Crane, R. K., Whitney, A. R., “Radar Meteorological Applications of Automatic Film Reading”, Proceedings of the 13th Annual SPIE Technical Symposium, pp. 399-405, 1968.

RADAR METEOROLOGICAL APPLICATIONS OF AUTOMATIC FILM READING

R. K. Crane and A. R. Whitney
Lincoln Laboratory*
Massachusetts Institute of Technology
Lexington, Massachusetts

Introduction

The data recording of storm information as detected by a weather radar has been customarily made on photographic film. Research radars and an occasional U. S. Weather Bureau radar are fitted with scope cameras to record the radar plan position indicator (PPI) display. Over the past 15 years a large sample of weather radar data has been accumulated in this fashion. The photographic technique provides an easy, quick, and inexpensive way to record weather radar data. The major drawback of this technique is data reduction. Information on storm shape, size, and intensity is normally extracted from the photographic images by hand. This means that only the most interesting aspects of individual storms are analyzed and the vast majority of the collected radar data is not analyzed. A vast amount of climatalogical information could be obtained from the existing store of weather radar data if an automatic technique of data retrieval were available. The first part of this report describes the use of a computer-controlled programmable film reader to process weather radar PPI photographs to obtain digital maps of rainfall intensity for use in climatalogical studies.

A second, related use of the film reader is the processing of raindrop shadowgraphs. The radar data describes the scattering properties of a volume of rain. This, however, is not of direct use to a meteorologist The information of most use is the liquid content of the rain. This is related to the scattering properties of a rain volume in a complicated way depending primarily on the distribution of raindrop sizes in the storm volume. One method of determining the sizes of raindrops is to photograph the shadows of the drops and to size the images on the film. The problem of automatically sizing raindrop shadows using the program-controlled film reader will be discussed in the second part of this presentation. Of importance in this application is the ability to recognize and reject irrelevant images and properly size drops of various shapes.

The work presented here has not been completely finished. What will follow is a description of an application of film reading to meteorological problems, a description of what has been done with some preliminary results, and conclusions about the usefulness of this approach.

PPI Film Reading

A PPI display presents position information for each target detected during an azimuthal scan of the radar antenna. A storm detected by the radar is displayed as a brightened area, often with indistinct edges, as shown in Fig. 1. The position and size of the storm on the display represents the position and size of the storm in the horizontal plane. Little of no information about storm intensity is provided by the display because of the limited dynamic range inherent in spot brightness. Intensity information is obtained by making a series of from five to ten PPI photographs each using a different receiver gain setting. Figure 1 was taken at the highest gain setting and represents the maximum areal coverage of the storm as detected by the radar. Successive frames for this storm show a reduced area of more intense rainfall. Figure 2 shows the fourth frame of

* Operated with support from the U. S. Air Force.

Fig. 1. Photograph of PPI display at highest receiver gain. Dark areas near center are ground clutter.

Fig. 2. Photograph of PPI display at fourth step of receiver gain

the sequence and represents the area of the storm with rainfall intensity greater than or equal to 14 times the minimum rainfall intensity displayed in Fig. 1.

The boundaries of brightened areas on successive frames are contour lines of equal storm intensity. All the frames taken together represent a contour map of rainfall intensity. The problem of automatically processing weather radar PPI data is that of generating a digital contour map from a set of maps representing the areas within particular contours. To do this, each frame must be read and remembered for comparison with the previous frame for the same data set.

A major difficulty in computer reading of this film is to correctly locate the exact position of each frame so that proper overlaying can be accomplished. This problem occurs in weather radar photography since no attempt is made to precisely position the film when the pictures are taken. One method of properly locating the image is with special feducial marks that can be recognized by the computer, such as display indicators. For the data presented here this was done by finding the display scope edges. Successive frames for a data set were scanned with a manually preset threshold. Figure 3 shows the overlay of data for a portion of the storm near the right edge of Figs. 1 and 2. Note that the contours of the storm as suggested by Figs. 1 and 2 are approximated closely by the film-scanning procedure. It was found, however, that the storm contours could vary significantly due to slight changes in the threshold. Such data as range rings can easily result in incorrect contours if the threshold is not carefully set.

The accuracy with which the PPI photographs can be read depends fundamentally on an accurate calibration of the radar, the PPI display, and the film reading threshold. The storm as displayed on the PPI photograph has indistinct edges. This is in part due to the method of operation of the display where, over a limited range of received signal levels, the intensity of the spot depends on the signal level. The uncertainty in threshold setting therefore becomes an error in the calibration of the radar, which can be removed only by calibrating the radar and the film reader by introducing targets of known size and cross-section or intensity. In processing radar data in which calibration targets are not present, which is true of most radar data currently stored on film, a fairly large uncertainty in the calibration constants must be tolerated. Experience has shown that efficient and accurate reduction of PPI weather radar data is possible only with a carefully maintained and calibrated system.

The authors are aware that other instruments, such as phase-contrast readers,

Fig. 3. Detail of digital contour map made from storm in Figs. 1 and 2. Outline shows contour level 4.

are able to efficiently generate contour maps from photographic data such as PPI photographs, but the output of most of these instruments is not in a form that is usable for further digital processing. The digital maps resulting from the film-reading procedure described here are placed directly onto magnetic tape so that further computer processing can be done without reference to the original PPI photographs.

Automatic Raindrop Sizing

Drop counting and sizing has been done manually for many years. However, most conventional methods are slow and cumbersome, and are usually limited to relatively small sample sizes. This is a distinct handicap to the radar meteorologist, who needs to have accurate drop-size statistics for many different storms. In an attempt to overcome these difficulties, a computer-controlled film scanner is being used in an attempt to increase the speed of analysis and guarantee uniformity of results. Although the computer program is not yet completely operational we have clearly demonstrated that there are no basic difficulties in automatic raindrop sizing.

The data which the film reader must process consists of shadowgraphs of raindrops in flight. Figure 4 shows the basic principles of the "raindrop camera," which was originally developed and used by the University of Illinois (Ref. 1). Raindrops are allowed to fall through a sampling volume approximately 3 ft. on a side and 14 inches deep. A high-speed strobe backlights the drops and images them as shadows on a photographic plate. A Cassegrainian optical system is used to eliminate any perspective effect through the 14-inch depth of the sample.

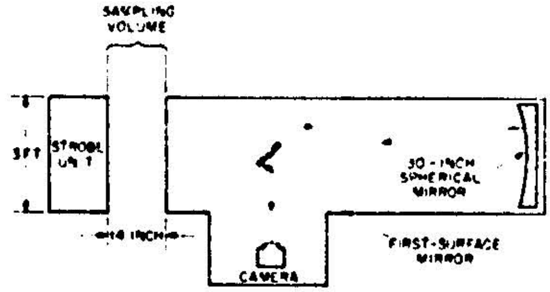

Fig. 4. Raindrop Camera Schematic

Photographic data is taken on 70 mm film, which gives a size reduction of approximately 15 times. Therefore, 500-micron-diameter raindrops, the smallest of interest, have an image diameter of about 30 microns on a 70-mm frame. The resolution of the film scanner, or light-spot size, is approximately 25 microns so that some care must be used when sizing these very small drops.

The film reader can be commanded to examine any point in a 4096 x 4096 point grid on the film frame, although only 1024 x 1024 points are actually resolvable. Use of the extra-fine spot position increments is valuable when sizing very small drops but is ordinarily not needed. The output of the scanner as it examines each commanded point is a gray scale level from 0 to 256 although the useful gray level discrimination is limited to about 16 levels.

A small area of a 70-mm sample frame is shown in Fig. 5. (This particular frame was made of artificially produced spray to simulate rain and has more drops per unit volume than most natural rain.) A large variety of shapes and sizes is evident,

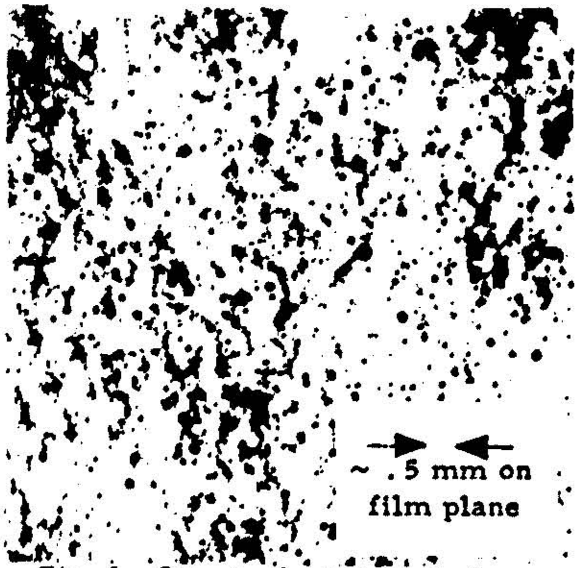

Fig. 5. Sample of raindrop shadowgraphs to be computer processed.

many with a characteristic "hole" in the center. This is due to the fact that a drop acts basically like a spherical lens and permits light passing through the center to remain relatively unrefracted. Drop diameters range from several hundred microns to several thousand microns. Only drops larger than approximately 500 microns are of interest in the present study.

Two quantities are of interest when reading these raindrop photographs. The first and most important is a measure of the size and number distribution of the drops. Secondly, and of somewhat lesser importance, is a measure of the ellipticity of the drops. Only two-dimensional information is, of course, available from the photograph and it is therefore necessary to make simple assumptions about the symmetry of the drops in order to gather any of this information. But, even ignoring this problem, a close examination of the sample data shown in Fig. 5 shows that some of the drops (or at least their images) are so mishapen that it is somewhat difficult to accurately interpret their size. Therefore, any automatic sizing procedure can only hope to size these sorts of drops consistently and rely on the basic random nature of raindrop sizes to smooth out errors by averaging over a very large sample. Also, any automatic sizing procedure must be able to recognize any very "badly" shaped drops or irrelevant images (such as borders and mirror supports) and either deal with them automatically or call for operator assistance.

The automatic sizing procedure is essentially divided into three steps: (1) A determination of the threshold that should be used over different subareas of the film. (2) A coarse scan of the film to find and tentatively isolate individual drops. (3) A fine scan of each individual drop to determine its size. These three steps are described in detail in the following sections.

Threshold Set and Drop Isolation

The first step in processing a frame of data is to scan the entire frame rapidly and perform a "dynamic" threshold level set. This allows the threshold level to be semi-continuously contoured over the entire frame to allow for variations in the gray scale of the background from one area of the frame to the next. During this dynamic threshold set, a preset threshold is used to detect whenever the spot has obviously encountered a drop and this point is rejected in setting the threshold for this particular area of the film frame.

After thresholding a coarse scan is made of the entire frame using the newly determined thresholds, in what might be called a "drop-isolation scan." Sample spot spacings are chosen so that each drop, down to the smallest diameter of interest, is "hit" at least once by the scanning spot. The coordinate location of each point that passes the threshold test is stored in core memory to await further processing. For large drops these points frequency occur as short line segments, and advantage can be taken of this to reduce the core storage space required to store this data. Instead of storing each "passed" point as a coordinate pair, these points are organized into horizontal line segments which can be stored as a starting pair of coordinates and a length. In practice this reduces required storage space by nearly an order of magnitude over a point-by-point storage method and makes subsequent processing faster and easier.

Immediately after the completion of the coarse scan all of the stored points are organized into connected areas, each such area being defined as an array of points (although actual storage is by lines) within which only adjacent points are included. Each of these connected areas tentatively identifies the location and approximate size of a drop. Two or more drops which are very close together may be included in a single connected area and hence be tentatively identified as a single drop. However, later processing will isolate each of these drops. Borders, mirror supports, etc., are identified by the large dimensions of the connected areas of which they are a part and are automatically deleted from further processing. Drops overlapping these large areas are deleted at the same time, but this is not serious because of the relatively infrequent occurrence of such events.

With all drop locations tentatively identified the computer enters the third phase of its processing of the frame. Each remaining connected area is chosen in succession and a slightly larger imaginary box drawn around it. This box is then scanned with a very fine increment scan in which one spot overlaps the next by the radius of the spot (~ 13 microns). This guarantees the highest possible "resolution" in this fine scan. The return from this scan is subjected to threshold and organized into connected areas in exactly the same manner as the coarse scan. Figure 6 shows an expanded display of the results of such a scan for a particular drop, superimposed on the field of the original coarse scan.

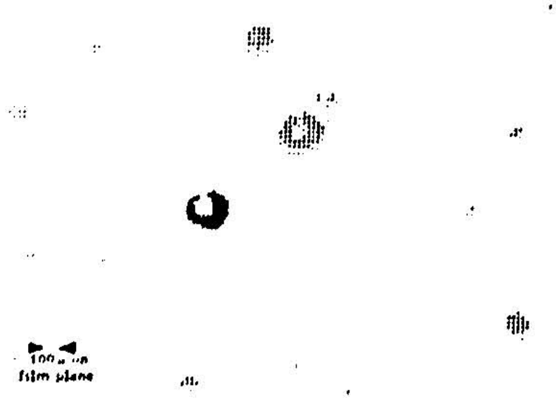

Fig. 6. Fine scan of single drop superimposed on coarse scan results. Box around drop shows the area that was fine scanned.

At this point there may be one or more complete drops within the finely scanned area and one or more partial drops which merge with the edge of the fine-scan area. Only connected areas which are totally included within the fine-scan area will be sized. Any which merge with the fine-scan box border are presumed to be the "main" drops in other fine-scan boxes and will be sized independently. These areas are automatically ignored and attention focused on the remaining connected areas. The number of points in each of the remaining connected areas is counted. The area with the largest count and any others with a count of at least one-fourth the largest are identified as separate drops. Any smaller areas are either small drops which have already been tentatively identified in the coarse scan and will be sized independently or very small off-shoots from the main drop(s) in this fine-scan box and not really drops in themselves. These are likewise ignored. The remaining connected areas are now identified as independent drops to be sized.

Figure 7 shows an example where three very closely spaced drops of approximately the same size form a single connected area after the coarse scan. The fine-scan box

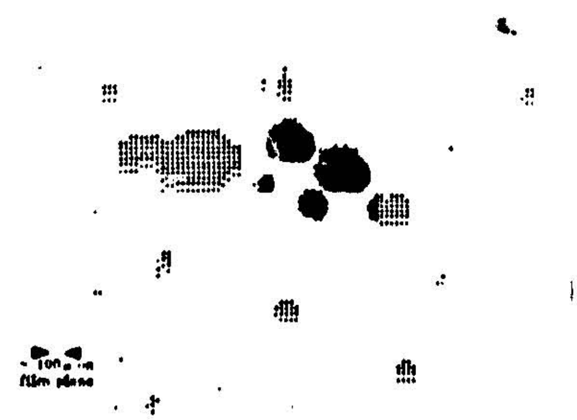

Fig. 7. Closely spaced drops. Square box defines fine-scanned area. Shaded area is result of fine scan. See text.

for this area is then automatically determined to completely enclose this connected area, as is shown. The fine-scan results, however, show the three drops as three separate connected areas of approximately the same size and entirely within the fine-scan box. These are then to be sized as individual drops. The areas touching the border of the fine-scan box are ignored since they form the main drops in other areas which will be fine-scanned later. Although the situation in this example is rarely encountered, it demonstrates how the drop-isolation procedure is capable of handling fairly complex data in an efficient manner.

Sizing

The actual sizing of the drops is now performed entirely from information gathered in the fine scan and stored in memory. After experimentation with several different methods of sizing a fairly simple and reliable algorithm was chosen. The image is simply enclosed in a closest-fit sixteen-sided polygon with all interior angles fixed to be equal but with sides allowed to shrink to zero length if necessary to achieve the best fit. The resulting polygon is unique and is taken to define the outline of the drop. In the vast majority of cases this gives a very close approximation to the size that a manual measurement would yield, which has to be set as a measure of the automatic sizing accuracy. With this type of sizing procedure the holes and gaps present in many of the drops are in effect ignored and sizing proceeds as if they were not present. Also, this method allows for some fairly simple checks on the validity of the sizing measurement. The "ellipticity" of the drop can be measured and used as one of the measured parameters as well as a signal that the sizing procedure is not sufficient. For example, two slightly overlapping drops (one behind the other in the sample volume) have a greater "ellipticity" than would normally be encountered in a single drop and hence this large ellipticity can be used as a signal as a need for operator assistance. The drop is then automatically displayed for the operator to size manually.

An example of the result of sizing a typical drop is shown in Fig. 8. Note that the sizing algorithm has effectively ignored any holes or gaps that might otherwise have upset the accurate determination of the drop size.

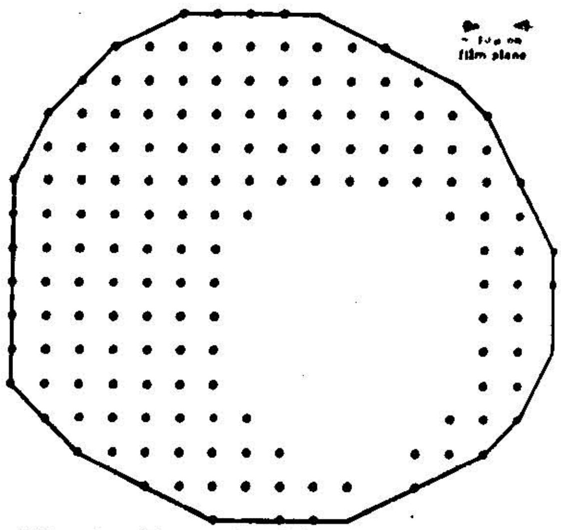

Fig. 8. Example of how sizing algorithm outlines a typical array of points from a fine scan of a drop image.

As we noted earlier one of the advantages of sizing with a machine is that the results are uniform to a very high degree. Reports of manual measurements of drop sizes from the same type of shadowgraph data have shown differences in the results of different operators, even on the same frame of data. This is mostly due to operator fatigue and to the somewhat qualitative judgment required to size some of the drops. The second and most important advantage is the saving in time. It is estimated that the speed-up in the analysis of the drop photographs is at least two orders of magnitude. Approximately four to five minutes of computer time is necessary to process a 70-mm frame of data containing several thousand drops, compared to several hours of manual processing time.

A natural question arises as to the accuracy of the measurements of individual drops. Unfortunately, as this is written the program has not progressed to the point where an accurate assessment of these capabilities can be made, although several comments are in order. Large drops (~ 2 mm or larger in diameter) have large enough images that the sanning spot of the film reader records many points within the drop and hence can size it to within a small percentage of its true diameter. Very small drops, however, have image sizes approximately the same size as the spot and hence one must be very careful with the sizing procedure. Use of an overlapping fine scan helps to reduce errors, but a quantitative judgment of sizing accuracy is not yet available. Preliminary experiments show, however, that sizing accuracies to 150 microns (corresponding to 10 microns on the film plane) should be readily attainable. In order to increase this accuracy it may be necessary to size very small drops several times with independent fine scans and average the results. Since nearly all of the small-drop images are nearly circular, this procedure should yield a good estimate of the true size.

Conclusion

We have found from our work that relatively simple applications of automatic reading can be of great value in some of the data reduction problems we have encountered. And though neither of the two applications that we have described is yet completely operational, we are convinced that both are workable examples and that simple extensions or modifications of the concepts presented here can be of value in the automatic processing of similar data in other fields.

(Ref. 1) Jones, D.M.A., and Dean, L.A., 1953, A Raindrop Camera; Research Report No. 3, DA 36-039, SC42446, Illinois State Water Survey.

Biographical Statement

Robert K. Crane was born in Worcester, Massachusetts on December 9, 1935. He received the B.S. and M.S. degrees in electrical engineering in 1957 and 1959, respectively from Worcester Polytechnic Institute, Worcester, Massachusetts.

From 1959 to 1964 he was employed by the MITRE Corporation in Bedford, Massachusetts. Since 1964 he has been employed by the M.I.T. Lincoln Laboratory, Lexington, Massachusetts.

Biographical Statement

Alan R. Whitney was born in Cheyenne, Wyoming on September 11, 1944. He received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology in 1967. His experience includes work in laser research at the Bell Telephone Laboratories in Murray Hill, New Jersey and in radar meteorology at the M.I.T., Lincoln Laboratory in Lexington, Massachusetts. He is currently completing work on a Ph.D at Massachusotts Institute of Massachusetts.