Novel Multiview Machine Learning Classification of Snowflakes: Harnessing Convolutional Neural Networks and Multiangle Multicamera Instruments

Hein Thant and Branislav M. Notaroš

Colorado State University, Electrical and Computer Engineering Department, Fort Collins, Colorado

Revised Manuscript, Second Revision

Corresponding author: Branislav M. Notaroš, notaros@colostate.edu

Early Online Release: This preliminary version has been accepted for publication in *Journal of Atmospheric and Oceanic Technology*, may be fully cited, and has been assigned DOI 10.1175/JTECH-D-24-0142.1. The final typeset copyedited article will replace the EOR at the above DOI when it is published.

ABSTRACT

Classification of snowflakes based on their geometric shape, degree of riming, and melt/dry state can improve understanding, characterization, and quantification of other geometrical, microphysical, and scattering properties of ice particles. For example, classification provides essential ground truth data for interpreting polarimetric radar signatures of snow while validating and advancing radar-based quantitative precipitation estimation. High-resolution photographs of snowflakes obtained by emerging multicamera instruments are well suited for snowflake classification, which, coupled with recent machine learning techniques based on Convolutional Neural Networks (CNNs), enable methods for accurate and fast automatic classification of snowflakes using images. Given that the appearance of a snowflake generally changes significantly with viewing angle, this work proposes and presents a novel multiview snowflake classification methodology based on the high-resolution photographs of frozen hydrometeors in freefall from multiple views collected by the multicamera instruments. The approach employs machine/deep learning algorithms leveraging multiangle camera systems and enhanced supervised CNN-based techniques to achieve precise classification of snowflakes based on their geometrical categories and accurate and reliable estimates of specific snowflake properties, such as riming degree and melt/dry state. This represents the first multiview snowflake-classification framework that takes full advantage of multiview camera systems. Presented multiview classification results show record accuracies of 98.57%, 98.22%, and 95.83% for geometric classes, riming degree, and melt/dry state, respectively.

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

SIGNIFICANCE STATEMENT

This work proposes and presents a novel multiview snowflake classification methodology leveraging recent developments in machine learning, multiview classification, and multiangle multicamera instruments for acquiring high-resolution photographs of frozen hydrometeors in freefall from multiple views. The results for multiview classification show record accuracies for snowflake geometric classification, riming degree estimation, and melt/dry state estimation, respectively, significantly outperforming other classification models in each of the same categories. Automatic multiview machine-learning-based winter hydrometeor classification enhances understanding, characterization, and quantification of geometrical, microphysical, and scattering properties of ice and snow hydrometeors. These improvements are essential for quantitative precipitation estimation algorithms and for microphysical parameterizations employed in numerical winter-weather forecast models and regional climate projections, with impacts on economy, safety, and everyday life.

1. Introduction

The complexities of snowflake formation and snowfall present considerable difficulties for observation, measurement, analysis, interpretation, and prediction (Dozier et al. 2016, Bukovčić et al. 2018). Snow and ice particles exhibit a wide natural variability in shape, size, internal composition, density, and habit, and this variability is amplified by their pronounced sensitivity to subtle environmental changes (Pruppacher & Klett, 2010, Ludlam and Mason 1957, Szyrmer and Zawadzki 2014, Notaros 2021). A comprehensive understanding of geometrical, microphysical, and scattering properties of ice and snow hydrometeors is essential for the development of radar-based quantitative precipitation estimation (QPE) algorithms and for microphysical interpretation of polarimetric radar signatures (Huang et al. 2015, Leinonen et al. 2013, Heymsfield et al. 2016, Ryzhkov and Zrnic 2019, Notaros 2022), and constitutes a major unresolved challenge in the field.

While polarimetric radars [e.g. NEXRAD (Heiss et al. 1990), NASA NPOL (Wolff et al. 2023), CSU-CHILL (Bunkrow et al. 2000), etc.] provide extensive coverage for snow observations and analyses, the integration of ground-based in situ instruments—which deliver high-detail measurements but possess limited spatial resolution—with radar data yields a more robust characterization of snow composition (Zhang et al. 2011), namely the type of hydrometeors inside a storm. This integration is essential for accurately representing snow microphysics within storms. Frequently used surface platforms for camera-based observation and validation include the Two-Dimensional Video Disdrometer (2DVD; Schönhuber et al. 2008) and the Precipitation Instrument Package [PIP; a more advanced version of the Snow Video Imager described in Newman et al. (2009), and Pettersen et al. (2020, 2021)]. Instruments specifically designed for Multiview snowflake imaging—such as the Multi-Angle Snowflake Camera (MASC; Garrett et al. 2012, Suefer & Bailey 2016), the Video In Situ Snowfall Sensor (VISSS; Maahn et al. 2023), and the Snowflake Measurement and Analysis System (SMAS; Thant et al. 2022)—provide additional characterization capabilities and can acquire thousands of snowflake images per hour.

Classification of precipitation, namely, assigning observed winter hydrometeors to predefined categories, can enhance understanding, characterization, and quantification of other geometrical, microphysical, and scattering properties of snowflakes, as well as enrich our understanding of polarimetric radar signatures of snow and advance radar-based QPE (e.g., Kennedy et al. 2018). It can significantly improve the radar-based estimation of liquid

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

equivalent snow rates and QPE by first classification of precipitation type followed by quantification, and thus implementation of the concept of classification first and then quantification (Straka et al. 2000) to winter precipitation.

The high-resolution photographs of snowflakes collected by multiangle multicamera instruments for snowflake research, such as the MASC and SMAS, are exceptionally well-suited for snowflake classification. Depending on visual inspection to classify particles is not practical for large-scale snowflake classification. Fortunately, recent advances in artificial intelligence and machine-learning techniques now enable accurate and fast automatic classification of snowflakes from these images.

Snowflake classification comprises categorization into geometric categories, degrees of riming (Praz et al. 2017, Hicks and Notaros 2019), and melt/dry states (Chase et al. 2022, Liao and Meneghini 2022), respectively. Snowflake geometric categories typically include Small Particle (S), Columnar Crystal (C), Planar Crystal (P), Graupel (G), Aggregate (A), and Columnar and Planar Combination (CP). Riming, as one of the crucial ice crystal growth processes that is based on the collection of supercooled water droplets onto an ice crystal's surface, altering its shape, size, orientation and density, is classified primarily into five categories: unrimed, rimed, densely rimed, graupel-like and graupel. Melt/dry state represents whether the snowflake has melted on its descent.

Praz et al. (2017) established the earliest benchmark for hydrometeor classification research, utilizing the MASC, and reported a classification accuracy of 95% over six classes of geometry specified above. They employed a multinomial logistic regression approach which required external feature extraction and optimization; this approach constrained the number of input features and number of classes due to computational inefficiencies and insufficient data for more classes. Moreover, the method introduced device-specific bias, as the extracted features reflected imperfections inherent to the particular instrument, reducing its generality and versatility and hindering its application across different measurement platforms.

In 2019, Hicks and Notaroš (2019) leveraged recent advancements in machine learning by applying Convolutional Neural Networks (CNNs) for fully automatic winter-hydrometeor classification. They achieved an overall accuracy of 93.4% for five classes of snowflake geometry with the CP category omitted due to insufficient data and 92.4% for riming degree estimation. Extending this approach, Key et al. (2021) incorporated more data and advanced training techniques to improve the accuracy of geometric classification to 96.2%. However, due to insufficient data availability and labeling, riming estimation was not performed in that study.

Leinonen and Berne (2020) investigated unsupervised learning techniques, using K-clustering and Generative Adversarial Networks (GAN) to carry out snowflake classification. Although these approaches demonstrate potential for discriminating and classifying hydrometeor images, they generate their own categories rather than mapping images pre-defined class labels. This study resulted in more riming degree discriminations, while geometric classification results were insufficient.

Multiview classification is a relatively novel classification concept based on utilizing multiview images, i.e., images of the object acquired from several viewing angles. Su et al. (2015) introduced the concept of Multiview Convolutional Neural Network architecture for recognizing 3D objects by aggregating information from multiple 2D views of computer-designed 3D objects into single and compact representations. Seeland and Mäder (2021) conducted a systematic study, exploring various previously reported fusing methods across different application domains, and proposed a comprehensive classification scheme that relies on an effective fusing of visual information captured through images depicting the same object from multiple perspectives.

This paper leverages the advantage of multiview snowflake photography and state-of-the-art machine learning techniques, specifically those based on CNNs and multiview CNN architectures. It proposes machine/deep learning algorithms with full utilization of multiview camera systems and enhanced supervised CNN-based techniques to achieve precise classification of snowflakes based on their geometrical categories and accurate and reliable estimations of specific snowflake properties, such as riming degree and melt/dry state.

While Seeland and Mäder (2021) applied multiview classification to objects very different from snowflakes, multiview image classification is especially suitable and beneficial for snowflakes, in conjunction with multicamera instruments, such as MASC and SMAS. Most snowflakes look vastly different from different viewing angles, and full utilization of multiview camera systems and multiview machine-learning-based classification enhances expert labeling accuracy and classification accuracy, as well as the reliability and versatility of the methodology.

The paper presents the development, implementation, and evaluation of a multiview CNN specifically designed to classify snowflakes using multiview high-resolution images from three views, combined with and enabled by techniques for image preprocessing, camera calibration, and snowflake matching, all created towards multicamera enabled instruments. The steps needed to develop a hydrometeor classifier using multiview CNNs are presented in detail and advantages to in situ research are highlighted. This approach benefits from the properties of CNNs such as the adaptability and generalizability over multiple instruments and transfer learning capability to adjust from one dataset to another in the same domain, namely, additional snowflake categories and different hydrometeor types, without expert knowledge (Donahue et al. 2014, Li et al. 2021). On the other hand, the concept of multiview preserves the dimensionality of the snowflakes better, improving the stability and reliability of snowflake manual labeling and classification algorithms. The algorithm is trained on tens of thousands of labeled data, and finetuned by hyperparameter tuning.

The paper is organized as follows. Section 2 describes the acquisition of raw snowflake images using MASC and SMAS and outlines various field deployments of instruments with data collections. Section 3 explains the preprocessing pipeline to transform raw images into data ready to be used for multiview classification, including image cleanup and cropping, calibration of cameras, and matching algorithm to form the final snowflake triplet (grouped images of the same snowflake from three camera views). Section 4 presents the proposed novel multiview snowflake classification methodology, describing the training dataset used, alongside the defined categories for the classification, the multiview CNN algorithm and architecture, the hyperparameter tuning performed, the testing methods, as well as hardware and software used in this work. Results and discussion appear in Section 5.

2. Instrumentation and Data Collection

a. Multi-Angle Snowflake Camera

The Multi-Angle Snowflake Camera (MASC; Figure 1), a commercially available instrument, uses three cameras in the horizontal plane separated by 36° to capture high-resolution photographs of snowflakes or other frozen hydrometeors in freefall from three views, while simultaneously measuring their fall speed. It was first introduced and tested in a field campaign in 2011 (Garrett et al. 2012). The pixel resolution is 50 μm, the virtual

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

measurement area is 30 cm^2, and the measurement volume rate is ~200 cm^3/s. The instrument includes two near-IR emitter-receiver pair arrays positioned one above the other, and as a particle falls through the upper array, it triggers all cameras. Additionally, the particle fall speed is obtained from the fall time between the two triggers. In the past, MASC observations have been employed to improve and constrain snowfall retrieval algorithms (Cooper et al. 2017, Schirle et al. 2019).

To enable better 3D shape reconstruction, two additional cameras were installed at a 55° angle with respect to the horizon to the MASC at Colorado State University (CSU), to provide nonplanar views (Kleinkort et al. 2017; Notaros 2021), as shown in Figure 1. Furthermore, the particle size distribution (PSD) can be estimated (Huang et al. 2017) and snowflake geometry, riming degree, and melt/dry states can be classified (Hicks and Notaroš 2019; Key et al. 2021; Thant et al. 2023b).

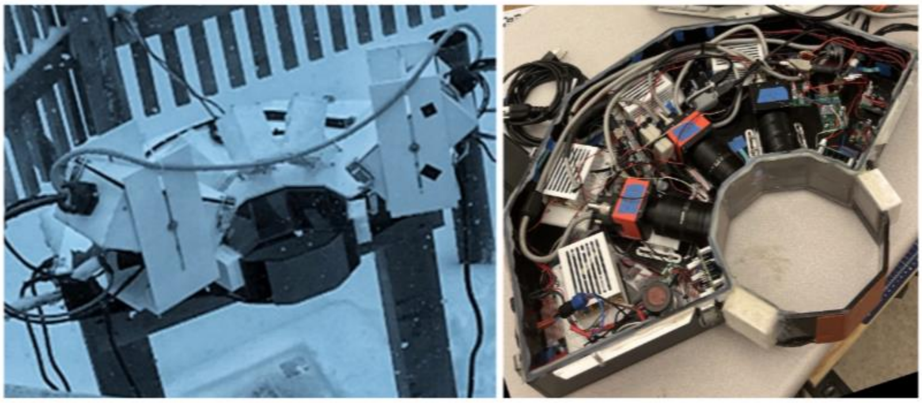

Fig. 1. Multi-Angle Snowflake Camera (MASC): photographs of the CSU-MASC exterior with two “external” cameras added (left) and interior (right).

However, image-based classification algorithms are constrained by the reduction of three-dimensional snowflakes to two-dimensional projections. This dimensionality loss can introduce labeling errors, misclassifications, and parallax distortions when a poor angle of the snowflake is captured. Humans typically integrate visual information from multiple views, whereas conventional machine-learning image classification models infer an object’s class from a single image. For challenging classification cases—such as classifying snowflakes—a single view often lacks sufficient information for accurate classification. Figure 2 illustrates images of the same snowflake (each row) taken from different views of the MASC (Figure 1), highlighting the mislabeling and misclassification that arise when images are classified using only a single viewing angle.

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

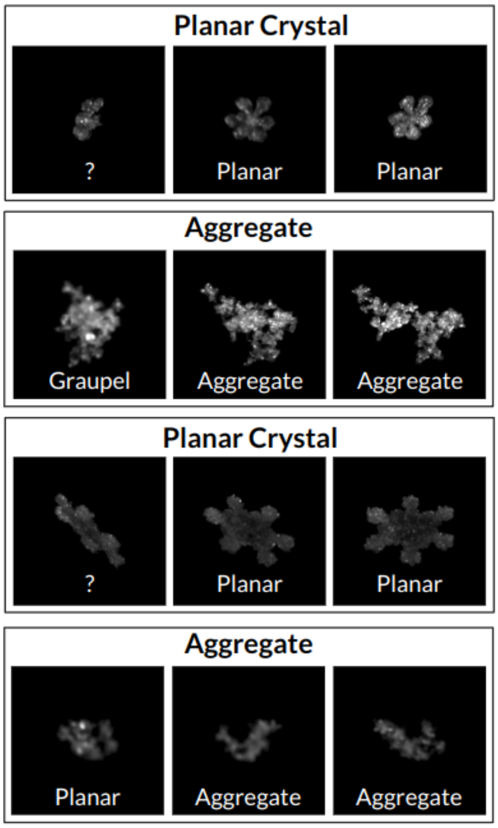

Fig. 2. Example images of snowflakes from multiple views where single-view classification or manual labeling of a snowflake from certain angles is erroneous. Each row represents the same snowflake viewed from three different angles.

b. Snowflake Measurement and Analysis System

The Snowflake Measurement and Analysis System (SMAS; Figure 3) is a novel system for measurement and analysis of snow particles in freefall, designed, developed, and built at CSU. The main features of the SMAS are: seven high-resolution cameras strategically arranged in three-dimensional configuration for reconstruction of shapes of snowflakes based on photographs of particles in freefall from multiple views (Kleinkort et al. 2017) and capability for measurement and analysis of multiple snowflakes at once, namely, processing of images with multiple snowflakes per frame (Thant et al. 2022).

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

Fig. 3. Snowflake Measurement and Analysis System (SMAS): photographs of the instrument exterior (left) and interior (right).

c. Deployments and Image Collection

Large quantities of unprocessed and unlabeled images have been collected by both the SMAS and MASC during the following field campaigns, with the MASC possessing a longer operational record. The CSU-MASC was utilized for several winters between 2014–2017 at the MASCRAD (MASC + Radar) project site in Greely, Colorado (Notaroš et al. 2016; Notaros 2021). The data collected during MASCRAD observations were used in the development of classification algorithms and studies of Hicks and Notaroš (2019) and Key et al. (2021). The MASC was also deployed during the International Collaborative Experiment for the Pyeongchang 2018 Olympic and Paralympic (ICE-POP) winter games, in Korea. Both SMAS and MASC were jointly deployed in NASA’s Wallops Flight Facility [WFF; Global Precipitation Measurement (GPM) site] in Virginia during the winter of 2021–2022 (Notaros et al. 2022; Thant et al. 2023a, 2024), and later, during the subsequent winters, 2022–2024, at the University of Connecticut as part of the Investigation of Microphysics and Precipitation for Atlantic Coast-Threatening Snowstorms (IMPACTS) field campaigns (McMurdie 2020; Filipiak et al. 2024). Figure 4 showcases photographs from each deployment.

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

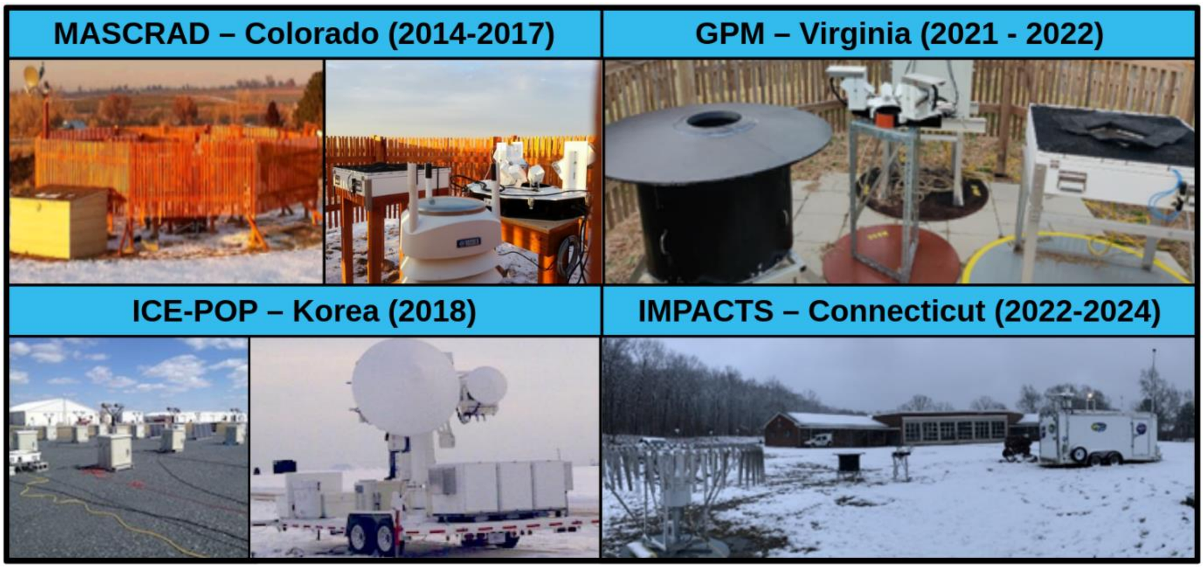

Fig. 4. Deployment sites of the MASC within MASCRAD and ICE-POP field campaigns, and MASC and SMAS within NASA WFF-GPM and IMPACTS experiments.

3. Image Preprocessing, Multiview Camera Calibration, and Snowflake Matching

As a precursor of snowflake classification, the first step of the image preprocessing workflow to transform raw images collected by the MASC and SMAS into data usable for multiview classification comprises image cleanup and cropping. These techniques were developed to operate robustly under noisy conditions, glares, and harsher environments where wind and humidity are heavy, promoting ice and snow buildup on the instrumentation. Background subtraction is performed using a four-image average, which eliminates both temporary and permanent artifacts, while ensuring parts of snowflakes in the foreground image are not mistakenly removed with the snowflakes in the background images while subtracting. Bounding boxes are then generated for each individual snowflake and cropped to 300x300 pixel PNG images. The software implementation is fully optimized and parallelized for fast real-time processing and can be implemented using MATLAB or Python. For each cropped image, the Spatial and Spectral Sharpness (S3; Vu and Chandler 2009) index is computed to assess image blur; images with an S3 value below 0.575 are categorized as blurry and excluded during quality-control filtering.

The second step is multiview camera calibration, performed to estimate the fundamental matrix, which characterizes the epipolar geometry between a pair of cameras. The CALTag

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

checkerboard (Atcheson et al., 2010) was employed because of its low reprojection error, providing fixed global coordinates that are simultaneously visible to all cameras. Detected checkerboard corners (Zhang 2000) served as correspondences for computing the fundamental matrix via the eight-point algorithm combined with RANdom SAmple Consensus (RANSAC) to find the fundamental matrix of each camera pair (Hartley and Zisserman 2004; Collins 2012). RANSAC parameters were set to a projection threshold of 0.1 and a confidence level of 0.99 to ensure accurate and reliable estimation of each camera-pair fundamental matrix.

The final step involves matching corresponding snowflake images across the three cameras to form a final triplet of images of the same snowflake from three camera views, ready for multiview classification. This was performed adapting the concept of epipolar matching (Hartley and Zisserman 2004), with pixel coordinates of all preprocessed images from one camera and the corresponding fundamental matrix for the camera pair being used to generate epipolar lines for each snowflake visible to the corresponding camera. The closest snowflakes on the corresponding camera that are under a 20-pixel distance to the epipolar lines are then assigned as a pair candidate. This was repeated for the pair candidates in reverse direction, forming a new set of epipolar lines on the first camera to confirm the snowflake pair, and such bidirectional verification was then repeated for the next camera pair. Finally, overlapping matches from the two camera pairs were merged to form triplets (Figure 5), as input to multiview snowflake classification.

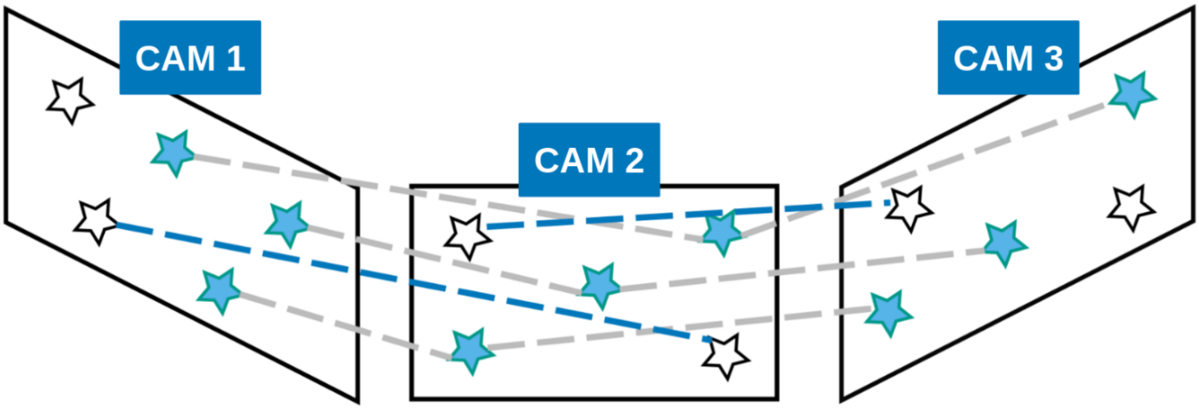

Fig. 5. Illustration of snowflake matching for three cameras. The snowflakes for which the procedure found successful matches across all three views are light-blue colored, with gray dashed lines connecting them. These are the snowflakes that are selected for further classification processing, whereas the remaining ones are discarded from multiview classification.

Data processing (image preprocessing, camera calibration, and snowflake matching) pipeline is fast and computationally efficient, designed to perform operations in approximately 30-50 milliseconds for each image on Nvidia single-board computers. This makes it suitable for deployment on computers typically running inside in situ measurement devices for real-time processing.

4. Novel Snowflake Classification Methodology

a. Multiview CNN Classification Scheme

Of importance for this work are further advancements of the concept of Multiview CNN architecture of Su et al. (2015) by Wang et al. (2015), Feichtenhofer et al. (2016), Setio et al. (2016), Dolata et al. (2017), Do et al. (2017), Feng et al. (2018), Geras et al. (2018), Lin and Kumar (2018), Lee et al. (2018), and Barbosa et al. (2020). In addition, notable and relevant applications of various versions of multiview classification are those in biological domains such as biometrics, plant biology, agronomy, etc. (Lin and Kumar 2018, Do et al. 2017, Lee et al. 2018).

Most importantly, Seeland and Mäder (2021) examined different fusion strategies for multiview classification and found that late fusion conducted through concatenation of the fully-connected layer resulted in drastic improvements over other strategies – consistently across all tested domains at various scales and amounts of data. Consequently, we adopt the late-fusion concatenation scheme as the backbone of our multiview snowflake classification methodology, leveraging the multiview capabilities of MASC and SMAS.

b. Training dataset and categories

For training the models, this study employed the MASCDB dataset, the only existing database of multiview images of snowflakes, collected by a MASC (Grazioli et al. 2022, 2023). The database consists of over 2.5 million snowflake images grouped into image triplets (three cameras/views of a conventional MASC). Hence, we concentrate on hydrometeor geometry classification, riming degree estimation, and identification of melting particles, from three camera views, using the grouped triplets of the MASCDB.

The MASCDB database features snowflakes labeled in six basic categories: Small Particle (S), Columnar Crystal (C), Planar Crystal (P), Graupel (G), Aggregate (A), and Columnar and Planar Combination (CP), as illustrated in Figure 6. The riming degree estimation includes five different degrees: unrimed (1), rimed (2), densely-rimed (3), graupel-like (4), and graupel (5), as shown in Figure 7. Figure 8 depicts the identification of the occurrence of melt/dry (M/D) particles.

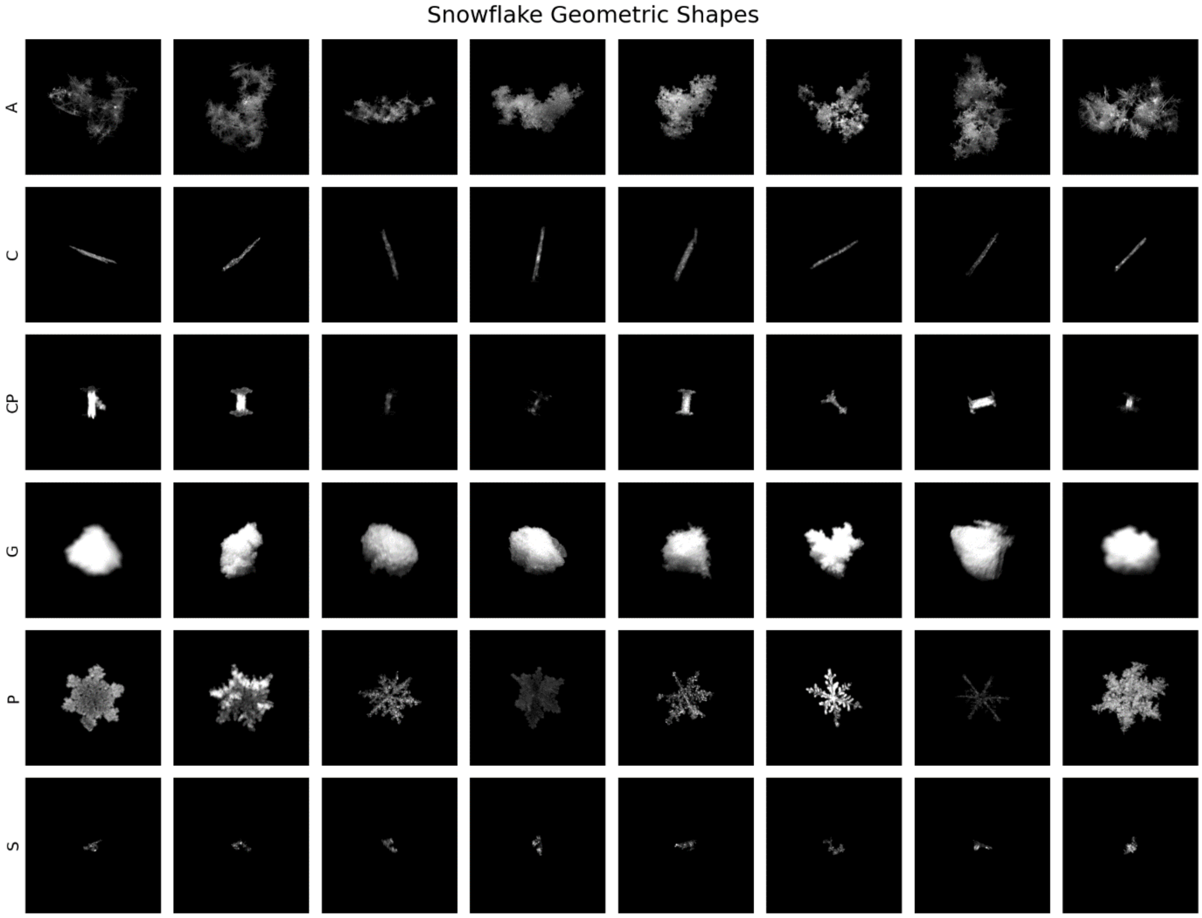

Fig. 6. Example images of snowflake geometric shapes from the MASCDB, showcasing (in the respective rows of images) Aggregate (A), Columnar Crystal (C), Columnar and Planar Combination (CP), Graupel (G), Planar Crystal (P), and Small Particle (S) classes.

Individual MASCDB campaigns were selected as test sets: namely, POPE-2020 (Ferrone and Berne 2022) was used as the test campaign for shape classification, with the data amount limited to 500, whereas riming and melt/dry estimation test campaigns were selected from ICEGENESIS-2021 (Billault-Roux et al. 2023), and the data was limited to 400 and 250, respectively. Image quality control filters (included in the MASCDB) were used to exclude lower quality images, and that simultaneously introduced an unbalanced number of images for each class, which, in turn, was mitigated by limiting the maximum number of images per class for training, validation, and testing, respectively. Then the dataset was split sequentially into 80% training and 20% validation, and per-class caps for the combination of training and

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

validation for shape, riming, and melt/dry classifications were set to 3600, 3000, and 1500, respectively. Finally, the training employed inverse class-frequency weighting (as implemented by the scikit-learn Python package), where underrepresented classes were assigned higher weights so that categories with more samples did not dominate.

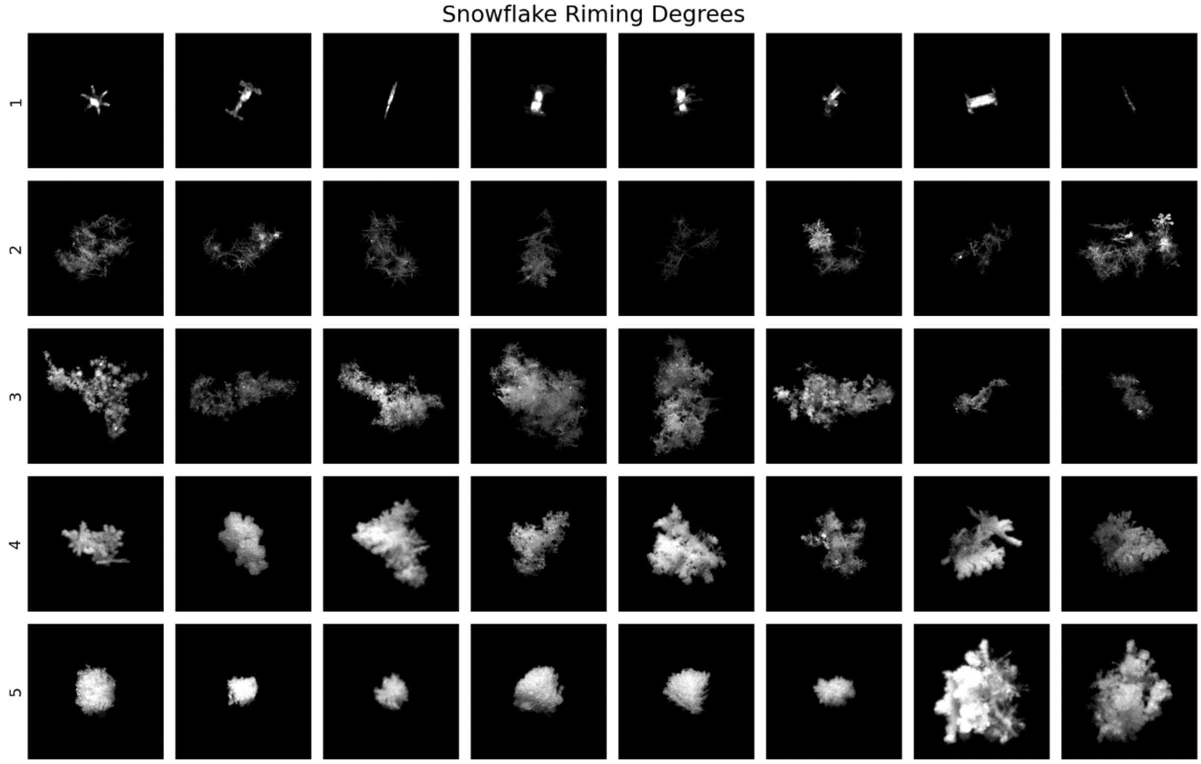

Fig. 7. Example images of snowflake riming degrees, from the MASCDB, exhibiting (in the respective rows of images) unrimed (1), rimed (2), densely-rimed (3), graupel-like (4), and graupel (5) classes.

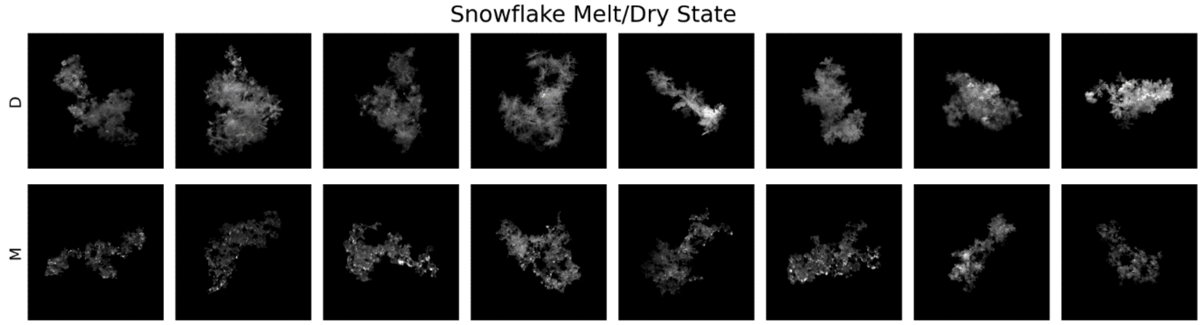

Fig. 8. Example images of snowflake dry state (D) and melted state (M), from the MASCDB.

c. Snowflake Classification Algorithm

Classification algorithm in this study combines two techniques: Multiview Late-Stage Fusing using concatenation and a Convolutional Neural Network based on the Inception-v3 architecture (Szegedy et al. 2015). Inception-v3 was selected for its favorable trade-off

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

between efficiency, in terms of computational resources and number of trainable parameters, and accuracy. Although the subsequent Inception-v4 performs slightly better, it requires substantially greater computational resources. Consequently, Inception-v3 was adopted to achieve balanced performance and efficiency when used with Multiview Late-Stage Fusing.

In particular, Inception-v3 architecture is designed to minimize the number of connections and parameters while preserving network efficiency (Szegedy et al. 2015; Brital 2021; Tsang 2019), which is accomplished through factorization into smaller convolutions and asymmetric convolutions reducing the number of parameters required. These techniques allow for promotion of high-dimensional representations without increasing computation costs, and appear in three main modules, labeled as Inception Modules A, B and C (Tsang 2019; Brital 2021). Inception Module A replaces a single 5x5 convolution with two successive 3x3 convolutions, decreasing the parameter count from 25 to 18 (a 28% reduction). Module B employs a 3x1 convolution, followed by a 1x3 one to emulate a 3x3 filter, reducing parameters from 9 to 6 (a 33% reduction). Module C utilizes a variety of asymmetric convolutions to further enrich high-dimensional representations. For down-sampling feature maps, Inception-v3 adopts a more efficient strategy than the max-pooling approaches of AlexNet and VGGNet. Namely, grid-size reduction Module A performs a convolution and max-pooling of stride 2, while Module B applies the same down-sampling principle combined with asymmetric convolutions, adding to the overall computational efficiency of the model.

This study implemented a pre-trained instance of Inception-v3 from the Keras library, omitting auxiliary classifiers and label-smoothing mechanisms that are present in the full architecture. The principal components of the particular Inception-v3 configuration used—including Inception Modules A, B, and C; Efficient Grid-Size Reduction Modules A and B; and the initial stem block—are depicted in Figure 9, which also illustrates how these elements are integrated to form the convolution stage.

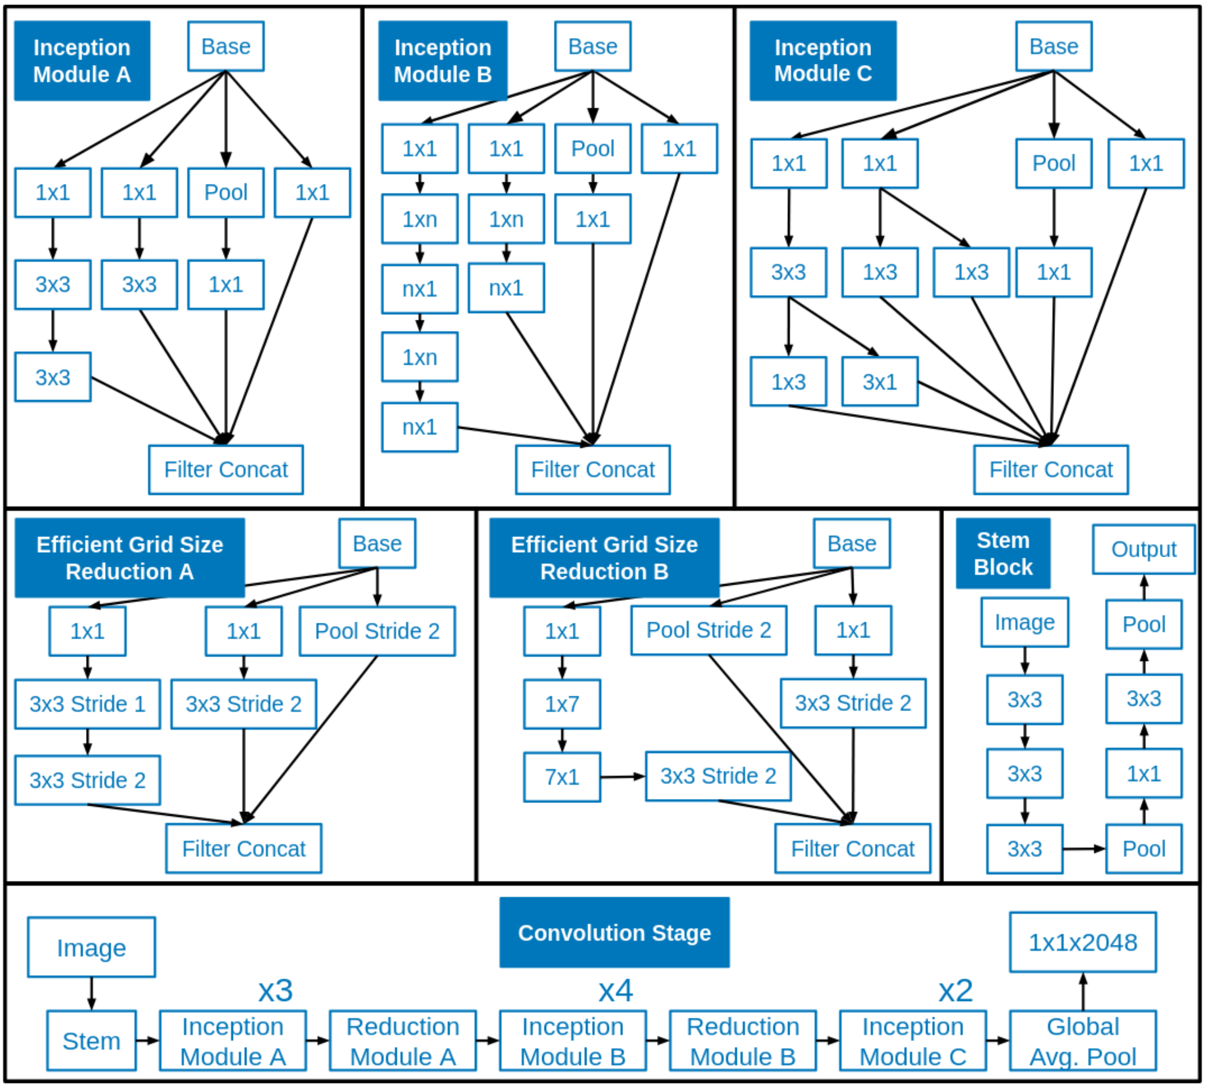

Fig. 9. Visual depictions of key components of Inception-v3 CNN implemented and used in this study, including three Inception Modules, two Efficient Grid Size Reduction Modules, and Stem Block. The bottom part of the graph shows the integration of all of the modules into the convolution stage of Inception-v3.

The classification architecture receives three separate image inputs, and each one is processed by an independent Inception-v3 convolutional feature extraction branch. The features generated by the three branches are then concatenated and fed into a fully connected block that comprises of a dense layer, batch normalization, activation function, and dropout regularization layer, resulting in a final classification softmax layer tailored to different scenarios (Geometry, Riming, Melt/Dry; Figure 10). All of the Inception-v3 convolutional feature extraction modules and the fully connected layers were fine-tuned and trained, with a categorical cross-entropy loss function being employed to optimize the model's performance in training.

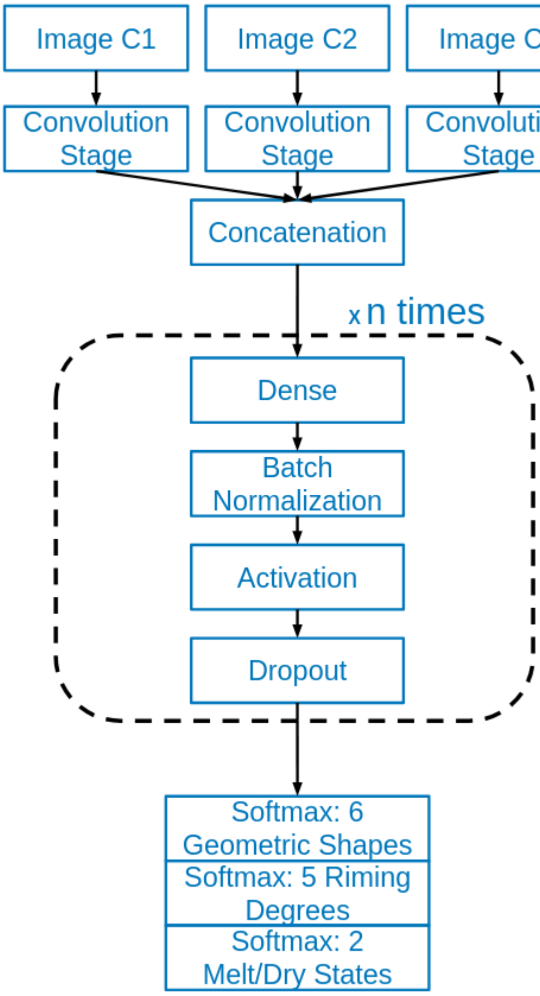

Fig. 10. The final multiview Inception-v3 architecture used to train all models for classifying snowflakes into six geometric shapes, five riming degrees, and two melt/dry states, respectively.

d. Hyperparameter Tuning

As a major part of the presented classification method, the algorithm was finetuned by hyperparameter tuning, namely, a two-phase coarse-to-fine hyperparameter optimization was performed. In the first phase, Bayesian Optimization, specifically the Tree-structured Parzen Estimator Approach (TPE; Bergstra et al. 2011, Akiba et al. 2019) via the Optuna framework, was employed on a reduced subset of data. Sixty trials were executed, with class-wise sample limits set to 1200 for the shape category, 960 for riming degrees, and 720 for melt/dry states. The second phase then re-evaluated ten most promising trial configurations from the first phase on the full training set described in Section 4b. This sequential approach allowed efficient exploration of the hyperparameter space while ensuring that the final model was tuned on the complete data distribution.

Hyperparameters optimized are activation function, optimizer, learning rate, batch size, fully-connected layer configuration, and dropout rate, and they can be described as follows. The activation function defines the non-linear transformation node that calculates the output based on its individual inputs and their weights. The optimizer governs the algorithmic update of the machine learning model parameters and weights. Batch size refers to the number of training samples processed concurrently in a single optimization step. The fully-connected layer configurations comprise different precompiled layers, with varying depth and width, based on which the model maps extracted feature vectors to different classes. The learning rate controls the magnitude of parameter adjustments per gradient step, with larger values accelerating convergence and smaller values providing finer-grained updates. Finally, the dropout rate denotes the fraction of neurons randomly set to zero during training to mitigate over-fitting. Table 1 enumerates ranges of hyperparameters, which are provided to Optuna.

Table 1. Hyperparameters used for finetuning the classification models and their respective ranges.

Hyperparameter	Parameter Range
Activation Function	relu, leaky_relu, gelu, swish
Optimizer	adam, rmsprop, sgd
Learning Rate	0.0001 – 0.1
Batch Size	16 – 64
Fully-connected Layer Config	[2048], [256, 16], [128, 8], [512, 128, 8], [256, 128, 16], [256, 64, 8], [128, 16, 8], [512, 256, 64, 16]
Dropout Rate	0.1 – 0.5

e. Testing Methodology

For the test set, we evaluated several additional performance metrics in comparison between the single-view and multiview configurations, namely, precision, recall, F1-score,

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

and Heidke Skill Score (HSS), which are described as follows. Precision quantifies the proportion of predicted positive instances that are true positives. Recall measures the fraction of actual positive instances correctly identified by the model. The F1-score is a metric obtained as the harmonic mean of precision and recall and attains values near 1 when both precision and recall are high. HSS assesses the improvement of the classifier relative to a random-guess baseline. Collectively, these metrics form a comprehensive assessment of the classification performance across both single-view and multiview models.

f. *Hardware and Software Used*

To efficiently train multiple models during hyperparameter tuning, three computing instances were employed, namely, two Google Colab L4 GPU instances and a local workstation equipped with two Nvidia Tesla P40 GPUs. The distributed setup enabled parallel training of multiple models and training time reductions, with the average training time per classification model of approximately 30 hours for shape and riming classification tasks, and 15 hours, given the smaller dataset, for melt/dry state estimation.

All training procedures were executed within a reproducible environment comprising Anaconda 3, TensorFlow, and Python, and the pretrained Inception-v3 model from Google's TensorFlow Keras repository, which was initially trained on ImageNet, served as the base network. This step is crucial, as training an Inception-v3-based classification model from scratch (random initial weights) requires substantial resources, both computing and time, in addition to the need for extensive datasets, to achieve decent performance. Hyperparameter optimization was conducted with the Optuna Python Library, employing Bayesian Optimization (TPE) to efficiently explore the hyperparameter search space.

5. Results and Discussion

Upon performing a systematic two-phase coarse-to-fine hyperparameter optimization, as outlined in Section 4d, Figure 11 displays the phase-two results of the top ten highest-performing reevaluations, and Table 2 lists the hyperparameter settings that yielded optimal performance for the shape classification, riming degree estimation, and melt/dry state determination, respectively, within the two-phase hyperparameter tuning scheme. Detailed plots of optimization history and hyperparameter importance as well as slice plots for phase one can be found in the Appendix (Figures A1, A2, and A3). The accuracies of geometric shape, riming degree, and melt/dry state classification results are 98.57%, 98.22%, and

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

95.83%, respectively. Additional performance metrics—including precision, recall, and F1-score—computed on the test sets for the optimal configurations are presented in Table 3.

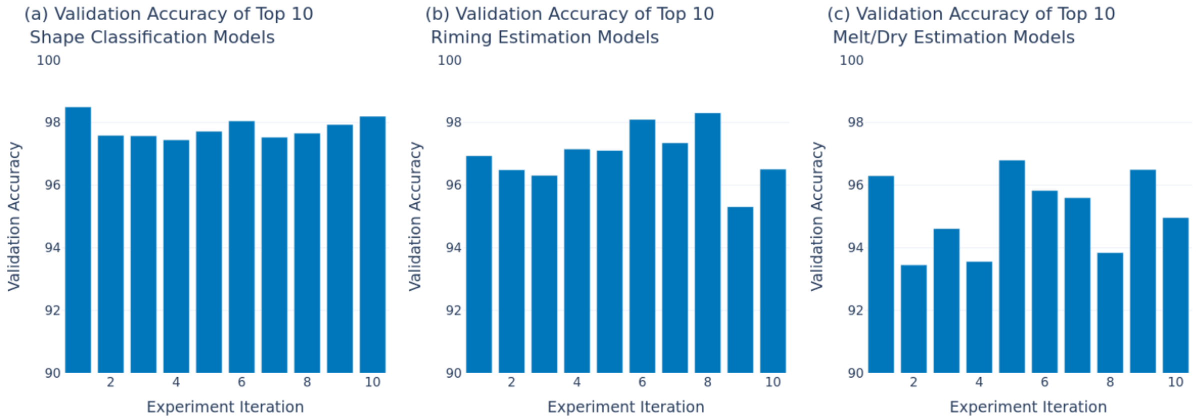

Fig. 11. Phase-two hyperparameter tuning results for geometric shape, riming degree, and melt/dry state models.

Table 2. Best-performing hyperparameter settings for geometric shape, riming degree, and melt/dry state classifications from the two-phase coarse-to-fine hyperparameter tuning.

Classification Type Hyperparameter	Geometric Shape	Riming Degree	Melt/Dry State
Activation Function	relu	relu	swish
Optimizer	rmsprop	adam	sgd
Learning Rate	0.00014	0.00029	0.0306
Batch Size	16	43	53
Fully-connected Layer Config	[256, 16]	[512, 128, 16]	[512, 256, 64, 16]
Dropout Rate	0.13058	0.31095	0.26358

Figure 12 shows accuracy and loss curves for both training and validation, generated across all geometric shapes, riming degrees, and melt/dry states, for the best accuracy found through hyperparameter tuning. The results indicate that the classification performances for the geometric shape and riming estimation are slightly better when compared to melt/dry classification. This can be attributed to the limited visual variability of melt/dry conditions

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

across multiple viewing angles, unlike geometric shape and riming degree, providing fewer discriminative features for multiview melt/dry classification.

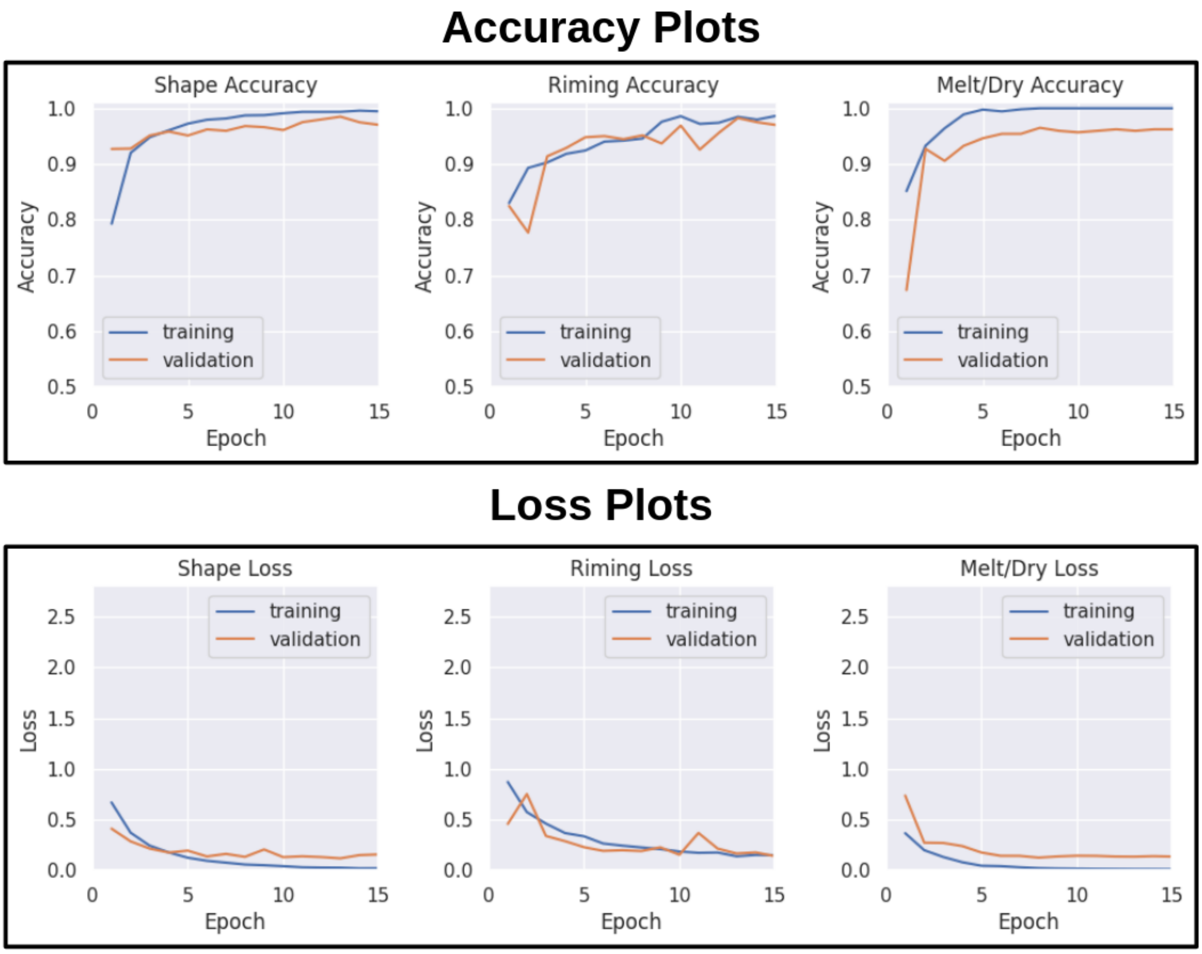

Fig. 12. Accuracy and loss plots for each epoch of training and validation for all categories of snowflake geometric shape, riming estimation, and melt/dry determination.

Additionally, when comparing the confusion matrices obtained with a single-view versus a multiview configuration of Inception-v3 – as shown in Figure 13 – we observe a significant improvement in the correct-class detection, across all classes in all categories, using the multiview approach and the same CNN.

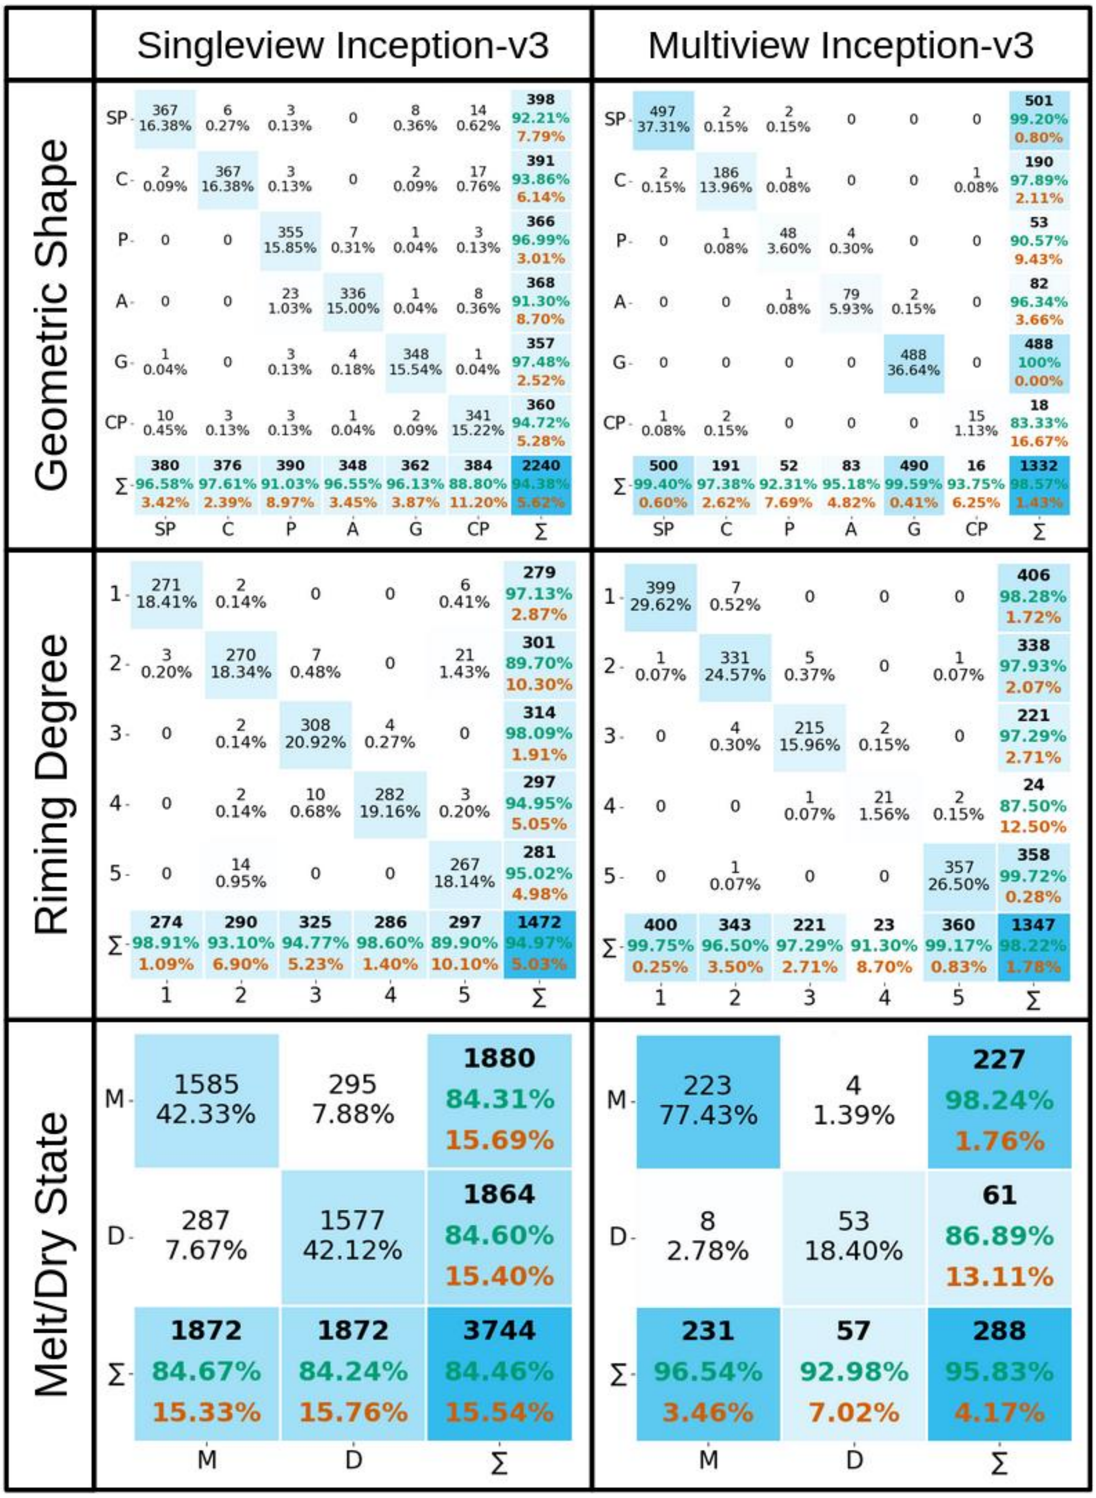

Fig. 13. Confusion matrices, with the predicted snowflake classifications (vertical axis) versus the actual correct snowflake properties (horizontal axis), for geometric shape, riming degree, and melt/dry state classifications: comparison of results obtained by single-view and multiview classifications using the same Inception-v3 CNN technique. In each of the six panels, cells in the bottom row give per-class classification accuracy (green) and error (red), with the last cell showing the overall network accuracy and error. The right-most column of each panel shows out of class accuracy rates, namely, the ability of the network to not confuse an image within a class with something else.

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

Table 3 shows additional performance metrics defined in Section 4e and evaluated on the test set, where a clear improvement from single-view to multiview classification can be observed. Specifically, the multiview model achieved an approximate 4% increase in geometric shape classification, a 3% increase in riming degree estimation, and an 11% increase for melt/dry state determination.

Table 3. Detailed comparison in additional performance metrics between single-view and multiview models for geometric shape, riming degree, and melt/dry state classifications.

Metric	**Single-view**			**Multiview**		
	Shape	**Riming**	**Melt/Dry**	**Shape**	**Riming**	**Melt/Dry**
Accuracy	94.38%	94.97%	84.46%	98.57%	98.22%	95.83%
Precision	0.9443	0.9498	0.8446	0.9455	0.9614	0.9256
Recall	0.9456	0.9506	0.8446	0.9627	0.9680	0.9476
F1-Score	0.9438	0.9498	0.8446	0.9536	0.9646	0.9361
Heidke Skill Score	0.9325	0.9371	0.6891	0.9796	0.9762	0.8721

We then compare results of this work with classification results from prior studies based on MASC images. In addition, we refer to an early study by Feind (2006) based on data collected and extracted using a Particle Measuring System 2-Dimensional Cloud (PMS 2D-C) probe mounted on a T-28 aircraft that traversed directly through storms. This study employed backpropagation neural networks (BPNNs) to classify eight categories of hydrometeors (drops, snow, hail, columns, needles, plates, dendrites, and holes) with an overall accuracy of 85%. However, because the PMS 2D-C records only binary silhouettes (black and white profiles), these images preclude the extraction of detailed microphysical characteristics of the particles. We also refer to a study by Lindqvist et al. (2012) on the Cloud Particle Imager, applying principal component analysis on seven particle types (plate, bullet, column, irregular, rosette aggregate, plate aggregate, and column aggregate) to extract features, which were coupled with Bayesian and weighted nearest-neighbor techniques to classify ice-cloud particles, resulting in overall accuracies between 80% and 90%.

When compared against prior studies, the enhanced CNN architecture and advanced processing techniques enabled multiview classification, which further improved the classification accuracy in all categories. Praz et al. (2017) set the initial standard by adopting and defining the general classes, which were utilized by Hicks and Notaroš (2019) and Key et al. (2021). The same classes were retained in this study, and a significant improvement was shown when using multiview over single-view classification, as well as all earlier works, as outlined in Table 4. It should be noted that the studies cited in Table 4 employed different datasets and data preprocessing pipelines including data preparation, splitting strategy, and instrumentation. Consequently, direct comparisons of classification accuracies are affected by these variations, giving rise to uncertainty in the evaluation of relative performance.

Table 4. Comparison of classification accuracies in all categories (geometric shape, riming degree, and melt/dry state) obtained by the novel multiview CNN snowflake classification methodology against previous published works and results. Also shown is comparison of single-view and multiview Inception-v3 CNN results.

Algorithm/paper; No. of Geometric classes	**Geometric Shape**	**Riming Degree**	**Melt/Dry State**
Feind (2006); 8	85%	-	-
Lindqvist et al. (2012); 8	81.1%	-	-
Praz et al. (2017); 6	94.7%	74.2%	93.9%
Hicks and Notaroš (2019); 5	93.4%	92.4%	-
Key et al. (2021); 5	96.2%	-	-
Single-view Inception-v3; 6	94.38%	94.97%	84.46%
Multiview Inception-v3; 6	**98.57%**	**98.22%**	**95.83%**

Figure 14 depicts representative snowflake triplet examples along with their predicted class labels. For instance, we can observe in the figure an interesting limitation arising from camera view angles and the decision-making capabilities of the multiview classification algorithm, where misclassification occurs between planar and columnar crystals. Specifically, Example Flake 8 was identified as a columnar crystal when observed from two viewpoints, whereas a third angle shows and labels a planar crystal from a very extreme angle. The machine learning model here is incapable of making the decision that this snowflake is a planar crystal. This is a limitation that needs to be addressed in further enhancements of multiview classification accuracy – in future work.

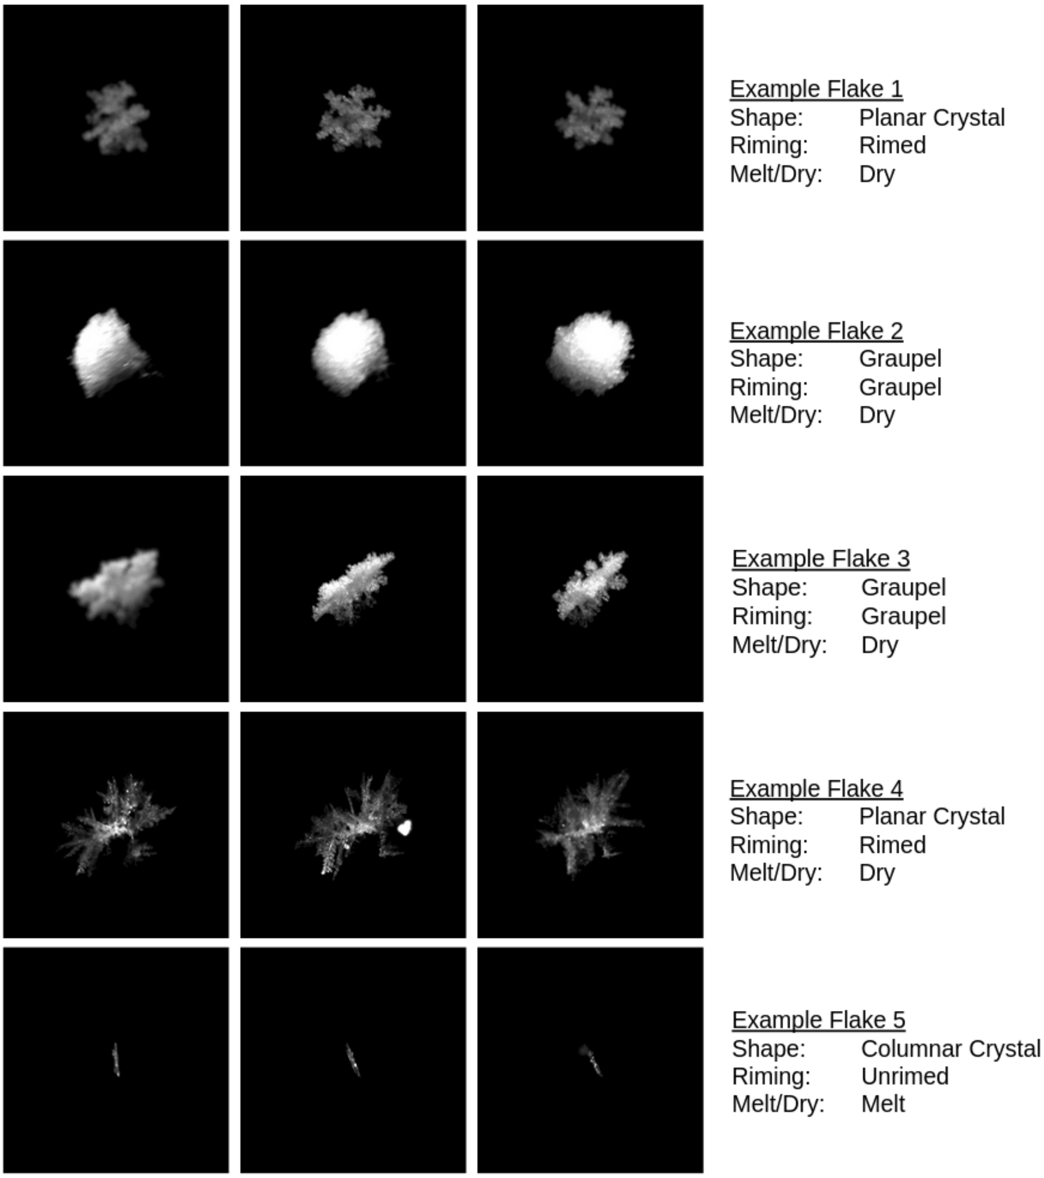

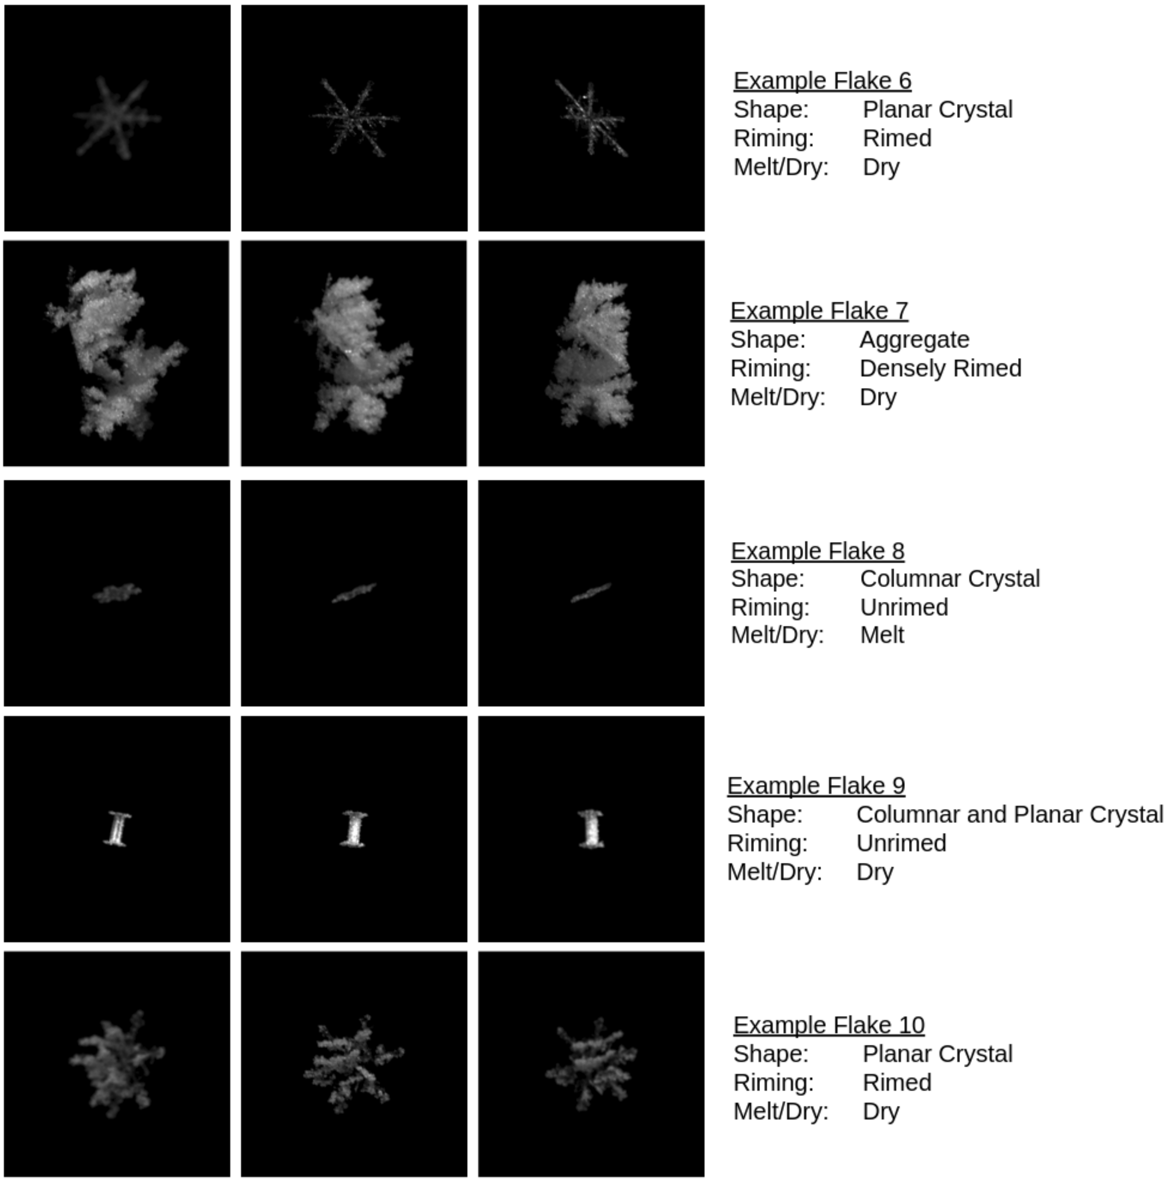

Fig. 14. Example snowflake triplets (images from three MASC views) with respective classification results in geometric shape, riming degree, and melt/dry state.

One constraint of the multiview pipeline stems from the hardware itself, under adverse weather conditions. In winter events with rain mixed-in with snow, image quality is degraded, due to severe glares that overlap with the snowflakes in certain scenarios. Additionally, moisture-induced adhesion of snowflakes to the camera lenses and illumination sources during snowstorms creates obstructions that compromise both preprocessing and classification.

Despite the increase in machine-learning model sizes, inference latency remains low. Benchmarks conducted on a very modest 2017 Intel Core i5 laptop equipped with 20 GB of

RAM yielded an average per-triplet inference time of approximately 1.5 seconds when processing samples individually (no batching). Field-deployment workstations, which typically include much faster CPUs and GPUs, and greater memory bandwidths, can therefore achieve true- or pseudo-real-time inference depending on the available computational resources.

Whereas, because exhaustive hyperparameter optimization entails a substantial computational cost, we employed a two-phase coarse-to-fine hyperparameter tuning strategy, future work will explore more extensive hyperparameter searches with additional computing resources being available. Moreover, applying model-compression techniques (at the expense of a slight reduction in accuracy) can substantially decrease the energy consumption of real-time inference systems, thereby reducing the carbon footprint and alleviating resource constraints in field deployments.

Future research should extend the classification to include additional subclasses of snowflake geometries, alongside the classification of blowing snow (generated by strong winds), where multiview imaging is expected to enhance class discrimination because it provides complementary features from multiple perspectives. Systematic studies should be conducted to determine the minimal number and optimal angles of cameras, required to achieve reliable classification and characterization while simultaneously handling the expanded subclass set and blowing-snow detection. Such investigations will increase the model’s generalizability across different instrumentation platforms. As with all machine-learning pipelines, classification performance is contingent on the quality of the annotated data; therefore, rigorous curation and validation of existing and new datasets are essential to maximize model reliability.

6. Conclusion

This paper has proposed and presented a novel multiview snowflake classification methodology leveraging recent developments in machine learning, multiview classification, and multiangle multicamera instruments for collecting high-resolution photographs of frozen hydrometeors in freefall from multiple views. By employing multiview Convolutional Neural Networks, this work has substantially enhanced automatic classification of snowflakes based on geometric characteristics, riming degree, and melt/dry state, relative to prior methods. The use of CNNs with multiview capabilities has proved to be particularly well-suited for image

Accepted for publication in *Journal of Atmospheric and Oceanic Technology.* DOI 10.1175/JTECH-D-24-0142.1.

classification of snowflakes, due to their efficient data handling, automatic feature extraction, versatility, and retention of dimensionality, coupled with the availability of advanced multiview multicamera snowflake observation instrumentation. To the best of our knowledge, this is the first multiview snowflake classification methodology, taking full advantage of multiview snowflake observation camera systems.

The results for multiview classification have shown record accuracies for snowflake geometric classification, riming degree estimation, and melt/dry state determination, respectively, significantly outperforming other classification models in each of the same categories. In addition, the paper has presented a comparative study of single-view and multiview classification using an Inception-v3 CNN technique. While the single-view method achieved accuracies of 94.4% for six classes of geometry, 95.0% for five degrees of riming, and 84.5% for melt/dry states, the multiview approach made significant improvements – to accuracies of 98.6%, 98.2%, and 95.8%, respectively. Additionally, the multiview pipeline enables higher confidence when manually labeling the data, as one can view the same snowflake from more angles when labeling.

As future work, the presented multiview classification results demonstrate excellent performance ready to attempt processing larger datasets from a variety of snowstorms, including discrimination of falling snowflakes from blowing snow. As more data is processed under diverse environmental conditions, the capability of the network to process and refine more data will be enhanced. With additional snowflake variety, geometric subclasses and rarer geometric classes can be introduced, expanding beyond the six classes currently utilized. The classification network developed in this research has broad applicability, not only for processing other snowstorm data but also as a processing frontend for any other multicamera instrument or system. Its modular design and straightforward architecture make it accessible to those with limited experience in image processing or machine learning.

Overall, automatic multiview CNN-based winter hydrometeor classification has potential to improve understanding, characterization, and quantification of geometrical, microphysical, and scattering properties of ice and snow hydrometeors, and hence, in a longer term, radar-based QPE algorithms, liquid equivalent snow rate retrievals, microphysical parameterizations used in numerical winter-weather forecast models, and numerical models for simulations of ice clouds and frozen precipitation and forecast models overall. This, in turn, can ultimately contribute to more accurate and reliable weather forecasts, hydrological assessment and predictions, and can as well impact regional climate modeling and simulations, and climate projections.

Acknowledgment

This work was supported by the National Science Foundation under grants AGS-2029806 and AGS-2445598.

Data Availability Statement

The triplet images dataset of MASC used in this study is provided in MASCDB database by Grazioli et al. (2023). Packaged machine learning training code for multiview snowflake classification, with the two-phase coarse-to-fine hyperparameter tuning, used in this study has been made publicly available at Thant and Notaroš (2025).

Appendix

This section provides the results for the optimization history, hyperparameter importance, and slice plots of the phase one of the two-phase coarse-to-fine hyperparameter tuning performed with Optuna as outlined in Sections 4d and 5. In particular, Figures A1, A2, and A3 present these results for the shape classification, riming degree estimation, and melt/dry state determination, respectively. All models show that the fully-connected layer configuration is the most important hyperparameter, where different configurations were selected and tested to identify the one yielding the highest validation accuracy for the models.

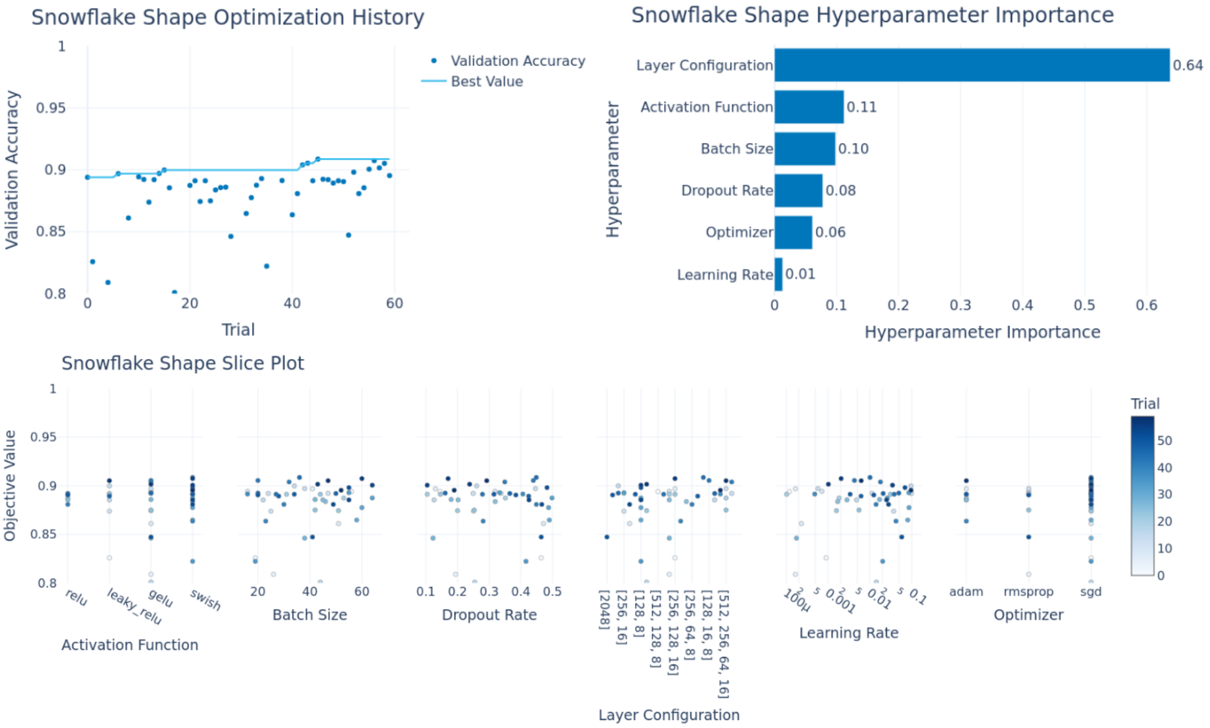

Fig. A1. Bayesian optimization plots for the phase one of the two-phase hyperparameter tuning of the geometric shape classification.

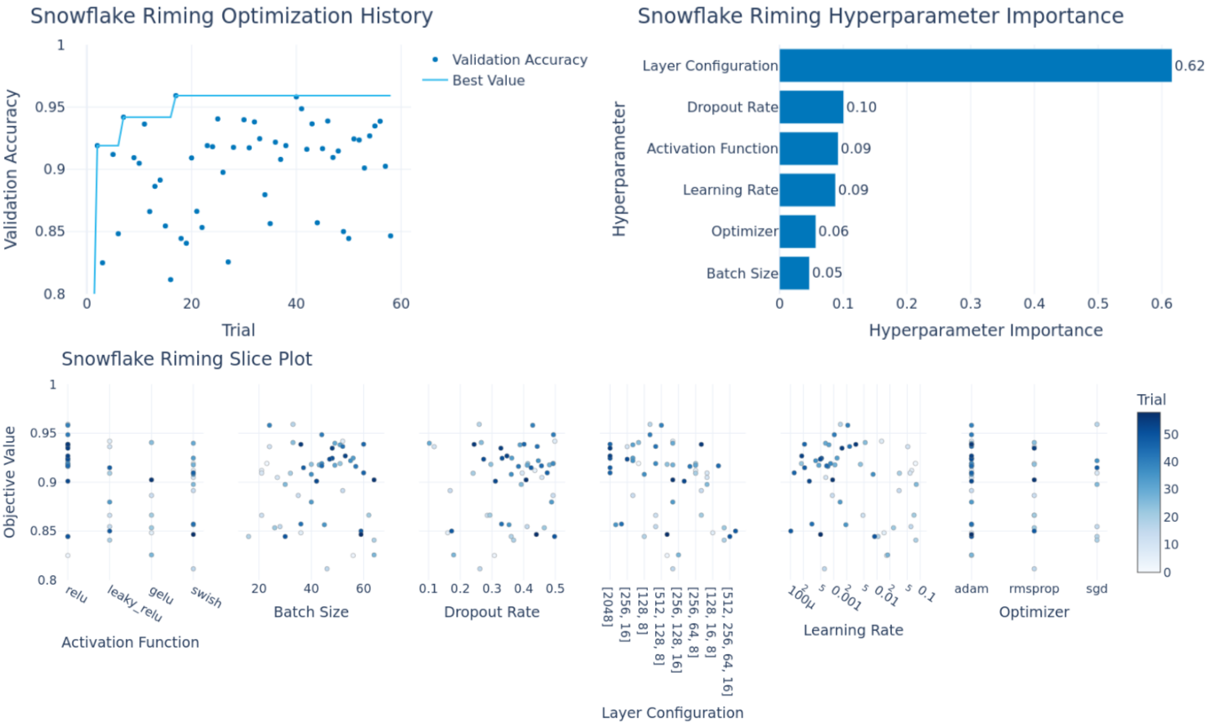

Fig. A2. Bayesian optimization plots for the phase-one hyperparameter tuning of the riming degree classification.

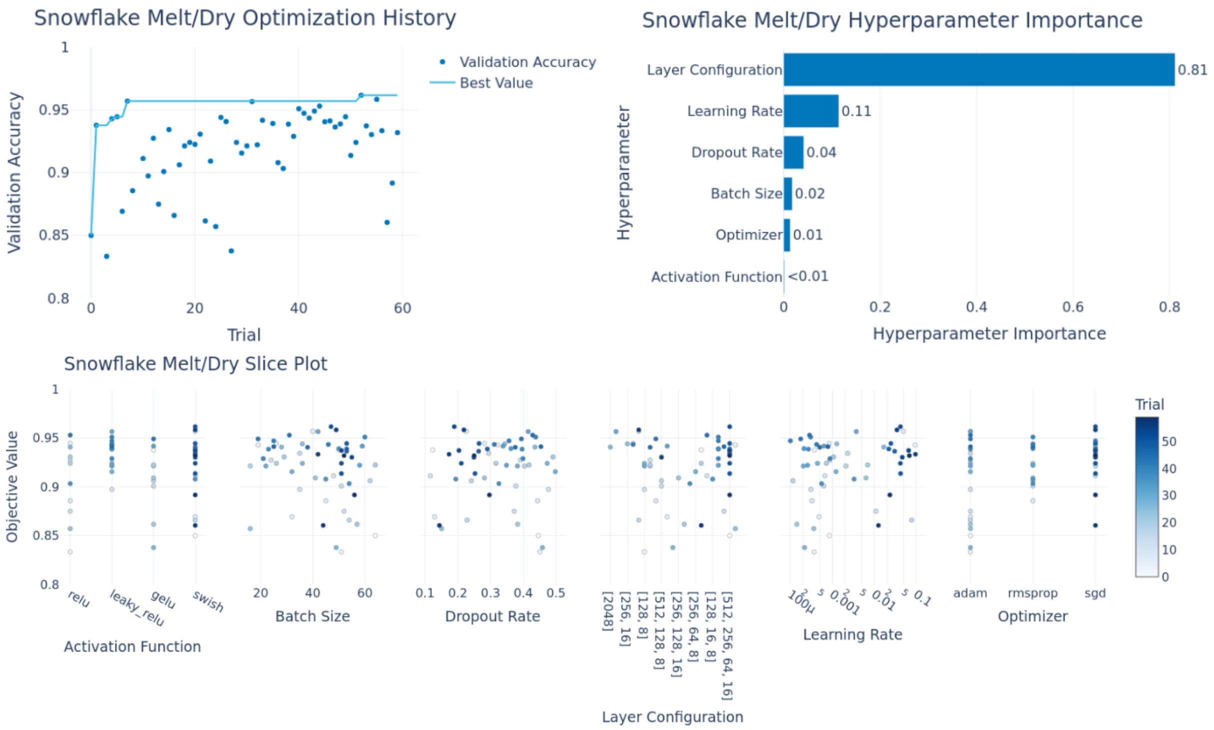

Fig. A3. Bayesian optimization plots for the phase-one hyperparameter tuning of the melt/dry state classification.

References

Akiba, T., S. Sano, T. Yanase, T. Ohta, and M. Koyama, 2019: Optuna: A next-generation hyperparameter optimization framework. *In Proceedings of the 25th ACM SIGKDD internation conference on knowledge discovery and data mining.* Anchorage, Alaska, Knowledge Discovery and Data Mining, 2623-2631, https://doi.org/10.1145/3292500.3330701.

Atcheson, B., F. Heide, and W. Heidrich, 2010: CALTag: High Precision Fiducial Markers for Camera Calibration. *Vision, Modeling and Visualiztion*, **10**, 41-48 https://doi.org/10.2312/PE/VMV/VMV10/041-048.

Barbosa, A., T. Marinho, N. Martin, and N. Hovakimyan, 2020: Multi-Stream CNN for Spatial Resource Allocation: A Crop Management Application. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, IEEE/CVF, 58-59.

Bergstra, J., R. Bardenet, Y. Bengio, and B. Kegl, 2011: Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, **24**.

Billault-Roux, A. C., and Coauthors, 2023: ICE GENESIS: Synergetic aircraft and ground-based remote sensing and in situ measurements of snowfall microphysical properties. *Bull. Amer. Meteor. Soc.*, **104**, 367-388, https://doi.org/10.1175/BAMS-D-21-0184.1

Brital, A., 2021: Inception V3 CNN Architecture Explained. *Medium*, Accessed August 5, 2024, https://medium.com/@AnasBrital98/**inception**-v3-cnn-architecture-explained-691cfb7bba08.

Brunkow, D., V. N. Bringi, P. C. Kennedy, S. A. Rutledge, V. Chandrasekar, E. A. Mueller, and R. K. Bowie, 2000: A description of the CSU-CHILL National Radar Facility. *J. Atmos. Oceanic Technol.*, **17**, 1596-1608, https://doi.org/10.1175/1520-0426(2000)017<1596:ADOTCC>2.0.CO;2

Bukovčić, P., A. Ryzhkov, D. Zrnić, and G. Zhang, 2018: Polarimetric Radar Relations for Quantification of Snow Based on Disdrometer Data. J. Appl. Meteor. Climatol., 57, 103–120, https://doi.org/10.1175/JAMC-D-17-0090.1.

Chase, R. J., S. W. Nesbitt, G. M. McFarquhar, N. B. Wood, and G. M. Heymsfield, 2022: Direct comparisons between GPM-DPR and CloudSat snowfall retrievals, *J. Appl. Meteorol. Climatol.*, **61**, 1257-1271, https://doi.org/10.1175/JAMC-D-21-0081.1

Collins, R., 2012: Lecture 20: The Eight-Point Algorithm, Accessed August 4, 2024, https://www.cse.psu.edu/~rtc12/CSE486/lecture20.pdf.

Cooper, S. J., N. B. Wood, and T. S. L'Ecuyer, 2017: A variational technique to estimate snowfall rate from coincident radar, snowflake, and fall-speed observations, *Atmos. Meas. Tech.*, **10**, 2557-2571, https://doi.org/10.5194/amt-10-2557-2017

Do, T. B., H. H. Nguyen, H. Vu, T. T. H. Tran, and T.-L. Le, 2017: Plant identification using score-based fusion of multi-organ images. *2017 9th International Conference on Knowledge and Systems Engineering (KSE)*, Hue, Vietnam, IEEE, 191–196, https://doi.org/10.1109/KSE.2017.8119457.

Dolata, P., M. Mrzygłód, and J. Reiner, 2017: Double-stream Convolutional Neural Networks for Machine Vision Inspection of Natural Products. *Applied Artificial Intelligence,* **31**, 643–659, https://doi.org/10.1080/08839514.2018.1428491.

Donahue, J., Y. Jia, O. Vinyals, N. Zhang, E. Tzeng, and T. Darrell, 2014: Decaf: A deep convolutional activation feature for generic visual recognition. *31st International* *conference on machine learning*, Beijing, China, Proceedings of Machine Learning Research, 647-655.

Dozier, J., E. H. Bair, and R. E. Davis, 2016: Estimating the spatial distribution of snow water equivalent in the world's mountains, *WIREs Water*, **3**, 461-474, https://doi.org/10.1002/wat2.1140

Feichtenhofer, C., A. Pinz, and A. Zisserman, 2016: Convolutional Two-Stream Network Fusion for Video Action Recognition. *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, IEEE/CVF, 1933-1941, https://doi.org/10.48550/arXiv.1604.06573.

Feind, R., 2006: Comparison of Three Classification Methodologies for 2D Probe Hydrometeor Images Obtained from the Armored T-28 Aircraft. Report SDSMT/IAS/R08-01, 61 pp. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=37a995890f06448a9fa9b3c108402ba7589841e5.

Feng, Y., Z. Zhang, X. Zhao, R. Ji, and Y. Gao, 2018: GVCNN: Group-View Convolutional Neural Networks for 3D Shape Recognition. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE/CVF, 264–272.

Ferrone, A., and A. Berne, 2022: Radar and ground-level measurements of clouds and precipitation collected during the POPE 2020 campaign at Princess Elisabeth Antarctica, *Earth System Science Data Discussions*, **15**, 1115-1132, https://doi.org/10.5194/essd-15-1115-2023.

Filipiak, B. C., and Coauthors, 2025: Winter Precipitation Measurements in New England: Results from the Global Precipitation Measurement Ground Validation Campaign in Connecticut. Preprint, Earth System Science Data, **2025**, 1-45, https://doi.org/10.5194/essd-2025-162.

Garrett, T. J., C. Fallgatter, K. Shkurko, and D. Howlett, 2012: Fall speed measurement and high-resolution multi-angle photography of hydrometeors in free fall. *Atmospheric Measurement Techniques Discussions*, **5**(4), 4827-4850, https://doi.org/10.5194/amt-5-2625-2012.

Geras, K. J., and Coauthors, 2017: High-Resolution Breast Cancer Screening with Multi-View Deep Convolutional Neural Networks. *arXiv preprint*, https://doi.org/10.48550/arXiv.1703.07047.

Grazioli, J., D. Tula, and A. Berne, 2015: Hydrometeor classification from polarimetric radar measurements: a clustering approach, *Atmospheric Measurement Techniques*, **8**, 149-170, https://doi.org/10.5194/amt-8-149-2015

Grazioli, J., G. Ghiggi, A.-C. Billault-Roux, and A. Berne, 2022: MASCDB, a database of images, descriptors and microphysical properties of individual snowflakes in free fall. *Sci Data*, **9**, 186, https://doi.org/10.1038/s41597-022-01269-7.

Grazioli, J., G. Ghiggi, and A. Berne, 2023: MASCDB, a database of images, descriptors and microphysical properties of individual snowflakes in free fall, Version 1.1, Zenodo, Accessed 10 June 2024, https://doi.org/10.5281/zenodo.8083133.

Hartley, R., and A. Zisserman, 2004: *Multiple View Geometry in Computer Vision Second Edition*. Cambridge University Press, 670 pp.

Heiss, W. H., D. L. McGrew, and D. Sirmans, 1990: NEXRAD: Next generation weather radar (WSR-88D). *Microwave Journal*, **33**, 79-89.

Heymsfield, A. J., S. Y. Matrosov, and N. B. Wood, 2016: Toward Improving Ice Water Content and Snow-Rate Retrievals from Radars. Part I: X and W Bands, Emphasizing CloudSat. *J. Appl. Meteor. Climatol.*, **55**, 2063-2090, https://doi.org/10.1175/JAMC-D-15-0290.1.

Hicks, A., and B. M. Notaroš, 2019: Method for Classification of Snowflakes Based on Images by a Multi-Angle Snowflake Camera Using Convolutional Neural Networks. *J. Appl. Oceanic Technol.*, **36**, 12, 2267-2282, https://doi.org/10.1175/JTECH-D-19-0055.1.

Huang, G.-J., V. N. Bringi, D. Moisseev, W. A. Petersen, L. Bliven, and D. Hudak, 2015: Use of 2D-video disdrometer to derive mean density–size and Ze–SR relations: Four snow cases from the light precipitation validation experiment. *Atmospheric Research*, **153**, 34–48, https://doi.org/10.1016/j.atmosres.2014.07.013.

Huang, G.-J., C. Kleinkort, V. N. Bringi, and B. M. Notaroš, 2017: Winter precipitation particle size distribution measurement by Multi-Angle Snowflake Camera. *Atmospheric Research*, **198**, 81–96, https://doi.org/10.1016/j.atmosres.2017.08.005.

Kennedy, P., M. Thurai, C. Praz, V. N. Bringi, A. Berne, and B. M. Notaroš, 2018: Variations in Snow Crystal Riming and ZDR: A Case Analysis. *J. Appl. Meteor. Climatol.*, **57**, 695-707, https://doi.org/10.1175/JAMC-D-17-0068.1.

Key, C., A. Hicks, and B. M. Notaroš, 2021: Advanced Deep Learning–Based Supervised Classification of Multi-Angle Snowflake Camera Images. *J. Atmos. Oceanic Technol.*, **38**, 1399-1414, https://doi.org/10.1175/JTECH-D-20-0189.1.

Kleinkort, C., G.-J. Huang, V. N. Bringi, and B. M. Notaroš, 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**, 679-702, https://doi.org/10.1175/JTECH-D-16-0099.1.

Lee, S. H., C. S. Chan, and P. Remagnino, 2018: Multi-Organ Plant Classification Based on Convolutional and Recurrent Neural Networks. *IEEE Transactions on Image Processing*, **27**, 4287–4301, https://doi.org/10.1109/TIP.2018.2836321.

Leinonen, J., and A. Berne, 2020: 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.

Leinonen, J., D. Moisseev, and T. Nousiainen, 2013: Linking snowflake microstructure to multi-frequency radar observations. *J. Geophys. Res: Atmos.*, **118**, 3259–3270, https://doi.org/10.1002/jgrd.50163.

Li, Z., F. Liu, W. Yang, S. Peng, and J. Zhou, 2021: A survey of convolutional neural networks: analysis, applications, and prospects. *IEEE transactions on neural networks and learning systems*, **33**, 6999-7019, https://doi.org/10.1109/TNNLS.2021.3084827.

Liao, L., and R. Meneghini, 2022: GPM DPR retrievals: Algorithm, evaluation, and validation, *Remote Sensing*, **14**,843, https://doi.org/10.3390/rs14040843

Lin, C., and A. Kumar, 2018: Contactless and partial 3D fingerprint recognition using multi-view deep representation. *Pattern Recognition*, **83**, 314–327, https://doi.org/10.1016/j.patcog.2018.05.004.

Lindqvist, H., K. Muinonen, T. Nousiainen, J. Um, G. M. McFarquhar, P. Haapanala, R. Makkonen, and H. Hakkarainen, 2012: Ice-cloud particle habit classification using principal components. *J. Geophys. Res.: Atmos.*, **117**, 16, https://doi.org/10.1029/2012JD017573.

Ludlam, F. H., and B. J. Mason, 1957: The Physics of Clouds. *Geophysik II / Geophysics II*, J. Bartels, Springer Berlin Heidelberg, 479–540.

Maahn, M., D. Moisseev, I. Steinke, N. Maherndl, and M. D. Shupe, 2023: Introducing the Video In Situ Snowfall Sensor (VISSS). *Atmos. Meas. Tech.*, **17**, 899-919, https://doi.org/10.5194/amt-17-899-2024.

McMurdie, L. A., G. Heymsfield, J. E. Yorks, S. A. Braun, 2019: Investigation of Microphysics and Precipitation for Atlantic Coast-Threatening Snowstorms (IMPACTS) Collection. Accessed 13 July 2025, https://doi.org/10.5067/IMPACTS/DATA101.

Newman, A. J., P. A. Kucera, and L. F. Bliven, 2009: Presenting the Snowflake Video Imager (SVI). *J. Atmos. Oceanic Technol.*, **26**, 167-179, https://doi.org/10.1175/2008JTECHA1148.1.

Notaros, B. M., 2021: Meteorological Electromagnetics: Optical and Radar Measurements, Modeling, and Characterization of Snowflakes and Snow. *IEEE Antennas and Propagation Magazine*, **63**, 14–27, https://doi.org/10.1109/MAP.2021.3054298.

Notaros, B. M., 2022: Polarimetric Weather Radar: Overview of principles and applications. *IEEE Antennas Propagation Magazine*, **64**, 43–54, https://doi.org/10.1109/MAP.2022.3143442.

Notaroš, B. M., and Coauthors, 2016: Accurate Characterization of Winter Precipitation Using Multi-Angle Snowflake Camera, Visual Hull, Advanced Scattering Methods and Polarimetric Radar. *Atmosphere*, **7**, 81, https://doi.org/10.3390/atmos7060081.

Notaros, B. M., V. N. Bringi, H. Thant, G.-J. Huang, and D. B. Wolff, 2022: Measurement and Characterization of Winter Precipitation at Wallops Island Snow Field Site. *2022 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting (AP-S/URSI),* Denver, Colorado, IEEE, 519–520, https://doi.org/10.1109/AP-S/USNC-URSI47032.2022.9886077.

Park, H. S., A. V. Ryzhkov, D. S. Zrnić, and K. Kim, 2009: The Hydrometeor Classification Algorithm for the Polarimetric WSR-88D: Description and Application to an MCS. Wea. Forecasting, 24, 730–748, https://doi.org/10.1175/2008WAF2222205.1.

Pettersen, C., L. F. Bliven, A. von Lerber, N. B. Wood, M. S. Kulie, M. E. Mateling, and D. B. Wolff, 202: The precipitation imaging package: Assessment of microphysical and bulk characteristics of snow. *Atmosphere*, **11**, 785, https://doi.org/10.3390/atmos11080785.

Pettersen, C., L. F. Bliven, M. S. Kulie, N. B. Wood, J. A. Shates, J. Anderson, M. E. Mateling, W. A. Petersen, A. von Lerber, and D. B. Wolff, 2020: The precipitation imaging package: Phase partitioning capabilities, *Remote Sensing*, **13**, 2183, https://doi.org/10.3390/rs13112183.

Praz, C., Y.-A. Roulet, and A. Berne, 2017: Solid hydrometeor classification and riming degree estimation from pictures collected with a Multi-Angle Snowflake Camera. *Atmospheric Measurement Techniques*, **10**, 1335–1357, https://doi.org/10.5194/amt-10-1335-2017.

Pruppacher, H. R., and J. D. Klett, 2010: Microstructure of Atmospheric Clouds and Precipitation. *Microphysics of Clouds and Precipitation*, H.R. Pruppacher and J.D. Klett, Eds., Springer Netherlands, 10–73.

Roberto, N., L. Baldini, E. Adirosi, L. Facheris, F. Cuccoli, A. Lupidi, and A. Garzelli, 2017: A Support Vector Machine Hydrometeor Classification Algorithm for Dual-Polarization Radar. *Atmosphere,* **8**, 134, https://doi.org/10.330/atmos8080134

Ryzhkov, A. V., and D. S. Zrnic, 2019: *Radar Polarimetry for Weather Observations*. Springer International Publishing, 497 pp.

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

Schönhuber, M., G. Lammer, and W. L. Randeu, 2008: The 2D-Video-Distrometer. *Precipitation: Advances in Measurement, Estimation and Prediction*, S. Michaelides, Springer Berlin Heidelberg, 3–31.

Seeland, M., and P. Mäder, 2021: Multi-view classification with convolutional neural networks. *Plos one*, **16**, e0245230, https://doi.org/10.1371/journal.pone.0245230.

Setio, A. A. A., and Coauthors, 2016: Pulmonary Nodule Detection in CT Images: False Positive Reduction Using Multi-View Convolutional Networks. *IEEE Transactions on Medical Imaging,* **35**, 1160–1169, https://doi.org/10.1109/TMI.2016.2536809.

Straka, J. M., D. S. Zrnić, and A. V. Ryzhkov, 2000: Bulk Hydrometeor Classification and Quantification Using Polarimetric Radar Data: Synthesis of Relations. *J. Appl. Meteor.*, **39**, 1341-1372, https://doi.org/10.1175/1520-0450(2000)039%3C1341:BHCAQU%3E2.0.CO;2.

Stuefer, M., and J. Bailey, 2016: Multi-Angle Snowflake Camera Instrument Handbook. DOE/SC-ARM-TR-158j, 26 pp, https://doi.org/10.2172/1261185.

Su, H., S. Maji, E. Kalogerakis, and E. Learned-Miller, 2015: Multi-view Convolutional Neural Networks for 3D Shape Recognition. *Proceedings of the IEEE International Conference on Computer Vision*, IEEE, 945–953.

Szegedy, C., V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, 2015: Rethinking the Inception Architecture for Computer Vision. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2818-2826.

Szyrmer, W., and I. Zawadzki, 2014: Snow Studies. Part IV: Ensemble Retrieval of Snow Microphysics from Dual-Wavelength Vertically Pointing Radars. *J. Atmos. Sci.*, **71**, 1171-1186, https://doi.org/10.1175/JAS-D-12-0286.1.

Thant, H., G.-J. Huang, V. N. Bringi, and B. M. Notaros, 2022: Snowflake Measurement and Analysis System (SMAS). *2022 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting (AP-S/URSI)*, Denver, Colorado, IEEE, 523–524, https://doi.org/10.1109/AP-S/USNC-URSI47032.2022.9886325.

Thant, H., D. B. Wolff, V. N. Bringi, and B. M. Notaros, 2023a: Wallops Flight Facility 2021-2022 Winter Snow Events: A Case Study Using Surface Instrumentation and Radars. *2023 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio Science Meeting (USNC-URSI),* Portland, Oregon, IEEE, 1707–1708, https://doi.org/10.1109/USNC-URSI52151.2023.10237857.

Thant, H., M. Zhizhin, and B. M. Notaros, 2023b: Machine Learning Classification of Snowflakes to Enhance Microphysical and Scattering Characterization of Snow. *2023 IEEE International Symposium on Antennas and Propagation and USNC-URSI Radio*

Science Meeting (USNC-URSI), Portland, Oregon, IEEE, 283–284, https://doi.org/10.1109/USNC-URSI52151.2023.10237821.

Thant, H., V. N. Bringi, and B. M. Notaros, 2024: Winter Event Observations at Wallops Flight Facility in 2022 and Ongoing Analyses of Collected Data. *2024 United States National Committee of URSI National Radio Science Meeting (USNC-URSI NRSM)*, Boulder, Colorado, IEEE, 355–355, https://doi.org/10.23919/USNC-URSINRSM60317.2024.10464690.

Thant, H., and B. M. Notaroš, 2025: Novel Multiview Snowflake Classification Using Convolutional Neural Networks and Multiangle Multicamera Instruments: Machine Learning Training Code, with Two-Phase Coarse-to-Fine Hyperparameter Tuning (Version 1.0). GitHub. https://github.com/hthant/snowflake_multiview.

Tsang, S.-H., 2019: Review: Inception-v3 — 1st Runner Up (Image Classification) in ILSVRC 2015. Medium, Accessed 5 August 2024, https://sh-tsang.medium.com/review-inception-v3-1st-runner-up-image-classification-in-ilsvrc-2015-17915421f77c.

Vu, C. T., and D. M. Chandler, 2009: S3: A Spectral and Spatial Sharpness Measure. *2009 First International Conference on Advances in Multimedia*, Colmar, France, IEEE, 37–43, https://doi.org/10.1109/MMEDIA.2009.15.

Wang, A., J. Cai, J. Lu, and T.-J. Cham, 2015: MMSS: Multi-Modal Sharable and Specific Feature Learning for RGB-D Object Recognition. *Proceedings of the IEEE International Conference on Computer Vision*, IEEE, 1125–1133.

Wolff, D. B., D. A. Marks, C. S. Pabla, J. L. Pippitt, A. Tokay, J. Wang, and M. Watson, 2023: The NASA Polarimetric (NPOL) weather radar facility and some applications, *Advances in Weather Radar. Volume 1: Precipitation sensing platforms*, V. N. Bringi, K. V. Mishra, and M. Thurai, The Institution of Engineering and Technology.

Zhang, G., S. Luchs, A. Ryzhkov, M. Xue, L. Ryzhkova, and Q. Cao, 2011: Winter Precipitation Microphysics Characterized by Polarimetric Radar and Video Disdrometer Observations in Central Oklahoma. *J. Appl. Meteor. Climatol.*, **50**, 1558-1570, https://doi.org/10.1175/2011JAMC2343.1.

Zhang, Z., 2000: A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **22**, 1330–1334, https://doi.org/10.1109/34.888718.