Response to the reviewers of AMT-2021-176: Reconstruction of the mass and geometry of snowfall particles from multi angle snowflake camera (MASC) images

We thank the reviewers for their critical assessment of our work, their very positive feedback and useful suggestions. In the following we address their concerns point by point.

italic text	*Comments of the reviewers*
roman text	Our answers or comments
blue text	Verbatim quotes from the manuscript

Reviewer 1: Dr. Davide Ori

Overall evaluation

The study presents a novel methodology to retrieve snowflake mass and geometrical properties from the simultaneous observation of the snow particle by multiple viewing angles. The authors developed a GAN algorithm to address this problem. It is worth noting that this is no trivial challenge since the shape of snowflakes is highly irregular.

Finding a methodology for the automatic estimation of snowflake mass, on a single particle level is still an area of active research since it poses significant challenges. The problem is also very relevant for many applications ranging from microphysical studies of snow to remote sensing simulations. For this reason, I consider the study of great interest for an AMT publication.

The paper is logically structured and easy to follow. The graphics are clear and the text is well written.

The methodology is well described and made easier to apply by other researchers thanks to the open code and data sharing.

The limitations of the proposed methodology are appropriately discussed and the application example helps to illustrate the relevance of the study.

I did not find any major flaw in the paper and I certainly recommend it for publication after minor revisions. I will take this opportunity to list a few minor comments that I hope will help the authors refining their work.

We thank Dr. Ori for the appreciation of our work and for the very interesting points of discussion and suggestions listed below.

General comments

Dr. Ori lists several interesting observations or remarks, that we address *point by point* in this section. Before to proceed, we believe it is worth to better explain or rephrase some aspects of our work that involve riming, as this is a key topic tackled in the comments.

- $\mathbf{R_c}$ **vs LWP**: the estimation of riming degree from MASC images follows the work of [Praz et al., 2017], where supervised classification is used on a input set of geometrical and textural features computed for each image. Initially, a categorical classification of riming into 5 discrete (but qualitative) classes is performed. This is then transformed, in a nonlinear manner, into a continuous value R_c ranging from 0 (unrimed) to 1 (fully developed graupel). It is important to note that this is very different with respect to the LWP parameter used in the models of, for example, [Leinonen and Szyrmer, 2015]. LWP describes the environment in which particles get rimed rather than the textural appearance of the rimed particle, which is what R_c does. If we imagine to be in a fully controlled scenario and we increase LWP, R_c will increase until reaching the value of 1 when the original shape of the particle is completely masked by rime accretion and graupel is generated. From this moment on, increasing LWP will still lead to an increase of mass and size, but R_c cannot increase further (the particle cannot get more rimed than fully rimed, in qualitative terms).

  We believe that relating LWP to R_c could be a very interesting and not trivial topic for future research and thus we rephrased the last sentence of our conclusions as:

  There are a number of future studies that can be conducted with this new tool, including ... as well as linking the qualitative riming degree as seen on images of rimed particles with the actual liquid water content of the rime accretion

  We also clarified that R_c is shown in the plots, in Sec.5:

  The same dataset is color-coded according to the apparent riming degree R_c (0 being unrimed particles and 1 fully developed graupel) and according to the classified hydrometeor type.

- R_c **vs density factor**: according to the, very relevant, literature suggestion of [Mason et al., 2018], the *density factor* described in their paper seems to be linked to LWP. Following the same train of thoughts of the previous point, also in this case we must keep in mind that R_c and its variability cannot be directly compared to a density factor without preliminary research in this direction. Even just intuitively, let's look at the variation of R_c for a few crystals (the image is adapted from [Praz et al., 2017]).

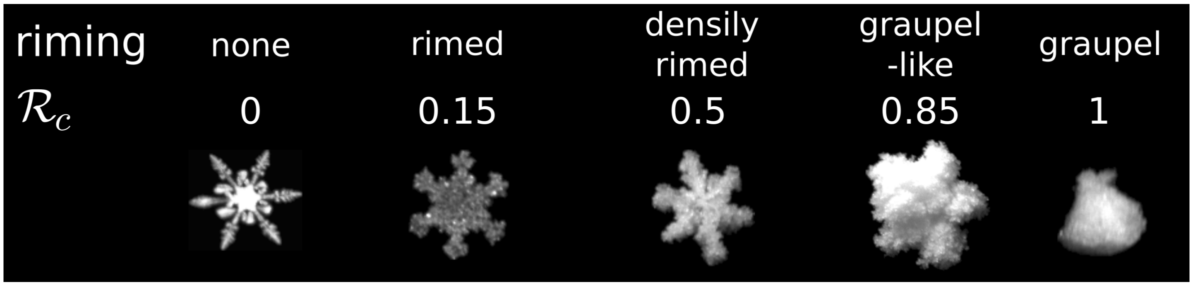

Here we can observe how R_c is really designed to capture the qualitative change from unrimed to graupel rather than the actual quantitative density change that these particles may experience.

- **Riming in the training set vs riming in MASC data**: we would like to recall that R_c is not among the inputs of 3D-GAN. 3D-GAN uses as input binary silhouettes of the particles, while in large part R_c is obtained from textural descriptors [Praz et al., 2017, Appendix].

1. *The thing that puzzles the most is the fact that the GAN is trained with model data. This point is appropriately discussed at line 155 and following, but I am still thinking about how this affects the results presented in table 4. It appears to me that the exponent of the power-law fit for the mass-size relation is always approximately 2 and only at very high degrees of riming (75-100%) it significantly deviates from this value. This is coherent to what was already shown in [Leinonen and Szyrmer, 2015] where the model predicts that riming mostly affects the prefactor of the mass-size relationship, while the exponent is not affected. In that study, a completely different growth model (i.e. rime growth, model C) is needed in order to affect the exponent of the mass-size relation. Because of that, I am suspicious about the fact that this result might be a consequence of the model employed.*

   Dr. Ori raises a very good point here. First and foremost, we realized that we should have better clarify the characteristics of the training set. In fact, model C of [Leinonen and Szyrmer, 2015] was not employed, but only the simultaneous (model A) or subsequent (model B) riming scenarios, at a given LWP. We clarified this important aspect in Sec. 3.3 *Training*:

   In order to create a training set, snowflakes are generated by randomly selecting a few input parameters. To cite the most important, LWP varies from 0.0 to 2.0 $kg\,m^{-2}$, the number of monomers varies between 1 and 50, the monomer type varies among dendrites, needles, rosettes, plates, columns; the riming process is chosen as either occurring at the same time with respect to aggregation (simultaneous) or only once aggregation is completed (subsequent).

The numerical details of this implementation are also available in the file *aggproj.py* in the open-access repository of 3D-GAN. The fact that model C is not used, together with the above-mentioned consideration on the difference between R_c and LWP as riming descriptors, should reduce the concerns about the model-dependence of the preliminary interpretations described in Sec. 5 of the manuscript.

2. *On the other hand, it seems that there is some evidence to support the idea that the exponent of the mass-size power law should vary more continuously between the unrimed aggregates b=2 and the 3D scaling b=3 (see e.g. [Mason et al., 2018] Retrievals of Riming and Snow Density from Vertically Pointing Doppler Radars, figure 1 and equation 8).*

A first mitigation of this effect, with an improvement in terms of continuity has been obtained by producing a better estimate of R_c. In the previous version of the paper, we just transformed the categorical riming degree (1 to 5) to the continuous riming degree R_c. In reality, the riming degree estimation of [Praz et al., 2017] provides the probability that the riming degree belongs to each of the classes and this information should be used. We use now the probability of each class before to compute R_c and thus obtain a_m and b_m values that vary in a smoother way especially near the edges (0 and 1). The revised manuscript contains the correct values; the changes are minimal but they go in the desired direction.

We overall believe that the transition is in reality not too abrupt. Let us try to compute a similar figure as in [Mason et al., 2018], but using R_c. Here, in Fig. 1, we take all the data collected by the MASC during various field campaigns (all together, for any particle type) and we stratify the computation of a_m (converted to *cgs* units as in [Mason et al., 2018]) and b_m couples into R_c bins of 0.05 size, ranging from $[0-0.05]$ to $[0.95-1]$. There is indeed a larger increase in b_m in the last bin (although in our view not too dissimilar to what shown in Fig.1 of [Mason et al., 2018]) getting towards a value of 3.

We must stress once again that R_c and its relations with the actual snowflake density is not yet investigated and at the same time that R_c does not have a Gaussian distribution, especially not near its boundaries. We prefer not to show yet a similar plot in this manuscript or to draw significant conclusions in this direction and we certainly agree that it is a topic of interest for future research. In our view, the goal of Sec. 5 of the manuscript is to provide a teaser, of course scientifically sound, of the possible applications of 3D-GAN. We clarified it at the beginning of the sections as:

We would like to provide the reader with examples and suggestions about possible applications and future research directions that could benefit from the output of 3D-GAN. We consider the retrieval of mass an immediate added value of 3D-GAN and we

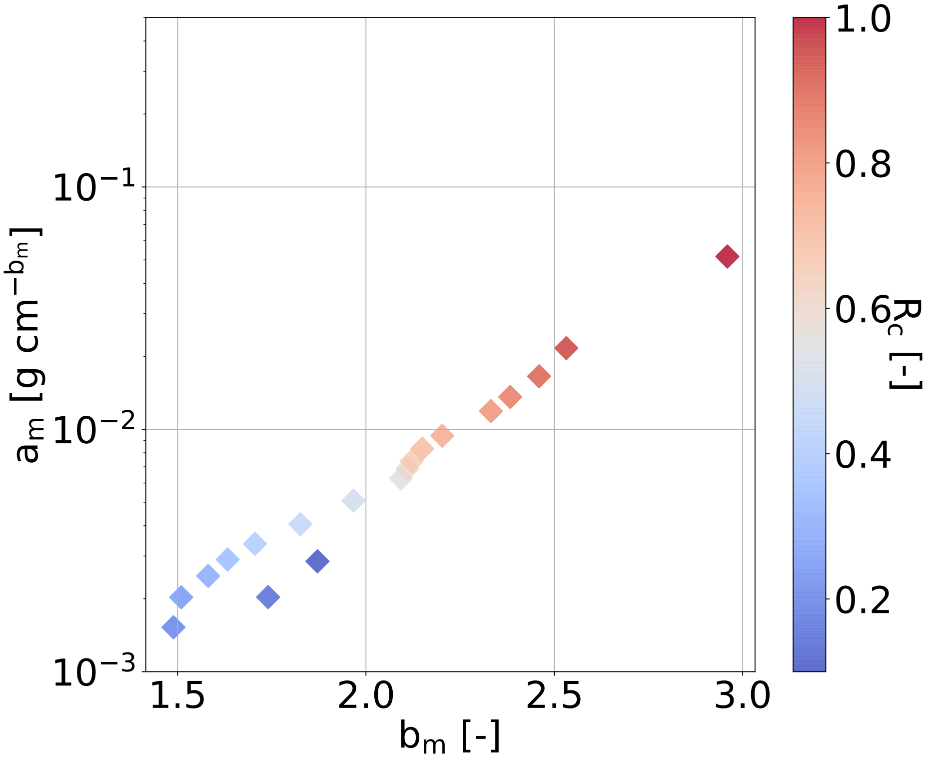

Figure 1: a_m vs b_m values calculated on a large database of MASC data during multiple field campaigns. Power laws are fitted stratifying data according to R_c (color coded) in R_c bins of size 0.05.

apply this retrieval here to datasets collected in the past years at various geographical locations

3. *In my view, when snowflakes grow by riming they increase mostly in mass (fill-in theory and [Seifert et al., 2019]) and will start increasing in size more and more while their rime mass fraction increases. In my view, the smallest particles should reach this limit first, hence their size would start increasing earlier than one of larger aggregates, thus the exponent b of the m-D power low starts increasing. This conceptual view comes from the idea that the riming degree of single particles is not constant inside a snowflake population.*

We completely agree with this view. In fact R_c estimates can vary from particle to particle. As personal view, I like to formulate it in a slightly different way: when the size starts to increase by riming (after filling is completed), the particles are better represented by different power laws of increasing b. We must underline that the minimum D_{max}, for snowflakes that allow for 3D-GAN reconstruction, is about $0.5\,mm$, so we miss lower sizes (and this may affect the power law fits). This was not clarified in the manuscript, so we added the following sentence (Sec. 3.3):

We then extracted these features from the dataset of [Praz et al., 2017], collected in Davos, Switzerland during 2016–2017. We excluded the particles classified as *small particles* by [Praz et al., 2017], as their size does not allow for any shape recognition or significant variability in the descriptors, and computed principal component analysis (PCA) of the feature distribution on the rest of the population. This excludes particles with maximum dimension roughly lower than half millimeter. We kept the three most important PCA components.

4. *If I understand correctly (please correct me if I am wrong), snowflakes coming from various particle populations (different times/ weather events) are stratified according to their riming degree in Table 4. If the logic of my previous paragraph is valid, the m-D fits derived in Table 4 are not well representing natural snowflake populations because, in general, one should assume a size-dependent riming degree.*

   Dr. Ori is correct in the premise: the values in Table 4 are calculated for a given field campaign, which may have lasted several months and experienced large meteorological variability and multiple snowfall events. However, as personal view, $m(D)$ relationships can be useful in two conceptually different ways:

   (a) They may aim, as Dr. Ori suggests, to describe actual populations of particles at the scale of individual snowfall events (or lower scales), as in [von Lerber et al., 2017]. This approach can be useful to better interpret radar data or other remote sensing data.

   (b) They may aim to provide general laws, independently from the actual population or actual snowfall event. In this sense, they can be stratified according to other parameters (hydrometeor type and/or riming degree in our case, but other combinations may be used) to help to get the best possible guess. This approach may be useful in numerical weather simulations, for example.

5. *My additional questions related to this point are: How do the results of Table 4 compare with those of [Mason et al., 2018]?*

   First of all, we updated the values in Table 4 and Table 5 after implementing two improvements in the pre-processing:

   - We implemented a more continuous estimation of R_c, as mentioned above in the response to Point 2.
   - As the MASC may be measuring blowing snow, alone or mixed with precipitation, we heavily pre-filtered the data using the method of [Schaer et al., 2020]. We added a sentence in Sec. 5:

We focus here exclusively on snowfall data and blowing snow images have been removed using the classification scheme of [Schaer et al., 2020].

The results change marginally in numerical terms but reduce significantly the size of the datasets for campaigns with relevant blowing snow. For illustration purposes, we show now in Fig. 5 a field campaign where the MASC was protected by a fence (and thus blowing snow was marginal).

We extended the discussion of how our values compare to the ones in [Mason et al., 2018], and other literature works in the middle of Sec. 5:

Considering the entire datasets of individual field campaigns, values of b_m between 1.80 and 2.04 are obtained, in agreement both with studies based on multi-sensor field measurements [von Lerber et al., 2017, e.g.] and on simulations [Leinonen and Szyrmer, 2015, Karrer et al., 2020]. Especially the work of [von Lerber et al., 2017] provides b_m values also lower than 1.7 and as low as 1.5, as occasionally estimated also by us. Other studies report b_m always larger than approximately 1.7 [Mason et al., 2018], 1.9 [Karrer et al., 2020] or 2 [Leinonen and Szyrmer, 2015].

The estimated prefactors a_m reproduce well the range of values that are documented in the literature. In *cgs* units, the values listed in Table 4 and 5 span roughly between 0.001 and 0.04 $g\,cm^{-b_m}$. This range of variation is similar to [von Lerber et al., 2017]. Also [Mason et al., 2018] reports values in this range, but occasionally higher: up to 0.08 for lump graupel, and larger than 0.1 only for hail or solid ice spheres.
[Leinonen and Szyrmer, 2015] obtains a maximum a_m value of approximately 0.09 $g\,cm^{-b_m}$, but only for a model aiming to reproduce the growth by riming of frozen droplets rather than ice crystals (called rime growth).

6. *How confident are you in the quality of the used riming model to represent real physics? Is it possible that the results are affected by a biased model?*

We believe that Dr. Ori refers here to the riming model used to produce the simulated snowflakes to train 3D-GAN and generate the 3D-printed flakes. Our confidence in this simulation tools comes mainly from the many studies that employed this tool, finding scientifically-sound results. We underline this point in Sec. 3.3:

it has been found to produce realistic mass–dimensional relations of both unrimed [Leinonen and Moisseev, 2015] and rimed [Leinonen and Szyrmer, 2015] snowflakes, and has been used successfully for modeling snowflake microphysics [Seifert et al., 2019] and remote sensing signals from snowflakes [Leinonen et al., 2018, Tridon et al., 2019, e.g.].

7. Can you elaborate on which improvements to the snowflake model might be needed in order to make it producing snowflakes whose mass scale with exponents in between 2 and 3?

The aggregate model (Model A or B in [Leinonen and Szyrmer, 2015]) produces exponents close to 2 while the rime-growth mode produces exponents close to 3. A possibly better combination of these models may prove useful to obtain an improved transition.

Specific comments

1. *Line 141: I guess also the orientation of the particle matters for the simulated silhouette. In other parts of the paper, the orientation of the particle is discussed as a source of uncertainty. I wonder if it would be possible to constrain orientation by means of hydrodynamic models and exploit it to constrain also the GANs retrieval*

   The orientation of the particles is mainly a source of uncertainty in the voxel-by-voxel evaluation with 3D-replicas rather than in this part of the manuscript devoted to the training of 3D-GAN. However, after training, the reconstruction may be potentially affected: from some angles the particle may be easier to reconstruct than from others. We added a statememnt in Sec. 4.1:

   Because the reconstruction is based on the silhouette of MASC images it follows that, for particles of irregular shape and size, the reconstructed output will vary to a certain extent with the orientation of the falling replicas. This is illustrated in Fig. 2 where one can observe how the reconstruction output varies over several consecutive experimental runs. At the same time, we expect also the reconstruction performance to vary: from some angles the particle may be easier to reconstruct than from others.

   We tried to better describe also the uncertainty in the voxel-by-voxel evaluation in Sec. 4.3:

   The orientations of the reconstructed snowflakes depend on the orientation of the printed replicas themselves, as they were falling in the MASC measurement area. The orientation of the reference model is instead fixed.

2. *Line 185. I guess that also the surface properties of ice (roughness mainly) are needed in order to simulate the interaction of light with the snowflake and that is also something that the used model does not provide.*

   We agree with Dr. Ori, and rephrased as:

   however, the radiative transfer of light inside snowflakes is highly complicated and, to our knowledge, no simulation tools exist that could be used to accurately model it and thereby generate proper simulated 2D images from our 3D models, which additionally does not provide surface properties of ice as roughness.

3. *Line 191. I guess there is a typo PhotoTOnic*

   Indeed. Corrected.

4. *Section 4.1 The used apparatus seems very expensive and has significant limitations nonetheless (max size, fragile material). Maybe the authors can give some indications on which are the technical specifications for a 3D printer suited to replicate the experiment.*

   We are not expert of the technical aspects of 3D-printing. However, the main condition is in our view to have a printer able to achieve a resolution of about 40 microns, i.e.

the one we used in the model. The printer used in our case, available in our institute, had a much finer resolution that was not strictly needed. We added a sentence in the conclusions that reads as:

When it will be feasible to 3d-print, at lower costs, a large number of snowflakes at a fine resolution (at least the 40 μm voxels used by the model presented here), it will be of interest to extend the validation to a larger and more variate sample.

5. *Line 202 Is the fall speed impacting the measurement capabilities of MASC? I guess the material used for 3D printing has a different density with respect to ice and thus a different fall speed.*

   Indeed the material used for 3D printing is denser than ice. However this does not impact our work as long as in-focus pictures are captured (i.e. the replicas are falling in the appropriate measurement area). Fall speed, though recorded, is not used by 3D-GAN and also is not used by the retrievals of [Praz et al., 2017] that we showed in this manuscript. Nevertheless, we paid attention and we used a mechanical support to always drop the snowflakes form a constant height of about 12 cm. As order of magnitude, the height required to reach a final fall speed for this replicas would be on the order of one meter (personal communication of colleagues using this dataset for aerodynamic simulations and modeling).

6. *Line 255. In Equation 2, if I got it right, m^i is the same quantity as m_i in Eq.1 just with the voxel identifier i shifted from subscript to superscript. If that is the case I would suggest using the same notation in both cases (I have a personal preference for subscripts)*

   Indeed. Corrected according to the suggestion.

Reviewer 2

Overall evaluation

The article presents a proposed method for hydrometeor mass retrieval and geometric 3D modelling by the application of a GAN trained for these purposes. The GAN is trained on simulated data and the reasoning for this is well explained, however this might warrant further work with captured data. 14 printed flakes for 198 images seems like the minimum (although I may be wrong).

The Reviewer raises a good point here. The GAN is **trained** on a large dataset of simulated data but **validated** on a relatively small physical sample of snowflakes of know geometry (the 3D-printed dataset). Unfortunately this is dictated at the present time by the technical difficulties and costs to generate and handle such small and fragile items. We tried to better stress this aspect in Sec. 4.1 with this sentence:

a total of 198 MASC triplets (and, accordingly, 198 GAN reconstructions) were obtained. Although a larger population of printed snowflakes would be desirable, we believe that, given the above mentioned limitations and technical difficulties, this training sample is a good starting point, including various snowflake habits as well as different riming degrees.

and in the Conclusions:

When it will be feasible to 3d-print, at lower costs, a large number of snowflakes at a fine resolution (at least the 40 μm voxels used by the model presented here), it will be of interest to extend the validation to a larger and more variate sample.

The paper is well written and does an admirable job of explaining a difficult topic. The authors make an effort to outline the limitations of their work and discuss their means to address them. The figures and tables are clean and support their work.

I have no concern publishing this document, although I do have some general questions for the authors.

We thank the Reviewer for the appreciation of our work and for the useful comments and suggestions listed here.

General comments

1. *One thought I have is if the snowflakes used (3D printed) for evaluation were part of the training? Or were they generated specifically for evaluation?*

   The snowflakes used in the evaluation (so, 3D printed) did not belong to the training set. As for the training set, they have been generated using the same simulation method but with different, randomly initialized, parameters and riming degree levels. For the training set we additionally discarded generated snowflakes with major dimension larger than 5 mm due to the technical limitations mentioned in the manuscript and we could not generate snowflake completely unrimed as they were not resistant enough to be manipulated.

2. *Is there any thought on how much error is introduced by using faux snowflakes in validation testing on a network trained on simulated snowflakes? I expect it to be minimal but wonder what your intuition is. This is a difficult problem to solve, and I commend your approach.*

   Regarding the possible error related to the usage of faux snowflakes in validation for a model trained on simulated particles, it is hard to state anything with a certain confidence at this stage as our validation approach is in a sense a pioneering effort. In our view the main sources of errors related to the usage of our replicas are, apart from the sample size issue, the ones mentioned in Sec. 4.1. It reads now, after some rephrasing and modifications:

   A few noteworthy limitations set the boundaries of what we could achieve with this approach:

   1 The maximum dimension of the printed snowflakes is in the range of $3 - 5$ mm. Smaller snowflakes could not be practically manipulated and larger ones could not be printed.

   2 We could not successfully generate completely unrimed particles ($LWP = 0$ kg m^{-2}) as they resulted in structures too fragile to be manipulated without breaking.

   3 Lightly rimed particles sometimes suffered damage while being handled in the MASC measurement area and could thus be used only for a limited number of times.

   also at the end of the same section, we discuss the effect of orientation and optical properties of printed snowflakes:

   we expect also the reconstruction performance to vary: from some angles the particle may be easier to reconstruct than from others and thus we performed multiple experiments with the same particles. An additional source of uncertainty may come from the fact that printed snowflakes are not made of ice: their color and optical properties may be different with respect to actual snowflakes. We assume this aspect to be of negligible importance in our case because only silhouettes are used as input.

3. *Did the printing allow for the introduction of air pockets? How solid were these printed flakes? Did any flakes have cavities?*

Micro air intrusions within individual crystals have been documented, for example in [Nelson and Swanson, 2019], but the snowflake generation model is not equipped to reproduce them. Some air pockets may be generated in aggregates if multiple crystals are combined in a way to leave empty spaces inside, although it is in principle rare to obtain a completely closed air pocket. The printing would allow for such type of cavities to be created.

From visual inspection of the *stl* files[1] and of the actual printed snowflake, we could not see any completely closed cavity. The structure of the snowflakes can be observed in the images of Table 1 of the manuscript. Although this table provides information about the modeled snowflakes (as if they were composed of ice) and not the printed replica themselves, we added in the revised version a column with information about their density (only mass was mentioned in the submitted version) and thus their solidity.

4. *How did the GAN perform with irregular shapes?*

   In principle all the snowflake shapes are irregular. However, in a first approximation, heavily rimed particles tend to have less complex/irregular shapes. In fact the relative errors of 3D-GAN are lower for regular shapes with respect to highly irregular ones. We rephrased a part of the discussion of Sec. 4.2.1:

   In our evaluation data set, the snowflakes having the largest mass are also the ones with the highest degree of riming (See Table 1). In this sense, 3D-GAN shows its ability to indirectly infer the riming degree and the related increase of mass by exploiting the information embedded in the silhouettes. At the same time, heavily rimed particles have more regular shapes and thus represent a less complex geometrical challenge for 3D-GAN. With this in mind it is also not surprising that BL06, which includes more information on particle geometry and compactness, outperforms a simple mass-size relation as M07.

   To better illustrate this aspect to the Reviewer, we prepared as example Fig. 2, here below. The image shows how the relative errors in the mass estimation of 3D-GAN are decreasing as the (color-coded) riming level increases.

5. *You cite [Kleinkort et al., 2017] with a "volume reconstruction (using a standard 3-camera MASC) is quantified to be 27% in terms of absolute error...". Kleinkort found improvement by introducing additional camera angles. Have you given thought to including additional angles for the GAN?*

   The method is actually modular enough to be adapted to include additional camera angles (see for example the functions in the file *aggproj.py* in the open-access code repository of 3D-GAN). The work of [Kleinkort et al., 2017] is however based on a modified MASC version, unique of its kind. We do not have such modifications on our

[1] Freely available for download at: [Grazioli et al., 2021, 10.5281/zenodo.4790962]

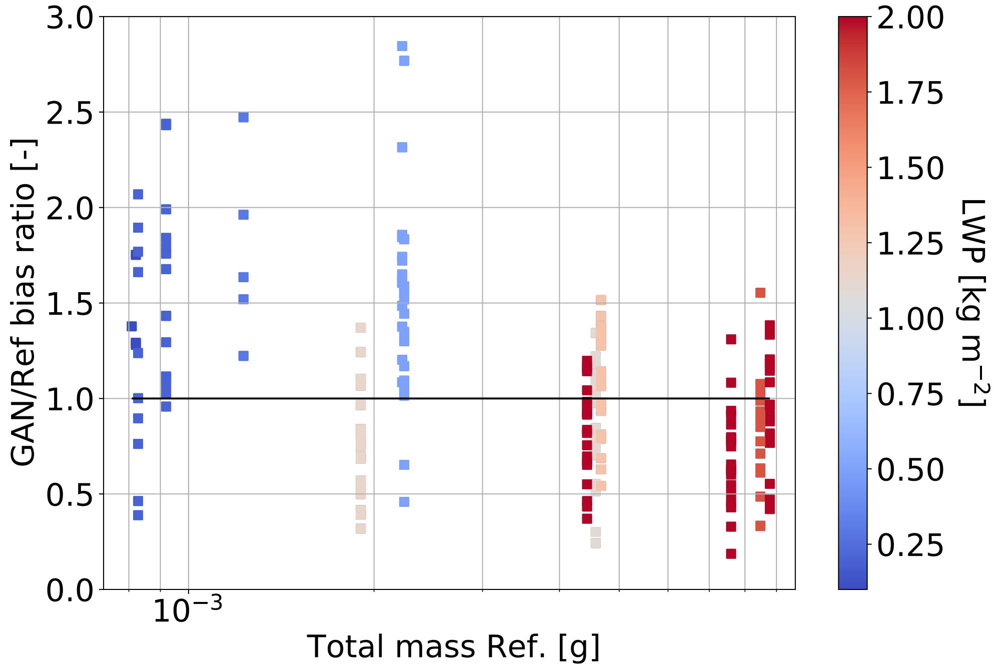

Figure 2: Relative estimation bias (3D-GAN / reference) for the mass of the snowflakes. Color-coded the LWP as proxy for the riming degree of the particles and thus for their complexity (more rimed being less complex). Each point represents a capture of a 3D-printed snowflake.

MASC system and therefore we cannot perform an evaluation as the one shown in the present manuscript.

Specific comments

1. *170- What were the 3 PCA components kept? Or what were they related to. Might help in reproducibility if we had that information.*

   The three most important components are kept (in terms of explained variance). We rephrased the sentence as:

   We kept the three most important PCA components, sorted in order of explained variance.

   The technical details of the PCA and the functions used can be found in the file *features.py* in the open-access code repository of 3D-GAN.

2. *192- Were there any experiments with different printer material? A larger validation set would be beneficial and perhaps a more durable material could assist in that.*

   We did not try different material and we relied on the support and availability of the facilities at EPFL with such equipment. We agree that a larger validation set would be beneficial. We therefore stress this aspect in the Conclusions with this rephrased sentence:

When it will be feasible to 3d-print, at lower costs, a large number of snowflakes at a fine resolution (at least the 40 μm voxels used by the model presented here), it will be of interest to extend the validation to a larger and more variate sample.

3. *289- Mean terms of mean NSE. Want to make sure this is not a typo.*

   It was a typo. Thanks for the correction. It reads now: in terms of mean NSE

4. *319 – Table 5 and 4... this causes me fits. I don't know if there is a rule of numbers being listed in order, but it certainly stands out. Repeated in Fig. 6 so at least the authors are consistent.*

   We agree with the Reviewer and we corrected, through the manuscript.

References

[Grazioli et al., 2021] Grazioli, J., Leinonen, J., and Berne, A. (2021). Support data and codes for the evaluation experiment section of the paper: "mass and geometry reconstruction of snowfall particles from multi angle snowflake camera (masc) images".

[Karrer et al., 2020] Karrer, M., Seifert, A., Siewert, C., Ori, D., von Lerber, A., and Kneifel, S. (2020). Ice Particle Properties Inferred From Aggregation Modelling. *J. Adv. Mod. Earth Sys.*, 12(8).

[Kleinkort et al., 2017] Kleinkort, C., Huang, G. J., Bringi, V. N., and Notaros, B. M. (2017). Visual Hull Method for Realistic 3D Particle Shape Reconstruction Based on High-Resolution Photographs of Snowflakes in Free Fall from Multiple Views. *J. Atmos. Oceanic Technol.*, 34(3):679–702.

[Leinonen et al., 2018] Leinonen, J., Lebsock, M. D., Tanelli, S., Sy, O. O., Dolan, B., Chase, R. J., Finlon, J. A., von Lerber, A., and Moisseev, D. (2018). Retrieval of snowflake microphysical properties from multifrequency radar observations. *Atmos. Meas. Tech.*, 11(10):5471–5488.

[Leinonen and Moisseev, 2015] Leinonen, J. and Moisseev, D. (2015). What do triple-frequency radar signatures reveal about aggregate snowflakes? *J. Geophys. Res. Atmos.*, 120(1):2014JD022072.

[Leinonen and Szyrmer, 2015] Leinonen, J. and Szyrmer, W. (2015). Radar signatures of snowflake riming: A modeling study. *Earth and Space Science*, 2(8):2015EA000102.

[Mason et al., 2018] Mason, S. L., Chiu, C. J., Hogan, R. J., Moisseev, D., and Kneifel, S. (2018). Retrievals of Riming and Snow Density From Vertically Pointing Doppler Radars. *J. Geophys. Res. Atmos.*, 123(24):13807–13834.

[Nelson and Swanson, 2019] Nelson, J. and Swanson, B. D. (2019). Lateral facet growth of ice and snow – Part 1: Observations and applications to secondary habits. *Atmos. Chem. Phys.*, 19(24):15285–15320.

[Praz et al., 2017] Praz, C., Roulet, Y.-A., and Berne, A. (2017). Solid hydrometeor classification and riming degree estimation from pictures collected with a multi-angle snowflake camera. *Atmos. Meas. Tech.*, 10(4):1335–1357.

[Schaer et al., 2020] Schaer, M., Praz, C., and Berne, A. (2020). Identification of blowing snow particles in images from a Multi-Angle Snowflake Camera. *The Cryosphere*, 14(1):367–384.

[Seifert et al., 2019] Seifert, A., Leinonen, J., Siewert, C., and Kneifel, S. (2019). The geometry of rimed aggregate snowflakes: A modeling study. *J. Adv. Mod. Earth Sys.*, 11(3):712–731.

[Tridon et al., 2019] Tridon, F., Battaglia, A., Chase, R. J., Turk, F. J., Leinonen, J., Kneifel, S., Mroz, K., Finlon, J., Bansemer, A., Tanelli, S., Heymsfield, A. J., and Nesbitt, S. W. (2019). The microphysics of stratiform precipitation during olympex: Compatibility between triple-frequency radar and airborne in situ observations. *J. Geophys. Res. Atmos.*, 124(15):8764–8792.

[von Lerber et al., 2017] von Lerber, A., Moisseev, D., Bliven, L. F., Petersen, W., Harri, A.-M., and Chandrasekar, V. (2017). Microphysical Properties of Snow and Their Link to Z(e)-S Relations during BAECC 2014. *J. Appl. Meteor. Clim.*, 56(6):1561–1582.