

Enhancing weather radar data by removing non-meteorological echoes, using neural networks trained on synthetic weather data

Richard Bölz^{a,1,*}, Tom Kirstein^{a,1}, Lukas Fuchs^a, Lukas Josipović^b,
Annette Böhm^c, Ulrich Blahak^b, Volker Schmidt^a

^a*Institute of Stochastics, Ulm University, Ulm, Germany*

^b*Deutscher Wetterdienst, Offenbach, Germany*

^c*Deutscher Wetterdienst, Meteorologisches Observatorium Hohenpeißenberg, Hohenpeißenberg, Germany*

Abstract

Meteorological weather radars are essential for atmospheric research, weather forecasting and aviation safety, but they often detect non-meteorological echoes from scatterers such as insects, birds, and ground clutter. These non-meteorological echoes can then lead to misinterpretations in quantitative precipitation estimation and hydrometeor classification, which cause difficulties for atmospheric research and weather forecasting. This paper introduces a novel AI-based approach to identify such non-meteorological echoes in polarimetric radar data using a convolutional neural network. More specifically, we utilize a so-called U-net, which relies on large amounts of labeled radar data for training. To address the challenge of acquiring labeled radar data consisting of meteorological and non-meteorological echoes, we generate synthetic training samples by combining preprocessed winter data (meteorological echoes) with cluttered summer data (non-meteorological echoes) provided by Deutscher Wetterdienst (DWD). After training on synthetic data, evaluation of the U-net approach on operationally measured radar data shows that it outperforms the state-of-the-art DWD classification algorithm overall. This is particularly evident in the preservation of precipitation signals at the boundaries of larger weather events.

Keywords and Phrases: machine learning, convolutional neural network, synthetic data generation, meteorological radar data, polarimetric radar, non-meteorological echo detection

¹RB and TK contributed equally.

*Corresponding author. E-mail address: richard.boelz@uni-ulm.de

1 Introduction

Since the late 1940s, meteorological radars have become indispensable tools in atmospheric and hydrological research and operational applications, e.g., research on clouds and precipitation formation processes, quantitative precipitation estimation (QPE), motion analysis and precipitation nowcasting, detection and tracking of convective storms, and radar data assimilation for numerical weather prediction (NWP), see [9] for a broad overview. These meteorological weather radars are ground-based remote sensing instruments that scan a conical volume originating from the radar site with high spatial and temporal resolution. Although this capability makes them highly suitable for atmospheric research, operational weather surveillance and forecasting, the data suffers from several sources of error. One particular problem is that radars not only detect meteorological particles, such as droplets, snowflakes, and hail, but also any other scatterers, i.e., objects in the conical volume that scatter electromagnetic waves. This can lead to misinterpretation of the measurements during subsequent analyzes. To address this problem, this paper proposes a novel AI-based method to identify signals corresponding to non-meteorological scatterers based on their spatial and statistical properties measured by meteorological radar products.

To explain the problem in more detail, it is useful to briefly summarize the basic measurement principle, as described in many textbooks; see, e.g., [6, 30]. A short and highly directed pulse of electromagnetic waves is sent out via an antenna in a certain direction, traveling at the speed of light in this direction along a so-called ray. During a subsequent listening period, a time series of returned signals, so-called echoes, from atmospheric scatterers is detected at the radar site. These echoes can be used to compute so-called radar moments, such as strength, phase, and frequency of the detected echoes, which contain information on the amount and nature of scatterers, as well as their radial movement towards or away from the radar site. The time elapsed since sending is a measure of the scatterers radial distance, or range, from the radar site. Thus, the sampling frequency of the time series defines the range resolution (“bin”) and the listening time defines the maximum measurement range. For typical operational applications, the radar antenna is rotated at a constant angular velocity and records data along ray after ray in changing directions. The data from one full azimuthal rotation of the antenna is also referred to as a sweep. If the antenna’s elevation angle is held constant for a sweep, it is called a plan position indicator (PPI) scan. Several successive PPI scans at different elevation angles are called volume scans, as they may sample large parts of the conical volume whose radius is given by the maximum range. Typical spatial resolutions of PPI scans are 250 m to 1 km along the range and 1° in azimuth. Typical volume scans consist of 5 to 20 PPI scans at different elevations, all of which are measured within five minutes.

Modern dual-polarization radars are capable of alternating the polarization plane of the electromagnetic radiation from pulse to pulse, commonly from horizontal to vertical and vice versa [4, 29]. Most precipitation particles are asymmetric (not spherical) and preferably fall down with their longer axis in horizontal orientation, which causes differences in the returned signal strength and phase between the alternating pulses. These differences and their correlations in time, which depend on the hydrometeor shapes and spatial orientations, led to the definition of other radar moments in addition to signal strength. These additional radar moments can be used to improve the echo interpretation with respect to the type and amount of scatterers [35], as will be further outlined in Section 2.1

A problem for correct interpretation is the fact that the type of scatterers is not known a priori, and non-meteorological scatterers like insects or turbulent fluctuations of the air’s refractive index (see[6, 34]) might be misinterpreted as, e.g., precipitation. Besides non-meteorological

77 echoes (which shall be the focus of the present paper) and attenuation, there are also many
78 other systematic and statistical error sources to be taken into account, e.g., beam bending,
79 beam blockage at obstacles, side lobe echoes, ground clutter, statistical signal fluctuations and
80 radar miscalibration [6, 9].

81 Biological scatterers, such as insects and birds, can provide valuable information to meteo-
82 rological forecasters, and to ornithologists and entomologists. For example, so-called "fine lines"
83 are sometimes visible in the radar data when thunderstorms expel cooler, denser air. Small air-
84 borne particles, such as dust and insects, are swept along the leading edge of this cooler air
85 mass and, together with echoes caused by turbulent fluctuations of the air's refractive index (see
86 below), contribute to the fine line (e.g., [34]), giving the forecaster valuable information about
87 the characteristics of the thunderstorm. Radar data can also be used to track bird migrations
88 [15] and monitor biodiversity [2]. However, for meteorological and hydrological applications such
89 as QPE [5, 29], precipitation nowcasting [26, 21] or hydrometeor classification [3, 33], it is essen-
90 tial to separate hydrometeor echoes from any other type of echo. This becomes more and more
91 important because many processing steps are being automated, and human intervention based
92 on individual experience is reduced.

93 Echoes from insects and turbulence in the boundary layer are predominantly observed during
94 the summer. The turbulence echoes are commonly attributed to Bragg scattering, which occurs
95 due to refractive index fluctuations at a length scale of half the radar wavelength [6, 34]. The
96 air's refractive index depends on pressure, temperature and humidity, where mainly fluctuations
97 in moisture and, to a lesser extent, temperature cause echoes. These echoes, together with those
98 from insects, are centered around the radar's location because they are limited to lower altitudes
99 within the planetary boundary layer (PBL) and the altitude of the radar ray increases with range.
100 Hence, they appear as annular stationary structures of weak but non-negligible echoes. Radar
101 moments prove to be very useful in distinguishing clutter from meteorological echoes because
102 they provide additional information on the shape and uniformity of the scatterers. However,
103 detecting and eliminating insect and turbulence echoes is still a challenge.

104 Conventional classification methods based on pre-defined thresholds of radar moments [18],
105 Bayesian classifiers [22], or fuzzy logic [7, 27] are difficult to configure such that they success-
106 fully filter out all echoes from insects in the boundary layer while preserving highly turbulent
107 regions, the edges of true precipitation regions and weak precipitation signals. Especially chal-
108 lenging situations include (i) cases with snow echoes or regions in the melting layer, because
109 they can exhibit polarimetric signatures similar to those of insects, (ii) regions far away from the
110 radar sites, where the signal-to-noise ratio is low, and (iii) "mixed" situations, e.g., summertime
111 convection surrounded by cloud-free regions with insects in the boundary layer. The classifiers
112 mentioned above are also limited by the fact that they can incorporate information only in close
113 proximity to the location of the bin to classify. This means that the classifier loses the overall
114 picture of the annual structure of the insects.

115 Artificial intelligence (AI) methods have been used in recent studies to detect and classify
116 echoes from biological scatterers in radar data for the purposes of biological studies. The focus
117 of these studies was on obtaining information about insects or birds and their subclasses, rather
118 than retaining all meteorologically useful data. For example, in [16], a ridge classifier was applied
119 to polarimetric radar data to distinguish between birds and insects, with meteorological signals
120 filtered out before classification. Moreover, random forest approaches based on polarimetric radar
121 data were used to discriminate between different biological scatterers [11, 15]. Besides this, in [1],
122 the problem of noisy or uncertain training labels was addressed using a deep learning approach
123 to improve the classification of biological scatterers in radar data.

124 Unlike the above-mentioned conventional approaches that rely on predefined thresholds,
125 Bayesian classifiers, or fuzzy logic, AI-methods can learn intricate patterns directly from data
126 without explicit feature engineering. This adaptability makes them particularly promising for
127 handling the complex and high-dimensional nature of radar data. Convolutional neural networks
128 (CNNs) specifically excel at extracting spatial features from image-like data structures, such as
129 radar sweeps, enabling tasks like object detection and segmentation, see [12]. Among CNNs, the
130 U-net introduced in [28] is particularly well-suited for segmentation tasks due to its encoder-
131 decoder architecture that captures both local features and global context via skip connections.
132 For further applications of CNNs in radar-based traffic and biomedical monitoring, see [17, 20].
133 In this context, the sweeps can be viewed as multi-channel images, where each pixel corresponds
134 to a radar gate and each image channel corresponds to a radar moment.

135 However, applying U-nets generally requires a large and diverse dataset of accurately labeled
136 training samples. Labeling radar data is especially challenging, as the labeling of even an in-
137 dividual range gate requires expert consideration of multiple-radar moments and their spatial
138 arrangement in multiple surrounding range gates. Given that each sweep consists of thousands
139 of radar gates and training requires hundreds or even thousands of labeled sweeps, relying on
140 manually labeled training data is infeasible. In the present paper, although a small number
141 of manually labeled radar sweeps was utilized for model validation and benchmarking, train-
142 ing was instead conducted on synthetically generated radar sweeps with corresponding ground
143 truth. More specifically, we first identified sweeps that feature either no meteorological echoes
144 or no non-meteorological echoes. Sweeps that have no meteorological echoes were identified by
145 considering time periods when no precipitation was measured. To identify sweeps featuring a
146 minimum amount of non-meteorological echoes, a time period was manually selected through ex-
147 pert evaluation. These initial sweeps were then combined to generate "mixed" synthetic training
148 sweeps featuring both meteorological and non-meteorological echoes with known ground truth.
149 In order to further enhance training data diversity and avoid overfitting, these synthetic sweeps
150 were augmented via scaling, rotation, and orientation inversion during training. Thus, the goal
151 of this paper is to develop and test a new method to identify non-meteorological echoes in sweeps
152 using a U-net and suitable radar moments as input.

153 The remainder of this paper is organized as follows. Section 2 describes the acquisition of radar
154 data and its preprocessing, as well as the architecture and training of the neural network that
155 is used for the detection of non-meteorological echoes. Section 3 describes the results obtained
156 by the presented method applied to a test dataset not used for training. Using manually labeled
157 radar sweeps, the network's results are compared to those obtained by the current state-of-the-art
158 classification method at DWD. The results are discussed in Section 4. Section 5 concludes.

159 2 Materials and methods

160 We now explain how the polarimetric radar data is acquired and represented, and introduce the
161 polarimetric radar moments in more detail, see Section 2.1. Section 2.2 describes the preprocess-
162 ing steps applied to the data. The preprocessed data is then used to generate the input for the
163 neural network, the architecture of which is described in Section 2.3. Section 2.4 explains the
164 synthetic training data generation. The network training itself is described in Section 2.5.

2.1 Acquisition of radar data

The radar data used in this paper is provided by DWD, which operates a radar network of 17 polarimetric weather radars, covering the entire landmass of Germany. Each of these radars sends out horizontally and vertically polarized electromagnetic waves through an antenna and records the power and phase of the returned signal [29]. The recorded signals are acquired in scans, each containing data from one 360° revolution of the antenna. These scans are referred to as sweeps, with the antenna first pointing north and then rotating clockwise. More precisely, every five minutes, each of the radars performs a volume scan consisting of 10 sweeps acquired at different elevation angles, as well as a terrain-following sweep [13], which measures precipitation as closely to the ground as possible. In the following, we utilize only these terrain-following sweeps. If the radar view is unobstructed, the elevation angle at which the electromagnetic waves are sent out for such sweeps is set to 0.8° . In cases where orography or obstacles obstruct the view, the elevation is raised accordingly.

The radars scan at a range resolution of 250 m and an azimuthal resolution of 1° , with a maximum range of 150 km, resulting in 360 rays of 600 radar gates each. Each of these radar gates is sampled by 50 pulses of electromagnetic waves, which serve as input to compute various radar moments for this radar gate. An important step after collecting the radar measurements is quality control of the radar data, which provides a solid data base for subsequent quantitative precipitation estimation, hydrometeor classification and thunderstorm detection. One main goal of this process is to identify and subsequently filter out, or correct for, non-meteorological echoes, that is, the signals returned by non-meteorological scatterers. Typical non-meteorological scatterers include ground clutter (trees, mountains), biological scatterers (birds, insects), wind turbines, towers, ships, airplanes, external signals (e.g. wifi), and chaff. In contrast, meteorological scatterers are any hydrometeors, e.g., rain, snow, hail, or drizzle. The signals returned by meteorological scatterers are called meteorological echoes.

The various radar moments are particularly useful for this purpose of identifying and filtering out non-meteorological echoes, as they provide insight into different qualities of the detected scatterers. The most widely used radar moment is the radar reflectivity factor [33, 11], which quantifies the strength of the returned signal. It is measured separately at horizontal and vertical polarization, yielding the horizontal and vertical reflectivity factors, denoted by DBZH and DBZV, respectively. Due to variations in raindrop size and number concentration, the radar reflectivity factor can vary by several orders of magnitude. Hence, it is usually expressed in logarithmic units (dBZ), where the minimum detectable value of the reflectivity factor increases with range. For example, at a distance of 1 km from the radar, the minimum detectable reflectivity factor is around -38 dBZ, whereas at ranges beyond 25 km, signals below -10 dBZ cannot be detected by DWD radars [10], i.e., they cannot be distinguished from noise.

The logarithmic ratio between the horizontal and vertical reflectivity factors (in linear scale) is called differential reflectivity, denoted by ZDR [32]. It contains implicit information about the shape of the scatterers and is commonly used to, e.g., discriminate between liquid and frozen hydrometeors. In combination with other radar moments, it is also useful for identifying clutter.

Finally, the depolarization ratio, denoted by UDR, is deduced from the differential reflectivity ZDR and the measurements of the cross-correlation coefficient (ρ_{HV}) [24]. This radar moment is a measure of how homogeneous and spherical the scatterers are within a range gate, and thus serves as a good discriminator between clutter and meteorological echoes [18, 25], with small values ($\ll 0$) indicating weather and values closer to 0 indicating clutter or biological signals. The DWD's data contains the radar reflectivity factors, also called radar reflectivities, for horizontal and vertical polarization, the differential reflectivity, and the depolarization ratio, among various

212 other radar moments. The reflectivity factors DBZH and DBZV, and the differential reflectivity
 213 ZDR have been filtered by a Doppler filter [6] aiming to remove any stationary non-meteorological
 214 scatterers, whereas no such filter has been applied to the depolarization ratio UDR. Figure 1
 215 shows examples of the radar moments DBZH, ZDR and UDR obtained from different sweeps.

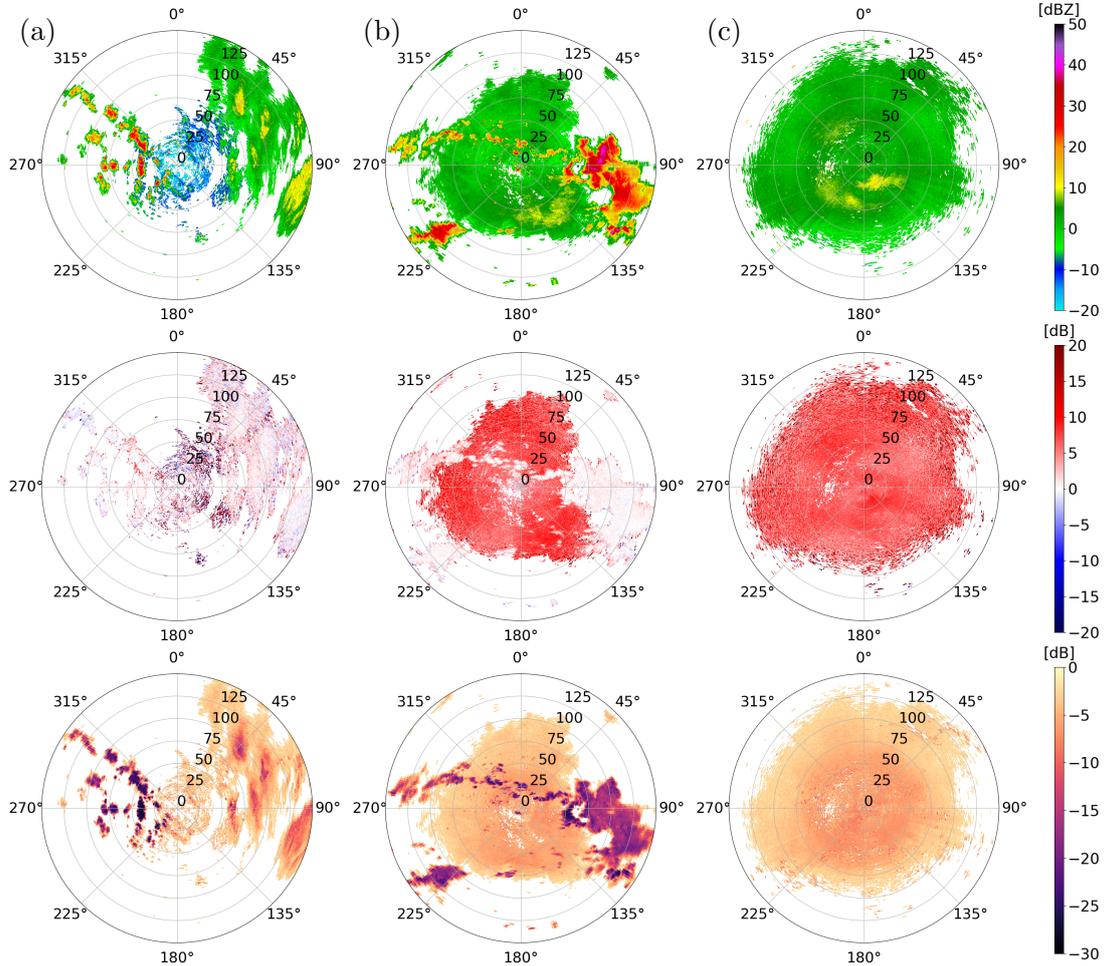


Figure 1: Horizontal reflectivity (first row), differential reflectivity (second row), and depolarization ratio (third row) for examples of (a) a winter sweep, (b) a mixed summer sweep, and (c) a cluttered summer sweep. White pixels are filtered by the initial preprocessing done by DWD.

216 The DWD provided all terrain-following sweeps measured at the 17 different polarimetric
 217 weather radars for three one-hour time periods: 09:00–10:00 UTC on June 16, 2022, 10:00–11:00
 218 UTC on June 18, 2022, and 17:00–18:00 UTC on January 9, 2025. Radar sweeps collected in
 219 January 2025 are referred to as winter sweeps, and radar sweeps collected in June 2022 are
 220 referred to as summer sweeps. Moreover, the radar sweeps collected on June 16 feature both
 221 meteorological and non-meteorological echoes and are referred to as mixed summer sweeps.
 222 In contrast, the radar sweeps collected on June 18 were collected during a time period when
 223 no precipitation was measured, indicating that only non-meteorological echoes were detected.
 224 Therefore, these sweeps are referred to as cluttered summer sweeps. For the winter sweeps, the
 225 specific time frame was chosen with the goal of selecting sweeps featuring a minimum amount
 226 of non-meteorological echoes. However, even for the winter sweeps, a certain amount of non-
 227 meteorological echoes cannot be avoided. Figure 1 shows an overview of the horizontal reflectivity
 228 (DBZH), differential reflectivity (ZDR), and depolarization ratio (UDR) for each of the three
 229 data categories.

2.2 Preprocessing of radar data

For each radar station, a single sweep consists of the measured values of the radar moments at the radar gates covered by the radar. These radar gates partition the scanned area into a polar grid defined by the radar's number of rays $n_a \in \mathbb{N} = \{1, 2, \dots\}$ and the number of range bins per ray $n_r \in \mathbb{N}$. More specifically, the radar gate corresponding to the i -th range bin and j -th ray is defined as the set

$$g_{i,j} = \left\{ wr \begin{pmatrix} \sin \theta \\ \cos \theta \end{pmatrix} : r \in (i-1, i], \theta \in \left(\frac{2\pi}{n_a}(j-1), \frac{2\pi}{n_a}j \right] \right\} \subset \mathbb{R}^2, \quad (1)$$

where $w = 250$ denotes the range resolution of the radar in meters.

Let \mathcal{M} denote the set of radar moments available for a given sweep. Then, for each radar moment $m \in \mathcal{M}$, its measured values from the sweep can be represented by a single-channel image $A^m = (A_{i,j}^m) \in \mathbb{R}^{n_r \times n_a}$. For each pixel $(i, j) \in \Omega = \{1, \dots, n_r\} \times \{1, \dots, n_a\}$, the entry $A_{i,j}^m$ corresponds to the value of the radar moment m at the associated radar gate $g_{i,j}$. Single-channel images can be combined into a multi-channel image $A = (A^m)_{m \in \mathcal{M}'} \in \mathbb{R}^{C \times n_r \times n_a}$, where $\mathcal{M}' \subseteq \mathcal{M}$ is the subset of included radar moments and $C = |\mathcal{M}'|$ is the number of channels. In the following, we refer to these multi-channel images as radar images. Moreover, the provided data also contains information on which radar gates had already been identified as non-meteorological during the preprocessing performed by DWD. For these radar gates, the corresponding entries of the DBZH image are assigned a placeholder value $p = -40$ dBZ, which lies far below the reflectivity range of meteorological echoes. Pixels with a horizontal reflectivity value of -40 dBZ are referred to as excluded pixels, as they conclusively do not correspond to meteorological echoes. All other pixels are referred to as included pixels. A Cartesian representation of a DBZH image is shown in Figure 2a. Most of the radar images still contain pixels that likely correspond to non-meteorological echoes, even for radar images corresponding to winter sweeps. To filter some of them out, we apply further preprocessing steps.

For example, as depicted in Figure 2a, the DBZH images show weak signals with horizontal reflectivity values below -10 dBZ near the radar. However, only very light drizzle could produce reflectivity values of that magnitude. That is, thresholding at -10 dBZ most likely removes noise and clutter, and does not remove any convectively significant weather. Therefore, pixels with lower horizontal reflectivity values are considered to correspond to non-meteorological echoes and are thus excluded as potential meteorological echoes, see Figure 2b. Their DBZH values are then also set to -40 dBZ to mark them as excluded pixels.

Moreover, meteorological echoes observed in the radar images typically form connected components consisting of several pixels, see Figure 1. In contrast, isolated included pixels are assumed to correspond to non-meteorological echoes. For this reason, each DBZH image is further processed by a despeckle filter consisting of two steps. First, we use the Hoshen-Kopelman algorithm [14] to identify all connected components of included pixels. More specifically, we apply the algorithm to the binarized image obtained by thresholding at -40 dBZ, separating the excluded pixels from the included ones. Here, we use the 4-neighborhood, which means that two pixels $(i_1, j_1), (i_2, j_2) \in \Omega$ with $(i_1, j_1) \neq (i_2, j_2)$ are neighboring if their associated radar gates $g_{i_1, j_1}, g_{i_2, j_2}$ share a common edge. Note that the DBZH image is periodic in the second component, see Eq. (1), since each index $j \in \{1, \dots, n_a\}$ represents the azimuthal angle of a ray, and the image corresponds to a sweep. Consequently, pixels corresponding to the first and last ray are neighboring if they have the same range index $i \in \{1, \dots, n_r\}$. More precisely, for each pixel $(i, j) \in \Omega$, the set $N_{(i,j)} \subset \Omega$ of neighboring pixels can be written as $N_{(i,j)} = \{(k, \ell) \in \Omega : |i - k| + d(j, \ell) = 1\}$,

273 where $d: \{1, \dots, n_a\} \times \{1, \dots, n_a\} \rightarrow \mathbb{N}_0$ denotes the periodic distance measure, given by

$$d(j, \ell) = \min\{|j - \ell|, n_a - |j - \ell|\}. \quad (2)$$

274 for each $\ell \in \{1, \dots, n_a\}$. After identifying the connected components, those consisting of less
 275 than six undetermined pixels are selected. Then, the DBZH values of pixels $(i, j) \in \Omega$ in these
 276 components are set to -40 dBZ, indicating that they do not correspond to meteorological echoes.
 277 The result of the despeckling process is shown in Figure 2c.

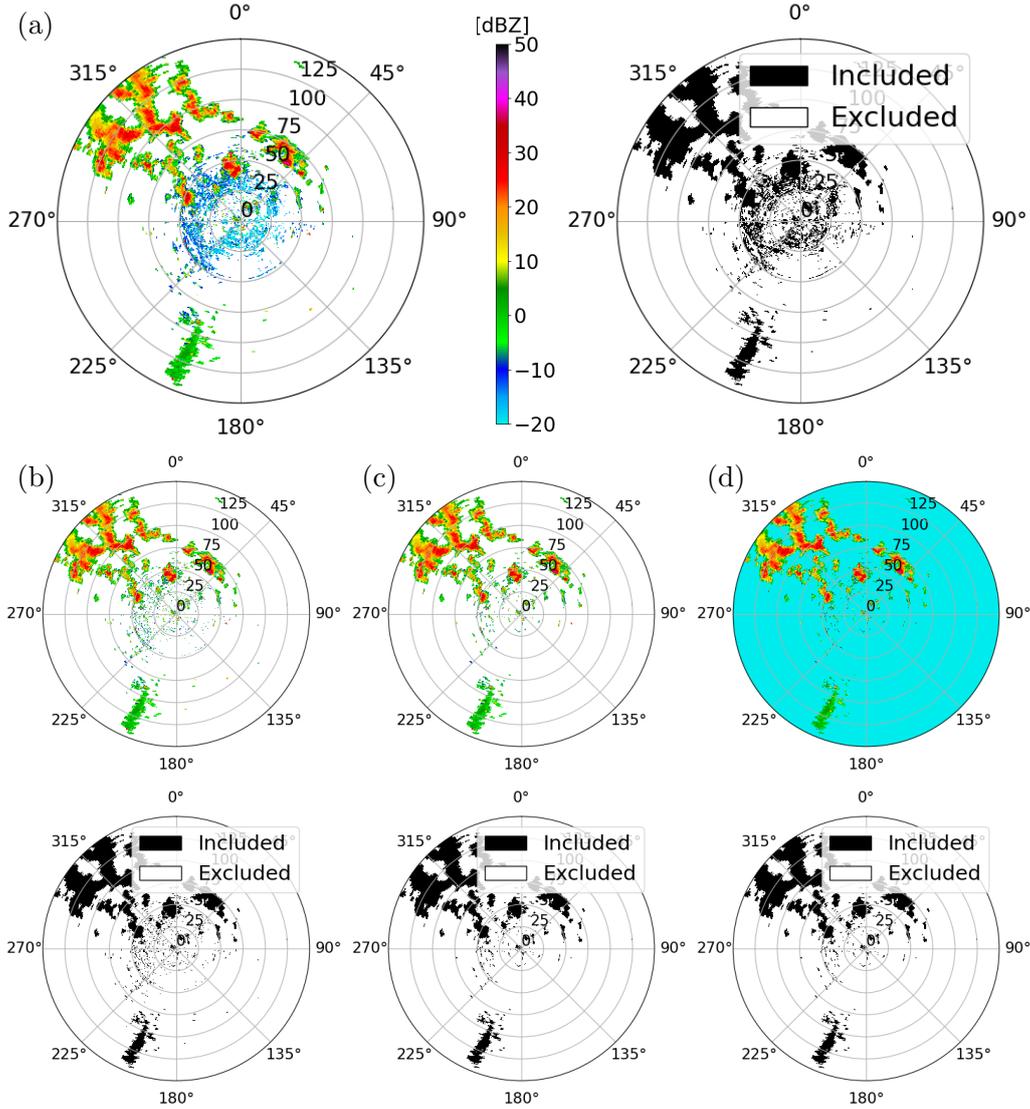


Figure 2: Cartesian representations of DBZH images: (a) DBZH image of a winter sweep, as provided by DWD, (b) after thresholding, (c) after despeckling, and (d) final DBZH image with placeholder values. White pixels in the DBZH images are either filtered out during the preprocessing by DWD or are assigned a placeholder value.

278 For the excluded pixels $(i, j) \in \Omega$, their corresponding entries in the DBZH image were set
 279 to -40 dBZ to indicate that they conclusively do not correspond to meteorological echoes, see
 280 Figure 2d. For the other radar moments, ZDR and UDR, we also assign specific placeholder
 281 values $p \in \mathbb{R}$ to the entries of their single-channel images. Similar to the horizontal reflectivity,
 282 these placeholder values are chosen to lie outside the range of typical values for meteorological
 283 echoes in our dataset. More specifically, for the differential reflectivity, the placeholder value is set

284 to $p = 20$ dB. This particular value was chosen since it lies above typical values for meteorological
 285 echoes (around 0 dB), but within the range for non-meteorological echoes (between 5 and 20 dB)
 286 in our data, as can be seen in Figures 1b and 1c. Furthermore, as described in [18], except for
 287 some special cases such as hail, the depolarization ratio of meteorological scatterers is typically
 288 below -12 dB. In contrast, the depolarization ratio of non-meteorological scatterers typically
 289 ranges between -12 and 0 dB. Therefore, we chose a placeholder value of $p = 0$ dB for the
 290 depolarization ratio.

291 The presented preprocessing steps remove most of the clutter with low reflectivity values.
 292 However, many mixed summer images contain much larger connected components corresponding
 293 to more reflective non-meteorological scatterers, as can be seen in Figure 1b. Detecting such
 294 echoes presents a challenging task for conventional methods such as fuzzy logic filtering [7] or
 295 depolarization ratio thresholding [18].

296 2.3 Description of the network architecture

297 The task of detecting non-meteorological echoes can be viewed as a segmentation task of multi-
 298 channel two-dimensional (2D) radar image data. Each pixel of the image is classified as belonging
 299 to meteorological or non-meteorological echoes, that is, each pixel is assigned a binary class label
 300 $c \in \{0, 1\}$. In [28], the so-called U-net architecture, a fully convolutional neural network, has
 301 been introduced to address similar segmentation tasks of 2D images.

302 A U-net has an encoder-decoder architecture with skip-connections, consisting of a so-called
 303 contracting path that downsamples the spatial resolution of the features and an expansive path
 304 that upsamples it again. Both the downsampling and the upsampling paths consist of several
 305 levels. The network considered in the present paper is an adapted version of the U-net architec-
 306 ture comprising five levels, with the fifth level acting as a bottleneck. Consequently, the network
 307 performs four downsampling steps and four upsampling steps. Each downsampling step consists
 308 of two convolution blocks and a 2×2 max-pooling layer for downsampling the spatial resolu-
 309 tion of the feature map [12]. A convolution block contains a padded 3×3 convolutional layer,
 310 followed by a batch normalization layer and a rectified linear unit (ReLU) activation function,
 311 which is given by $\text{ReLU}: \mathbb{R} \rightarrow \mathbb{R}_0^+ = [0, \infty)$, $\text{ReLU}(x) = \max(x, 0)$. The bottleneck at the fifth
 312 layer consists of two such convolution blocks. For each of the subsequent upsampling steps, the
 313 spatial resolution of the feature channels is increased using bilinear interpolation [8]. The re-
 314 sulting channels are then concatenated with the feature channels of the corresponding level of
 315 the downsampling path. Two convolutional blocks following the concatenation of those feature
 316 channels complete the upsampling step. After the final upsampling step, the sigmoid activation
 317 function $\sigma: \mathbb{R} \rightarrow [0, 1]$, given by

$$\sigma(x) = \frac{1}{1 + e^{-x}},$$

318 is applied to the output.

319 Compared to the original U-net architecture presented in [28], we use padded convolutional
 320 layers with a padding of 1 in each step, preserving the image size throughout the network.
 321 Since the network input is periodic in the second component, we use periodic padding [31] in
 322 this direction, while the padding in the direction of the first component is reflective. Moreover,
 323 the spatial resolution of the features is upsampled using bilinear interpolation instead of an
 324 up-convolution. Furthermore, the network does not include a final 1×1 convolution. Instead,
 325 the output of the final convolution block is already a single-channel map to which the sigmoid
 326 activation function is applied directly. A sketch of the network architecture described above is
 327 shown in Figure 3.

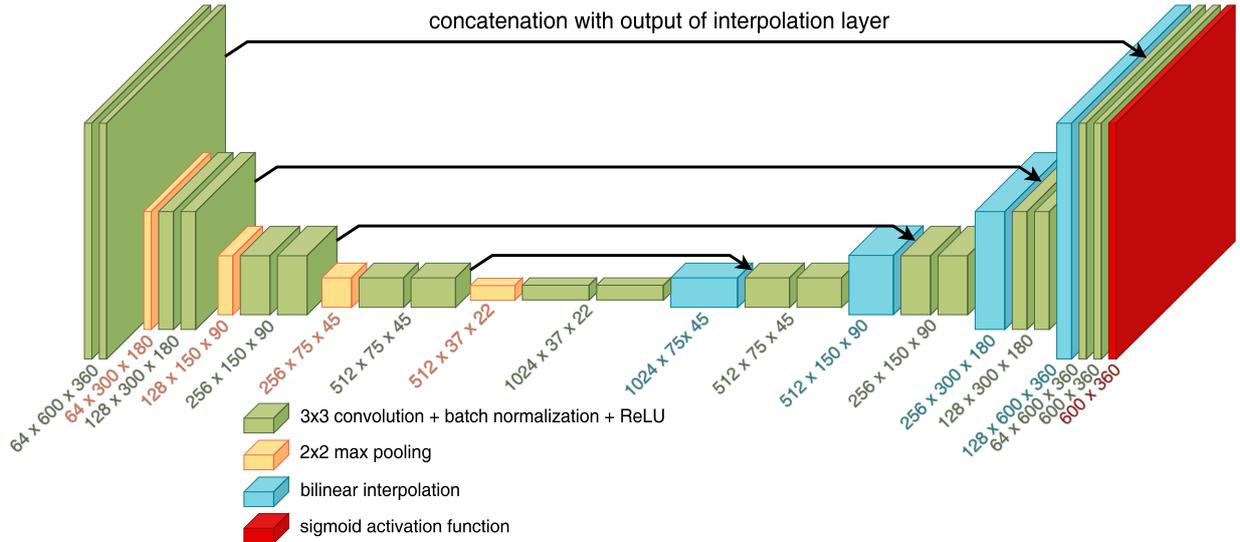


Figure 3: Schematic representation of the network architecture. Below each layer, the respective output dimensions are given.

328 The network outputs are images of the same spatial dimensions as each of the input channels.
 329 More specifically, the network architecture stated above with parameter constellation $\theta \in \Theta \subset \mathbb{R}^q$
 330 is given by a function $f_\theta: \mathbb{R}^{k \times h \times w} \rightarrow [0, 1]^{h \times w}$, with spatial dimensions $h, w \in \mathbb{N}$ and the number
 331 of input channels $k = 3$. Here, θ denotes the vector of trainable parameters of the network, and
 332 $\Theta \subset \mathbb{R}^q$ is the network's parameter space for some $q \in \mathbb{N}$. Each input channel corresponds to
 333 a single-channel image of a radar moment, as described in Section 2.1, allowing the network to
 334 combine information about different radar moments to detect non-meteorological echoes.

335 The goal is then to perform supervised training of the network to assign binary labels to each
 336 pixel $(i, j) \in \Omega$, describing which pixels correspond to meteorological and non-meteorological
 337 echoes, respectively. This training process requires pixel-wise labeled training data, that is, each
 338 pixel $(i, j) \in \Omega$ must be assigned a ground truth class label $c_{i,j} \in \{0, 1\}$. However, mixed radar
 339 images containing both meteorological and non-meteorological echoes lack pixel-wise ground
 340 truth labels, and generating such labels by hand is very costly. Therefore, we propose a method
 341 to generate synthetic multi-channel images and corresponding pixel-wise ground truth labels for
 342 training the network.

343 2.4 Generation of synthetic training data

344 As outlined in Section 2.3, the network architecture considered in this paper can utilize radar
 345 images as input, where each channel corresponds to a single-channel image of a specific radar mo-
 346 ment. The training samples consist of such radar images and corresponding binary ground truth
 347 images indicating whether each pixel corresponds to a meteorological or a non-meteorological
 348 echo. We generate synthetic radar images by combining radar images corresponding to winter
 349 sweeps (featuring meteorological echoes) and radar images corresponding to cluttered summer
 350 sweeps (featuring non-meteorological echoes). For brevity, we refer to these as winter radar
 351 images and cluttered summer radar images, respectively. This approach is motivated by the
 352 observation that the horizontal reflectivity (DBZH) values in mixed summer radar images often
 353 have the appearance of a combination of (scaled) DBZH values from winter and of cluttered
 354 summer radar images, as illustrated in Figure 1b. In this way, we can generate labeled training
 355 samples where the contributions of meteorological and non-meteorological echoes are explicitly

356 known, addressing the challenge of difficult labeling in measured mixed radar images.

357 More specifically, we generate a realistic mixed (synthetic) DBZH image by combining two
 358 DBZH images $A^{H,w}, A^{H,s} \in \mathbb{R}^{n_r \times n_a}$ from a winter and a cluttered summer radar image in the
 359 following manner. Since the horizontal reflectivity is measured in decibels, which is a logarithmic
 360 unit of measurement, we first transform the entries of the DBZH images into linear scale using
 361 the transformation $\Phi: \mathbb{R} \rightarrow \mathbb{R}^+ = (0, \infty)$ [29], which is given by

$$\Phi(x) = 10^{x/10}.$$

362 The transformed reflectivity values are scaled by suitable weighting factors $\lambda_w > 0$ and $\lambda_s > 0$
 363 for the DBZH images from the winter and the cluttered summer radar image, respectively. For
 364 each pixel, the scaled reflectivity values are summed up and transformed back into decibels to
 365 obtain the synthetic DBZH image $A^{H,m} \in \mathbb{R}^{n_r \times n_a}$ with entries

$$A_{i,j}^{H,m} = \begin{cases} \Phi^{-1}(\lambda_w \Phi(A_{i,j}^{H,w}) + \lambda_s \Phi(A_{i,j}^{H,s})), & \text{if } A_{i,j}^{H,w} > -40 \text{ dBZ or } A_{i,j}^{H,s} > -40 \text{ dBZ,} \\ -40, & \text{otherwise,} \end{cases}$$

366 for each $(i, j) \in \Omega$.

367 Similarly to the placeholder values described in Section 2.2, we assign the entry of the syn-
 368 thetic DBZH image $A_{i,j}^{H,m}$ a placeholder value $p = -40$ dBZ if the pixel is excluded for both
 369 radar images. To find appropriate weighting factors λ_w and λ_s , we consider the difference in
 370 horizontal reflectivity values observed in Figure 1. Since the range of the horizontal reflectiv-
 371 ity values corresponding to non-meteorological echoes is almost identical for cluttered summer
 372 radar images (Figure 1c) and mixed summer radar images (Figure 1b), the choice of $\lambda_s = 1$
 373 is suitable. However, the horizontal reflectivity values corresponding to meteorological echoes
 374 are significantly lower for winter radar images (Figure 1a) compared to mixed summer radar
 375 images. Since the resulting synthetic DBZH image should resemble the typical DBZH images
 376 from mixed summer radar images, a scaling factor $\lambda_w > 1$ is necessary. In order to determine a
 377 suitable choice, four representative mixed summer radar images were chosen. For each of them,
 378 a winter radar image with precipitation areas of similar size was assigned. Then, for different
 379 weighting factors $\lambda_w \in \{1.5, 3, 5, 8, 15\}$, we compared the resulting synthetic DBZH image with
 380 that of the respective mixed summer radar image. As a result, a weighting factor λ_w between
 381 5 and 8 was deemed appropriate. Figure 4 shows how a synthetic DBZH image emerges from
 382 a winter and a cluttered summer radar image using a weighting factor $\lambda_w = 6.5$. The resulting
 383 synthetic DBZH image (Figure 4c) is then compared with one from a mixed summer radar image
 384 (Figure 4d).

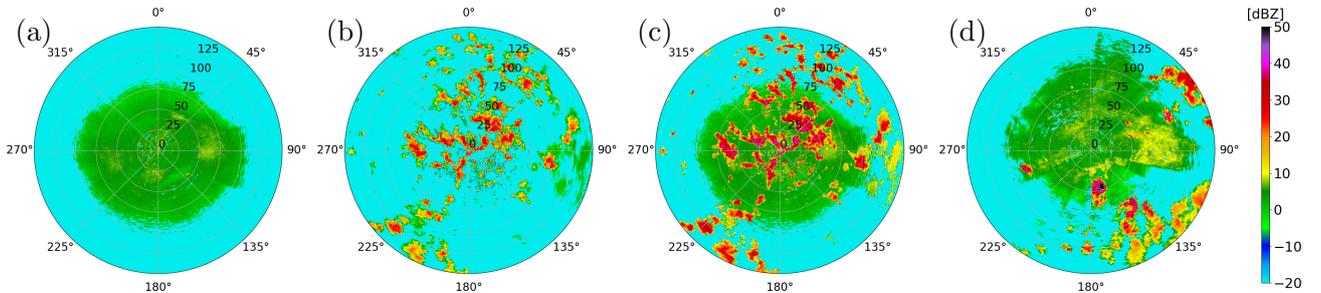


Figure 4: Comparison of Cartesian representations of different DBZH images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image with weighting factors $\lambda_w = 6.5$, and $\lambda_s = 1$ and (d) mixed summer radar image.

385 The two chosen winter and cluttered summer radar images are then also used to generate
 386 synthetic single-channel images for the other radar moments considered in Section 2.1. In ad-
 387 dition to the horizontal reflectivity, which is proportional to the particle concentration and the
 388 particle's horizontal cross-section within the radar gate, the differential reflectivity (ZDR) pri-
 389 marily reflects the mean axis ratio of the particles [33]. Since differential reflectivity is defined
 390 as the difference between horizontal and vertical reflectivity in decibels, synthetic values can be
 391 generated by taking the difference between synthetic horizontal and vertical reflectivity values.
 392 Similarly to the synthetic DBZH image $A^{H,m}$, a synthetic DBZV image $A^{V,m} \in \mathbb{R}^{n_r \times n_a}$ is also
 393 generated. Then, the corresponding synthetic ZDR image is denoted by $A^{ZDR,m} \in \mathbb{R}^{n_r \times n_a}$, and
 394 its entries are given by

$$A_{i,j}^{ZDR,m} = \begin{cases} A_{i,j}^{H,m} - A_{i,j}^{V,m}, & \text{if } A_{i,j}^{H,m} > -40 \text{ and } A_{i,j}^{V,m} > -40, \\ 20, & \text{if } A_{i,j}^{H,m} = -40 \text{ or } A_{i,j}^{V,m} = -40, \end{cases}$$

395 for each $(i, j) \in \Omega$, see Figure 5. Following Section 2.2, a placeholder value of $p = 20$ dB is
 396 assigned to all entries of the synthetic ZDR image for which at least one of the corresponding
 397 synthetic DBZH and DBZV values is a placeholder value.

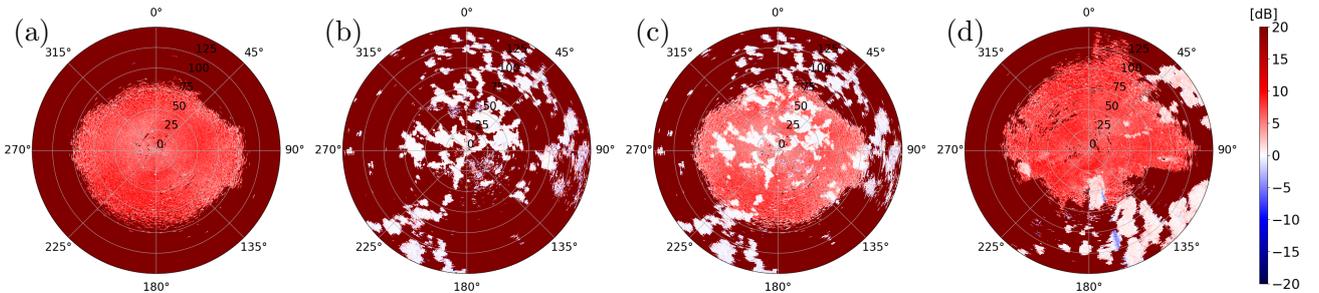


Figure 5: Comparison of Cartesian representations of different ZDR images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image, and (d) mixed summer radar image.

398 In addition to describing the size and shape of particles using DBZH and ZDR, we also use
 399 the depolarization ratio (UDR) to quantify the homogeneity of the detected particles' shapes
 400 within a radar gate. To generate a synthetic UDR image $A^{UDR,m} \in (-\infty, 0]^{n_r \times n_a}$, we use the
 401 property that pixels $(i, j) \in \Omega$ included in the winter radar image are assumed to correspond to
 402 meteorological echoes. The corresponding entries from the synthetic UDR image are assigned the
 403 depolarization ratio values $A^{UDR,w} \in (-\infty, 0]^{n_r \times n_a}$ from the winter radar image. Pixels excluded
 404 from the winter radar image are assigned the values $A^{UDR,s} \in (-\infty, 0]^{n_r \times n_a}$ from the cluttered
 405 summer radar image if they are included in this radar image. Otherwise, they are assigned a
 406 placeholder value. More precisely, the entries of $A^{UDR,m}$ are given by

$$A_{i,j}^{UDR,m} = \begin{cases} A_{i,j}^{UDR,w}, & \text{if } A^{H,w} > -40 \text{ dBZ}, \\ A_{i,j}^{UDR,s}, & \text{if } A^{H,w} = -40 \text{ dBZ and } A^{H,s} > -40 \text{ dBZ}, \\ 0, & \text{otherwise,} \end{cases}$$

407 for each $(i, j) \in \Omega$, see Figure 6.

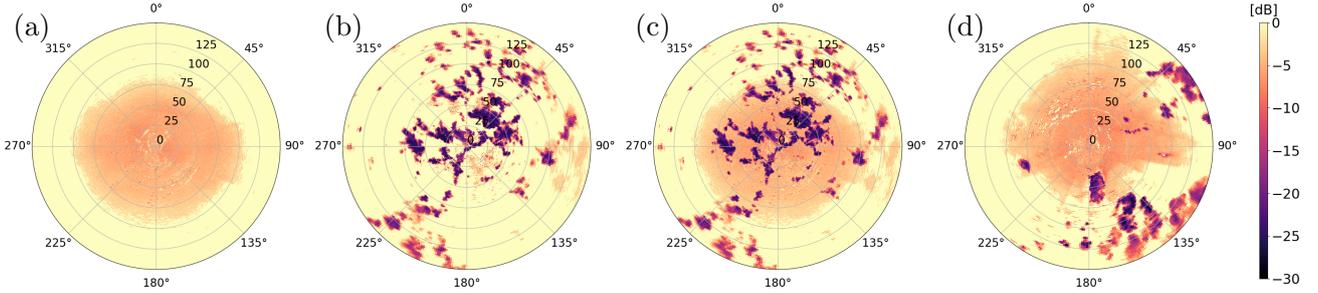


Figure 6: Comparison of Cartesian representations of different UDR images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image, and (d) mixed summer radar image.

408 Alongside the synthetic radar image $A^m = (A^{H,m}, A^{ZDR,m}, A^{UDR,m}) \in \mathbb{R}^{3 \times n_r \times n_a}$, we generate
 409 a corresponding ground truth image $G \in \{0, 1\}^{n_r \times n_a}$. Its values indicate which pixels of the
 410 synthetic radar image correspond to meteorological ($G_{i,j} = 0$) and non-meteorological echoes
 411 ($G_{i,j} = 1$), respectively. Pixels included in the winter radar image $A^{H,w}$ are assumed to correspond
 412 to meteorological echoes, as described in Section 2.2. Since the synthetic image is generated in
 413 such a way that these pixels also correspond to meteorological echoes in the synthetic image,
 414 they are assigned a ground truth label of 0. All of the other pixels are labeled 1, as they are
 415 excluded from the winter radar image. More precisely, for each pixel $(i, j) \in \Omega$, the ground truth
 416 label $G_{i,j}$ is given by

$$G_{i,j} = \begin{cases} 0, & \text{if } A^{H,w} > -40 \text{ dBZ}, \\ 1, & \text{otherwise.} \end{cases}$$

417 Generating synthetic radar images in this way provides mixed radar images with pixel-wise
 418 ground truth labeling, assigning each pixel as either meteorological or non-meteorological echoes.
 419 This ground truth allows the three synthetic channels to serve as training samples for the neural
 420 network.

421 2.5 Network training

422 The available winter dataset $\mathcal{D}^w = \{A_n^w\}_{n=1}^{180} \subset D = \mathbb{R}^{3 \times n_r \times n_a}$ consists of 180 radar images
 423 $A_n^w = (A_n^{H,w}, A_n^{V,w}, A_n^{UDR,w})$, where the radar moments present in the data are horizontal reflectivity,
 424 vertical reflectivity, and the depolarization ratio. Similarly, the available cluttered summer
 425 dataset is given by $\mathcal{D}^s = \{A_n^s\}_{n=1}^{180} \subset D$. We partition each of these datasets into two disjoint
 426 subsets. One subset of radar images is used for generating the training samples, while the other
 427 set is reserved exclusively for network evaluation. For this purpose, we partition the index set
 428 $\mathcal{I} = \{1, \dots, 180\}$ into two disjoint subsets $\mathcal{I}_{\text{train}}, \mathcal{I}_{\text{test}} \subset \mathcal{I}$, i.e.,

$$\mathcal{I}_{\text{train}} \cup \mathcal{I}_{\text{test}} = \mathcal{I} \quad \text{and} \quad \mathcal{I}_{\text{train}} \cap \mathcal{I}_{\text{test}} = \emptyset.$$

429 The index sets $\mathcal{I}_{\text{train}}, \mathcal{I}_{\text{test}} \subset \mathcal{I}$ are chosen so that all winter and cluttered summer radar images
 430 from a single radar station are used for network evaluation, while the radar images from other
 431 radar stations are used for the generation of training samples.

432 We use the methodology described in Section 2.4 to generate synthetic training samples
 433 by combining radar images from the winter and cluttered summer datasets. Each synthetic
 434 training sample consists of a synthetic radar image $A \in D$ and a corresponding ground truth
 435 image $G \in \{0, 1\}^{n_r \times n_a}$. The synthetic radar image comprises the radar moments of horizontal
 436 reflectivity, differential reflectivity, and the depolarization ratio, i.e., $A = (A^H, A^{ZDR}, A^{UDR})$.

437 Before generating the synthetic training samples, we apply various data augmentation steps
 438 to increase their diversity. More precisely, as described in Section 2.4, the DBZH and DBZV
 439 images from the winter and cluttered summer datasets are scaled by reflectivity weighting factors
 440 $\lambda_w, \lambda_s > 0$, respectively. In addition, we apply a rotation and a possible inversion of orientation.

441 Specifically, we introduce a parametric mapping $g_{\psi, \psi'}: D \times D \rightarrow D \times \{0, 1\}^{n_r \times n_a}$
 442 that combines two radar images to generate a synthetic radar image and a corresponding
 443 ground truth image. The space of admissible parameters for data augmentation is given by
 444 $\Psi = \mathbb{R}^+ \times \{1, \dots, n_a\} \times \{-1, 1\}$. Each parameter vector $\psi = (\lambda, \varphi, o) \in \Psi$ contains the reflectivity
 445 weighting factor $\lambda > 0$, the rotation angle $\varphi \in \{1, \dots, n_a\}$, and the orientation flag
 446 $o \in \{-1, 1\}$. Here, a rotation angle of $\varphi \in \{1, \dots, n_a\}$ indicates a rotation by φ rays. Hence,
 447 a shift of φ rows is applied to the radar image. An orientation flag of $o = -1$ indicates an
 448 inverted orientation. In that case, the radar image is flipped along the x -axis. To distinguish the
 449 transformations we apply to winter and cluttered summer radar images, we introduce restricted
 450 admissible parameter spaces $\Psi^w, \Psi^s \subset \Psi$. They only differ in the admissible range of the reflectivity
 451 weighting factor, which is given by the interval $[5, 8]$ for Ψ^w and by $[0.9, 1.1]$ for Ψ^s . The
 452 set of synthetic training samples generated by combining two radar images is then denoted by

$$\mathcal{D}_{\text{train}}^m = \{g_{\psi_w, \psi_s}(A_{n_w}^w, A_{n_s}^s) : \psi_w \in \Psi^w, \psi_s \in \Psi^s, n_w, n_s \in \mathcal{I}_{\text{train}}\}.$$

453 These synthetic training samples feature both meteorological and non-meteorological echoes.
 454 However, in the training process, the network should be provided with training samples from
 455 various scenarios. Since the radar images from the winter and cluttered summer datasets, \mathcal{D}^w
 456 and \mathcal{D}^s , are used to represent meteorological and non-meteorological echoes, respectively, it
 457 is possible to generate synthetic radar images that feature only one type of echo. To this end,
 458 we introduce a radar image $A^{\text{exc}} \in D$ in which all pixels are excluded, i.e., all of its entries
 459 are assigned the corresponding placeholder values. A synthetic radar image that features only
 460 meteorological echoes is then generated by combining a winter radar image and the radar image
 461 A^{exc} . In particular, the set of training samples featuring only meteorological echoes is given by

$$\mathcal{D}_{\text{train}}^w = \{g_{\psi_w, \psi}(A_{n_w}^w, A^{\text{exc}}) : \psi_w \in \Psi^w, \psi \in \Psi, n_w \in \mathcal{I}_{\text{train}}\}.$$

462 Analogously, the set of synthetic radar images featuring only non-meteorological echoes is given
 463 by

$$\mathcal{D}_{\text{train}}^s = \{g_{\psi, \psi_s}(A^{\text{exc}}, A_{n_s}^s) : \psi \in \Psi, \psi_s \in \Psi^s, n_s \in \mathcal{I}_{\text{train}}\}.$$

464 Note that the test dataset $\mathcal{D}_{\text{test}}^m$, comprising radar images featuring both meteorological and non-
 465 meteorological echoes, and the test datasets $\mathcal{D}_{\text{test}}^w$ and $\mathcal{D}_{\text{test}}^s$, comprising radar images featuring
 466 only one type of echo, are generated analogously using the index set $\mathcal{I}_{\text{test}}$ instead of $\mathcal{I}_{\text{train}}$.

467 As stated in Section 2.3, the network can be described by its vector of trainable parameters
 468 $\theta \in \Theta$ using a function $f_\theta: \mathbb{R}^{k \times h \times w} \rightarrow [0, 1]^{h \times w}$. Here, $h = n_r$ and $w = n_a$ denote the image
 469 dimensions of a single-channel image, and $k = 3$ is the number of input channels. Then, for a
 470 given radar image $A \in D$, the output $f_\theta(A)$ of the network represents the matrix of probabilities
 471 $(f_\theta(A))_{i,j}$ that a pixel (i, j) corresponds to a non-meteorological echo for each $(i, j) \in \Omega$.

472 The training process of the network involves iterating over 5600 mini-batches, and the train-
 473 able parameters of the network are adapted at each iteration. More precisely, in each training
 474 step, a mini-batch $B = \{(A_n, G_n)\}_{n=1}^4$ is built by generating a total of four synthetic radar
 475 images and corresponding ground truth images. Each training sample is chosen as follows: With
 476 a probability of 0.75, it is uniformly sampled from the training dataset $\mathcal{D}_{\text{train}}^m$ and therefore fea-
 477 tures both meteorological and non-meteorological echoes. Otherwise, we either sample it from

478 $\mathcal{D}_{\text{train}}^w$ or $\mathcal{D}_{\text{train}}^s$, with a probability of 0.125 for each case. For each training sample (A_n, G_n)
 479 of the mini-batch B , with $n \in \{1, 2, 3, 4\}$, the network output $f_\theta(A_n)$ is computed and com-
 480 pared to the ground truth image G_n using the binary cross-entropy loss, which is a measure of
 481 how well the network outputs match the ground truth images. More specifically, for any batch
 482 $B = \{(A_n, G_n)\}_{n=1}^4$ and vector of trainable network parameters $\theta \in \Theta$, the cross-entropy loss
 483 $L_\theta(B)$ of the batch is given by

$$L_\theta(B) = -\frac{1}{4} \sum_{n=1}^4 \sum_{(i,j) \in \Omega} \ell(f_\theta(A_n)_{i,j}, (G_n)_{i,j}),$$

484 where $\ell(y, x) = x \log(y) + (1 - x)(1 - \log(y))$, with $f_\theta(A_n)_{i,j}$ and $(G_n)_{i,j}$ denoting the entries
 485 of $f_\theta(A_n)$ and G_n , respectively. After computing the cross-entropy loss $L_\theta(B)$, its gradient is
 486 computed with respect to the vector of trainable network parameters θ . Based on that, this
 487 parameter vector θ is then adjusted using the Adam optimizer [19]. Furthermore, the performance
 488 of the trained network is compared to that of the current state-of-the-art classification method
 489 used at DWD, see Section 3.

490 3 Results

491 In this section, we evaluate the performance of the trained network f_θ through both visual
 492 inspection and quantitative analysis of the predicted segmentations. Although all mixed radar
 493 images used for training were synthetically generated, we also validate network performance on
 494 experimentally measured mixed radar images using expert-labeled ground truth. Furthermore,
 495 we benchmark these results against those achieved by a state-of-the-art (SotA) method developed
 496 by DWD.

497 3.1 Metrics for quantitative evaluation

498 To quantitatively evaluate the segmentation performance, we compare the network output
 499 against corresponding ground truth. For that purpose, pixels $(i, j) \in \Omega$ with a horizontal re-
 500 flectivity value of $A_{i,j}^H = -40$ dBZ are excluded from the evaluation, as they conclusively do
 501 not correspond to meteorological echoes. Furthermore, the physical area of a radar gate $g_{i,j}$ in-
 502 creases linearly with its distance from the radar. Therefore, to obtain metrics with a meaningful
 503 physical interpretation regarding the spatial extent of the echoes, we introduce a weight matrix
 504 $W = (W_{i,j}) \in \mathbb{R}^{n_r \times n_a}$. For each pixel $(i, j) \in \Omega$, its weight $W_{i,j}$ is given by

$$W_{i,j} = \begin{cases} w_i, & \text{if } A_{i,j}^H > -40 \text{ dBZ,} \\ 0, & \text{otherwise.} \end{cases}$$

505 Here, $w_i = \frac{\pi}{n_a}(2r_i - d)d$ denotes the physical area of radar gate $g_{i,j}$, where $d = 250$ m is the
 506 range resolution, and $r_i = id$ is the outer radius of the radar gate. Using the weights of the
 507 individual pixels, we define the weighted area $\mathcal{A}(\Omega')$ of any subset of pixels $\Omega' \subset \Omega$ as the sum
 508 of the weights of the pixels in that set, given by

$$\mathcal{A}(\Omega', W) = \sum_{(i,j) \in \Omega'} W_{i,j}.$$

509 Using a radar image $A \in D = \mathbb{R}^{3 \times n_r \times n_a}$ as input, the network output $f_\theta(A) \in [0, 1]^{n_r \times n_a}$ is
 510 transformed into a binary prediction $P = (P_{i,j}) \in \{0, 1\}^{n_r \times n_a}$ by thresholding at 0.5. Specifically,
 511 the prediction $P_{i,j}$ at pixel $(i, j) \in \Omega$ is given by

$$P_{i,j} = \begin{cases} 0, & \text{if } (f_\theta(A))_{i,j} < 0.5, \\ 1, & \text{if } (f_\theta(A))_{i,j} \geq 0.5, \end{cases}$$

512 where prediction $P_{i,j} = 0$ represents meteorological echoes, and $P_{i,j} = 1$ represents non-
 513 meteorological echoes.

514 To quantify segmentation performance, we compute the area-weighted intersection over union
 515 (IOU), also referred to as the critical success index [23]. For class $c \in \{0, 1\}$, the class-specific
 516 IOU is defined as the ratio of the area of the intersection of pixels classified as class c in the
 517 prediction and pixels labeled as class c in the ground truth, to the area of their union, given by

$$\text{IOU}^c(G, P, W) = \frac{\mathcal{A}(\Omega|_{P=c} \cap \Omega|_{G=c}, W)}{\mathcal{A}(\Omega|_{P=c} \cup \Omega|_{G=c}, W)}.$$

518 Here, the set of pixels assigned to class c by the prediction P and the set of pixels labeled as
 519 class c in the ground truth G are given by

$$\Omega|_{P=c} = \{(i, j) \in \Omega : P_{i,j} = c\} \quad \text{and} \quad \Omega|_{G=c} = \{(i, j) \in \Omega : G_{i,j} = c\},$$

520 respectively. We then obtain a single performance metric for the entire radar image by computing
 521 the combined IOU as the area-weighted sum of the class-specific IOUs, given by

$$\text{IOU}(G, P, W) = \sum_{c \in \{0,1\}} \frac{\mathcal{A}(\Omega|_{G=c})}{\mathcal{A}(\Omega)} \text{IOU}^c(G, P, W).$$

522 While the area-weighted IOU is as an effective aggregate metric for overall segmentation
 523 quality, it does not distinguish between incorrectly classifying pixels as meteorological and in-
 524 correctly classifying pixels as non-meteorological echoes. Therefore, it is helpful to quantify the
 525 portion of correct predictions within specific regions, e.g., regions labeled as class $c = 0$ in the
 526 ground truth. To this end, we define the area-weighted accuracy on a subset $\Omega' \subseteq \Omega$. This metric
 527 calculates the portion of the weighted area within Ω' that is correctly classified and is given by

$$\text{ACC}(G, P, W; \Omega') = \frac{\mathcal{A}(\Omega' \cap \Omega|_{P=G})}{\mathcal{A}(\Omega')},$$

528 where $\Omega|_{P=G} = \{(i, j) \in \Omega : P_{i,j} = G_{i,j}\}$ denotes the set of all correctly classified pixels. Setting
 529 $\Omega' = \Omega$ yields the area-weighted accuracy for the entire radar image.

530 This general formulation allows us to derive the class-specific precision and recall, which pro-
 531 vide insight into specific error types. The precision PRE^c corresponds to the accuracy evaluated
 532 on $\Omega|_{P=c}$, i.e., the set of pixels classified as class c . It quantifies the correctness of the prediction,
 533 where a low precision implies a high rate of false positives for class c , i.e., a large portion of
 534 pixels classified as class c are incorrectly classified. Conversely, the recall REC^c corresponds to
 535 the accuracy evaluated on $\Omega|_{G=c}$, i.e., the set of pixels labeled as class c in the ground truth.
 536 It measures the completeness of the detection, where a low recall implies a high rate of false
 537 negatives, i.e., a large number of pixels with ground truth label c are incorrectly classified. These
 538 metrics are given by

$$\text{PRE}^c(G, P, W) = \text{ACC}(G, P, W; \Omega|_{P=c}) \quad \text{and} \quad \text{REC}^c(G, P, W) = \text{ACC}(G, P, W; \Omega|_{G=c}).$$

539 While false positives for meteorological echoes (lower REC^1 and PRE^0) introduce undesirable
 540 noise, missing weather events (lower REC^0 and PRE^1) constitute a severe loss of information.
 541 From a meteorological perspective, a high recall REC^0 for meteorological echoes is essential to
 542 ensure that weather events are correctly identified. This is particularly important for meteorolog-
 543 ical echoes with high reflectivity values, as they correspond to severe weather events. Therefore,
 544 we also investigate the network performance with respect to the horizontal reflectivity values
 545 A^{H} . For this, we evaluate the segmentation performance on subsets of pixels with horizontal
 546 reflectivity values $A_{i,j}^{\text{H}}$ in the semi-open interval $[z, z + 5)$ for $z \in \mathbb{R}$, given by

$$\Omega|_{A^{\text{H}} \in [z, z+5)} = \{(i, j) \in \Omega : A_{i,j}^{\text{H}} \in [z, z + 5)\}.$$

547 Similarly, we analyze the segmentation performance with respect to the distance from the radar,
 548 i.e., on subsets of pixels

$$\Omega|_{i \in [r, r+19]} = \{(i, j) \in \Omega : r \leq i \leq r + 19\},$$

549 for range indices $r \in \{1, \dots, n_r - 19\}$. This is motivated by the fact that non-meteorological echoes
 550 in winter radar images, which are rare but can appear close to the radar station, as shown in
 551 Figure 2d, result in incorrect ground truth labels in the synthetically generated training data.

552 Using these subsets of pixels, we compute the area-weighted accuracy depending on reflec-
 553 tivity and distance from the radar, given by

$$\begin{aligned} \text{ACC}^{\text{H}}(G, P, W; z) &= \text{ACC}(G, P, W; \Omega|_{A^{\text{H}} \in [z, z+5)}), \\ \text{ACC}^{\text{D}}(G, P, W; r) &= \text{ACC}(G, P, W; \Omega|_{i \in [r, r+19]}), \end{aligned}$$

554 respectively, and the recall for class $c \in \{0, 1\}$ depending on reflectivity and distance from the
 555 radar, given by

$$\begin{aligned} \text{REC}^{c, \text{H}}(G, P, W; z) &= \text{ACC}(G, P, W; \Omega|_{G=c} \cap \Omega|_{A^{\text{H}} \in [z, z+5)}), \\ \text{REC}^{c, \text{D}}(G, P, W; r) &= \text{ACC}(G, P, W; \Omega|_{G=c} \cap \Omega|_{i \in [r, r+19]}), \end{aligned}$$

556 respectively.

557 Note that the introduced metrics can be computed for individual radar images and cor-
 558 responding ground truth images, as well as for datasets $\mathcal{D} = \{(A_n, G_n)\}_{n=1}^N \subset D \times \{0, 1\}^{n_r \times n_a}$
 559 consisting of $N \in \mathbb{N}$ such image pairs. In this case, the metric is first computed for each radar
 560 image in the dataset individually and then aggregated as an area-weighted sum. For example,
 561 the area-weighted accuracy ACC on subsets $\Omega' = (\Omega'_n)_{n=1}^N \in \Omega^N$ for the dataset \mathcal{D} is given by

$$\text{ACC}(\mathcal{D}; \Omega') = \frac{1}{\sum_{n=1}^N \mathcal{A}(\Omega'_n, W_n)} \sum_{n=1}^N \mathcal{A}(\Omega'_n, W_n) \text{ACC}(G_n, P_n, W_n; \Omega'_n),$$

562 where $P_n \in \{0, 1\}^{n_r \times n_a}$ is the prediction and $W_n \in \mathbb{R}^{n_r \times n_a}$ is the weight matrix for the n -th radar
 563 image. As before, setting $\Omega'_n = \Omega$ for all $n \in \{1, \dots, N\}$ yields the area-weighted accuracy for
 564 the dataset \mathcal{D} , while the evaluation on $\Omega'_n = \Omega|_{P_n=c}$ and $\Omega'_n = \Omega|_{G_n=c}$ for all $n \in \{1, \dots, N\}$
 565 and $c \in \{0, 1\}$ yields the class-specific precision and recall for the dataset \mathcal{D} , respectively.
 566 Analogously, the intersection over union IOU for the dataset \mathcal{D} is given by

$$\text{IOU}(\mathcal{D}) = \frac{1}{\sum_{n=1}^N \mathcal{A}(\Omega, W_n)} \sum_{n=1}^N \mathcal{A}(\Omega, W_n) \text{IOU}(G_n, P_n, W_n).$$

567 3.2 State-of-the-art method

568 The SotA method used by DWD comprises a series of processing steps designed to identify and
 569 filter non-meteorological echoes, as well as a correction for attenuation, as described in [13].
 570 These processing steps include a threshold filter, a filter for spokes and rings, a blacklist filter, a
 571 polarimetric clutter filter based on fuzzy logic, a second-trip filter, and a speckle filter. Notably,
 572 the polarimetric clutter filter has been extended in 2025 to include an additional so-called insect
 573 filter, which is used operationally during the summer months (April – October).

574 Echoes from insects, which are referred to as insect echoes, are typically characterized by
 575 the following features in the radar moments: reduced cross-correlation coefficient ρ_{HV} (< 0.9),
 576 increased differential reflectivity ZDR (> 2 dB), and low reflectivity factors DBZH and DBZV
 577 (< 20 dBZ). In principle, pixels corresponding to insect echoes can be correctly classified based
 578 on these features. However, there are also weather events, such as snowfall, that exhibit similar
 579 features, particularly at greater distances from the radar. Moreover, pixels corresponding to
 580 a mixture of meteorological echoes and clutter can also exhibit these features but should be
 581 classified as meteorological echoes.

582 The goal of the DWD insect filter is to filter out insect echoes, i.e., identify the corresponding
 583 pixels and classify them as non-meteorological echoes. It utilizes the radar moments introduced
 584 previously, as well as the so-called clutter correction moment (CCORH), which quantifies the
 585 portion of the reflectivity filtered out by the Doppler filter in dBZ. The insect filter is a decision
 586 tree-based pixel classifier that considers the pixel’s radar moment values and the statistical
 587 properties of its local neighborhood.

588 More precisely, the newly implemented insect filter operates in two steps: First, all pixels
 589 eligible for filtering are identified. Then, for each pixel identified in the first step, it is tested
 590 whether the pixel’s UDR value, or the UDR value in its local neighborhood, allows for filtering.
 591 For this purpose, the so-called $q \times q$ neighborhood $N_{(i,j)}^{q \times q} \subset \Omega$ of pixel $(i, j) \in \Omega$ for an odd integer
 592 $q \in \{1, 3, 5, \dots\}$ is considered, which is given by

$$N_{(i,j)}^{q \times q} = \left\{ (k, \ell) \in \Omega : |i - k| \leq \frac{q-1}{2} \text{ and } d(j, \ell) \leq \frac{q-1}{2} \right\},$$

593 where d denotes the periodic distance measure as defined in Eq. (2). To incorporate information
 594 from the surroundings of a pixel into the filtering process, it has proven helpful to use percentile
 595 values. In contrast to using means, maxima, or minima, percentiles are insensitive to individual
 596 extreme values.

597 A pixel $(i, j) \in \Omega$ is eligible for filtering if (a) its reflectivity value is below 20 dBZ, (b) the
 598 CCORH value in its 3×3 neighborhood $N_{(i,j)}^{3 \times 3}$ is not significant, and (c) the ZDR value in $N_{(i,j)}^{3 \times 3}$ is
 599 sufficiently different from 0 dB. All three conditions must be met simultaneously to prevent the
 600 misclassification of meteorological echoes. In particular, the CCORH condition was introduced
 601 to retain pixels within precipitation areas that, due to a mixture of meteorological echoes and
 602 clutter, exhibit a reduced quality of the polarimetric moments. Otherwise, they would be filtered
 603 out based on their high UDR value. Furthermore, the considered radar moments of snow and
 604 insects are very similar, except for ZDR. Therefore, the ZDR condition is necessary to retain
 605 pixels whose associated radar gates contain weak snowfall, especially at greater distances from
 606 the radar.

607 To avoid filtering pixels corresponding to meteorological echoes at greater distances, where
 608 the reduced signal-to-noise ratio results in lower-quality polarimetric moments, an additional
 609 distance-dependent condition is implemented. Pixels at greater distances from the radar are

610 only eligible for filtering if a certain percentage of pixels closer to the radar have already been
 611 filtered out by the insect filter. This condition accounts for the fact that insects usually occur in
 612 near-surface layers, and if insects are present at higher altitudes, they are typically also found
 613 closer to the ground.

614 An eligible pixel is filtered out if either the pixel itself exhibits a high UDR value (> -5 dB),
 615 or if high UDR values dominate in the local neighborhood (20th-percentile UDR in its 5×5
 616 neighborhood > -11 dB) and the pixel itself does not show a particularly low value (< -18 dB).
 617 These two conditions ensure that pixels with a very high likelihood of non-meteorological echoes
 618 (based on their own UDR values), as well as pixels mainly surrounded by pixels corresponding
 619 to non-meteorological echoes with moderate UDR values, are filtered out.

620 Note that the distance-dependent condition and the thresholds for conditions (b) and (c)
 621 have to be chosen for each radar station individually, as they depend on the radar wavelength,
 622 the scan strategy and the surrounding topography, which determine the exact characteristics of
 623 non-meteorological echoes. We therefore do not specify any explicit values in this paper.

624 3.3 Network evaluation

625 We evaluate the trained network using the metrics stated in Section 3.1 and compare its perfor-
 626 mance to the SotA method detailed in Section 3.2. Segmentation performance is evaluated on
 627 training datasets $\mathcal{D}_{\text{train}}^m$, $\mathcal{D}_{\text{train}}^w$, and $\mathcal{D}_{\text{train}}^s$, and test datasets $\mathcal{D}_{\text{test}}^m$, $\mathcal{D}_{\text{test}}^w$, and $\mathcal{D}_{\text{test}}^s$, as introduced
 628 in Section 2.5. Moreover, both methods are also evaluated on a non-synthetic mixed dataset
 629 $\mathcal{D}_{\text{exp}}^m$, given by

$$\mathcal{D}_{\text{exp}}^m = \{(A_n, G_n) : n \in \{1, \dots, 5\}\},$$

630 where $A_n \in D$ are radar images measured in June 2022 featuring both meteorological and non-
 631 meteorological echoes, and $G_n \in \{0, 1\}^{n_r \times n_a}$ are the corresponding ground truth images, which
 632 were manually labeled by a subject expert.

633 Table 1 shows a comparison of the resulting metrics averaged over the respective datasets as
 634 described in Section 3.1. The overall performance metrics on the test datasets are slightly lower
 635 than those on the training datasets. The accuracy ACC and intersection over union IOU on the
 636 mixed test dataset $\mathcal{D}_{\text{test}}^m$ are 0.964 and 0.932, respectively, compared to 0.982 and 0.966 on the
 637 mixed training dataset $\mathcal{D}_{\text{train}}^m$.

Table 1: Area-weighted performance metrics for the evaluation of the trained network and the SotA method. Precision PRE^c and recall REC^c are computed for meteorological ($c = 0$) and non-meteorological echoes ($c = 1$).

Method	Type of data	PRE^0	PRE^1	REC^0	REC^1	ACC	IOU
network	mixed synthetic images (training)	0.970	0.959	0.962	0.930	0.982	0.966
	cluttered summer images (training)	-	-	-	0.989	0.989	0.989
	winter images (training)	-	-	0.937	-	0.937	0.937
	mixed synthetic images (test)	0.904	0.991	0.979	0.958	0.964	0.932
	cluttered summer images (test)	-	-	-	0.977	0.977	0.977
	winter images (test)	-	-	0.962	-	0.962	0.962
	mixed measured images	0.957	0.969	0.977	0.906	0.960	0.923
SotA	mixed measured images	0.988	0.775	0.821	0.983	0.893	0.810

638 Figure 7 visualizes the echoes that were correctly and incorrectly classified by the network
639 for one synthetic test image in dataset $\mathcal{D}_{\text{test}}^m$. Most of the meteorological echoes, particularly
640 those with high reflectivity, were correctly identified. This aligns with the recall value REC^0
641 for meteorological echoes of 0.979 on dataset $\mathcal{D}_{\text{test}}^m$. Moreover, the resulting segmentation does not
642 show any notable holes within areas surrounded by meteorological echoes. However, as can be
643 seen in Figure 7c, misclassified non-meteorological echoes are clearly visible as small speckles,
644 leading to a lower precision value PRE^0 for meteorological echoes of 0.904.

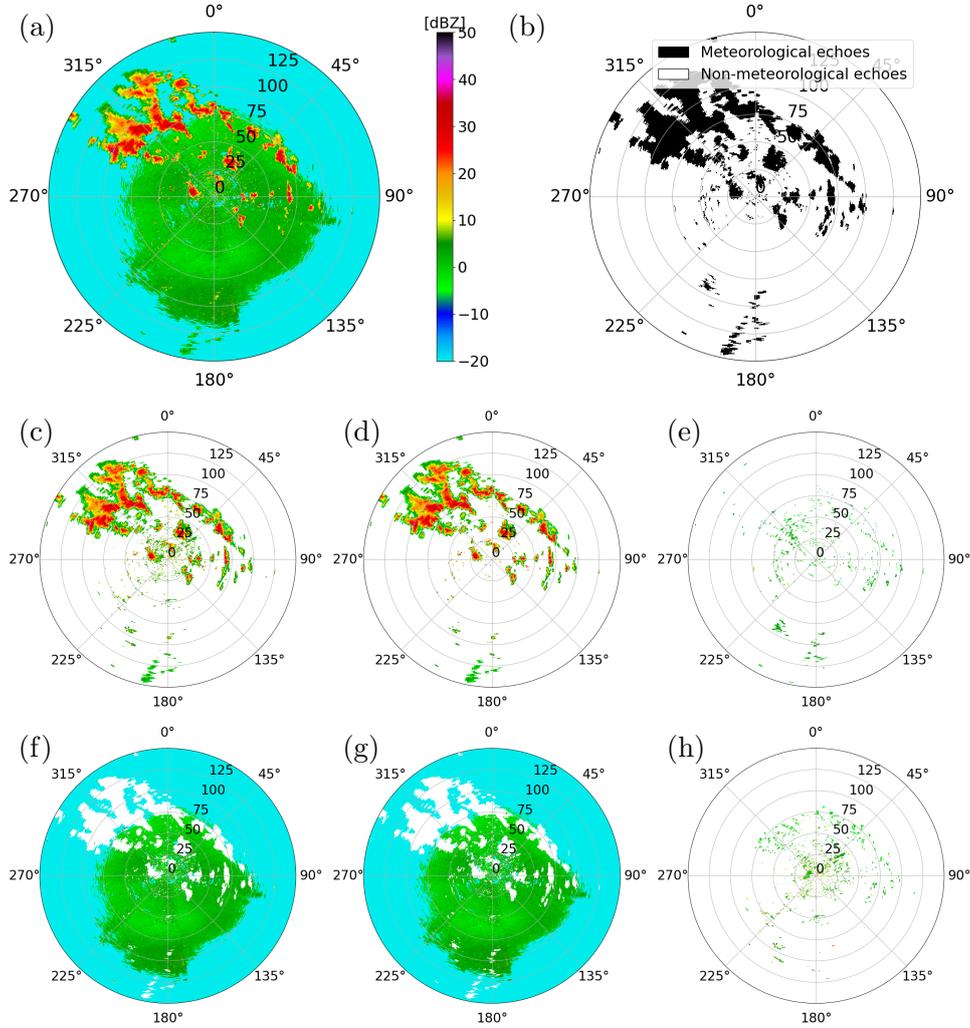


Figure 7: Evaluation of the network on a synthetic test data: (a) input data and (b) corresponding ground truth, (c) pixels classified as meteorological echoes by the network, (d) correctly classified meteorological echoes, and (e) meteorological echoes misclassified as non-meteorological echoes. The bottom row shows (f) pixels classified as non-meteorological echoes, (g) correctly classified non-meteorological echoes, and (h) non-meteorological echoes misclassified as meteorological echoes.

645 When evaluated on the experimentally measured mixed dataset $\mathcal{D}_{\text{exp}}^m$, the network perfor-
646 mance is comparable, with accuracy ACC and intersection over union IOU of 0.960 and 0.923,
647 respectively. However, when evaluating the SotA method on the same mixed dataset $\mathcal{D}_{\text{exp}}^m$, it
648 performs significantly worse, with accuracy ACC and intersection over union IOU of 0.893 and
649 0.810, respectively. These overall differences are caused by larger differences in precision PRE^1
650 for non-meteorological echoes and recall REC^0 for meteorological echoes, where the SotA method
651 achieves scores of 0.775 and 0.821, respectively, compared to the neural network with scores of

652 0.969 and 0.977. The SotA method achieves better scores for recall REC^1 for non-meteorological
653 echoes and precision PRE^0 for meteorological echoes, with values of 0.983 and 0.988, respectively,
654 compared to the neural network with scores of 0.906 and 0.957.

655 Figure 8 shows the incorrectly segmented echoes of the network (Figures 8c and 8f) and of
656 the SotA method (Figures 8b and 8e) for one experimentally measured mixed radar image in
657 dataset \mathcal{D}_{exp}^m . The top row shows the incorrect segmentations of meteorological echoes for both
658 methods, with the neural network making significantly fewer errors. Here, the SotA method
659 tends to incorrectly classify the pixels located at the boundary of weather events. However, the
660 SotA method performs significantly better in correctly identifying non-meteorological echoes
661 near the radar station.

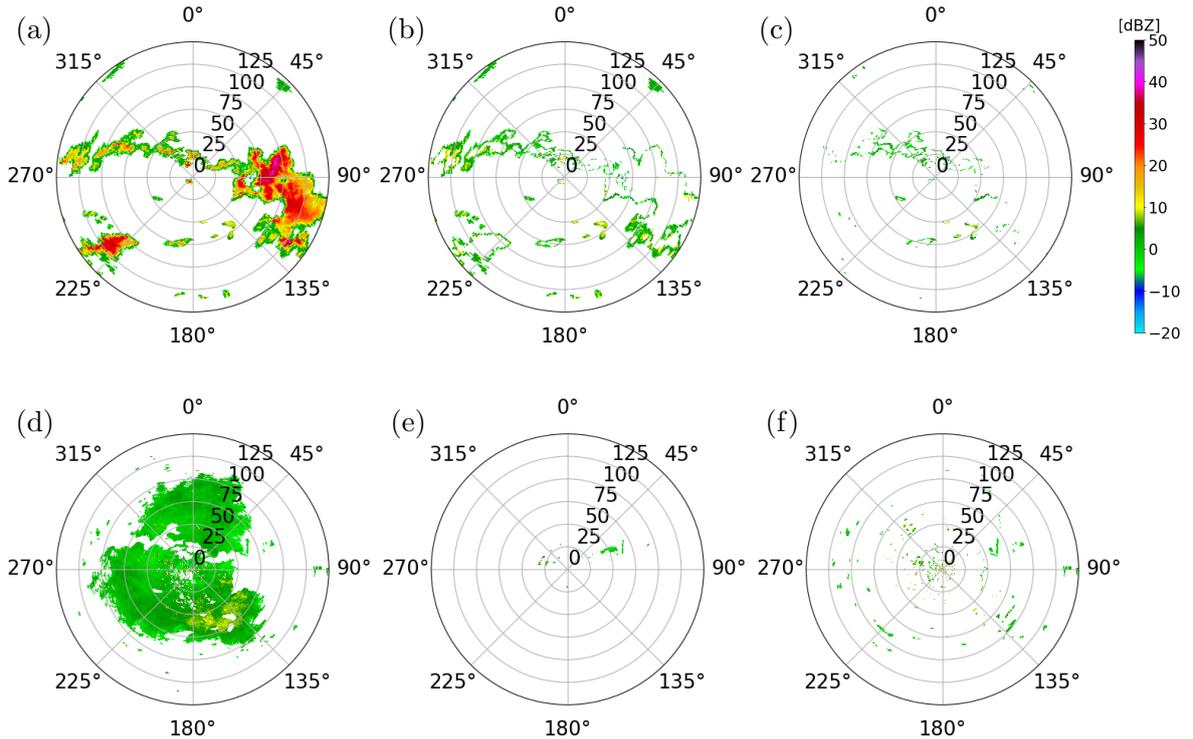


Figure 8: Evaluation of the SotA method and the trained network on the experimentally measured radar image shown in Figure 1b: The first row shows (a) the ground truth for the meteorological echoes and the meteorological echoes misclassified as non-meteorological echoes by (b) the SotA method and (c) the trained network. Analogously, the second row shows (d) the hand-labeled ground truth for the non-meteorological echoes and the non-meteorological echoes misclassified as meteorological echoes by (e) the SotA method and (f) the trained network.

662 In order to further investigate this, Figure 9a shows the accuracy depending on the dis-
663 tance, while Figures 9b and 9c show the class-specific recall scores for meteorological and non-
664 meteorological echoes, respectively. The networks accuracy ACC^D on the experimentally mea-
665 sured mixed dataset \mathcal{D}_{exp}^m increases with distance from the radar, rising from below 0.9 within
666 25 km to nearly 1.0 beyond 75 km. The same trend can be observed for the test dataset \mathcal{D}_{test}^m ,
667 for which only slightly higher scores are achieved by the network. When compared to the SotA
668 method, the recall $REC^{0,D}$ for meteorological echoes depending on distance is higher for the net-
669 work than for the SotA method for all distances, and the opposite is true for non-meteorological
670 echoes. The most significant differences occur close to the radar, at a distance of less than
671 25 km, where the neural network achieves a recall $REC^{1,D}$ of 0.877, compared to that of the
672 SotA method with 0.981. This relatively lower recall close to the radar quantifies the errors we

673 see in Figure 8c, where a significant number of non-meteorological echoes near the radar are
674 incorrectly classified. When considering the recall $\text{REC}^{0,D}$ for meteorological echoes depending
675 on distance, the network clearly outperforms the SotA method at distances further from the
676 radar, with a recall value of 0.989 compared to that of the SotA method with 0.831 beyond
677 75 km. The echoes further away from the radar represent a proportionally larger area, as shown
678 in the histograms, and thus the network outperforms the SotA method overall.

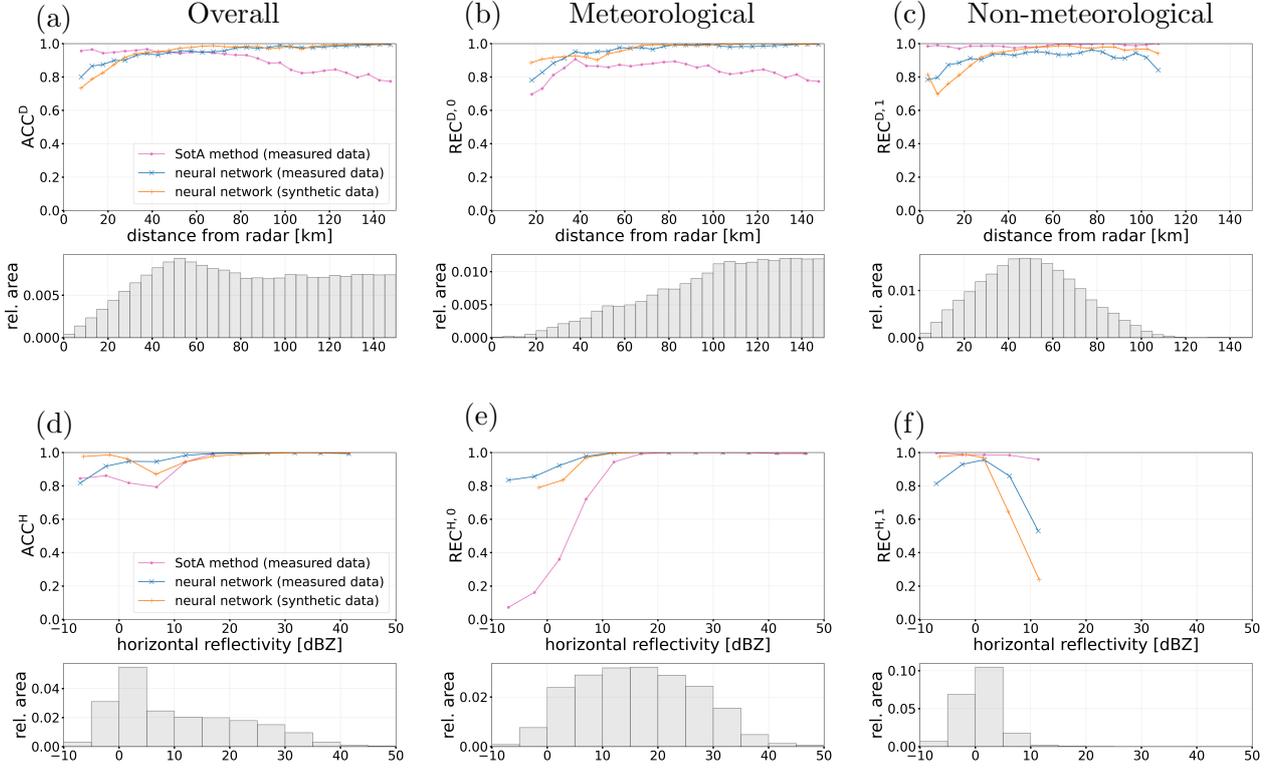


Figure 9: First row: Accuracy in dependence on the distance from the radar for (a) meteorological and non-meteorological echoes combined, (b) meteorological echoes, and (c) non-meteorological echoes, as classified in the ground truth. Second row: Accuracy in dependence on the horizontal reflectivity values for the same three sets of pixels. The metrics are only shown for distances and horizontal reflectivities with a relative area greater than 0.0005, respectively. The histograms below show the area distribution for each of the three sets of pixels, respectively. The SotA method was evaluated on the experimentally measured mixed dataset $\mathcal{D}_{\text{exp}}^m$ and the neural network was evaluated on the same experimentally measured dataset and additionally on the mixed test dataset $\mathcal{D}_{\text{test}}^m$.

679 Analogously, we evaluate the accuracy and class-specific recall in dependence on the horizon-
680 tal reflectivity values, as shown in Figures 9d to 9f. For meteorological echoes with a reflectivity
681 higher than 10 dBZ, both the network and the SotA method achieve a recall $\text{REC}^{0,H}$ near 1.0.
682 However, for horizontal reflectivities below 10 dBZ, the SotA performance degrades significantly,
683 whereas the network maintains a recall of 0.936. While meteorological echoes with reflectivity
684 values below 0 dBZ are rare in the dataset, those with reflectivity values between 0 and 10 dBZ are
685 much more frequent, as shown in the histogram of Figure 9e. For example, the edges of weather
686 events are typically less reflective and are often incorrectly classified as non-meteorological echoes
687 by the SotA method, as shown in Figure 8b. In the case of non-meteorological echoes, the network
688 struggles to classify echoes with high reflectivity (> 10 dBZ), achieving a lower recall $\text{REC}^{1,H}$ of
689 0.430 compared to the SotA method, which achieves a recall value of 0.961. Nevertheless, such

high-reflectivity non-meteorological echoes are rare in the dataset.

4 Discussion

The results presented in Section 3 demonstrate that a neural network is capable of effectively identifying non-meteorological echoes in radar images. One of the most significant findings is the network’s ability to generalize from synthetically generated radar images to experimentally measured mixed radar images $\mathcal{D}_{\text{exp}}^m$. Although the training process relied exclusively on synthetically generated radar images, network performance on experimentally measured radar images was in line with, or exceeding, the performance of a SotA method currently used at DWD.

A key advantage of the proposed network-based method over the SotA method used at DWD lies in the preservation of meteorological echoes, particularly at the boundaries of weather events and in regions of weak precipitation. As illustrated in Figure 8b, and quantified in Figure 9e, the SotA method tends to incorrectly classify the boundaries of weather events as non-meteorological echoes. This is likely due to the SotA method’s reliance on heuristically chosen thresholds and local neighborhood statistics (e.g., 3×3 or 5×5 neighborhoods). In contrast, the network utilizes context from a larger neighborhood, allowing it to better recognize the spatial structures of weather events.

However, the evaluation also highlighted limitations that should be addressed in future work. The network showed a lower recall REC^1 for non-meteorological echoes compared to the SotA method, particularly for radar gates close to the radar at distances of less than 25 km, as shown in Figure 9c. The network’s errors are likely a consequence of the underlying winter radar images used for training. For the generation of mixed synthetic radar images, we assign echoes in winter radar images exclusively as meteorological echoes. However, in reality, the winter radar images also contain non-meteorological echoes, particularly in areas close to the radar, as can be seen in Figure 2. Consequently, the corresponding pixels are incorrectly labeled as meteorological in the ground truth. Therefore, the network is disincentivized from correctly identifying non-meteorological echoes near the radar during training.

Regarding the model’s generalizability, the evaluation results indicate that the network is robust and did not suffer from significant overfitting during training. The overall performance metrics on the synthetic test dataset $\mathcal{D}_{\text{test}}^m$ are very similar to those observed on the synthetic training dataset $\mathcal{D}_{\text{train}}^m$. Furthermore, the network demonstrated robustness across different times of day. The mixed training data $\mathcal{D}_{\text{train}}^m$ was generated by combining cluttered summer radar images (10:00–11:00 UTC) with winter radar images (17:00–18:00 UTC). The scaling factors λ_w and λ_s used to combine these were explicitly calibrated so that the generated images resembled mixed summer radar images measured at 09:00–10:00 UTC. Despite this specific calibration, the network performed well on the experimentally measured dataset $\mathcal{D}_{\text{exp}}^m$, which includes radar images collected in the afternoon (16:00 UTC). Moreover, the network was trained using radar images from only three one-hour time periods. To further enhance the network’s generalization capabilities, future work should incorporate data from a broader range of meteorological situations.

There are several improvements to the proposed method. First, as seen in Figure 7, the network output includes isolated misclassified pixels. Applying a post-processing despeckle filter, similar to the one described in Section 2.2 for the input data, could remove these small artifacts and further improve precision without adding methodological complexity. Second, the current approach processes radar images from a single elevation angle and time step. Incorporating additional elevation angles or time steps could provide the network with additional spatially-resolved

735 or time-resolved structural information. Although this would add significant complexity to the
 736 proposed methodology, it represents a very promising adjustment. For one, non-meteorological
 737 echoes are less common at higher elevation angles, whereas weather events often extend verti-
 738 cally. Furthermore, incorporating time-resolved information would allow the network to exploit
 739 differences in movement patterns. For example, insects are typically either stationary or move
 740 irregularly, whereas weather events exhibit coherent motion, see Figure 10.

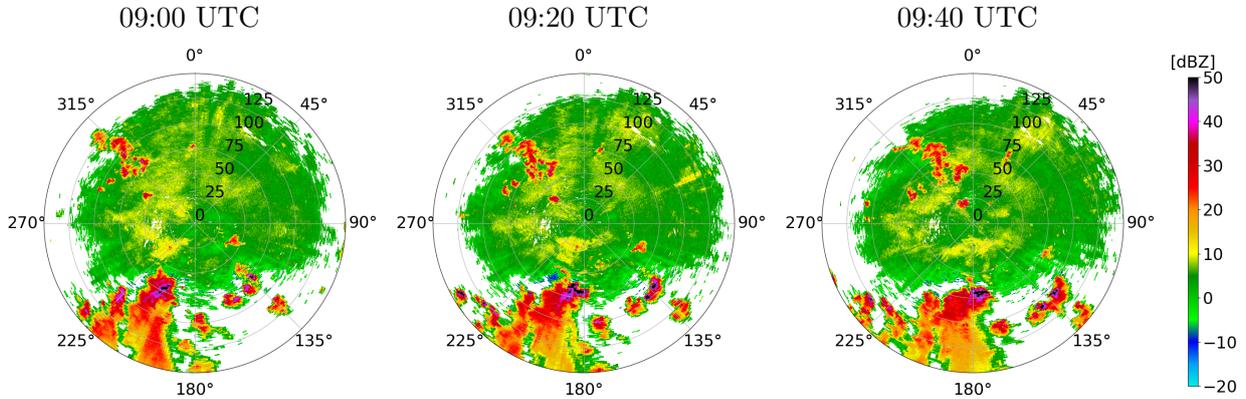


Figure 10: Time series of mixed summer radar images collected at the same radar station at 20-minute intervals.

741 5 Conclusions

742 Accurate identification of non-meteorological echoes in radar images is a prerequisite for reli-
 743 able quantitative precipitation estimation, hydrometeor classification and aviation safety. In this
 744 work, we introduced a deep learning approach using a U-net architecture to segment polarimetric
 745 weather radar data. To overcome the scarcity of pixel-wise labeled training data, we developed
 746 a novel method for generating synthetic training data by combining winter data and cluttered
 747 summer data.

748 The proposed method was evaluated against a state-of-the-art (SotA) method currently used
 749 at Deutscher Wetterdienst. Our results show that the neural network outperforms the SotA
 750 method in preserving weak precipitation and the boundaries of weather events, achieving an
 751 overall accuracy of 0.960 and an intersection over union of 0.923, compared to 0.893 and 0.810
 752 for the SotA method on expert-labeled data.

753 Most importantly, this study demonstrates that synthetically generated radar moments can
 754 be used for training deep neural networks in order to accurately identify non-meteorological
 755 echoes. Despite being trained exclusively on synthetically generated mixed radar data, the net-
 756 work accurately identified non-meteorological echoes in experimentally measured mixed radar
 757 data without any further fine-tuning.

758 To ensure robustness across the entire year, future work should incorporate data from spring
 759 and autumn to capture seasonal variations. Furthermore, the segmentation performance could
 760 be enhanced by extending the neural network input to include temporal information from con-
 761 secutive time steps or volumetric data from multiple elevation angles. Furthermore, the neural
 762 network’s capabilities could be extended beyond 2D spatial features. Utilizing consecutive time
 763 steps or multiple elevation angles as input could provide the time-resolved and volumetric context
 764 needed for even more precise segmentation.

References

- [1] Atanbori, J., Frantzidis, C. A., Al-Khafajiy, M., Aliyu, A., Sohani, B., Appiah, K., Moore, H., Sanders, C., and Ward, A. I. (2025). Learning with noisy labels for classifying biological echoes in polarimetric weather radar observations using artificial neural networks. *Neurocomputing*, 634:129892.
- [2] Bauer, S., Tielens, E. K., and Haest, B. (2024). Monitoring aerial insect biodiversity: A radar perspective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 379:20230113.
- [3] Besic, N., Figueras i Ventura, J., Grazioli, J., Gabella, M., Germann, U., and Berne, A. (2016). Hydrometeor classification through statistical clustering of polarimetric radar measurements: A semi-supervised approach. *Atmospheric Measurement Techniques*, 9:4425–4445.
- [4] Bringi, V. N. and Chandrasekar, V. (2001). *Polarimetric Doppler Weather Radar: Principles and Applications*. Cambridge University Press.
- [5] Chen, J.-Y., Trömel, S., Ryzhkov, A., and Simmer, C. (2021). Assessing the benefits of specific attenuation for quantitative precipitation estimation with a C-band radar network. *Journal of Hydrometeorology*, 22:2617–2631.
- [6] Doviak, R. J. and Zrnić, D. S. (1993). *Doppler Radar and Weather Observations*. Academic Press, 2nd edition.
- [7] Dufton, D. R. L. and Collier, C. G. (2015). Fuzzy logic filtering of radar reflectivity to remove non-meteorological echoes using dual polarization radar moments. *Atmospheric Measurement Techniques*, 8:3985–4000.
- [8] Dumitrescu, D. and Boiangiu, C.-A. (2019). A study of image upsampling and downsampling filters. *Computers*, 8:30.
- [9] Fabry, F. (2015). *Radar Meteorology: Principles and Practice*. Cambridge University Press.
- [10] Frech, M., Hagen, M., and Mammen, T. (2017). Monitoring the absolute calibration of a polarimetric weather radar. *Journal of Atmospheric and Oceanic Technology*, 34(3):599–615.
- [11] Gauthreaux, S. and Diehl, R. (2020). Discrimination of biological scatterers in polarimetric weather radar data: Opportunities and challenges. *Remote Sensing*, 12:545.
- [12] Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- [13] Helmert, K., Tracksdorf, P., Steinert, J., Werner, M., Frech, M., Rathmann, N., Hengstebeck, T., Mott, M., Schumann, S., and Mammen, T. (2014). DWDs new radar network and post-processing algorithm chain. In *8th European Conference on Radar in Meteorology and Hydrology*, volume 4, pages 1–5. https://www.pa.op.dlr.de/erad2014/programme/ExtendedAbstracts/237_Helmert.pdf.
- [14] Hoshen, J. and Kopelman, R. (1976). Percolation and cluster distribution. I. Cluster multiple labeling technique and critical concentration algorithm. *Physical Review B*, 14:3438–3445.
- [15] Hu, C., Sun, Z., Cui, K., Mao, H., Wang, R., Kou, X., Wu, D., and Xia, F. (2024). Classification of biological scatters using polarimetric weather radar. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17:7436–7447.

- 804 [16] Jatau, P., Melnikov, V., and Yu, T.-Y. (2021). A machine learning approach for classifying
805 bird and insect radar echoes with S-band polarimetric weather radar. *Journal of Atmospheric*
806 *and Oceanic Technology*, 38:1797–1812.
- 807 [17] Kern, N., Aguilar, J., Grebner, T., Meinecke, B., and Waldschmidt, C. (2022). Learning on
808 multistatic simulation data for radar-based automotive gesture recognition. *IEEE Transac-*
809 *tions on Microwave Theory and Techniques*, 70:5039–5050.
- 810 [18] Kilambi, A., Fabry, F., and Meunier, V. (2018). A simple and effective method for sepa-
811 rating meteorological from nonmeteorological targets using dual-polarization data. *Journal of*
812 *Atmospheric and Oceanic Technology*, 35:1415–1424.
- 813 [19] Kingma, D. P. and Ba, J. (2017). Adam: A method for stochastic optimization.
814 arXiv:1412.6980.
- 815 [20] Krauss, D., Engel, L., Ott, T., Braunig, J., Richer, R., Gambietz, M., Albrecht, N., Hille, E.,
816 Ullmann, I., Braun, M., Dabrock, P., Kolpin, A., Koelewijn, A., Eskofier, B., and Vossiek, M.
817 (2024). A review and tutorial on machine learning-enabled radar-based biomedical monitoring.
818 *IEEE Open Journal of Engineering in Medicine and Biology*, 6:680–699.
- 819 [21] Leinonen, J., Hamann, U., Nerini, D., Germann, U., and Franch, G. (2023). Latent diffu-
820 sion models for generative precipitation nowcasting with accurate uncertainty quantification.
821 arXiv:2304.12891.
- 822 [22] Mäkinen, T., Ritvanen, J., Pulkkinen, S., Weissaupt, N., and Koistinen, J. (2022). Bayesian
823 classification of nonmeteorological targets in polarimetric doppler radar measurements. *Jour-*
824 *nal of Atmospheric and Oceanic Technology*, 39:1561–1578.
- 825 [23] Mbizvo, G. K. and Larner, A. J. (2024). On the dependence of the critical success index
826 (CSI) on prevalence. *Diagnostics*, 14(5).
- 827 [24] Melnikov, V. and Matrosov, S. Y. (2013). Estimations of aspect ratios of ice cloud particles
828 with the WSR-88D radar. In *36rd AMS Conference on Radar Meteorology*, pages 1–10.
- 829 [25] Michelson, D., Hansen, B., Jacques, D., Lemay, F., and Rodriguez, P. (2020). Monitoring the
830 impacts of weather radar data quality control for quantitative application at the continental
831 scale. *Meteorological Applications*, 27:e1929.
- 832 [26] Pulkkinen, S., Chandrasekar, V., von Lerber, A., and Harri, A.-M. (2020). Nowcasting of
833 convective rainfall using volumetric radar observations. *IEEE Transactions on Geoscience*
834 *and Remote Sensing*, 58(11):7845–7859.
- 835 [27] Radhakrishna, B., Fabry, F., and Kilambi, A. (2019). Fuzzy logic algorithms to identify
836 birds, precipitation, and ground clutter in S-band radar data using polarimetric and nonpo-
837 larimetric variables. *Journal of Atmospheric and Oceanic Technology*, 36:2401–2414.
- 838 [28] Ronneberger, O., Fischer, P., and Brox, T. (2015). U-Net: Convolutional networks for
839 biomedical image segmentation. arXiv:1505.04597.
- 840 [29] Ryzhkov, A. V. and Zrnić, D. S. (2019). *Radar Polarimetry for Weather Observations*.
841 Springer.
- 842 [30] Sauvageot, H. (1992). *Radar Meteorology*. Artech House, Inc.

- 843 [31] Schubert, S., Neubert, P., Pöschmann, J., and Protzel, P. (2019). Circular convolutional
844 neural networks for panoramic images and laser data. In *2019 IEEE Intelligent Vehicles*
845 *Symposium (IV)*, pages 653–660.
- 846 [32] Seliga, T. and Bringi, V. (1976). Potential use of radar differential reflectivity measurements
847 at orthogonal polarizations for measuring precipitation. *Journal of Applied Meteorology*, 15:69–
848 76.
- 849 [33] Straka, J. M., Zrnić, D. S., and Ryzhkov, A. V. (2000). Bulk hydrometeor classification
850 and quantification using polarimetric radar data: Synthesis of relations. *Journal of Applied*
851 *Meteorology*, 39:1341–1372.
- 852 [34] Wilson, J. W., Weckwerth, T. M., Vivekanandan, J., Wakimoto, R. M., and Russell, R. W.
853 (1994). Boundary layer clear-air radar echoes: Origin of echoes and accuracy of derived winds.
854 *Journal of Atmospheric and Oceanic Technology*, 11:1184–1206.
- 855 [35] Zhang, G. (2016). *Weather Radar Polarimetry*. CRC Press.