# Rapid Maximum Likelihood Classification

*Paul V. Bolstad** and *T. M. Lillesand*
Environmental Remote Sensing Center, 1225 West Dayton Street, 12th Floor, University of Wisconsin-Madison, Madison, WI 53706

ABSTRACT: We describe an improved table look-up technique for performing rapid maximum likelihood classification on large images. The method provided a more than 20-fold reduction in classification time relative to standard algorithms in a three-band classification of a full Landsat Thematic Mapper (TM) scene. While powerful, the improved algorithm is also simple, portable, and can run in limited memory desktop computer environments. The described algorithm significantly improves the practicality of large area land-cover classifications, such as those required for statewide and regional analyses.

## INTRODUCTION

A VARIETY OF AUTOMATED METHODS have been developed over the past two decades in order to classify multispectral satellite image data into distinct feature categories (e.g., Swain and Davis, 1978; Richards, 1986). These methods convert multispectral image data into information about the location and extent of earth surface features, and they have been applied extensively in such diverse fields as geology, hydrology, water quality assessment, the agricultural sciences, forestry, and land-use management (e.g., Jensen and Toll, 1982; Ryerson *et al.*, 1985; Franklin, 1986; Irons and Kennard, 1986; Lathrop and Lillesand, 1986; Hopkins *et al.*, 1988).

Supervised classification methods are those which require significant pre-classification input by the image analyst, and the maximum likelihood decision rule is by far the most common supervised classification method used for analyzing satellite image data (Richards, 1986). The standard implementation of supervised maximum likelihood classification requires the selection of "training" samples representing the feature types to be mapped (Lillesand and Kiefer, 1987). The sampled data normally include several "spectral classes" to adequately represent each "information class" or feature type. These training data are used to estimate the parameters of a probability density function for each spectral class. Generally, a multivariate normal probability model is chosen, with the training sample mean vectors and variance/covariance matrices defining the probability density functions. These density functions are then used to calculate the likelihoods of spectral class membership for each pixel in the image. The class with the highest computed likelihood is assigned to the output classified image.

The popularity of the maximum likelihood classifier is due to a number of characteristics (Swain and Davis, 1978; Schowengerdt, 1983; Richards, 1986). First, the maximum likelihood decision rule is intuitively appealing because the most "likely" outcome among candidate outcomes in chosen. Second, the decision rule has a well-developed theoretical foundation, and, for normally distributed data, is mathematically tractable and by many measures statistically desirable. Third, a maximum likelihood classification can readily accommodate covarying data, a common occurrence with satellite image data. Finally, maximum likelihood classifiers have been proven to perform well over a range of cover types, conditions, and satellite systems (Swain and Davis, 1978; Richards, 1986; Lillesand and Kiefer, 1987).

Despite the advantages of maximum likelihood classifiers, most implementations have exhibited at least one serious drawback,

namely, long classification times. For $N$ spectral bands and $T$ training sets, computing the maximum likelihood for each $N$-tuple (pixel measurement vector) of image data requires at least $(N^2 + N) * T$ multiplications and $(N-1)*(N+1) + 2*N$ additions in the most commonly implemented form of the maximum likelihood decision rule (Richards, 1986). Accordingly, per-pixel maximum likelihood classification requires billions of calculations when applied to large-area high-resolution satellite image data, such as those provided by the Landsat Thematic Mapper (TM) and SPOT High Resolution Visible (HRV) multispectral scanners. Full-scene classification times have been limited to date by computational speed, even when standard numeric coprocessors are used. Run-times of from 5 to 70 days can be expected for classifying a full-scene Landsat TM image with current pipeline-architecture desktop computers (Westman, 1989).

There are several potential means by which classification speed can be improved. First, increased processor clock speed will improve numeric computations approximately proportionally. Clock speeds have been doubling every three to four years in many computing environments and should continue to do so for at least the next several years. However, clock speed improvements foreseeable in the next decade will still result in impractical full-scene or multiple scene classification times in the desktop environment.

Enhanced numeric coprocessors are a second potential source of classification speed improvement. Commercially available specialized coprocessors can yield two- to four-fold increases in classification speeds when compared to the use of "standard" coprocessors. Unfortunately, enhanced numeric processors are often expensive, and can require special languages or algorithms, thereby reducing code portability.

A third option for improving classification run-times involves the adoption of parallel processing technology. System bottlenecks inherent in pipeline architecture are avoided by performing the numeric calculations in parallel, and very high throughput can be achieved with array processors (Westman, 1989). However, as with enhanced numeric coprocessors, array processors can be very expensive, and usually require hardware specific code.

As an alternative to the above "hardware" approaches to improving classification speed, this paper reports on a simple, portable table look-up algorithm which dramatically reduces processing times for large scenes while retaining the advantages of the maximum likelihood classifier. As described below, this algorithm derives its speed by avoiding most of the redundant computation which characterizes many current implementations of the maximum likelihood classifier. The procedure also circumvents many of the disadvantages of previous table look-up methods.

---

*Presently with the Department of Forestry, 324 Cheatham Hall, Virginia Polytechnic Institute and State University, Blacksburg, VA 24061

---

## BASIC CONCEPTS

Most current maximum likelihood classifiers calculate relative class membership "likelihoods" with respect to all training sets for each pixel in an image. They then assign the resultant most likely class identity to the output image and discard the intermediate computational results. Thus, an identical set of multiplications and additions will be computed each time a frequently occurring $N$-tuple is present in the image. This method is inefficient when there are several instances of the same $N$-tuple, because of the redundant calculations involved. There are at most $L^N$ unique $N$-tuples with $N$ bands measure over $L$ possible gray levels, and in practice the number of distinct $N$-tuples is much smaller (Shlien, 1975). Thus, when there are many more pixels than distinct $N$-tuples to classify, it becomes quicker to calculate the likelihood for the $N$-tuples once, save them in a suitably indexed look-up table, and simply recall the result the next time each $N$-tuple is encountered. Most current implementations of the maximum likelihood algorithm do not utilize this approach, apparently because the potential data space occupied by the set of possible $N$-tuple combinations is considerably larger than the data space available in most computing environments. For example, both SPOT HRV and Landsat TM data are quantized in eight bits, affording 256 unique data values per band. This allows a total of $256^3$, or 16,777,216 unique $N$-tuples with SPOT multispectral data, and $256^6$ (discounting the thermal band), or approximately $2.8 * 10^{14}$ unique $N$-tuples with Landsat TM data. This is considerably larger than memory available with current hardware. In spite of these large potential data spaces, two observations suggest that a look-up table approach may still be used to improve classification speeds in the context of computer memory limitations.

First, current high resolution satellite data sets contain three or fewer fundamental dimensions of variability for many applications, and much of the information in the image data can be captured by using three (or fewer) image bands (Goodenough et al., 1978; Nelson et al., 1984; Horler and Ahern, 1986; Lathrop and Lillesand, 1987). For example, SPOT-1 HRV data are only available in three bands, and a majority of the variability in the data is often contained in just two principal dimensions (Blohm, 1989). For Landsat TM data, a large part of the variability usually occurs in just three principal dimensions. For many vegetative cover types, much of the variability is represented by one band from each of the visible, near-infrared (NIR), and mid-infrared (MIR) portions of the spectrum.

A second factor suggesting the potential utility of the table look-up method involves the actual versus the potential frequency distributions characteristic of most satellite image data. Although $256^3$ unique combinations are possible with three-band eight-bit images, most of the data are usually concentrated in a small region of the entire potential data space (Shlien, 1975; Malila, 1985; Metzler and Malila, 1985; Murphy et al., 1985). As the observation of typical single-band image histograms indicates, over 90 percent of the data for 1-tuples are often contained in a relatively limited subset of digital values. Similar conditions also commonly exist for higher order $N$-tuples. This means that regions of data concentration can be effectively identified, and the look-up table can be limited to only the frequently occurring $N$-tuples. In this way, the size of the look-up table can be greatly reduced, yet a large portion of the image data can be represented by the look-up table. Hence, during image classification the look-up table can be used only for those $N$-tuples which occur most frequently. Likelihoods for only the relatively infrequent $N$-tuples not in the look-up table can be computed separately each time they occur.

The advantages of a table look-up approach have been previously recognized and demonstrated for Landsat Multispectral Scanner (MSS) data (Eppler, 1974; Shlien, 1975; Mather, 1985).

However, these approaches have not been widely adopted, particularly for Landsat TM and SPOT HRV data. The algorithm described herein is distinct from the previous approaches in at least three respects: data quantitization independence, improved portability, and run-time dimensioning to improve resource utilization and increase efficiency.

Previous table look-up methods have either lacked generality, required specialized, non-portable code to implement (Shlien, 1975; Mather, 1985), or depended on the order of training data input (Eppler, 1974; Jones, 1974). The approach adopted by Shlien (1975) and Mather (1985) involves storing calculated maximum likelihoods in a hash table for later access. In both cases, a fixed-size hash table was adopted. A prime division hash function (Knuth, 1973) and machine-dependent bit shuffling were implemented to calculate table addresses, while collisions were resolved by repetitive probing. This approach has low portability because rapid hash table addressing results mainly from machine-dependent subroutines. Further, pixel concatenation for address computation depends on machine integer dimension, further reducing portability to other computing platforms. The algorithm was developed for six-bit representations of MSS data, so applicability to eight-bit data from other sensors, such as SPOT HRV and Landsat TM sensors, is uncertain. Finally, the fixed, 10,000 element hash table may be too small for eight-bit TM and SPOT data, resulting in frequent collisions and rehashing, or may limit representation to a small percentage of the image data, reducing efficiency in either case. The algorithm described herein does not use a hash function for table address calculation; rather, it accesses table elements directly. Direct accessing is possible because the look-up table is restricted only to the "data rich" portion of the spectral domain.

The adopted algorithm identifies regions of high data density by using histograms of individual spectral bands to select digital numbers for inclusion in the look-up table. These histograms correspond to the marginal frequency distributions of the $N$-dimensional data. This method requires the compilation of an image histogram for each band prior to classification, but this can be accomplished at minor computational expense. Sampling theory suggests accurate histograms can be compiled quickly for very large images by using a fixed sample size and subsampling systematically along rows and columns of the image data (Williams, 1978). Although single-band histograms do not guarantee an optimal selection of digital numbers for inclusion in the look-up table, they offer considerable computational time savings relative to $N$-way histogram compilation while providing suitable selections.

Finally, another factor influencing the design and development of our algorithm is the fact that many current maximum likelihood classifiers compute the likelihoods for all possible spectral classes, irrespective of their statistical distance from the current $N$-tuple to be classified. Many $N$-tuples, particularly those very near a training set mean, have very high likelihoods for one spectral class and relatively low likelihoods for the remainder of the spectral classes. Accordingly, less computationally intensive methods (e.g., parallelepiped techniques) can be used to first classify these "highly certain" $N$-tuples, and maximum likelihood classification can be used only for the remaining observations which fall further from class means. Provided this approach leads to class assignment identical to the maximum likelihood method, this stratified classification procedure can result in significant time savings relative to traditional maximum likelihood implementations with no degradation in classification accuracy.

Thus, the success of our improved algorithm depends on (1) how well the single-band histograms identify regions of high data density, (2) the size of the resultant look-up table, and (3)

the accurate specification and classification of "near" spectral classes.

## ALGORITHM DESCRIPTION

The basic concepts described above have been developed and tested in a stratified, three-band table look-up maximum likelihood classifier (Figure 1). The adopted approach calculates the marginal (one-dimensional) histograms of the three-dimensional data space. A large number of pixels (greater than 10,00) are sampled systematically in a uniformly spaced grid for each band in each image, and the percent frequency of occurrence is calculated for each of the 256 (0 to 255) possible digital numbers. Each frequency histogram is held in a separate vector, used initially to calculate histograms and subsequently to index the look-up table. The number and location of digital counts above an analyst-specified minimum frequency threshold are flagged for each band (Figures 2 to 4). Thus, only that portion of the spectral domain which satisfies the marginal frequency threshold criteria for all bands is included in the look-up table (Figure 5). The number of digital counts selected for the three bands, plus one additional count for each band (described below), define the three dimensions of the look-up table. If sufficient memory space is available, the look-up table is then allocated dynamically; if space is not available, the threshold is successively increased and a smaller set of qualifying digital numbers is reselected for each band until an array of the maximum manageable size is successfully allocated. Categories are then assigned according to the appropriate maximum likelihood value.

Run-time array dimensioning and allocation offer several advantages, chief among them being a balanced selection of digital numbers and full utilization of available memory resources. Selected histogram vector elements are then numbered sequentially to index corresponding elements in the three-dimensional look-up table array. Unselected elements are set to values pointing to the lower "faces" of the look-up table, with an index value set to a null value which does not occur or occurs infrequently in the input data (e.g., −1 or 0). This null value is the

"additional" count mentioned above (Figure 6). This selection and indexing of the histogram vectors reduces the size of the look-up table by eliminating sparsely populated spectral regions. This structuring also allows the image data to be used as nested indices during classification, in accordance with the following C syntax:

Feature Type ijk = Look up table [Vect1[i]] [Vect2[j]] [Vect3[k]] , where i, j , and k are digital data observed in the images; Vect1, Vect2 , and Vect3 are histogram vectors; Look up table is the three-dimensional table; and Feature Type ijk is the feature type with the maximum likelihood for the given N-tuple. Thus, once calculated, pixel N-tuples represented in the table can be assigned the appropriate class at the computational cost of accessing three one-dimensional arrays and one
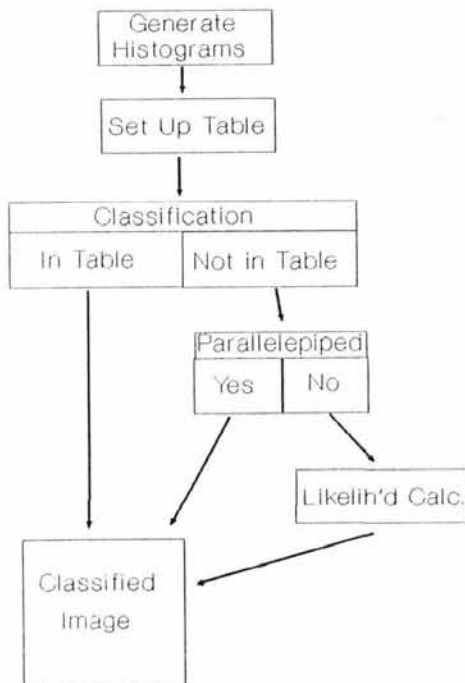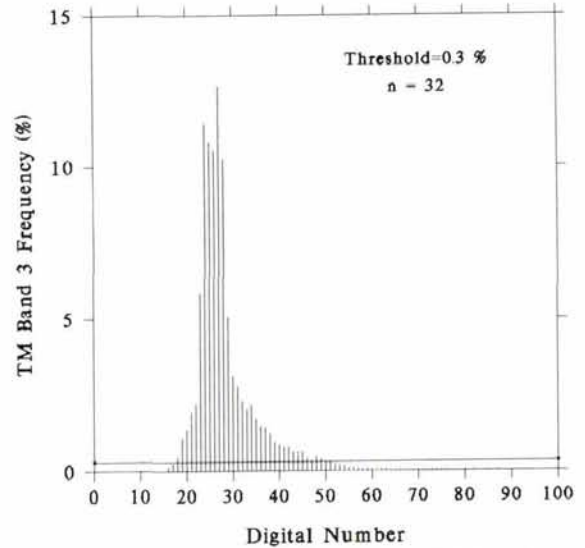


FIG. 2. Frequency histogram and number of digital numbers with frequency greater than 0.3 percent for band 3 of a Landsat Thematic Mapper image of a forested region in northern Wisconsin.
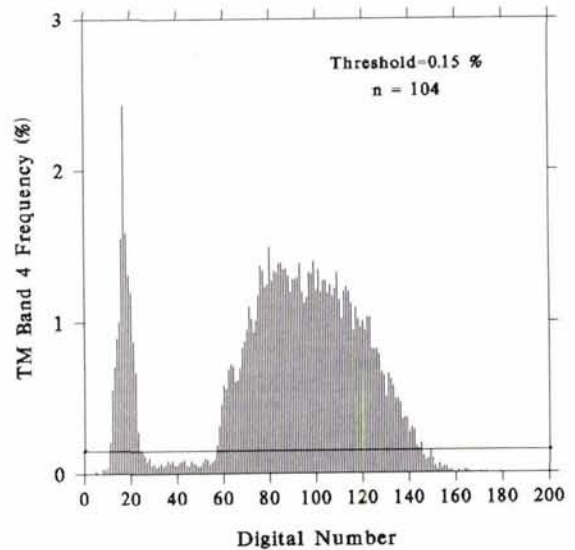


FIG. 3. Frequency histogram and number of digital numbers with frequency greater than 0.15 percent for band 4 of a Landsat Thematic Mapper image of a forested region in northern Wisconsin.
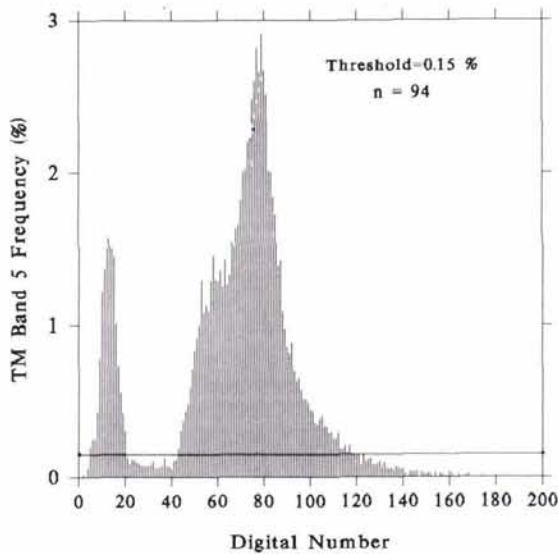


FIG. 1. Overview of the table look-up classifier.

FIG. 4. Frequency histogram and number of digital numbers with frequency greater than 0.15 percent for band 5 of a Landsat Thematic Mapper image of a forested region in northern Wisconsin.
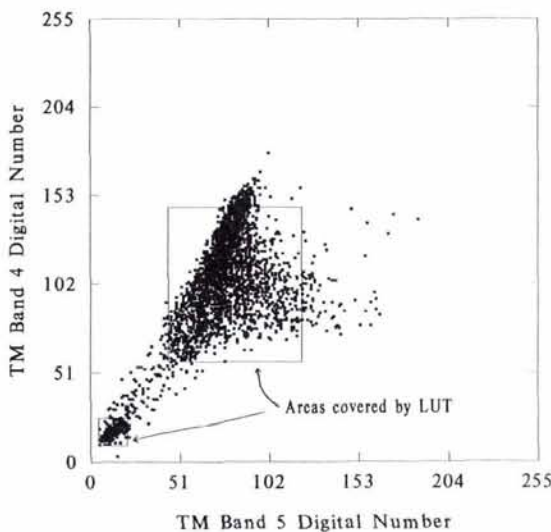


FIG. 5. A two-dimensional example of the portion of the total spectral domain covered by the look-up table. Those areas inside the rectangles satisfy a marginal frequency threshold criteria of 0.15 percent for both bands, illustrated in Figures 3 and 4. The feature type with the maximum likelihood is calculated for each distinct N-tuple within the selected regions and stored in the look-up table (LUT).
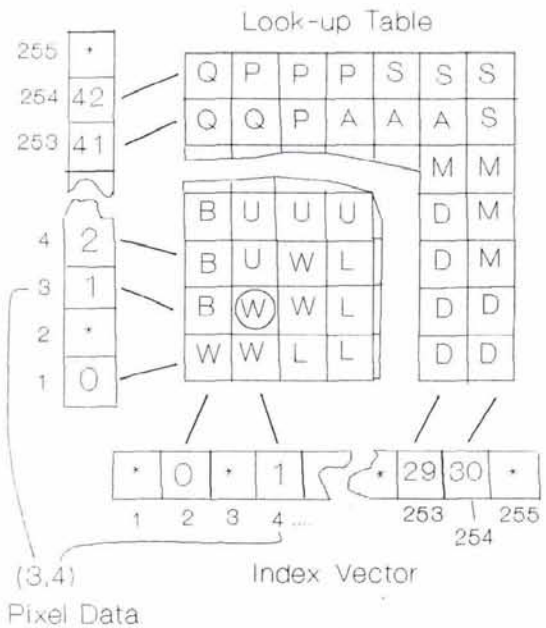


FIG. 6. A two-dimensional example of the sequential table look-up indexing scheme used. Image data (e.g., observation vector digital number combination 3,4) are use to index 256 element vectors, which in turn index the look-up table. Asterisks in the index vectors indicate unselected digital numbers. The look-up table contains the identification number of the most likely feature category, in this case, category "W."

three-dimensional array, a significant time savings when compared to repeating the likelihood calculations or to hash table address calculation.

The proposed method of look-up table access exhibits several advantages when compared to previously described schemes (Eppler, 1974; Shlien, 1975; Mather, 1985). First, the described method is more general, in that no assumptions about the number of selected N-tuples are required nor is complex training set structuring. Second, the indexing scheme is simple, resulting in code parsimony. Finally, the indexing method is inherently rapid when using standard higher level languages, and optimization by means of lower level programming is not required.

Thus, the algorithm is easily ported across a broad range of computing platforms.

As mentioned previously, those N-tuples falling outside the range of the look-up table could be assigned by means of the usual repetitive likelihood calculation method. However, this requires the calculation of all likelihoods for each pixel, including many distant classes with correspondingly low likelihoods. In many cases only one or a few likelihoods are viable candidates. There are several potential methods for screening candidates; we chose a parallelepiped procedure to identify the classes of high potential likelihood (Lillesand and Kiefer, 1987). Although not sensitive to data covariance structure, the parallelepiped classifier is a rapid, simple method for multi-dimensional distance discrimination. Only classes which uniquely fall within class-specific distances of the pixel to be classified are considered as candidates. That is, if the pixel is within a set distance (e.g., one standard deviation) of the mean in all spectral dimensions for one and only one class, then this class identity is assigned to the pixel. Although not a maximum likelihood classification, practice suggests that, under most circumstances, setting the distances "close" to the spectral mean values results in assignments for qualifying pixels which are identical to those of the maximum likelihood classification. If the distances are chosen conservatively, the resultant set of likelihood calculations will be greatly reduced, yet still result in proper class selection.

Finally, for N-tuples which fall outside the range of the look-up table and do not satisfy the parallelepiped criterion, the usual pixel-by-pixel maximum likelihood classification is performed. This three-phase approach to classification yields the combined efficiency and accuracy of the proposed algorithm.

## MATERIALS AND METHODS

The above algorithm was written at the University of Wisconsin-Madison Environmental Remote Sensing Center (ERSC)

and has been initally tested using Landsat TM data collected as part of ERSC's participation in the National Science Foundation supported Long-Term Ecological Research Program (Lillesand *et al.*, 1989). The classifier was written with the C programming language to operate on an Intel 80386 based microcomputer under the DOS operating system. Although code was developed on a microcomputer platform, we attempted to adhere to the ANSI C standard to ensure portability.

Initial testing of the algorithm was performed using images that were subset from a Landsat TM scene located in the forested region of northeastern Wisconsin (Path 25, Row 28, acquired 9 June 1988). Vegetation in this region is dominated by northern hardwood and boreal forests (Kotar *et al.*, 1988). A representative sub-image of approximately 500 by 600 pixels was selected from a cloud-free portion of the full TM scene. Training sets had been developed previously for a much larger portion of the image using standard supervised maximum likelihood classification techniques. A total of 79 spectral classes were used to represent 14 land-cover classes (Northern Hardwoods, Mixed Hardwood/Conifer, Red/White Pine, Jack Pine, White Pine/ Norway Spruce, Upland Brush, Lowland Brush, Lowland Conifer, Mixed Lowland Vegetation, Grass and Sedge, Herbaceous Vegetation, Aquatic Vegetation, Urban/Bare Soil, and Water). The scene was initially classified using the traditional supervised maximum likelihood pixel-by-pixel method and algorithm. Classification accuracies averaged 89 percent (Benson *et al.*, 1989).

A set of benchmark classifications using both the traditional and new approach were conducted on the sub-image, using three TM image bands: a visible (band 3), the near infrared (band 4), and a mid-infrared (band 5). Benchmark tests compared classifier characteristics and run-times at a number of histogram selection threshold levels. Measurements included time spent on input/output, histogram generation and count selection, look-up table elaboration (time spent calculating and assigning class maximum likelihoods to $N$-tuples represented in the table), table look-up during classification, and calculation of likelihoods during classification for $N$-tuples not covered by the table. Times for each respective activity were reduced to a per-pixel basis, and extrapolated using appropriate multipliers to cover a range of image sizes. Total classification times were compared to those of a similar traditional maximum likelihood classifier. This traditional classifier was developed by removing the histogram generation, look-up, and parallelepiped portions of the table look-up classifier; essentially the lower right-hand box in Figure 1. This traditional method employs the "standard" form of the maximum likelihood decision rule, wherein only the exponent of the likelihood function is used as a discriminant (Richards, 1986). All classifications were compared digitally, pixel by pixel, to the classified image derived using the traditional maximum likelihood classifier.

All benchmark classifications were performed on a 25mHz Intel 80386 computer (CompuAdd model 325) equipped with 1 mbyte of RAM and running the DOS 3.3 operating system. Approximately 585 kbytes of RAM were available for program use on the system after device drivers, DOS environment space, and other required system space had been allocated.

## RESULTS AND DISCUSSION

As shown in Figure 7, classification times were dramatically improved using the look-up table method on the three-band data set. Relative improvement over the traditional method was a function of image size, and increased with increasing image size. The relative advantage increased because there are two stages of the look-up table method which account for a major portion of run-times on large images. The first stage, table elaboration, depends on the size of the look-up table, and thus the number of digital number values resulting from the frequency
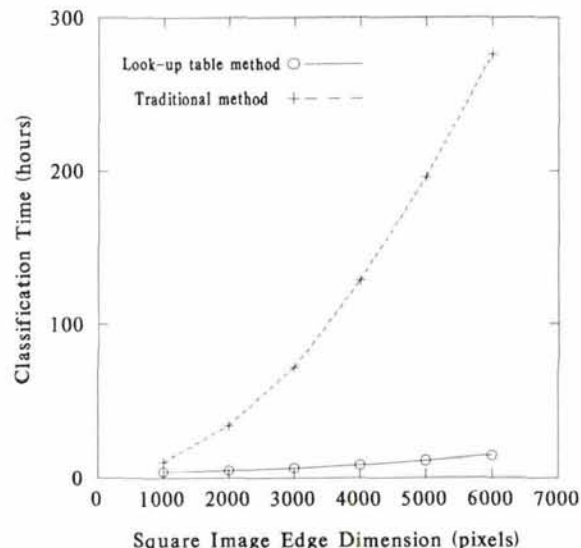


FIG. 7. Time required for classification of square three-band images using both the traditional and look-up table maximum likelihood classification algorithms. This and all benchmark classifications were performed on a 25 mHz 80386 system equipped with a math coprocessor. Seventy-nine spectral classes were used for 14 information classes, with an overall classification accuracy or 89 percent.

threshold. At the lowest threshold tested (0.1 percent), tables contained on the order of 470,000 elements, and likelihood calculations were performed for all elements. However, not all the 470,000 $N$-tuples occurred in the image, so some of the elements were computed unnecessarily. Time spent calculating these likelihoods can be considered as wasted, because they are never accessed. An alternate strategy was tested which calculated the table elements when they were initially encountered in the image, thereby avoiding likelihood calculations for $N$-tuples which do not occur in the image. This opportunistic table filling method was considerably faster than the adopted table look-up algorithm on small to mid-sized images, because table generation overhead was reduced. Also, this opportunistic approach was faster than a traditional maximum likelihood classification for all scene sizes tested larger than approximately 60 pixels on an edge. This compares to the adopted look-up table approach, which fills the table completely regardless of image size, and was not faster than the traditional maximum likelihood classifier until image edge dimension surpassed approximately 600 pixels. However, for full-scene Landsat TM images, the opportunistic approach proved slower than the adopted approach, due to the overhead incurred in checking the look-up table for prior assignment while classifying each image pixel.

Run-times for the image classification and input/output are the second major phase and are added to table generation overhead, resulting in total classification times. As the size of the image increases, the proportionate amount of time spent allocating and elaborating the look-up table compared to classifying the image decreases. For square images 3000 pixels on a side (approximately equivalent to full SPOT multi-spectral scenes or TM quarter scenes), total classification times using the look-up table method required 6.2 hours, approximately 9.5 times faster than traditional methods. For a TM full-scene size image, the look-up table method would require 14 hours, approximately 21 times faster than traditional methods. These dramatic improvements in classification times were initially verified for only one satellite image of forested lands in northeastern Wisconsin. Although there are many scene-dependent factors which affect

classification times, such as cover types, spectral variability, number of training sets, and time of year, this work provides an example of the substantial reduction in classification run-times possible with the method.

The look-up table classification times were a function of threshold level as well as image size (Figure 8). For small images, classification times were relatively constant across a range of selection thresholds. This is because time saved by reducing the size of the look-up table was offset by fewer $N$-tuples being contained in the table, and hence more time spent calculating likelihoods during image classification. Lower thresholds provide shorter times for larger images. This trend was true down to thresholds of 0.1 percent, the smallest tested due to memory limitation imposed by the DOS operating system. At this threshold, classifying a full scene TM would require 14 hours on an inexpensive, DOS-based desktop computer. Although the classification times are likely to be a function of the scene-specific characteristics noted above, the look-up table algorithm provides a substantial savings in classification time over a broad range of histogram thresholds and image sizes.

Extrapolating the TM benchmark classification results to SPOT multispectral full scene data (3000 by 3000 pixels) indicates classification times of approximately 6 hours. However, this is probably an over-estimate, for two reasons. First, the benchmark tests were performed using Landsat TM visible (band 3), NIR (band 4), and MIR (band 5) images. The data ranges and hence look-up table dimensions for the infrared bands were much greater than for the visible bands of the study region, both on SPOT multispectral and Landsat TM imagery. Because SPOT multispectral data contain two visible and only one IR band, an optimal portion of the spectral data space could be included in the look-up table, thereby decreasing classification times. Second, the 6-hour estimate is probably conservative because of the steeper gain and hence smaller ranges for SPOT multispectral bands when compared with corresponding TM bands (Blohm, 1989). Histogram ranges for all three SPOT multispectral bands have been observed to be narrower than the corresponding bands of concurrent TM imagery, both for the described

study area and for a similar region in northwestern Wisconsin. Preliminary tests classifying approximately one-half of a SPOT HRV image indicate full-scene classification times of approximately 4 hours when using 71 spectral classes.

Both look-up table dimension and the percent of image data $N$-tuples included in the look-up table increased with decreasing frequency thresholds (Figure 9). Look-up table size ranged from 14 to 474,000 elements over a range of thresholds from 1.5 percent to 0.1 percent, while the percent of observed image $N$-tuples found in the table increased from 1 to 93 percent over the same range. Thus 93 percent of the image area was included in a look-up table comprising 2.8 percent of the potential spectral domain (474,000 / 16,770,000) at the lowest histogram threshold level tested (0.1 percent).

The 2.8 percent of the spectral domain represented in the look-up table was sparsely populated in that only 48,247 of the table $N$-tuples were observed in the test images. Hence, 10.2 percent of the look-up table elements were used during image classification. More complex indexing schemes could be used to decrease the storage requirements of the look-up table, e.g., hashing schemes, complex linked lists, binary, quad, or higher order trees, or some combination of these and indexed arrays. However, these schemes, while reducing storage requirements, are more difficult to code; further, guidelines need to be established to identify a judicious mix of smaller, slower access structures (e.g., portable has functions) and faster, larger structures (look-up tables). These methods are topics for further investigation and show promise for data storage reduction without sacrificing performance.

During the image classification stage, approximately 5 percent of run-time was spent classifying the 93 percent of the image covered by the look-up table, while the remaining 95 percent of classification stage run-time was spent assigning classes to $N$-tuples not represented in the table. Of those $N$-tuples outside the table, approximately 30 percent could be assigned using the parallelepiped classification rule (2.1 percent of the total), while 70 percent were assigned using the traditional classification method (4.9 percent). Thus, with a frequency threshold set at 0.1%, a total of 98 percent of the image pixels were assigned using the true maximum likelihood algorithm, while 2 percent were assigned using the parallelepiped algorithm. Because a majority of the full scene classification times are spent on classifying the "outlying" $N$-tuples not covered by the look-up table, time savings by the parallelepiped classification rule are
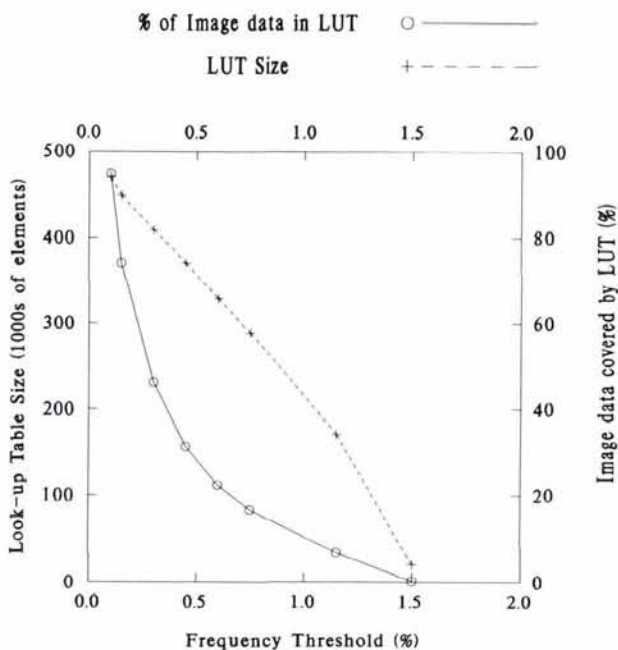


FIG. 8. The effect of image size and frequency threshold level on classification times when using the look-up table algorithm in a maximum likelihood classification.
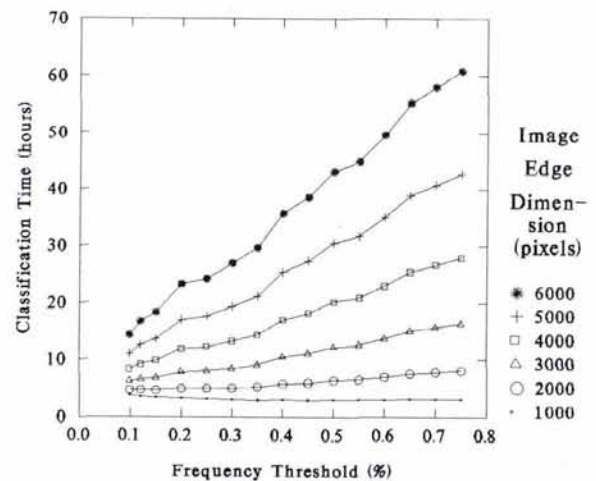


FIG. 9. The relationships among look-up table (LUT) size, percent of image data covered by the table, and the frequency threshold chosen.

still significant. Tests without the parallelepiped sub-section indicate TM full-scene classification times were reduced approximately 28 percent by the inclusion of the parallelepiped screening.

Image differencing results show the table look-up classification to be identical to the traditional method. While those *N*-tuples represented in the table will result in equivalent class assignment, the parallelepiped part of the algorithm could result in class assignment different from that of a maximum likelihood classifier for the small percentage of image data not contained in the look-up table. The small one standard deviation bounds chosen for the parallelepiped classifier, when coupled with the restriction of selecting one and only one class, resulted in identical class assignment as the maximum likelihood algorithm under the conditions of this study. This complete agreement between the parallelepiped and maximum likelihood classifications, while not surprising in light of the small bounds chosen, is not guaranteed. Clearly, agreement depends on the nature of the image data, the training sets chosen, and the size of the parallelepipeds. However, this study indicates that, by choosing conservative bounds, high agreement should result. In short, the look-up table classifier covered a large percentage of the image data and significantly reduced classification times by avoiding redundant calculations. For most of the remaining data, the parallelepiped algorithm improved classification times without sacrificing accuracy.

Because of the positive relationship between look-up table size, table elaboration time, and percent of image data covered by the table, the percentage of total run-time for table generation decrease with increasing threshold (Figure 10). Approximately equal parts were spent on look-up table creation and image classification for a 0.1 percent threshold. Although large-image classification times continued to decrease with decreasing frequency threshold down to the smallest threshold level tested (0.1 percent), extrapolation indicates convergence at some threshold/table-size combination. The optimum threshold under the conditions of this study is below 0.1 percent, but could not be identified because of memory limitations imposed by the DOS operating system. Below this minimum threshold, classification times would actually increase with decreasing threshold level. In general, the minimum threshold level is a function of many variables, including scene spectral diversity, image characteristics (e.g., clarity, time of year), scene size (Figure 8), number of training sets, the number of image bands used, and the relative times of numeric computations and look-up table access. Nonetheless, substantial improvements in classification times can be realized over a broad range of thresholds.

The described algorithm was developed and tested for classification of three-band images, but there are no threoretical reasons the algorithm could not be applied to images having larger numbers of bands. However, there are practical constraints imposed by available RAM memory on the number of input image bands. Approximately 25 mbytes of RAM would be required when employing four bands under the conditions of this study (two TM visible and two TM infrared bands). This is near the limit of what is generally allowable or affordable with current desktop computer technology. Table size with five or more TM bands and low histogram selection criteria would likely be larger than available memory. Table sizes for five and higher band classifications could be accommodated by raising the frequency selection threshold, although this would also tend to reduce performance because a reduced percentage of the spectral domain would be represented in the look-up table. As mentioned earlier, much image data are provided with only three image bands (e.g., SPOT multispectral data), and for many regions of the world most of the variation in TM data is contained in three principal dimensions. If the additional discrimination from more than four bands is required, the information could be maintained utilizing a dimensionality reduction technique such as a principal components transformation. Thus, information from five or more bands could be incorporated by raising the frequency selection thresholds, a feature reduction transformation, or the application of a hybrid indexing scheme, e.g., through a combination of table look-up and a portable hashing scheme.

## CONCLUSIONS

A stratified table look-up algorithm can significantly reduce classification times for a maximum likelihood classification with three satellite image bands. With this method full-scene Landsat TM classification can be accomplished overnight using *presently* available, inexpensive, modestly equipped desktop computer systems. Current prices of the computer system used for the described work total less than $5,000. Current moderately priced systems, available for less than $15,000, exhibit an 8- to 12-fold increase in throughput under scientific/engineering workloads when compared to the system used in the present work (Varhol, 1989). This suggests TM full-scene, three-band classifications of approximately 3 to 4 hours on such systems.

Table look-up classification performance is a function of the frequency selection threshold chosen; however, improvements in classification speed are large over a range of thresholds. Time savings increase with increasing image size. The table look-up algorithm significantly enhances the practicality of land-cover classification over large areas, particularly for state-wide and regional analysis where multiple satellite images are to be classified.

## ACKNOWLEDGMENTS

Legend:
- Likelihood calculation time during classification
- Look-up table creation and access time
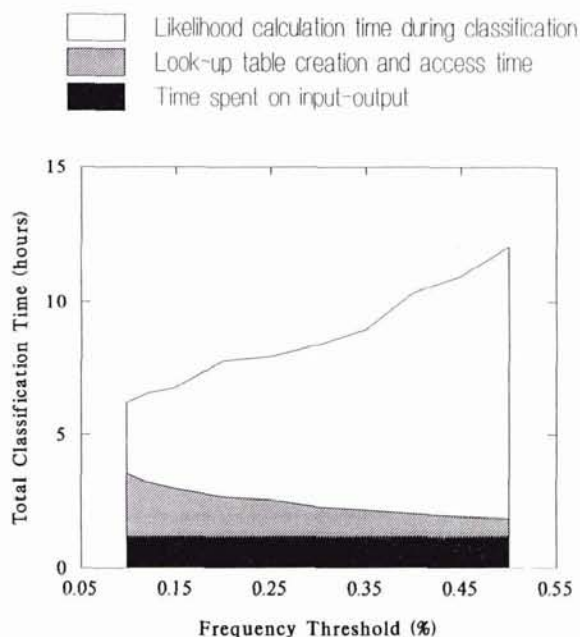- Time spent on input-output

FIG. 10. The effect of frequency threshold on the relative proportion of time spent on input/output, look-up table generation, and likelihood calculations for pixels which fall outside the look-up table during classification. A square image 3000 pixels on an edge was used with 79 spectral classes.

## REFERENCES

Benson, J. L., P. V. Bolstad, L. M. Christel-Rose, P. E. Gessler, M. A. Mohamed, L. A. Morrison, N. E. Podger, S. Ribanszky, A. R. Selvan, J. A. Westman, and L. R. Whiting, 1989. *The Development and Demonstration of a Combined Remote Sensing/Geographic Information System for the North Temperate Lakes Long-Term Ecological Research Site.* Unpublished Report, Environmental Remote Sensing Center, University of Wisconsin-Madison, 181 p.

Blohm, J. D., 1989. *A Comparison of SPOT and Thematic Mapper Satellite Data for Forestry Applications,* Unpublished Master of Science Thesis, University of Wisconsin-Madison, 297 p.

Eppler, W. G., 1974. An improved version of the table look-up algorithm for pattern recogntion, *Proc. 9th International Symposium on Remote Sensing of Environment,* pp. 793–812.

Franklin, J., 1986. Thematic Mapper analysis of coniferous forest structure and composition, *International Journal of Remote sensing,* 7(10) pp. 1287–1301.

Goodenough, D. G., P. M. Narendra, and K. O'Neill, 1978. Feature subset selection in remote sensing, *Canadian Journal of Remote Sensing,* 4, pp. 143–149.

Hopkins, P. J., A. L. Maclean, and T. M. Lillesand, 1988. Assessment of Thematic Mapper imagery for forestry applications under Lakes States conditions, *Photogrammetric Engineering & Remote Sensing,* 54(1), pp. 61–70.

Horler, D. N. H., and F. J. Ahern, 1986. Forestry information content of Thematic Mapper data, *Int J. Remote Sensing,* 7(3), pp. 405–428.

Irons, J. R., and R. L., Kennard, 1986. The utility of Thematic Mapper senor characteristics for surface mine monitoring, *Photogrammetric Engineering & Remote Sensing,* 52(3), pp. 389–396.

Jensen, J. R., and D. L. Toll, 1982. Detecting residential land-use development at the urban fringe, *Photogrammetric Engineering & Remote Sensing,* 48(4), pp. 629–643.

Jones, C., 1974. Implementation for an advanced table look-up classifier for large area land-use classification, *Proc. 9th International Symposium on Remote Sensing of Environment,* pp. 813–824.

Knuth, D. E., 1973. *The Art of Computer Programming: Sorting and Searching.* Addison-Wesley, Reading, Massachusetts, 723 p.

Kotar, J., J. A. Kovatch, and C. T. Locey, 1988. *A Field Guide to the Forest Habitat Types of Northern Wisconsin,* Department of Forestry, University of Wisconsin-Madison, not paginated.

Lathrop, R. G., Jr., and T. M. Lillesand, 1986. Use of Thematic Mapper data to assess water quality in Green Bay and Central Lake Mich-

igan, *Photogrammetric Engineering & Remote Sensing,* 52(5), pp. 671–680.

———, 1987. An evaluation of Thematic Mapper data for forest cover mapping in northern Wisconsin, *Proceedings: 11th Pecora Remote Sensing Symposium,* pp. 386–395.

Lillesand, T. M., and R. W. Kiefer, 1987. *Remote Sensing and Image Interpretation,* 2nd Ed., John Wiley and Sons, New York, 721 p.

Lillesand, T. M., M. D. MacKenzie, J. R. Vande Castle, and J. J. Magnuson, 1989. Incorporating remote sensing and GIS technology in long-term and large-scale ecological research, *Proceedings, GIS/LIS '89,* pp. 228–232.

Malila, W. A., 1985. Comparison of the information content of Landsat TM and MSS data, *Photogrammetric Engineering & Remote Sensing,* 51(9), pp. 1449–1457.

Mather, P. M., 1985. A computationally-efficient maximum-likelihood classifier employing prior probabilities for remotely-sensed data, *International Journal of Remote Sensing,* 6(2), pp. 369–376.

Metzler, M. D., and W. A. Malila, 1985. Characterization and comparison of Landsat-4 and Landsat-5 Thematic Mapper data, *Photogrammetric Engineering & Remote Sensing,* 51(9), pp. 1315–1330.

Murphy, J. M., F. J. Ahern, P. F. Duff, and A. J. Fitzgerald, 1985. Assessment of radiometric accuracy of Landsat-4 and Landsat-5 Thematic Mapper data products from canadian production systems, *Photogrammetric Engineering & Remote Sensing,* 51(9), pp. 1359–1371.

Nelson, R. F., R. S. Latty, and G. Mott. 1984. Classifying northern forests with Thematic Mapper Simulator data, *Photogrammetric Engineering & Remote Sensing,* 50(5), pp. 607–617.

Richards, J. A. 1986. *Remote Sensing Digital Image Analysis,* Springer-Verlag, 292 p.

Ryerson, R. A., R. N. Dobbins, and C. Thibault, 1985. Timely crop area estimates from Landsat, *Photogrammetric Engineering & Remote Sensing,* 51(11), pp. 1735–1743.

Schowengerdt, R. A. 1983. *Techniques for Image Processing and Classification in Remote Sensing,* Academic Press, 249 p.

Shlien, S. 1975. *Practical Aspects Related to Automatic Classification of Landsat Imagery Using Look-up Tables.* Canadian Centre for Remote Sensing, Report 75–2, Ottowa, Canada.

Swain, P. H., and S. Davis. 1978. *Remote Sensing: A Quantitative Approach,* McGraw-Hill Book Co., 396 p.

Varhol, P. D., 1989. Great performers, *MIPS,* 2(1), pp. 56–57.

Westman, J. A., 1989. *Applications of Cray X-MP Supercomputer to Digital Image Classification.* Unpublished Master of Science Thesis, University of Wisconsin-Madison, 262 p.

Williams, B., 1978. *A Sampler on Sampling,* John Wiley and Sons, New York, 254 p.

---

### 43rd Photogrammetric Week
### Stuttgart, 9-14 September 1991

This internationally-recognized "vacation course in photogrammetry" has been held at Stuttgart University since 1973. Because Professor Dr.-Ing. Friedrich Ackermann, one of those responsible for the scientific program, is to retire soon, this 43rd Photogrammetric Week will be his farewell seminar. Essential lines of his work have been chosen as the main topics for the meeting:

• **GPS for Photogrammetry** • **Digital Photogrammetric Image Processing** • **Photogrammetry and Geo-Information Systems** •

Lectures and discussions will be held in the mornings. Technical interpreters will be available for simultaneous translations into German or English. Demonstrations are scheduled for the afternoons.

For further information and applications, contact: Universitat Stuttgart, Institut fur Photogrammetrie, Keplerstrasse 11, D-7000 Stuttgart 1, FRG, telephone 0711/121-3386 or FAX 0711/121-3500.