Presenting the Snowflake Video Imager (SVI)

ANDREW J. NEWMAN* AND PAUL A. KUCERA+

Department of Atmospheric Sciences, University of North Dakota, Grand Forks, North Dakota

LARRY F. BLIVEN

Instrumentation Sciences Branch, NASA Goddard Space Flight Center, Greenbelt, Maryland

(Manuscript received 28 March 2008, in final form 10 July 2008)

ABSTRACT

Herein the authors introduce the Snowflake Video Imager (SVI), which is a new instrument for characterizing frozen precipitation. An SVI utilizes a video camera with sufficient frame rate, pixels, and shutter speed to record thousands of snowflake images. The camera housing and lighting produce little airflow distortion, so SVI data are quite representative of natural conditions, which is important for volumetric data products such as snowflake size distributions. Long-duration, unattended operation of an SVI is feasible because datalogging software provides data compression and the hardware can operate for months in harsh winter conditions. Details of SVI hardware and field operation are given. Snowflake size distributions (SSDs) from a storm near Boulder, Colorado, are computed. An SVI is an imaging system, so SVI data can be utilized to compute diverse data products for various applications. In this paper, the authors present visualizations of frozen particles (i.e., snowflake aggregates as well as individual crystals), which provide insight into the weather conditions such as temperature, humidity, and winds.

1. Introduction

The Snowflake Video Imager (SVI) is a new instrument for characterizing frozen precipitation. We were motivated to develop this system by the need to quantify falling frozen precipitation over large geographic regions because of the importance of snowfall to drinking water supplies, agricultural production, floods, and source material for glaciers. To address these topics, the National Aeronautics and Space Administration's (NASA's) Global Precipitation Measurement (GPM) Mission plans to expand its satellite remote sensing capability to cover high-latitude regions as well as tropical regions. Many factors contribute to snowflake geometry, which is important for both active and passive retrievals (i.e., Bringi and Chandrasekar 2001; Matrosov et al. 1995; Meirold-Mautner et al. 2007; Liu 2004). Knowledge of snowflake classification, as well as particle distribution characteristics, ought to reduce uncertainty in remote sensing inversion algorithms. To develop reliable inversion algorithms, snowflake characteristics need to be documented at many diverse locations over extended periods of time to develop an understanding of the variability of snowflake characteristics for a given climatological regime.

The SVI is a new automated sampling system that can be deployed to obtain such data. The following sections contain descriptions of the SVI field instrument, of the SVI data processing algorithms, and of SVI data products. The last section summarizes the main features of an SVI.

2. SVI imaging system

a. Field unit

The field unit is a video system inside heated housing, plus a halogen lamp that is located 3 m from the camera. The video system is a camera and a lens. The camera is a Supercircuits PC28C monochrome C-mount camera

* Current affiliation: Department of Atmospheric Science, Colorado State University, Fort Collins, Colorado.

+ Current affiliation: National Center for Atmospheric Research, Boulder, Colorado.

Corresponding author address: Andrew J. Newman, Department of Atmospheric Science, Colorado State University, Fort Collins, CO 80523.
E-mail: anewman@atmos.colostate.edu

DOI: 10.1175/2008JTECHA1148.1

© 2009 American Meteorological Society

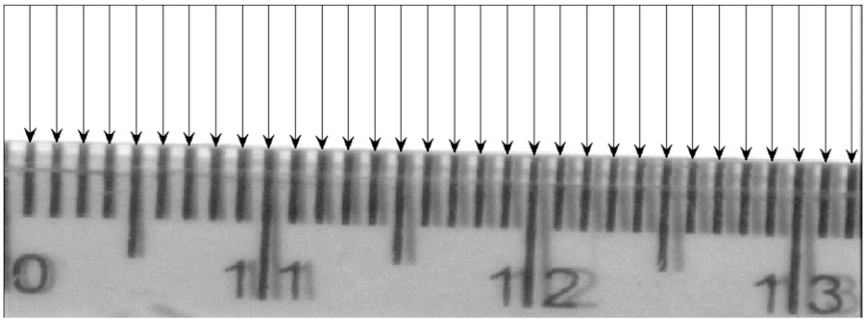

FIG. 1. The SVI calibration procedure uses a ruler in the focal plane and adjustment of the lens. The SVI calibration is straightforward and the images are acceptable for the intended applications.

with a charge-coupled device (CCD) image sensor. The sensor has 640 pixels by 480 pixels, however, we operate it in 640 × 240 noninterlaced mode so that the frame rate is 60 frames per second. A CCD camera was chosen rather than a line scan camera so that images are obtained almost instantaneously and thus no correction is needed for particle movement. The exposure time is 1/100 000 s so blurring due to particle motion is insignificant. A 100–300-mm lens is attached to the camera. During calibration, the video system is adjusted so that the focal plane is 2 m from the end of the lens and the field of view (FOV) is 32 mm by 24 mm, hence the nominal pixel size of an image is 0.05 mm by 0.1 mm. The calibration steps are to (i) set the f-stop to 8, (ii) place a ruler 2 m from the end of the lens, and (iii) adjust the zoom and focus so that the ruler is in sharp focus with the desired FOV.

Figure 1 shows an example of a calibration image. The goal is to have a sharp image with32 mm uniformly spanning the 640 horizontal pixels of the image. This image shows that the overall length is within ~0.5 mm of the desired length. The arrows were generated by software and accounts for the 0- and 32-mm locations on the horizontal axis, as well as the slight tilt of the ruler relative to the major axis of the image. There is adequate correspondence between all the arrows and the ruler ticks, and similar results are obtained for vertical scaling. For the calibration configuration, the bright light source produces shadows of the ruler markings on the wall behind the ruler. The calibration method can be refined by placing a grid in the FOV to create a lookup table that can be used (in software) to correct for errors. Furthermore, distortions can be minimized by using a telecentric lens. For this study, this economical lens and simple calibration procedure produce images and results suitable for the intended applications.

The light source is a halogen flood lamp with a 300-W bulb, which is located approximately 3 m away from the end of the lens. The frosted window on the flood lamp is 50 cm × 80 cm and it diffuses the light to help provide a uniform background.

b. Datalogger

The datalogger uses hardware and software to record images from the field unit. Analog video images (RS170 format) from the camera are routed to the datalogger by a coax cable. The datalogger is a personal computer (PC) with a Windows-based operating system and a National Instruments 1409 video acquisition card, which converts the incoming video stream to digital format. The PC acquires images at a maximum frequency of 60 Hz, with typical operational frequencies of 55–58 Hz. The data rate is too great to effectively store raw images for an entire winter, so we developed a data compression algorithm using LabVIEW. Our acquisition program has several steps. First it adjusts the acquisition card setup so that it utilizes its full 8-bit dynamic range (i.e., the maximum brightness in an 8-bit image of the lamp is nearly 255.) In the absence of particles, the raw images are bright because the FOV only spans the halogen lamp. Yet due to the bulb and reflector design, there is about a factor of 2 in the range of pixel intensity. A more uniform light source could be used; however that would increase costs. To minimize costs, we use software rather than hardware to obtain more uniform background brightness. We wrote an automatic gain control (AGC) subroutine that minimizes the brightness variability. The AGC subroutine produces images of the lamp with brightness levels within a 5% range. To do this, the AGC uses pixel brightness from an average image (from 128 images) to adjust each pixel in subsequent images. When an object is in the measurement volume, it produces dark pixels in the images. The number of snowflake particle pixels is a small fraction of the total number of pixels. Thus we can reduce the data storage significantly by compressing the data. For each

FIG. 2. SVI deployed in North Dakota. The camera is located in the heated housing on the right of the image, while the flood lamp is on the left.

dark pixel, the compression subroutine records the (i) time, (ii) brightness level, and (iii) the pixel location. Although only a small portion of the total data stream is recorded, all the pertinent snowflake data are archived. The threshold value is set to assure no loss of useful data.

c. *Field operation*

An SVI field unit is shown in Fig. 2. During calm winds, an SVI records in situ conditions well because the measurement volume is separated from the instrument and snowflakes fall freely through the measurement volume. Furthermore, most winds have negligible effects on snowflakes within the SVI measurement volume. So sensor induced turbulence is usually not an issue and volume density measurements are reliable. However, when winds blow along the optical axis, airflow around the instrument might cause sampling anomalies. To assess this potential problem, we conducted airflow simulations using the computational fluid dynamics (CFD) software package FLUENT, developed by Fluent, Inc. (Newman and Kucera 2005). Wind blowing parallel to the optical axis is the worst-case scenario for instrument interference, which is what we simulated. Figure 3a displays vertical velocity magnitudes and reveals that nonzero vertical velocities dissipate within approximately 1 m of the end of the camera housing. Because snow has a downward velocity component and the vertical velocity near the focal plane is close to zero, the shielding effect of the camera housing is not amplified. Figure 3b shows ray paths for passive tracers flowing past the SVI. The rays indicate nearly horizontal flow throughout the sample volume. The SVI causes acceleration of the wind around the camera housing, which leads to areas of higher horizontal wind speeds above and below the housing. There is also a region of lower velocity downwind from the housing. Lower wind speeds extend the full length of the instrument, which means that there is some modification of the wind field. Overall, the simulation indicates relatively minor changes to the wind field in the SVI sampling volume.

From the simulations and Fig. 2, we conclude that the SVI housing has minimal effect on its sample volume and the housing creates insignificant interference for most wind directions. In the worst-case scenario of winds parallel to the optical axis, the camera housing induced vertical velocities that dissipated outside of the image volume; hence shadow effects are almost inconsequential. The simulation results help us to interpret the snowpack in Fig. 2, which shows that there is a minor ripple in the snowpack downwind and close to the camera housing. However, in the region around the focal plane, which is 2 m from the camera housing, the snowpack seems to be a regular part of the surrounding snowpack. The SVI in Fig. 2 is orientated to assess system performance for a worst-case setup and it shows that the snowpack under the SVI measurement volume is relatively unperturbed. This is an excellent result

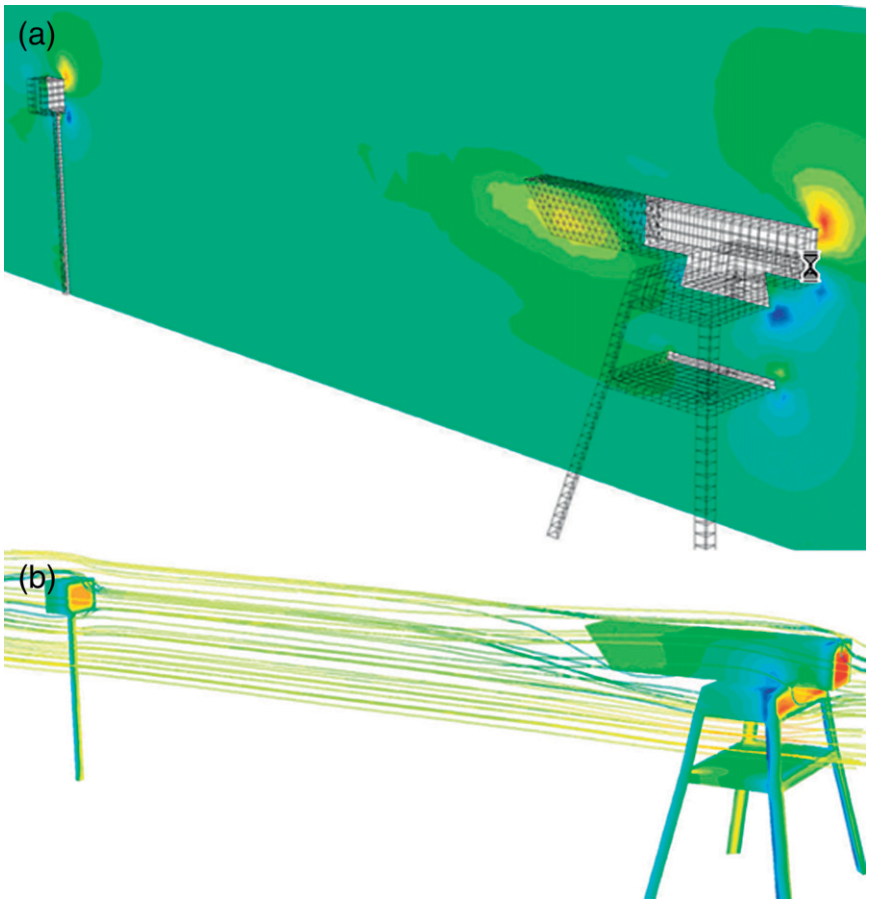

FIG. 3. (a) Cross section displaying vertical velocities. The scale ranges from −2 (dark blue) to 2.5 m s^{-1} (red) with contour intervals of 22 cm s^{-1}. (b) Cross section displaying passive tracers with static pressure contoured on the SVI. The tracers are shaded by velocity magnitude, which ranges from 0 (dark blue) to 4 m s^{-1} (yellow).

because it shows that the SVI functions well in a natural setting, There is no need for snow fences that make an artificially calm wind environment so that snowflakes fall into the measurement volume. SVIs can function with any orientation in a natural setting. However, to reduce the potential of wind effects in datasets, SVIs should be orientated with regard to storm winds (i.e., with the optical axis perpendicular to the climatological mean wind).

The SVI camera is orientated with the 24-mm axis in the vertical direction to minimize the likelihood of particles appearing in multiple images during calm conditions. Yet particles falling slower than 0.72 m s^{-1} may appear in consecutive images. At the North Dakota field site, it was usually windy during frozen precipitation events, so multiple images of particles were rare. Independent sampling is desirable for the computation of distributions. Consequently we recommend that anemometer data be collected in conjunction with SVI data to identify calm wind conditions. If anemometer data are not available, SVI images can be viewed to detect calm conditions or the sampling rate can be reduced during the analysis to ensure independent sampling.

3. SVI data processing software

a. Detection

SVI images are predominately white, with an occasional gray shadowgram of a snowflake. The detection software finds snowflakes in SVI images in the compressed data files and outputs a record of basic information for each snowflake into a summary file. These snowflakes have a myriad of shapes when they fall somewhere between the camera and the lamp. The detection software needs to be robust enough to handle a variety of shapes well and to limit detection to snowflakes that are suitably imaged (i.e., not artifacts of out of focus snowflakes). First, we examine the detection algorithm by assessing its performance with

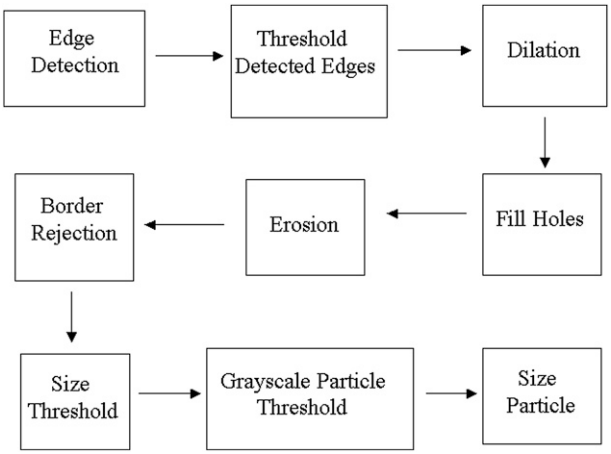

FIG. 4. Flowchart of the SVI particle detection algorithm.

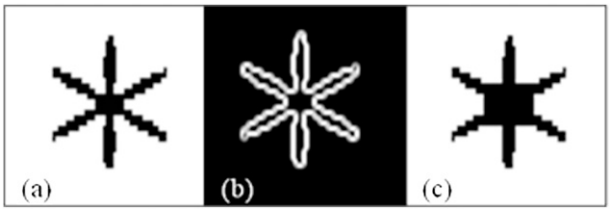

FIG. 5. Processing of a simulated dendrite by the SVI particle detection algorithm: (a) the simulated crystal, (b) after application of the Sobel edge detection kernel, and (c) the final image after the dilations and erosions.

simulated objects that are uniformly black and in focus. Then we discuss implementation of the algorithm with SVI images of natural snow.

The SVI detection algorithm is concisely summarized here. [See Newman (2007) for further details, and to learn about image processing techniques, see texts such as Russ (2002) and Seul et al. (2000).] Figure 4 shows a flowchart of the SVI algorithm. The processing begins with a Sobel edge detection routine (step 1) that detects both vertical and horizontal edges. Next an edge threshold (step 2) is applied. If the threshold selected is too close to white (255), then many snowflake fragments are detected due to snowflakes far from the focal plane. On the other hand, if the threshold value is too close to black (0), then only a few snowflakes very close to the focal plane are detected. We obtained the operational threshold value by examination of simulated snowflakes and natural snowflakes. We optimized the threshold value to yield as many snowflakes as possible, with a minimal number of fragments from blurred snowflakes. Dilation (step 4) is performed in an attempt to fill any holes in the perimeter of the hydrometeor, and then the hole filling subroutine (step 5) darkens the interior of any closed outline of an object in the image. Erosion (step 6) removes extra pixels around the edges of the particle from the dilation step (i.e., returns the maximum dimensions of the object to the correct length). Any particle touching the border is rejected (step 6) because it cannot be known how much of the hydrometeor is missing. Size filtering (step 7) rejects any object with fewer than 30 pixels (approximately 0.3-mm equivalent diameter) because of the great uncertainty for those small particles. Step 8 is used to filter out other processing artifacts and hydrometeors that are very far out of the focal plane. Size and shape information are produced in step 9, then characteristics for each snowflake are output to a summary file.

We validated the detection algorithm using (i) images of water droplets and (ii) images of randomly sized computer-generated objects. Images of water drops (five sizes ranging in diameter from about 2 to 5 mm) were recorded for drops falling close to the focal plane. Analysis of the data shows that for each size, the standard error was about 1%. Likewise, the difference between the SVI equivalent diameters and the drop mass diameters is less than 1%. So the SVI data system and analysis software package provide good results for water droplets near the focal plane. These results validate the acquisition and sizing software for ideal conditions.

To assess other aspects of the SVI sizing software, we generated images of squares, circles, and dendrites. By simulating objects, the exact size is known and direct comparisons to the SVI output can be made for shapes that resemble snowflakes. The metrics we use to determine the accuracy of the SVI algorithm varied depending on the type of object generated. Yet for each metric, the bias and root-mean-square error (RMSE) was computed. One thousand objects per class were used to compute the statistics. Rectangles were simulated using integer multiples of pixel length and the computed bias and RMSE of maximum length are zero, which indicates that the software functions properly. Circles were simulated using real numbers for diameters, and so although the bias is almost zero, the RMSE is 0.02 mm. The small RMSE shows that the inherent uncertainty of pixel occultation is negligible for most applications. We simulated dendrites following Wang and Denzer (1983), and an example is shown in Fig. 5a. The main problem with retrieving some dendritic metrics is the inner region where the arms are near each other. The edge detection subroutine places the edge outside the dark pixels as shown in Fig. 5b, so that when dilation is performed, the arms are expanded outward and some of the space between them is filled (Fig. 5c). Figure 5 indicates that the edge detection algorithm does not produce broken edges for the simulated particles,

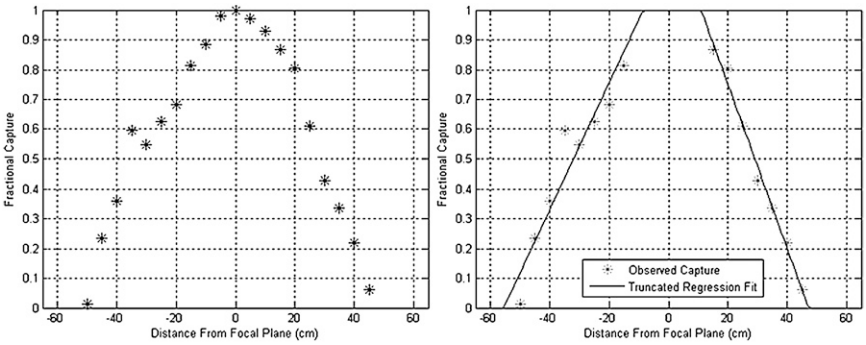

FIG. 6. The F values for an object of 6 mm in maximum length (left) without and (right) with linear fits to the data.

which raises the question of why dilations and erosions are necessary. For our idealized particles, the maximum possible gradient is used (i.e., a white background with perfectly black objects). SVI images of natural snowflakes have various shades of gray resulting in less than optimal gradients, which creates edges of various intensities. When the edge threshold is performed, some parts of the outer edge may be missing, which creates boundary breaks that the dilations and erosions repair. Consequently, analysis of SVI snowflake images benefits from including dilations and erosions in the particle detection algorithm. So if one computes total pixel count or equivalent diameter, there will be a bias from the analysis algorithm. Fortunately, snowflake size is generally characterized by maximum length, so the filling effect is small. In fact, the errors for maximum length estimates are −0.15-mm bias and 0.2-mm RMSE. We attribute the bulk of these errors to quantization effects associated with use of the equations to construct digital images. Overall, these simulations show that the sizing algorithm is robust and suitable for many applications.

b. Volumetric effects

We formulated the SVI depth of field (DOF) as a function of particle size from data obtained during laboratory experiments. DOF is determined by the camera system, the object, and the analysis software. The camera systems are the same for all SVIs, so only the analysis software is examined here. Imaged objects affect the DOF through their optical properties and their size, which contribute through contrast and blurring. Contrast is a function of optical density and ranges from opaque to translucent. Blurring occurs when an imaged point resides outside of the focal plane. When this occurs, the light acquired by the lens from the point will be imaged as a circle on the detector. This circle size depends on the distance from the focal plane. As the distance from the focal plane is increased, the blurring increases as well. With increasing blurring, the contrast of the object imaged decreases. At some distance from the focal plane, the contrast is sufficiently decreased such that the analysis software does not detect the particle. This location is dependent on the image analysis techniques and user-defined parameters. For an in-depth discussion on photography and the issues of blurring, see photography texts and/or an optical physics text [such as *University Physics* by Young and Freedman (2000)].

For analysis of SVI data, we follow the method presented by Frank et al. (1994) that uses laboratory experiments to derive the DOF relationship from

$$\mathrm{DOF} = L(D), \tag{1}$$

where L is the distance from the focal plane and D a length of the object. We used objects with maximum lengths from 2 to 10 mm and we dropped them through the sample volume at varying locations. The fractional capture F is defined as

$$F = C_{\mathrm{act}} \Big/ C_{\mathrm{all}}, \tag{2}$$

where C_{all} is the number of objects captured by the SVI for objects in the focal plane and C_{act} is the actual number of objects captured at the various locations along the optical axis. Figure 6 displays the experimentally determined F values for an object of 6 mm. Linear regression fit lines are also shown. We assume that for the region very close to the focal plane, F is 1 and Fig. 6b shows that a linear fit to the truncated data is reasonable. Next the DOF was determined through discrete trapezoidal integration, given as

$$\mathrm{DOF} = \sum_{i}^{i=N-1} \left\{ \frac{[F(D_i) + F(D_{i+1})]}{2} \Delta D \right\}, \tag{3}$$

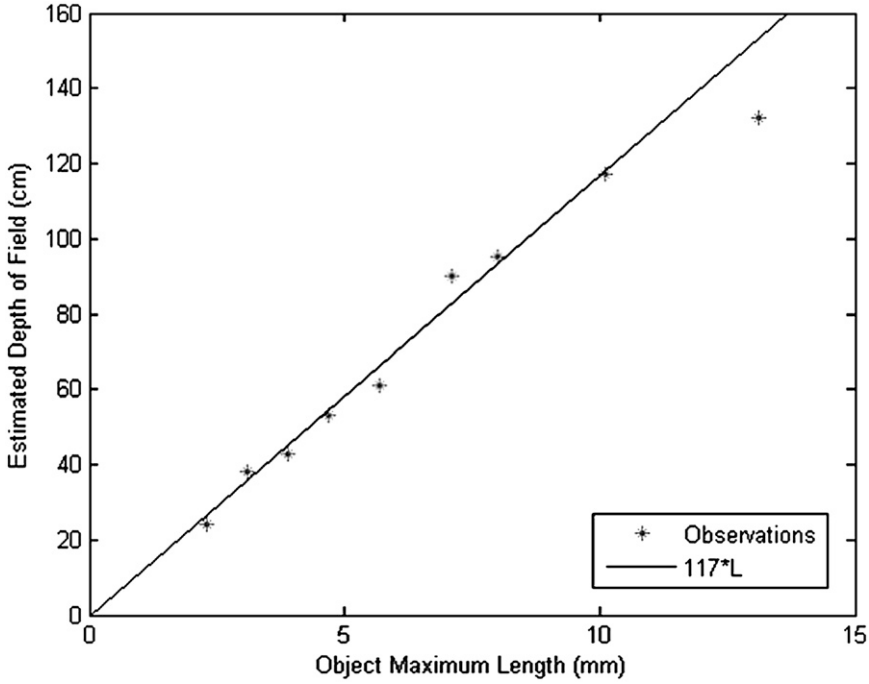

FIG. 7. The DOF values along with a linear DOF relationship plotted to the first eight points.

where N is the total number of drop points, $F(D_i)$ is the fractional capture at the ith point, $F(D_{i+1})$ is the fractional capture at the $i+1$th point, and ΔD is the distance between the two points. At last the DOF relationship for the SVI can be determined from the results from the eight size classes that ranged from 2.3 to 10 mm. Figure 7 shows DOF values for each object size along with a linear fit to data, which indicates that the DOF is approximately 117 times the object maximum length. Newman (2007) estimates the DOF uncertainty to be around 15% for laboratory conditions. It is possible that extreme differences in snowflake densities and shapes would result in a larger DOF uncertainty. Further work is needed to quantify DOF uncertainty in actual measurement conditions. This DOF uncertainty will result in snowflake size distribution (SSD) uncertainty, as can be seen in section 4a.

We used translucent objects to estimate the DOF. Snowflakes have complex optical properties due to their crystalline parts and their aggregate structure, so they have texture that ranges from transparent to translucent. Our laboratory simulations do not account for optical texture, so there is some uncertainty in applying the laboratory DOF formulation to natural conditions. In summary, we have developed an estimate of DOF for the SVI from laboratory experiments, but further research is needed to assess its reliability in natural conditions.

c. *Blurring effects*

To estimate sizing error, a water drop of known size was dropped through the sample volume. Figure 8 shows that the imaged object size generally increases with distance from the focal plan. With regards to SVI data products, this trend is diminished due to fractional capture. We computed the distribution of particles

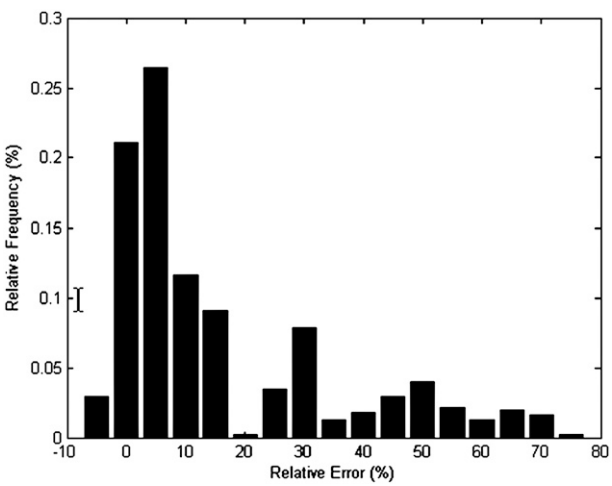

FIG. 8. The L sizing error distribution shows that although relatively large measurement errors do occur, they are unlikely. Consequently, the standard error of L sizing for this SVI hardware and software configuration is ~18%.

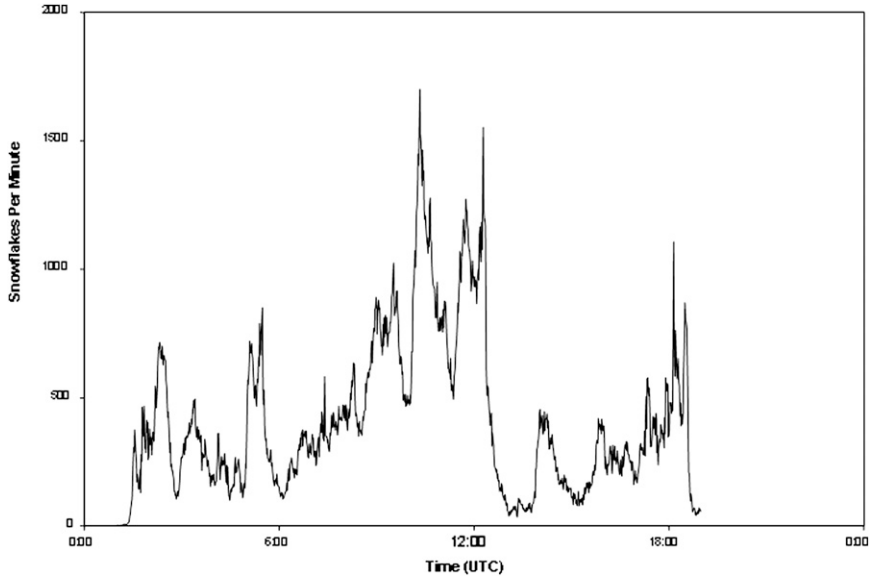

FIG. 9. History of snowflake observations during an 18-h storm observed near Boulder.

sizing errors, shown in Fig. 8, by using fractional capture data. Recall that fractional capture is the ratio of objects the processing algorithm accepts to the total number of objects imaged. Because fractional capture decreases with distance from the focal plane, the sizing error distribution has a somewhat exponential form as shown in Fig. 8. So although relatively large errors are possible, they occur infrequently. Using the sizing error distribution, we estimate the sizing standard error to be 18%, which is suitable for our applications. Note that if we set the threshold in the data analysis to reduce the DOF, the sizing error would decrease and the number of observed snowflakes would decrease. On the other hand, if we set the threshold to increase the DOF, the sizing error would increase and the number of observed snowflakes would increase. We selected a threshold to provide as many snowflakes as possible, with a sizing error that is suitable for routine operations.

4. SVI data products

a. Snowflake size distribution

SVIs record cross-sectional images, so a variety of metrics can be computed to characterize particles. The SVI analysis software provides some representative sizing metrics, such as the bounding rectangle lengths, the equivalent rectangle lengths, and the feret diameter. The bounding rectangle lengths are the maximum x- and y-axis dimensions of a snowflake; the equivalent rectangle lengths are the mean x- and y-axis lengths; and the feret diameter is the maximum length (L) between any two points on the perimeter of the snowflake. A key feature of L is that it is invariant to coordinate axis rotation. Thus, we present snowflake size distributions using L as the metric; L is commonly reported in the literature and the presentation of SVI data products in terms of L lays the groundwork for future studies.

An SVI was located near Boulder, Colorado, during December 2007. About 170 mm of snow fell during an event observed on 11 December 2007. The SVI recorded snowfall from 0100 to 1900 UTC, with the time history of the number of snowflakes per minute shown in Fig. 9. During the event, the SVI detected approximately 446 000 snowflakes; hence the average is about 400 snowflakes per minute. The maximum number of snowflakes per minute was 1701 and that occurred during 1020 UTC. The SVI frame rate is nearly 60 frames per second, so about one in nine frames had a snowflake. Yet during the peak minute, about half of the frames had snowflakes. On the other hand, the hour with the maximum number of snowflakes occurred between 1000 and 1100 UTC, during which ~65 000 snowflakes were detected, or about one in three frames have snowflakes. For this snow storm, the SVI detected almost 70 000 snowflakes for each 25 mm of snow on the ground. We conclude that the SVI hardware and software packages yield a sufficient number of snowflakes for many snowflake sizing studies.

Each segment of the snow size distribution is computed as follows:

$$\mathrm{SSD}_i = N_i \Big/ (V_i dD). \tag{4}$$

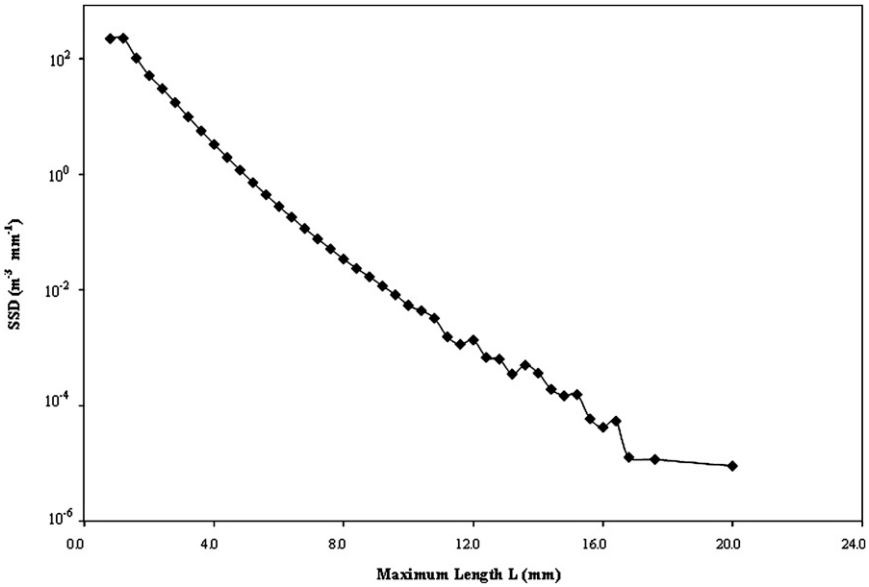

FIG. 10. SSD(L) from a snowstorm observed near Boulder.

SSD_i is the number of particles per unit volume per unit size, so the units of SSD_i are (m^{-3} mm^{-1}); dD is the particle interval size, which is 0.4 mm for 60 bins spanning 0.0–24 mm; N_i is the number of particles within each bin detected by analysis of the SVI data; V_i is the volume corresponding to N_i; N_i and V_i are obtained by counting particles and summing volume increments for the appropriate time interval, which in this study is either 1 min or the duration of the snow event, and V_i is computed as follows:

$$V_i = n_t \mathrm{DOF}_{\mathrm{Lm}} \mathrm{FOV}_{\mathrm{Lm}}, \tag{5}$$

where n_t is the number of frames accumulated by the SVI during the time interval; Lm is the length at the middle of each size interval; and volume is the product of area and length (i.e., depth of field times field of view). We use the laboratory results, so $\mathrm{DOF}_{\mathrm{Lm}}$ (m) = 0.117 Lm (mm). We compute $\mathrm{FOV}_{\mathrm{Lm}}$, the field of view, based upon the 32 mm × 24 mm calibration image, with an adjustment for edge effects, that is,

$$\mathrm{FOV}_i = 10^{-6}(32 - \mathrm{Lm})(24 - \mathrm{Lm}). \tag{6}$$

SSD for the entire snow event is shown in Fig. 10, and it has a typical exponential distribution with an abundance of small particles and relatively few large ones. For this event, the size distribution spans more than seven orders of magnitude. The largest snowflake has the longest axis of 20 mm. The SSD shows that there are about 100 1-mm particles for each 4-mm snowflake. The presence of snowflakes larger than 10 mm implies that there was some period(s) with high aggregation efficiencies.

A time history of snow size distributions for the snow storm observed near Boulder is shown in Fig. 11, which displays the SSD for each minute. For the 1-min records, the size distribution covers less than four decades because the SVI records fewer than 3600 images per minute. However, large snowflakes are present and it is interesting to observe when they occur because that implies calmer winds. Figure 9 shows that there is a spike in the snowflake count after 1200 UTC and that after about 1300 UTC, the snowflake count decreases to a minimum. Figure 11 indicates that the first period has an abundance of small particles and a limited number of aggregates, whereas the second period has fewer small particles and an abundance of large aggregates. This pattern suggests that differing aggregation conditions occurred within a brief period of time, most likely when the wind subsided.

b. *Snowflake visualizations*

An SVI can operate for months at a field site or be set up to obtain data during a storm; either way, thousands of snowflakes are recorded during significant snowstorms. Several snowflakes detected by the particle detection algorithm are illustrated in Fig. 12, which shows that some snowflakes are blurred and others are well focused. Blurring due to particle motion is not an issue because an SVI uses a two-dimensional detector that takes a snapshot of an entire snowflake at a very brief instant in time. On the other hand, line scan cameras

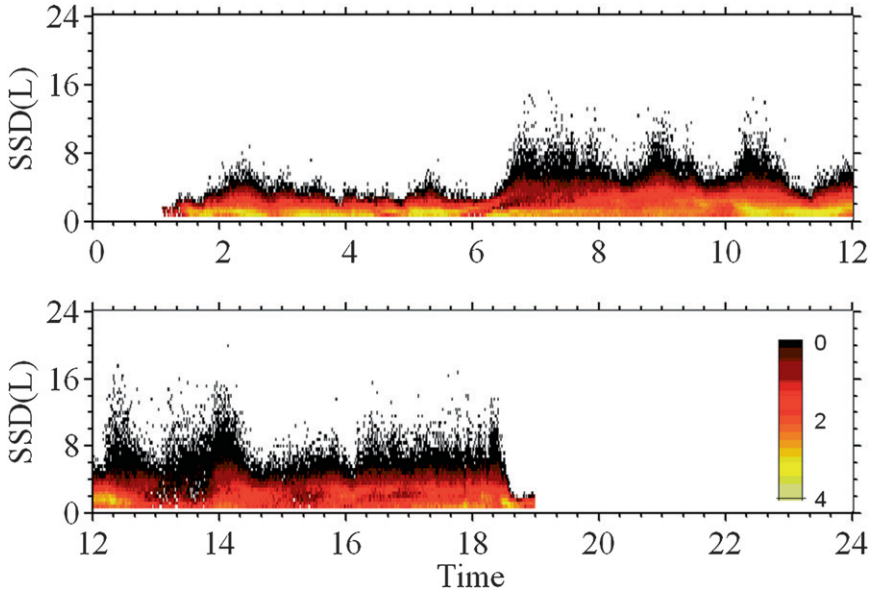

FIG. 11. Minute time history SSD of a snowstorm observed near Boulder. The data span 4 decades of SSD(L) values. Yellow bands show periods of enhanced numbers of small particles. Black towers show periods with large aggregates. Color scale is log[SSD (mm^{-3} mm^{-1})].

produce snowflake images by assembling line scans taken sequentially in time. So for windy conditions, snowflake images from a line scan camera are distorted due to horizontal motion. To compensate for this effect, those images must be altered, which requires fall velocity and horizontal velocity information for each snowflake. On the other hand, compensation for wind is simply not needed for SVI data. However, blurring attributable to the location of a snowflake along the optical axis is an issue. Positional blurring occurs due to the optical phenomenon known as the "circle of confusion." Ideally, a point object in the focal plane produces a point image on the detector. However, if a point object is along the optical axis but not in the focal plan, it casts a spot on the detector rather than a point. The spot perimeter is the circle of confusion. Positional blurring hinders observation of crystal structure in some SVI images. Figure 12 shows that snowflakes far from the focal plane appear as light gray blobs; closer to the focal plane, snowflakes have some structure and are darker; and last, snowflakes near the focal plane reveal crystal structure and holes (if present) and they have well-defined borders with sharp contrast. All the particles in Fig. 12 meet the requirements specified in section 3 for particle sizing studies; however, for snowflake visualizations, the selection process needs to be refined.

For snowflake visualizations, we developed software that quickly selects snowflakes that are usually well focused. SVI images are grayscale and snowflakes closer to the focal plane tend to produce darker images, so there are a number of ways to select particles. For example, to accept or reject a particle, one could use any of the following intensity metrics: maximum, mean, minimum, or standard deviation. Because of the variability of snowflakes, a single parameter selection process will not reject all blurred particles. So if a more robust algorithm is required, a multilevel selection algorithm would be appropriate. We wanted a simple algorithm that works efficiently with the SVI operational particle detection algorithm. The particle detection algorithm has a runtime approximately equal to the observational

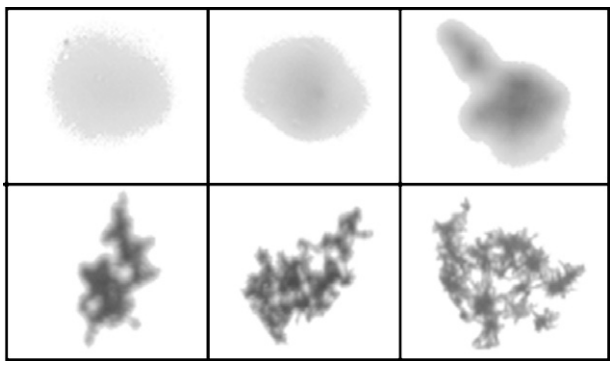

FIG. 12. Snowflake images from an SVI showing various degrees of blurring. The amount of blurring is related to camera and analysis software settings. These results were obtained with settings that provide DOF $\simeq 120 \times$ (particle size), hence the SVI records a large population of snowflakes. Although all these snowflakes are acceptable for sizing studies, the blurring obscures visualization of crystal structure.

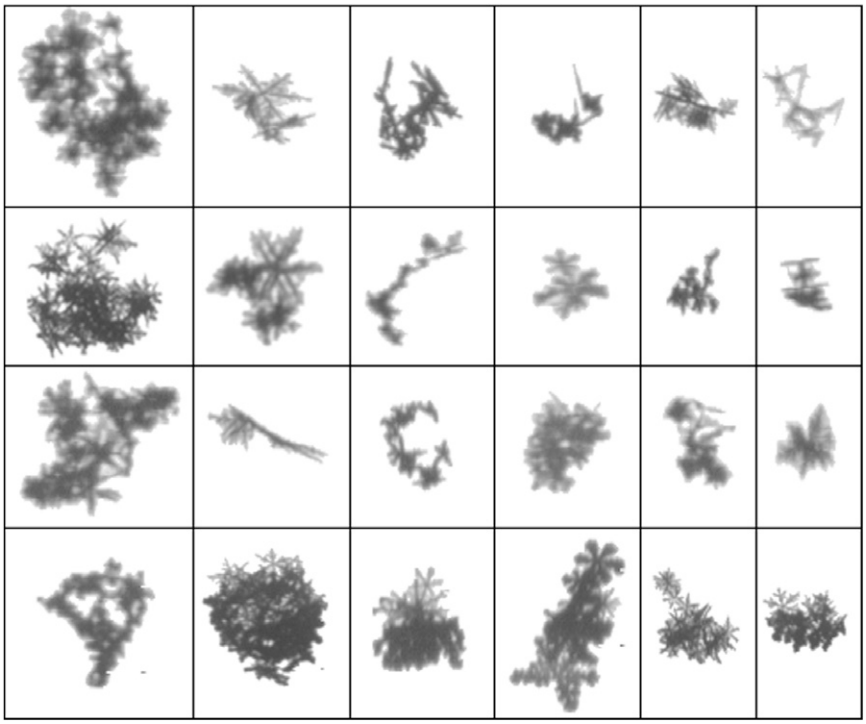

FIG. 13. Aggregate snowflakes selected by SVI visualization software.

time of data. We want the particle selection algorithm to run much faster. So we choose to sieve through the particles from the detection algorithm to find well-focused snowflakes. Our selection process uses two inputs: definition of size classes, and specification of the number of particles for each class. The SVI visualization program (i) sieves the particle file by size class, (ii) sorts size–class particles by minimum intensity, that is, the darkest pixel, and (iii) outputs the desired number of particles in darkest pixel order. The selected particles are sorted by size and displayed on composite images. The number of rows and columns in a composite image depends upon the size of the particles; each composite image has a unique title and file name for easy identification. To facilitate the viewing of composite images, we construct a digital video for each snow day.

To assemble snowflakes shown in Fig. 13, we have selected particles from SVI data recorded during the winter of 2005–06 in North Dakota, as well as the Boulder snowstorm. These aggregates are composed of whole and fragmented snow crystals and the images show that the aggregates have both solid areas and voids (i.e., holes). From a particle analysis perspective, a hole is any contiguous set of pixels within an object that have intensity values outside the selected range for objects; holes means the number of holes inside an object; hole area means the area of holes within an object; and hole ratio means the ratio of the object area excluding holes to the total area of the object, as determined by area/(area + holes area). Software exists that fills holes so that all particles have a hole ratio value of zero, which is useful for sizing particles. The aggregates in Fig. 13 are composed of mainly needles and dendrites.

SVI data also record individual snow crystals, as shown in Fig. 14. Environmental factors affect the development to snow crystals types. Temperature and humidity are especially important and various combinations of these variables result in numerous types of crystals. This complicates interpretation of remote sensing data from frozen precipitation because microwave scattering is highly dependent upon particle size and shape characteristics. So snow crystal classification is likely to be a key factor for developing robust inversion algorithms for space applications. SVI data can contribute to interpretation of remote sensing data because it can be used to classify snow crystals. The International Snowflake Commission defines snow crystal types. Figure 14 lists the crystal types and shows examples of SVI images that range from plates to ice pellets. The variety of snow crystals shows that a broad range of temperature and humidity conditions contributed to the formation of frozen precipitation in North Dakota during the winter of 2005/06. The SVI did not, however, record any hail, which is not unexpected

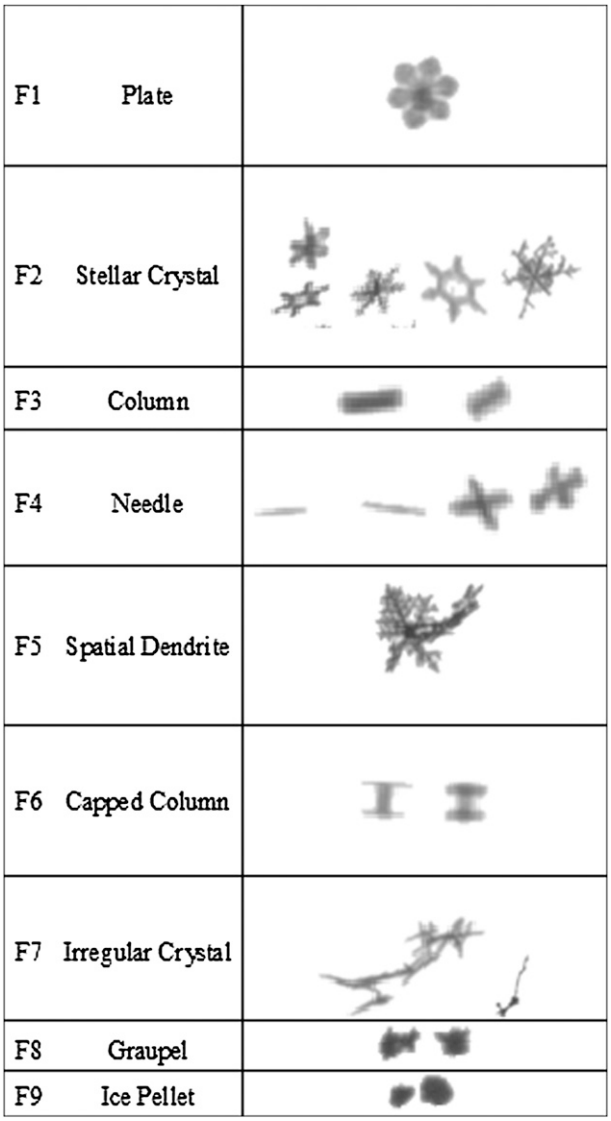

FIG. 14. Snow crystals from SVI images can be classified according to crystal type as defined by the International Snowflake Commission.

because hail is usually observed during warmer months. SVI visualizations are useful for assessing snow crystal type.

5. Summary

An SVI can be installed at a field site in about an hour. Its operation requires minimal maintenance. In fact, several systems have operated unattended for months in harsh winter conditions with winds up to 50 kts and temperatures as low as −40°C. For most wind directions, instrument interference with the sample volume is inconsequential. The worst case is when winds rate is almost 10 megapixels s^{-1}. Our data compression algorithm compresses ~0.5 billion images from a 3-month season to about 50 gigabytes, so an entire season of data fits on an ordinary disc. The optical depth of field is ~120 times the particle size, so for a 5-mm snowflake the DOF is ~0.6 m, and consequently the measurement volume is ~0.0003 m^3, or ~1.1 m^3 min^{-1}. Consequently, a 24-h snowstorm yields several hundred thousand snowflake images. The SVI snowflake visualization software selects well-focused images that we use to classify the snow crystal type as defined by the International Snowflake Classification. The SVI analysis software also produces snowflake size distributions, which are the standard data product for frozen precipitation disdrometers. An SVI obtains images of snowflakes, so the data can be analyzed to compute particle orientation distributions, which are important for interpreting dual-polarization radar data. Thus an SVI system is able to produce classical data products, as well as refined data products.

Intercomparison between the SVI and other disdrometers is beyond the scope of this investigation. In the future, comparisons can be made with (i) the two-dimensional video disdrometer (2DVD), (ii) the particle size and velocity optical disdrometer (PARSIVEL), and (iii) the hydrometeor velocity and shape detector (HVSD) (Kruger and Krajewski 2002; Brandes et al. 2007; Löffler-Mang and Joss 2000; Löffler-Mang and Blahak 2001; Barthazy et al. 2004). Certainly each instrument has its strengths and limitations, so it would be useful to publish comparisons of costs, operations, and data products.

In summary, an SVI system enables (i) measurement of snowflake size distributions, (ii) snowflake classification from grayscale images, and (iii) further analysis as desired because snowflake images are digitally archived. An SVI does not measure fall velocity; nevertheless snowflake fall velocity could be obtained by refinement of the camera and lighting systems.

Acknowledgments. This work was supported by NASA Headquarters under the NASA Earth and Space Science Fellowship Program, Grant NNG04GQ08H. The authors would like to thank Dr. Forrest Ames for help with the Fluent simulations and the three anonymous reviewers who gave valuable input to this manuscript.

REFERENCES

Barthazy, E., S. Göke, R. Schefold, and D. Högl, 2004: An optical array instrument for shape and fall velocity measurements of hydrometeors. *J. Atmos. Oceanic Technol.,* **21,** 1400–1416.

Brandes, E. A., K. Ikeda, G. Zhang, M. Schönhuber, and R. Rasmussen, 2007: A statistical and physical description of hydrometeor distributions in Colorado snow storms using a video disdrometer. *J. Appl. Meteor. Climatol.,* **46,** 634–650.

Bringi, V. N., and V. Chandrasekar, 2001: *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press, 636 pp.

Frank, G., T. Härtl, and J. Tschiersch, 1994: The pluviospectrometer: Classification of falling hydrometeors via digital image processing. *Atmos. Res.,* **34,** 367–378.

Kruger, A., and W. F. Krajewski, 2002: Two-dimensional video disdrometer: A description. *J. Atmos. Oceanic Technol.,* **19,** 602–617.

Liu, G., 2004: Approximation of single scattering properties of ice and snow particles for high microwave frequencies. *J. Atmos. Sci.,* **61,** 2441–2456.

Löffler-Mang, M., and J. Joss, 2000: An optical disdrometer for measuring size and velocity of hydrometeors. *J. Atmos. Oceanic Technol.,* **17,** 130–139.

——, and U. Blahak, 2001: Estimation of the equivalent radar reflectivity factor from measured snow size spectra. *J. Appl. Meteor.,* **40,** 843–849.

Matrosov, S. Y., R. F. Reinking, R. A. Kropfli, and B. W. Bartram, 1996: Estimation of ice hydrometeor types and shapes from radar polarization measurements. *J. Atmos. Oceanic Technol.,* **13,** 85–96.

Meirold-Mautner, I., C. Prigent, E. Defer, J. R. Pardo, J. P. Chaboureau, J. P. Pinty, M. Mech, and S. Crewell, 2007: Radiative transfer simulations using mesoscale cloud model outputs: Comparisons with passive microwave and infrared satellite observations for midlatitudes. *J. Atmos. Sci.,* **64,** 1550–1568.

Newman, A. J., 2007: Surface and vertical retrievals of snowfall using a video disdrometer and a 915 MHz vertical profiler. M.S. thesis, University of North Dakota, 147 pp.

——, and P. A. Kucera, 2005: Gauging rainfall. *Fluent News Fall Newsletter,* No. 14. The Fluent Group, Lebanon, NH, 12–13.

Russ, J. C., 2002: *The Image Processing Handbook*. 4th ed. CRC Press, 732 pp.

Seul, M., L. O'Gorman, and M. J. Sammon, 2000: *Practical Algorithms for Image Analysis: Description, Examples, and Code*. Cambridge University Press, 295 pp.

Wang, P. K., and S. M. Denzer, 1983: Mathematical description of the shape of plane hexagonal snow crystals. *J. Atmos. Sci.,* **40,** 1024–1028.

Young, H. D., and R. A. Freedman, 1999: *Sears and Zemansky's: University Physics.* 10th ed. Addison-Wesley, 1274 pp.