Hough Transform Extraction of Cartographic Calibration Marks from Aerial Photography

Abstract

Cartographic compilation requires precision mensuration. The calibration of mensuration processes is based on specific fiducials. External fiducials, around the exterior frame of the image, must be precisely measured to establish the overall sensor geometry. Premarked ground points are provided within the image by placement of panels on the ground at locations whose position is precisely known. Both types of registration marks must be known within the pixel space of a digitized image in order for the feature extraction process to be able to be accurate with respect to delineated features. Classical mensuration of these targets requires that a photogrammetrist view the image on a display and use pointing devices, such as a mouse, to pick exact points. We have developed interactive tools which eliminate the precise pointing action for the operator. In this paper we will discuss the theory of the analysis used by the tool and present examples of operation.

Automated Cartography

Background

The volume of demand for maps and related cartographic products has begun to grow at all levels of both government and business. Governments have recognized that maps and related spatially organized databases are a direct method to depict and control the status of many situations. Business demand for maps has come in equal measure from responding to governmental requirements as well as organizing the spatial knowledge that affects a business. The volume of map demands has occurred, fortunately, at a time when the growth of computing power has emerged to ease the burden of handling large amounts of data. Consequently, there is a similar expansion of interest in the creation of digital mapping systems, systems which will integrate computer and database techniques for the generation and management of map data in a digital computer-compatible format.

The cartographic feature extraction process is characterized by the extraction of three distinctly different types of features: point features, line features, and area features. A point feature is one which enters the cartographic data as a single set of coordinate values along with any intrinsic attributes necessary to describe the feature. The fiducial marks employed in the mensuration of computer images can be classified as point features. The fiducial marks are typically identified by a single set of coordinates and a cross. As applied to image mensuration, the goal is to combine the capabilities of the computer and photogrammetrist to efficiently and accurately identify, precisely locate, and delineate the fiducial marks. As applied to fiducials or marked ground points, the photogrammetrist can quickly and easily identify the area containing the points and "box" them. By employing the algorithms described in this paper, the software, knowing the properties of the feature of interest and being provided with a restricted search region, can precisely locate the coordinates of the fiducial or the marked ground point. The algorithm discussed herein was developed and operated in mono photography, but it is completely applicable to stereo imagery, with the obvious modifications. Because stereo feature extraction is becoming dominant, the algorithm can be applied to stereo extraction as required.

The objective of an interactive tool is to use the identification capabilities of the human operator, capabilities not yet duplicated by methods such as computer vision and artificial intelligence, and the processing abilities of the computer to quickly, accurately, and consistently identify the fiducial and ground points. This tool relieves the operator of boring and time consuming activities, allowing him to devote his attention to tasks that exceed computer capabilities.

The GHOUGH (Fischler, 1986; Fischler, 1989) algorithm is a complex and computationally intensive algorithm that primarily addresses the detection of objects whereas the Hough algorithm presented herein is much simpler and is focused towards location precision.

Hough Transform Detection

Detection of Straight Lines

The theory and algorithm for using the Hough transform to detect lines in images is presented by Duda and Hart (1972). The image pixels are mapped into a parameter space (ρ, Θ) by the polar Hough transform

$$\rho = x * \cos(\Theta) + y * \sin(\Theta) \tag{1}$$

where Θ is the angle in polar coordinates and ρ is the radius from the origin in polar coordinates. This point-to-curve transformation has the property that collinear points in the image correspond to curves through a single common point in the parameter space.

Modification of the Hough Line Detection Algorithm: Overview

The extension of the Hough transform algorithm from detection of a single line to a fiducial or marked ground point is

Photogrammetric Engineering & Remote Sensing,
Vol. 59, No. 7, July 1993, pp. 1161–1167.

0099-1112/93/5907–1161$03.00/0

B.R. Hunt
T.W. Ryan
E.A. Gifford
Department of Electrical and Computer Engineering,
University of Arizona, Tucson, AZ 85721.

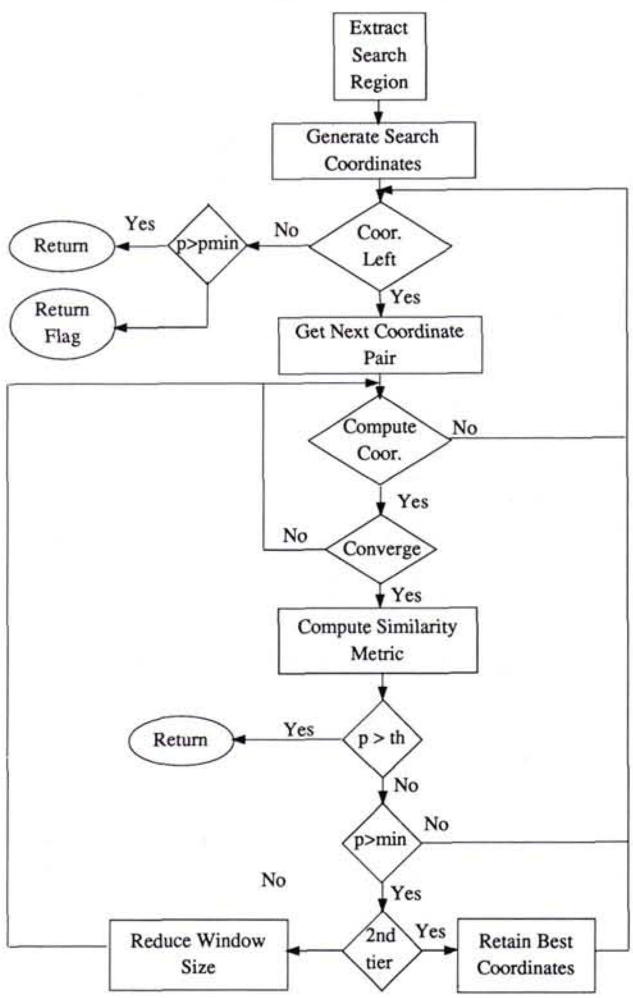

Figure 1. Fiducial detection algorithm.

relatively simple. Herein, we use the following simple model: that a fiducial mark or the image of a marked ground point creates in the image two orthogonal lines intersecting to form a structure similar to a plus sign, i.e., "+". We recognize that not all fiducials or ground points are orthogonal in the form of a plus sign, but the modification to nonorthogonal cases is a simple modification to our algorithm. The position of the fiducial mark is defined as the coordinates of the point of intersection. Therefore, a Hough transform of the fiducial or ground point produces two peaks or clusters separated by 90° in the parameter space (or separated by a different value, for nonorthogonal cases). The parameterization only requires two parameters ρ, Θ. The nature of the fiducial or ground point and the detection algorithm employed embody scale and rotation invariance.

The standard Hough algorithm is modified so that a projection onto the Θ axis of the accumulator array is correlated with a sliding window to detect an orthogonal pair of peaks. The sliding window consists of two tapered peaks separated by 90°. The two ρ, Θ pairs are used to identify the position of the fiducial mark in the image utilizing the inverse transform of Equation 2: i.e.,

$$\begin{bmatrix} x \\ y \end{bmatrix} = \frac{1}{\cos \Theta_1 \times \sin \Theta_2 - \cos \Theta_2 \times \sin \Theta_1} \times \begin{bmatrix} \sin \Theta_2 & -\sin \Theta_1 \\ -\cos \Theta_2 & \cos \Theta_1 \end{bmatrix} \times \begin{bmatrix} \rho_1 \\ \rho_2 \end{bmatrix} \quad (2)$$

Because Θ_1 and Θ_2 are separated by 90°, the equation can be simplified to

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} \sin \Theta_2 & -\sin \Theta_1 \\ -\cos \Theta_2 & \cos \Theta_1 \end{bmatrix} \times \begin{bmatrix} \rho_1 \\ \rho_2 \end{bmatrix} \quad (3)$$

Once the values of ρ and Θ are estimated by correlation with the detection window, the inverse of Equation 3 is used to transform the pair of polar coordinates into the x, y coordinates of the image plane.

Assumptions for the Algorithm

Assumptions made during algorithm development include

- The fiducials or marked ground points are constructed from two line segments intersecting orthogonally at their respective midpoints.
- The fiducials or marked ground points are positioned in a relatively homogeneous background darker than the mark itself, the contrast being greater for the external fiducials than for the marked ground points.
- The size of the fiducial and marked points is approximately known and fairly standard throughout the image.
- The external fiducials are "ideal": a bright, perfectly shaped cross in a black background. (These fiducials are placed on the film by the aerial photography process.)
- The marked ground points are cloth panels placed on the ground in precisely known locations and may not be ideal with respect to contrast, size, shape, and/or orthogonality.

Processing Techniques Used to Augment the Hough Transform

To implement the Hough transform algorithm for marked ground points on "real" imagery, the basic Hough algorithm is augmented with a variety of processing techniques. The final algorithm is broken down into five categories which function together to perform the precise detection of the marks. The five categories are

- Partitioning the Search Region;
- A two-tiered structure for windowing the image before processing;
- A recursive structure for increasing accuracy;
- Techniques for identifying marks, performing the Hough transform, and extracting/decoding the coordinates; and
- The computation of a similarity measure between a projection of the Hough transform for candidate coordinates and models of projections or of fiducials or marked ground points.

The first four elements extract features from the imagery having orthogonal line structure and approximately the same size as the fiducials or ground points, e.g., bright spots in dark backgrounds, etc. The similarity metric selects the feature most resembling a fiducial or ground point.

Algorithm Overview Details

Figure 1 is a flowchart of the detection algorithm. The software extracts the search region from the "box" drawn by the operator around the mark of interest, and generates a set of search coordinates which define a sequence of overlapping windows. The coordinates are processed until the fiducial or marked ground point is identified, or the set is exhausted. The portion of the image corresponding to each coordinate pair is processed using the Hough transform algorithm until either a set of candidate coordinates converge, or "no coordinates" are identified, whereafter the next pair of search coordinates are retrieved. Once a set of candidate coordinates are found, a similarity metric comparing the candidate feature to a known fiducial or marked ground point is computed. If the correlation is high enough, the program returns the coordinates as the coordinates of the fiducial. If the correlation is very low, the algorithm retrieves the next pair of search coordinates. If the correlation lies somewhere in the middle, the window size is reduced and a portion of the image centered at the candidate coordinates is extracted and processed similarly. The algorithm retains the best coordinates, those with the highest associated correlation, and terminates when the search region is completed. If the best coordinates exceed a minimum threshold, they are returned as the coordinates of the mark; otherwise, a flag indicating "no fiducial identified" is returned.

Partitioning the Search Region

The "box" selected by the operator defines the search region for the algorithm. However, it is not desirable to perform a Hough transform on the entire search region. The array sizes may be large, thereby increasing the computation time. Therefore, the algorithm partitions each search region into a number of overlapping windows and applies the Hough transform to each window. This partitioning reduces the problem of determining which pixels have a high probability of lying on the fiducial to a simple thresholding problem, because it is assumed that the fiducial or marked ground point is relatively brighter than the imagery in its immediate neighborhood.

Prior to execution, the algorithm extracts the operator-delineated search region of the image from the image file. Next, the algorithm generates a sequence of search coordinates in an outward spiral, initiating at the center of the user-defined box and extending to encompass the entire search region. The points are computed so that the windowing operations of the first stage will overlap, e.g., 30 percent. This is done so that the entire candidate mark will fall entirely within one of the windowed segments of the image.

The Two-Tiered Windowing Structure

The algorithm could proceed using only a single fixed window size, but this method has a number of drawbacks. If the window is too large, extraneous features may be included in the window and be identified as fiducial pixels, thereby reducing the accuracy of the generated coordinates or possibly preventing detection altogether. If the window is too small, not enough of the fiducial or marked ground point and contrasting background may be captured to establish an optimum threshold and perform the Hough transform. Using a smaller window also increases the total number of windowing operations required to cover the entire search region, thereby increasing computation time. Consequently, the algorithm employs a two-tiered windowing structure in which a relatively large window is used to partition the search region and generate candidate coordinates, and a second smaller window is applied to the candidate coordinates.

In the first stage, the algorithm employs each of the search coordinates, extracts a partition of the image centered about the search coordinates having a size equal to the first stage window, thresholds the windowed partition to identify fiducial pixels, computes the Hough transform on each of the prospective fiducial pixels, and recursively generates a pair of candidate coordinates for each window where the image structure suggests a possible fiducial mark. If the evidence of a fiducial mark is strong enough, the algorithm will return those coordinates and terminate. If the similarity metric exceeds a minimum threshold, the algorithm proceeds to the second windowing stage after each pair of candidate coordinates is identified. Otherwise, the algorithm gets the next pair of search coordinates.

In the second stage, the algorithm reduces the window size by half and calculates the Hough transform, substituting the candidate coordinates for the search coordinates employed in the first stage. As coordinates are generated by the second stage, presumably being very similar to their corresponding candidate coordinates, those coordinates having the strongest similarity metric are retained. If a pair of coordinates has a sufficiently high similarity metric, the algorithm returns those coordinates and terminates.

The first tier of the windowing structure identifies features in the image that exhibit characteristics in their Hough transforms similar to those of a fiducial mark. The second tier, centering a smaller window around the candidate coordinates, filters out most of the remaining imagery and improves the accuracy of the coordinates. The smaller window also provides a more accurate Hough transform and a better similarity metric, facilitating the selection of the correct coordinates.

Recursion

The fiducial detection algorithm is implemented recursively to improve the accuracy of the fiducial coordinates. When the algorithm, operating on search coordinates in stage 1 and candidate coordinates in stage 2, returns a pair of candidate coordinates, those coordinates are recursively fed back into the function until the candidate coordinates output by the function do not change. When convergence occurs, the final coordinates are returned. Recursion improves the accuracy of the algorithm because the windowing, and thus the Hough transform, is centered around the actual candidate coordinates.

Optimum Thresholding

In order to determine a set of candidate points to be applied to the Hough transform, the algorithm partitions the search region by sequentially windowing small portions of the image. An optimum thresholding function computes a threshold for each block of windowed imagery by maximizing the interclass variance between dark and bright regions (Reddi *et al.*, 1984). The interclass variance as a function of the threshold k is defined as

$$\sigma^2(k) = p_D(m_D(k) - m_o(k))^2 + p_B(m_B(k) - m_o(k))^2 \quad (4)$$

where

p_D is the probability of a dark pixel whose value is less than k,
p_B is the probability of a bright pixel whose value is greater than k,
m_D is the mean of the dark pixels,
m_B is the mean of the bright pixels,
m_O is the total means, and
$\sigma^2(k)$ is the interclass variance for threshold k.

A small partition of the image containing a fiducial mark should contain a relatively few bright pixels and a large number of darker pixels. The thresholding procedure should set the threshold sufficiently high such that only fiducial or marked ground point pixels are transformed. In the implementation, if the number of pixels identified as fiducial or marked pixels exceeds 30 percent, the program rejects the candidate region and moves on to the next set of search coordinates. The 30 percent figure is an approximate upper bound on the percentage of a window a fiducial could occupy. The maximum percentage occurs for a small window having the same width as the a fiducial. For example, a 6 by 6 window (36 pixels) encompassing a fiducial having a width of 6 pixels (11 pixels) gives the upper bound of approximately 30 percent. Obviously, the size of the windows, the width of fiducials or marked ground points, and the pixel size must be used to set this threshold for a class of imagery.

For windowed imagery not containing a fiducial mark, the thresholding function returns either a threshold if there is sufficient variation of grey levels in the imagery, e.g., some type of feature, or returns a flag causing the algorithm to move on to the next pair of search coordinates. This flag is returned if the windowed imagery is substantially homogenous, all black or all white. Using the 30 percent model and the thresholding function, the algorithm eliminates a substantial portion of the imagery before processing, thereby saving computational time and reducing detection errors.

Hough Transform

The algorithm maps those pixels in the windowed image whose grey levels exceed the threshold into the parameter space (ρ, Θ) using the Hough transform equation

$$\rho = x_i \cos(\Theta) + y_i \sin(\Theta) \qquad (1)$$

This mapping is performed for each *x,y* pair by computing ρ for each value of Θ between 0° and 175°. The parameters are quantized such that the angle Θ has 5° of resolution and ρ has a resolution of 0.5 pixels, e.g., each change of 0.5 in ρ is mapped into a separate cell in parameter space.

Detection/Decoding

Because the fiducial mark is identified by two orthogonal points in the parameter space, the parameter space is scanned for the largest pair of orthogonal peaks. This is accomplished by first setting a threshold, e.g., 30 percent of the maximum peak, and thresholding the parameter space: setting values less than the threshold to zero and maintaining those values which exceed the threshold. Higher thresholds may filter out more erroneous features, thereby reducing the number of candidate coordinates, but may also risk removing the correct coordinates.

The algorithm projects the thresholded transform onto the Θ axis by summing the remaining accumulator magnitudes (AMs) for each Θ and dividing by the square of the maximum distance between any two of the non-zero peaks: i.e.,

$$\text{projection}(\Theta) = \Sigma(\text{AM}(\rho))/(\rho_{max} - \rho_{min})^2 \text{ sum over } \rho, \qquad (5)$$

where ρ_{min} and ρ_{max} are the minimum and maximum values of ρ, respectively, that have non-zero accumulator values for a given Θ. This projection emphasizes sharp, high peaks in the accumulator and deemphasizes broad peaks.

The one-dimensional projection is then scanned for the maximum orthogonal pair of peaks, e.g., 0° and 90°, and concurrently retains the coordinates in ρ,Θ that contributed to those peaks. A matched filter having 90° peaks is correlated with the projection to account for any parameter quantization effects and Hough transform uncertainty. Ideally, each peak in the projection would correspond to only one point in the accumulation array, but in practice the peak will most likely have two or three associated points. The algorithm eliminates as extraneous peaks those having ρ values significantly different from the majority, and computes an average ρ for each of the orthogonal Θ values. The algorithm then takes these two ρ,Θ pairs and decodes them using the inverse Hough transform (Equation 3) to produce the candidate coordinates in *x,y* space.

Similarity Measure

The final major element of the algorithm is the similarity metric. The similarity metric is employed to select the "best" pair of candidate coordinates identified by either the first or second tiers of the windowing system. Ideally, the computer identifies only the correct coordinates as candidates, but some partitions of the image may contain features which have transforms and projections that exhibit orthogonal structure. The similarity metric is a correlation of the projections of the candidate coordinates with models of projections or signatures of known fiducial marked ground points. The program employs two models for profiles; one for an ideal external fiducial and one for a marked ground point. The ideal fiducial projection is two very sharp peaks separated by 90°. The marked ground point is modeled by a raised cosine function with a period of 90°, i.e.,

$$f(\Theta) = 1 + \cos(4(\Theta - \Theta_0)) \qquad (6)$$

where Θ_0 is the rotation of the fiducial. The choice of the raised cosine function to model the fiducial projection was made by visually observing the projection of a fiducial.

The similarity metric is computed for each set of candidate coordinates in both the first and second windowing stages. If in stage one the correlation is very high, the program returns the coordinates and terminates. In stage two, the program retains those coordinates with the highest metric, returning them when the search region is exhausted. Again, if the metric is sufficiently high, the program returns the coordinates and terminates before all the imagery has been evaluated. If, in either stage, the metric does not exceed a minimum value, the algorithm moves to the next partition. Note finally that correlation with the model effectively solves the problem of a fiducial or marked ground point being several pixels in width, because the several adjacent ρ bins will still tend to a single peak in the ρ coordinate after correlation is completed.

Resolution

The goal of the algorithm is to identify the *x,y* coordinates of the fiducial mark with one pixel resolution. The quantization of the parameter space (ρ,Θ) affects the performance of the algorithm. If ρ,Θ are quantized too coarsely, the array size and thus the computation time diminishes and good peaks are formed but the resolution of the fiducial coordinates in *x,y* space fails to meet the one-pixel tolerance. Transform redundancy may also cause significant random peaks to be formed in the parameter space. Quantizing too finely increases computation time and causes the peaks to be broad which also reduces the accuracy of the final coordinates.

Testing was used to find ranges of quantization levels for ρ and Θ that produce one-pixel accuracy. If ρ is quantized approximately between 0.33 and 2 pixels and Θ between 1° and 30°, the fiducial coordinates are identified within the one-pixel tolerance having a similarity metric exceeding 0.4. Because the ρ,Θ detection is by the correlation on Hough projections discussed above, subpixel detection accuracy can

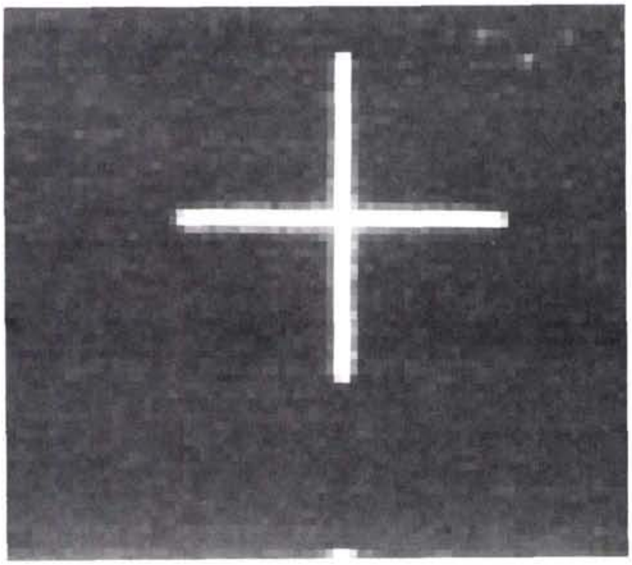

Figure 2. External fiducial 1.

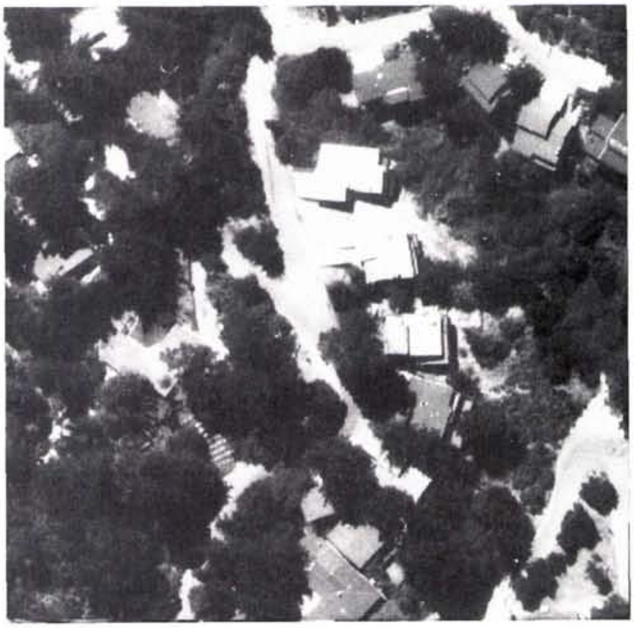

Figure 3. Internal fiducial 1.

be possible if the usual techniques of interpolating the correlation function are employed.

Testing

The algorithm has been tested on aerial photography images of residential areas having marked ground points and external fiducials. The external fiducial, shown in Figure 2, is a very bright mark in a black background having a size of 40 pixels and was measured to have coordinates 43,28. The marked ground points in each of two 404- by 404-pixel images are fairly light cloth panels forming an orthogonal cross in a relatively darker background having a size of approximately 8 pixels and known coordinates of 200,200. See Figure 3.

The External Fiducial

The computer correctly identified the external fiducial mark (Figure 2) at coordinates 43,28 in the image plane with a correlation of 0.86. A Hough space quantization of 0.5 pixels and 5° was employed. Figure 4 is a surface plot of the thresholded Hough transform. Only two pairs of very sharp peaks remain in the accumulation array after thresholding. The two large peaks associated with each Θ are attributable to the construction of the external fiducial. The external fiducial mark is two pixels wide, and thus the accumulator array quantized at 0.5-pixel resolution will contain two peaks for each Θ; one for each of the adjacent parallel lines. Because the lines are one pixel apart in the x,y plane, the transform peaks should be one pixel or two cells apart in parameter space. If the quantization is changed to two pixels, both peaks should occur in the same cell.

Marked Ground Point (Cloth Panel)

Figure 3 is a 404- by 404-pixel aerial photograph of a suburban area in northern California. The image comprises houses, roads, trees, and a marked ground point (cloth panel) at coordinates 200,200. (Note: due to printing effects, it may be very difficult for the reader to identify the cloth panel.)

Table 1 enumerates the candidate coordinates generated during the first and second windowing stages for initial window sizes of 16 and 12 pixels for the marked ground point. For an initial window size of 16, the computer identifies only 11 candidate coordinates out of approximately 160,000 image pixels. The minimum thresholding criterion and second stage reduce this number to only three coordinates, the correct coordinates and two other features. The similarity metric for the correct coordinates is 0.79 compared to 0.34 and 0.33 for the other two coordinates. For the initial window size of 12, the computer identifies only nine coordinates in the first stage and three in the second. The similarity metric of the correct coordinates is 0.79 compared to 0.11 and 0.17 for the other coordinates.

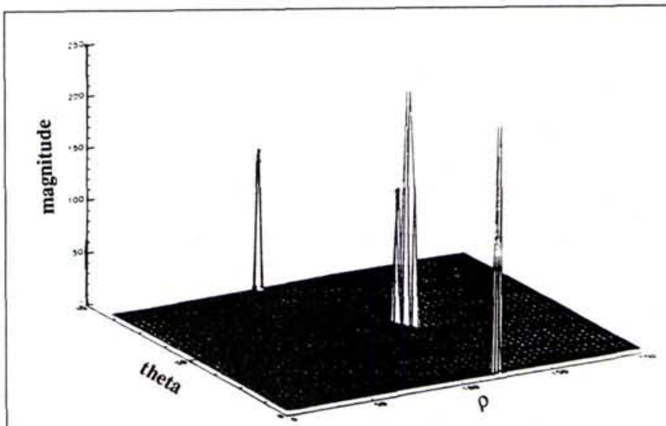

Figure 4. Thresholded Hough transform.

Testing the Effects of Varying the Window Size

The size of the window employed to segment the image into blocks for processing is an important aspect and constraint of the algorithm. Unfortunately, the larger the window, the more difficult it is to isolate the fiducial pixels within the windowed image. Employing the optimum thresholding function to distinguish the fiducial pixels from the neighboring imagery is only effective if the surrounding area is fairly homogeneous and darker than the fiducial pixels. Increasing the window size may capture imagery that changes the threshold such that actual fiducial pixels are not identified or may identify extraneous pixels as fiducial pixels which are thereafter included in the Hough mapping. Large windows may prohibit fiducial detection, increase the false alarm rate, and degrade coordinate resolution by including non-fiducial features in the window. The fiducial algorithm typically operates in a computing environment where the operator may literally spend hundreds of hours producing a map. Thus, a few seconds of background computation is of secondary importance. For first stage window sizes ranging from 8 pixels to approximately 24 pixels, the coordinates are identified within the 1-pixel tolerance with correlations exceeding 0.4. For window sizes smaller than the above, the algorithm identified the coordinates and terminated before moving to the second stage of windowing.

Testing the Effects of Varying the Parameter Space Quantization

The quantization of the parameter space affects the resolution and confidence of coordinate identification. If the quantization is too coarse, random points may be mapped into the same bin, causing false peaks. If the quantization is too fine, sharp peaks will not build up and a fiducial will not be detected. Additionally, too fine a resolution may cause nearly collinear points not to be identified as a single line. Such a smoothing capability is necessary to compensate for the distortion of the fiducial caused by the digitization process.

Adequate identification (correlation>0.5) was realized for $0.33<\Delta\rho<2$ with $\Delta\Theta = 5°$ and for $1°<\Delta\Theta<30°$ with $\Delta\rho = 0.5$. Maximum performance (correlation = 0.83), greatest detection confidence and coordinate resolution, occurred for $\Delta\rho=0.5$ and $\Delta\Theta = 10°$.

Testing the Effects of Varying the Hough Threshold

The Hough transform threshold is employed to remove small peaks in the accumulation array prior to projecting the data onto the theta axis. Ideally, only the two orthogonal peaks would remain after thresholding, thereby facilitating a large correlation, as is the case for the external fiducial. However, for marked ground points, which are much smaller and less ideal, the orthogonal peaks are not as pronounced, and relatively large random peaks occur.

The actual threshold used for gleaning the transform array is arbitrary as long as it is less than the peak value. The criteria for selecting the threshold are high probability of detection, accuracy of identification, and high correlation. If the threshold is too high, the orthogonal peaks may be removed from the transform. Additionally, if the threshold is too high, peaks that actually represent lines in non-fiducial features may be removed, possibly causing the algorithm to misidentify the feature as a fiducial. If the threshold is too low, the small peaks in the array may not be filtered, thus widening the peaks associated with the orthogonal components, flattening the projection, and reducing the correlation.

The algorithm identifies the fiducial within one pixel, with a correlation exceeding 0.4, for a wide range of thresholds, from 0 to 90 percent of the peak value in the parameter space. The optimal performance is found at 30 to 40 percent of the highest peak. A threshold of approximately 30 percent is large enough to filter out noise and bins containing single votes while retaining all the significant peaks in the parameter space.

It should be possible to find an optimal value for thresholding the Hough transform. Computing the optimal threshold would require constructing a probability density function for the parameter space representation of a fiducial mark from a large number of instances of fiducials. The density function could then be employed to determine an optimum threshold. A very large amount of test data would be necessary to construct a probability density function. The simple thresholding procedure provided good results.

Table 1. The Candidate Coordinates and Associated Correlations Produced by the First and Second Windowing Stages for Initial Window Sizes of 16 and 12 Pixels, Respectively, Computed for the *Entire* Image.

Stage 1			Stage 2		
x	y	ρ	x	y	ρ
Window = 16					
85	87	0.3			
101	67	0.48	101	67	0.34
55	118	0.12			
56	133	0.08			
60	140	0.26			
18	106	0.31	18	106	0.33
200	200	0.76	200	200	0.79
239	199	0.31			
136	225	0.28			
76	212	0.0			
95	227	0.13			
Window = 12					
200	200	0.81	200	200	0.79
191	189	0.25			
101	67	0.43	101	67	0.11
260	243	−0.66			
137	225	0.32			
227	284	0.47			
83	80	0.44	83	80	0.17
54	119	0.10			
18	106	0.18			

Conclusions

An automated algorithm combining the feature identification abilities of the photogrammetrist and the processing capabilities of the computer for extraction of cartographic calibration data from aerial photography has been demonstrated. The algorithm accomplishes the two tasks of correctly identifying the calibration marks and accurately positioning the mark. The algorithm correctly selects the calibration mark with a high level of confidence. Using this algorithm, the computer is able to quickly and accurately identify the coordinates of either external fiducial marks or marked ground points within one pixel of resolution, while requiring the photogrammetrist to only "box-in" the mark and select the algorithm from an appropriate menu.

Acknowledgments

This research was performed under a contract with International Imaging Systems of Milpitas, California.

References

Duda, Richard O., and Peter E. Hart, 1972. "Use of the Hough Transformation To Detect Lines and Curves in Pictures", *Communications of the ACM*, 15(1):11–15.

Fischler, Martin, 1986. *Computer Vision Research and its Application to Automated Cartography*, Technical Report, SRI International, Menlo Park, California.

———, 1989. *Image Understanding Research and its Applications to Cartography and Computer-Based Analysis of Aerial Imagery*, Technical Report, SRI International, Menlo Park, California.

Reddi, S. S., S. F. Rudin, and H. R. Keshavan, 1984. An Optimal Multiple Threshold Scheme for Image Segmentation, *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-14(4):661–665.

(Received 3 September 1991; Revised and Accepted 6 October 1992; Revised 30 December 1992)