Review of Introducing the Video In Situ Snowfall Sensor (VISSS)
Reviewed by Charles Helms

This manuscript describes the VISSS instrument. VISSS is a new video-based precipitation microphysics probe designed to capture high-resolution images of snowflakes using a pair of orthogonally-pointing high-speed cameras. The manuscript also compares the VISSS measurements to those of two other precipitation microphysics probes: PIP and Parsivel. Overall, I found the manuscript to be of high quality, although there are some minor improvements that would further improve upon this quality.

As I only feel minor revisions are necessary, I've opted to put all my comments (some of which are simply small typo corrections) in the order they appear in the text.

Line 6 (also Lines 317 and 409): Is VISSS observing up to 100,000 unique particles per minute (i.e., a falling particle is only counted once during its transit across the domain) or is it making up to 100,000 particle observations per minute (i.e., each measurement of a particle is counted, even if that particle has been measured in a previous frame)? If it's the latter, I suggest changing the wording to "100,000 particle observations per minute"

Line 6: This is the first time PIP is mentioned; suggest moving the definition of the PIP acronym from Line 8 to here

Line 52 (also Line 378): The Del Guasta (2022) reference is inserted parenthetically here instead of being in-line (similar for the Battaglia et al. reference on line 378).

Line 57: The PIP acronym was already defined above, although, personally, I don't see any problem with it being defined a second time so I leave changing this up to author discretion.

Line 60 (and elsewhere): Maybe a bit pedantic, but 100 microns per pixel is the pixel size rather than the resolution (i.e., the minimum resolvable particle size).

Paragraph starting on Line 89: I really appreciate that the authors include the details about the camera in the text (and in table 1). The only additional piece of information the authors might consider adding is the type of camera (i.e., CCD, CMOS Global Shutter, CMOS Rolling Shutter, etc).

Line 113: suggest removing the word "also" to improve readability (authors' discretion)

Table 1: The pixel sizes are inconsistent in their use of "." or "," as a decimal point.

Table 1: If there is room for it in the table, I suggest changing "Exposure time" to "Effective exposure time" to make it clear that this is the duration of the LED being on rather than the actual exposure time of the camera itself

Figure 3: It looks like the blue ellipse only appears in the final annotated image. If this is intentional, it might be worth adding a note that the fitEllipse shape is only annotated on that image otherwise it would be helpful to either note that the blue ellipse is obscured by another ellipse (and indicate which ellipse this is) or use a dashed line for the blue ellipse and put it on top of the other ellipses.

Line 159: "sphere" should be changed to "circle"

Line 175: Should "vertical position" be "horizontal position"? I would think the vertical position information would be known from the leader camera. If not, this discrepancy might need a sentence or two of brief explanation.

Paragraph starting on Line 183: It's still not clear to me how the frame matching works. Is this a case of matching up the first frame that the particle appears in? If so, wouldn't this require the alignment of the two cameras to be extremely good (on the order of half a pixel or less, presumably)? Or is this more manually intensive than I'm thinking it is and the matching is based on matching up how the particle tumbles as it passes through the domain? Regardless, more details would be helpful. (A side thought that occurred to me while reading this: have you considered matching the frames up by using a camera flash while the LEDs are obscured? I'm not sure if that'll work or not, but figured I'd mention it anyway)

Line 213: "reader" should be "leader"

Section 3.4: I suggest removing the proof of concept tracking (and the related figure 4). It doesn't really add much to the manuscript and nearest-neighbor-based particle tracking is really bad outside of very low winds and/or very light snowfall. Depending on how much progress has been made, it might be worth adding some more information on the actual particle tracking algorithm the authors are developing.

Line 269 – 271: If multiple observations of a single particle are included in the PSD, wouldn't this bias the PSD towards slower falling particles?

Line 313: It took me several times reading these sentences to realize the "Fig 6. *c*" did not refer to panel c of figure 6, but that these are two completely separate thoughts; suggest replacing "c" with "Particle complexity" or "Particle complexity *c*" to avoid this issue.

Line 314: insert "particles" after "weighted to smaller"

Lines 336 – 354 (regarding the PIP underestimation of N0*): I don't think the dilation is the issue here. PIP's processing applies an edge detection filter, dilates the resulting image twice (using the 3x3 kernel), fills in any holes, and then erodes the hole-filled image twice to (theoretically) undo the dilation step (using the same 3x3 kernel). A few possibilities that come to my mind to explain the discrepancies between VISSS and PIP are: 1) the dilation is merging nearby needles into a single particle, thus decreasing the number of small particles (presumably this would be paired with an increase in larger particles); 2) the image compression (which averages vertically-adjacent pixels to reduce the data rates) is essentially destroying the smaller needles; and 3) that the dilation and hole-filling of higher complexity particles is artificially inflating the equivalent diameter and this is introducing a bias into the PSD moments used to compute N0* (presumably this one is less relevant for the needles, but later you mention PIP having issues with high complexity particles, so I included this as well). Unfortunately, I don't have any deeper insight as to which one of these might be the culprit (if it even is one of these), but I'll add this to the top of my list of PIP behaviors to look into. Either way, I don't think it would be the dilation itself as that adds to the particles.

From the VISSS side of things, if the PSD is being biased towards slower falling particles by including all particle observations, this might produce a bias at small particles in the VISSS PSD. It should be relatively easy to test this by comparing the VISSS PSD as it appears here to the PSD that would be produced if only particles appearing in every 25th frame are included in the PSD (this is how PIP computes its PSD, for reference). Under normal conditions, 25 frames should be more than enough time for any particles observed in a frame to exit the measurement volume.

Line 337: should “width” be “length”? In my mind the width of a particle is more closely aligned with (one of) the shorter axis of a particle, but when reading this it feels like the authors are referring to the longer axis of the needles.

Lines 342 – 343 and 352 – 353: As I mentioned above, dilation shouldn’t result in the removal of any parts of a snowflake as dilation expands the particles. Erosion could, in theory, remove parts, but I doubt it since the erosion step occurs after the particles have already been dilated twice.

Fig. 6: It might be helpful to also include the observed particles per minute for the PIP and Parsivel just to give a point of comparison. For PIP at least, the PSD is computed using the *_a_p_60.dat files (which only includes particles observed during every 25th frame to avoid double counting).

Additionally, if the authors’ wish to, the particle complexity can be computed from the PIP files by dividing the particle area by the hydraulic radius (Hy_Rad), which is the ratio between area and perimeter to get the particle perimeter and then plugging the relevant values into Eq 1. That said, as discussed in Helms et al. (2022) [section 3.1], the PIP software takes some potentially questionable (when applied to a snowflake) shortcuts when computing the perimeter, so comparing particle complexity between VISSS and PIP may not be particularly informative as to the accuracy of VISSS. I am less familiar with the specifics of the Parsivel output files, but similar methods may be possible there as well. Either way, I leave choice of adding this up to the authors’ discretion.

Fig. 7: It’s hard to make anything out on the image due to the small size of each frame. If might be beneficial to include a subset of these to make each frame larger so readers can better appreciate the resolution of the cameras.

Fig. 8: It would be helpful to have the time period over which the PSDs are computed for each of the instruments either in the caption and/or the text (apologies if this is in the text and I missed it, I added this comment after having read through the paper and didn’t see it mentioned when looking back through again).

Line 432: I certainly appreciate VISSS being open source! We (collectively) shouldn’t have to pull teeth to understand how instruments produce their measurements.

Appendix A: It took a couple attempts, but I was able to replicate the derivation in Appendix A.