A Comparative Evaluation of Snowflake Particle Size and Shape Estimation Techniques used by the Precipitation Imaging Package (PIP), Multi-Angle Snowflake Camera (MASC), and Two-Dimensional Video Disdrometer (2DVD)

Charles N. Helms[1,2], S. Joseph Munchak[1], Ali Tokay[1,3], and Claire Pettersen[4]

[1]Mesoscale Atmospheric Processes Laboratory, NASA Goddard Space Flight Center, Greenbelt, MD, USA
[2]NASA Postdoctoral Program—Universities Space Research Association, Columbia, MD, USA
[3]Joint Center for Earth Systems Technology, University of Maryland, Baltimore County, Baltimore, MD, USA
[4]Climate and Space Sciences and Engineering Department, University of Michigan, Ann Arbor, MI, USA

Correspondence: Charles N. Helms (charles.n.helms@nasa.gov)

Abstract. Measurements of snowflake particle size and shape are important for studying the snow microphysics. While a number of instruments exist that are designed to measure these important parameters, this study focuses on the measurement techniques of three digital video disdrometers: the Precipitation Imaging Package (PIP), the Multi-Angle Snowflake Camera (MASC) and the Two-Dimensional Video Disdrometer (2DVD). To gain a better understanding of the relative strengths and weaknesses of these instruments and to provide a foundation upon which comparisons can be made between studies using data from different instruments, we perform a comparative analysis of the measurement algorithms employed by each of the three instruments by applying the algorithms to snowflake images captured by PIP during the ICEP-POP 2018 field campaign.

Our analysis primarily focuses on the measurements of area, equivalent diameter, and aspect ratio. Our findings indicate that area and equi-area diameter measurements using the 2DVD camera setup ~~should be the most accurate, followed by MASC, which is slightly more accurate than PIP. In terms of the precision of the area and equi-area diameter measurements, however, MASC is considerably more precise than PIP or 2DVD, which provide similar precision once the effects of the PIP image compression algorithm are taken into account~~ have the greatest accuracy of the three instruments due to its resilience to motion blurring effects. Both PIP and MASC use shape-fitting algorithms to measure aspect ratio. While our analysis of the MASC aspect ratio suggests that the measurements are reliable, our findings indicate that both the ellipse and rectangle aspect ratios produced by PIP under-performed considerably due to the shortcomings of the PIP shape-fitting techniques. That said, we also demonstrate that reliable measurements of aspect ratio can be retrieved from PIP by reprocessing the PIP images using either the MASC ~~shape-fitting technique~~ ellipse-fitting algorithm or a tensor-based ellipse-fitting ~~technique~~ algorithm. Because of differences in instrument design, 2DVD produces measurements of particle horizontal and vertical extent rather than length and width. Furthermore, the 2DVD measurements of particle horizontal extent can be contaminated by horizontal particle motion. Our findings indicate that, although the correction technique used to remove the horizontal motion contamination performs remarkably well with snowflakes despite being designed for use with rain drops, the 2DVD measurements of particle horizontal extent are potentially unreliable.

1 Introduction

Digital video disdrometers have been in use since the 1990's and there exist several instruments designed around using digital cameras to observe precipitation properties. The measurements from these instruments are not only critical for understanding precipitation microphysics but are also critical for developing algorithms to retrieve microphysical properties using radar and passive microwave sensor data. Due to the variety of instruments available, issues can arise when making comparisons between studies and algorithms that are based on data from different instruments. To this end, the present study will perform a comparison of the particle area, equivalent diameter, and aspect ratio measurements produced by three common ground-based digital video disdrometers: the Precipitation Imaging Package (PIP; Pettersen et al., 2020, 2021), the Multi-Angle Snowflake Camera (MASC; Garrett et al., 2012), and the Two-Dimensional Video Disdrometer (2DVD; Kruger and Krajewski, 2002; Schönhuber et al., 2007, 2008).

Studies invoking snowflake particle size frequently rely on either the maximum dimension length of a particle (e.g., Locatelli and Hobbs, 1974; Heymsfield and Westbrook, 2010), typically defined as either the major-axis length of a fitted ellipse or as the diameter of the smallest circumscribing circle, or on the equivalent diameter (e.g., Han and Braun, 2021; Pettersen et al., 2020, 2021), defined as the diameter of a circle or sphere with equal projected area (for the equi-area diameter) or volume (for the equi-volume diameter) to a particle. 2DVD combines orthogonal measurements of a particle to estimate the particle volume, from which the equi-volume diameter is computed while both PIP and MASC compute the equi-area diameter from their respective area measurements, both of which are determined by counting the number of pixels within the particle. That said, there have been efforts to use the multiple cameras mounted by MASC to compute particle volume (Leinonen et al., 2021).

~~There are numerous examples of studies which rely heavily on either of these measures of particle size. For example, Han and Braun (2021)~~In the case of Han and Braun (2021), the authors characterize the global three-dimensional distribution of precipitation mean particle sizes using satellite radar data; they used a form of the equivalent diameter as their metric of particle size, specifically the mass-weighted mean liquid-water-equivalent diameter. As with many studies that invoke particle size distributions, their choice of metric came down to the choice of metric used in the original particle size distribution parameterization. For Han and Braun (2021), the authors used data produced following Grecu et al. (2016), which used the particle size distribution parameterization described in Testud et al. (2001). Testud et al. (2001) chose to base their parameterization on equivalent diameter, specifically the volume-weighted mean liquid-water-equivalent diameter, for two key reasons: the authors argue that this variable has less physical ambiguity than other metrics and is relatively easy to compute.

The equivalent diameter measurements have also been used to determine a number of bulk snowflake characteristics. Pettersen et al. (2020) used PIP-observed equivalent diameter measurements to compare a number of microphysical and bulk characteristics of snow, including the liquid-water-equivalent snow rate and the bulk effective density of the snow particles, to collocated observations. Their study found that the liquid-water-equivalent snow rate compared very well to both established retrievals as well as observations. Using these bulk characteristics, Pettersen et al. (2021) was able to discriminate between different precipitation phases using the PIP observations.

Empirical relationships between snowflake or ice crystal size and terminal fall speed frequently use the maximum dimension of the particle as their independent variable (e.g., Locatelli and Hobbs, 1974). The more advanced relationships also incorporate the Reynolds number, which is dependent on the particle area exposed to the flow via the Best number; in some cases this area is replaced by an area ratio that is defined as the ratio between the area exposed to the flow and the area of the circumscribing circle, which is a function of the particle maximum dimension (Heymsfield and Westbrook, 2010). In this way, the terminal fall speed relationships can depend on both the equivalent diameter, which is directly tied to particle area, and the maximum dimension of the particle. Similarly, combinations of equivalent diameter and maximum dimension have also been used to estimate particle mass (e.g., von Lerber et al., 2017).

Snowflake shape , separate from the snowflake size, is generally characterized by the aspect ratio of the snowflakes. Korolev and Isaac (2003) found that the mean aspect ratio of in-cloud ice particles varied as a function of the particle maximum dimension for particles smaller than ∼65 μm whereas the aspect ratio of larger particles (i.e., those typical of snow) tended to vary inversely with temperature; the mean aspect ratios observed by the study were between 0.6 and 0.8 with the snowflake-sized particles tending to have mean aspect ratios closer to 0.6. To constrain radar backscatter calculations, aspect ratio is frequently prescribed, often with a mean value of 0.6 assumed (e.g., Matrosov et al., 2005), although a number of studies have performed backscatter calculations using a range of prescribed aspect ratios (e.g., Tyynelä et al., 2011; Westbrook, 2014). The reverse of this process has also been attempted: ? Munchak et al. (2022) used an optimal estimation technique to retrieve a number of snow microphysical properties, including aspect ratio, from a dual-frequency dual-polarization ground-based radar, albeit with mixed results.

As will be discussed in greater detail below, both PIP and MASC produce their measurements by fitting simple two-dimensional shapes (ellipses and rectangles, specifically) to two-dimensional projections of the three-dimensional snowflakes that the instruments are observing. Because snowflakes come in a large variety of shapes, especially when taking aggregate snowflakes into consideration, any attempt to use a simple shape, such as an ellipse or rectangle, to represent these particles suffers from the inherent limitation of under-representing the complexity of the snowflakes. Furthermore, the use of two-dimensional shapes to represent three-dimensional snowflakes adds an additional layer of limitations revolving around the degree to which the two-dimensional projection accurately represents the dimensions and orientations of the three-dimensional particle (e.g., Jiang et al., 2017). Despite these shortcomings, the measurement of snowflake dimensions based on shape fitting has proven to be a useful tool for studying snow microphysics and understanding the relative capabilities of these measurements is critical to their successful use in research and applications.

The goal of the present study is to evaluate the techniques used by each of the three ground-based digital video disdrometers to determine the size and shape of precipitation particles in terms of the area, equivalent diameter, and aspect ratio. In doing so, we will provide a point of comparison for these three sensors. The next section will briefly discuss the data sets used for this study. In section 3, we will provide a detailed description of each sensor and their respective processing algorithms as they relate to the particle shape. This will be followed by a description of our analysis methods in section 4. Section 5 will present the results of our study. Section 6 will compare the results of the present study to the findings of previous studies. Finally, section 7 will present our conclusions alongside a brief summary of the study.
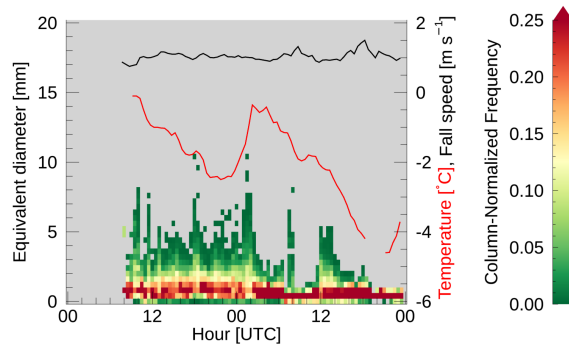
Figure 1. Time series of (shading) particle size distribution (in terms of equivalent diameter [mm]), (black) mean particle fall speed [m s^{-1}], and (red) collocated air temperature [°C] at the MHS site on 7–8 March 2018. Equivalent diameter and fall speed measurements come from the PIP data and the air temperature was measured using a Vaisala WXT520. The particle size distribution time series histogram has been normalized in the y direction.

2 Data

The data sets used to perform the analyses central to this study were collected during the 2018 International Collaborative Experiment for PyeongChang Olympic and Paralympics (ICE-POP; Lee and Kim, 2019; Lim et al., 2020; Tapiador et al., 2021). ICE-POP, which took place between 1 November 2017 and 17 March 2018, was an international collaboration between a number of programs and agencies, including the NASA Global Precipitation Mission (GPM) Ground Validation program and the Korean Meteorological Administration (KMA). The field campaign was designed to study winter precipitation in complex terrain and involved a number of measurement sites featuring a mixture of in-situ and remote-sensing instrumentation. Of interest to the present study is the May Hills supersite (MHS; 37.6652°N, 128.6996°E, 789 m) as this site has collocated PIP, MASC, and 2DVD instruments and experienced frequent snow events. This study will focus on data collected during a snow event that took place on 7–8 March 2018. This event was selected as it contains both aggregate snow particles, which provide a large variety of shapes (and, therefore, aspect ratios), and lump graupel particles, which tend to be relatively spherical with high aspect ratios. Although the general presence of these habits were primarily identified through visual inspection of the PIP data, further support of their presence can be found by examining the time series of the particle size distribution, particle fall speed, and air temperature, as depicted in Fig. 1. Periods when aggregates are present can be identified by the larger equivalent diameters and slower fall speeds and periods with lump graupel can be identified by the smaller equivalent diameters and faster fall speeds; the lump graupel can be discerned from liquid precipitation based on the below freezing temperatures, which extend over a deep layer according to a nearby thermodynamic sounding (not shown). Data collected during a snow event on 9 January 2018 and 22 January 2018 were also examined, but are not included here as their inclusion did not produce any notable changes in the results.

3 Instruments

This study will evaluate the measurement techniques of three video disdrometers: PIP, MASC, and 2DVD. The instruments and their processing algorithms, as they relate to measuring the shape and orientation of snow particles, will be covered in this section.

3.1 Precipitation Imaging Package (PIP)

PIP (Pettersen et al., 2020, 2021) is a video disdrometer developed by NASA as a successor to the Snowflake Video Imager (SVI; Newman et al., 2009). PIP is composed of two parts: a high-speed video camera and a light source to backlight the precipitation particles that pass through the open sampling volume. The camera, of the 'charge-coupled device' (CCD) variety and located 2 m from the light source, continuously records images at a rate of 380 frames per second as 640 by 480 pixel images with a ~~resolution~~ pixel size of 0.1 mm by 0.1 mm using a focal plane located 1.33 m away from the camera. In order to reduce the data rate produced by the high speed camera, the PIP images are compressed before being processed, resulting in an effective ~~resolution~~ pixel size of 0.1 mm in the horizontal by 0.2 mm in the vertical. The present study uses the ICE-POP PIP data produced using version 1403 of the PIP processing code.

The PIP processing algorithms are written using the IMAQ Visual software package (National Instruments, 2000), a software package designed for use in performing automated quality control of manufacturing processes. Of particularly interest to the present study are the methods the algorithm uses to compute the dimensions of the precipitation particle. While particle area is determined by counting the number of pixels within the particle, the particle dimensions are determined by fitting an ellipse or a rectangle to the particle as viewed by PIP, which sees the two-dimensional projection of the particle, after the PIP image has been compressed to reduce data rates. The IMAQ software package performs both the ellipse and rectangle fitting by constructing an ellipse or rectangle that has an equal area and an equal perimeter to the target particle; a derivation of the equations used to make these fits are presented in Appendix A to supplement the limited information given in the IMAQ Vision manual. The IMAQ Vision software package computes the particle perimeter by subsampling the boundary points to produce a smoother representation of the perimeter; for this purpose, the boundary points are located at the corners of the pixels that make up the particle perimeter. The long and short axes of the fitted shapes are then used as the length and width of the particle. Specifically, the ellipse fit produces the variables 'Ellip_Maj' and 'Ellip_Min' for the length and width, respectively, while the rectangle fit produces the variables 'Rec_BS' and 'Rec_SS' for the length and width, respectively. Note, in both the ellipse and rectangle fits, the perimeter and area are defined such that any holes in the particle are not taken into account. Because of this, the PIP-determined area and perimeter of a particle with no holes will be the same as an otherwise-identical particle with a hole when determining the shape fits.

3.2 Multi-Angle Snowflake Camera (MASC)

MASC (Garrett et al., 2012) is another video disdrometer and was developed at the University of Utah; the setup described below is that used during the ICE-POP field campaign. Unlike PIP, MASC employs three cameras and only records images

when triggered by a particle falling through the field of view of a pair of near-infrared sensors. The near-infrared sensors are also used to trigger a flash, which illuminates the particle from the side facing the cameras. The three cameras, also of the CCD variety and separated from one another by $36°$, are mounted looking into an enclosed sampling volume, open only at the top and bottom, and have a common focal point 10 cm away from each lens. Typically, the camera system is artificially restricted to only being triggered once per second in order to reduce excessive flashes. The sampling volume is defined by the intersection of the 35-mm fields of view and the 10-mm depths of field of the three cameras. While the cameras have a pixel size of 9 μm by 9 μm, the system is only triggered when particles with a maximum dimension of greater than 0.1 mm are detected by the near-infrared sensors.

Although the MASC processing algorithm was originally written in MATLAB, a version written in Python also exists and it is this version of the ellipse fitting method that will be analyzed here (Shkurko et al., 2016). To simplify matching particles between the three cameras, MASC only processes images that contain a single precipitation particle and, of these, only particles that are in focus are actually analyzed. The method used to determine whether a particle is in focus is detailed in Garrett et al. (2012). Similar to PIP, MASC also fits an ellipse to the precipitation particles to determine the particle shape properties. In the Python version of the MASC algorithm, this is accomplished using the fitEllipse function within the OpenCV package (https://docs.opencv.org/3.4/, accessed 15 December 2021), which uses the 'LIN' algebraic distance ellipse-fitting algorithm described in (Fitzgibbon and Fisher, 1995) and returns both the major and minor axis lengths as well as the orientation angle of the fitted ellipse.

3.3 Two-Dimensional Video Disdrometer (2DVD)

2DVD (Kruger and Krajewski, 2002; Schönhuber et al., 2007) uses two orthogonally oriented horizontal line-scan cameras to take measurements of backlit precipitation particles. Since the viewing planes of the two cameras are separated by 6–7 mm, the particle fall speed can be measured via the difference in the arrival time of the particle in each viewing plane. These horizontally scanning line-scan cameras capture a series of one-dimensional images of the particle as it passes through the viewing plane. Each camera captures a one dimensional image every 18μs and these one-dimensional images are then pieced together by stacking each scan on top of one another with a thickness dependent on the observed fall speed to produce a two-dimensional image of the precipitation particle for each camera. ~~Furthermore, since the viewing planes of the two cameras are separated by 6–7 mm, the particle fall speed can be measured via the difference in the arrival time of the particle in each viewing plane~~Each of the 2DVD line-scan cameras has 512 photodetectors that are calibrated to observe a 100-mm wide field of view (Kruger and Krajewski, 2002), from which we can infer that the pixel size is $\sim$0.195 mm (195 μm). Although 2DVD was designed to study rain drops, the instrument has been used to study snow as well (e.g., Huang et al., 2010).

Because the 2DVD cameras do not capture the entire shape of the particle at once, consideration must be given for how the horizontal motion impacts the perceived particle shape. For a particle with no horizontal motion, ~~the result of~~ piecing together the one-dimensional images will result in a particle image that depicts a close approximation of the actual particle. If the particle is moving horizontally, however, then the particle will appear to have been shifted horizontally in each subsequent one-dimensional image. The end result of imaging a horizontally moving particle is that the particle will appear to have undergone
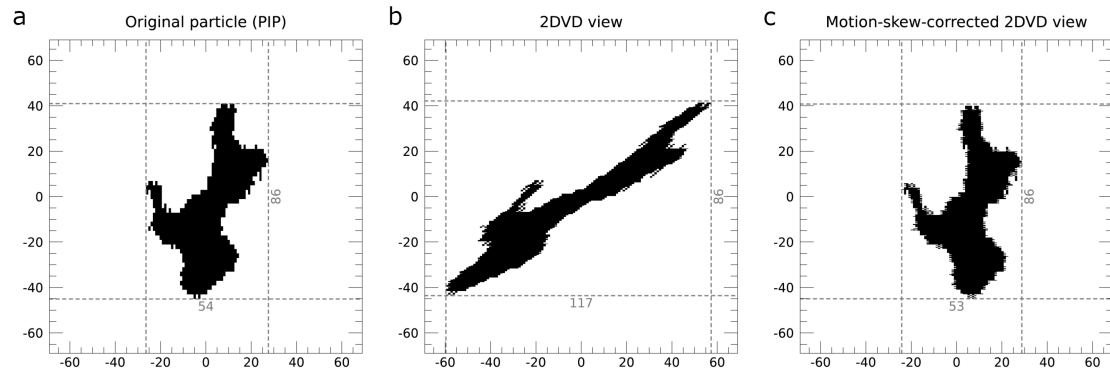
Figure 2. An example of the emulated 2DVD measurements and our unskewing algorithm including (a) the original particle as captured by the PIP camera, (b) the emulated 2DVD view of the particle before the skew correction, and (c) the emulated 2DVD view of the particle after applying our skew correction. The coordinates are in units of PIP pixels from the particle center. The gray dashed lines indicate the extent of the bounding boxes for each view of the particle, with the length of each bounding box edge marked on the bottom and right edges in units of PIP pixels.

a shear transformation relative to the actual particle. Figure 2a–b depicts an example of this shearing effect by sampling a PIP image as though a 2DVD instrument were measuring the particle (see section 4.2).

180 The image skewing produced by the horizontal movement of particles, as evident in Fig. 2b, can invalidate certain 2DVD measurements. While an algorithm designed to remove these effects exists, the algorithm is, unfortunately, in the form of a proprietary software package and, as such, its exact inner workings are unclear. That said, Kruger and Krajewski (2002) state that the algorithm works by assuming the top-most and bottom-most pixels of the particle should be vertically aligned with one another. The algorithm then skews the image in the opposite direction to a sufficient degree to place the top- and bottom-most pixels into vertical alignment. An approximation of the unskewing algorithm has been applied to the particle in Fig. 2c; this

185 approximation of the unskewing algorithm is described in section 4.2.

Unlike MASC and PIP, the equivalent diameter that is actually recorded in the 2DVD data sets is an equi-volume diameter. The details of how this equi-volume diameter is computed from the raw 2DVD data can be found in Schönhuber et al. (2008). Briefly, the calculation works by assuming the shape observed in each pair of scan lines, captured by the orthogonal pair of line-scan cameras, is that of an elliptical cylinder whose major and minor axes are entirely captured by the pair of line-scan

190 cameras. The volumes of these elliptical cylinders are then summed to determine the total volume of the particle. While this method is reasonable for approximately ellipsoidal particles (e.g., raindrops or hailstones), the complex surface structures present in snowflakes and snowflake aggregates can invalidate the volume calculations.

4   Method

The ideal method for comparing the measurements made of snowflakes by each of these three instruments would be to have all
three instruments measure the same particles simultaneously. Due to the construction of each of these instruments, however,
this is not physically possible. Instead, the algorithms used by each instrument to determine particle shape will be analyzed
by emulating each instrument's analysis algorithms using PIP-recorded images as an input. PIP images are used as the basis
of the measurement emulation rather than MASC images because the PIP captures considerably more particles in its sampling
volume than does MASC.

The PIP images used for this study come in the form of Audio Video Interleave (AVI) video files. PIP produces these video
files every 10 minutes and each video file contains the first 2000 frames in which PIP identified a particle during that 10 minute
period. If fewer than 2000 frames contained particles during a given 10 minute period, the video file will contain fewer than
2000 frames. A description of how we emulate the MASC and 2DVD observations is provided in sections 4.1 and 4.2 and an
overview of the method by which we process the PIP AVI video files in order to apply our emulated algorithms is presented in
Appendix B. Additionally, we will test and evaluate an alternative ellipse-fitting method, which employs a mass tensor and is
hereafter referred to as the tensor method, for retrieving particle length and width from the PIP images; this method is described
in section 4.3.

Ideally, we would include all PIP images from the 7–8 March 2018 period in our study rather than just those images that
are included in the AVI files. Since the raw PIP image files are archived and contain every frame with precipitation, this is
theoretically possible. That said, due to the very high data rates involved with the PIP data, the processing code would need to
be highly optimized or use specialized software packages in order to process even a couple days of the raw data in a reasonable
time frame. Thankfully, the AVI files contain a sufficiently large collection of precipitation particles to enable us to perform
our study.

For the emulation of both MASC and 2DVD as well as the ~~tensor ellipse-fitting~~ implementation of the tensor method, the
particle image pixels have values of either zero (non-particle pixel) or one (particle pixel). Any pixel that does not have the
background value (i.e., white) is included as part of the particle. Particles in the AVI files are matched to the PIP measurements
by using a region grow approach, whereby a coherent region is grown from a single point by checking its eight neighboring
pixels for inclusion in the region and then repeating the process with each newly included pixel. The region grow approach is
initially attempted at the PIP-determined particle centroid. For some particles with either large concave regions or holes, the
PIP-measured particle center can correspond to a non-particle pixel. In these instances, an attempt is made to locate a new
starting pixel for the region grow technique by searching the following (x,y) coordinate offsets in the given order: (-1,0), (1,0),
(0,-1), (-1,-1), (1,-1), (0,1), (-1,1), (1,1). If a particle pixel is not located at any of these center offset positions, the process is
repeated after multiplying the offsets by a factor of two; in cases where a particle pixel remains elusive, the process is repeated
with multiples of three and then four. If no particle pixels are located after checking these 32 offset positions, the particle is
considered invalid and is ignored. Additionally, only particles that appear both in the particle tables (i.e., the PIP files ending in

'_a_p.dat') and in one of the two velocity tables (i.e., the PIP files ending in either '_a_v_1.dat' or '_a_v_2.dat') are considered here as the variables contained in these files are necessary for performing the emulations.

4.1 MASC emulation

Our emulation of the MASC processing algorithm is based on the algorithm used for the U.S. Department of Energy Atmospheric Radiation Measurement (ARM) program MASC instruments (Stuefer and Bailey, 2016; Shkurko et al., 2016). Because both the PIP and MASC ultimately ~~make~~ perform their measurements using two-dimensional images, no special processing is performed to prepare the images before ~~the emulated MASC measurementsare made~~emulating the MASC measurements. The emulation of the actual measurements ~~will be~~ is performed by using the same ellipse fitting function as is used by the ARM MASC instrument algorithm, specifically the fitEllipse function from the OpenCV Python package (see section 3.2). Although the MASC cameras have a considerably higher resolution than the PIP camera, it is worth noting that the effects of the camera resolution difference are expected to be negligible in terms of shape fitting. One exception to this is that the fitEllipse function effectively requires at least five pixels within the target particle; for the actual MASC measurements, this represents a particle far smaller than the threshold minimum measured size of 0.1 mm. In the case of the emulated MASC, a five ~~pixel~~ PIP-pixel particle would have an area of 0.5 mm^2 and a maximum dimension of at least 0.3 mm as the PIP pixels are calibrated to be 0.1 mm of each side. To avoid having particles that can only be analyzed by a subset of the measurement algorithms, our comparisons here will only consider particles which have an equivalent area of at least 0.5 mm^2.

4.2 2DVD emulation

The 2DVD instrument captures a series of one-dimensional images as a particle falls through the observing volume. As such, the PIP images must be resampled in a way that mimics the 2DVD instrument itself. To this end, we use snapshots of the PIP-observed particles and replicate the horizontal and vertical motions of the particles, as measured by PIP, using bilinear interpolation. The vertical motion is replicated by shifting the vertical coordinates of the bilinear interpolation upward by an amount equal to the ~~product of the~~ particle fall speed ~~and~~ divided by the camera observation frequency; for 2DVD, the camera observation frequency is 68200 Hz. The horizontal motion is replicated by shifting the particle snapshot pixel locations horizontally in the direction of particle motion by an amount equal to the distance traveled by the particle during the elapsed time of the emulated observations, to the nearest 0.1 mm. This elapsed time is equal to the product of the inverse of the camera observation frequency and the number of times the particle has been sampled in the vertical up until that point (i.e., the number of emulated 2DVD line scans performed). When performing the interpolations, only the interpolated pixels with a value of one are considered part of the emulated 2DVD observation.

An example of an emulated 2DVD observation is depicted in Fig. 2 with the original, PIP-observed, particle image depicted in Fig. 2a and the emulated 2DVD particle image depicted in Fig. 2b. Note the considerable skewing that has occurred in the emulated 2DVD image relative to the PIP image due to the horizontal movement of the particle towards the left as it passed through the virtual line scan camera. As previously mentioned, 2DVD applies an unskewing algorithm to correct for the distortions introduced by the horizontal motion of the particle. Because this algorithm is part of a proprietary software package

and not openly available, we have chosen to approximate the effects of the algorithm based on the conceptual description provided by Kruger and Krajewski (2002).

Recall, from section 3.3, that the 2DVD unskewing algorithm works by assuming that the top-most and bottom-most points on an observed particle are actually vertically aligned with one another (Kruger and Krajewski, 2002). In the event of multiple particle pixels appearing on the top-most or bottom-most scan line, it is unclear how the actual algorithm selects the top-most or bottom-most point of the particle. To this end, we have decided to define the top-most and bottom-most points of the particle as the mean horizontal location of the particle pixels on the top-most and bottom-most scan lines, respectively.

The actual unskewing process in our algorithm works by linearly interpolating the horizontal offset, as implied by the motion-skewed locations of the top-most and bottom-most points of the particle, to the vertical position of each scan line. Each scan line is then horizontally shifted by the appropriate offset via another linear interpolation. The value of the interpolated pixels are then rounded to either zero, indicating a non-particle pixel, or one, indicating a particle pixel. Figure 2c depicts the result of applying the unskewing algorithm; note the slight differences in the apparent particle orientation between the original image (Fig. 2a) and the unskewed image (Fig. 2c). We will examine the effects of these differences on the accuracy of the 2DVD bounding-box measurements in section 5.3.

4.3 ~~Tensor-fitted ellipse~~Tensor method

To provide an additional point of comparison for the particle measurements, we also implemented an alternative ellipse fitting strategy, referred to here as the tensor method, that uses the method implemented in the fit_ellipse program in the Coyote IDL Library (accessible at http://www.idlcoyote.com, accessed ~~15 December 2021~~27 March 2022). This method works by computing the eigenvalues, e_1 and e_2, of a two-by-two mass distribution tensor matrix, which is defined as

$$\tau = \begin{bmatrix} \overline{(\Delta y)^2} & \overline{\Delta x \Delta y} \\ \overline{\Delta x \Delta y} & \overline{(\Delta y)^2} \end{bmatrix}, \tag{1}$$

where Δx and Δy are the distances from the particle centroid in the horizontal and vertical directions, respectively, and the overbars indicate averaging. The semimajor and semiminor axes, a and b, respectively, are then calculated from the ~~eingenvalues~~eigenvalues via

$$a = 2\sqrt{e_1} \quad \text{and} \quad b = 2\sqrt{e_2}. \tag{2}$$

The major and minor axes of the fitted ellipse can then be obtained simply by doubling the semimajor and semiminor axes. Finally, in the interest of completeness, the tensor-fitted ellipse orientation, measured counterclockwise from the positive x axis, is computed as $\arctan(e_2/e_1) - 90°$. As will be discussed in section 5.2, the ~~tensor-fitted ellipse~~tensor method produces similar, although not identical, fits to that of the emulated MASC method.

5 Results

5.1 Area and equivalent diameter

~~There are a number of factors of varying importance that can affect the measurement of the particle area and, therefore, the equivalent diameter. For this study, however,~~

Here we will examine the theoretical accuracy ~~and precision~~ of the MASC, PIP, and 2DVD area and equivalent diameter measurements in terms of ~~motion blur as determined from the pixel resolution~~ the susceptibility of each instrument to motion blurring, which is a function of the pixel size, exposure time (i.e., shutter speed), and particle ~~fall~~ vertical and horizontal movement speed. For both PIP and MASC, the equivalent diameter is determined from the area of the two-dimensional projection of the particle, which is measured by counting the number of pixels within the particle and then multiplying that number by the calibrated pixel size (i.e., the area each pixel represents in the image). As discussed in section 3.3, 2DVD, however, reports an equi-volume diameter rather than an equi-area diameter by summing the ~~values~~ volumes of stacked elliptical cylinders. For simplicity, we will only consider the effects of motion blur on the two-dimensional area and equi-area diameter from a single camera for 2DVD rather than the estimated volume and equi-volume diameter that the 2DVD software computes. ~~It should be noted that~~ Additionally, there are a number of factors that are not being considered in the present study. ~~For simplicity, we~~ We have not included the effects of image thresholding on the area measurements nor have we considered the effects of each camera's point spread function, which describes how light is distributed amongst each of the imaging pixels of the camera sensor. We also do not consider the effects of the circle of confusion (i.e., the blurring of out-of-focus particles) as this will only be an issue if the particle is located outside of the instrument's depth of field.

To understand the expected impacts of motion blurring on area and equivalent diameter, we perform a statistical analysis of simulated particle motion blurring. To simulate the particle motion blurring, we generate 196 one-dimensional particles, whose lengths are uniformly distributed between 0.5 mm and 20 mm, and then add the expected motion blurring for a particle moving at 1 m/s, 4 m/s, and 10 m/s (in the along-particle direction) for each instrument. The expected motion blurring is computed by multiplying the particle movement speed by the camera exposure time. The particle motion speeds are chosen as follows (e.g., Locatelli and Hobbs, 1974; Garrett et al., 2012; Vázquez-Martín et al., 2021): 1 m s^{-1} represents a relatively fast falling snow particle, 4 m s^{-1} represents an excessively fast fall speed for snow particles, and 10 m s^{-1} represents relatively fast horizontal motion (albeit below the U.S. National Weather Service wind threshold for blizzard conditions, $\sim$15 m s^{-1}). The motion-blurred particles are then randomly moved a fractional distance between zero and two pixel lengths before being pixelated as per the pixel sizes of each instrument. In the case of PIP, the image compression is also taken into account. This process is repeated 100000 times in total for each particle size and each fall speed. For computational simplicity, these calculations assume that a pixel is part of the particle if it has at least 50% coverage.

The MASC setup used during ICE-POP had a pixel ~~resolution~~ size of 33.5 μm and an exposure time of 40 μs or 1/25000th of a second (Jacopo Grazioli, 2021, personal communication). With this exposure time, a single point on an arbitrary precipitation particle falling directly downwards at 1 m s^{-1} would appear as a 40 μm long vertical line. With a pixel ~~resolution~~ size of 33.5 μm, there would be minor motion blurring as a single point on the particle would contribute to at least two pixels and possibly

three, although ~~in the three pixel case the contribution to two of the pixels would be very minor. This blurring would effectively add an extra pixel below all of the lower-edge pixels and above all the upper-edge pixels of a particle, thereby increasing the area and, therefore, the equivalent diameter~~our statistical analysis indicates an average 1.19 extra particle pixels are added. Note, the actual number of pixels added in any instance will always be a whole number. For a 4 m s^{-1} fall speed, which would be an extremely fast fall speed for a snowflake, the point would appear as a 160 μm line, appearing in at least five pixels (average of 4.78 extra pixels) in the ICE-POP MASC. For a particle moving at 10 m s^{-1}, the point appears as a 400 μm line, contributing to at least 12 pixels (average of 11.94 extra pixels), although it should be noted that the enclosed sampling volume used by MASC would make such a particle motion less likely than it would be with an open sampling volume.

PIP has an exposure time of 28 μs or ∼1/35700th of a second. With this exposure time, a ~~single~~ point on a ~~precipitation particle falling vertically at~~ particle moving at either 1 m s^{-1}, 4 m s^{-1}, or 10 m s^{-1} would appear as ~~a 28 μmlong vertical line; however, because PIP~~ either a 28 μm, 112 μm, or 280 μm line, respectively. Unlike MASC, however, PIP uses image compression in the vertical and, as such, the effects of horizontal and vertical motions must be examined separately. For horizontal motion, PIP has a pixel size of 0.1 mm (100 μm). As such, the motion blurring ~~would be almost negligible . In the case of a particle falling at~~ effects from 1 m s^{-1} of motion would be negligible under our assumptions with only 0.28 pixels added on average. At 4 m s^{-1} ~~, the single point would produce a 112 μm long line for the given exposure time and , as such, would contribute to at least two pixels . Complicating this, however, is that~~ and 10 m s^{-1}, the motion blurring has a large impact with an average of 1.12 and 2.8 pixels added, respectively.

As mentioned previously, PIP compresses the images prior to transmission to the computer where the image analysis is performed. As part of the PIP compression algorithm, pairs of vertically adjacent pixels are averaged together, resulting in a pixel size of 0.1 mm by 0.2 mm while the image is compressed. When the image is decompressed, this averaged value is assigned to the individual pixels making up the pixel pair (L. Bliven, 2021, personal communication), producing an effective ~~resolution~~ pixel size of 0.1 mm by 0.2 mm. Because the motion blurring effects apply to the pre-compression pixels rather than the post-compression pixels, the interaction between image compression and motion blurring warrants a closer examination.

Figure 3 depicts schematic representations of how the effects of motion blurring and image compression interact for a single column of PIP pixels when viewing a 0.4-mm tall particle. ~~Note, for the purposes of this demonstration,~~ As with our statistical analysis, we have arbitrarily decided that a pixel will count as being part of the particle when it appears to be at least 50% covered by the particle, taking both motion blurring and compression into account. In actuality, the assignment of ~~a~~ an 8-bit PIP pixel to a particle is more complicated and is described in Newman et al. (2009). The simplest example is that of a particle suspended in air (i.e., with zero fall speed), which will experience no motion blurring (Fig. 3a). The left column of pixels in the panel depicts the pre-compression view, which sees the particle as being four pixels tall (i.e., 0.4 mm) as there is no motion blurring or compression. Because a particle can be located anywhere relative to the compression algorithm's vertically paired pixels, we must examine the two extreme possible alignments to understand the range of possible effects. In the first of these possibilities, the top edge of the particle is aligned with the top of the upper pixel in the pixel pair, depicted in the center column of pixels in the panel. This alignment results in the post-compression image also depicting a particle that is four pixels tall. In the second of these possibilities, the top edge of the particle is aligned with the top of the lower pixel in the pixel pair,
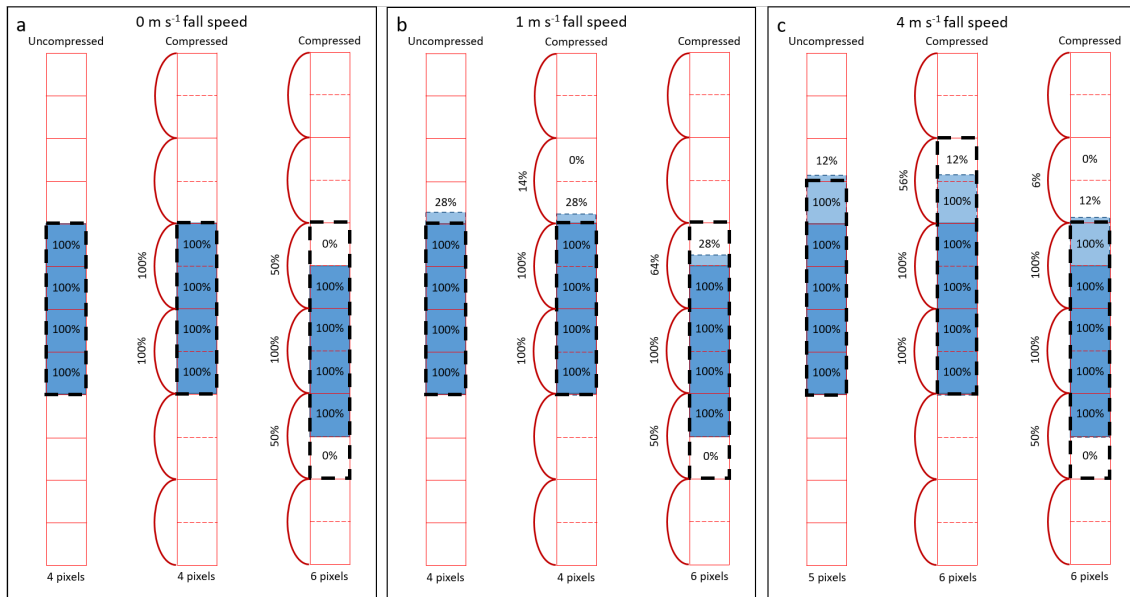
Figure 3. Schematic depictions of the combined effects of PIP image compression and motion blurring for a 0.4 mm tall particle (a) suspended in air, (b) falling at 1 m s^{-1}, and (c) falling at 4 m s^{-1}. The red boxes indicate individual PIP pixels and the dark red braces indicate the pairs of pixels that are averaged together during the image compression. The darker blue shading indicates the area covered by the actual particle at the end of the image exposure period and the lighter blue shading indicates the additional area that appears to be covered by the particle due to motion blurring. The dashed black box indicates the pixels which appear to be contained within the particle (i.e., having ge50% coverage), taking into account both motion blur and compression; for demonstration purposes, we count a pixel as being in the particle if it appears to be at least 50% covered after motion blur and compression effects are applied. The value listed inside each pixel indicates the percent coverage of that pixel by the apparent particle prior to compression and the value listed along each brace indicates the percent coverage of the paired pixels post-compression (this value is assigned to both of the paired pixels when the image is decompressed for analysis). In each panel, the left column of pixels represents the view before compression and the right two columns of pixels represent the view after compresscompression, corresponding to the two extremes of how a particle could align to the compression pixel pairs.

depicted in the right column of pixels in the panel. This alignment results in the post-compression image depicting the particle as being six pixels tall instead of four pixels tall because the completely covered top- and bottom-most pixels of the particle are averaged with completely uncovered non-particle pixels resulting in a post-compression pair of pixels that each appear to be 50% covered by the particle (as indicated by the number next to the dark red brace in Fig. 3a). As such, the image compression can add between zero and two additional pixels to the height of a suspended particle.

For a PIP-viewed particle with a non-zero fall speed, both both motion blurring and image compression must be taken into account; schematic representations for particles falling a 1 m s^{-1} and 4 m s^{-1} are depicted in Fig. 3b and Fig. 3c, respectively. Because of our prescribed threshold of 50% coverage, the 1-m-s^{-1} fall speed case ends up producing the same result as very similar results to the suspended particle case: the addition of between zero and two additional pixels to the height with an

average of 1.28 extra pixels at 1 m s^{-1} compared to 1.01 pixels when suspended. For a particle falling at 4 m s^{-1}, however, the motion blurring adds an extra ~~two pixels~~ pixel to the height prior to compression and the image compression adds an ~~additional zero to two pixels~~another pixel; in total, the motion blurring and compression effects add between ~~two and four~~ one and two pixels to the height of the particle ~~.~~ (average of 2.12 pixels). Since the image compression only impacts vertical motion blurring, the 10 m s^{-1} particle motion is not included in Fig. 3, but were a snowflake to fall at such an extreme speed, motion blurring would, on average, add 3.80 pixels after accounting for the effects of image compression.

Of the three instruments considered in this study, 2DVD has the shortest exposure time at 18μs or ~1/55500th of a second (Schönhuber et al., 2008). Each 2DVD line-scan camera has 512 photodetectors that are calibrated to observe a 100-mm wide field of view (Kruger and Krajewski, 2002), from which we can infer that the pixel size is ~0.195 mm (195 μm). With the shortest exposure time and largest pixel size, it is unsurprising that 2DVD should be the most resilient to motion blur effects according to our analysis of the camera setups. Even for a snowflake with an extremely fast fall speed of 4 m s^{-1}, a single point on the snowflake would only move 72 μm during the exposure time, which is only ~37% of the pixel size assuming the camera elements gather light isotropically in the vertical and horizontal directions. As with MASC, 2DVD uses an enclosed sampling volume, which reduces the opportunities for particles to have large horizontal motions; however, a particle moving at 10 m s^{-1} would still only add 0.92 pixels on average to the along-motion particle dimension.

The extent to which the motion blurring discussed above might impact the measurement of particle area and, by extension, the equivalent diameter, is highly dependent on the shape and orientation of the particle. The area errors introduced by motion blur will increase as the number of particle pixels ~~with an adjacent~~ located directly below a non-particle pixel ~~either above or below increases; these particle pixels are~~ (hereafter referred to as top-edge ~~or bottom-edge pixels, respectively. Motion blur of the~~ pixels) increases. The motion blur of these top-edge pixels occurs when the particle leaves those pixels during the image exposure period~~while motion blur for the bottom-edge pixels occurs when the particle enters those pixels during the image exposure period. The motion blur effects will impact the area calculation at both the top-edge pixels and bottom-edge pixels simultaneously. .~~ For a particle with an entirely convex edge, the total number of ~~top- and bottom-edge~~ top-edge pixels will be equal to ~~twice~~ the horizontal width of the particle, in pixels; for non-convex particles, the ~~combined number of top- and bottom-edge~~ number of top-edge pixels will depend on how many additional ~~top- and bottom-edge~~ top-edge pixels, if any, the concave edge regions add to the particle.

~~Figure 4 depicts the impacts~~ Using the statistical analysis detailed above, we determined the average effects of motion blurring ~~on~~ (and image compression) on area by multiplying the number of extra particle pixels by the length in pixels of the original particle (i.e., we assume the particles are circular for the area calculations). In the interest of computational simplicity, coverage in the along-motion and across-motion directions are considered independently (i.e., a pixel located such that the across-motion particle width produces 50% coverage of the pixel and the ~~area and equivalent diameter measurements for a perfectly circular particle being observed by MASC and PIP; 2DVD is not included in this figure as the motion blurring effects are negligible by comparison. Note that the PIP values also account for the effects of the image compression algorithm, using the average number of additional pixels for each fall speed.~~ along-motion particle width produces 50% coverage of the pixel would count as a particle pixel even though such a pixel would actually only be 25% covered by the particle). Additionally, we

400 assume that the particle is positioned such that its across-motion extent begins at the edge of a pixel and that the cumulative effects of pixelating a circular particle on apparent pixel coverage are negligible (in actuality, the coverage for each edge pixel would vary across the edge of a circular particle). Finally, we assume that all particles have a non-oblique motion relative to the pixel grid. The results of this analysis are depicted in Fig. 4 with the fractional biases being calculated against the diameter and area that would be observed by each instrument in the absence of motion blurring but with pixelation effects.

405 While the relative impact of the area and equivalent diameter errors introduced by motion blur depend on the application for which the measurements are being used, we can draw a number of general conclusions from our analysis. Key among these is that very small particles will experience considerably greater ~~impact~~ relative impacts from the motion blur errors than will larger particles falling at the same rate; this is particularly evident in the fractional bias plots (Fig. 4c,d). Generally speaking, however, smaller particles of a given ice crystal habit will tend to have a slower fall speed than larger particles of that same

410 habit and will, therefore, typically experience less motion blur ~~. Even with~~ in the vertical. For a very fast fall speed of 4 m s^{-1}, the overestimation of the equivalent diameter for very large circular particles (diameter $\gtrsim$ 10 mm) is approximately two orders of magnitude ~~or more~~ smaller than the actual equivalent diameter ~~.~~ (i.e., a fractional bias of 0.01). Even with a relatively fast horizontal motion of 10 m s^{-1}, the equivalent diameter fractional bias still remains below 0.03 for these very large circular particles. Critical to interpreting and applying these results is the fact that we assumed that the particles are

415 perfectly circular. As the ~~the horizontal dimension of the particle grows relative to the~~ particles under consideration become more oblate, with their horizontal dimension growing relative to their vertical dimension, ~~the relative (and absolute)~~ both the absolute and relative effects of motion blurring on the measured area and equivalent diameter ~~measurements will also grow.~~ ~~Additionally, as the particle outline becomes more complex and the number of top- and bottom-edge pixels in the particle~~ ~~increases, the measurement errors~~ will increase relative to circular particles of the same area. An oblate particle of equal

420 area to a circular particle will have more top-edge pixels than the circular particle. Increasing the number of top-edge pixels increases the number of pixels that the motion blur is applied to, thereby increasing the absolute effects of the motion blurring. Furthermore, because the particle area is remaining constant, the relative effects of the motion blurring are also increasing. For similar reasons, increasing the complexity of the particle outline such that the number of top-edge pixels increases will also result in larger biases due to motion blur~~will also increase~~.

425 ~~Approximate effects of motion blur on (a,c) measured area mm^{2}and (b,d) measured equivalent diameter mmfor hypothetical~~ ~~perfectly circular particles falling at (dashed line) 1 m s^{-1} and (solid line) 4 m s^{-1} as (red) MASC and (black) PIP would~~ ~~measure them. The effects are portrayed both in terms of (a,b) the bias, computed as the difference between the motion-blurred~~ ~~quantity and the actual particle quantity, as well as (c,d) the fractional bias, computed as the bias divided by the actual particle~~ ~~quantity. Calculations are performed based on the number of additional full pixels added by motion blur to a circular particle.~~

430 ~~Specifically, for each additional pixel of diameter, the motion blur will add two (eight) pixels to the MASC measurement of a~~ ~~particle falling at 1 m s^{-1} (4 m s^{-1}) and one (three) pixels to the PIP measurement of a particle falling at 1 m s^{-1} (4 m s^{-1}).~~ ~~Gray vertical lines indicate the corresponding quantity for particles with actual diameters of 1, 5, 10, and 15 mm.~~

~~The differences between PIP, MASC, and 2DVD measurements of area and equi-area diameter are best considered~~ Rephrasing this discussion in terms of ~~precision and accuracy. As demonstrated in Fig. 4~~accuracy, PIP will produce a slightly less accurate
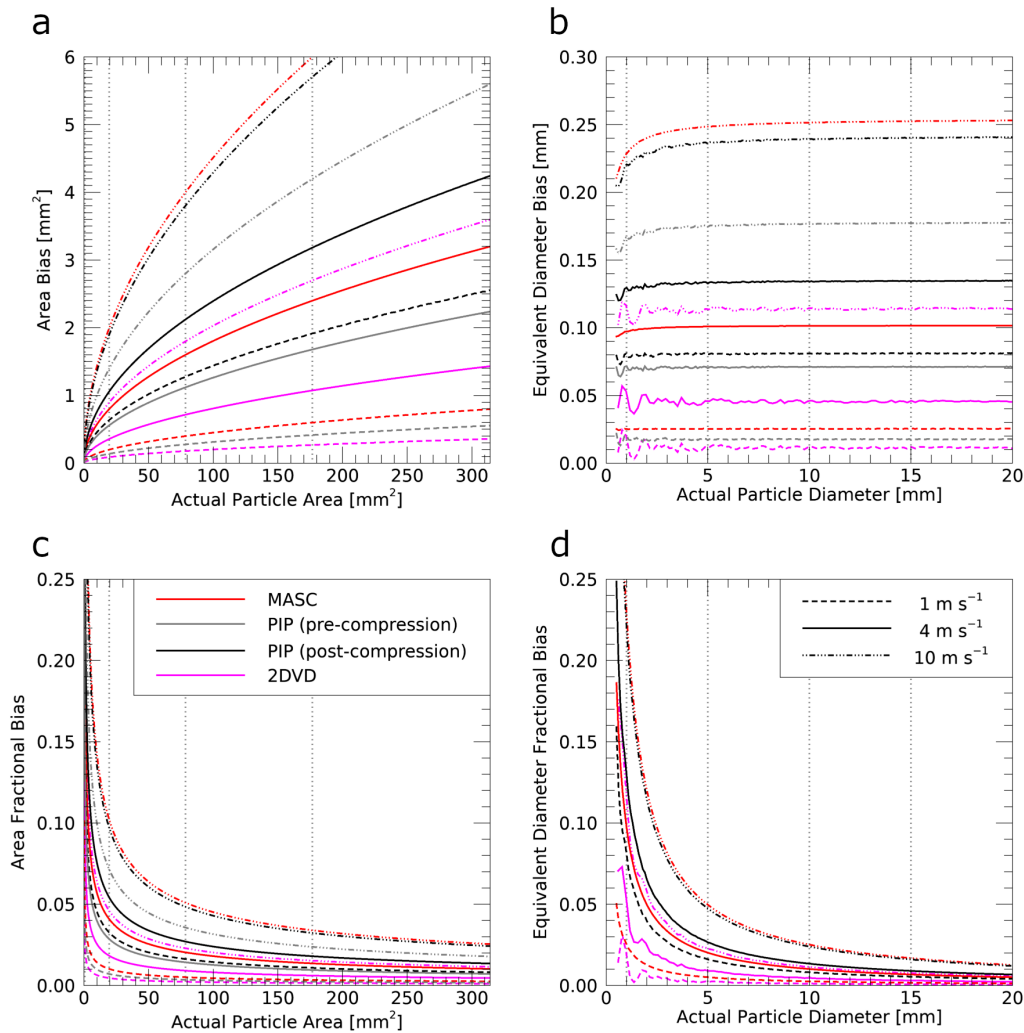
Figure 4. Approximate effects of motion blur on (a,c) measured area [mm^2] and (b,d) measured equivalent diameter [mm] for hypothetical perfectly circular particles moving at (dashed line) 1 m s^{-1}, (solid line) 4 m s^{-1}, and (dash-dot line) 10 m s^{-1} as (red) MASC, (gray) PIP before accounting for image compression, (black) PIP after accounting for image compression, and (magenta) 2DVD would measure them. The effects are portrayed both in terms of (a,b) the bias, computed as the difference between the motion-blurred quantity and the actual particle quantity, as well as (c,d) the fractional bias, computed as the bias divided by the non-blurred, non-compressed particle quantity. Gray vertical lines indicate the corresponding quantity for particles with diameters of 1, 5, 10, and 15 mm.

435 measurement of area and equivalent diameter than MASC in the presence of motion blurring—, as demonstrated in Fig. 4. Interestingly, were the PIP images not compressed prior to their analysis, the PIP would produce slight more accurate measurements than MASC(not shown). That said, MASC will produce a considerably more precise measurement than either PIP or 2DVD

due to the very high resolution of the MASC cameras relative to the other cameras examined in this study.. As we established above, the 2DVD should be the least effected is the least affected by motion blur and, as a result, should have has the most accurate area measurements. Due to the complicating factor of the PIP image compression algorithm, it is unclear whether the precision of the PIP, which has 0.1 mm tall camera pixels but 0.2 mm tall effective pixels, would have greater or approximately equal precision as 2DVD, which has ~0.195 mm pixels, as this would likely depend on the alignment between the particle edges and the vertically paired pixels used in the PIP image compression algorithm. For a worst-case scenario, however, the precision of the PIP would likely be similar to that of the 2DVD as the post-compression effective PIP pixel height is only slightly greater than that of the 2DVD, although the large pixel size may impact the precision of those measurements.

5.2 MASC and PIP shape-fitting measurements

As discussed in section 4, the PIP data for the period between 0000 UTC 7 March and 0000 UTC 9 March 2018 at the ICE-POP MHS observation site was processed using algorithms designed to emulate both the MASC and 2DVD measurements of particle dimension. The present section will focus on the shape-fitting techniques of measurements made using the MASC and PIP shape-fitting methods as well as an alternative ellipse-fitting technique, referred to here as the tensor method. A comparison of the bounding box measurement techniques used by PIP and 2DVD will be the subject of a separate section (section 5.3) rather than being included here. Although the tensor-fitted ellipse measurements will be used here as a point of comparison between the various instrument algorithms, it should be noted that the tensor-fitted ellipse measurements are not a 'ground truth' and are subject to errors of their own.

Comparing the PIP-fitted ellipse particle dimensions to those of the tensor-fitted ellipse, we found that the PIP-fitted ellipse tends to overestimate the long dimension of the particle (Fig. 5a), measured as the major axis of the fitted ellipse, and underestimate the short dimension (Fig. 5d), measured as the minor axis of the fitted ellipse, relative to the tensor-fitted ellipse. As a result, the aspect ratio, computed as the short dimension divided by the long dimension, is almost always underestimated by the PIP-fitted ellipse relative to the tensor-fitted ellipse. This underestimation is particularly noticeable for larger tensor-fitted ellipse aspect ratios with the PIP-fitted ellipse aspect ratio rarely exceeding 0.6 and never reaching 0.7 during this event, despite there being a period of lump graupel, which should have a relatively high aspect ratio due to its almost spherical shape, on 8 March (Fig. 6a).

Using the aspect ratio computed from the

Using the PIP-fitted rectangle instead of the as the basis of the particle dimension measurements produces a sizable improvement to the agreement between the PIP-fitted long dimension and the tensor-fitted long dimension (Fig. 5a,b) but almost no change in agreement for the short dimension (Fig. 5d,e). To quantify these changes in the dimension measurement agreement, we use both the mean absolute difference, defined as the average magnitude of the difference between two sets of measurements, and the mean absolute fractional difference, defined as the average magnitude of the normalized difference between two sets of measurements; for Fig. 5, the normalization is done relative to the measurements made using the tensor method.

Comparing the mean absolute differences for the the PIP-fitted ellipse appears to increase the spread in both long and short dimensions,producing more agreementwith the ellipse and PIP-fitted rectangle dimensions, the PIP-fitted rectangle long dimension

has a mean absolute difference from the tensor-fitted ~~ellipse measurements~~ long dimension that is approximately three-fifths that of the PIP-fitted ellipse long dimension (Fig. 5~~b,~~a,b). The PIP-fitted ellipse and rectangle short dimensions, however, show almost no change in mean absolute difference (Fig. 5d,e). ~~As a result~~

In terms of aspect ratios, the PIP-fitted rectangle aspect ratio (Fig. 5h) does not suffer from the artificial cap that was present with the PIP-fitted ellipse aspect ratio (Fig. 5~~h)~~ g) and much of the reduction in the mean absolute difference (0.254 versus 0.184) is likely due to this change. That said, the PIP-fitted rectangle aspect ratio still has a tendency to greatly underestimate the aspect ratio relative to the tensor-fitted ellipse aspect ratio (Fig. 6b). Interestingly, the PIP-fitted rectangular aspect ratio frequently has a value of one but very rarely has a value between 0.9 and one; this peculiarity will be revisited later in this section. It should be noted, however, that the PIP-fitted rectangle aspect ~~ratios do capture the expected~~ ratio does capture the increase in aspect ratio ~~that would be expected during the period~~ associated with the periods of lump graupel precipitation on 8 March around 0900 UTC and after 1400 UTC.

In contrast to the two sets of PIP-fitted measurements, the MASC-fitted ellipse measurements show very good agreement with the tensor-fitted ellipse measurements of both the long and short dimension of the fitted particles (Fig. 5c,f) with mean absolute differences ranging between approximately one-fifth and two-fifths those for the PIP-fitted shapes. As a result, the aspect ratios computed from ~~the MASC-fitted and tensor-fitted ellipses~~ shapes fitted using the emulated MASC method and the tensor method are in fairly good agreement as well (Fig. 5i). This general agreement between the particle dimensions measured from the MASC-fitted and tensor-fitted ellipses suggests that these two shape fitting algorithms may ~~be giving~~ give more reliable results than the PIP-fitted ellipse or rectangle algorithms. To gain insight into whether the ~~MASC- and tensor-fitted ellipse~~ MASC and tensor method algorithms are, in fact, performing better than the PIP algorithms, we will take a closer look at individual particles.

Figure 7a-d depicts snapshots of particles taken from the PIP AVI files with the fitted shapes from the various fitting algorithms depicted in Fig. 7e-h. For simplicity, these particles will be referred to by the Fig. 7 panel letter of the unannotated panels (i.e. panels a–d). Particles (a) and (b) are both likely some type of aggregate frozen precipitation based on their odd shapes. Based on the relatively circular shapes of the remaining two particles, relatively high fall speeds ~~, and lack of a warm nose~~ (black line, Fig. 1), subfreezing near-surface temperatures (red line, Fig. 1), and the lack of an above freezing temperature layer in a nearby thermodynamic sounding (not shown), particles (c) and (d) are likely both examples of lump graupel. Qualitative and quantitative inspection of the particles in Fig. 7 and Table 1 supports the patterns identified by the scatter-plot analysis performed above. Specifically, the PIP-fitted ellipse appears to always overestimate the particle long dimension; the PIP-fitted rectangle appears to sometimes overestimate, sometimes underestimate, and sometimes accurately capture the particle long dimension; and, finally, the emulated MASC-fitted and the tensor-fitted ellipses tend to produce fairly similar particle dimensions. Additionally, visual examination of these and other (not shown) individual particles suggests that the emulated MASC-fitted and the tensor-fitted ellipses tend to provide more reasonable estimates of particle dimension than either of the PIP-fitted shapes.

The primary cause of the discrepancies between the actual particle shape and the PIP-fitted ellipse major and minor axes is rooted in the reliance on only the particle area and perimeter for making the PIP fits, as discussed in section 3.1. For ~~particles~~

Table 1. Dimensions and aspect ratios of the fitted shapes for various particle fits depicted in Fig. 7, referenced by the figure panel letter for the unannotated particle images. The perimeter stretching factor is computed by dividing the actual particle perimeter by the perimeter of a circle of equal area to the particle.

Fitted-shape measurement	Fig. 7 panel			
	(a)	(b)	(c)	(d)
Equivalent diameter [mm]	1.680	1.950	1.130	1.180
Perimeter stretching factor	1.148	1.308	1.144	1.068
PIP-fitted ellipse				
Major-axis length [mm]	2.478	3.434	1.660	1.529
Minor-axis length [mm]	1.141	1.105	0.767	0.916
Aspect ratio	0.460	0.322	0.462	0.599
PIP-fitted rectangle				
Long-side length [mm]	1.790	3.020	1.190	0.990
Short-axis length [mm]	1.240	0.990	0.840	0.990
Aspect ratio	0.693	0.328	0.706	1.000
Emulated MASC-fitted ellipse				
Major-axis length [mm]	2.093	2.121	1.238	1.132
Minor-axis length [mm]	1.288	1.831	0.926	1.065
Aspect ratio	0.615	0.863	0.748	0.941
Tensor-fitted ellipse				
Major-axis length [mm]	2.178	2.311	1.331	1.227
Minor-axis length [mm]	1.363	1.831	0.990	1.150
Aspect ratio	0.626	0.792	0.744	0.938

a particle with complicated outlines, such as particle (b), the ~~perimeter of the particle~~ particle perimeter is far greater than the perimeter ~~would be for~~ of an ellipse or rectangle of either equal area or equal ~~length and width or of equal area~~dimensions. To make up ~~for~~ the extra perimeter for a given area, the PIP-fitted ellipse or rectangle must have a larger long dimension at the expense of the short dimension~~; note, extending the short dimension at the expense of the long dimension would result in a smaller perimeter for a given area.~~.

The relationship between excess perimeter for a given area, relative to a perfect ~~ellipsoid~~ellipse, and error in the PIP-fitted ellipse or rectangle dimension lengths is demonstrated in Fig. 8 for a hypothetical particle that we will ~~treat as being~~ assume is generally circular with ~~an edge that undulates about a~~ a mean radius of 0.5 mm, but whose edge is not a perfect circle and is allowed to vary about the mean radius of 0.5 mm in such a way that the area, center position, and area remain unchanged from that of a 0.5-mm-radius circle. The amplitude of this ~~undulation~~ variation in radius is represented in Fig. 8 by the perimeter stretching factor, computed as the actual perimeter of a particle divided by the perimeter of a circle of equal area to the particle.

~~For our hypothetical particle, the perimeter stretching factor is prescribed and the actual perimeter is computed by multiplying the perimeter of a 0.5-mm-radius circle by the perimeter stretching factor.~~ A perimeter stretching factor of one will result in a

520   perfect circle while a larger perimeter stretching ~~factors would be associated with particles with increasingly rough edges.For clarity, we do not construct a particle with undulating edges to make the measurements, but instead merely parameterize its existence using the~~ factor indicates the presence of larger variations in radius (i.e., a more complicated shape). For our hypothetical particle, we prescribe a value between one and two for the perimeter stretching factor and set the particle area to be equal to that of a 0.5-mm-radius circle. As the PIP-fitting method only requires perimeter and area, our parameterization

525   of the hypothetical particle is sufficient to perform an ellipse or rectangle fit without constructing the hypothetical particle itself. The equations used to make the PIP fits are derived in appendix A. Note that nothing in these equations relates to the actual shape of the particle itself (e.g., circular or elliptical); this missing information is, in fact, the core issue with both the PIP ellipse and rectangle fits. Without knowledge of the underlying distribution of pixels within the particle, the PIP algorithm assumes all particles are either smooth-edged ellipses or rectangles, depending on which fit is used.

530   The lack of shape information in the PIP shape-fitting equations does not entirely explain the artificial cap in aspect ratio at ~0.6 for the PIP-fitted ellipses (as in Fig. 6a), however. If the shape issues were the sole cause of the artificial cap, the circular presentation of particle (d) in Fig. 7 would produce a very good fit with an aspect ratio approaching one. Instead, the aspect ratio cap is a result of where the aspect ratio is most sensitive to the perimeter length for a given particle area. Based on Fig. 8a, the PIP-fitted ellipse major and minor axis lengths and aspect ratio show the greatest sensitivity to perimeter length when the

535   perimeter stretching factor is almost one (i.e., the particle is almost a perfect circle). As a result, any deviation from the smooth edge of a perfect circle will produce relatively rapid changes in the ellipse axis ~~length~~ lengths and the aspect ratio. It was noted above that the PIP-fitted ellipse measurements (Fig. 6a) appear to produce an artificial cap on aspect ratio of approximately 0.6. Based on our calculations for Fig. 8, a PIP-fitted ellipse aspect ratio of 0.6 is produced when the perimeter stretching factor is approximately 1.065 ~~despite our hypothetical particlebeing circular (i. e., having an aspect ratio of 1.0). For a circular particle~~

540   ~~with a radius of 0.5 mm (giving a~~ regardless of the actual shape of the particle. For reference, a perimeter stretching factor of 1.065 corresponds to taking a 10-mm diameter circular particle, which would have a perimeter of ~~~3.14 mm), a perimeter stretching factor of 1.065 corresponds to the addition of approximately 0.2 mm to the particle perimeter when performing a PIP-ellipse fit.~~ 31.4 mm, and stretching its perimeter by 2 mm. Small increases in perimeter, such as this, can be introduced by a few very small deviations of the particle edge from a perfect circle as well as by the inability to perfectly represent a circle

545   using square pixels (i.e., pixelation effects).

     The PIP-fitted rectangle long- and short-side lengths (Fig. 8b) have a similar sensitivity to changes in perimeter stretching factor when the perimeter stretching factor is relatively small. A key difference from the PIP-fitted ellipses, however, is that a rectangle cannot be fit to a particle whose perimeter stretching factor is less than $\sqrt{4/\pi}$ (~1.128), below which the particle perimeter is not long enough to construct a rectangle of the required area. Particle (d) in Fig. 7 is an example of a particle

550   whose actual perimeter is too short to produce a physical rectangle fit using the PIP method. The perimeter stretching factor for particle (d) is approximately 1.068, as listed in Table 1. Based on this and similar particles, the PIP processing algorithm appears to automatically assign an aspect ratio of one to any particle where a rectangle fit is not physically possible~~.~~, the

effects of which can be seen in Fig. 6b. How the PIP algorithm determines a dimension length in these cases, however, remains unclear.

555   The artificial cap in aspect ratio that appears for the PIP-fitted ellipses (Fig. 6a) does not appear to be present in the aspect ratio distributions for the PIP-fitted rectangles (Fig. 6b). A key factor that leads to this difference is demonstrated in Fig. 8b: a rectangle aspect ratio of 0.6 corresponds to a perimeter stretching factor of $\sim$1.165 compared to 1.065 for a PIP-fitted ellipse (Fig. 8a). The importance of this perimeter stretching factor difference of 0.1 comes down to the observed distribution of particle perimeters relative to their areas (i.e., the distribution of particle stretching factors). Figure 9 depicts the distribution

560   of perimeter stretching factors for the 7–8 March 2018 case. Only a very small portion of particles have a perimeter stretching factor smaller than 1.065 (left red dashed line, Fig. 9) whereas the most frequently observed perimeter stretching factors are smaller than the 1.165 perimeter stretching factor (right red dashed line, Fig. 9) that corresponds to the 0.6 aspect ratio for a PIP-fitted rectangle.

5.3   2DVD and PIP bounding-box measurements

565   Because the 2DVD does not instantaneously capture full images of a particle, a 'long dimension versus short dimension' aspect ratio is not computed by the instrument. Instead, 2DVD reports the 'oblateness' of a particle, defined as the aspect ratio of the horizontal width of a particle versus its vertical height. For each camera, the oblateness is computed as the bounding-box height divided by the bounding-box width. An overall oblateness is then computed from the two single-camera oblateness values by their geometric mean. Here we will focus on the measurements of the bounding box height and width rather than the oblateness

570   itself.

Both PIP and 2DVD record measurements of the horizontal and vertical extent of a particle in the form of bounding-box measurements, computed as the horizontal distance between the left-most and right-most particle pixels and the vertical distance between the bottom-most and top-most particle pixels. Because PIP takes measurements using two-dimensional images, the accuracy of the PIP bounding-box measurements is directly tied to the resolution of the camera and the degree of motion

575   blurring. 2DVD, however, uses a line-scan camera and, as such, the 2DVD bounding-box measurements are derived quantities. The bounding-box height is derived from the time required for the particle to pass through the sampling region while the bounding-box width is computed by compositing the individual line scans to rebuild the particle. As discussed in section 3.3, 2DVD uses an unskewing algorithm to correct the bounding-box width measurements for horizontal particle motion. As the actual 2DVD unskewing algorithm is part of a proprietary software package, we use a conceptually similar unskewing

580   algorithm, described in section 4.2.

The vertical sampling frequency of the emulated 2DVD images is usually much higher than that of the underlying PIP images due to the high temporal sampling frequency of 2DVD. In order for the vertical sampling frequency of the emulated 2DVD images to be lower than that of the underlying PIP images, the precipitation particle would need to have a fall speed greater than $\sim$68.2 m s^{-1}, which is unrealistic for any of the particles of interest. Because of this, we would expect the emulated

585   2DVD bounding-box height to almost exactly match the PIP-measured bounding-box height and this expectation is borne out in

the data as depicted in Fig. 10a : the mean absolute difference between the PIP and 2DVD bounding-box heights is only 0.016 mm (a mean absolute fractional difference of 0.013).

Of greater interest is how well the emulated 2DVD bounding-box width measurements perform, especially after we apply our approximation of the 2DVD unskewing algorithm. Figure 10b a compares the uncorrected emulated 2DVD bounding-box width with the PIP-measured bounding-box width. Unsurprisingly, the uncorrected 2DVD measurements tend to greatly overestimate the width of the bounding box due to the skewing introduced by any horizontal motion of the particle. One point of interest, however, is that the uncorrected 2DVD measurements can sometimes underestimate the bounding box width. This can occur when a particle of sufficiently low aspect ratio is moving in the opposite direction of the tilt of the top of the particle (e.g., a needle crystal whose top is to the left of the particle centroid moving towards the right); this will result in the particle being compressed in the horizontal because the particle is scanned from the bottom upwards as it falls through the plane of the line scan camera.

The emulated 2DVD measurements of bounding-box width are surprisingly accurate once our approximation of the 2DVD unskewing algorithm is applied (Fig. 10c). b). Quantitatively, the unskewing algorithm produces a reduction in mean absolute difference from 0.58 mm to 0.187 mm. After the correction, the overestimation of bounding-box width is considerably reduced compared to the uncorrected emulated 2DVD measurements. The underestimation of bounding-box width that was present with the uncorrected measurements, however, is more common after the correction. The increase in underestimation is due to the assumption that the top-most and bottom-most pixels of the 2DVD-imaged particle are supposed to be vertically aligned with one another. This assumption is frequently broken by snow particles and an , leaving the 2DVD measurements of snowflake horizontal extend potentially unreliable. An example of this assumption resulting in a very minor underestimation of bounding-box width can be seen in Fig. 2.

6 Discussion

Previous studies have examined a number of error sources that impact the accuracy of optical disdrometer measurements. To better understand the magnitude of the differences examined in the present study, it is worth a brief comparison to the findings of past studies.

Simulating two-dimensional observations of three-dimensional particles for the SVI (the predecessor of PIP) and 2DVD, Wood et al. (2013) found that, under reasonable conditions, using a two-dimensional projected long dimension length would result in the observed long dimension being approximately 82% of the actual particle long dimension, thereby underestimating reflectivity by 3.2 dB. Using three-dimensional snowflake replicas, Leinonen et al. (2021) determined that the two-dimensional measurements of particle long dimension made by MASC have a normalized root-mean-square error of ∼6% relative to the actual particle long dimension. Leinonen et al. (2021) only examined measurements of snowflakes with a long dimension between 3 mm and 5 mm and the study used the largest long dimension measured among the three MASC cameras rather than using a single-camera measurement. By comparison, we found the normalized root mean square difference between the long dimension of the tensor-fitted ellipse and the PIP-fitted ellipse, PIP-fitted rectangle, and MASC-fitted ellipse to be

~26%, ~17%, and ~7%, respectively, and these numbers do not change much when only particles with a tensor-fitted ellipse long dimension between 3 mm and 5 mm are included in the calculation. This suggests that the PIP shape-fitting issues have a much greater impact on the determination of the particle long dimension than does the loss of information caused by projecting a three-dimensional snowflake into a two-dimensional view. Interestingly, it also suggests that differences between the MASC-fitted ellipse and tensor-fitted ellipse are slightly larger than the error introduced by using the two-dimensional projection of the particle to measure long dimension length, although this could be due to Leinonen et al. (2021) using the maximum long dimension from among the three MASC cameras rather than just using a single camera.

Fitch et al. (2021) found that the largest snowflakes are less likely to be captured by MASC in strong winds (> 5 m s^{-1}) than other snowflakes. This bias towards collecting smaller particles was especially prevalent when the MASC did not have any form of wind shielding. For the original tall-form 2DVD, however, Nešpor et al. (2000) found that small raindrops were more likely to be caught in the turbulence produced by the wind passing over the instrument and that, under certain conditions, this could prevent these small raindrops from reaching the sampling volume. It is not clear from their study, however, whether this sampling bias would extend to frozen precipitation nor the extent to which the newer low-profile 2DVD addresses this issue. Newman et al. (2009) demonstrated that the SVI instrument setup produces negligible wind shielding effects at the sampling volume. As the PIP setup is virtually identical to the SVI setup, this finding should also apply to PIP. As detailed in the previous section, the present study examined the impact of wind on the snowflake measurements in terms of motion blurring. While we have not specifically examined the expected horizontal particle motions within the sampling volumes of each instrument, it is expected that the open sampling volume of the PIP would make it more prone to horizontal motion blurring than would the closed sampling volumes of 2DVD and MASC. Recent work using observations from a PIP located in Haukelister, Norway, demonstrated reliable results during moderate to high wind conditions for multiple snowfall events over a winter season (Shates et al., 2021). Additionally, these PIP observations were used in conjunction with a MASC to retrieve liquid water equivalent snowfall rates that corresponded well with accumulation observations from a collocated double fence automated reference standard (Schirle et al., 2019; Wolff et al., 2015). The findings of these studies suggest that the greater potential for horizontal motion blurring of an open sampling volume compared to a closed sampling volume may be relatively minor, at least for snow accumulation applications.

7 Conclusions

The present study set out to evaluate the measurement techniques used by PIP, MASC, and 2DVD to characterize the size and shape of snowflakes in terms of area, equivalent diameter, and aspect ratio. This was accomplished by using imagery collected by a PIP instrument during the ICE-POP 2018 field campaign to emulate the relevant measurements made by MASC and 2DVD. More specifically, a comparison was made between the ellipse- and rectangle-fitting algorithms used by PIP and the ellipse-fitting algorithm used by MASC; a tensor-based ellipse-fitting algorithm was also evaluated ~~along-side~~ alongside the PIP and MASC shape-fitting algorithms. Additionally, measurements of the bounding-box width and height made by PIP

were compared to the emulated 2DVD width and height measurement in order to evaluate the performance of 2DVD when measuring snow particles. The key findings of our comparative evaluation are listed below.

- Even when exposed to extreme motion blurring (10 m s^{-1} particle motion), all instruments had an area bias of less than 10% and a equivalent diameter bias of less than 5% for particles with an actual equivalent diameter larger than 5 mm (section 5.1).

- The 2DVD camera setup ~~should provide~~ provides the most accurate measure of area (and equi-area diameter) due to its resilience to motion blurring~~; that said,~~, although the relatively low ~~pixel resolution means that the 2DVD area and equi-area diameter measurements should also have low precision~~ resolution may impact the precision of those measurements (section 5.1).

- The ~~very high resolution MASC cameras should provide the greatest precision in area measurements; despite being susceptible to motion blurring, the MASC cameras should also produce the most accurate measurements of~~ accuracy of PIP area and equivalent diameter ~~(section 5.1).~~

- ~~Although the PIP camera setup is slightly~~ measurements is dependant on the direction of motion blurring due to the vertical image compression used by PIP; PIP is more resilient to ~~motion blurring than MASC, the application of an image compression algorithm that involves vertically averaging pixels means that the PIP measurements of area and equivalent diameter are slightly less accurate than MASC and can potentially have a relatively low precision similar to that of 2DVD~~ horizontal motion blurring and less resilient to vertical motion blurring in comparison to MASC (section 5.1).

- The MASC ellipse-fitting algorithm produces reasonable ellipse fits of the particles and, as such, produces reliable aspect ratios (section 5.2).

- The PIP shape-fitting algorithms do not perform well due to their reliance on only the area and perimeter of a particle; this renders the following variables unreliable: 'Ellip_Maj', 'Ellip_Min', 'Rec_BS', and 'Rec_SS' (section 5.2); other variables reported by the PIP are unaffected by the shape-fitting issues.

- The ~~tensor-based ellipse fitting algorithm~~ tensor method performed comparably to the MASC ellipse-fitting algorithm ~~; as such, any aspect ratio measurements using this algorithm should be reliable~~ (section 5.2).

- Our implementation of the 2DVD unskewing algorithm, designed to remove the effects of horizontal motion from the 2DVD bounding-box-width measurements, performed surprisingly well with snow particles considering the instrument was designed to measure rain drops~~(section 5.3)~~. That said, the corrected bounding-box-width measurements are still prone to error due to the motion skewing effects. We expect that the actual, proprietary implementation of this algorithm ~~should~~ would have a similar level of performance (section 5.3).

These results suggest that all three instruments have applications for studying snow, albeit with differing degrees of effectiveness. The ~~two-dimensional images collected by~~ MASC and PIP cameras, which capture two-dimensional images directly, are better suited for measuring snowflake shape than the line-scan cameras used by 2DVD as the MASC and PIP cameras are ~~ideal for measuring snow as they are~~ not susceptible to the image skewing produced by the horizontal motion of a particle that is experienced by 2DVD. That said, the very short exposure time of the 2DVD camera ~~effectively results in complete immunity~~ results in high resilience to motion blur, making any measurements of two-dimensional projected particle area highly accurate. ~~Despite the high accuracy of the 2DVD area measurements, its inability to measure the maximum dimension or aspect ratio of a particle (outside of horizontally or vertically oriented particles) limits the usefulness of its measurements.~~

~~The appropriateness of PIP and MASC for characterizing snowflakes largely comes down to the planned application for the data. The very high resolution imagery produced by MASC is ideal for determining ice particle habit, something that is impractical using the low-resolution PIP imagery. PIP, however, captures considerably more particles than MASC does and this greater particle count makes PIP better suited to constructing particle distributions than MASC. Unfortunately, the PIP~~ PIP suffers from the issue that the shape-fitting routines do not perform well on precipitation particles and, as a result, the ~~maximum dimension and aspect ratio~~ PIP-fitted ellipse and rectangle dimension (and, therefore, aspect ratio) measurements are unreliable~~; however, the PIP variables other than the ellipse and rectangle dimensions~~. Worth noting, however, is that the other PIP variables appear to be reliable. As a result, any non-shape-fitting-derived values (e.g., Pettersen et al., 2020) will also remain reliable. As the present study has demonstrated, the PIP imagery can be reprocessed and reliable measurements of ~~maximum~~ the long dimension and aspect ratio can be made via the application of an alternative ellipse-fitting algorithm, such as the MASC or tensor-based algorithms. While not demonstrated here, it may be possible to also implement a shape-fitting algorithm for 2DVD using the reconstructed images captured by the line-scan camera, although the reliability of the resulting shape measurements from such an algorithm would need further investigation to test the impacts of the image skewing.

Data availability. Bliven, Larry. 2020. GPM Ground Validation Precipitation Imaging Package (PIP) ICE POP. Dataset available online from the NASA Global Hydrology Resource Center DAAC, Huntsville, Alabama, U.S.A. DOI: http://dx.doi.org/10.5067/GPMGV/ICEPOP/PIP/DATA101

Appendix A: Derivation of shape-fitting equations for PIP

According to the IMAQ Vision concepts manual (National Instruments, 2000), the equations to fit either an ellipse or a rectangle can be derived from the equations for area and perimeter. The area and perimeter equations for a rectangle are, respectively,

$$A = ab \quad \text{and} \tag{A1}$$

$$P = 2(a + b), \tag{A2}$$

where A is the actual particle area, P is the actual particle perimeter, and a and b are the lengths of the long and short sides of the fitted rectangle. For an ellipse, PIP uses the following equations for area and perimeter, respectively:

$$A = \pi ab \quad \text{and} \tag{A3}$$

$$P = \pi \sqrt{2\left(a^2 + b^2\right)}, \tag{A4}$$

where the major axis length is equal to $2a$ and the minor axis length is equal to $2b$. Note, the equation PIP uses for the perimeter of an ellipse is actually an upper bound on the perimeter of an ellipse rather than an actual estimate of the perimeter (Jameson, 2014).

The derivation for the equations used to compute the dimensions of both the PIP-fitted ellipse and rectangle proceed similarly: solving the area equation for either a or b and substituting the result into the perimeter equation. This substitution is followed by putting the resulting equation into the standard form for a quadratic equation for the rectangle and a quartic equation for the ellipse. The standard form quadratic equation for the rectangle fit is

$$0 = -x^2 + \frac{P}{2}x - A, \tag{A5}$$

where x is either a or b depending on the variable for which the area equation was solved. Similarly, the standard form quartic equation for the ellipse fit is

$$0 = -\pi^2 x^4 + \frac{P^2}{2}x^2 - A^2. \tag{A6}$$

Note that the IMAQ Vision concepts manual incorrectly records the standard form equation for the ellipse fit as being identical to that of the rectangle fit.

The standard form equations for both the rectangle- and ellipse-fits can then be solved for their roots. As the ellipse-fit equation is a biquadratic special case of a quartic equation, it can be solved using the same method as a quadratic equation by substituting $z = x^2$, determining the roots, and then substituting $x = \pm\sqrt{z}$ into those roots. Solving both equations via the these methods gives roots of

$$x = \frac{P}{4} \pm \frac{\sqrt{\frac{P^2}{4} - 4A}}{2} \tag{A7}$$

for the PIP-fitted rectangle and

$$x = \pm\sqrt{\frac{P^2}{4\pi^2} \pm \frac{\sqrt{\frac{P^4}{4} - 4\pi^2 A^2}}{2\pi^2}} \tag{A8}$$

for the PIP-fitted ellipse. The only physically meaningful roots will be the positive roots. In the case of the quartic ellipse equation, one of the positive roots will give the value of a and the other will give the value of b. The other dimension for the rectangle fit can be computed by plugging the positive root into the equation for area and solving for the unknown dimension length (either a or b).

Table B1. List of relevant variables within PIP data files for the reprocessing PIP AVI video files.

Variable	~~file~~ File	Description
recnum	_a_p.dat, _a_v_1.dat, _a_v_2.dat	Record number of particle measurement, unique within each 10-minute period
Q, R	_a_p.dat	Frame index split into two numbers, unique within each 10-minute period.
x_cent, y_cent	_a_p.dat	PIP-measured particle centroid location in pixels
vel_v_1, vel_v_2	_a_v_1.dat, _a_v_2.dat	PIP-measured particle vertical fall speed in m s^{-1}

Appendix B: PIP video file processing

740   The following builds on section 4 by further discussing the steps taken to reprocess the PIP video files in order to produce the emulated MASC and tensor method fits as well as the emulated 2DVD measurements. Reprocessing the PIP video files for the alternative fits requires four sets of files: the AVI video files, the '_a_p.dat' particle table files, the '_a_v_1.dat' velocity table files, and the '_a_v_2.dat' velocity table files. The two velocity table files cover different sets of particles: the '_a_v_1.dat' files contain data for particles that were only observed in two consecutive frames while the '_a_v_2.dat' files contain data for
745   all particles that were observed in three or more consecutive frames. Using both files maximizes the number of particles that we can use to make emulated 2DVD measurements. Each AVI file corresponds to one of each of the '.dat' files, although the '.dat' files typically cover a much greater period of time than what is depicted in the AVI files. The key '.dat' file variables for the reprocessing algorithm are listed in Table B1.

The PIP AVI videos include both the PIP view as well as a header that indicates which PIP unit recorded the observations,
750   the time stamp, a combined 'Q_R' frame index, and the number of the frame within the AVI video file (Fig. B1). The key to matching a single frame of the AVI file to the data contained within the '.dat' files is matching the single-number 'Q_R' frame index depicted in the video to the two-number 'Q' and 'R' frame index listed in the '_a_p.dat' files; 'Q_R' can be computed as 'Q' times 32767 plus 'R'. To obtain the 'Q_R' frame index from the video image, we constructed a simple image matching algorithm. The image matching algorithm compares, via subtraction, the image of each digit of the 'Q_R' frame index to a set
755   of images of individual digits that were previously cropped from PIP AVI videos. A matched digit should result in a difference of zero, although we selected the digit that produces the smallest difference and allowed for a difference of up to 100, noting that the 'Q_R' frame index can be between one and six digits long and each additional digit shifts the position of all the numbers in the frame by several pixels.

Once the 'Q_R' frame index is determined from the video file and it has been matched to the 'Q' and 'R' frame index in the
760   '_a_p.dat' file, the record number ('recnum') in the '_a_p.dat' file can be matched to the record number in the '_a_v_1.dat' and '_a_v_2.dat' velocity table files. Note, only a subset of particles listed in the '_a_p.dat' file will have matching entries in the velocity table files. Following the above method, we should now have the centroid location and fall speed of each particle in a given frame of the PIP AVI video file. The centroid is used to identify all pixels within a particle, as described in section 4, and the fall speed is used to emulate the 2DVD measurements, as described in section 4.2.

Author contributions. Dr. Helms led this research by developing and running the code to reprocess the PIP video files, implementing the various measurement algorithms, and performing the various analyses. Dr. Munchak and Dr. Tokay provided input on the methods used in the study and, critically, tested the reprocessed PIP data for quality purposes. Dr. Pettersen provided data support and advice on reprocessing the PIP images.

Competing interests. Some authors are members of the editorial board of Atmospheric Measurement Techniques. The peer-review process was guided by an independent editor. The authors have no other competing interests to declare.

Acknowledgements. The authors would like to thank ~~Dr.~~ David Wolff for the fruitful discussions regarding various aspects of this manuscript. Additionally, the authors would ~~also~~ like to thank ~~Drs.~~ Larry Bliven and Jacopo Grazioli for their insights into the inner workings of the PIP and MASC instruments, respectively. The authors are also indebted to the ICE-POP 2018 field campaign participants for their data collection and processing efforts. Dr. Helms' work was supported by an appointment to the NASA Postdoctoral Program at NASA Goddard Space Flight Center, administered by Universities Space Research Association under contract with NASA. Dr. Pettersen's work was supported via NASA grants NASA 80NSSC19K0712 and NASA 80NSSC21K0931.

References

Fitch, K. E., Hang, C., Talaei, A., and Garrett, T. J.: Arctic observations and numerical simulations of surface wind effects on Multi-Angle Snowflake Camera measurements, Atmos. Meas. Tech., 14, 1127–1142, https://doi.org/10.5194/amt-14-1127-2021, 2021.

780 Fitzgibbon, A. W. and Fisher, R. B.: A Buyer's Guide to Conic Fitting, in: Proceedings of the British Machine Vision Conference, pp. 513–522, BMVA Press, https://doi.org/10.5244/C.9.51, 1995.

Garrett, T. J., Fallgatter, C., Shkurko, K., and Howlett, D.: Fall speed measurement and high-resolution multi-angle photography of hydrometeors in free fall, Atmos. Meas. Tech., 5, 2625–2633, https://doi.org/10.5194/amt-5-2625-2012, 2012.

Grecu, M., Olson, W. S., Munchak, S. J., Ringerud, S., Liao, L., Haddad, Z., Kelly, B. L., and McLaughlin, S. F.: The GPM Combined
785 Algorithm, J. Atmos. Oceanic Technol., 33, 2225–2245, https://doi.org/10.1175/JTECH-D-16-0019.1, 2016.

Han, M. and Braun, S. A.: Understanding the Global Three-dimensional Distribution of Precipitation Mean Particle Size with the Global Precipitation Measurement Mission, J. Climate, 34, 9775–9796, https://doi.org/10.1175/JCLI-D-21-0134.1, 2021.

Heymsfield, A. J. and Westbrook, C. D.: Advances in the estimation of ice particle fall speeds using laboratory and field measurements, J. Atmos. Sci., 67, 2469–2482, https://doi.org/10.1175/2010JAS3379.1, 2010.

790 Huang, G.-J., Bringi, V. N., Cifelli, R., Hudak, D., and Petersen, W. A.: A methodology to derive radar reflectivity-liquid equivalent snow rate relations using C-band radar and a 2D video disdrometer, J. Atmos. Oceanic Technol., 27, 637–651, https://doi.org/10.1175/2009JTECHA1284.1, 2010.

Jameson, G.: Inequalities for the perimeter of an ellipse, The Mathematical Gazette, 98, 227–234, https://doi.org/10.1017/S002555720000125X, 2014.

795 Jiang, Z., Oue, M., Verlinde, J., Clothiaux, E. E., Aydin, K., Botta, G., and Lu, Y.: What can we conclude about the real aspect ratios of ice particle aggregates from two-dimensional images?, J. Appl. Meteor. Climatol., 56, 725–734, https://doi.org/10.1175/JAMC-D-16-0248.1, 2017.

Korolev, A. and Isaac, G.: Roundness and Aspect Ratio of Particles in Ice Clouds, J. Atmos. Sci., 60, 1795–1808, https://doi.org/10.1175/1520-0469(2003)060<1795:RAAROP>2.0.CO;2, 2003.

800 Kruger, A. and Krajewski, W. F.: Two-Dimensional Video Disdrometer: A description, J. Atmos. Oceanic Technol., 19, 602–617, https://doi.org/10.1175/1520-0426(2002)019<0602:TDVDAD>2.0.CO;2, 2002.

Lee, G. and Kim, K.: International Collaborative Experiments for PyeongChang 2018 Olympic and Paralympic winter games (ICE-POP 2018), in: AGU Fall Meeting, American Geophysical Union, San Francisco, CA, 2019.

Leinonen, J., Grazioli, J., and Berne, A.: Reconstruction of the mass and geometry of snowfall particles from multi angle snowflake camera
805 (MASC) images, Atmos. Meas. Tech., 14, 6851–6866, https://doi.org/10.5194/amt-14-6851-2021, 2021.

Lim, K.-S. S., Chang, E.-C., Sun, R., Kim, K., Tapiador, F. J., and Lee, G.: Evaluation of simulated winter precipitation using WRF-ARW during the ICE-POP 2018 field campaign, Weather and Forecasting, 35, 2199–2213, https://doi.org/10.1175/WAF-D-19-0236.1, 2020.

Locatelli, J. D. and Hobbs, P. V.: Fall speeds and masses of solid precipitation particles, J. Geophys. Res., 79, 2185–2197, https://doi.org/10.1029/JC079i015p02185, 1974.

810 Matrosov, S. Y., Heymsfield, A. J., and Wang, Z.: Dual-frequency radar ratio of nonspherical atmopsheric hydrometeors, Geophys. Res. Lett., 32, L13 816, https://doi.org/10.1029/2005GL023210, 2005.

Munchak, S. J., Schrom, R. S., Helms, C. N., and Tokay, A.: Snow Microphysical Retrieval from the NASA D3R Radar During ICE-POP 2018, Atmos. Meas. Tech., 15, 1439–1464, https://doi.org/10.5194/amt-15-1439-2022, 2022.

National Instruments: IMAQ Vision Concepts Manual, National Instruments, [Available online at https://www.ni.com/pdf/manuals/322916a.pdf], 2000.

Nešpor, V., Krajewski, W. F., and Kruger, A.: Wind-induced error of raindrop size distribution measurement using a two-dimensional video disdrometer, J. Atmos. Oceanic Technol., 17, 1483–1492, https://doi.org/10.1175/1520-0426(2000)017<1483:WIEORS>2.0.CO;2, 2000.

Newman, A. J., Kucera, P. A., and Bliven, L. F.: Presenting the snowflake video imager (SVI), J. Atmos. Oceanic Technol., 26, 167–179, https://doi.org/10.1175/2008JTECHA1148.1, 2009.

Pettersen, C., Bliven, L. F., von Lerber, A., Wood, N. B., Kulie, M. S., Mateling, M. E., Moisseev, D. N., Munchak, S. J., Petersen, W. A., and Wolff, D. B.: The Precipitation Imaging Package: Assessment of microphysical and bulk characteristics of snow, Atmosphere, 11, 785, https://doi.org/10.3390/atmos11080785, 2020.

Pettersen, C., Bliven, L. F., Kulie, M. S., Wood, N. B., Shates, J. A., Anderson, J., Mateling, M. E., Petersen, W. A., von Lerber, A., and Wolff, D. B.: The precipitation imaging package: Phase partitioning capabilities, Remote Sensing, 13, 2183, https://doi.org/10.3390/rs13112183, 2021.

Schirle, C. E., Cooper, S. J., Wolff, M. A., Pettersen, C., Wood, N. B., L'Ecuyer, T. S., Ilmo, T., and Nygård, K.: Estimation of snowfall properties at a mountainous site in Norway using combined radar and in situ mircrophysical observations, J. Appl. Meteor. Climatol., 58, 1337–1352, https://doi.org/10.1175/JAMC-D-18-0281.1, 2019.

Schönhuber, M., Lammer, G., and Raneu, W. L.: One decade of imaging precipitation measurement by 2D-video-distrometer, Adv. Geosci., 10, 85–90, 2007.

Schönhuber, M., Lammer, G., and Randeu, W. L.: The 2D-video-distrometer, in: Precipitation: Advances in measurement, estimation and prediction, edited by Michaelides, S., pp. 3–31, Springer, 2008.

Shates, J. A., Pettersen, C., L'Ecuyer, T. S., Cooper, S. J., Kulie, M. S., and Wood, N. B.: High-latitude precipitation: Snowfall regimes at two distinct sites in Scandinavia, J. Appl. Meteor. Climatol., 60, 1127–1148, https://doi.org/10.1175/JAMC-D-20-0248.1, 2021.

Shkurko, K., Garrett, T., and Gaustad, K.: Multi-Angle Snowflake Camera Value-Added Product, DOE/ARM Tech. Rep. DOE/SC-ARM-TR-187, DOE Office of Science Atmospheric Radiation Measurement (ARM) Program, https://doi.org/10.2172/1342901, 2016.

Stuefer, M. and Bailey, J.: Multi-Angle Snowflake Camera Instrument Handbook, DOE/ARM Tech. Rep. DOE/SC-ARM-TR-158, DOE Office of Science Atmospheric Radiation Measurement (ARM) Program, https://doi.org/10.2172/1261185, 2016.

Tapiador, F. J., Villalba-Pradas, A., Navarro, A., García-Ortega, E., Lim, K.-S. S., Ahn, K. D., and Lee, G.: Future directions in precipitation science, Remote Sensing, 13, 1074, https://doi.org/10.3390/rs13061074, 2021.

Testud, J., Oury, S., Black, R. A., Amayenc, P., and Dou, X.: The concept of "normalized" distribution to describe raindrop spectra: A tool for cloud physcics and cloud remote sensing, J. Appl. Meteor., 40, 1118–1140, https://doi.org/10.1175/1520-0450(2001)040<1118:TCONDT>2.0.CO;2, 2001.

Tyynelä, J., Leinonen, J., Moisseev, D., and Nousiainen, T.: Radar Backscattering from Snowflakes: Comparison of Fractal, Aggregate, and Soft Spheroid Models, J. Atmos. Oceanic Technol., 28, 1365–1372, https://doi.org/10.1175/JTECH-D-11-00004.1, 2011.

Vázquez-Martín, S., Kuhn, T., and Eliasson, S.: Shape dependence of snow crystal fall speed, Atmos. Phys. Chem., 21, 7545–7565, https://doi.org/10.5194/acp-21-7545-2021, 2021.

von Lerber, A., Moisseev, D., Bliven, L. F., Petersen, W., Harri, A.-M., and Chandrasekar, V.: Microphysical properties of snow and their link to Ze-S relations during BAECC 2014, J. Appl. Meteor. Climatol., 56, 1561–1582, https://doi.org/10.1175/JAMC-D-16-0379.1, 2017.

Westbrook, C. D.: Rayleigh scattering by hexagonal ice crystals and the interpretation of dual-polarization radar measurements, Quart. J. Roy. Meteor. Soc., 140, 2090–2096, https://doi.org/10.1002/qj.2262, 2014.

Wolff, M. A., Isaksen, K., Petersen-Øverleir, A., Ødemark, K., Reitan, T., and Brækkan, R.: Derivation of a new continuous adjustment function for correcting wind-induced loss of solid precipitation: Results of a Norwegian field study, Hydrol. Earth Syst. Sci., 19, 951–967, https://doi.org/10.5194/hess-19-951-2015, 2015.

855   Wood, N. B., L'Ecuyer, T. S., Bliven, F. L., and Stephens, G. L.: Characterization of video disdrometer uncertainties and impacts on estimates of snowfall rate and radar reflectivity, Atmos. Meas. Tech., 6, 3635–3648, https://doi.org/10.5194/amt-6-3635-2013, 2013.
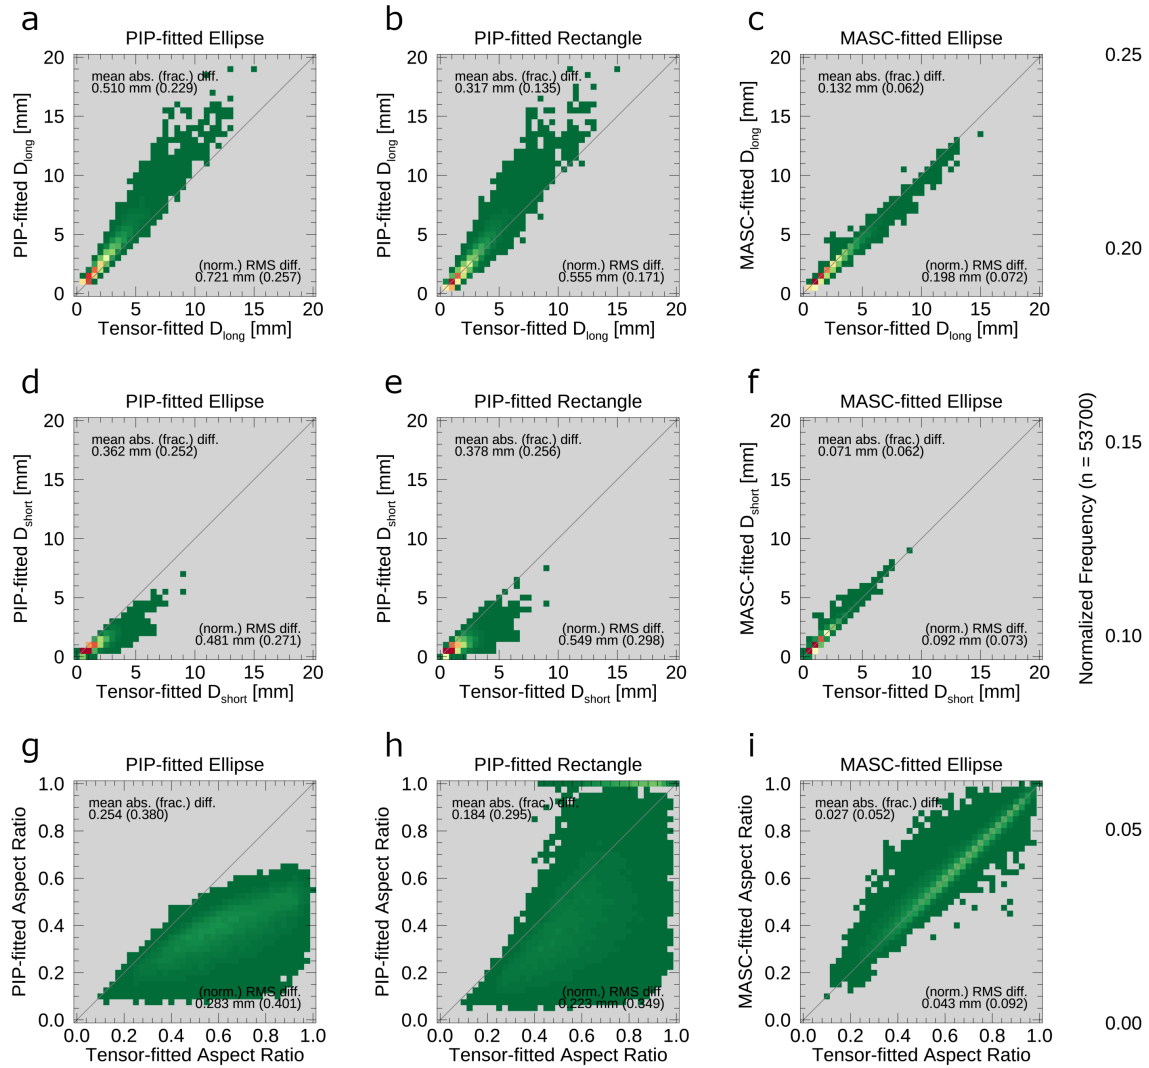
Figure 5. ~~Scatterplots~~ Two-dimensional histograms of (a-c) particle long dimension in mm, (d-f) particle short dimension in mm, and (g-i) particle aspect ratio for (a,d,g) the PIP-fitted ellipse, (b,e,h) the PIP-fitted rectangle, and (c,f,i) the emulated MASC-fitted ellipse as a function of the tensor-fitted ellipse values. ~~Data points are indicated by open circles~~ The diagonal grey lines indicate where the x value is equal to the y value. The mean absolute difference and ~~only data~~ root-mean-square difference for each pair of measurements is indicated on each panel with the mean absolute fractional difference and normalized root-mean-square difference in parentheses. Only particles with a PIP-measured area greater than 0.5 mm^2 are included. ~~PIP data was collected between 0000 UTC March 2018 and 0000 UTC 9 March 2018 at~~ Note, a logarithmic color table has been used to highlight the ~~ICE-POP MHS site~~ frequency distribution at lower values.
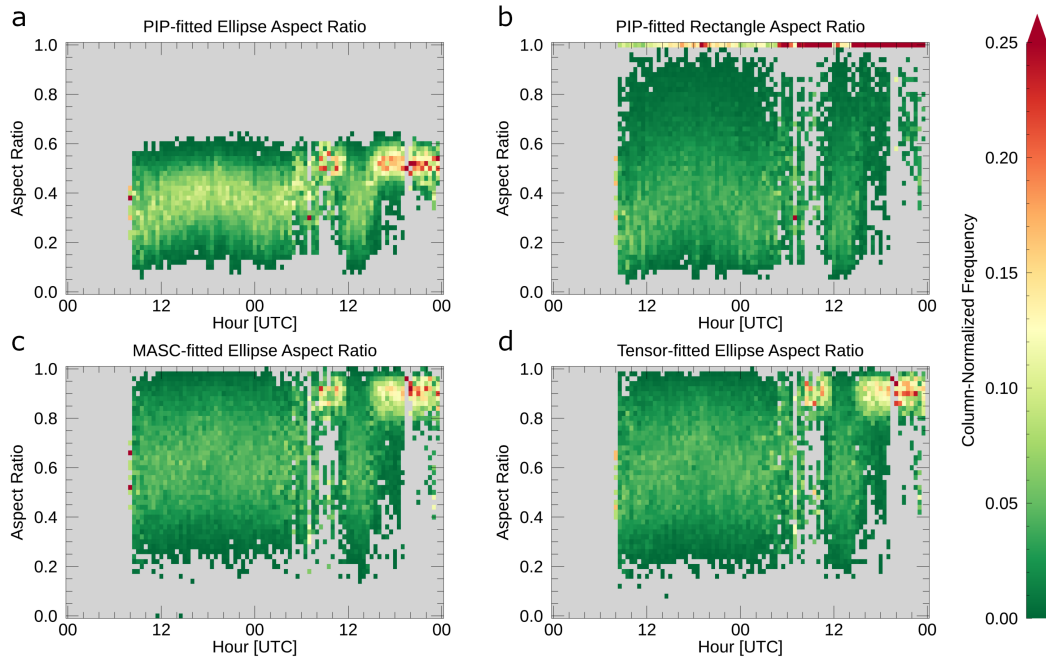
Figure 6. Time series of the (a) PIP-fitted ellipse and (b) PIP-fitted rectangle, (c) emulated MASC-fitted ellipse, and (d) tensor-fitted ellipse aspect ratio distributions for the 7–8 March 2018 case. Only particles with an area of at least 0.5 mm^2 were included in the histograms. The bin counts of each column have been normalized such that they add up to one.
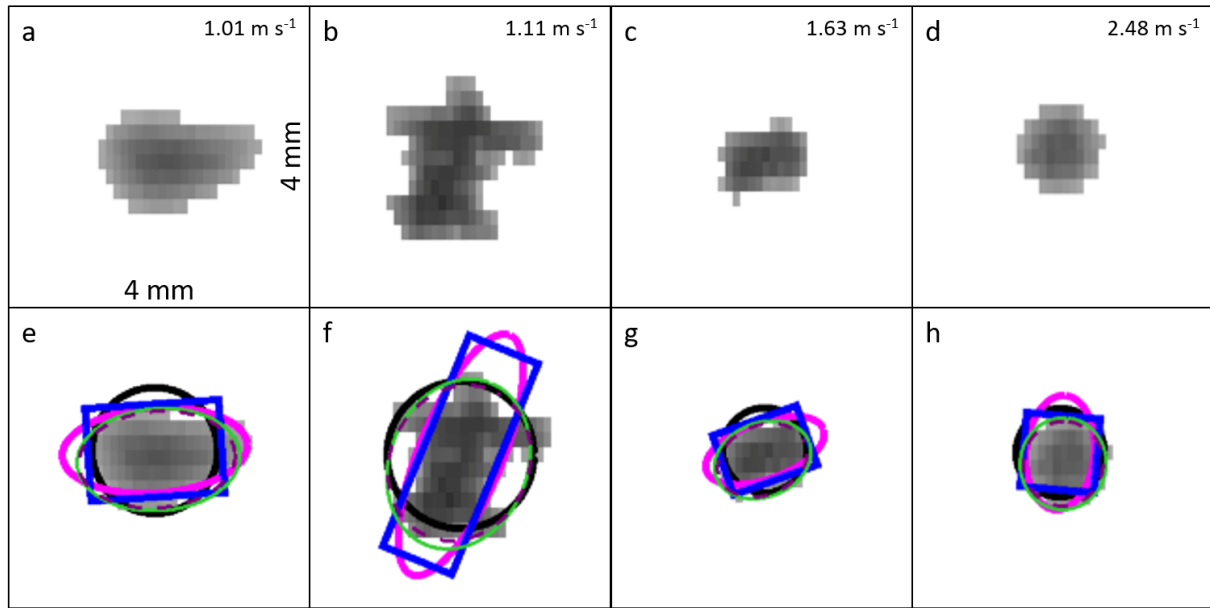
Figure 7. Snapshots of ~~three~~ four particles as viewed by PIP collected by PIP at approximately (a,b,e,f) 1904 UTC 7 March 2018 and (c,d,g,h) 1756 UTC 8 March 2018 with (e-h) overlaid particle shape estimates as determined by the (black) area-equivalent circle, (magenta) PIP-fitted ellipse, (blue) PIP-fitted rectangle, (purple dashed) emulated MASC-fitted ellipse, and (green) tensor-fitted ellipse. Each panel represents a square that is 4-mm long on each side. The dimensions of the fitted shapes for each particle are presented in Table 1. Fall speeds are listed in the top right corner of (a-d); the fall speeds for particles (a) and (b) are computed by PIP while the fall speeds for particles (c) and (d) are manually calculated as the average frame-to-frame center position motion as these specific particles were not included in the PIP particle velocity files.
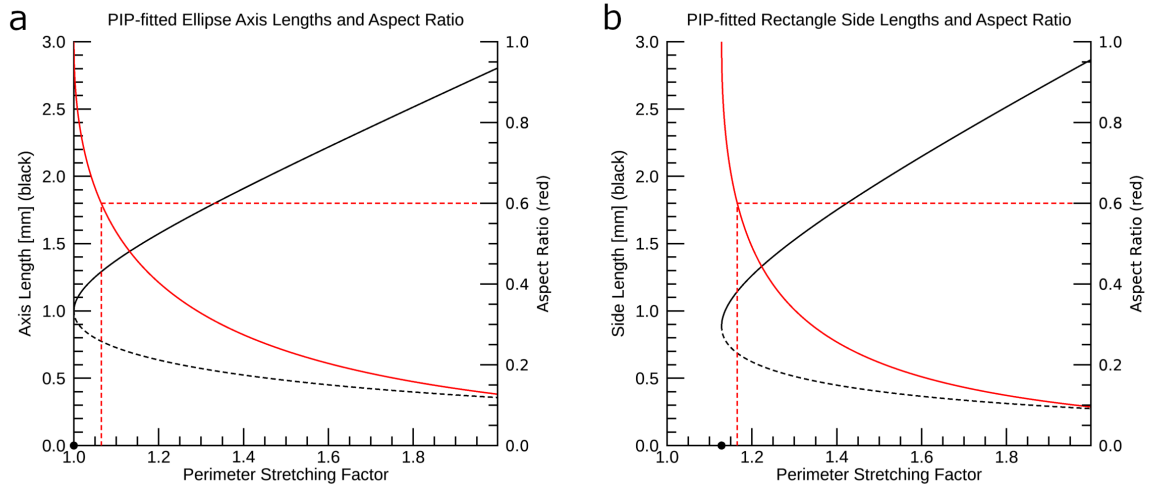
Figure 8. Demonstration of the impact of excess perimeter for a given area on the (a) PIP-fitted ellipse and (b) PIP-fitted rectangle long-axis length (black solid), short-axis length (black dashed), and aspect ratio (red solid). This example uses a circle with a radius of 0.5 mm to compute the base perimeter and area. The ~~perimeter of this circle is then multiplied by the perimeter stretching factor before the ellipse or rectangle fitting is performed according to the method used by PIP. The~~ red dashed lines highlight the perimeter stretching factor that corresponds to an aspect ratio of 0.6 and the black dot indicates the minimum perimeter stretching factor for which a physically meaningful shape can be fitted.
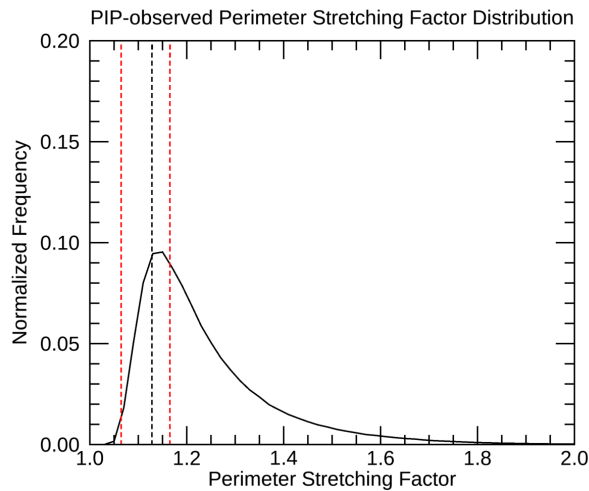
Figure 9. Distribution of perimeter stretching factors computed from PIP observations collected during the 7–8 March 2018 snow event. The two red vertical dashed lines indicate the perimeter stretching factor corresponding to an aspect ratio of 0.6 for a (left red dashed line) PIP-fitted ellipse and (right red dashed line) PIP-fitted rectangle. The black vertical dashed line indicates the smallest perimeter stretching factor that will permit a physically meaningful rectangle fit. Only particles with an area of at least 0.5 mm^2 are included in the histogram.
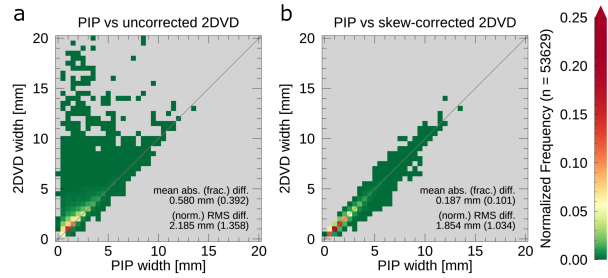
Figure 10. ~~Scatter plots~~ Two-dimensional histograms comparing the PIP ~~and 2DVD measurements of~~ bounding-box width to the (a) ~~bounding-box height~~ uncorrected and (b,~~c~~) skew-corrected 2DVD bounding-box width. The mean absolute difference and root-mean-square difference for each pair of measurements is indicated on each panel with the ~~uncorrected 2DVD-measured width depicted in (b)~~ mean absolute fractional difference and ~~the corrected 2DVD-measured width depicted~~ normalized root-mean-square difference in ~~(c)~~parentheses. Only particles with a PIP-measured area of at least ~~5~~ 0.5 mm^2 are included in the plots. ~~All data points are plotted with open circles~~Note, a logarithmic color table has been used to highlight the frequency distribution at lower values.
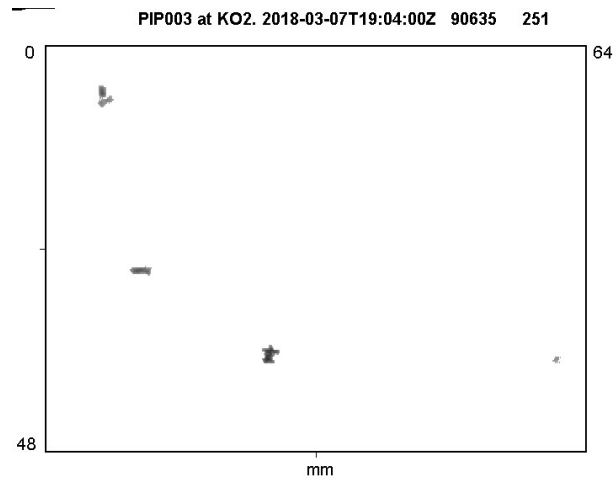
Figure B1. Snapshot taken from a PIP AVI video file collected by the PIP located at the MHS site during ICE-POP. The header can be decoded as follows: 'PIP003 at KO2' indicates which PIP recorded the file, '2018-03-07T12:04:00Z' is the time stamp, '90635' is the 'Q_R' frame index within the current 10-minute period, and '251' is the frame number within this specific AVI file. Note, the black mark in the top-left corner is a digital tag that can be used to link the quicklook image in the AVI file back to the raw data file.