Responses to Reviewer 3 Revised Manuscript Comments

The manuscript has improved significantly as a result of the authors responses to the referee comments. In particular the analysis has improved and details like density plots (2D histograms) in Figs 5, 10, as well as using mean difference as measure for comparison.
Some of the modifications include major changes and editing. In the following I have several remarks, comments, and questions regarding some of the changes. Please have a look and change and improve/clarify the manuscript.
Unless otherwise noted, i refer to Fig. and Lines as they are numbered in the version2 manuscript.

RE Major 3) (Precision)

L404-405: "As we established above, the 2DVD is the least affected by motion blur and, as a result, has the most accurate area measurements, although the large pixel size may impact the precision of those measurements."
I still beleive that "precision" is the wrong term to use here.
Larger pixels ultimately affect accuracy of determined dimensions (picel size will contribute to the error and this contribution will scale with the pixel size). Furthermore, the negative effects of pixelation will be intensified with larger pixels (with consequences on perimeter stretching, see below about "cap").
While I'm sure we could come to a consensus on this eventually, I've opted to remove it since it's not particularly critical to the paper. The new sentence reads:

"As we established above, the 2DVD is the least affected by motion blur and, as a result, has the most accurate area measurements, although it should be noted that 2DVD also has a relatively large pixel size." Lines 409 – 411

RE Major 4) Motion blur

Corrected.
Added horizontal blur at 10m/s?
Check new Fig 4 and discussion.

MASC 33.5um one pixel; 40us exposure
=> at 1m/s 40um blurred "streak" for any arbitrary point on particle
This blur streak is linearly (if motion along one pixel axis) on average on 40um/(33.5um/pixel)= 1.19 pixels.
This expected value is equal to that found by statistical analysis of 100000 random particles.
New L311-312: "would contribute to at least two pixels and possibly three, although our

statistical analysis indicates an average 1.19 extra particle pixels are added."
Regarding contributions to three pixels, it may be worth commenting: ...three: in that case the contributions to two of these pixels is minor.
Mention that 1.19 is exactly as expected!
Corresponds to the average extra pixels you show in brackets for 4m/s and 10m/s

Added note that the 1.19 is as expected and note about the contributions being minor in the three pixel case:

“With a pixel 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 the contribution to two of the pixels in the three-pixel case would be minor. Our statistical analysis indicates an average 1.19 extra particle pixels are added, which matches the expected value that can be calculated from the motion and pixel size.” Lines 314 – 317

L315: "For a particle moving at 10 m s−1..."
You seem to mean "For a particle moving horizontally at 10 m s−1..."
Note, the 400um blurred line is true for a vertically or horizontally moving particle (for normal camera orientation), not for any other oblique angle (i.e. the most likely motion). Nevertheless, I would not change the analysis by including oblique angles (would make it unnecessarily complicated).

I added “horizontally” as suggested (technically it also applies vertically, but that shouldn’t be happening under normal circumstances) as well as a note that the analysis doesn’t apply to oblique angles:

“For computational simplicity, these calculations assume that a pixel is part of the particle if it has at least 50% coverage and that the motion is not oblique to the pixels (i.e., only completely vertical or horizontal motion is considered).” Lines 309 – 311

“For a particle moving horizontally 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.” Lines 320 – 322

Is motion blur in horizontal and vertical directions treated separately for PIP (L319-321)?

Yes. The horizontal motion blur will be the pre-compression blur while the vertical blur will be the post-compression blur. To clarify this, I added a note:

“For horizontal motion, corresponding to the pre-compression blurring effects, PIP has a pixel size of 0.1 mm (100 μm).” 326 – 327

L353-356: "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."
"the 10 m s−1 particle motion" => "the horizontal 10 m s−1 particle motion"

I feel like this might cause confusion as the rest of the sentence is dealing with vertical motion.

"add 3.80 pixels after accounting for ..." => "add 3.80 pixels after accounting for ... and X pixels without compression (for horizontal motion)"

To match the wording in the rest of the paragraph, I change the wording of the sentence to read:

"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 2.80 pixels before accounting for the effects of image compression and 3.80 pixels after accounting for image compression." Lines 360 – 362

Fig 3 caption: "The darker blue shading indicates the area covered by the actual particle at the end of the image exposure period" => "... the start of the image exposure period"?

Since the particle is falling, the actual particle (darker blue shading) is located at the bottom, so I think "at the end of the image exposure period" is correct.

L 351, PIP: for suspended particles an average of 1.01 pixels added??
The 1.01 have not been mentioned or explained. (Maybe thsi refers to L345-346 ("the image compression can add between zero and two additional pixels to the height of a suspended particle.")
Comment: Even without pixel compression: a four pixel particle may result in a five pixel particle on picture (50% on lowest and 50% on highest pixel.

It appears that there was a minor statistical fluke; there should only be an average of one additional pixel. The compressed pixel values are the primary reason I went with a statistical approach rather than just computing an expected value as the compression effects are not straight forward to account for. The rest of the additional post-compression PIP pixel counts appear to be accurate. The sentence now reads:

"Because of our prescribed threshold of 50% coverage, the 1-m-s^{-1} fall speed case ends up producing the 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 1.0 pixels when suspended based on our statistical analysis" Lines 355 – 357

L366-372: The area error also depends on particle orientaion. A needle or column is more affected if oriented horizontally.

This is correct (see lines 372 – 373 of newly revised manuscript). I also considered adding a parenthetical note after shape to suggest habit as an example, but I decided against it as aggregate shape is also very important and I don't want to mislead the reader into focusing on habit alone.

RE Major 5) "Cap" Major 6) gap

I think you should discuss in appropriate places the problematic definitions of the PIP ellipse and rectangle and the consequences of these: long dimension of ellipse can be far longer than maximum dimension; pixelation leads to lowered aspect ratio for elipse; many particles for which rectangle fit is not possible. These problems explain the "cap" in elipse fitted aspect ratios and the apparent gap (in rectangle-fitted aspect ratios).

While I do discuss these topics throughout the paper (Appendix A goes in depth into the precise definition of the PIP ellipse and rectangle; Section 5.2 covers the problems with the shape fitting that result from how the PIP shapes are defined), I do agree that they should also appear in a brief summary at some point to remind the reader. To this end, the relevant bullet point in the conclusions section now reads:

"The PIP shape-fitting algorithms do not perform well due to their reliance on only the area and perimeter of a particle leading to a tendency towards overestimating the long dimension and underestimating the short dimension that is highly sensitive to small deviations from the smooth-edged fitted shape; 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." Lines 615 – 619

L621 "PIP suffers from the issue that the shape-fitting routines do not perform well on precipitation particles..."Where was this discussed in paper?
Does it refer to the "cap"/pixelation discussion (see below)? Related to my comments on problematic definitions of PIP ellipse and rectangle (see above)? In any case, it is unclear to me why you talk about precipitaion particles here. Do you mean ice/snow?

The performance of the PIP shape fitted routine is covered in Section 5.2. Regarding "precipitation" I would argue that the PIP shape-fitting will not do well with fitting an ellipse to the slightly distorted circle of a rain drop either (see Fig. 8a), but since this paper is explicitly about ice/snow, I've added the word "frozen" before "precipitation" to clarify this. Line 631

L623: "PIP variables"?

I've reworded a couple sentence to clarify this:

"Worth noting, however, is that the PIP shape-fitting issues only impact the four variables listed in the relevant bullet point above; the other PIP reported measures of particle size (e.g., particle bounding box dimensions, area, equivalent diameter) remain reliable. As such, any products that use a non-shape-fitting measure of particle size (e.g., Pettersen et al., 2020) remain unaffected by the shape-fitting issue." Lines 632 – 636

Please explain what you mean with "cap" when you first use this expression!

New L432-434: "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. 5g) and much of the reduction in the mean absolute difference (0.254 versus 0.184) is likely due to this change."

I've added the following sentence to the end of the paragraph where we point out the artificial cap in the PIP-fitted ellipse aspect ratio (although we didn't refer to it as an artificial cap there):

"This effectively represents an artificial cap in PIP-fitted ellipse aspect ratio." Line 427

L432: Cap? Fig 5h shows rectangle which does not have the "cap", refer to Fig 5g) instead!

I think the figure references are correct as they stand. I refer to 5h as the PIP-fitted rectangle aspect ratio, which I point out does not suffer from the cap, and 5g as the PIP-fitted ellipse aspect ratio, which I point out does suffer from the cap.

L 478: refer to Fig 6a): I would refer to Fig 5g) and Fig 6a).

Added reference to Fig 5g:

"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 Figs. 5g and 6a), however." Lines 485 – 486

L434: What is "this change"?

Poor wording on my part. Here's the new sentence:

"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. 5g) and much of the reduction in the mean absolute difference (0.254 versus 0.184) is likely tied to the lack of said artificial cap." Lines 438 – 440

L434-435 "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)"
Cannot be seen from Fig 6b (alone) => Refer to Fig 5h!

Added reference to Fig 5h and Fig 6d:

"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 (Figs. 5h and 5b,d).” Lines 440 – 441

L436-438: "It should be noted, however, that the PIP-fitted rectangle aspect ratio does capture the increase in aspect ratio associated with the periods of lump graupel precipitation on 8 March around 0900 UTC and after 1400 UTC."
The text suggests that you are looking at Fig 6b. It is not obvious at first waht you are referring to exactly: the higher density at higher aspect ratios. But then you should say the same also for aspect ratios from PIP-fitted ellipses and refer to both Fig.s 6a and 6b).

The intent here was to point out that the PIP-fitted rectangles capture the large aspect ratio while the PIP-fitted ellipses do not; now that you point it out, I see the choice of wording I used was suboptimal since both show an increase in the aspect ratio. As such, I’ve reworded the sentence to clarify this and to make it clear that the periods of lump graupel correspond to periods where the 1.0 aspect ratio is dominating. I also swapped the order of this sentence and the previous one to improve the flow:

“In contrast to the PIP-fitted ellipse aspect ratios, the PIP-fitted rectangle aspect ratio does capture the large aspect ratios associated with the periods of lump graupel precipitation on 8 March around 0900 UTC and after 1400 UTC, although these aspect ratios are almost entirely reported as a value of one. 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.” Lines 441 – 445

L447-454 Discussion of Fig 7: example particles.
You should state that you now refer to "MASC-fitted and the tensor-fitted ellipses" as correct or reference.

I’m hesitant to declare one (or both) of these as the “correct” measurement, hence the “more reasonable estimate” wording in the final sentence of this paragraph:

“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.” Lines 463 – 465

As for being used as a reference, I feel like the last sentence of the first paragraph of section 5.2 covers that sufficiently:

“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.“ Lines 417 – 419

Panel c and d:
L449-452 "Particles (a) and (b) are both likely some type of aggregate frozen

precipitation based on their odd shapes.
450 Based on the relatively circular shapes of the remaining two particles, relatively high fall speeds (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."

The particles are named after the panel letter of the unannotated view of the particle, hence “particle (a)” is the particle that appears in panels a and e, “particle (b)” is the particle that appears in panels b and f, etc.:

“For simplicity, these particles will be referred to by the Fig. 7 panel letter of the unannotated panels (i.e. panels a--d).” Lines 455 – 456

L460-463: "For a particle with complicated outlines, such as particle (b), the particle perimeter is far greater than the perimeter of an ellipse or rectangle of either equal area or equal dimensions."
This is wrong: the definition of the fitted ellipse is that it has the same perimeter and area, so you cannot say that the perimeter of the particle is far greater than that of the ellipse. Further, it is unclear what "equal dimensions" refer to.

Good catch! Not sure what “equal area” is doing in there (presumably there was a reason, but it’s escaping me now; best guess is I was editing this sentence after editing the following paragraph where we keep the area constant). It should just be “of equal dimensions.” On the topic of dimensions, I’ve adjusted the wording to clarify things:

“For particles with complicated outlines, such as particle (b), the particle perimeter is far greater than the perimeter of an ellipse or rectangle of equal dimensions (i.e., length and width).” Lines 467 – 469

L464: "perfect ellipse"??
"The relationship between excess perimeter for a given area, relative to a perfect ellipse..."
It is the excess with respect to a circle (not "perfect ellipse".

The choice of “perfect ellipse” here was intended to clarify that we are not talking about an ellipse that has been pixelated nor something that is very closely approximated by an ellipse but deviates slightly from a “perfect ellipse”. Ellipse is preferable to circle as it’s the more general case and the concept that the sentence is introducing can be readily applied to a non-circular ellipse as well as to the special case of the circle. We use a circle here instead of an ellipse of some arbitrary aspect ratio as there is no analytical solution for the perimeter of an ellipse in terms of major and minor axis length.

L475-476: "this missing information is, in fact, the core issue with both the PIP ellipse and rectangle fits."
This is one issue, i.e. the definition of ellipse and rectangular fits. The other issue is pixelation.

circle should result in perimeter stretching of 1 and aspect ratio of 1, but it does not due to pixelation.

The pixilation issue, while important, is secondary to the lack of shape information in the fitting equations. If the relationship between perimeter stretching factor and aspect ratio (Fig 8a) were linear instead of exponential (i.e., the fitting equations were better), then pixilation would not have as large of an impact on the ellipse fit. I would argue that having the shape information in the fitting process acts to reduce the sensitivity of the fit to pixilation errors. Evidence of this can been seen in the various comparisons between the PIP-fitted shapes (especially the ellipse fit) and the MASC and tensor fitted ellipses (both of which do take shape information into account). Since all the shape fits use the same images, they all encounter the same level of pixilation, but the fits that include shape information produce far more reasonable estimates of the particle shapes. Looking at the opposite situation, even with zero pixilation, slight deviations from a perfect ellipse will still result in relatively large changes to the ellipse for all but the smallest aspect ratio particles (as per Fig 8a).

L489-490: "...very small deviations of the particle 490 edge from a perfect circle as well as by the inability to perfectly represent a circle using square pixels (i.e., pixelation effects)."
Cannot separate these effects. However, pixelation alone explains the "cap" regardless if there are any small natural dieviations (that are not resolved due to the same pixelation.

I agree that you cannot separate the small deviations from the pixilation effects in practice, but they are not synonymous. The pixilation effects create small deviations, but small deviations are also introduce by the actual true shape of the particle being observed. Whether or not these small deviations are captured in the pixelated image would depend on how they line up with the pixels, the size of the deviation relative to the pixel size, and how the deviation interacts with the image compression (which introduces a directional dependency to the pixilation effects).

Take a circle and determine the measurements after emulating PIP images. Then check the perimetr stretching factor. You will likely get a perimeter stretching factor larger than 1 as a consequence of pixelation and determination of the perimeter of the pixelized particle image. The pixelation itself creates a variation of perimeter location from any defined centre point around the radius of the equivalent circle.
Cfr. Fig. 9.

Outside of some very unique circumstances involving perfect alignment of the pixels and the circle, I agree with this statement.

Rectangle fit: perimeter stretching below 1.128: apparently the algorithm does a compromise and accepts a square with too long perimeter and too small area (and aspect ratio of 1).
The rectangle is less susceptible to pixelation: the "cap" therefore does not exist.

The lowest perimeter stretching factor (~1.05, see below) is below the limit where rectangular fits are possible. Therefore, no cap.

Discussion in L499-507 can be simplified, it is all due to the definition of rectangle fit.

Figure 8 indicates that the ellipse and rectangle fit are similarly sensitive to small changes in perimeter stretching factor (i.e., small deviations from a perfect circle, or, more generally, a perfect ellipse), but, critically, this sensitivity occurs at different perimeter stretching factors. The lack of a cap is primarily due to where the distribution of observed particle perimeter stretching factors peaks relative to that sensitivity (see Fig. 9 and related discussion). If the observed perimeter stretching factor distribution were to be maximized at one, then we wouldn't see the artificial cap in the PIP-fit ellipse (it wouldn't make the fits any better, it would just mean that the distribution would look deceptively reasonable and potentially mislead researchers to trust the data).

Fig 9 shows again that rectangle fit is not well defined: for half the particles the fit is not possible and the PIP algorithm assigns an aspect ratio of 1.

This is precisely why there does not appear to be an artificial cap in aspect ratio and why we argue that, while the distribution of PIP-fitted rectangle aspect ratio might look reasonable at first glance, the truth of the matter is the PIP-fitted rectangles do not produce reliable measures of shape either.

There is nothing special about the aspect ratio 0.6. Instead you could look at tha maximum aspect ratio for PIP-fitted ellipse that you found (~0.65) and relate that to the minimum stretching factor found for PIP (~1.05).

You are correct that the aspect ratio of 0.6 is more or less arbitrary (at least among values that are close to the maximum PIP-fitted ellipse aspect ratio). There's no real narrative difference between using 0.6, using the maximum PIP-fitted ellipse aspect ratio (which is approximately 0.65, as you pointed out), or using a value of 0.525, which approximates the peak of the distribution during the lump graupel periods.

Fig 7 improved in New Fig8; check discussion!
Check discussion new Fig 8 and 9 (tracked changes new L 464-507: perimeter stretching)!

I appreciate the new sentence about PIP perimeter calculation (Lines 127 – 129), but it is not clear:
1) boundary points are located at the corners of the pixels that make up the particle perimeter
2) subsampling the boundary points to produce a smoother representation of the perimeter
What does subsampling mean?? How does smoothening work?
This is crucial: the smoothening and perimeter calculation affects the perimeter stretching factor and enlarges it more or less (the pixelation effect).

Here’s the entirety of information I have on the perimeter measurement that the IMAQ software provides (from the 2003 version of the IMAQ manual):

“Perimeter: Length of a boundary of a region. Because the boundary of a binary image is comprised of discrete pixels, IMAQ Vision subsamples the boundary points to approximate a smoother, more accurate perimeter. Boundary points are the corners of the pixels that make up the boundary.”

Personally, I interpret this to mean that the perimeter of a corner pixel will be taken along the diagonal of the pixel rather than around the edge of the pixel. Either way, I’ve clarified this by tweaking the sentence and adding another (as well as pointing to the 2003 manual since the 2000 manual doesn’t have this information):

“According to the 2003 IMAQ Vision software manual (National Instruments, 2003), 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. Our interpretation of this is that a corner pixel (i.e., a particle pixel sharing exactly two adjacent sides with other particle pixels) will contribute only its diagonal length to the perimeter rather than the length of the exterior two sides. How this subsampling behaves for a particle pixel that only shares a single side with other particle pixels is unclear.” Lines 127 – 131

MINOR:

L 299: "...particle moving at ... 10 m/s (in the along-particle direction)"
Is the "along-particle direction" the direction defined by the particle motion? Then that is not needed. In any way "along-particle direction" is confusing.
The along-particle direction is the direction along the long-axis of the particle. Reworded sentence to avoid confusion:

“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^{-1}, 4 m s^{-1}, and 10 m s^{-1} in the direction of the particle's long axis for each instrument.” Lines 300 – 302

Sect 5.1:
movement speed
motion speed
fall speed
I think "movement speed" and "motion speed" sound somewhat odd. I would use "speed" and/or "motiion".
Further, "horizontal movement of particles" => "horizontal movement of particles" sounds better.

I think movement speed is more explicit than motion, although velocity is an even better term. I've gone through and made the terminology more consistent by using velocity. The goal behind "particle movement speed" was to reduce the chances that the reader misinterprets "particle speed" to be only the fall speed (which is only one component of the particle velocity, the other being the horizontal speed)

Sect 5.2
measurements made using
=>
measurements emulated

I was originally going to swap these out, but doing so makes the sentence inaccurate as neither the PIP measurements nor the measurements based on the tensor-fit ellipse are emulated (i.e., the tensor-fit ellipse is not emulating another instrument). Upon further consideration, we are making measurements of the particles, those measurements are emulating the measurements of the MASC and 2DVD. Thus, both "measurements made using" and "measurements emulated" are technically correct for the MASC and 2DVD methods although "measurements emulated" give a more specific description for the MASC and 2DVD methods here.

L467 duplicate "area... and area"

Fixed

L601: "...although the relatively low resolution may impact the precision of those measurements (section 5.1)."
Also here, I would make clear what you mean with "resolution":
"...although the relatively low resolution (large pixel size) may impact the precision of those measurements (section 5.1)."
See also the comments on precision above.

See my response to your first comment. Changed wording to read:

"The 2DVD camera setup provides the most accurate measure of area (and equi-area diameter) due to its resilience to motion blurring, although it should be noted that 2DVD also has a relatively large pixel size (section 5.1)." Lines 608 – 609