a translation vector t such that the mean squared error ϵ^2, defined by $\epsilon^2 = E[((x-t)-P(x-t))' ((x-t)-P(x-t))]$, is minimized. The minimization is achieved when $t=E[x]$ and P projects onto the plane spanned by the two eigenvectors, v_1, v_2, with largest eigenvalues of $E[(x-E[x])(x-E[x])']$, the covariance matrix. The coordinates of any data vector relative to the principal axes' plane can be found easily as $(x-E[x])'v_1$ for the first component and $(x-E[x])'v_2$ for the second component.

Bayes Classification

Although there are many ways of classifying data, such as: linear decision theory, which constructs hyperplanes boundaries or nearest-neighbor search procedures, they are statistically admissible only under severely restrictive conditions. In addition to allowing the use of a loss function which can weigh a correct classification more heavily on the more important categories, a Bayes decision rule is almost always statistically admissible and optimal. Therefore, to get information regarding the best possible classification which can be made, a Bayes decision rule may be employed.

We used a Bayes decision rule (see Middleton, 1960) to classify each resolution cell of the radar imagery data independently of its neighboring cell. This was done because, at the time of this study, we did not have a Bayes decision program taking into account the similarity of neighboring resolution cells. In order for the Bayes decision rule to indicate the classificatory structure of the data rather than merely to "memorize" the data, the data set can be randomly divided in half: the first half can be usedto generate the statistics for the Bayes rule, and the second half can be classified on the basis of this rule. A "goodness criterion" may be adopted, such as the total percentage of correct classification or the average percentage of overall categories of correct classification in which the percentages are properly weighed by the loss function adopted.

Clustering

To find the natural classes in the data structure as seen through the sensor's eyes is perhaps the hardest of the three operations suggested, because the problem is harder to define. Certainly some information about the natural classes will be obtained from the principal axes representation and the Bayes classification. Categories which are confused will definitely appear. However, it may be that parts of the environment which are categorized by the investigator as the same will occur in different natural classes, indicating that further distinctions may be made.

We will approach the natural-class problem of image data by two clustering methods. One sequentially forms clusters from data measurements, called centers, which have the highest probability in any small spatial region in the image (a spatial conditional probability approach), and includes in each cluster the points which are sufficiently close to the center in measurement space and whose probabilities do not differ too much from the probability of the center point. The other method forms clusters taxonomically on the basis of Euclidean distance in measurement space, using the multiple-linkage model.

Clustering by Spatial Conditional Probability

The conditional probability approach first quantizes the data in the following way. Let the density in the ijth resolution cell be $X^{ij}=(x_1^{ij}, x_2^{ij}, \ldots, x_n^{ij})$ where x_k^{ij} is the output from the kth sensor when the kth sensor is observing the area represented by the ijth resolution cell. Let R_k be the range for the kth coordinate. The quantizing function Q is defined as $Q(X)=(i_1, i_2, \ldots, i_n)$, where $(i_{j-1}/10)\ R_j \leq x_j < i_j R_j/10$.

In forming clusters, priority is given to the quantized measurement vectors for which the conditional probability of occurrence in an image spatial region is higher than the joint probability of its occurrence over the whole image. For example, a multispectral image may be partitioned into a number of connected spatial regions, as in Fig. 3. The conditional probability distribution for each region is computed and compared with the joint probability distribution for the entire image. All the measurement vectors which have the highest conditional probability in the multispectral images are given priority in forming clusters. Each cluster is formed around a center, that is, a

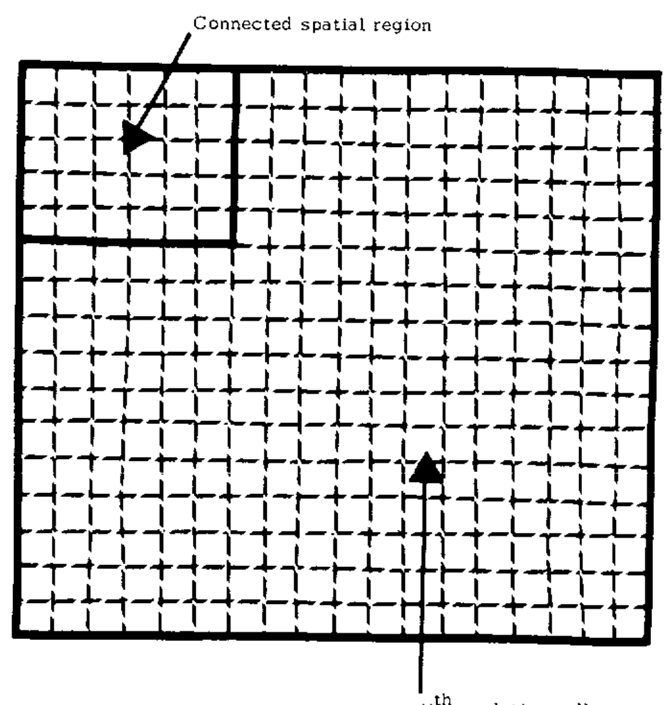

Fig. 3. Multispectral image.

quantized measurement vector which is not yet in a cluster and which has the highest conditional probability for some spatial region. At any iteration, for a point X to be included in the cluster just being formed, there must exist a point Y already in the cluster such that:

1. $|x_i - y_i| \leq |1$, for each i.
2. $P(X) \leq \epsilon_1 P(C)$, where C is the center and ϵ_1 is a prechosen parameter.
3. $P(X) \leq \epsilon_2 P(Z)$, where Z has the highest probability of any point already in the cluster and ϵ_2 is a prechosen parameter.
4. $P(X) \leq \frac{1}{\epsilon_2} P(W)$, where W has the lowest probability of any point already in the cluster.

If there are any points which do not meet these criteria, they are put into the nearest cluster. A fuller account of the conditional probability approach is given in Haralick and Kelly (1969).

Clustering by Multiple-Linkage Techniques

Multiple-linkage cluster analysis has been widely used as a numerical taxonomic technique in the biological sciences (see Sokal and Sneath, 1963). However, many of the assumptions made in classifying biota cannot be made in analyzing radar images, because there is no evolutionary relationship between the objects being categorized. Nevertheless, since the measurement vectors which can be clustered together are in some sense similar, the technique is worthwhile in finding the natural classes of a data set.

The clustering procedure takes the following form.

1. A matrix of the similarity coefficients between each pair of measurement vectors is constructed. (For the multispectral images, these are the measurement vectors in each resolution cell.)

2. The matrix is scanned and all pairs of measurement vectors which are mutually close are linked together.

3. A new matrix is computed on the basis of the paired vectors and single vectors. This new matrix is scanned, and all single vectors or paired vectors which are mutually close are linked together. This step is continued iteratively until all clusters are linked together.

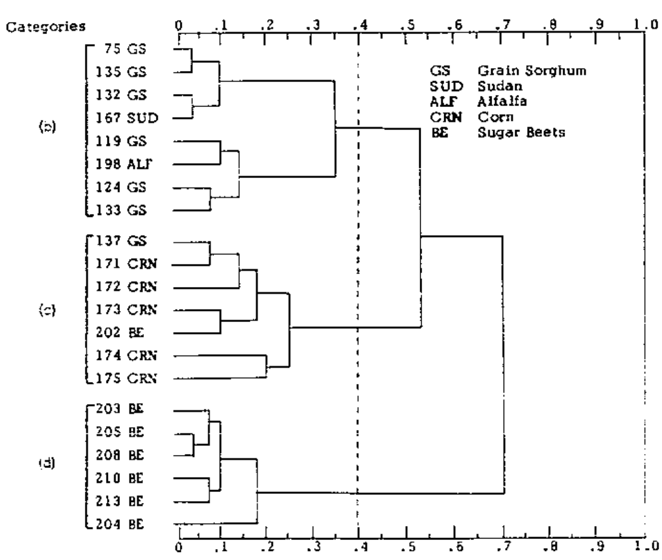

Fig. 4. Section of a dendogram showing graphically the results of cluster analysis. This is a section of the actual dendogram used in categorizing the August and September data. The section of 21 data points was selected from the original 231 because it shows portions of three of the four major categories clustered. The dashed line shows the break point.

Figure 4 illustrates the dendrogram obtained from the multiple-linkage clustering.

The investigator may determine the natural-class structure of the data by truncating the dendrogram at some distance, as illustrated by the vertical dashed line in Fig. 4. Here a distance coefficient of 0.4 gives a reasonable picture and the natural classes are (b), (c), and (d), as indicated in the left margin. The dendrogram may also be used as a guide for further analysis. For example, the sugar beet field (202 BE) which linked with the corn field should be investigated in greater detail to determine why it is so different from the rest of the sugar beet fields.

There are some important differences between the conditional probability clustering technique and the multiple-linkage cluster analysis. The multiple-linkage cluster technique analyzes the data microscopically, linking each observed measurement vector with similar observed measurement vectors and simultaneously forming a set of clusters. In order to do this, an $N \times N$ matrix must be stored if there are N observed measurement vectors. The conditional probability clustering, a macroscopic technique, starts with the most important cluster, forms the clusters sequentially, and needs a storage of $2N$ places. These differences reflect themselves in the amount of data the techniques can handle. The conditional probability clustering technique is quick and is easily used with data sets containing thousands of measurement vectors; the multiple-linkage clustering is slower, the data sets containing even hundreds of measurement vectors are at the limit of the procedure. However, the multiple-linkage clustering compensates for these disadvantages by giving a more detailed analysis of the data set.

Statistical Crop Discrimination

The region where the crop information is collected lies a few miles west of Garden City, in western Kansas (Fig. 5). In this area there are over 400 fields for which data were collected at the time of radar overflights. The area was chosen because fields are generally large, slopes are extremely gentle, soils are uniform, both dry-land and irrigated cropping is practiced, and there is some diversity of crop types, agricultural techniques, and moisture conditions. In this test area we can hold constant in an analysis such factors as topography, soil, and field size, and can obtain the maximum variations in soil and crop states, moisture levels, and so on, in order to explore quantitatively both crop discrimination and the contribution of different parameters to the radar backscatter.

We shall illustrate the threefold approach of data representation, data classification, and data clustering to two sets of data obtained from K-band radar imagery taken over the Garden City area. At the time the analysis was done, image digitization equipment was not available at the University of Kansas, and we had to satisfy ourselves with only the average radar return in each farming field. This field average was obtained by taking a few random scans over each field in the images with a microdensitometer and using the average of the scan lines as an estimate of the average radar return for the field. It should be noted that, since we are working with averages and the variance of averages about their true values are less than the variances of the gray tones from the resolution cells which make up the

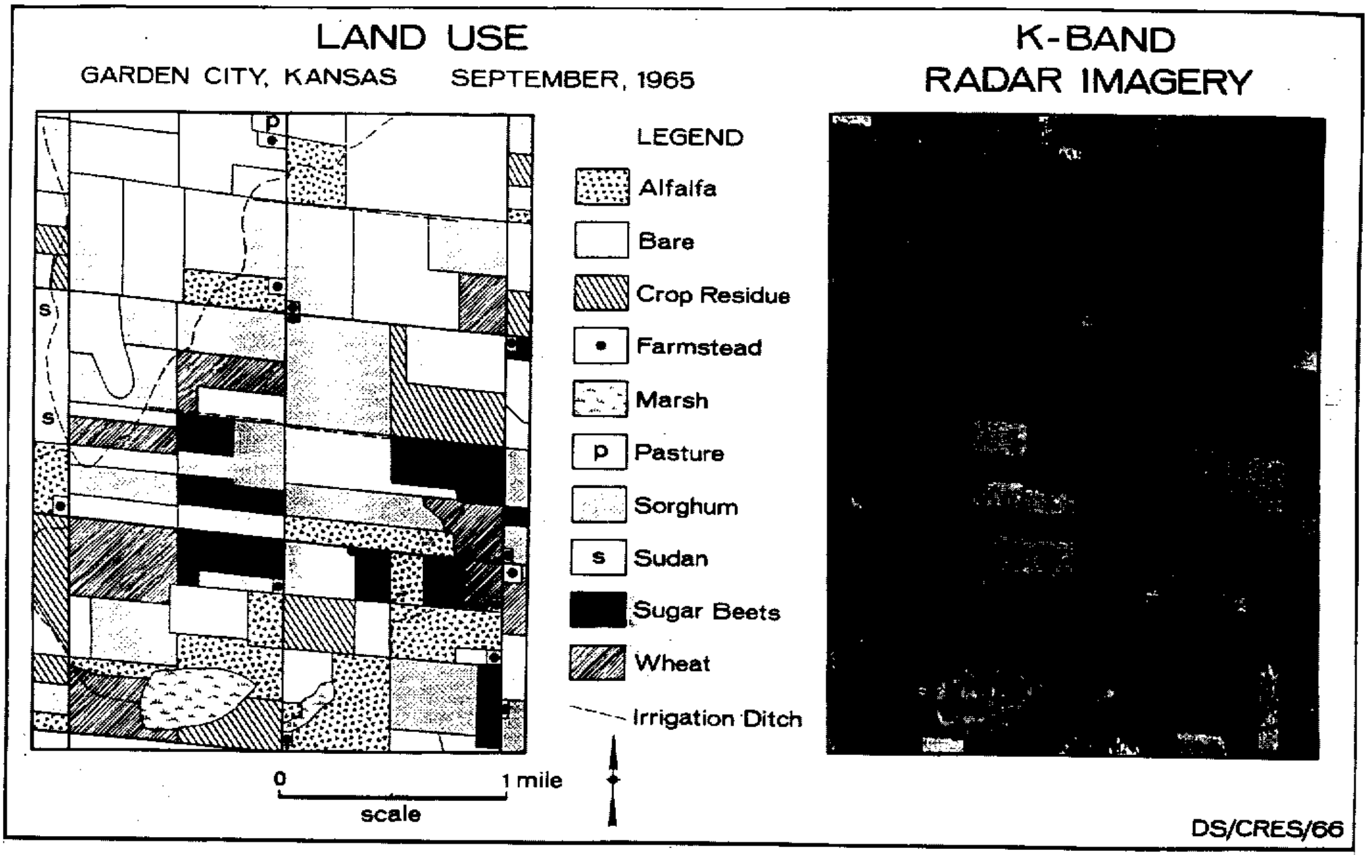

Fig. 5. Radar Imagery.

averages, the results obtained are better than those which would be obtained on the basis of the gray tones of the individual resolution cells.

The first set of images was obtained in a time-sequential fashion during the months of August and September in 1965. The August image had only one polarization, being horizontally transmitted and received, while the September images had two polarizations, being horizontally transmitted and both horizontally and vertically received. All three images were taken over the same test site near Garden City, Kansas. The second set of images, also taken over this area, was obtained during the month of July in 1966 and had all four polarization combinations. We wanted to find out:

1. the environmental structures related to crop type which the radar imagery system preserves,
2. while recognizing the paucity of our data, whether more environmental structure was preserved
   (a) if the data were taken in time sequence with few polarization combinations, or
   (b) if the data were taken at the same time with all the polarization combinations.

Figures 6 and 7 illustrate scattergrams with the principal axes representation for these data. It may be noticed from the general pattern that, as the agriculture becomes more intense (surface-cover or land-usage categories going from bare ground to sugar beets), the radar return becomes greater. This correlates most directly with the crop percent of cover. For both sets of data we immediately see the overlapping of categories. In the July data the bare ground and wheat stubble overlap. Then the other part of wheat stubble overlaps with alfalfa, while the alfalfa overlaps with the corn, and grain sorghum overlaps and covers both alfalfa and corn. Sugar beets stand almost isolated by themselves. In the time-sequential data, bare ground stands out by itself, while wheat stubble overlaps with alfalfa and both overlap with the grain sorghum. The corn and sugar beets stand by themselves. This indicates that, for the July data, we would expect the Bayes decision rule to confuse wheat stubble, alfalfa, corn, and grain sorghum with each other, and to confuse bare ground and wheat stubble with each other while classifying

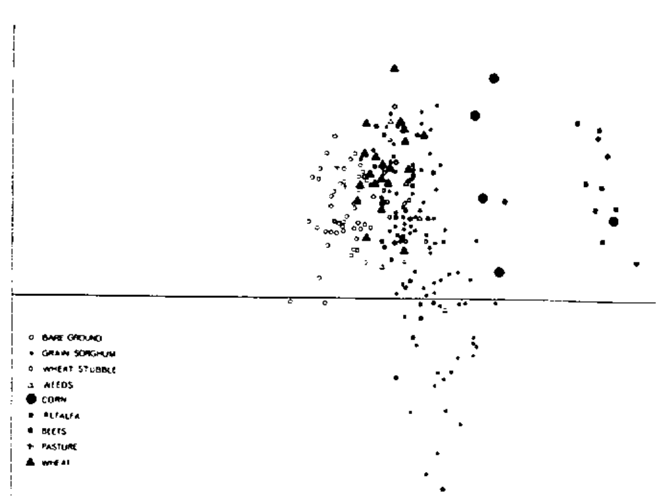

Fig. 6. Scattergram of August–September 1965 radar imagery field averages projected onto first two principle axes.

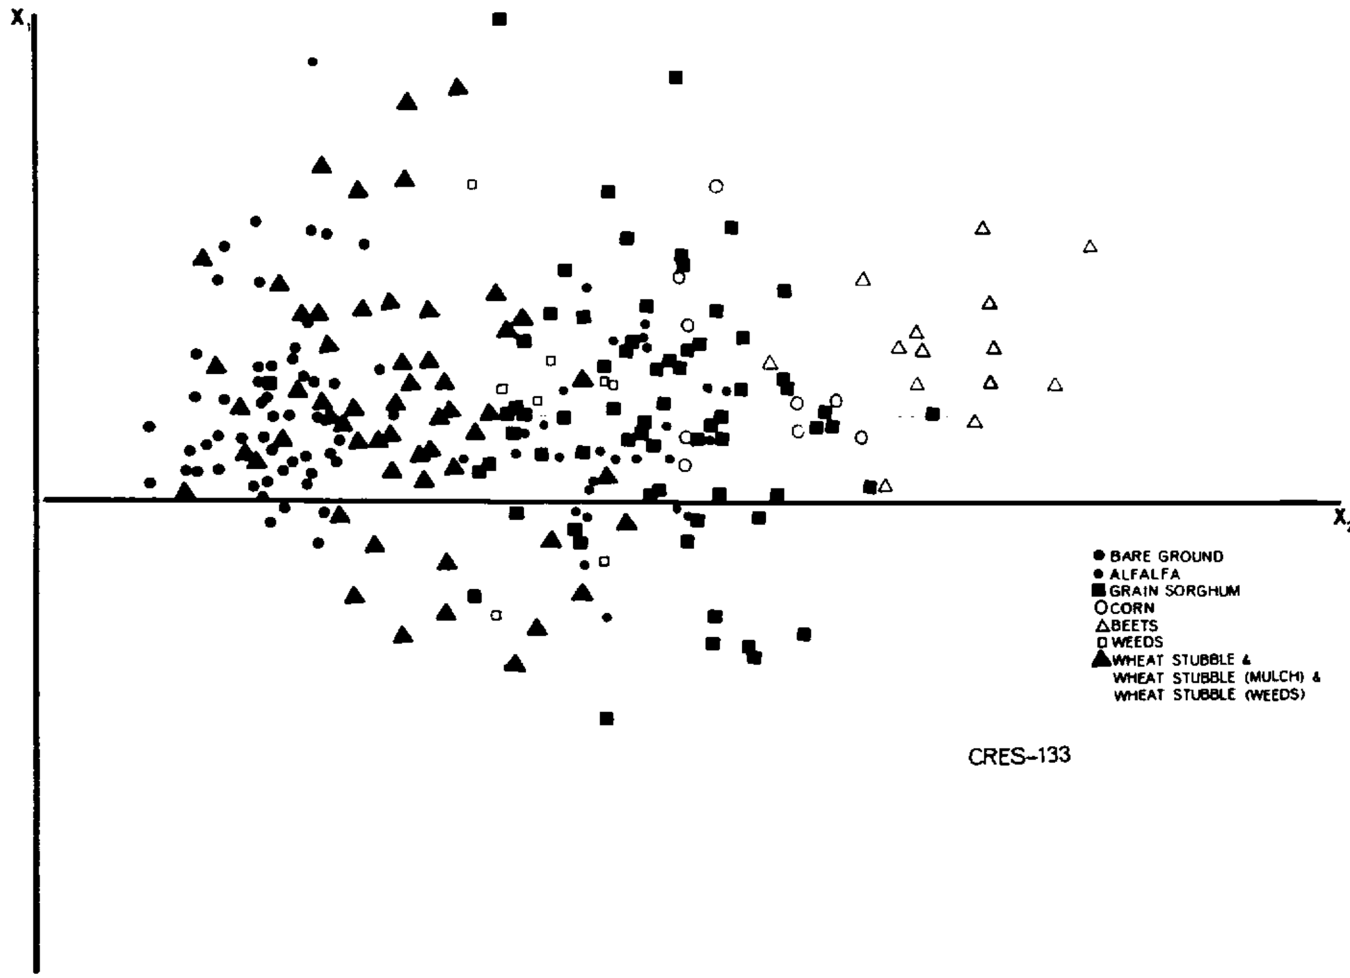

Fig. 7. Scattergram of July 1966 radar imagery field averages projected onto first two prinsiple axes.

sugar beets fairly well. For the time-sequential data we would expect the Bayes decision rule to confuse wheat stubble, alfalfa, and grain sorghum with each other, while classifying bare ground, corn, and sugar beets fairly well. This is indeed true if we examine Tables I and II, which indicate the results of the Bayes classification. For this problem the loss function was chosen to be an equal loss function; the same finite negative amount was lost for each correct classification, and nothing was lost for an incorrect classification. Notice that 90% of the data grouped as shown were correctly classified for August-September, while only 78% of the July data were correctly classified. We would expect a natural-class approach to show that the categories confused in the Bayes classifying technique all belonged to the same natural class or cluster.

The spatial conditional probability approach could not be fully utilized because of the lack of image-digitization equipment. However, a suitable approximation was tried. All field averages which were measurements of the same land usage category were lined up together. In this way the spatial closeness of similar types of fields in the data sequence could be utilized exactly as it would have been if the data sequence represented resolution cells of an image. Each spatial region consisted of a set of 10 sequential vectors in the sequence. The error from this approximation results mainly from the fact that clustering procedures of this type require large amounts of data to work well, and using field averages for resolution cells reduces the data size by two or three orders of magnitude. Despite this handicap, the conditional probability clustering procedure indicates that the categories of vegetation which were confused in the Bayes classification do form fairly isolated clusters or similarity sets, as shown in Tables III and IV. Although the Bayes procedure categorized the August-September data better than the July data, the clusters found in the July data correspond better to the surface cover categories than those clusters found in the August-September data. The reason was that there was no physical separation between the confused surface cover categories, as may be seen from the scattergrams. We should note that these clusters were formed without any information regarding "ground truth," i.e., knowledge concerning the vegetation category which each measurement vector actually represents.

The Euclidean distance-similarity coefficient from the multiple-linkage clustering technique was found to discriminate crop types nearly as well as the other methods of categorization (see Tables V and VI). However, little discriminatory potential was indicated for the correlation coefficient. This suggests that depolarization of the radar signal from crop to crop is sufficiently variable and perhaps modest enough in degree that good clusters cannot be formed with the data. The correlation similarity coefficient was not used in later analyses.

The question of the operation of reciprocity may also be raised in this connection, as indeed it should also be when both HV and VH images are used for measures with Euclidean distance-similarity coefficients. The principle of reciprocity states that, when the full polarization matrix (HH, VV, HV, VH) is obtained at the same instant of time, the two cross-polarization terms (HV, VH) will be identical. Thus they should not be used in clustering with either Euclidean distance or correlation coefficient measures, for adding an extra channel unduly emphasizes the cross-polarized return and, to the extent that it does not contribute crop discriminatory information, it will add

Table I
Identification Accuracy for August-September 1965 Radar Imagery[a]

Surface Actually Covered	Number of Measurements Identified as: (a) Bare ground pasture	(b) Grain sorghum, bare ground-wheat, alfalfa wheat stubble-weeds	(c) Corn	(d) Sugar beets
(a) Bare ground pasture	22 (91 %cor.id.)	2	0	0
(b) Grain sorghum, wheat stubble-weeds, alfalfa, bare ground-wheat	5	73 (91 %cor.id.)[b]	2	0
(c) Corn	0	0	2 (67 %cor.id.)	1
(d) Sugar beets	0	0	1	5 (83 %cor.id.)

[a]A random sample of 113 measurements of field averages from a total set of 226 were used to train a Bayes decision classifier, and the other 113 measurements of field averages were classified on the basis of the Bayes decision rule. Ninety percent of the types of surface cover, grouped as shown, were correctly identified. Each field had the indicated type of surface cover for the months of both August and September, the exception being bare ground-wheat, which was bare ground in August and wheat in September.

[b]The 91 % correct identification in groups (a) and (b) means that 91 % of the measurements in each of these groups were actually identified as being in that group. In general the individual types of surface cover within each group were completely confused with one another.

Table II
Identification Accuracy for July 1966 Radar Imagery[a]

Surface Cover Actually Measured	Number of Measurements Identified as: (a) Bare ground	(b) Corn, alfalfa, grain sorghum, weeds, pasture, wheat stubble	(c) Sugar beets
(a) Bare ground	14 (52 % cor. id.)	13	0
(b) Corn, alfalfa, grain sorghum, weeds, pasture, wheat stub.	11	75 (85 % cor. id.)[b]	2
(c) Sugar beets	0	1	6 (86 % cor. id.)

[a]A random sample of 129 measurements of field averages from a total set of 251 were used to train a Bayes decision classifier, and the other 122 measurements of field averages were classified on the basis of the Bayes decision rule. Of the types of surface cover shown, 78 % were correctly identified in the above groupings.

[b]The 85 % correct identification means that 85 % of the measurements in group (b) were actually identified as being in group (b). In general, the individual types of surface cover in group (b) were completely confused with one another.

Table III
Conditional Probability Clustering for August-September 1965 Radar Imagery[a]

Surface Cover Actually Measured	Number of measurements Identified as: I	II	III
(a) Bare ground, sudan, grain sorghum, alfalfa, bare ground-wheat, wheat stubble-weeds	205	4	0
(b) Corn	2	3	0
(c) Sugar beets	0	1	11

[a]Six clusters were asked for, of which four were collapsed together since they had the most similar category distributions. $\epsilon_1=0.27$, $\epsilon_2=0.36$.

Table IV
Conditional Probability Clustering for July 1966 Radar Imagery[a]

Surface Cover Actually Measured	Number of Measurements Identified as: I	II	III
(a) Bare ground, wheat stubble	99	17	0
(b) Alfalfa, grain sorghum, corn, weeds, pasture	10	115	0
(c) Sugar beets	0	3	11

[a]Four clusters were asked for, two of which were collapsed together since they had the most similar category distributions. $\epsilon_1=0.14$, $\epsilon_2=0.36$

Table V
Identification Accuracy for August-September Radar Imagery Using Complete Linkage Cluster Analysis as the Grouping Technique[a]

Surface Cover Actually Measured	Number of Measurements Identified as: (a) Bare ground, pasture	(b) Grain sorghum, bare ground-wheat, alfalfa wheat stubble-weeds	(c) Corn	(d) Sugar beets
(a) Bare ground, pasture	47 (98 % correct identification)	1	0	0
(b) Grain sorghum, wheat stubble-weeds, alfalfa, bare ground-wheat	14	151 (90 % correct identification)[b]	1	0
(c) Corn	0	0	5 (100 % correct identification)	0
(d) Sugar beets	0	0	1	11 (92 % correct identification)

[a]The data used were average image densities for monopolarization (HH) August 1965 imagery and for multiple polarization (HH, HV) September 1965 imagery. Multiple-linkage clustering was applied using Euclidean distance as the similarity coefficient (see Sokal and Sneath, 1963, for details of the clustering procedure).

[b]The majority of these fields were young wheat fields very short and lacking complete cover.

Table VI
Identification Accuracy for July 1966 Radar Imagery Using Complete Linkage Cluster Analysis as the Grouping Technique[a]

Surface Cover Actually Measured	Number of Measurements Identified as (a) Bare ground	(b) Corn, alfalfa, grain sorghum, weeds, pasture, wheat stubble	(c) Sugar beets
(a) Bare ground	53 (94% correct identification)	3	0
(b) Corn, alfalfa, grain sorghum, weeds, pasture, wheat stubble	25	159 (84% correct identification)[b]	5
(c) Sugar beets	0	1	13 (93% correct identification)

[a]The data used were 259 measurement vectors each consisting of four multipolarization (HH, HV, VV, VH) image densities. Multiple-linkage cluster analysis was used as the categorizing technique; Euclidean distance was employed as the similarity coefficient (see Sokal and Sneath, 1963, for details of the clustering procedure).

[b]Mostly dry stubble.

noise to the analysis. We included both HV and VH images in this analysis because they were obtained on separate aircraft passes on the same day, and the strict operation of reciprocity will not hold. However, because visual comparison of the images shows in fact that they are much more closely similar than the like images (HH, VV), in practice it is almost as if reciprocity were observed.

In an attempt to glean more information from the analysis of of August-September data, category (b) (see Table V) was subdivided into three separate categories (Table VII). Note that category (b_3) contains mostly grain sorghum, category (b_2) contains mostly grain sorghum and alfalfa, and category (b_1) contains mostly grain sorghum and fall wheat, indicating that there is some discrimination even in the very complex portion of the data set. Comparing this grouping with Fig. 6 shows that it is hard to discriminate such categories in the scattergram. Partial discrimination of this nature leads to the belief that, if data were available in other wavelengths or through time, better discrimination could be achieved.

Temporal Radar Data during the Growing Season

Obtaining data through time leads to the question of the best time to collect the data if better discrimination is desired. Looking again at Table VII, we see that the problem lies in the b_1, b_2, and b_3 categories. To arrive at a solution to the problem of finding the optimum time to gather data from crop-type discrimination, the curves in Fig. 8 were constructed. The solid lines on the curves are based on actual averages of the total data available at the University of Kansas (July, August, September, October, and November). We have based the dotted lines on crop growth patterns and maturity dates.

Using the curves to determine when to collect data which would allow better discrimination of grain sorghum and winter wheat (category b_1), we see that data taken in April or May would allow us to achieve much better discrimination. During these spring months the grain sorghum fields are being prepared for the planting of the crop. On the other hand, winter wheat is in the process of rapid growth, and therefore would give a much higher radar return. Likewise, the same time period would allow much better discrimination between grain sorghum and alfalfa (category b_2). Alfalfa, like winter wheat, is in the early rapid growth stage during the early spring months. The discrimination of alfalfa from winter wheat appears difficult at first glance. However, the curve for winter wheat in Fig. 8 is for the entire winter wheat growing season, which starts in August. Many times, winter wheat will not be immediately replanted in the crop rotation system; that is, winter wheat will be planted on a fallow field or will be followed by a fallow field. Even when a field of winter wheat is replanted after harvest, discrimination may be achieved by concentrated sensing in July and early August, during the bare-ground phase of the cycle. In addition, repetitive sensing during the summer will reveal the cycle of growth and cutting of alfalfa.

In the various tables showing the discriminations between crops achieved with radar it was apparent that the categories in each group were generally indistinguishable; thus the identification of individual crops is notably poorer than for the lumped categories. As we have just discussed, the use of radar at different times will greatly improve crop discrimination (indeed panchromatic black and white photography, say six times, through a growing season is demonstrably capable of discrimination (Schepis, 1968)). It is appropriate also to caution that it would be unwise to infer from the modest results above that radar is of low discriminatory power in distinguishing

Table VII
Six Category Breakdown Using Multiple-Linkage Cluster Analysis Based on the Euclidean Distance Coefficient

Category	(a) Bare Ground and Wheat	(b_1) Grain Sorghum and Wheat	(b_2) Grain Sorghum and Alfalfa	(b_3) Grain Sorghum	(c) Corn	(d) Sugar beets
Bare ground	45	1				
Grain sorghum	1	28	24	37	1	
Weeds and wheat stubble	2	6	8			
Fall wheat	9	18	4			
Alfalfa	1	6	16	2		
Sudan	1			2		
Corn					5	
Pasture	2					
Sugar beets					1	11
No. Correct/Total	54/61	46/59	40/52	37/41	5/7	11/11
Percent Correct	89%	78%	77%	90%	71%	100%

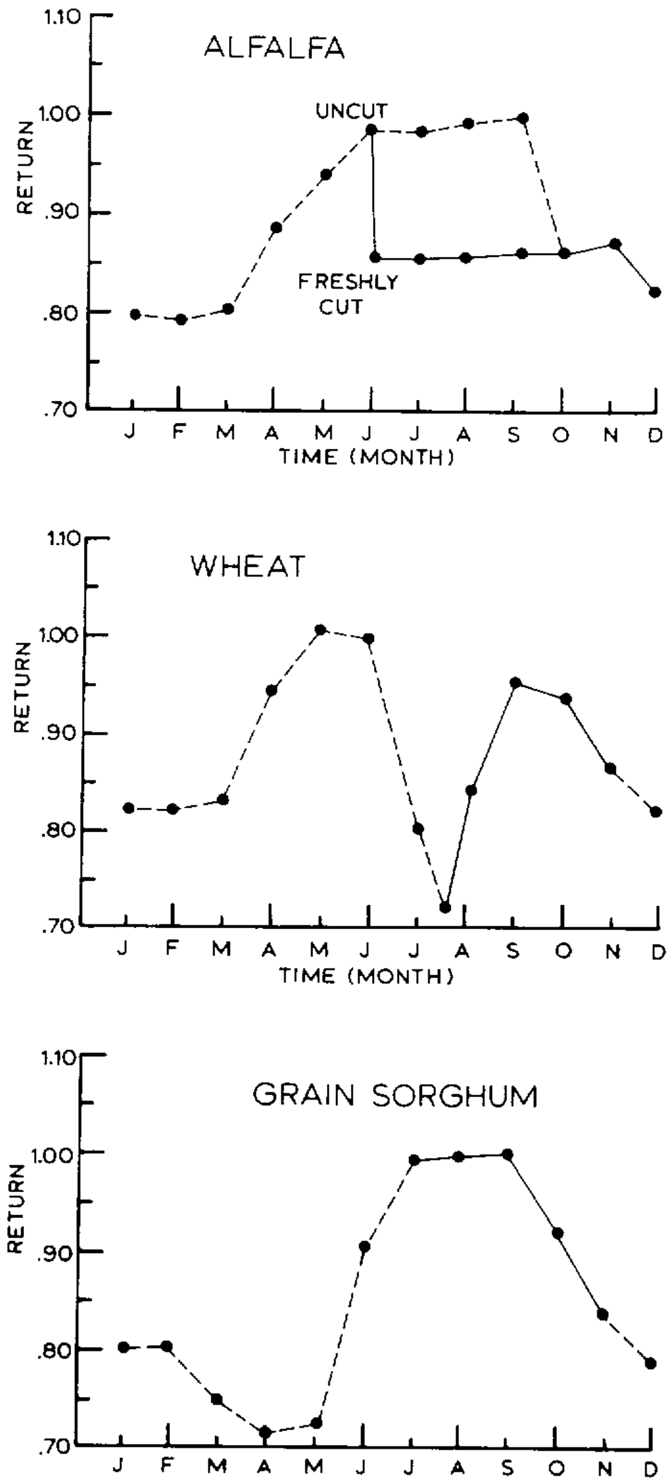

Fig. 8. Relative radar return curves: The relative densities are based on HH polarization K-Band radar image densities. The solid lines are averages of the return for each crop type based upon the temporal data available. The images were obtained over a three year time period and therefore do not represent the average return for a given set of fields. The dashed lines were interpolated using known crop growth cycles, maturity dates, the average return from categories such as bare ground and wheat stubble, and the condition of the crop, (e.g., wheat is bent over and dried somewhat during the winter months, hence the lower return during that time. The upper curve for alfalfa represents growth to a stage suitablefor cutting in summer until the final cutting late in September or early October. The lower curve for alfalfa represents the value of the radar return from freshly cut alfalfa.

crops. Would any single channel (with polarizations) or several channels do much better in the visible or thermal regions?

A Comparison of Radar and Photography

A preliminary test of the last observation by Mr. W. G. Brooner of CRES using color photography and HH and HV radar at Garden City showed to our surprise that, with data for September, radar was notably superior to color photography for discriminating crop types. Although there is no guarantee that comparable results would be obtained elsewhere or at other times, the data do demonstrate that additional research is both needed and justified with imaging radar, in comparison with other sensors in crop discrimination.

Figures 9 and 10 were prepared to compare radar and photographic crop discriminations. Figure 9 is derived from K-band radar (AN/APQ-97) imagery September 1965 at Garden City. The average density values for each field are plotted for the HH and HV polarized images. Figure 10 is based on Ektachrome aerial photography (Kodak 8442) at Garden City, September 27, 1968. The average film densities for each field are plotted for the green-sensitive emulsion layer (490–590 nm) and the red-sensitive emulsion layer (590–690 nm) of the photography. No photography and radar were available for the same dates; these were the closest match in our files.

On both figures the boundaries between crop groups were located by eye and positioned in order to achieve maximum separation of individual crops and crop groups.

In Fig. 9 the y-axis is the like density HH polarization; the x-axis the cross density, HV polarization. Notice that the hyperplanes erected lie at about 45°, implying that both polarizations are contributing information.

In the upper group bounded by the upper hyperplane, there are 30 sugar beets and 1 corn, suggesting that, if this relation-

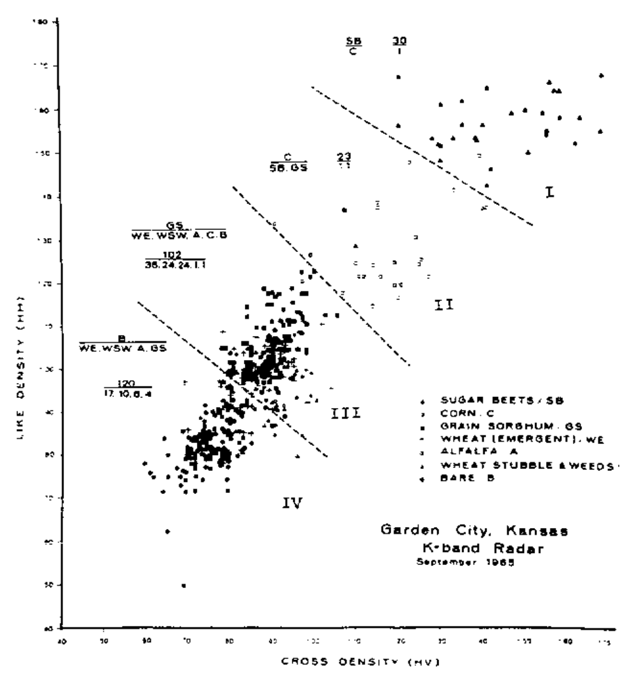

Fig. 9. Crop discrimination with HH and HV polarized radar imagery, Garden City, Kansas, September, 15 1965. Each point represents the average density value for a single field.

Explanation: $\frac{SB}{C} = \frac{30}{1}$ represents $\frac{\text{Sugar Beets}}{\text{Corn}} = \frac{\text{30 fields.}}{\text{1 field}}$

The primary crop in the various categories are: I, Sugar Beets, II, Corn; III, Grain Sorghum, Wheat, and Alfalfa; IV, Bare.

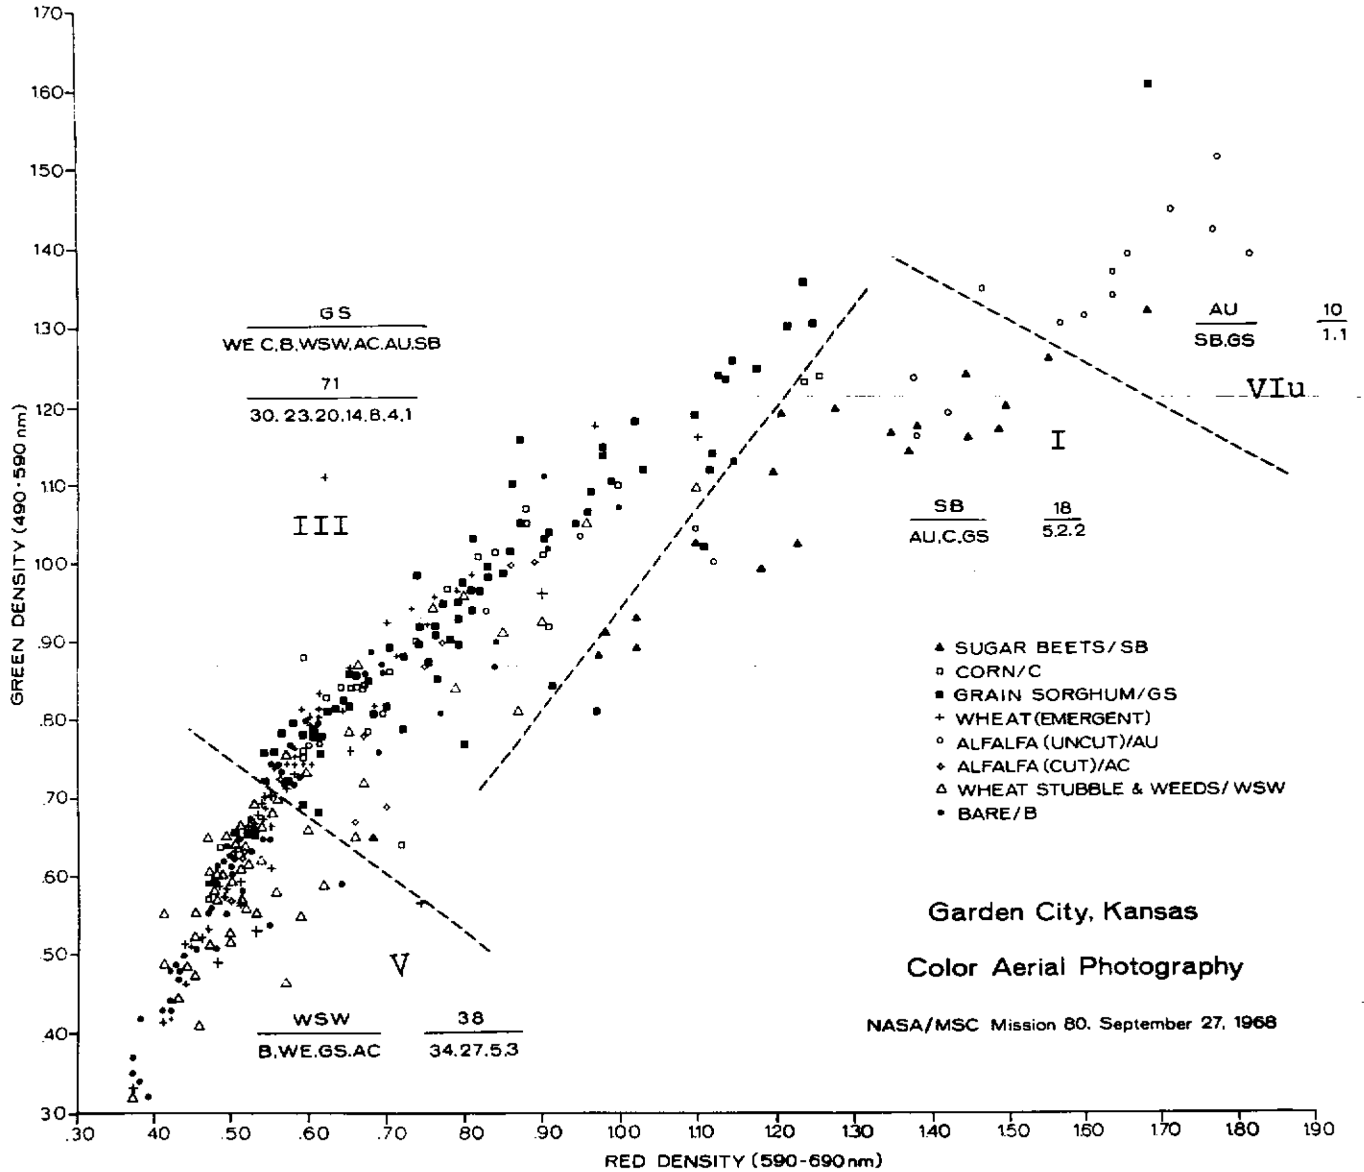

Fig. 10. Crop discrimination with color aerial photography, Garden City, Kansas, September 27, 1968. Each point represents the average density value for a single field.

Explanation: $\frac{\text{AU}}{\text{SB, GS}} = \frac{10}{1,\,1}$ represents $\frac{\text{Uncut Alfalfa}}{\text{Sugar Beets, Grain Sorghum}} = \frac{\text{10 fields}}{\text{1 field of each}}$

The primary crop in the various categories are: VIu, Uncut Alfalfa; I, Sugar Beets; III, Grain Sorghum, Wheat, and Alfalfa; V, Wheat Stubble and Weeds, Bare, Wheat.

ship held in adjacent areas, it would form a useful basis for extrapolation and discrimination. The next lowest region has 23 cases of corn, 1 of sugar beets, and 1 of grain sorghum. If the same relationship also holds in adjoining areas, it seems to be promising for projection. The lowermost group has 120 cases of bare ground; 17 of emergent wheat, in which the wheat occupies less than 5% of the total area (for all intents and purposes it is really bare ground); 10 of wheat stubble and weeds in which the wheat stubble was dry and made little contribution to the return; 6 of alfalfa; and 4 of grain sorghum. It will be seen that, with the exception of a few cases of alfalfa and grain sorghum, the lowermost category represents bare ground or, practically speaking, bare ground in the form of barely emergent wheat. Notice that in the next category above there is only one example of bare ground. Consequently, the lowermost group also looks reasonably encouraging for projection by extrapolation to adjacent areas.

The next to lowest category is a grab bag of emergent wheat, wheat stubble and weeds, alfalfa, corn, bare ground, and grain sorghum, though it is predominantly grain sorghum. It is obvious that the two polarization single-frequency radar achieves inadequate discrimination in this region. However, the uppermost portion of the region is dominated by grain sorghum, and a subset may be discriminated which could be used as a basis for prediction for grain sorghum in the lower portion of the mixture. Radar imagery at a different time optimized for the detection and discrimination of wheat (say, April) would aid in improving predictions in this area.

The scatter diagram of density measurements from color aerial photography (Fig. 10) is designed to simulate Fig. 9 as closely as possible. Since the photography took place a week later in September than the radar imagery, there will have been some cutting of corn, grain sorghum, and alfalfa and some emergence of wheat in what would earlier have been bare ground. A new category is added: cut alfalfa as distinguished from uncut alfalfa.

The linearity of the photography scatter diagram indicates that there is a high correlation between the two photographic layers (red- and green-sensitive). A comparable linearity exists between the blue and red layers (not reproduced here), and no greater discrimination was found between crops.

Table VIII
Radar versus Photographic Discrimination of Similar Categories

Category Indicated crop type	Percent of Category of Indicated Crop Type		Percent of Indicated Crop Type within Category	
	Radar	Photo	Radar	Photo
Sugar beets	97	67	97	90
Grain sorghum	54	42	95	90
Fallow and bare	83	67	83	67

Table VIII, prepared by W. G. Brooner, summarizes the comparison between radar and photographic system crop discrimination. Crop regions are more clearly delimited on the radar plot than on the photo plot, both in terms of within-category crop homogeneity and between-category crop differences.

In finishing this comparison, it should be noted that no test has yet been run of multispectral, multipolarization radar systems outside K-band, although we currently have under study X-band multiple polarization radar data at Garden City. For crops with a large physical structure such as corn, wavelengths greater than 3 cm (X-band) may also be useful, even out to wavelengths of 10 to 20 cm. Even longer wave-lengths may prove valuable in discrimination of soil moisture states or between tree densities in forest land. Clearly, much work remains to determine both the spread and number of frequencies which may effectively be used.

Concluding Remarks

The following conclusions emerge from the preceding discussion.

1. In any remote sensing system concerned with the detection of and discrimination between crops, there is a severe fall-off in the ability of the system to provide the answers requested of it, when constraints of increasing severity are placed on the sensor, the timeliness of data collection, the congruencing of data sets through time, or the spatial or data discrimination requirements of the system. Thus, rather than asking too much of a sensor system and suffering disappointment because it cannot provide the required information, it is appropriate to define for each remote sensing problem various alternative constraints, strategies, and loss functions so that rational choices may be made.

2. Though only a small amount of data were available, we have clearly shown that identification accuracies in the ninetieth percentile can be attained by using the methods illustrated. Further work needs to be done on clustering techniques which can analyze data in as detailed a fashion as the multiple-linkage clustering procedure, yet with the ability of the conditional probability clustering method to handle large amounts of data quickly. No attempt has been made here to associate crop type with the spatial texture distribution it produces on the image, yet this might well allow for identification accuracies approaching 99%.

3. Sequential radar flights throughout a number of growing seasons are required to find the nature of the time dependencies and internal consistencies in the data for different crops, regions, and years. It may be that year-to-year variations will be greater in the visible and infrared regions than with radar; hence the latter may prove more valuable in crop discrimination studies than we would expect at first sight. We believe that this hunch should be carefully evaluated in future experiments.

4. A great deal of work will be required to define the radar frequencies, polarizations, and bandwidths most suitable for crop discrimination studies.

Acknowledgments

We wish to acknowledge helpful criticisms by Dr. R. K. Moore. The bulk of the section comparing radar and photography at Garden City was drawn from an unpublished memorandum by Mr. W. G. Brooner.

Financial support for this study was provided by the following contracts: Project Themis (DOD), DAAK02-68-C-0089; USGS Contract 14-08-0001-12077; and NASA Contract NAS 9-7175. The support of the University of Kansas Computation Center is also gratefully acknowledged.

References

Haralick, R. M., and G. L. Kelly (1969), Pattern recognition with measurement space and spatial clustering for multiple images, *Proc. IEEE* 57 (4) 654–665.

Hoffer, R. M. (1967), Interpretation of remote multispectral imagery of agricultural crops, Res. Bull. 831, Laboratory for Agricultural Remote Sensing, Vol. 1, Purdue University, Agricultural Experiment Station, Lafayette, Ind.

Kendall, M. G., and A. Stuart (1966), *The Advanced Theory of Statistics*, Vol. 3, Hafner Publishing Co., London and New York.

Meyer, W. D. (1967), Analysis of radar calibration data, Final Report, Goodyear Aerospace Corporation, Litchfield Park, Ariz. 1–98.

Middleton, D. (1960), *Statistical Communication Theory*. McGraw-Hill Book Co., New York.

Schwarz, D. E., and F. Caspall (1968), The use of radar in the discrimination and identification of agricultural land use, *The Fifth International Symposium on Remote Sensing of Environment*. University of Michigan, Ann Arbor, Mich., pp. 233–248.

Schepis, E. L. (1968), Time lapse remote sensing in agriculture, *Photogrammetric Eng.* 34 (11), 1166–1179.

Simonett, D. S., J. E. Eagleman, A. B. Erhart, D. C. Rhodes, and D. E. Schwarz (1967), The potential of radar as a remote sensor in agriculture: a study with K-band imagery in western Kansas, CRES Report 61-21, Center for Research in Engineering Science, University of Kansas, Lawrence, Kansas.

Sokal, R. R., and P. H. A. Sneath (1963), *Principles of Numerical Taxonomy*. W. H. Freeman, London.

Wilks, S. S. (1962), *Mathematical Statistics*. John Wiley and Sons, New York.

Received July 8, 1969
Revised version received December 30, 1969