

Reconstruction of the mass and geometry of snowfall particles from multi angle snowflake camera (MASC) images

Jussi Leinonen^{*,1,2}, Jacopo Grazioli^{*,1}, and Alexis Berne¹

¹Environmental Remote Sensing Laboratory, École polytechnique fédérale de Lausanne, Lausanne, Switzerland

²Federal Office of Meteorology and Climatology MeteoSwiss, Locarno-Monti, Switzerland

^{*}Equal contribution to this work

Correspondence: Jacopo Grazioli (jacopo.grazioli@epfl.ch)

Abstract. This paper presents a method named 3D-GAN, based on a generative adversarial network (GAN), to retrieve the total mass, 3D structure and the internal mass distribution of snowflakes. The method uses as input a triplet of binary silhouettes of particles, corresponding to the triplet of stereoscopic images of snowflakes in free fall captured by a Multi-Angle Snowflake Camera (MASC). 3D-GAN is trained on simulated snowflakes of known characteristics whose silhouettes are statistically similar to real MASC observations and it is evaluated by means of snowflake replicas printed in 3D at 1 : 1 scale. The estimation of mass obtained by 3D-GAN has a normalized RMSE (NRMSE) of 40%, a mean normalized bias (MNB) of 8% and largely outperforms standard relationships based on maximum size and compactness. The volume of the convex hull of the particles is retrieved with MNRSE of 35% and MNB of +19%. In order to illustrate the potential of 3D-GAN to study snowfall microphysics and highlight its complementarity with existing retrieval algorithms, some application examples and ideas are provided, using as showcases the large available datasets of MASC images collected worldwide during various field campaigns. The combination of mass estimates (from 3D-GAN) and hydrometeor classification or riming degree estimation (from independent methods) allows for example to obtain mass-to-size power law parameters stratified on hydrometeor type or riming degree. The parameters obtained in this way are consistent with previous findings, with exponents overall around 2 and increasing with the degree of riming.

1 Introduction

Cloud and precipitation microphysics refer to the interactions and processes that are relevant at the scale of individual particles. Microphysics and microstructure, namely the distribution of particle properties like size, shape, number density and mass, define together the state and the evolution of clouds and precipitation at this scale (Pruppacher and Klett, 2000). The parametrization of microphysics in numerical weather models has a major impact on how accurately the links within the water cycle are depicted (Thompson et al., 2008; Liu et al., 2011; Morrison et al., 2020). Research on ice-phase precipitation (snowfall) microphysics and microstructure is complicated by the complex geometry of individual or aggregate crystals (see Magono and Lee, 1966; Ford, 2014, for visual examples) and by the multitude of processes influencing the density, size and fall speed of the hydrometeors: riming, aggregation, melting, vapor deposition, sublimation, secondary ice production, only to cite a few. Not only is the complex three dimensional structure of snowflakes poorly understood, but there is also an active

25 debate on the range of validity and applicability of mass-size relations (Dunnavan et al., 2019; Karrer et al., 2020) as well as on the appropriate shape approximation of snowflakes as oblate or prolate spheroids or ellipsoids (Jiang et al., 2017, 2019; Dunnavan et al., 2019). It seems established that the model of snowflakes as spheroids of constant density is oversimplified and outdated (Dunnavan et al., 2019) as this assumption affects significantly the interpretation or retrievals of active remote-sensing measurements of snowfall (Leinonen et al., 2013; Leinonen and Szyrmer, 2015).

30 Among the challenges in the field of snowfall microphysics, a key role is played by the difficulty to conduct undisturbed observations of snowflakes in free fall. Garrett and Yuter (2014) underlined how currently used size-fall speed relations still rely on experiments performed on a limited number of snowflakes (Locatelli and Hobbs, 1974). Particle habits and shapes are even more complicated to study and much of the current knowledge on individual ice crystals is based on controlled laboratory experiments rather than outdoor *real-world* observations (Takahashi, 2014)(Takahashi, 2014; Weitzel et al., 2020). The commercialization of a few ground-based instruments designed to collect actual images of falling hydrometeors and estimate at the same time their fall speed has given a noticeable momentum to this field of research. An example is the two-dimensional video disdrometer (2DVD, Kruger and Krajewski, 2002) providing orthogonal silhouettes and speed of falling particles. More recently, accurate and high-resolution depictions of snowflakes could be obtained with imagers like the Snow Video Imager/Particle Image Probe (Newman et al., 2009) or with the Multi-Angle Snowflake Camera (MASC Garrett et al., 2012). The availability of actual images has promoted the development and rapid improvement of several automatic hydrometeor classification techniques (Grazioli et al., 2014; Gavrilov et al., 2015; Praz et al., 2017; Leinonen and Berne, 2020) adapted to the data of these sensors. While the accuracy of the measurements of fall velocity provided by those instruments is often hampered by wind and turbulence (Nešpor et al., 2000; Garrett and Yuter, 2014; Fitch et al., 2021), the added value in terms of microphysical characterization is significant.

45 Unlike other instruments, MASC captures simultaneously three pictures of falling hydrometeors from three distinct coplanar viewpoints, opening the conceptual possibility to perform a 3D reconstruction of the observed snowflakes. To date, the only documented effort to perform this retrieval has been proposed by the visual hull (VH hereafter) approach of Kleinkort et al. (2017). VH has been shown to produce good retrievals of multi-dimensional shapes from the combination of single-view cameras, especially for a modified version of MASC mounting 5 cameras instead of 3. VH performs an accurate retrieval of volumes and it does not tackle directly the retrieval of the mass of individual snowflakes, which is one of the focuses and motivations of the method proposed in the present paper. The need to obtain simultaneous measurements or estimates of mass and shape of individual particles has been in fact declared *urgent for the scientific community* by Jiang et al. (2019).

In this article we present a method, based on a generative adversarial network (GAN), to retrieve the three-dimensional distribution of mass of individual snowflakes using as input the two-dimensional triplet of images collected by a MASC.

55 GANs are nowadays finding application in the field of environmental and atmospheric sciences (e.g. Leinonen and Berne, 2020; Leinonen et al., 2021), thanks to their versatility and their ability to perform 3D reconstruction of images has been already explored, for example, in the medical field (Yang et al., 2017). The GAN presented here is trained on a set of simulated snowflakes (generated using the technique of Leinonen et al., 2013; Leinonen and Moisseev, 2015; Leinonen and Szyrmer,

2015; Karrer et al., 2020) and evaluated on 3D-printed 1 : 1 scale snowflake replicas repeatedly dropped into the MASC
60 sampling area.

This article is organized as follows. Section 2 describes the MASC, the instrument used in the present study. Section 3 details the methods, data and the novel mass and shape estimation algorithm. Section 4 is devoted to the evaluation of the retrieval, while Sec. 5 provides examples of applications and potential future studies. Section 6 draws the main conclusions of this work.

2 The multi-angle snowflake camera (MASC)

65 The method presented here is built and designed for the data collected by the multi-angle snowflake camera (MASC). We briefly recall here the most important technical characteristics of the instrument and the known limitations, and we refer the interested reader to more detailed literature on the subject at the end of this section.

A MASC is composed of three high resolution co-planar cameras pointed to a common focal point. Each camera is separated by 36° with respect to the next one (rotation around the vertical axis) such that a picture of a given snowflake can be obtained
70 simultaneously at an angle of 0° and $\pm 36^\circ$. Two infrared (IR) emitter-detector pairs are triggering the cameras and three associated spotlights. The IR arrays are separated vertically by 32 mm, providing in this way an estimate of particle fall velocity. For the data shown here, the MASC system is composed of three 2448×2048 pixels cameras and the maximum acquisition rate is about 2 Hz (as in Praz et al., 2017).

The data processing steps employed in this study are the same as in Praz et al. (2017), although only a minor part of the
75 information generated is used as actual input of the method described in the following section, while another part can be used to interpret and complement the output (as illustrated in Sec. 5). The preprocessing involves snowflake identification (and matching) in the three images, calculation of geometrical and textural descriptors, image quality evaluation, hydrometeor classification and riming degree estimation.

Although the MASC is a relatively recent instrument, the interested reader can find a fair amount of relevant literature about
80 it. Its measurement principle is detailed in Garrett et al. (2012). Several works exploited MASC data to investigate geometry and fall speed characteristics of hydrometeors (Garrett and Yuter, 2014; Garrett et al., 2015; Jiang et al., 2019), and others were devoted to hydrometeor classification techniques as Praz et al. (2017); Hicks and Notaros (2019); Leinonen and Berne (2020). Recent work tackled the challenge of automatic discrimination of precipitation, wind-blown snow and their mixtures (Schaer et al., 2020). The limitations of the instrument and noteworthy wind-related data degradation issues are well summarized in
85 Garrett and Yuter (2014); Fitch et al. (2021). Finally, an upgraded version of the MASC equipped with 5 cameras is described in Kleinkort et al. (2017).

3 Methods

3.1 Generative adversarial networks

GANs, belonging to the family of deep learning techniques (Alom et al., 2019), are generative models that are trained as a combination of two neural networks: the discriminator and the generator. The discriminator is trained to distinguish samples that belong to the training dataset from those that do not, while at the same time the generator is trained to produce outputs that the discriminator considers to be real. This results in the two training processes competing against each other, which is referred to as adversarial training. Since the discriminator is a powerful image recognition network, the generator must learn to produce highly realistic outputs in order to successfully “fool” the discriminator. The generator is able to produce diverse outputs because it is fed random noise as an input, and the generator learns to map the distribution of the noise to the distribution of the input data. In a conditional GAN, both the discriminator and the generator additionally receive a condition as input data, and therefore the generator learns the conditional probability distribution of the input data.

In the original GAN formulation of Goodfellow et al. (2014), the discriminator is a binary classifier, but it was found by Arjovsky et al. (2017) that some of the instability problems of GANs are remedied by reformulating the objective using a dual of the Wasserstein distance of probability distributions. Gulrajani et al. (2017) then combined this approach with a constraint on the gradient of the weights with respect to the training objective; this combination is referred to as a Wasserstein GAN with gradient penalty (WGAN-GP). Given its superior stability with respect to the original GAN, a WGAN-GP is employed in the present study.

3.2 3D reconstruction GAN

Our 3D reconstruction GAN, named 3D-GAN hereafter, is formulated as a conditional WGAN-GP, where the desired data is the 3D structure of the snowflake and the condition is a set of three binary images (silhouettes) captured from the angles at which the MASC sees the snowflake. The objective for the generator is thus to generate a 3D structure that the discriminator considers as appropriate for the image triplet.

The generator network is shown in Fig. 1a. The inputs are three snowflake silhouettes of 128×128 pixels. The first part of the processing passes the inputs through a series of residual downsampling blocks followed by a fully connected layer, resulting in a set of descriptors for each image. Following the “Siamese network” approach (Chicco, 2021), this step is implemented using the same weights for each image. The descriptors are then concatenated and processed through several fully connected layers, resulting in a set of descriptors for the image triplet. At this stage, the noise is also included in the model by multiplying the input of the second fully connected layer with the noise vector. These descriptors are then passed through one more fully connected layer to produce 2048 variables, which are then reshaped to $32 \times 32 \times 32$ 3D feature maps of $4 \times 4 \times 4$ pixel size. In the final stage of processing, the 3D feature maps are passed through upsampling blocks, eventually producing a 3D grid of $32 \times 32 \times 32$ grid volume elements (voxels). The size of the produced grid was selected as a compromise between resolution and computational requirements.

The inputs to the discriminator (Fig. 1b) are a 3D grid (either from the training dataset or from the generator) and a triplet of
 120 images. The images are processed to descriptors using a Siamese network in a manner identical to the generator. Meanwhile, the grid is passed through a set of downsampling 3D convolution blocks, the result of which is flattened into descriptors. The descriptors for both the 3D grid and the images are processed through multiple fully connected blocks. The descriptors for the grid and the images are then combined by multiplying them with each other. The result of this is passed through more fully connected blocks, eventually producing a single scalar as the discriminator output.

125 While the silhouettes are binary images, the value of each voxel in the 3D grid is proportional to the average density of the ice-air mixture within that voxel, scaled such that the mean density of the nonzero voxels is approximately 1. It is therefore possible to compute the snowflake mass from the outputs of the GAN. We however found that we can achieve better mass estimation with a separate neural network trained specifically to predict the mass. For this, we used a network architecture similar to that of the discriminator (Fig. 1b) except without the 3D grid input and processing branch. This network gives us the
 130 total mass m ; to estimate the mass m_i in each voxel i in the 3D grid output of the generator, we scale the voxel value as

$$m_i = \begin{cases} y_i \frac{m}{\sum_{j: y_j > 0} y_j}, & y_i > 0 \\ 0, & y_i \leq 0 \end{cases} \quad (1)$$

where y_i is the generator output at voxel i .

3.3 Training

Training the 3D reconstruction GAN requires large training datasets of 3D structures and MASC imagery. As it is extremely
 135 difficult to accurately map the 3D structure of a snowflake, such datasets are currently not available from measurements of real snowflakes. Thus, we train the GAN using synthetic observations from modeled snowflakes created with the snowflake generation model described in Leinonen et al. (2013), Leinonen and Moisseev (2015) and Leinonen and Szyrmer (2015). This model creates volumetric 3D models of snowflakes, and is capable of modeling single crystals, aggregation and riming. The degree of riming is indirectly prescribed by the liquid water path (LWP, in kg m^{-2}) parameter (Leinonen and Szyrmer, 2015).

140 The generated snowflake models are defined by a set of volume elements of $40 \mu\text{m}$ size, each either entirely filled with solid ice of density $\rho_{ice} = 917 \text{ kg m}^{-3}$ or empty. 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).

145 For each snowflake generated with the model, we calculated the silhouettes that would be seen by the MASC from the three different camera angles; the silhouettes were artificially blurred by a randomized amount to simulate conditions where the snowflakes are out of focus.

In order to fully utilize the 3D grid and the projection image in the training process and at the same time operate with data of fixed dimensions, the voxels and the projection pixels can correspond to different physical sizes for different snowflakes.

150 Thus, for example, a snowflake of 5 mm maximum dimension would have a grid element size of approximately $5 \text{ mm}/32 = 0.156 \text{ mm}$ and a silhouette pixel size of approximately $5 \text{ mm}/128 = 0.0391 \text{ mm}$.

The training samples of the GAN are loaded from data files that contain, for each snowflake: the 3D grid, the grid voxel size, 12 simulated projection silhouettes, and the projection pixel size. The 12 silhouettes comprise four image triplets; the images in each triplets are 36° apart corresponding to the MASC camera separation, while the four triplets are spaced 90° from each other. When training the GAN, we increase sample diversity by selecting one of the four triplets at random for each training sample and training step and then rotating the grid correspondingly. We also randomly apply mirroring for further data augmentation.

As mentioned above, we adopt the approach of using a model instead of real observations in the training process out of necessity, while acknowledging that it has a number of potential drawbacks and uncertainties:

- 160 1. The model algorithms may not be representative of the physics of real snowflake formation.
2. Although the model can accept any input parameters, the distribution of the model parameters may not match that of real conditions in nature.
3. The simulation of the image formation is not necessarily accurate.
4. By using the silhouettes instead of the gray-scale images captured by the real MASC, we lose the texture information
165 contained in the real MASC images.

For point 1, regarding the realism of the physics of the model, we note that although the model does not implement a fully physical simulation of snowflake formation, 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 (e.g. Leinonen et al., 2018; Tridon et al., 2019).

To mitigate issue 2, we forced the distribution of parameters closer to that found in nature using the following strategy. First, we identified a selection of morphological image features that Praz et al. (2017) found important for identifying snowflakes, and which did not use texture information and therefore could be calculated also for the silhouettes. 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, sorted in order of explained variance. Then, while generating snowflakes, we applied a realistic range of parameters such as snow crystal type, number of crystals per aggregate, and amount of riming, and calculated the same features and
175 PCA components. After generating a large number of crystals, we then accepted the generated snowflakes to the final dataset
180 only when they made the distribution of the PCA components for the generated snowflakes closer to that of the real snowflakes,

rejecting the others. Thus, we obtained a distribution of snowflake samples that is close to real ones at least in terms of visual descriptors. After this filtering step, the final training set included 20472 samples.

For issue 3, we attempted to simulate the main features of image formation such as randomly blurring the images to simulate situations where they are out of focus. The primary manner in which we attempted to determine if our simulation of the MASC silhouettes is adequate was to use MASC observations of artificial snowflakes 3D-printed from our models, thus using the real MASC rather than a simulation to produce the images. This method is far too labor-intensive and expensive to create large training datasets, but we use it for evaluating the model, as described in Sect. 4.

As for issue 4, we accept the lack of texture identification as a current shortcoming of the model. This is unfortunate because the availability of high-resolution texture is one of the greatest advantages of the MASC; 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 for example. On the other hand, using only the silhouette images may make our approach easier to adapt to silhouette-only instruments such as the 2DVD.

4 Evaluation

4.1 Experiment with snowflake replicas

In order to evaluate the performance of 3D-GAN with real MASC data, we used a set a snowflakes printed in 3D. The snowflake shape models were computer-generated with the technique described in Section 3.3.

The printer used to generate the particles is a Nanoscribe ~~Phototonic~~ Photonic Professional GT+ (PPGT+)¹ and the material used is a polymer (IP-Q) supplied by Nanoscribe (Bagheri and Jin, 2019). Once polymerized, the material is similar to polymethyl methacrylate (PMMA). The resolution used to generate the flakes is the 3D laser spot size of 1.5μm diameter (horizontal plane) and 8μm height (vertical axis).

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

1. The maximum dimension of the printed snowflakes ~~could not exceed 5~~ 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 ($LWP = 0 \text{ kg m}^{-2}$) 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.

14 printed snowflakes were used in the evaluation; an overview of their characteristics is shown in Table 1. We dropped each particle several times through the MASC measurement and after discarding physically damaged particles, a total of 198 MASC

¹ See <https://www.nanoscribe.com/en/products/photonic-professional-gt2>

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. 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 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.

4.2 1D descriptors

A first evaluation of the ability of 3D-GAN to reconstruct realistic snowflakes can be obtained by looking at one-dimensional descriptors. We selected for this purpose the total snowflake mass m , gyration radius r_g , maximum size D_{max} and volume of convex hull V_{CH} . D_{max} and V_{CH} are geometric quantities that define exactly the spatial extent of a snowflake. However, the mass distribution of the GAN output is continuously varying and therefore it is not straightforward to define the exact boundaries of a hydrometeor. The way we tackled this limitation and obtained exact estimates of size and volume is detailed in Appendix A. The evaluation of the descriptors discussed in this section is also summarized in Table 2.

4.2.1 Mass estimation

Mass estimation is a major added value of the proposed method or at least, in response to the current needs of the scientific community (Jiang et al., 2019), a readily usable product. Figure 3 shows that mass is overall well reconstructed. As a reference, 3D-GAN reconstruction is compared with the methods of Matrosov et al., 2007, M07 and Baker and Lawson, 2006, BL06. ML07 and BL06 are retrieval formulas designed for 2-dimensional images and we use them here as a benchmark. A large number of mass-size relations exist in the literature, either fine-tuned for specific types of crystals, aggregates and riming degree (e.g. Leinonen and Moisseev, 2015; Karrer et al., 2020) or obtained by combining the information of several sensors in dedicated field campaigns (e.g. von Lerber et al., 2017). M07 and BL06 are chosen because, like 3D-GAN, they do not require prior knowledge about hydrometeor type and can be readily calculated from the 2-D views of the MASC from silhouette-type images without exploiting textural information. M07 is an adaptive mass-size relation where the exponent and prefactor take different values as the particle dimension (D_{max}) increases and it is a relation in principle valid for unrimed snowflakes. BL06 includes more advanced geometrical considerations and it uses the maximum dimensions in two orthogonal directions, projected area and perimeter.

3D-GAN largely outperforms both of these estimation approaches, as summarized in Table 2. The Normalized Root Mean Square Error (NRMSE) is roughly 40% for 3D-GAN, 70% for BL06 and 103% for M07 while the Mean Normalized Bias (MNB) is close to 10%, -40% and -72% respectively. BL06 is able to provide better estimates than M07 although they are

245 both affected by significant negative biases that become mostly evident for the heaviest snowflakes. 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 ~~fully~~ 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
250 information on particle geometry and compactness, outperforms a simple mass-size relation as M07.

4.2.2 Geometry estimation

We evaluate here two geometrical quantities: D_{max} and V_{CH} (Fig. 3, bottom panels). Both quantities are reconstructed in a satisfactory manner, with NRMSE of 12% and 35% respectively and MNB of 7% and 19%. The estimation of D_{max} is compared with what can be achieved using individual 2D images, selecting the maximum of the three estimates of D_{max} ,
255 one for each camera view, as for example in Praz et al. (2017). D_{max} is slightly better reconstructed using the 2D images directly due to the fact that the 3D-GAN mass distribution output varies smoothly and the exact boundaries can only be approximated with the approach detailed in Appendix A. The retrieval of D_{max} from 2D images is practically unbiased: a result in itself interesting for MASC users. Riming has no major impact on the quality of D_{max} , while it affects the retrieval of V_{CH} : particles with LWP greater than 1 are overall better reconstructed (improvements of 15% in terms of NRMSE while no
260 significant differences in terms of bias, not shown). It is not surprising that heavily rimed particles, are better reconstructed in terms of geometry because their geometry is significantly less complex. In Kleinkort et al. (2017), the performance of the VH reconstruction algorithm for what concerns volume reconstruction (using a standard 3-camera MASC) is quantified to be 27% in terms of absolute error, for a simple spherical test object. The mean absolute error of 3D-GAN for all the printed replicas, thus for significantly more complex shapes, is 30%. If only heavily rimed particles, thus less complex shapes, are considered
265 (LWP > 1 kg m⁻²), the error is further reduced down to 26%. The 3D structure of even heavily rimed particles is certainly more complex than a sphere and thus it is reasonable and conservative to assume that 3D-GAN performs at least as good as VH for what concerns volume reconstruction, with the significant added value to provide at the same time an estimate of mass m .

The effect of the smooth variation of mass of the 3D-GAN output, without sharp edges, is evident when looking at the
270 gyration radius: r_g (Fig. 3), defined in this case as:

$$r_g = \sqrt{\frac{\sum_{i=1}^N (d_{CM}^i)^2 m^i}{m}} \quad \sqrt{\frac{\sum_{i=1}^N (d_{CM}^i)^2 m_i}{m}} \quad (2)$$

where N is the number of voxels, d_{CM}^i is the distance of each voxel with respect to the center of mass of the snowflake and ~~m^i~~ m_i is the voxel mass content. r_g is overestimated by 3D-GAN (overall by 13%), indicating that the mass contents of the reconstructed snowflakes have a larger spread around the respective centers of mass in comparison to the structure of the
275 reference snowflakes.

4.3 3D mass distribution evaluation

With the evaluation setup described above, the 3D distribution of mass is available. In principle this allows one to compare the reconstructed and reference snowflakes with a *voxel by voxel* approach. Although this 1 : 1 comparison is undoubtedly ambitious and not straightforward, it is worth to show here some results in this direction. There are two main preliminary issues to be considered:

1. 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. The grid resolution of reference snowflakes is fixed at 40 μm while the grid resolution of the GAN output varies from flake to flake, as mentioned in Sec. 3.3 and it is generally lower (100 μm or more).

In order to address point 1, a preliminary alignment of each snowflake pair (reconstructed vs modeled reference) is performed. The snowflakes are considered as point clouds and their best alignment is found with the (rigid-body) point cloud alignment technique implemented in the OPEN3D package of Zhou et al. (2018). Issue 2, grid resolution, is tackled by computing *voxel by voxel* performance indicators of mass distribution at various grid resolutions, by first down-scaling the data of both snowflakes into a common grid.

Several performance descriptors can be used to evaluate the reconstruction in terms of overlap or quantitative error. We can define here the following two descriptors. Given a pair of 3-dimensional snowflakes, one being the 3D-GAN reconstruction and one the reference, let the Matched Mass Ratio (MMR) be:

$$\text{MMR}(\Delta V) = \frac{\sum_{i=1}^{N'} (m_{3\text{D-GAN}}^i + m_{\text{REF}}^i)}{m_{3\text{D-GAN}} + m_{\text{REF}}} \quad (3)$$

where ΔV is a given voxel size (resolution of the regular grid), $m_{3\text{D-GAN}}^i$ (m_{REF}^i) is the content of mass of the i^{th} voxel of the GAN reconstructed snowflake (reference true snowflake). $m_{3\text{D-GAN}}$ (m_{REF}) is the total mass, invariant across scales, of 3D-GAN (reference) snowflake and N' defines the set of voxels where the mass content is both nonzero for the GAN and the reference. MMR varies between 0 (worse) and 1 (best) and it evaluates how well the mass of the reconstructed snowflake and the mass of the reference overlap, independently whether the total mass itself is correctly estimated. A MMR close to 1 indicates that *as a whole* the combined mass of the two snowflakes occupies the same voxels. A second, significantly more severe and quantitative, indicator is the normalized sum of errors (NSE):

$$\text{NSE} = 100 \times \frac{\sum_{i=1}^M |m_{3\text{D-GAN}}^i - m_{\text{REF}}^i|}{m_{\text{REF}}} \quad (4)$$

where M is the entire set of voxels where the mass content of 3D-GAN or the reference is nonzero. NSE does not allow for error compensation and it can in principle be as low as 0% only if the estimate of total mass of the GAN is perfect.

Figure 4 illustrates the behavior of MMR and NSE across scales. The distribution of mass is overall well matched (mean MRR above 0.8), while the sum of individual errors accumulates, in terms of mean NSE, from 100% to about 50% over the

range of grid scales. It must be underlined, however, that (i) the best achievable results of NSE are limited to a minimum mean NSE of about 40%, due to the error in the mass estimation itself and (ii) the alignment of the two objects is assumed to be optimal. No real adversary method exists to compute the 3D mass distribution of snowflakes from MASC images, so we decided to show a comparison with an idealized reference as illustrated in the red curves of Fig. 4. This reference correspond
 310 to an ellipsoidal approximation of the reference snowflake with two major competitive advantages with respect to 3D-GAN:

- The ellipsoid is fitted directly on the original 3D model of the reference snowflake and not from the MASC captures. The orientation and overlap is thus optimal and no complications and uncertainties due to the realization of an actual measurement can play a role here.
- The density of the ellipsoid is adapted in order to perfectly match the total mass of the reference snowflake.

315 The *Ellipsoid* method is an idealization of the best possible approximation of the snowflakes by means of an ellipsoid of constant density, thus in principle largely superior to any ellipsoidal approximation that can be obtained using actual MASC measurements. The performance of 3D-GAN is close to the one of this idealized retrieval both in terms of MMR and of NSE across all the scales and, most importantly, it is superior at the small scales: up to about 1.25 mm for MMR and up to 0.75 mm for NSE. Regarding NSE, as the scale of the comparison approaches the dimension of the snowflake, the ellipsoid
 320 approximation obviously exploits the advantage of “knowing” the exact total mass. Given the idealized nature of the benchmark and the complexity of the retrieval itself we consider the performance of 3D-GAN satisfactory.

5 Examples of application

The information provided by 3D-GAN is an important complement to what can be calculated or retrieved from MASC data (for example size, shape, complexity, orientation, hydrometeor type or riming degree, as in Praz et al., 2017). We would like to
 325 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 ~~want to provide in this section a few examples of application on~~ apply this retrieval here to datasets collected in the past years at various geographical locations. We focus here exclusively on snowfall data and blowing snow images have been removed using the classification scheme of Schaer et al. (2020).

330 The availability of both mass and size estimates can be used to construct $m(D_{max})$ relationships using the measurements of a single instrument, the MASC. These relations can then be stratified according to the identified hydrometeor type or as a function of the apparent riming degree, taking advantage of previous work in this direction (Praz et al., 2017). An example is shown in the scatter plots of Fig. 5, for data collected in Switzerland in 2016 and 2017. 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
 335 classified hydrometeor type. Keeping in mind that 3D-GAN does not have access to textural information other than binary particle silhouettes, it is reassuring to observe in these plots several features that make physical sense. For example: for a given

particle maximum size, the riming degree increases the mass content; graupel has the largest mass content (at a given size) while columns the lowest, except for the largest observed sizes that can only be reached by aggregate snowflakes.

Table 5 and 4 and 5 provide the parameters of the $m(D_{max})$ power laws calculated for various field campaigns conducted in the Alps and in Antarctica over several years. While an in-depth microphysical interpretation of these results and their differences linked to season and geographical location is beyond the scope of this study, it is worth to briefly discuss these results and hypothesize how they will be useful to support future research in this direction. Considering the entire datasets of individual field campaigns, values of b_m between 1.91 and 2.19 1.80 and 2.04 are obtained, in agreement both with studies based on multi-sensor field measurements (e.g. von Lerber et al., 2017, and references therein) and on simulations (Leinonen and Szyrmer, 2015; Karrer et al., 2020). Also 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 values reproduce well the range of values that are documented in the literature². 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).

As shown in Fig. 3 and previously discussed, globally-representative $m(D_{max})$ relations can generate very large errors for particles of similar maximum size but different riming degree (i.e. moving along the y dimension of Fig. 5 for a given value on the x axis). For this reason, the possibility to stratify and combine the output of 3D-GAN with hydrometeor classification and riming degree information as shown in this chapter is very relevant for microphysical studies. As detailed in Leinonen and Szyrmer (2015) and von Lerber et al. (2017), both a_m and b_m increase with particle density (riming degree) and b_m approaches values of 3 for fully-developed graupel. The sharp increase in a_m and b_m for fully rimed particles is a good indication of the change of dominant growth mechanism, switching from aggregation to rime accretion. It is also interesting to observe, in Fig. 6, how the $m(D_{max})$ relations stratified on riming degree and hydrometeor type are relatively consistent from one field campaign to the other, showing however a certain level of variability that may leave room to microphysical interpretations. In summary, the availability of an estimate of particle mass from 3D-GAN, in combination with existing retrievals based on MASC data provides new possibilities for future studies. For example:

- Investigate how to relate estimates of riming degree based on the appearance of the particles, like R_c of (Praz et al., 2017), with physically-based riming degree descriptors based on the liquid path encountered like LWP of von Lerber et al. (2017) or Leinonen and Szyrmer (2015).
- Exploit the availability of mass estimates to find and explain the observed relations with size, shape, fall velocity and vertical structure of precipitation.

²Note that a_m is often given in cgs units in published research, while the mkgs convention is used in the tables of the present paper

370 – Exploit the 3D mass distribution for scattering simulations and remote-sensing applications.

6 Summary, Conclusions, Outlook

The MASC instrument is a state-of-the art device to investigate and describe the habits and microphysical properties of solid-phase precipitation particles. Large datasets of triplets of hydrometeor images have been gathered worldwide, with more to be collected during present and future field campaigns. MASC data provided already noteworthy contributions to studies of snowfall microphysics and recent algorithms exist to estimate the hydrometeor type, riming degree as well as volume properties of the particle pictured by the masc. With one exception (Kleinkort et al., 2017), limited effort has been conducted so far to exploit the multi-dimensionality of MASC images to retrieve three-dimensional properties of the hydrometeors. We presented here a method, based on machine learning and trained on synthetic data (with verified realistic properties), to retrieve the three-dimensional distribution of mass of individual snowflakes using a triplet of silhouettes as input, corresponding to the MASC images. Unlike the pioneering work of Kleinkort et al. (2017), mass estimation is provided as a key output and not merely shape and volume.

We have conducted a validation of 3D-GAN by means of 3D-printed replicas of realistic snowflakes of known characteristics. Due to technical limitations and difficulties to handle small and fragile particles, the evaluation is limited to a range of values of sizes and masses that does not fully overlap the one of naturally occurring snowflakes. The mass content is estimated with low bias (10% mean overestimation) and with a normalized RMSE of 40%. Concerning geometrical features, D_{max} is reconstructed with a mean overestimation of 7% (NRMSE 12%), the volume of the convex hull V_{CH} is overestimated by 19% (NRMSE 35%) and the gyration radius r_g by 13% (NRMSE 16%). The evaluation of 3D-GAN reconstructions was also conducted on a voxel-by-voxel basis, after alignment of the reconstructed snowflakes and the original model by means of point cloud alignment (technically called ‘registration’). We have additionally shown that, in order to provide results comparable to 3D-GAN with an ellipsoidal approximation one would need both be able to achieve the best possible 3D fit and to exactly retrieve the mass of the snowflake: two requirements that are extremely unlikely to be fulfilled using MASC data as input.

The 3D-GAN method has still margin for improvement. For example, about the input (black-white silhouettes): future studies may employ image simulation techniques in order to add the missing textural information (including lights and shadows) to the training set and thus to the input. Although we are aware that this is not a straightforward step, it would allow one to fully exploit the MASC data. Our evaluation highlighted a positive bias for r_g , V_{CH} and D_{max} suggesting that 3D-GAN could be improved in terms of particle compactness. 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.

We have shown some examples of application of the novel method, by combining the retrieved mass with dimensional information as well as hydrometeor type and riming degree and then fitting the coefficients of mass-to-size power laws. We obtained a set of values for the prefactor a and exponent b that are in line with previous theoretical or experimental studies, both in absolute terms and in terms of variation according to riming degree: the increase of exponent with riming degree is well

observed in the data collected during various field campaigns. There are a number of future studies that can be conducted with this new tool, including improved scattering simulations, microphysical characterization of the snowfall measured in various locations worldwide as well as linking the qualitative riming degree as seen on ~~rimed particle images~~ images of rimed particles with the actual liquid water content of the rime accretion.

Code and data availability. The code to generate simulated snowflakes is openly available at: <https://github.com/jleinonen/aggregation>. The 3D-GAN code is available at: <https://github.com/jleinonen/masc3dgan>. The codes and data to support the evaluation of the performances of 3D-GAN, including the models and shapefiles of the replica snowflakes are published and available at (10.5281/zenodo.4790962 Grazioli et al., 2021). Raw or processed MASC data for any campaign mentioned in the paper, as well as a MATLAB code to pre-process the data according to the method of Praz et al. (2017) are available upon request to the authors.

Appendix A: Approximation of geometrical features

As discussed in the manuscript, the GAN output consists of a three-dimensional distribution of mass. These values vary smoothly and do not generate a clear cutoff at the edges of the reconstructed snowflakes, artificially expanding their apparent size. For practical purposes, quantities such as D_{max} or the volume of the convex hull may be of interest and thus we propose a simple but conceptually-sound method to define the geometrical extent of each snowflake by means of an adaptive minimum density threshold.

The goal is to obtain, for each individual snowflake, an optimal density threshold ρ_{th}^{opt} [kg m^{-3}] such that only voxels having $\rho \geq \rho_{th}^{opt}$ are used to define the spatial extent of the particle. Let P be the three dimensional matrix defining the density of each voxel of a given snowflake. As the maximum density can vary for each snowflake, it is normalized between 0 and 1. Let P_{th} be the same matrix, censored by zeroing the voxels having a lower density content than an arbitrary threshold ρ_{th} . Two scalar quantities can be defined, given P and ρ_{th} :

1. $\overline{\rho_{th}}$ defines the mean density of non-zero voxels.
2. m_{th}/m_0 the residual total mass of the censored matrix with respect to the uncensored total mass.

The first scalar quantity increases with increasing threshold levels (as voxels of low density are removed), while the second one decreases. Let us then multiply the two scalars and define a simple weight W^* as a trade-off between increase of mean density (and compactness), which we want to reward, and the associated loss of mass, which should be penalized.

$$W^* = \overline{\rho_{th}} \frac{m_{th}}{m_0} \quad (A1)$$

The evolution of W^* as a function of the threshold level has a behavior as illustrated in Fig. A1. We choose an optimal threshold corresponding to the location of maximum W^* in order to balance the two errors. The final threshold is then applied to define the spatial extent assigned to the snowflake. In this work we used this approach to be able to evaluate the GAN output against

the 3D-printed replicas, exclusively for quantities as D_{max} or the convex hull volume. The total mass as well as the gyration radius are calculated on unfiltered data.

435 *Author contributions.* JL and AB formulated the project and developed the methodology used in this study. JG performed the evaluation and the applications showed here. JL wrote the software code needed to implement the method. JG wrote the article with contributions from AB and JL.

Competing interests. Alexis Berne is a member of the editorial board of the journal (associate editor).

Financial support. This research has been supported by the Swiss National Science Foundation (grant no. 200020_175700)

440 *Acknowledgements.* We would like to thank Julien Dorsaz and the Center of MicroNano Technology at EPFL for the useful exchanges and the 3D-printing of snowflakes. We thank Alfonso Ferrone, EPFL-LTE, for the ideas, support and help to conduct the measurements with the replicas. We thank MeteoSwiss and the past and present members of EPFL-LTE for the MASC data of various field campaigns. We thank Boris Aguilar to share his implementation of the BL06 mass-size retrieval and several participants of the project ICEGENESIS (<https://www.ice-genesis.eu/>) for the useful exchanges during the past months.

References

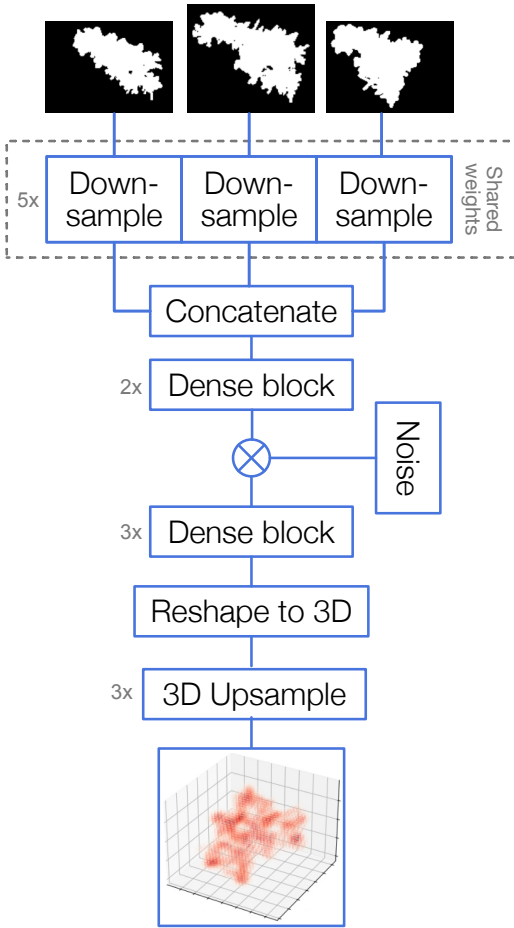
- 445 Alom, M. Z., Taha, T. M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M. S., Hasan, M., Van Essen, B. C., Awwal, A. A. S., and Asari, V. K.: A State-of-the-Art Survey on Deep Learning Theory and Architectures, *Electronics*, 8, <https://doi.org/10.3390/electronics8030292>, 2019.
- Arjovsky, M., Chintala, S., and Bottou, L.: Wasserstein GAN, arXiv preprint arXiv:1701.07875, <https://arxiv.org/abs/1701.07875>, 2017.
- Bagheri, A. and Jin, J.: Photopolymerization in 3D Printing, *ACS Appl. Polym. Mat.*, 1, 593–611, <https://doi.org/10.1021/acsapm.8b00165>, 450 2019.
- Baker, B. and Lawson, R. P.: Improvement in determination of ice water content from two-dimensional particle imagery. Part I: Image-to-mass relationships, *J. Appl. Meteor. Clim.*, 45, 1282–1290, <https://doi.org/10.1175/JAM2398.1>, 2006.
- Chicco, D.: Siamese Neural Networks: An Overview, in: *Artificial Neural Networks*, edited by Cartwright, H., pp. 73–94, Springer, New York, New York, USA, https://doi.org/10.1007/978-1-0716-0826-5_3, 2021.
- 455 Dunnavan, E. L., Jiang, Z., Harrington, J. Y., Verlinde, J., Fitch, K., and Garrett, T. J.: The Shape and Density Evolution of Snow Aggregates, *J. Atmos. Sci.*, 76, 3919–3940, <https://doi.org/10.1175/JAS-D-19-0066.1>, 2019.
- 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>, <https://amt.copernicus.org/articles/14/1127/2021/>, 2021.
- 460 Ford, B.: The Hidden Secrets of Snowflakes, *The Microscope*, 62, 171 – 181, <http://www.mccroneinstitute.org/v/1021/The-Microscope-Volume-62-Fourth-Quarter-2014>, 2014.
- Garrett, T. J. and Yuter, S. E.: Observed influence of riming, temperature, and turbulence on the fallspeed of solid precipitation, *Geophys. Res. Lett.*, 41, 6515–6522, <https://doi.org/10.1002/2014GL061016>, <http://onlinelibrary.wiley.com/doi/10.1002/2014GL061016/abstract>, 2014.
- 465 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.
- Garrett, T. J., Yuter, S. E., Fallgatter, C., Shkurko, K., Rhodes, S. R., and Endries, J. L.: Orientations and aspect ratios of falling snow, *Geophys. Res. Lett.*, 42, 4617–4622, 2015.
- Gavrilov, S., Kubo, M., Tran, V. A., Ngo, D. L., Nguyen, N. G., Nguyen, L. A. T., Lumbanraja, F. R., Phan, D., Satou, K., et al.: Feature 470 Analysis and Classification of Particle Data from Two-Dimensional Video Disdrometer, *Advances in Remote Sensing*, 4, 1, 2015.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y.: Generative Adversarial Nets, in: *Advances in Neural Information Processing Systems 27*, edited by Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q., pp. 2672–2680, Curran Associates, Inc., <https://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>, 2014.
- Grazioli, J., Tuia, D., Monhart, S., Schneebeli, M., Raupach, T., and Berne, A.: Hydrometeor classification from two-dimensional video 475 disdrometer data, *Atmos. Meas. Tech.*, 7, 2869–2882, <https://doi.org/10.5194/amt-7-2869-2014>, <http://www.atmos-meas-tech.net/7/2869/2014/>, 2014.
- Grazioli, J., Leinonen, J., and Berne, A.: 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", <https://doi.org/10.5281/zenodo.4790962>, 2021.

- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C.: Improved Training of Wasserstein GANs, in: *Advances in Neural Information Processing Systems 30*, edited by Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R., pp. 5767–5777, Curran Associates, Inc., <https://papers.nips.cc/paper/7159-improved-training-of-wasserstein-gans.pdf>, 2017.
- Hicks, A. and Notaros, B. M.: Method for Classification of Snowflakes Based on Images by a Multi-Angle Snowflake Camera Using Convolutional Neural Networks, *J. Atmos. Oceanic Technol.*, 36, 2267–2282, <https://doi.org/10.1175/JTECH-D-19-0055.1>, 2019.
- 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. Clim.*, 56, 725–734, 2017.
- Jiang, Z., Verlinde, J., Clothiaux, E. E., Aydin, K., and Schmitt, C.: Shapes and Fall Orientations of Ice Particle Aggregates, *J. Atmos. Sci.*, 76, 1903–1916, <https://doi.org/10.1175/JAS-D-18-0251.1>, 2019.
- Karrer, M., Seifert, A., Siewert, C., Ori, D., von Lerber, A., and Kneifel, S.: Ice Particle Properties Inferred From Aggregation Modelling, *J. Adv. Mod. Earth Sys.*, 12, <https://doi.org/10.1029/2020MS002066>, 2020.
- Kleinkort, C., Huang, G. J., Bringi, V. N., and Notaros, B. M.: 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, 679–702, <https://doi.org/10.1175/JTECH-D-16-0099.1>, 2017.
- Kruger, A. and Krajewski, W. F.: Two-dimensional video disdrometer: a description, *J. Atmos. Oceanic Technol.*, 19, 602–617, 2002.
- Leinonen, J. and Berne, A.: Unsupervised classification of snowflake images using a generative adversarial network and K -medoids classification, *Atmospheric Measurement Techniques*, 13, 2949–2964, <https://doi.org/10.5194/amt-13-2949-2020>, <https://amt.copernicus.org/articles/13/2949/2020/>, 2020.
- Leinonen, J. and Moiseev, D.: What do triple-frequency radar signatures reveal about aggregate snowflakes?, *J. Geophys. Res. Atmos.*, 120, 2014JD022 072, <https://doi.org/10.1002/2014JD022072>, <http://onlinelibrary.wiley.com/doi/10.1002/2014JD022072/abstract>, 2015.
- Leinonen, J. and Szyrmer, W.: Radar signatures of snowflake riming: A modeling study, *Earth and Space Science*, 2, 2015EA000 102, <https://doi.org/10.1002/2015EA000102>, <http://onlinelibrary.wiley.com/doi/10.1002/2015EA000102/abstract>, 2015.
- Leinonen, J., Lebsock, M. D., Tanelli, S., Sy, O. O., Dolan, B., Chase, R. J., Finlon, J. A., von Lerber, A., and Moiseev, D.: Retrieval of snowflake microphysical properties from multifrequency radar observations, *Atmos. Meas. Tech.*, 11, 5471–5488, <https://doi.org/10.5194/amt-11-5471-2018>, 2018.
- Leinonen, J., Nerini, D., and Berne, A.: Stochastic Super-Resolution for Downscaling Time-Evolving Atmospheric Fields With a Generative Adversarial Network, *IEEE T. Geosci. Remote Sens.*, 59, 7211–7223, <https://doi.org/10.1109/tgrs.2020.3032790>, <http://dx.doi.org/10.1109/TGRS.2020.3032790>, 2021.
- Leinonen, J., Moiseev, D., and Nousiainen, T.: Linking snowflake microstructure to multi-frequency radar observations, *J. Geophys. Res. Atmos.*, 118, 3259–3270, <https://doi.org/10.1002/jgrd.50163>, 2013.
- Liu, C., Ikeda, K., Thompson, G., Rasmussen, R., and Dudhia, J.: High-Resolution Simulations of Wintertime Precipitation in the Colorado Headwaters Region: Sensitivity to Physics Parameterizations, *Mon. Weather Rev.*, 139, 3533–3553, <https://doi.org/10.1175/MWR-D-11-00009.1>, 2011.
- 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>, <http://onlinelibrary.wiley.com/doi/10.1029/JC079i015p02185/abstract>, 1974.
- Magono, C. and Lee, C. W.: Meteorological classification of natural snow crystals, *J. Fac. Sci., Hokkaido Univ., Series VII*, 2, 321–335, 1966.

- Mason, S. L., Chiu, C. J., Hogan, R. J., Moisseev, D., and Kneifel, S.: Retrievals of Riming and Snow Density From Vertically Pointing Doppler Radars, *J. Geophys. Res. Atmos.*, 123, 13 807–13 834, <https://doi.org/10.1029/2018JD028603>, 2018.
- Matrosov, S. Y., Clark, K. A., and Kingsmill, D. E.: A polarimetric radar approach to identify rain, melting-layer, and snow regions for applying corrections to vertical profiles of reflectivity, *J. Appl. Meteor. Clim.*, 46, 154–166, 2007.
- 520 Morrison, H., van Lier-Walqui, M., Fridlind, A. M., Grabowski, W. W., Harrington, J. Y., Hoose, C., Korolev, A., Kumjian, M. R., Milbrandt, J. A., Pawlowska, H., Posselt, D. J., Prat, O. P., Reimel, K. J., Shima, S.-I., van Diedenhoven, B., and Xue, L.: Confronting the Challenge of Modeling Cloud and Precipitation Microphysics, *Journal of Advances in Modeling Earth Systems*, 12, e2019MS001 689, <https://doi.org/https://doi.org/10.1029/2019MS001689>, <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001689>, e2019MS001689 2019MS001689, 2020.
- 525 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, 2000.
- Newman, A. J., Kucera, P. A., and Bliven, L. F.: Presenting the Snowflake Video Imager (SVI), *JAOT*, 26, 167–179, <https://doi.org/10.1175/2008JTECHA1148.1>, <http://journals.ametsoc.org/doi/abs/10.1175/2008JTECHA1148.1>, 2009.
- Praz, C., Roulet, Y.-A., and Berne, A.: Solid hydrometeor classification and riming degree estimation from pictures collected with a Multi-Angle Snowflake Camera, *Atmos. Meas. Tech.*, 10, 1335–1357, <https://doi.org/10.5194/amt-10-1335-2017>, 2017.
- 530 Pruppacher, H. R. and Klett, J. D.: *Microphysics of clouds and precipitation*, Kluwer Academic Publishers, Dordrecht, 2nd rev. and enl. ed., with an introduction to cloud chemistry and cloud electricity edn., 2000.
- Schaer, M., Praz, C., and Berne, A.: Identification of blowing snow particles in images from a Multi-Angle Snowflake Camera, *The Cryosphere*, 14, 367–384, <https://doi.org/10.5194/tc-14-367-2020>, <https://tc.copernicus.org/articles/14/367/2020/>, 2020.
- 535 Seifert, A., Leinonen, J., Siewert, C., and Kneifel, S.: The Geometry of Rimed Aggregate Snowflakes: A Modeling Study, *J. Adv. Mod. Earth Sys.*, 11, 712–731, <https://doi.org/10.1029/2018MS001519>, 2019.
- Takahashi, T.: Influence of Liquid Water Content and Temperature on the Form and Growth of Branched Planar Snow Crystals in a Cloud, *J. Atmos. Sci.*, 71, 4127–4142, <https://doi.org/10.1175/JAS-D-14-0043.1>, <http://journals.ametsoc.org/doi/abs/10.1175/JAS-D-14-0043.1>, 2014.
- 540 Thompson, G., Field, P. R., Rasmussen, R. M., and Hall, W. D.: Explicit Forecasts of Winter Precipitation Using an Improved Bulk Microphysics Scheme. Part II: Implementation of a New Snow Parameterization, *Mon. Weather Rev.*, 136, 5095–5115, <https://doi.org/10.1175/2008MWR2387.1>, <http://journals.ametsoc.org/doi/abs/10.1175/2008MWR2387.1>, 2008.
- 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.: The Microphysics of Stratiform Precipitation During OLYMPEX: Compatibility Between Triple-Frequency Radar and
- 545 Airborne In Situ Observations, *J. Geophys. Res. Atmos.*, 124, 8764–8792, <https://doi.org/https://doi.org/10.1029/2018JD029858>, 2019.
- von Lerber, A., Moisseev, D., Bliven, L. F., Petersen, W., Harri, A.-M., and Chandrasekar, V.: Microphysical Properties of Snow and Their Link to Z(e)-S Relations during BAECC 2014, *J. Appl. Meteor. Clim.*, 56, 1561–1582, <https://doi.org/10.1175/JAMC-D-16-0379.1>, 2017.
- Weitzel, M., Mitra, S. K., Szakáll, M., Fugal, J. P., and Borrmann, S.: Application of holography and automated image processing for laboratory experiments on mass and fall speed of small cloud ice crystals, *Atmos. Chem. Phys.*, 20, 14 889–14 901, [https://doi.org/10.5194/acp-](https://doi.org/10.5194/acp-20-14889-2020)
- 550 20-14889-2020, <https://acp.copernicus.org/articles/20/14889/2020/>, 2020.
- Yang, B., Wen, H., Wang, S., Clark, R., Markham, A., and Trigoni, N.: 3D Object Reconstruction from a Single Depth View with Adversarial Learning, in: 2017 IEEE INTERNATIONAL CONFERENCE ON COMPUTER VISION WORKSHOPS (ICCVW 2017), IEEE Interna-

- tional Conference on Computer Vision Workshops, pp. 679–688, IEEE; IEEE Comp Soc, <https://doi.org/10.1109/ICCVW.2017.86>, 16th IEEE International Conference on Computer Vision (ICCV), Venice, ITALY, OCT 22-29, 2017, 2017.
- 555 Zhou, Q.-Y., Park, J., and Koltun, V.: Open3D: A Modern Library for 3D Data Processing, <https://arxiv.org/abs/1801.09847>, 2018.

(a) Generator



(b) Discriminator

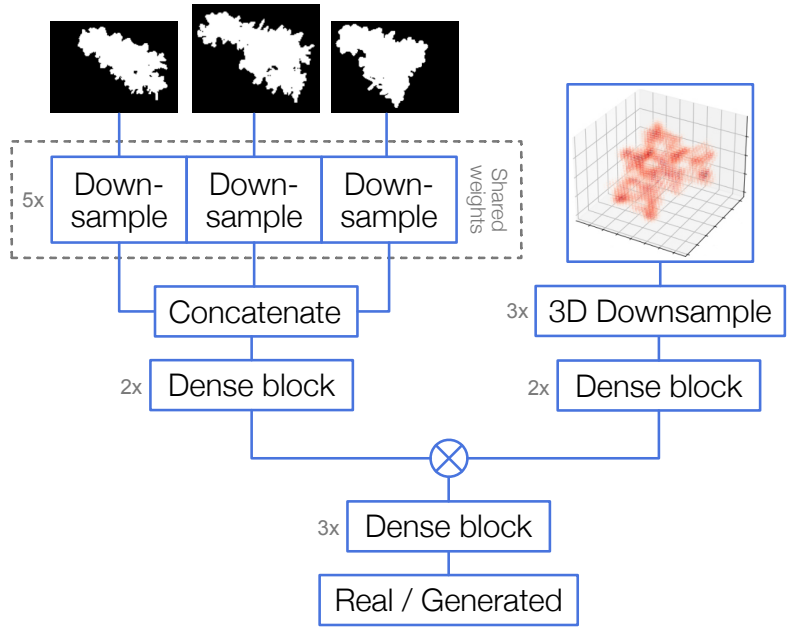


Figure 1. Illustration of the architectures of the (a) generator and (b) discriminator of the GAN.

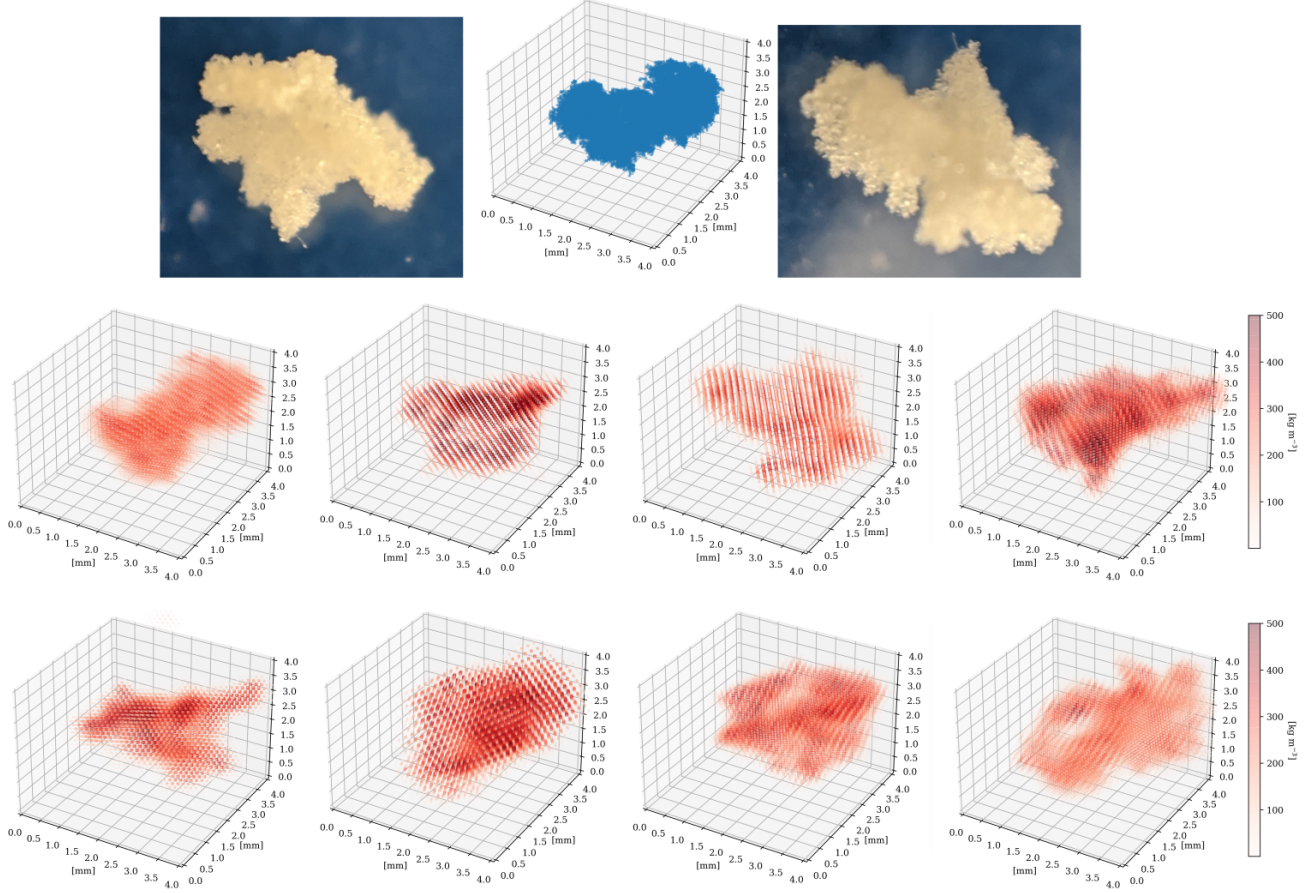


Figure 2. Example of the reconstruction outcome obtained while releasing eight consecutive times the same snowflake replica in the MASC measurement area. Top rows: actual photos of the replica and a pseudo-3D representation (every blue point represents a voxel filled with ice). Bottom rows: reconstructions obtained with 3D-GAN. Every point is color-coded according to the density (ice mass content) of each voxel.

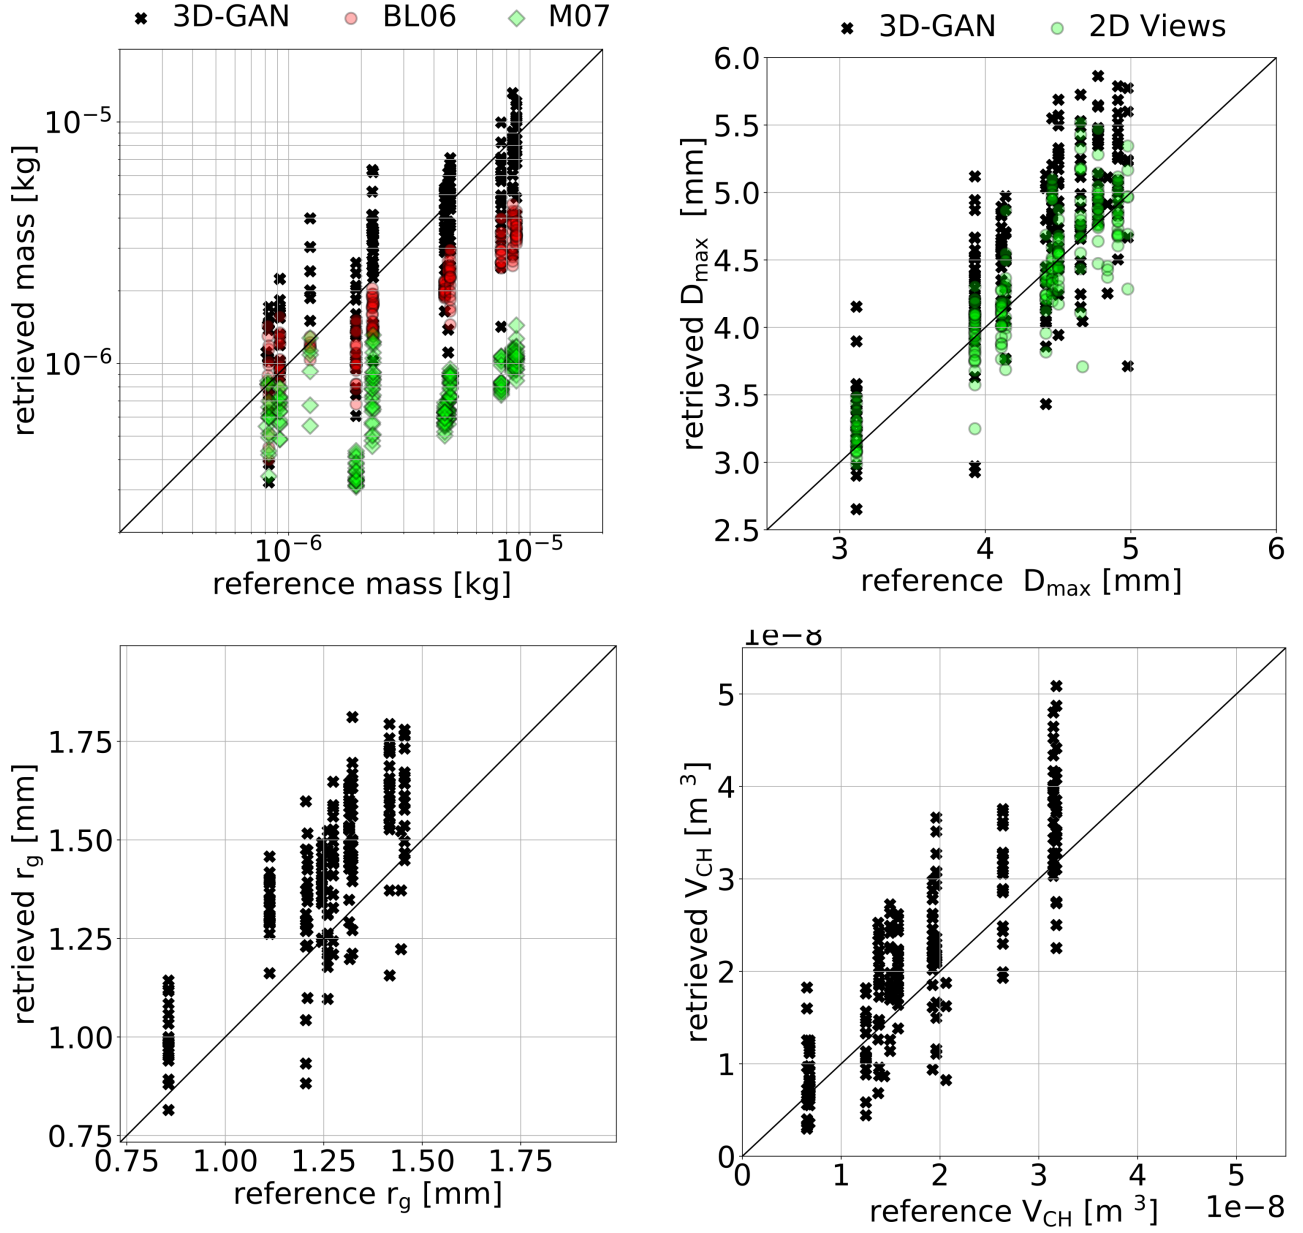


Figure 3. Scatter plot of reference (3D printed replicas) and reconstructed characteristics of the snowflakes used in the evaluation. (top-left): mass, (top-right) D_{\max} , (bottom-left): r_g , (bottom-right): V_{CH} . For the mass we display 3D-GAN reconstructions as well as mass-size reconstructions of Matrosov et al., 2007, ML07 and Baker and Lawson, 2006, BL06. M07 and BL06 are calculated on individual MASC views and the mean value over the three views is shown here in each marker. For D_{\max} the estimation obtained using the 2D MASC images is highlighted in green.




Figure 4. Distribution of matched mass ratio MMR (left) and normalized sum of errors NSE (right) values as a function of the grid volume size at which the evaluation is conducted. Median values over the evaluation sample are indicated by the markers while the vertical lines overlap the 25-75% percentile range. An artificial horizontal displacement is added to the data series to enhance readability. The reference method (*Ellipsoid*) corresponds to an optimal ellipsoidal fit of the reference snowflake with perfect mass match, as described in the text.

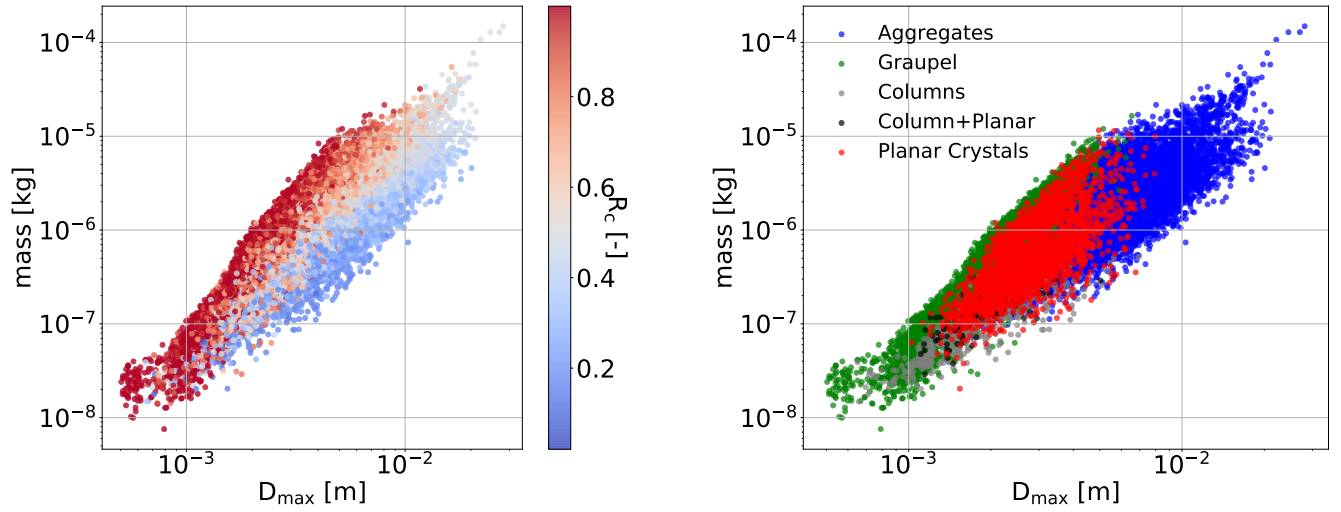


Figure 5. $m(D_{\max})$ scatter plots for the data of a field campaign, named [Valais-2016](#)[Davos 2015](#), taking place in the Swiss Alps in [2016](#)[2015](#) and [2017-2016](#). Left: data color-coded according to riming index R_c of Praz et al. (2017). Right: data color-coded according to the hydrometeor classes, also of Praz et al. (2017).




Figure 6. $m(D_{max})$ power law relations for the data of various field campaigns, as listed in Table 5-4 and 45. Left: data color-coded according to riming index R_c of Praz et al. (2017). Right: data color-coded according to the hydrometeor classes, also of Praz et al. (2017). Different curves of the same color correspond to different field campaigns.



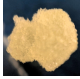

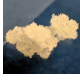




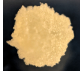







Figure A1. Example, for a reconstructed snowflake, of the distribution of the weight W^* as a function of the threshold on the voxel density content ρ . The weight is displayed here as normalized between 0 and 1. The threshold maximizing W^* is considered as optimal and it is used to censor the data to calculate geometric descriptors. A maximum could always be obtained for all the reconstructed snowflakes.

Table 1. 3D printed snowflakes used for the evaluation of the 3D reconstruction

ID	LWP	# captures	D_{max}	Mass*	<u>Density**</u>	Image
1	1.8 kg m^{-2}	17	4.91 mm	8.47 mg	<u>267</u> kg m^{-3}	
2	2.0 kg m^{-2}	19	4.78 mm	8.77 mg	<u>279</u> kg m^{-3}	
3	1.15 kg m^{-2}	19	3.12 mm	1.9 mg	<u>280</u> kg m^{-3}	
4	0.5 kg m^{-2}	16	4.65 mm	2.24 mg	<u>150</u> kg m^{-3}	
5	0.5 kg m^{-2}	14	4.14 mm	2.22 mg	<u>160</u> kg m^{-3}	
6	0.1 kg m^{-2}	3	4.84 mm	0.82 mg	<u>40</u> kg m^{-3}	
7	0.2 kg m^{-2}	15	4.42 mm	0.92 mg	<u>74</u> kg m^{-3}	
8	0.1 kg m^{-2}	1	4.67 mm	0.81 mg	<u>56</u> kg m^{-3}	
9	2.0 kg m^{-2}	19	4.46 mm	7.6 mg	<u>288</u> kg m^{-3}	
10	1.3 kg m^{-2}	19	4.50 mm	4.67 mg	<u>238</u> kg m^{-3}	
11	1.1 kg m^{-2}	19	4.11 mm	4.58 mg	<u>238</u> kg m^{-3}	
12	0.3 kg m^{-2}	6	4.96 mm	1.23 mg	<u>89</u> kg m^{-3}	
13	2 kg m^{-2}	20	3.95 mm	4.44 mg	<u>281</u> kg m^{-3}	
14	0.2 kg m^{-2}	11	3.91 mm	0.83 mg	<u>127</u> kg m^{-3}	

* This is the (ice) mass of the simulated snowflake and not the actual mass of the (polymer) replica.

** Calculated using the volume of the convex hull

Table 2. Summary of Root Mean Square Error (RMSE), Normalized-RMSE (NRMSE, normalized on the mean value of reference data) and Mean Normalized Bias (MNB) resulting from the comparison with the 3D-printed snowflake replicas.

	RMSE	NRMSE [%]	MNB [%]
Mass [$\text{kg} \times 10^{-6}$]			
3D-GAN	1.8	42	8
BL06	2.9	68	-39
M07	4.4	103	-72
D_{max} [$\text{m} \times 10^{-3}$]			
3D-GAN	0.5	12	7
2D views	0.3	6	1
V_{CH} [$\text{m}^3 \times 10^{-9}$]			
3D-GAN	6.5	35	19
r_g [$\text{m} \times 10^{-3}$]			
3D-GAN	0.2	16	13

Table 3. List of field installations of the MASC instrument, for which data are shown in Fig 6 and 5.

Name	Location	Period
Icegenesis 2020	Swiss Jura mountains	Dec 2020 - Mar 2021
Davos 2015	Swiss Alps	Oct 2015 - Jun 2016
APRES3 2015	East Antarctica	Nov 2015 - Jan 2016
Valais 2016	Swiss Alps	Dec 2016 - Apr 2017
Jura 2019	Swiss Jura mountains	Nov 2019 - Apr 2020

Table 4. Values of the parameters of the relation $m = a_m D_{max}^{b_m}$ estimated on the datasets of different field campaigns for various degrees of riming. m is estimated with 3D-GAN, while R_c is the normalized riming index as in Praz et al. (2017), averaged over the three camera views. D_{max} is the maximum dimension obtained from the triplet of images of the MASC.

R_c [-]	0.0–0.25	0.25–0.50	0.50–0.75	0.75– 1.0 <u>0.99</u>	1.0 <u>>0.99</u>	All
Icegenesis 2020						
samples	705 <u>183</u>	409 <u>552</u>	591 <u>386</u>	304 <u>392</u>	571 <u>257</u>	
a_m [kg m ^{-b_m}]	0.045 <u>0.014</u>	0.018 <u>0.017</u>	0.052 <u>0.09</u>	0.492 <u>1.457</u>	22.622 <u>49.582</u>	0.103 <u>0.037</u>
b_m [-]	2.05 <u>1.86</u>	1.84 <u>1.83</u>	1.97 <u>2.05</u>	2.26 <u>2.46</u>	2.86 <u>2.97</u>	2.11 <u>1.91</u>
Davos 2015						
samples	3584 <u>1864</u>	4766 <u>7509</u>	10288 <u>8532</u>	6865 <u>8973</u>	4965 <u>3144</u>	
a_m [kg m ^{-b_m}]	0.007 <u>0.003</u>	0.013 <u>0.014</u>	0.050 <u>0.135</u>	0.728 <u>1.282</u>	21.152 <u>135</u>	0.042 <u>0.025</u>
b_m [-]	1.72 <u>1.58</u>	1.78 <u>1.80</u>	1.95 <u>2.13</u>	2.32 <u>2.43</u>	2.84 <u>3.13</u>	1.91 <u>1.83</u>
APRES3 2015						
samples	1009 <u>285</u>	987 <u>674</u>	1761 <u>642</u>	1080 <u>715</u>	1761 <u>257</u>	
a_m [kg m ^{-b_m}]	0.029 <u>0.012</u>	0.045 <u>0.021</u>	0.193 <u>0.158</u>	1.335 <u>2.156</u>	4.972 <u>16.199</u>	0.197 <u>0.079</u>
b_m [-]	1.95 <u>1.82</u>	1.99 <u>1.85</u>	2.19 <u>2.15</u>	2.45 <u>2.53</u>	2.64 <u>2.81</u>	2.19 <u>2.04</u>
Valais 2016						
samples	3331 <u>444</u>	3036 <u>2175</u>	4962 <u>1928</u>	3417 <u>2214</u>	5508 <u>921</u>	
a_m [kg m ^{-b_m}]	0.02 <u>0.05</u>	0.024 <u>0.014</u>	0.073 <u>0.133</u>	0.375 <u>0.869</u>	13.409 <u>183.131</u>	0.077 <u>0.02</u>
b_m [-]	1.89 <u>1.66</u>	1.88 <u>1.79</u>	2.03 <u>2.13</u>	2.22 <u>2.37</u>	2.77 <u>3.17</u>	2.02 <u>1.80</u>
Jura 2019						
samples	520 <u>164</u>	705 <u>1313</u>	1797 <u>1074</u>	1150 <u>1026</u>	2762 <u>1204</u>	
a_m [kg m ^{-b_m}]	0.026 <u>0.272</u>	0.025 <u>0.0234</u>	0.095 <u>0.169</u>	0.627 <u>0.786</u>	122.652 <u>155.926</u>	0.150 <u>0.032</u>
b_m [-]	1.93 <u>2.35</u>	1.89 <u>1.94</u>	2.07 <u>2.17</u>	2.31 <u>2.35</u>	3.13 <u>3.16</u>	2.11 <u>1.84</u>

Table 5. Values of the parameters of the relation $m = a_m D_{max}^{b_m}$ estimated on the datasets of different field campaigns for various hydrometeor types. m and is estimated with GAN-3D, while the hydrometeor type is obtained with the method of Praz et al. (2017). D_{max} is the maximum dimension obtained from the triplet of images of the MASC.

<i>Type</i>	AG	GR	COL	CPC	PC
Icegenesis					
samples	1148 <u>1060</u>	762 <u>458</u>	397 <u>110</u>	8 <u>5</u>	167 <u>137</u>
a_m [kg m ^{-b_m}]	0.016 <u>0.017</u>	22.913 <u>34.5</u>	0.005 <u>0.011</u>	-	0.379 <u>0.354</u>
b_m [-]	1.77 <u>1.78</u>	2.86 <u>2.93</u>	1.73 <u>1.84</u>	-	2.29 <u>2.27</u>
Davos 2015					
samples	18130 <u>19343</u>	7789 <u>7404</u>	883 <u>619</u>	94 <u>130</u>	2204 <u>2526</u>
a_m [kg m ^{-b_m}]	0.024 <u>0.023</u>	27.073 <u>58.795</u>	0.011 <u>0.004</u>	0.002	0.148 <u>0.158</u>
b_m [-]	1.83 <u>1.82</u>	2.89 <u>3.02</u>	1.83 <u>1.67</u>	1.48 <u>1.54</u>	2.12 <u>2.14</u>
APRES3 2015					
samples	2184 <u>1341</u>	2841 <u>787</u>	541 <u>160</u>	301 <u>120</u>	317 <u>165</u>
a_m [kg m ^{-b_m}]	0.059 <u>0.03</u>	4.158 <u>10.8</u>	0.0147 <u>0.026</u>	0.023 <u>0.016</u>	0.213 <u>0.277</u>
b_m [-]	1.99 <u>1.88</u>	2.62 <u>2.77</u>	1.87 <u>1.95</u>	1.88 <u>1.82</u>	2.2 <u>2.24</u>
Valais 2016					
samples	9695 <u>5100</u>	7450 <u>1946</u>	959 <u>172</u>	98 <u>18</u>	1129 <u>446</u>
a_m [kg m ^{-b_m}]	0.035 <u>0.017</u>	13.560 <u>87.326</u>	0.009 <u>0.007</u>	0.002 <u>-</u>	0.213 <u>0.621</u>
b_m [-]	1.9 <u>1.77</u>	2.78 <u>3.07</u>	1.8 <u>1.76</u>	1.55 <u>-</u>	2.21 <u>2.37</u>
Jura 2019					
samples	2980 <u>2940</u>	3399 <u>1667</u>	191 <u>47</u>	14 <u>8</u>	169 <u>119</u>
a_m [kg m ^{-b_m}]	0.033 <u>0.030</u>	121.319 <u>126.612</u>	0.062 <u>-</u>	-	0.071 <u>0.465</u>
b_m [-]	1.87 <u>1.86</u>	3.13 <u>3.14</u>	2.08 <u>-</u>	-	2.01 <u>2.33</u>

AG: aggregates, GR: graupel, COL: columns, CPC: combination of planar crystals and columns, PC: planar crystals. Results are reported only if at least 80 samples are available.