Automated Digital Image Analysis of Video Ice Crystal Data*

Jane Niehues-Brooks
Frederick C. Harris, Jr.
Department of Computer Science
University of Nevada,
Reno, NV 89557
fredh@cs.unr.edu

Abstract

This paper presents a procedure for automating digital image analysis of cloud particle images recorded on video tape using the Cloudscope.[1] The Cloudscope is mounted on the wing of an aircraft and records cloud crystals as they impact the external lens and sublimate. The procedure presented breaks the analysis into several parts: frequency domain filtering, frame to frame movement correction, and finally new particle detection, counting, and sizing. These procedures were automated using LabVIEW and Concept V.i.

keywords: image analysis, Cloudscope, fft, ift, frequency mask, threshold, replicator

1 Introduction

This paper discusses an analytic tool for Cloudscope ice measurement. We have implemented a technique for automating the digital image analysis of cloud particle data recorded using the Cloudscope instrument. The Cloudscope is a specialized data collection device operated for atmospheric research by the Atmospheric Science Center at Desert Research Institute. This instrument mounts on the wing of an aircraft and is essentially a microscope with an attached video camera. Ice crystals and cloud droplets impact on the small window of the Cloudscope and are recorded by the video system (Hi8 video tape). A line drawing of this instrument is shown in Figure 1 [1]. The goal of this project was to automate the counting and sizing of the cloud crystals recorded and in future work add the ability to compute their the sublimation rates.

The automated processing system is made up of a programmable Sony-Visca Hi-8 VCR and an Apple Macintosh Centris 650 with a Scion LG-3 Frame Grabber Card. The frame grabber is loaded with 64 MBytes of RAM and is capable of grabbing up to 128 frames of video, which is equivalent to 4 seconds of flight time at a 30 frames per second capture rate. We created a virtual instrument to perform the automation and image analysis using National Instrument's LabVIEW instrument control software and Graftek's Concept V.i image analysis software. A similar program for the Replicator instrument [6] was used as a guide in implementing the LabVIEW and Concept V.i tools. The CloudScope Autoanalysis Vi is also made up of several of our own more complex image analysis modules. A detailed description of these can be found in [4].

The rest of this paper is structured as follows: Section 2 discusses the actual techniques used to perform the image processing. The Image Processing is broken down into three areas: noise removal, image alignment between frames, and new particle detection, sizing and counting. Results follow in Section 3.

2 Image Processing Algorithm

To succeed in analyzing the Cloudscope video data there are three basic tasks that have to be accomplished. These tasks include: filtering noise, correcting for image jitter, and finally detecting, counting, and sizing new particles. Once captured, each image goes through a sequence of steps to clean it up. After the image is cleaned up, a threshold is used to change it

*This work was partially supported by a NASA Ames under contract NAG 2-1104

[1]This instrument is operated by the Atmospheric Science Center at Desert Research Institute.

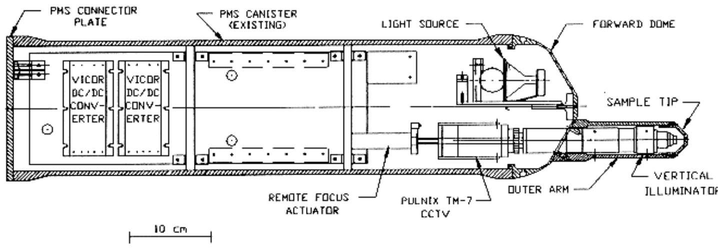

Figure 1: Cloudscope Line Drawing

to a binary image of white particles on a black background. A comparison with the prior image is performed and new particles present are identified and sized. In principle this is a very simple process. Unfortunately, the Cloudscope experiences both electrical interference and physical instability. These issues result in flight video recordings containing excessive movements from frame to frame, color variations due to inconsistent lighting, considerable high frequency signal interference in the images, along with false particles caused by tiny scratches and foreign matter on both the CCD lens and the impact lens. The ability to successfully threshold an image into particles and background is the key to automating the counting and sizing of the crystals. Unless preprocessing is performed on each image, any of the above problems will contribute to unsuccessfully thresholding the image, thereby making it impossible to automate the particle count.

For our processing purposes we use a 256 by 256 image. When captured, each image is 640 by 480, but we choose to transfer in every other line in both the X and Y direction for two reasons. The most obvious being that a smaller image will process faster, and the second, to handle a problem caused by the interlacing of the image. If a crystal impacts on the lens after the odd rows have been recorded, it will have a striped appearance due to just the even rows being recorded and would not be detected as a single particle. Using every other line of the image alleviates this problem. An adjustment to the image size must also be made. Since our Fast Fourier Transform (FFT) and Inverse Fourier Transform (IFT) routines require dimensions in powers of two, we extract a 256 by 240 portion of the image and expand the image in the Y direction from 240 to 256. This expansion causes an aspect ratio change, which we can correct for later.

The noise on the videos along with the wide color variations result in images that are impossible to threshold properly. Figure 2[2] is an example of a noisy low contrast image. The noise in this image is more apparent when using an enhancement technique called contrast stretching. This is a method of increasing the image dynamic range by spreading the images narrow range of gray levels over the available range of 256 colors. This technique is used here for better viewing.

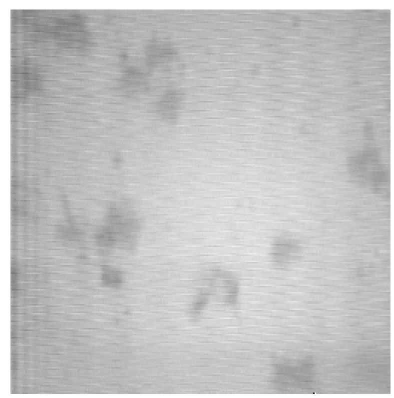

Figure 2: Unprocessed noisy image with poor contrast.

To investigate the noise present, we looked at the Fourier spectrum of the images in the frequency domain. The vertical lines on either side of the center line in the spectrum showed the high frequency components corrupting the image. From this it was apparent that some of the preprocessing must be performed in the frequency domain. We used an FFT to transform the image from the spatial domain to the frequency domain, and chose to use a frequency filter mask to remove the undesirable frequency components from the image. The lowest frequency color variation is due in a large part to the shape of the impact lens (a property known as vignetting [5], page 11). Other contributors

[2] All images are from the NASA SUCCESS Project, May 3, 1996 Flight

to the light variation are our artificial light source, daylight, flying into the sun, and night flights. This slowly varying color can be corrected by blocking out the lowest frequency, also known as the central–order component, located at the center of our spectrum. This technique is confirmed by discussions in [5] pages 10 to 12 and [3] pages 144 to 145. By placing a horizontal center line in the mask we can remove the vertical wave-like ridges that traverse horizontally from left to right across the image. This filtering method was adapted from the optical grid example in [3] pages 142–144, where a vertical slit passes only the vertical components of the grid and a horizontal slit passes only the horizontal components. The resulting high contrast image with the corrupting noise removed is given in Figure 3.

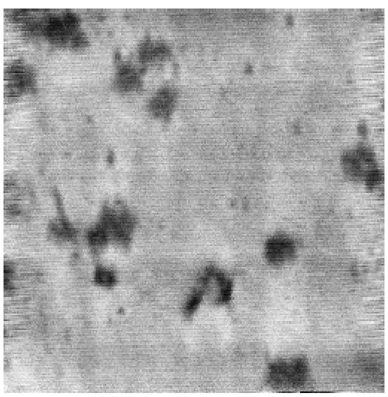

Figure 3: High contrast image after processing.

Another hurdle to overcome before any reliable data analysis can be performed is to correct for the image movement from frame to frame. Plane vibration from turbulence and electro-mechanical problems are more than likely the cause of this jitter. Image movement in this section of our video tape has been measured in excess of +/- 14 pixels. This is considerable since it lies within the size range of the crystals we hope to measure. In order to successfully detect a new particle, the two consecutive images must have the same reference point.

Several techniques for aligning the images were tried with various levels of success, such as: using the first image as a reference and aligning all following images relative to a fixed particle on that first image; selecting a target particle that is present on two consecutive images and aligning them based on the difference in the center of mass coordinates of the two matched particles; and performing a cross-correlation between the two images (the discrete version of this computation

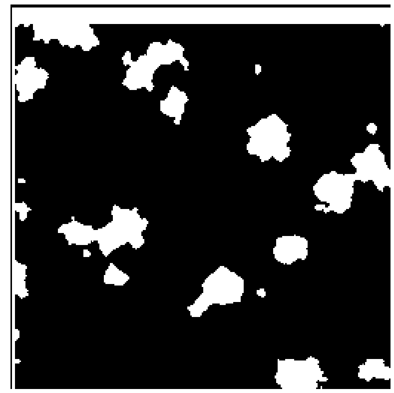

Figure 4: Binary version of image 2 after alignment.

is presented in [2]).

In aligning the two images, the prior image is held fixed and the current image is realigned by extracting all or a portion of the image based on the offset values, and then pasting it back onto a white background. A white background is used so that any particle lying on, or touching the boarder, will still be considered a boarder particle when the image is aligned. Figure 4 shows the result of aligning image 2 with image 1. A correction of +3 in the X–direction and +12 in the Y–direction is made. For most particles in the image, this is a fairly accurate correction, but this method is not without exceptions. There are always a few specks or scratches present on the CCD lens that show up in the images. Because of their physical location, they do not jump from frame to frame as do the particles on the impact lens. This introduces an error back into the analysis because these particles have just been misaligned causing them to be counted as new particles. It is possible that as a second stage to this procedure, we could reverse the alignment and subtract out the CCD lens noise and hope that no valid data gets removed. An alternate plan would be to leave it to some post processing number crunching where erroneous data can be identified and deleted from the data set.

Given two clean aligned consecutive images, we are then able to count and size the particles. When counting, we are only concerned with the new particles in an image. A new particle is one that was not present on the previous frame. The size of a particle is the area it occupies within the image. After thresholding this image is still not as clean as we would like so we add a median filtering step. This is a spatial convolution process in which we keep the median value of all neighboring pixels.

For a particle to be counted it must be surrounded by black background. If a particle touches or overlaps the image border, it is considered outside the field of view. Since its actual size is not known, it is removed from the image. Notice that both of the particles located at the middle of the right border of Figure 4 will be removed because they are touching each other and the border. In this case it is debatable whether all border particles should be rejected. Future advancements in processing techniques may minimize the loss of valid data such as this.

Finally, the *new* particles can be counted and sized. The size of a particle is based on the number of pixels that make up its contiguous white area. The parameters selected for our analysis include; area, max intercept length, mean perpendicular intercept, perimeter, center of mass X and Y coordinates, and the particles bounding box coordinates. Once the measurement is complete, the data is written to a file after the set of frames have been analyzed. This data will then be used to compute ice crystal concentrations

The actual dimensions of the 640 by 480 pixel image is 420 μm by 314 μm. Correcting for expansion of the original image in the Y dimension and the extraction of only 256 pixels of the 320 in the X dimension our actual per pixel dimension in the X dimension is 1.313 μm and 1.227 μm in the Y dimension. Our smallest particle size of three by three pixels square (3.94μm x 3.68μm) is 14.5 μm^2. Realistically, only those particles with an area of more than fifty pixels (80 μm^2) should be counted as real with any accuracy. For our measurements we have chosen to count all particles remaining after erosion and post filter the data based on size once it's been stored to a file.

3 Results

The time to process the ten frames of data is approximately 6.5 minutes, and to process a maximum of 128 frames (4 seconds of flight) is 87 minutes which includes the time to initialize and and grab the frames. This time will increase in future runs when we use the whole 320 by 240 image expanded to 512 by 256 or a 640 by 480 contracted and expanded to 512 by 512. Ideally, we would have preferred a faster process time, but due to the compute time necessary for the FFT and IFT, this was not possible. However, considering the alternative of a manual count, which for a dense particle concentration would be extremely tedious and time consuming, time is not the main issue.

There are some physical properties of the CloudScope data that we may never be able to handle. One example of this would be: a particle landing on its edge in one frame then falling over in the next frame. Even to the human eye this appears to be two particles and as expected our program will count it as two particles. Crystals overlapping either partially or completely is another uncontrollable occurrence. With our current program version this will be counted as one large particle if they both landed at the same time. If on the other hand, as discussed in the previous section, they touched down on two different frames the second particle may be detected as a much smaller particle, as two separate smaller particles, or as a particle with a whole or cutout in it.

Being able to detect and size very small ice crystals is of great interest to the atmospheric scientists who want to use this kind of analysis tool. We foresee this as an on going project of analysis enhancements and improved capabilities.

References

[1] P. Arnott, Y. Dong, R. Purcell, and J. Hallet. Direct airborne sampling of small ice crystals and the concentration and phase of haze particles. In *Ninth Symposium on Meteorological Oberservations and Instrumentation*, pages 415–420, 1995.

[2] R. Gonzalez and R. Woods. *Digital Image Processing*. Addison-Wesley Publishing Company, 1992.

[3] Joseph W. Goodman. *Introduction to Fourier Optics*. McGraw-Hill, Inc, 1968.

[4] Jane Niehues-Brooks. Automated digital image analysis of video ice crystal data. Master's thesis, Department of Computer Science, University of Nevada, Reno, December 1997.

[5] John C. Russ. *The Image Processing Handbook*. CRC Press, 1995.

[6] Mark Turner. Digital image analysis of replicator ice crystal data for detection of small particles in cirrus clouds. Master's thesis, Atmospheric Physics, University of Nevada, Reno, May 1996.