



1 Enhancing weather radar data by removing 2 non-meteorological echoes, using neural 3 networks trained on synthetic weather data

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

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

7 ^b*Deutscher Wetterdienst, Offenbach, Germany*

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

10 11 Abstract

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

27
28 *Keywords and Phrases:* machine learning, convolutional neural network, synthetic data gen-
29 eration, meteorological radar data, polarimetric radar, non-meteorological echo detection

¹RB and TK contributed equally.

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



30 1 Introduction

31 Since the late 1940s, meteorological radars have become indispensable tools in atmospheric and
32 hydrological research and operational applications, e.g., research on clouds and precipitation
33 formation processes, quantitative precipitation estimation (QPE), motion analysis and precip-
34 itation nowcasting, detection and tracking of convective storms, and radar data assimilation
35 for numerical weather prediction (NWP), see Fabry (2015) for a broad overview. These me-
36 teorological weather radars are ground-based remote sensing instruments that scan a conical
37 volume originating from the radar site with high spatial and temporal resolution. Although this
38 capability makes them highly suitable for atmospheric research, operational weather surveil-
39 lance and forecasting, the data suffers from several sources of error. One particular problem is
40 that radars not only detect meteorological particles, such as droplets, snowflakes, and hail, but
41 also any other scatterers, i.e., objects in the conical volume that scatter electromagnetic waves.
42 This can lead to misinterpretation of the measurements during subsequent analyzes. To address
43 this problem, this paper proposes a novel AI-based method to identify signals corresponding
44 to non-meteorological scatterers based on their spatial and statistical properties measured by
45 meteorological radar products.

46 To explain the problem in more detail, it is useful to briefly summarize the basic measure-
47 ment principle, as described in many textbooks; see, e.g., Doviak and Zrnić (1993); Sauvageot
48 (1992). A short and highly directed pulse of electromagnetic waves is sent out via an antenna in
49 a certain direction, traveling at the speed of light in this direction along a so-called ray. During a
50 subsequent listening period, a time series of returned signals, so-called echoes, from atmospheric
51 scatterers is detected at the radar site. These echoes can be used to compute so-called radar mo-
52 ments, such as strength, phase, and frequency of the detected echoes, which contain information
53 on the amount and nature of scatterers, as well as their radial movement towards or away from
54 the radar site. The time elapsed since sending is a measure of the scatterers radial distance, or
55 range, from the radar site. Thus, the sampling frequency of the time series defines the range
56 resolution (“bin”) and the listening time defines the maximum measurement range. For typical
57 operational applications, the radar antenna is rotated at a constant angular velocity and records
58 data along ray after ray in changing directions. The data from one full azimuthal rotation of
59 the antenna is also referred to as a sweep. If the antenna’s elevation angle is held constant for a
60 sweep, it is called a plan position indicator (PPI) scan. Several successive PPI scans at different
61 elevation angles are called volume scans, as they may sample large parts of the conical volume
62 whose radius is given by the maximum range. Typical spatial resolutions of PPI scans are 250 m
63 to 1 km along the range and 1° in azimuth. Typical volume scans consist of 5 to 20 PPI scans
64 at different elevations, all of which are measured within five minutes.

65 Modern dual-polarization radars are capable of alternating the polarization plane of the
66 electromagnetic radiation from pulse to pulse, commonly from horizontal to vertical and vice
67 versa (Bringi and Chandrasekar, 2001; Ryzhkov and Zrnić, 2019). Most precipitation parti-
68 cles are asymmetric (not spherical) and preferably fall down with their longer axis in horizontal
69 orientation, which causes differences in the returned signal strength and phase between the alter-
70 nating pulses. These differences and their correlations in time, which depend on the hydrometeor
71 shapes and spatial orientations, led to the definition of other radar moments in addition to sig-
72 nal strength. These additional radar moments can be used to improve the echo interpretation
73 with respect to the type and amount of scatterers (Zhang, 2016), as will be further outlined in
74 Section 2.1

75 A problem for correct interpretation is the fact that the type of scatterers is not known
76 a priori, and non-meteorological scatterers like insects or turbulent fluctuations of the air’s



77 refractive index (see Doviak and Zrnić (1993); Wilson et al. (1994)) might be misinterpreted as,
78 e.g., precipitation. Besides non-meteorological echoes (which shall be the focus of the present
79 paper) and attenuation, there are also many other systematic and statistical error sources to be
80 taken into account, e.g., beam bending, beam blockage at obstacles, side lobe echoes, ground
81 clutter, statistical signal fluctuations and radar miscalibration (Doviak and Zrnić, 1993; Fabry,
82 2015).

83 Biological scatterers, such as insects and birds, can provide valuable information to meteoro-
84 logical forecasters, and to ornithologists and entomologists. For example, so-called "fine lines" are
85 sometimes visible in the radar data when thunderstorms expel cooler, denser air. Small airborne
86 particles, such as dust and insects, are swept along the leading edge of this cooler air mass and,
87 together with echoes caused by turbulent fluctuations of the air's refractive index (see below),
88 contribute to the fine line (e.g., Wilson et al. (1994)), giving the forecaster valuable information
89 about the characteristics of the thunderstorm. Radar data can also be used to track bird migra-
90 tions (Hu et al., 2024) and monitor biodiversity (Bauer et al., 2024). However, for meteorological
91 and hydrological applications such as QPE (Chen et al., 2021; Ryzhkov and Zrnić, 2019), precip-
92 itation nowcasting (Pulkkinen et al., 2020; Leinonen et al., 2023) or hydrometeor classification
93 (Besic et al., 2016; Straka et al., 2000), it is essential to separate hydrometeor echoes from any
94 other type of echo. This becomes more and more important because many processing steps are
95 being automated, and human intervention based on individual experience is reduced.

96 Echoes from insects and turbulence in the boundary layer are predominantly observed during
97 the summer. The turbulence echoes are commonly attributed to Bragg scattering, which occurs
98 due to refractive index fluctuations at a length scale of half the radar wavelength (Doviak and
99 Zrnić, 1993; Wilson et al., 1994). The air's refractive index depends on pressure, temperature
100 and humidity, where mainly fluctuations in moisture and, to a lesser extent, temperature cause
101 echoes. These echoes, together with those from insects, are centered around the radar's location
102 because they are limited to lower altitudes within the planetary boundary layer (PBL) and the
103 altitude of the radar ray increases with range. Hence, they appear as annular stationary struc-
104 tures of weak but non-negligible echoes. Radar moments prove to be very useful in distinguishing
105 clutter from meteorological echoes because they provide additional information on the shape and
106 uniformity of the scatterers. However, detecting and eliminating insect and turbulence echoes is
107 still a challenge.

108 Conventional classification methods based on pre-defined thresholds of radar moments (Kil-
109 ambi et al., 2018), Bayesian classifiers (Mäkinen et al., 2022), or fuzzy logic (Dufton and Collier,
110 2015; Radhakrishna et al., 2019) are difficult to configure such that they successfully filter out all
111 echoes from insects in the boundary layer while preserving highly turbulent regions, the edges
112 of true precipitation regions and weak precipitation signals. Especially challenging situations
113 include (i) cases with snow echoes or regions in the melting layer, because they can exhibit
114 polarimetric signatures similar to those of insects, (ii) regions far away from the radar sites,
115 where the signal-to-noise ratio is low, and (iii) "mixed" situations, e.g., summertime convection
116 surrounded by cloud-free regions with insects in the boundary layer. The classifiers mentioned
117 above are also limited by the fact that they can incorporate information only in close proximity
118 to the location of the bin to classify. This means that the classifier loses the overall picture of
119 the annual structure of the insects.

120 Artificial intelligence (AI) methods have been used in recent studies to detect and classify
121 echoes from biological scatterers in radar data for the purposes of biological studies. The fo-
122 cus of these studies was on obtaining information about insects or birds and their subclasses,
123 rather than retaining all meteorologically useful data. For example, in Jatau et al. (2021), a



124 ridge classifier was applied to polarimetric radar data to distinguish between birds and insects,
125 with meteorological signals filtered out before classification. Moreover, random forest approaches
126 based on polarimetric radar data were used to discriminate between different biological scatter-
127 ers (Gauthreaux and Diehl, 2020; Hu et al., 2024). Besides this, in Atanbori et al. (2025), the
128 problem of noisy or uncertain training labels was addressed using a deep learning approach to
129 improve the classification of biological scatterers in radar data.

130 Unlike the above-mentioned conventional approaches that rely on predefined thresholds,
131 Bayesian classifiers, or fuzzy logic, AI-methods can learn intricate patterns directly from data
132 without explicit feature engineering. This adaptability makes them particularly promising for
133 handling the complex and high-dimensional nature of radar data. Convolutional neural networks
134 (CNNs) specifically excel at extracting spatial features from image-like data structures, such as
135 radar sweeps, enabling tasks like object detection and segmentation, see Goodfellow et al. (2016).
136 Among CNNs, the U-net introduced in Ronneberger et al. (2015) is particularly well-suited for
137 segmentation tasks due to its encoder-decoder architecture that captures both local features and
138 global context via skip connections. For further applications of CNNs in radar-based traffic and
139 biomedical monitoring, see Kern et al. (2022); Krauss et al. (2024). In this context, the sweeps
140 can be viewed as multi-channel images, where each pixel corresponds to a radar gate and each
141 image channel corresponds to a radar moment.

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

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

166 2 Materials and methods

167 We now explain how the polarimetric radar data is acquired and represented, and introduce the
168 polarimetric radar moments in more detail, see Section 2.1. Section 2.2 describes the preprocess-



169 ing steps applied to the data. The preprocessed data is then used to generate the input for the
170 neural network, the architecture of which is described in Section 2.3. Section 2.4 explains the
171 synthetic training data generation. The network training itself is described in Section 2.5.

172 2.1 Acquisition of radar data

173 The radar data used in this paper is provided by DWD, which operates a radar network of 17
174 polarimetric weather radars, covering the entire landmass of Germany. Each of these radars sends
175 out horizontally and vertically polarized electromagnetic waves through an antenna and records
176 the power and phase of the returned signal (Ryzhkov and Zrnić, 2019). The recorded signals are
177 acquired in scans, each containing data from one 360° revolution of the antenna. These scans
178 are referred to as sweeps, with the antenna first pointing north and then rotating clockwise.
179 More precisely, every five minutes, each of the radars performs a volume scan consisting of 10
180 sweeps acquired at different elevation angles, as well as a terrain-following sweep (Helmert et al.,
181 2014), which measures precipitation as closely to the ground as possible. In the following, we
182 utilize only these terrain-following sweeps. If the radar view is unobstructed, the elevation angle
183 at which the electromagnetic waves are sent out for such sweeps is set to 0.8°. In cases where
184 orography or obstacles obstruct the view, the elevation is raised accordingly.

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

197 The various radar moments are particularly useful for this purpose of identifying and filtering
198 out non-meteorological echoes, as they provide insight into different qualities of the detected
199 scatterers. The most widely used radar moment is the radar reflectivity factor (Straka et al.,
200 2000; Gauthreaux and Diehl, 2020), which quantifies the strength of the returned signal. It is
201 measured separately at horizontal and vertical polarization, yielding the horizontal and vertical
202 reflectivity factors, denoted by DBZH and DBZV, respectively. Due to variations in raindrop size
203 and number concentration, the radar reflectivity factor can vary by several orders of magnitude.
204 Hence, it is usually expressed in logarithmic units (dBZ), where the minimum detectable value
205 of the reflectivity factor increases with range. For example, at a distance of 1 km from the radar,
206 the minimum detectable reflectivity factor is around -38 dBZ, whereas at ranges beyond 25 km,
207 signals below -10 dBZ cannot be detected by DWD radars (Frech et al., 2017), i.e., they cannot
208 be distinguished from noise.

209 The logarithmic ratio between the horizontal and vertical reflectivity factors (in linear scale)
210 is called differential reflectivity, denoted by ZDR (Seliga and Bringi, 1976). It contains implicit
211 information about the shape of the scatterers and is commonly used to, e.g., discriminate between
212 liquid and frozen hydrometeors. In combination with other radar moments, it is also useful for
213 identifying clutter.



214 Finally, the depolarization ratio, denoted by UDR, is deduced from the differential reflectivity
 215 ZDR and the measurements of the cross-correlation coefficient (ρ_{HV}) (Melnikov and Matrosov,
 216 2013). This radar moment is a measure of how homogeneous and spherical the scatterers are
 217 within a range gate, and thus serves as a good discriminator between clutter and meteorological
 218 echoes (Kilambi et al., 2018; Michelson et al., 2020), with small values ($\ll 0$) indicating weather
 219 and values closer to 0 indicating clutter or biological signals. The DWD's data contains the radar
 220 reflectivity factors, also called radar reflectivities, for horizontal and vertical polarization, the
 221 differential reflectivity, and the depolarization ratio, among various other radar moments. The
 222 reflectivity factors DBZH and DBZV, and the differential reflectivity ZDR have been filtered by
 223 a Doppler filter (Doviak and Zrnić, 1993) aiming to remove any stationary non-meteorological
 224 scatterers, whereas no such filter has been applied to the depolarization ratio UDR. Figure 1
 225 shows examples of the radar moments DBZH, ZDR and UDR obtained from different sweeps.

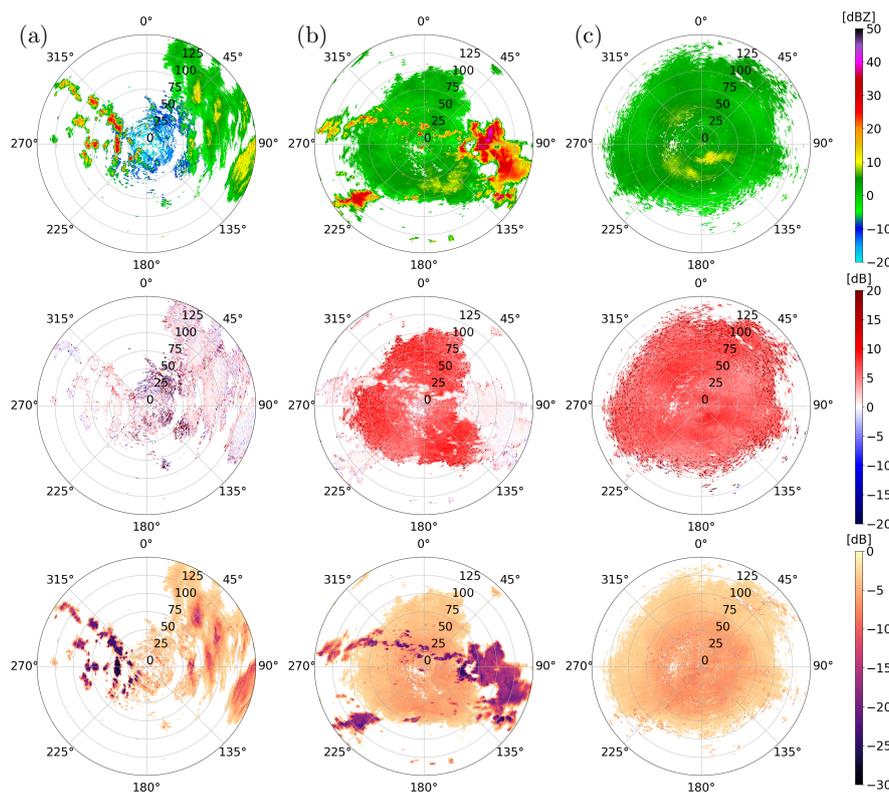


Figure 1: Horizontal reflectivity (first row), differential reflectivity (second row), and depolarization ratio (third row) for examples of (a) a winter sweep, (b) a mixed summer sweep, and (c) a cluttered summer sweep. White pixels are filtered by the initial preprocessing done by DWD.

226 The DWD provided all terrain-following sweeps measured at the 17 different polarimetric
 227 weather radars for three one-hour time periods: 09:00–10:00 UTC on June 16, 2022, 10:00–11:00
 228 UTC on June 18, 2022, and 17:00–18:00 UTC on January 9, 2025. Radar sweeps collected in
 229 January 2025 are referred to as winter sweeps, and radar sweeps collected in June 2022 are
 230 referred to as summer sweeps. Moreover, the radar sweeps collected on June 16 feature both
 231 meteorological and non-meteorological echoes and are referred to as mixed summer sweeps.



232 In contrast, the radar sweeps collected on June 18 were collected during a time period when
 233 no precipitation was measured, indicating that only non-meteorological echoes were detected.
 234 Therefore, these sweeps are referred to as cluttered summer sweeps. For the winter sweeps, the
 235 specific time frame was chosen with the goal of selecting sweeps featuring a minimum amount
 236 of non-meteorological echoes. However, even for the winter sweeps, a certain amount of non-
 237 meteorological echoes cannot be avoided. Figure 1 shows an overview of the horizontal reflectivity
 238 (DBZH), differential reflectivity (ZDR), and depolarization ratio (UDR) for each of the three
 239 data categories.

240 2.2 Preprocessing of radar data

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

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

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

247 Let \mathcal{M} denote the set of radar moments available for a given sweep. Then, for each radar
 248 moment $m \in \mathcal{M}$, its measured values from the sweep can be represented by a single-channel
 249 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
 250 $A_{i,j}^m$ corresponds to the value of the radar moment m at the associated radar gate $g_{i,j}$. Single-
 251 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
 252 $\mathcal{M}' \subseteq \mathcal{M}$ is the subset of included radar moments and $C = |\mathcal{M}'|$ is the number of channels. In the
 253 following, we refer to these multi-channel images as radar images. Moreover, the provided data
 254 also contains information on which radar gates had already been identified as non-meteorological
 255 during the preprocessing performed by DWD. For these radar gates, the corresponding entries
 256 of the DBZH image are assigned a placeholder value $p = -40$ dBZ, which lies far below the
 257 reflectivity range of meteorological echoes. Pixels with a horizontal reflectivity value of -40 dBZ
 258 are referred to as excluded pixels, as they conclusively do not correspond to meteorological
 259 echoes. All other pixels are referred to as included pixels. A Cartesian representation of a DBZH
 260 image is shown in Figure 2a. Most of the radar images still contain pixels that likely correspond
 261 to non-meteorological echoes, even for radar images corresponding to winter sweeps. To filter
 262 some of them out, we apply further preprocessing steps.

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

270 Moreover, meteorological echoes observed in the radar images typically form connected com-
 271 ponents consisting of several pixels, see Figure 1. In contrast, isolated included pixels are as-
 272 sumed to correspond to non-meteorological echoes. For this reason, each DBZH image is further
 273 processed by a despeckle filter consisting of two steps. First, we use the Hoshen-Kopelman



274 algorithm (Hoshen and Kopelman, 1976) to identify all connected components of included pix-
275 els. More specifically, we apply the algorithm to the binarized image obtained by threshold-
276 ing at -40 dBZ, separating the excluded pixels from the included ones. Here, we use the 4-
277 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 neigh-
278 boring if their associated radar gates $g_{i_1, j_1}, g_{i_2, j_2}$ share a common edge. Note that the DBZH image
279 is periodic in the second component, see Eq. (1), since each index $j \in \{1, \dots, n_a\}$ represents the
280 azimuthal angle of a ray, and the image corresponds to a sweep. Consequently, pixels correspond-
281 ing to the first and last ray are neighboring if they have the same range index $i \in \{1, \dots, n_r\}$.
282 More precisely, for each pixel $(i, j) \in \Omega$, the set $N_{(i, j)} \subset \Omega$ of neighboring pixels can be written
283 as $N_{(i, j)} = \{(k, \ell) \in \Omega : |i - k| + d(j, \ell) = 1\}$, where $d: \{1, \dots, n_a\} \times \{1, \dots, n_a\} \rightarrow \mathbb{N}_0$ denotes
284 the periodic distance measure, given by

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

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

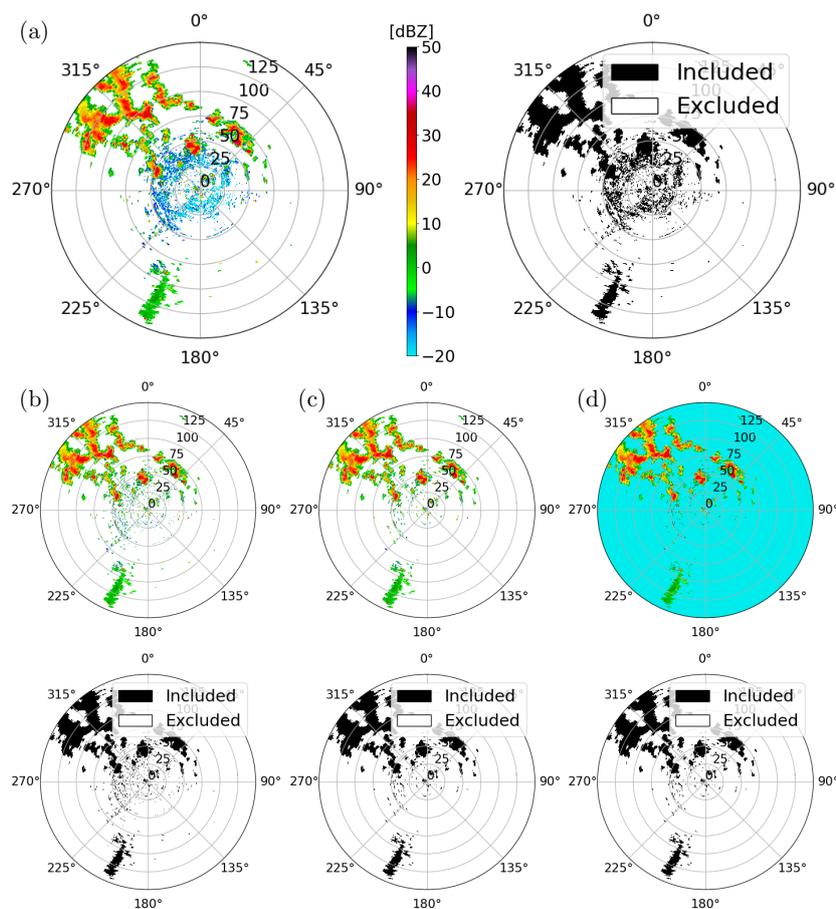


Figure 2: Cartesian representations of DBZH images: (a) DBZH image of a winter sweep, as provided by DWD, (b) after thresholding, (c) after despeckling, and (d) final DBZH image with placeholder values. White pixels in the DBZH images are either filtered out during the preprocessing by DWD or are assigned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

289 For the excluded pixels $(i, j) \in \Omega$, their corresponding entries in the DBZH image were set
 290 to -40 dBZ to indicate that they conclusively do not correspond to meteorological echoes, see
 291 Figure 2d. For the other radar moments, ZDR and UDR, we also assign specific placeholder
 292 values $p \in \mathbb{R}$ to the entries of their single-channel images. Similar to the horizontal reflectivity,
 293 these placeholder values are chosen to lie outside the range of typical values for meteorological
 294 echoes in our dataset. More specifically, for the differential reflectivity, the placeholder value is set
 295 to $p = 20$ dB. This particular value was chosen since it lies above typical values for meteorological
 296 echoes (around 0 dB), but within the range for non-meteorological echoes (between 5 and 20 dB)
 297 in our data, as can be seen in Figures 1b and 1c. Furthermore, as described in Kilambi et al.
 298 (2018), except for some special cases such as hail, the depolarization ratio of meteorological
 299 scatterers is typically below -12 dB. In contrast, the depolarization ratio of non-meteorological
 300 scatterers typically ranges between -12 and 0 dB. Therefore, we chose a placeholder value of
 301 $p = 0$ dB for the depolarization ratio.

302 The presented preprocessing steps remove most of the clutter with low reflectivity values.



303 However, many mixed summer images contain much larger connected components corresponding
304 to more reflective non-meteorological scatterers, as can be seen in Figure 1b. Detecting such
305 echoes presents a challenging task for conventional methods such as fuzzy logic filtering (Dufton
306 and Collier, 2015) or depolarization ratio thresholding (Kilambi et al., 2018).

307 **2.3 Description of the network architecture**

308 The task of detecting non-meteorological echoes can be viewed as a segmentation task of multi-
309 channel two-dimensional (2D) radar image data. Each pixel of the image is classified as belonging
310 to meteorological or non-meteorological echoes, that is, each pixel is assigned a binary class label
311 $c \in \{0, 1\}$. In Ronneberger et al. (2015), the so-called U-net architecture, a fully convolutional
312 neural network, has been introduced to address similar segmentation tasks of 2D images.

313 A U-net has an encoder-decoder architecture with skip connections, consisting of a so-called
314 contracting path that downsamples the spatial resolution of the features and an expansive path
315 that upsamples it again. Both the downsampling and the upsampling paths consist of several
316 levels. The network considered in the present paper is an adapted version of the U-net architec-
317 ture comprising five levels, with the fifth level acting as a bottleneck. Consequently, the network
318 performs four downsampling steps and four upsampling steps. Each downsampling step consists
319 of two convolution blocks and a 2×2 max-pooling layer for downsampling the spatial resolu-
320 tion of the feature map (Goodfellow et al., 2016). A convolution block contains a padded 3×3
321 convolutional layer, followed by a batch normalization layer and a rectified linear unit (ReLU)
322 activation function, which is given by $\text{ReLU}: \mathbb{R} \rightarrow \mathbb{R}_0^+ = [0, \infty)$, $\text{ReLU}(x) = \max(x, 0)$. The bot-
323 tleneck at the fifth layer consists of two such convolution blocks. For each of the subsequent
324 upsampling steps, the spatial resolution of the feature channels is increased using bilinear inter-
325 polation (Dumitrescu and Boiangiu, 2019). The resulting channels are then concatenated with
326 the feature channels of the corresponding level of the downsampling path. Two convolutional
327 blocks following the concatenation of those feature channels complete the upsampling step. After
328 the final upsampling step, the sigmoid activation function $\sigma: \mathbb{R} \rightarrow [0, 1]$, given by

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

329 is applied to the output.

330 Compared to the original U-net architecture presented in Ronneberger et al. (2015), we
331 use padded convolutional layers with a padding of 1 in each step, preserving the image size
332 throughout the network. Since the network input is periodic in the second component, we use
333 periodic padding (Schubert et al., 2019) in this direction, while the padding in the direction of
334 the first component is reflective. Moreover, the spatial resolution of the features is upsampled
335 using bilinear interpolation instead of an up-convolution. Furthermore, the network does not
336 include a final 1×1 convolution. Instead, the output of the final convolution block is already a
337 single-channel map to which the sigmoid activation function is applied directly. A sketch of the
338 network architecture described above is shown in Figure 3.

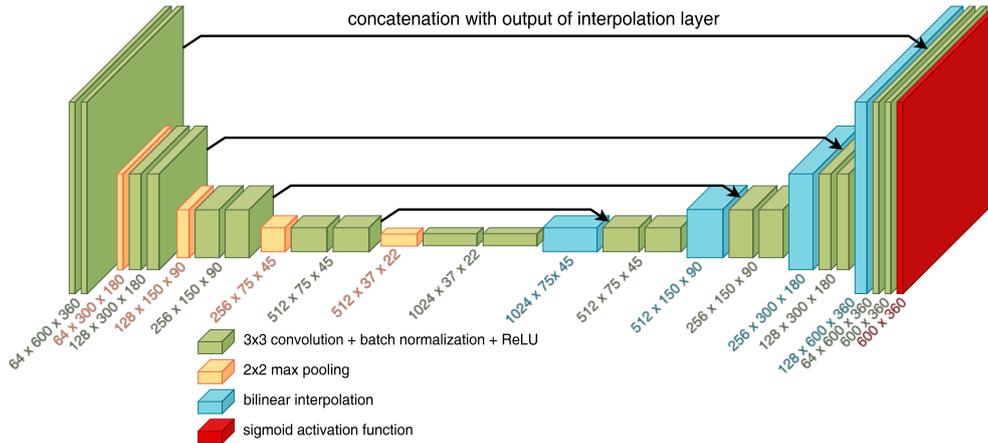


Figure 3: Schematic representation of the network architecture. Below each layer, the respective output dimensions are given.

339 The network outputs are images of the same spatial dimensions as each of the input channels.
 340 More specifically, the network architecture stated above with parameter constellation $\theta \in \Theta \subset \mathbb{R}^q$
 341 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
 342 of input channels $k = 3$. Here, θ denotes the vector of trainable parameters of the network, and
 343 $\Theta \subset \mathbb{R}^q$ is the network's parameter space for some $q \in \mathbb{N}$. Each input channel corresponds to
 344 a single-channel image of a radar moment, as described in Section 2.1, allowing the network to
 345 combine information about different radar moments to detect non-meteorological echoes.

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

354 2.4 Generation of synthetic training data

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



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

368 More specifically, we generate a realistic mixed (synthetic) DBZH image by combining two
 369 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
 370 following manner. Since the horizontal reflectivity is measured in decibels, which is a logarithmic
 371 unit of measurement, we first transform the entries of the DBZH images into linear scale using
 372 the transformation $\Phi: \mathbb{R} \rightarrow \mathbb{R}^+ = (0, \infty)$ (Ryzhkov and Zrnić, 2019), which is given by

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

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

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

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

378 Similarly to the placeholder values described in Section 2.2, we assign the entry of the syn-
 379 thetic DBZH image $A_{i,j}^{H,m}$ a placeholder value $p = -40$ dBZ if the pixel is excluded for both
 380 radar images. To find appropriate weighting factors λ_w and λ_s , we consider the difference in
 381 horizontal reflectivity values observed in Figure 1. Since the range of the horizontal reflectiv-
 382 ity values corresponding to non-meteorological echoes is almost identical for cluttered summer
 383 radar images (Figure 1c) and mixed summer radar images (Figure 1b), the choice of $\lambda_s = 1$
 384 is suitable. However, the horizontal reflectivity values corresponding to meteorological echoes
 385 are significantly lower for winter radar images (Figure 1a) compared to mixed summer radar
 386 images. Since the resulting synthetic DBZH image should resemble the typical DBZH images
 387 from mixed summer radar images, a scaling factor $\lambda_w > 1$ is necessary. In order to determine a
 388 suitable choice, four representative mixed summer radar images were chosen. For each of them,
 389 a winter radar image with precipitation areas of similar size was assigned. Then, for different
 390 weighting factors $\lambda_w \in \{1.5, 3, 5, 8, 15\}$, we compared the resulting synthetic DBZH image with
 391 that of the respective mixed summer radar image. As a result, a weighting factor λ_w between
 392 5 and 8 was deemed appropriate. Figure 4 shows how a synthetic DBZH image emerges from
 393 a winter and a cluttered summer radar image using a weighting factor $\lambda_w = 6.5$. The resulting
 394 synthetic DBZH image (Figure 4c) is then compared with one from a mixed summer radar image
 395 (Figure 4d).

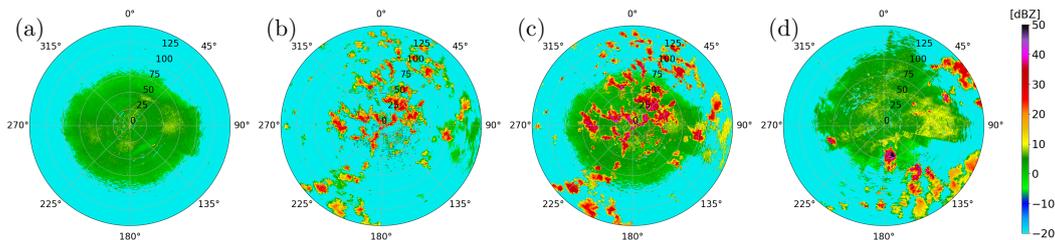


Figure 4: Comparison of Cartesian representations of different DBZH images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image with weighting factors $\lambda_w = 6.5$, and $\lambda_s = 1$ and (d) mixed summer radar image.



396 The two chosen winter and cluttered summer radar images are then also used to generate syn-
 397 thetic single-channel images for the other radar moments considered in Section 2.1. In addition
 398 to the horizontal reflectivity, which is proportional to the particle concentration and the parti-
 399 cle’s horizontal cross-section within the radar gate, the differential reflectivity (ZDR) primarily
 400 reflects the mean axis ratio of the particles (Straka et al., 2000). Since differential reflectivity is
 401 defined as the difference between horizontal and vertical reflectivity in decibels, synthetic values
 402 can be generated by taking the difference between synthetic horizontal and vertical reflectivity
 403 values. Similarly to the synthetic DBZH image $A^{H,m}$, a synthetic DBZV image $A^{V,m} \in \mathbb{R}^{n_r \times n_a}$
 404 is also generated. Then, the corresponding synthetic ZDR image is denoted by $A^{ZDR,m} \in \mathbb{R}^{n_r \times n_a}$,
 405 and its entries are given by

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

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

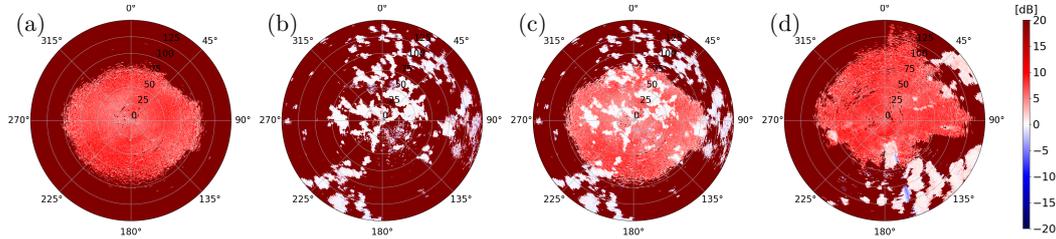


Figure 5: Comparison of Cartesian representations of different ZDR images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image, and (d) mixed summer radar image.

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

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

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

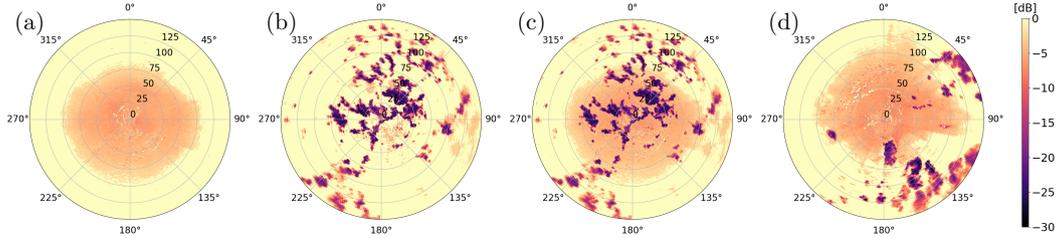


Figure 6: Comparison of Cartesian representations of different UDR images: (a) cluttered summer radar image, (b) winter radar image, (c) their corresponding synthetic radar image, and (d) mixed summer radar image.

419 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
 420 a corresponding ground truth image $G \in \{0, 1\}^{n_r \times n_a}$. Its values indicate which pixels of the
 421 synthetic radar image correspond to meteorological ($G_{i,j} = 0$) and non-meteorological echoes
 422 ($G_{i,j} = 1$), respectively. Pixels included in the winter radar image $A^{H,w}$ are assumed to correspond
 423 to meteorological echoes, as described in Section 2.2. Since the synthetic image is generated in
 424 such a way that these pixels also correspond to meteorological echoes in the synthetic image,
 425 they are assigned a ground truth label of 0. All of the other pixels are labeled 1, as they are
 426 excluded from the winter radar image. More precisely, for each pixel $(i, j) \in \Omega$, the ground truth
 427 label $G_{i,j}$ is given by

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

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

432 2.5 Network training

433 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
 434 $A_n^w = (A_n^{H,w}, A_n^{V,w}, A_n^{UDR,w})$, where the radar moments present in the data are horizontal reflectivity,
 435 vertical reflectivity, and the depolarization ratio. Similarly, the available cluttered summer
 436 dataset is given by $\mathcal{D}^s = \{A_n^s\}_{n=1}^{180} \subset D$. We partition each of these datasets into two disjoint
 437 subsets. One subset of radar images is used for generating the training samples, while the other
 438 set is reserved exclusively for network evaluation. For this purpose, we partition the index set
 439 $\mathcal{I} = \{1, \dots, 180\}$ into two disjoint subsets $\mathcal{I}_{\text{train}}, \mathcal{I}_{\text{test}} \subset \mathcal{I}$, i.e.,

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

440 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
 441 from a single radar station are used for network evaluation, while the radar images from other
 442 radar stations are used for the generation of training samples.

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



448 Before generating the synthetic training samples, we apply various data augmentation steps
 449 to increase their diversity. More precisely, as described in Section 2.4, the DBZH and DBZV
 450 images from the winter and cluttered summer datasets are scaled by reflectivity weighting factors
 451 $\lambda_w, \lambda_s > 0$, respectively. In addition, we apply a rotation and a possible inversion of orientation.

452 Specifically, we introduce a parametric mapping $g_{\psi, \psi'}: D \times D \rightarrow D \times \{0, 1\}^{n_r \times n_a}$
 453 that combines two radar images to generate a synthetic radar image and a corresponding
 454 ground truth image. The space of admissible parameters for data augmentation is given by
 455 $\Psi = \mathbb{R}^+ \times \{1, \dots, n_a\} \times \{-1, 1\}$. Each parameter vector $\psi = (\lambda, \varphi, o) \in \Psi$ contains the reflectivity
 456 weighting factor $\lambda > 0$, the rotation angle $\varphi \in \{1, \dots, n_a\}$, and the orientation flag
 457 $o \in \{-1, 1\}$. Here, a rotation angle of $\varphi \in \{1, \dots, n_a\}$ indicates a rotation by φ rays. Hence,
 458 a shift of φ rows is applied to the radar image. An orientation flag of $o = -1$ indicates an
 459 inverted orientation. In that case, the radar image is flipped along the x -axis. To distinguish the
 460 transformations we apply to winter and cluttered summer radar images, we introduce restricted
 461 admissible parameter spaces $\Psi^w, \Psi^s \subset \Psi$. They only differ in the admissible range of the reflectivity
 462 weighting factor, which is given by the interval $[5, 8]$ for Ψ^w and by $[0.9, 1.1]$ for Ψ^s . The
 463 set of synthetic training samples generated by combining two radar images is then denoted by

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

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

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

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

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

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

478 As stated in Section 2.3, the network can be described by its vector of trainable parameters
 479 $\theta \in \Theta$ 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
 480 dimensions of a single-channel image, and $k = 3$ is the number of input channels. Then, for a
 481 given radar image $A \in D$, the output $f_\theta(A)$ of the network represents the matrix of probabilities
 482 $(f_\theta(A))_{i,j}$ that a pixel (i, j) corresponds to a non-meteorological echo for each $(i, j) \in \Omega$.

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



489 $\mathcal{D}_{\text{train}}^w$ or $\mathcal{D}_{\text{train}}^s$, with a probability of 0.125 for each case. For each training sample (A_n, G_n)
 490 of the mini-batch B , with $n \in \{1, 2, 3, 4\}$, the network output $f_\theta(A_n)$ is computed and com-
 491 pared to the ground truth image G_n using the binary cross-entropy loss, which is a measure of
 492 how well the network outputs match the ground truth images. More specifically, for any batch
 493 $B = \{(A_n, G_n)\}_{n=1}^4$ and vector of trainable network parameters $\theta \in \Theta$, the cross-entropy loss
 494 $L_\theta(B)$ of the batch is given by

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

495 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
 496 of $f_\theta(A_n)$ and G_n , respectively. After computing the cross-entropy loss $L_\theta(B)$, its gradient is
 497 computed with respect to the vector of trainable network parameters θ . Based on that, this
 498 parameter vector θ is then adjusted using the Adam optimizer (Kingma and Ba, 2017). Further-
 499 more, the performance of the trained network is compared to that of the current state-of-the-art
 500 classification method used at DWD, see Section 3.

501 3 Results

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

508 3.1 Metrics for quantitative evaluation

509 To quantitatively evaluate the segmentation performance, we compare the network output
 510 against corresponding ground truth. For that purpose, pixels $(i, j) \in \Omega$ with a horizontal re-
 511 flectivity value of $A_{i,j}^H = -40$ dBZ are excluded from the evaluation, as they conclusively do
 512 not correspond to meteorological echoes. Furthermore, the physical area of a radar gate $g_{i,j}$ in-
 513 creases linearly with its distance from the radar. Therefore, to obtain metrics with a meaningful
 514 physical interpretation regarding the spatial extent of the echoes, we introduce a weight matrix
 515 $W = (W_{i,j}) \in \mathbb{R}^{n_r \times n_a}$. For each pixel $(i, j) \in \Omega$, its weight $W_{i,j}$ is given by

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

516 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
 517 range resolution, and $r_i = id$ is the outer radius of the radar gate. Using the weights of the
 518 individual pixels, we define the weighted area $\mathcal{A}(\Omega')$ of any subset of pixels $\Omega' \subset \Omega$ as the sum
 519 of the weights of the pixels in that set, given by

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



520 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
 521 transformed into a binary prediction $P = (P_{i,j}) \in \{0, 1\}^{n_r \times n_a}$ by thresholding at 0.5. Specifically,
 522 the prediction $P_{i,j}$ at pixel $(i, j) \in \Omega$ is given by

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

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

525 To quantify segmentation performance, we compute the area-weighted intersection over union
 526 (IOU), also referred to as the critical success index (Mbizvo and Larner, 2024). For class $c \in$
 527 $\{0, 1\}$, the class-specific IOU is defined as the ratio of the area of the intersection of pixels
 528 classified as class c in the prediction and pixels labeled as class c in the ground truth, to the
 529 area of their union, given by

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

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

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

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

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

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

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

540 where $\Omega|_{P=G} = \{(i, j) \in \Omega : P_{i,j} = G_{i,j}\}$ denotes the set of all correctly classified pixels. Setting
 541 $\Omega' = \Omega$ yields the area-weighted accuracy for the entire radar image.

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

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



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

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

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

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

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

564 Using these subsets of pixels, we compute the area-weighted accuracy depending on reflectivity
 565 and distance from the radar, given by

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

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

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

568 respectively.

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

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

574 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
 575 image. As before, setting $\Omega'_n = \Omega$ for all $n \in \{1, \dots, N\}$ yields the area-weighted accuracy for
 576 the dataset \mathcal{D} , while the evaluation on $\Omega'_n = \Omega|_{P_n=c}$ and $\Omega'_n = \Omega|_{G_n=c}$ for all $n \in \{1, \dots, N\}$
 577 and $c \in \{0, 1\}$ yields the class-specific precision and recall for the dataset \mathcal{D} , respectively.
 578 Analogously, the intersection over union IOU for the dataset \mathcal{D} is given by

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



579 3.2 State-of-the-art method

580 The SotA method used by DWD comprises a series of processing steps designed to identify and
 581 filter non-meteorological echoes, as well as a correction for attenuation, as described in Helmert
 582 et al. (2014). These processing steps include a threshold filter, a filter for spokes and rings, a
 583 blacklist filter, a polarimetric clutter filter based on fuzzy logic, a second-trip filter, and a speckle
 584 filter. Notably, the polarimetric clutter filter has been extended in 2025 to include an additional
 585 so-called insect filter, which is used operationally during the summer months (April – October).

586 Echoes from insects, which are referred to as insect echoes, are typically characterized by
 587 the following features in the radar moments: reduced cross-correlation coefficient ρ_{HV} (< 0.9),
 588 increased differential reflectivity ZDR (> 2 dB), and low reflectivity factors DBZH and DBZV
 589 (< 20 dBZ). In principle, pixels corresponding to insect echoes can be correctly classified based
 590 on these features. However, there are also weather events, such as snowfall, that exhibit similar
 591 features, particularly at greater distances from the radar. Moreover, pixels corresponding to
 592 a mixture of meteorological echoes and clutter can also exhibit these features but should be
 593 classified as meteorological echoes.

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

600 More precisely, the newly implemented insect filter operates in two steps: First, all pixels
 601 eligible for filtering are identified. Then, for each pixel identified in the first step, it is tested
 602 whether the pixel’s UDR value, or the UDR value in its local neighborhood, allows for filtering.
 603 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
 604 $q \in \{1, 3, 5, \dots\}$ is considered, which is given by

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

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

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

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



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

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

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

636 3.3 Network evaluation

637 We evaluate the trained network using the metrics stated in Section 3.1 and compare its perfor-
 638 mance to the SotA method detailed in Section 3.2. Segmentation performance is evaluated on
 639 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
 640 in Section 2.5. Moreover, both methods are also evaluated on a non-synthetic mixed dataset
 641 $\mathcal{D}_{\text{exp}}^m$, given by

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

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

645 Table 1 shows a comparison of the resulting metrics averaged over the respective datasets as
 646 described in Section 3.1. The overall performance metrics on the test datasets are slightly lower
 647 than those on the training datasets. The accuracy ACC and intersection over union IOU on the
 648 mixed test dataset $\mathcal{D}_{\text{test}}^m$ are 0.964 and 0.932, respectively, compared to 0.982 and 0.966 on the
 649 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



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

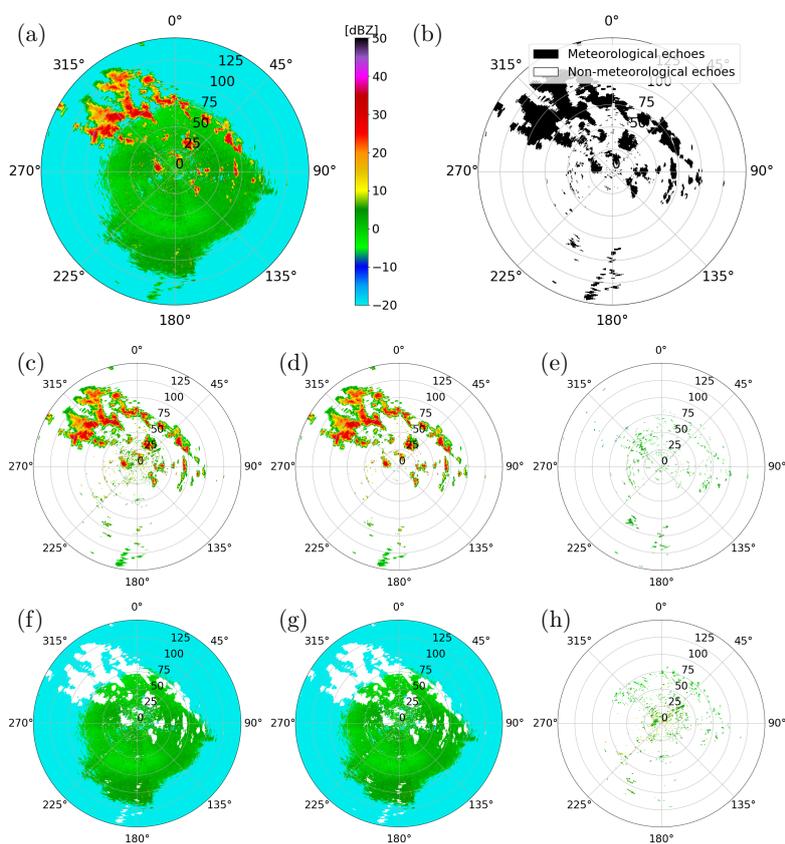


Figure 7: Evaluation of the network on a synthetic test data: (a) input data and (b) corresponding ground truth, (c) pixels classified as meteorological echoes by the network, (d) correctly classified meteorological echoes, and (e) meteorological echoes misclassified as non-meteorological echoes. The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

657 When evaluated on the experimentally measured mixed dataset $\mathcal{D}_{\text{exp}}^m$, the network perfor-
 658 mance is comparable, with accuracy ACC and intersection over union IOU of 0.960 and 0.923,
 659 respectively. However, when evaluating the SotA method on the same mixed dataset $\mathcal{D}_{\text{exp}}^m$, it
 660 performs significantly worse, with accuracy ACC and intersection over union IOU of 0.893 and
 661 0.810, respectively. These overall differences are caused by larger differences in precision PRE^1
 662 for non-meteorological echoes and recall REC^0 for meteorological echoes, where the SotA method
 663 achieves scores of 0.775 and 0.821, respectively, compared to the neural network with scores of



664 0.969 and 0.977. The SotA method achieves better scores for recall REC^1 for non-meteorological
 665 echoes and precision PRE^0 for meteorological echoes, with values of 0.983 and 0.988, respectively,
 666 compared to the neural network with scores of 0.906 and 0.957.

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

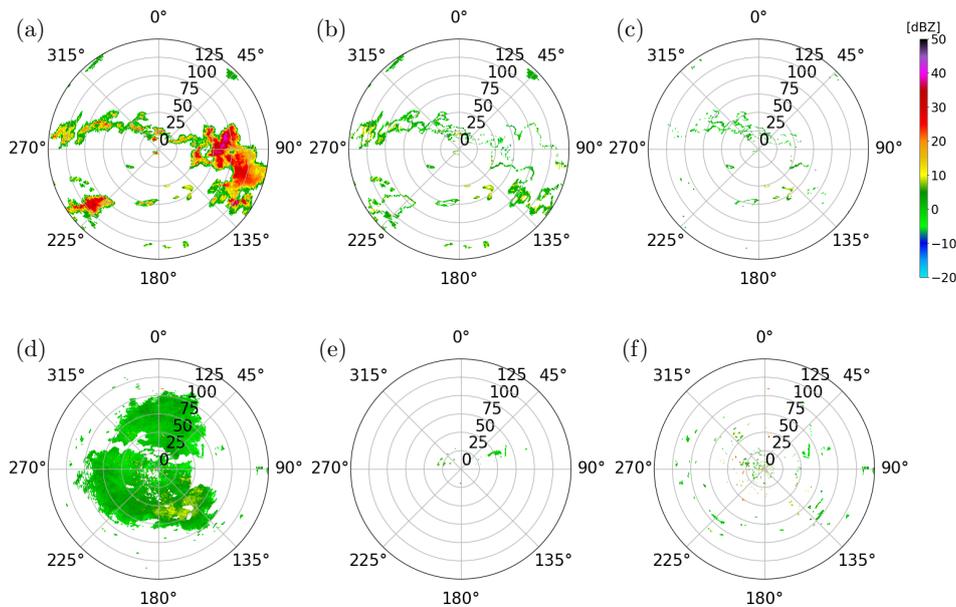


Figure 8: Evaluation of the SotA method and the trained network on the experimentally measured radar image shown in Figure 1b: The first row shows (a) the ground truth for the meteorological echoes and the meteorological echoes misclassified as non-meteorological echoes by (b) the SotA method and (c)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

674 In order to further investigate this, Figure 9a shows the accuracy depending on the dis-
 675 tance, while Figures 9b and 9c show the class-specific recall scores for meteorological and non-
 676 meteorological echoes, respectively. The networks accuracy ACC^D on the experimentally mea-
 677 sured mixed dataset \mathcal{D}_{exp}^m increases with distance from the radar, rising from below 0.9 within
 678 25 km to nearly 1.0 beyond 75 km. The same trend can be observed for the test dataset \mathcal{D}_{test}^m ,
 679 for which only slightly higher scores are achieved by the network. When compared to the SotA
 680 method, the recall $REC^{0,D}$ for meteorological echoes depending on distance is higher for the net-
 681 work than for the SotA method for all distances, and the opposite is true for non-meteorological
 682 echoes. The most significant differences occur close to the radar, at a distance of less than
 683 25 km, where the neural network achieves a recall $REC^{1,D}$ of 0.877, compared to that of the
 684 SotA method with 0.981. This relatively lower recall close to the radar quantifies the errors we



685 see in Figure 8c, where a significant number of non-meteorological echoes near the radar are
 686 incorrectly classified. When considering the recall $REC^{0,D}$ for meteorological echoes depending
 687 on distance, the network clearly outperforms the SotA method at distances further from the
 688 radar, with a recall value of 0.989 compared to that of the SotA method with 0.831 beyond
 689 75 km. The echoes further away from the radar represent a proportionally larger area, as shown
 690 in the histograms, and thus the network outperforms the SotA method overall.

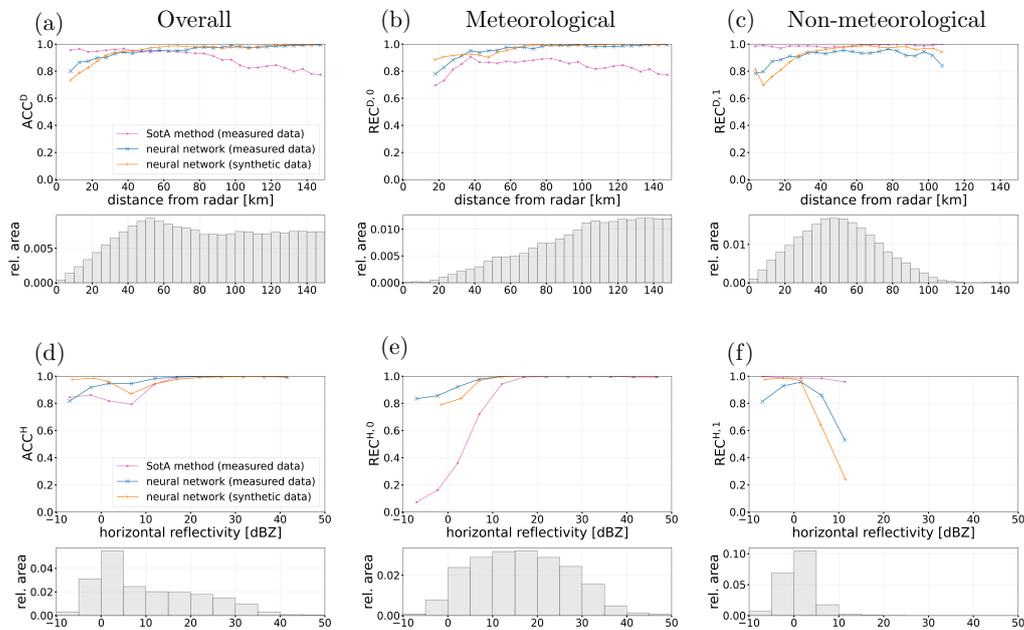


Figure 9: First row: Accuracy in dependence on the distance from the radar for (a) meteorological and non-meteorological echoes combined, (b) meteorological echoes, and (c) non-meteorological echoes, as classified in the ground truth. Second row: Accuracy in dependence on the horizontal reflectivity 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}_{exp}^m and the neural network was evaluated on the same experimentally measured dataset and additionally on the mixed test dataset \mathcal{D}_{test}^m .

691 Analogously, we evaluate the accuracy and class-specific recall in dependence on the horizontal
 692 reflectivity values, as shown in Figures 9d to 9f. For meteorological echoes with a reflectivity
 693 higher than 10 dBZ, both the network and the SotA method achieve a recall $REC^{0,H}$ near 1.0.
 694 However, for horizontal reflectivities below 10 dBZ, the SotA performance degrades significantly,
 695 whereas the network maintains a recall of 0.936. While meteorological echoes with reflectivity
 696 values below 0 dBZ are rare in the dataset, those with reflectivity values between 0 and 10 dBZ are
 697 much more frequent, as shown in the histogram of Figure 9e. For example, the edges of weather
 698 events are typically less reflective and are often incorrectly classified as non-meteorological echoes
 699 by the SotA method, as shown in Figure 8b. In the case of non-meteorological echoes, the network
 700 struggles to classify echoes with high reflectivity (> 10 dBZ), achieving a lower recall $REC^{1,H}$ of
 701 0.430 compared to the SotA method, which achieves a recall value of 0.961. Nevertheless, such



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

703 4 Discussion

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

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

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

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

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



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

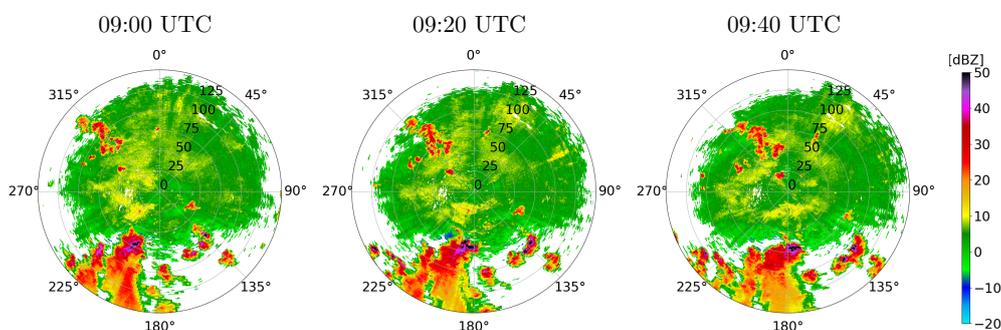


Figure 10: Time series of mixed summer radar images collected at the same radar station at 20-minute intervals.

753 5 Conclusions

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

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

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

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



777 References

- 778 Atanbori, J., Frantzidis, C. A., Al-Khafajiy, M., Aliyu, A., Sohani, B., Appiah, K., Moore, H.,
779 Sanders, C., and Ward, A. I. (2025). Learning with noisy labels for classifying biological echoes
780 in polarimetric weather radar observations using artificial neural networks. *Neurocomputing*,
781 634:129892.
- 782 Bauer, S., Tielens, E. K., and Haest, B. (2024). Monitoring aerial insect biodiversity: A radar per-
783 spective. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 379:20230113.
- 784 Besic, N., Figueras i Ventura, J., Grazioli, J., Gabella, M., Germann, U., and Berne, A. (2016).
785 Hydrometeor classification through statistical clustering of polarimetric radar measurements:
786 A semi-supervised approach. *Atmospheric Measurement Techniques*, 9:4425–4445.
- 787 Bringi, V. N. and Chandrasekar, V. (2001). *Polarimetric Doppler Weather Radar: Principles*
788 *and Applications*. Cambridge University Press.
- 789 Chen, J.-Y., Trömel, S., Ryzhkov, A., and Simmer, C. (2021). Assessing the benefits of specific
790 attenuation for quantitative precipitation estimation with a C-band radar network. *Journal*
791 *of Hydrometeorology*, 22:2617–2631.
- 792 Doviak, R. J. and Zrnić, D. S. (1993). *Doppler Radar and Weather Observations*. Academic
793 Press, 2nd edition.
- 794 Dufton, D. R. L. and Collier, C. G. (2015). Fuzzy logic filtering of radar reflectivity to remove
795 non-meteorological echoes using dual polarization radar moments. *Atmospheric Measurement*
796 *Techniques*, 8:3985–4000.
- 797 Dumitrescu, D. and Boiangiu, C.-A. (2019). A study of image upsampling and downsampling
798 filters. *Computers*, 8:30.
- 799 Fabry, F. (2015). *Radar Meteorology: Principles and Practice*. Cambridge University Press.
- 800 Frech, M., Hagen, M., and Mammen, T. (2017). Monitoring the absolute calibration of a polari-
801 metric weather radar. *Journal of Atmospheric and Oceanic Technology*, 34(3):599–615.
- 802 Gauthreaux, S. and Diehl, R. (2020). Discrimination of biological scatterers in polarimetric
803 weather radar data: Opportunities and challenges. *Remote Sensing*, 12:545.
- 804 Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press.
- 805 Helmert, K., Tracksdorf, P., Steinert, J., Werner, M., Frech, M., Rathmann, N., Hengste-
806 beck, T., Mott, M., Schumann, S., and Mammen, T. (2014). DWDs new radar network
807 and post-processing algorithm chain. In *8th European Conference on Radar in Meteorology*
808 *and Hydrology*, volume 4, pages 1–5. [https://www.pa.op.dlr.de/erad2014/programme/
809 ExtendedAbstracts/237_Helmert.pdf](https://www.pa.op.dlr.de/erad2014/programme/ExtendedAbstracts/237_Helmert.pdf).
- 810 Hoshen, J. and Kopelman, R. (1976). Percolation and cluster distribution. I. Cluster multiple
811 labeling technique and critical concentration algorithm. *Physical Review B*, 14:3438–3445.
- 812 Hu, C., Sun, Z., Cui, K., Mao, H., Wang, R., Kou, X., Wu, D., and Xia, F. (2024). Classification
813 of biological scatters using polarimetric weather radar. *IEEE Journal of Selected Topics in*
814 *Applied Earth Observations and Remote Sensing*, 17:7436–7447.



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



- 852 Schubert, S., Neubert, P., Pöschmann, J., and Protzel, P. (2019). Circular convolutional neural
853 networks for panoramic images and laser data. In *2019 IEEE Intelligent Vehicles Symposium*
854 *(IV)*, pages 653–660.
- 855 Seliga, T. and Bringi, V. (1976). Potential use of radar differential reflectivity measurements at
856 orthogonal polarizations for measuring precipitation. *Journal of Applied Meteorology*, 15:69–
857 76.
- 858 Straka, J. M., Zrnić, D. S., and Ryzhkov, A. V. (2000). Bulk hydrometeor classification and
859 quantification using polarimetric radar data: Synthesis of relations. *Journal of Applied Mete-*
860 *orology*, 39:1341–1372.
- 861 Wilson, J. W., Weckwerth, T. M., Vivekanandan, J., Wakimoto, R. M., and Russell, R. W.
862 (1994). Boundary layer clear-air radar echoes: Origin of echoes and accuracy of derived winds.
863 *Journal of Atmospheric and Oceanic Technology*, 11:1184–1206.
- 864 Zhang, G. (2016). *Weather Radar Polarimetry*. CRC Press.