



HAL
open science

Using Deep Learning for Restoration of Precipitation Echoes in Radar Data

Pierre Lepetit, Laurent Barthès, Cécile Mallet, Camille Ly, Nicolas Viltard,
Yvon Lemaître

► **To cite this version:**

Pierre Lepetit, Laurent Barthès, Cécile Mallet, Camille Ly, Nicolas Viltard, et al.. Using Deep Learning for Restoration of Precipitation Echoes in Radar Data. *IEEE Transactions on Geoscience and Remote Sensing*, 2022, 60, pp.5100914. 10.1109/TGRS.2021.3052582 . insu-03151106

HAL Id: insu-03151106

<https://insu.hal.science/insu-03151106v1>

Submitted on 11 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Using Deep Learning for Restoration of Precipitation Echoes in Radar Data

Pierre Lepetit ^{* ‡}, Laurent Barthès [‡], Cécile Mallet [‡], Camille Ly [‡], Nicolas Viltard[‡] Yvon Lemaître[‡],

^{*} CNRM, Météo-France

[‡] LATMOS, Guyancourt, France

contact:

pierre.lepetit@latmos.ipsl.fr,

Abstract—Raw data issued from meteorological radars are often corrupted by unwanted signals generically called clutter (non-precipitating echoes). Hills and tall buildings, atmospheric turbulence, birds and insects yield patterns that complicate the interpretation of radar images and might add bias in the Quantitative Precipitation Estimates (QPE). These non precipitating echoes differ from meteorological information by their polarimetric signatures and by their particular shapes. Many methods have been developed to detect and remove clutter in radar data. In order to deal with the two tasks, respectively referred to as segmentation and restoration, we present an approach based on machine learning technics and more specifically based on deep learning methods inspired from recent works developed in various domain [1], [2], [3], [4]. Most deep learning algorithms use supervised learning. This implies that it is necessary to have not only raw radar images but also ground truth (precipitation echoes images). This assumption is rarely met for meteorological radars because of the lack of ground truth. In this study, we show that a semi-supervised approach can be successfully used to mitigate this problem.

Index Terms—meteorological radar, polarimetry, weak supervision, fully convolutional networks, deep learning, image segmentation, blind inpainting, restoration.

I. INTRODUCTION

Quantitative precipitation estimates (QPE) using weather radar requires filtering of non-precipitating echoes such as ground clutter, clear air echoes, interferences and anomalous propagation. The reflectivity factor of these non-precipitating echoes are of the same order of magnitude as those of precipitation and can cause substantial errors in rain accumulations. Non-precipitating echoes mitigation has been an ongoing problem for many years and is the subject of an extensive literature. Various approaches have been used such as thresholding, decision trees, [5], multilayer perceptron [6], [7] and fuzzy logic [8], [9], [10], [11]. The frequency of the radars (in S band or C band, generally) and the type of unwanted signals can vary, but most of the existing methods predict the class of each pixel from its immediate neighbors through handcrafted filters, as either texture indices [9] or gradients [10]. Restoration algorithms, generally applied to the binary image provided by the segmentation step, consists in a completion of masked regions (non-blind restoration).

In the early 2000's, segmentation of unwanted signals on meteorological polarimetric radar data was already dominated by data-driven methods, and more specifically by fuzzy logic methods [12].

Nowadays, segmentation algorithms based on fuzzy logic are used by national weather agencies such as the National Weather Service (NWS) [14] or Météo-France [9], and are implemented in operational radars [15]. Nevertheless, segmentation performances are typically lower in case of light precipitations and when non-precipitating echoes overlap rainy areas. In addition, fuzzy logic algorithms, which are data-driven supervised approaches, require the labelling of a large number of pixels to properly estimate conditional empirical probability density functions (pdf) that induce time-consuming manual tasks. Moreover this task has to be renewed for each new setting. However in practice, these pdf rely on a collection of individual pixels that have been manually selected from a relatively small amount of radar images, inside isolated and typical unwanted echoes. In these collections, the contextual information contained in the vicinity of the pixel is lost or at best coded in a basic manner in the form of a texture index or a gradient.

Recently, in the computer vision community, segmentation, non-blind and blind restoration tasks have been addressed with Full Convolutional Networks (FCNs) [1], [2]. Thanks to these deep architectures, the predicted class (in the case of pixel segmentation) or intensity of a pixel (in the case of pixel restoration) relies on a broad neighbourhood through nonlinear filters that have been learned through a data-driven optimisation. For instance, the pixel classification of Wang et al. [16] implies a Convolutional Neural Network (CNN) used as a preliminary, coarse classification step. The prediction is then refined by a fuzzy logic classifier. The authors notice an improvement of unwanted signal detection and attribute it to the fact that the CNN output depends on a wider spatial window around the pixel.

Deep neural networks yielded interesting results on the rain/no-rain classification problem in a supervised framework [17], [18]. Some studies went beyond that framework: Bechini and Chandrasekar [19] and Besic et al. [20] adapted

and developed semi-supervised clustering methods for classification of hydrometeors where clusters are initialised thanks to a fuzzy logic approach. However, these works mainly deal with the discrimination of precipitation echoes (rain, snow, hail, etc) as our task is to discriminate precipitation echoes from unwanted echoes.

In this study, we propose an approach meant to deal with the lack of ground truth : labelled radar echoes (segmentation task) or reference images of precipitation echoes (restoration task). We present a weakly supervised method for the restoration of radar horizontal reflectivity maps for precipitation estimation. Our method is based on FCNs and to our knowledge this is the first time FCNs are employed to improve the quality of weather radar images.

The paper is organized as follows: section 2 describes the FCN architecture and presents the general framework for weakly-supervised machine learning. Section 3 provides information on the dataset and explain how the different databases were built. In section 4 we present the learning process we develop while section 5 shows the performances of the method in different contexts. Finally section 6 summarizes the results of the study and provide some suggestions to improve the method.

II. FULL CONVOLUTIONAL NETWORKS AND WEAKLY-SUPERVISED FRAMEWORK

Segmentation and restauration with FCNs

FCNs are deep neural networks built from sets of convolution filters. The input image is conveyed through a cascade of such sets, alternating non-linear renormalisation operations and down-sampling operations.

The figure 1 represents one of the first FCN, the U-net [13]), developed for semantic segmentation in a biomedical context and is now widely used. The non-linear operations generally consist in positive part functions (or ReLU functions). Their combination with the convolution filters generalises the texture indices as computed in [9]. The down-sampling operations decrease the size of the successive outputs, termed feature maps, while the size of the filter's sets is growing. During a second step, the size of the feature maps increases thanks to specific linear filters (transposed convolution layers). In the U-net, skip connections (concatenations) allow to mix local and global information in the last steps (see fig. 1).

During the training phase, the network's parameters θ_k are upgraded through a gradient descent. Increments are proportional to $\nabla_{\theta_k} L(f_{\{\theta_k\}}(x), y)$ where (x, y) is a input-target pair and $L(\dots)$ is a cost function that measures the discrepancy between the output $f_{\{\theta_k\}}(x)$ and the target. FCNs offer a better alternative than previous architectures on several tasks. These networks are used for semantic segmentation [13], [1], [2], [14] and blind restoration [2], [3]. In [4], FCN are even trained simultaneously on restoration and segmentation of unwanted signals following a multi-task approach [15].

For blind restoration, FCNs dominate the state of the art on classical problems in Machine Learning as text and watermark

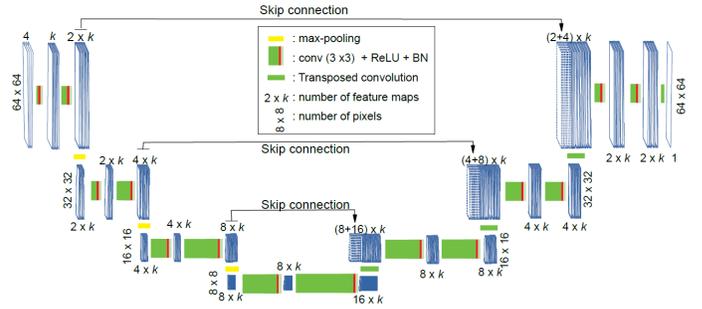


Fig. 1. FCN architecture (U-net) schema modified from [13]. The 4-polarimetric variables used as input go through a series of linear filters (convolutions, max-pooling) and non-linear filters (ReLU). Concatenation of the feature maps are represented by "skip connections". Here, the shapes of the feature maps are given for an input of 64×64 pixels. The number of feature maps are parameterised by k which is set here to 64.

removal. These network are already applied for the post-processing of remote sensing data. For example, Xiao et al. [16] clean strip noise by FCN on infrared satellite images of clouds.

In the aforementioned studies, the training is perfectly supervised: for blind restoration tasks, two datasets are given: the first one contains unwanted signals (z_i), that could be text of various colours and fonts, watermarks, etc. The other is made of clean pictures (y_i). The FCN is trained on input-target pairs of the form: $(x_i = y_i \oplus z_i, y_i)$, where the operator \oplus stands for a matting operator. These notations are kept in the remainder of the text.

Weakly-supervised restoration

For blind restoration, deep learning methods have recently been developed in the case where no clean images are available.

In a recent work, Halperin et al. [17] address the case where no clean picture is available to train the model. Two datasets are used: The first one contains corrupted pictures y'_i that mix clean pictures y_i with unknown unwanted patterns z'_i ($y'_i = y_i \oplus z'_i$). The second contains images with unwanted signals only (z_i) that are different, but of the same nature as the z'_i . The y_i are estimated through an iterative algorithm reported as sensitive to the initialisation step (NES: Neural Egg Separation): at each step, an FCN is trained to separate z_i from the $\tilde{y}_i^k \oplus z_i$ where the \tilde{y}_i^k are the current estimations of y_i .

As in [17], we assume that the corruption process could be modelled by a superposition of unwanted signals z'_i . Moreover, we consider that images z_i , distributed as z'_i , could be sampled during non precipitating periods. The present algorithm looks like an improvement over the first step of the NES: the FCN is trained to restore images that have been corrupted on input-target pairs of the form $([y_i \oplus z'_i] \oplus z_i, [y_i \oplus z'_i])$. In these pairs, the target $y'_i = y_i \oplus z'_i$ is corrupted by the unknown unwanted signals z'_i . Through this

design, the developed method should be robust to target noise.

In parallel, the problem of segmentation of unwanted signals is addressed. Again, corrupted images are filled with known unwanted signals. The support of these known unwanted signals defines the two classes: a pixel is “corrupted” or not. The pixels that are in the support of the unknown unwanted signals are hence generally associated to the “bad” label. This is a situation of noisy target.

Noisy targets

A problem of restoration under noisy target conditions has recently been studied by Lehtinen et al. [18]. From a formal approach of the training process, the authors adapt the cost function to the noise statistics and use FCNs on Machine Learning benchmarks to show that performances are identical or even better than in the supervised framework. However, in that study, the noise on targets is independent from the noise on inputs. Their conclusions could not be directly applied in the case where inputs partly contain the same unwanted signals as the targets. By comparison, analysis and correction of the label noise impacts on classification and segmentation algorithms have been thoroughly studied [19]. When a fraction of the labels are switched randomly (“flip noise”), CNN remain relatively robust [20], [21], [22]. Rolnick et al. [21] show for example that the size of the learning database may even mitigate the impacts of a massive flip noise.

Petit et al. [23] used a FCN for a binary segmentation task in a semi-supervised framework. In their “baseline” approach, all the available images are seen by the network during the training. On these images, whole regions have been mislabelled, as are our unknown unwanted signals. The authors show that the effect of the label noise is significant. A second approach, where the dataset is iteratively completed, allows to recover the scores obtained without label noise. However, it relies on the use of a perfectly labeled sub-dataset.

Numerous less demanding methods exist. In Brodley et al. [24], on an image-wise classification task, a first training pass involves all the available images. Training scores allow to detect and reject suspicious data. A second training pass is then made on the remaining images. Interesting cases may be bypassed however, but this method leads to better scores on both synthetic and real data. This approach that was adapted for the present approach to perform binary segmentation and restoration.

III. DATA

A. Data from the meteorological radar of Trappes

The method was trained and tested on raw data produced by the polarimetric ground radar of Trappes (Météo-France), which covers the Parisian metropolitan area. It is part of the ARAMIS network of french weather radars [25]. The radar rotates around a vertical axis so that its beam describes a whole 360° at a given elevation before moving to the next

one. The beam itself is split into so-called “radar bins”. The radar bins are located by a set of three coordinates: elevation, azimuth and range from the radar. The full exploration sequence lasts fifteen minutes and includes eighteen 360° -rotation of the antenna, three of which are performed at 0.4° elevation with a five minutes lag. The images in the present work are thus made of pixels corresponding to radar bins over a 360° rotation of the antenna at a constant elevation of 0.4° . All four polarimetric variables: Z_h , Z_{dr} , ρ_{hv} and ϕ_{dp} , are measured simultaneously by the radar and used as inputs but the objective was restricted to the cleaning of the horizontal reflectivity only. The restored images are thus one-channel images.

These choices were guided by the predominant role of the lowest elevation and the horizontal reflectivity in hydrological applications.

Projected on a map, the data display precipitation echoes, but also ground echoes of variable intensity possibly caused by anomalous propagation of the electromagnetic waves, clear air echoes, radial strips due to interferences, echoes of planes, etc. Following Doviak and Zrnica [26], all those unwanted signals will be hereafter called with the generic term of “clutter”. The permanent echoes: relief, buildings, antennas, will be referred to as “ground clutter” when a more specific term is required. Similarly, the transient echoes: interferences, anomalous propagation, planes, birds, insects, clear-air will be referred to as “transient clutter”. Some of those clutters (ground echoes and interferences) are shown on fig. 2.

Ten months of the 2017 radar archives were used to build the dataset. These data are given at a time step of 5 minutes and consist of $4 \times 720 \times 1066$ pixel matrices. The first dimension corresponds to the four polarimetric variables: horizontal reflectivity (Z_h), correlation coefficient (ρ_{hv}), radial gradient of the differential phase (K_{dp}) and differential reflectivity (Z_{dr}). The last two dimensions are: the azimuth (resolution of 0.5°) and the range (1066 240 m-radar bin spanning over 250 km).

In the archives, the horizontal reflectivity were already corrected for calibration and thresholded. The correction ensures the temporal homogeneity of the measurements [27]. The thresholding is range-dependent and ensures that the SNR is sufficient to maintain a high data quality. For instance, at 100 km from the radar, the reflectivities below $Z_{min} = 3 \text{ dBZ}$ and the corresponding three additional polarimetric variables are set to missing value.

The polarimetric variable are generally used to infer the nature of the echoes [12], [11], [9], [5] in most of the operational algorithms. On figure 2, some clutter such as the visible strip and some ground clutter can be seen and display specific textures and values.

Inputs are extracted (see below) from these multi-channel images without re-projection: the radar coordinate system has been privileged. It allows a uniform processing for the parts of the polarimetric fields that are equidistant from the radar.

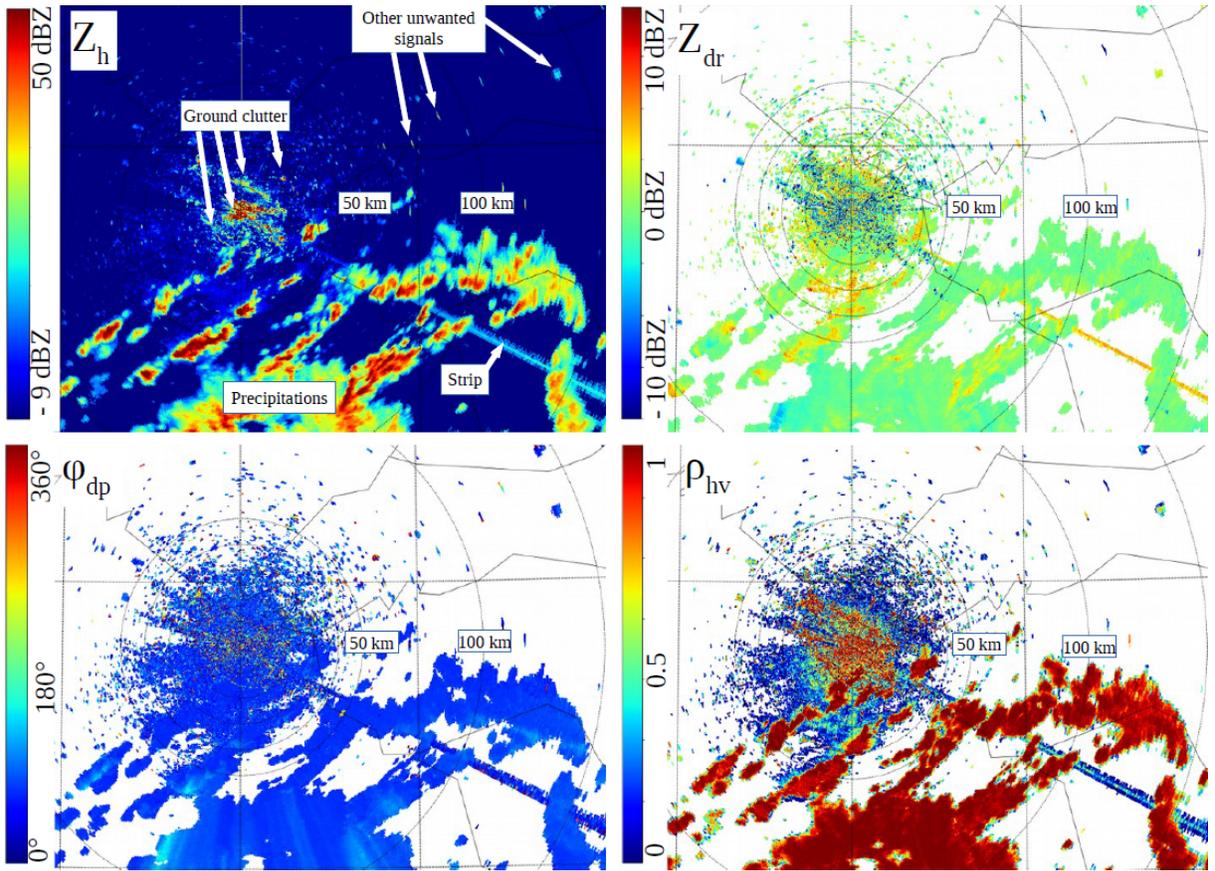


Fig. 2. Partial view of Plan Position Indicator (PPI) of the four polarimetric variables measured by the Trappes radar on 08/15/2017 at 14h20 UTC. Top left: horizontal reflectivity Z_h , top right: differential reflectivity Z_{dr} , bottom left: differential phase ϕ_{dp} and bottom right: correlation coefficient ρ_{hv} . The elevation angle is 0.4° . The radar is located at the center of the circles. The five inner-most circles are spaced 10 km apart. The black lines represent coarse positions of the main rivers. On the images, a mixture of rain echoes and clutter can be seen.

B. Building the Datasets

Two datasets have been prepared. The first, B_p , is extracted during precipitating events. It contains four-channel images (y'_i) of precipitation echoes (y_i) mixed with clutter (z'_i). The second dataset, B_c , is extracted during non-precipitating periods. It contains four-channel images of clutter only (z_i). "Known" clutter will refer to those of B_c .

In both cases, the images have been cropped down to 128×128 pixels along the same grid. This grid is made of concentric rings in which the physical parameters of the pixels (altitude of the measurement, volume per pixel, SNR mean ratio) are considered homogeneous.

Precipitating events were detected using the time series of 152 rain gauges located in the first 150 km around the radar. Following the work of Dilmi et al. [28], a precipitating event begins as soon as the sum of the 5 min rain accumulations over all rain gauges exceeds an empirical threshold of 1 mm. To define non-precipitating periods, one hour margins were applied before and after the time sequence of a precipitating

events.

The extraction grid ranges from 11 km (48^{th} gate) to 148 km (296^{th} gate) from the radar. It contains five slightly overlapping rings spreading over 128 gates each (see figure 3). This width is an empirical trade-off between the available memory of the GPUs, the homogeneity of the physical parameters and the characteristic size of the precipitation cells.

Each ring is itself divided in 4 adjacent extraction windows, each one spreading over 128 successive azimuths (64°). Hence the images extracted in these windows size 128×128 pixels.

In the internal ring, ground clutter is widespread and masks precipitating echoes. Images from this ring are kept in the B_u dataset only, where they are grouped with the images from the second ring.

On the four external rings, images extracted during non-precipitating periods were often lacking significant clutter and sometimes presented weak rain echoes. Even the images of the B_p dataset were often devoid of precipitation cells.

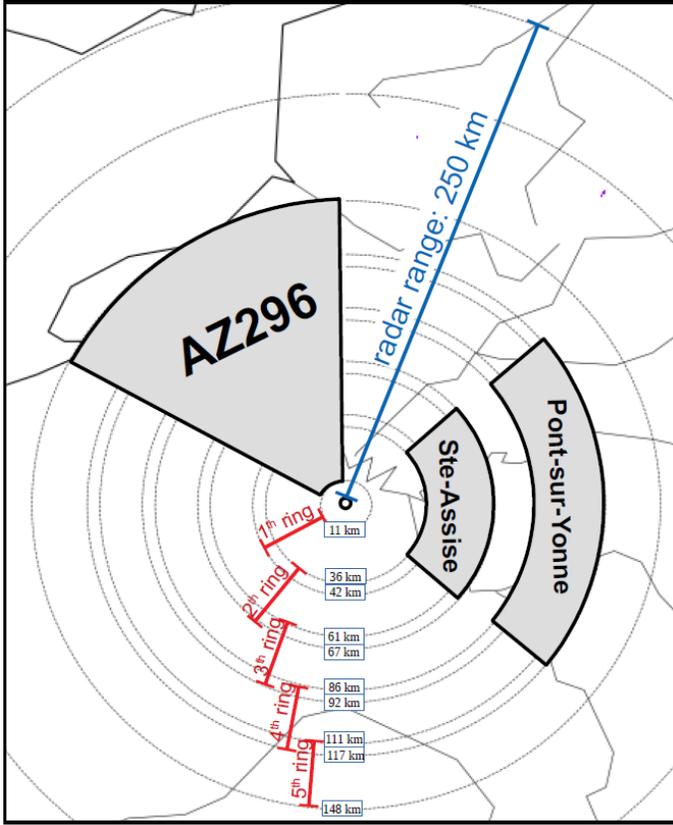


Fig. 3. Rings of extraction and relative positions of AZ296, Sainte-Assise and Pont-sur-Yonne areas.

To avoid uninteresting cases and limit the target noise during the training phase, an additional selection scheme was applied for each dataset and in each ring separately. This three-steps selection scheme relies on the fact that precipitation echoes are generally spatially smoother than clutter. First, the Haar wavelet transform is applied on the 128×128 horizontal reflectivity. Then the resulting Haar coefficients are squared and summed scale by scale, so each image is mapped to a 7-dimensional vector. These vectors are normalised to sum 1 and finally classified by K-means into seven clusters. This scheme was repeated eight times, one time for each ring and for each dataset. The homogeneity of the clusters were manually checked. The clusters without precipitation or with weak precipitation were removed from B_p and clusters without significant clutter were removed from B_u .

C. Splitting for training and testing

The training was performed on each of the four outer rings. To test the method however, the images extracted in the sector between 296°N and 360°N (AZ296), spanning over 128 successive azimuths (see fig. 3), were kept apart.

Since we want to clean all the images of our dataset, it is also necessary to verify that the method performs well on images of corrupted precipitations of the training set (images of B_p).

For that purpose, the Sainte-Assise and Pont-Sur-Yonne area were used for both the learning and the test phase.

The Sainte-Assise area corresponds to one of the extraction windows of the second ring (see fig. 3 and fig. 5). It displays different types of ground clutter. Several pylons of more than 200 m height generate a spot of high reflectivity on the mean reflectivity field (see in the lower left corner of fig 5.a). Numerous other pixels are affected by ground echoes to various degrees, depending on the atmospheric conditions. In that area, three strips are also observed. The main strip is broader, more frequent and much stronger. Along this strip, and disregarding the pixels corrupted by superposed ground clutter, the horizontal reflectivity may reach 20 dBZ. It is probably caused by an interference with the airport surveillance radar of Orly located in Palaiseau. It will further be referred to as the "Palaiseau strip".

The area of Pont-sur-Yonne (see figure 3, situated in the fourth ring of extraction, shows often the same strips but contrary to Saint-Assise, ground clutter and clear sky-echoes are infrequent.

D. Final form of the training datasets

Finally, both training datasets contain about 150,000 images of size $4 \times 128 \times 128$ pixels, spread over the four outer rings (see table I).

The archive data of the Trappes Meteorological radar will soon be available on the AERIS¹ web portal. The image's references of both our datasets² will also be available.

IV. METHOD

A. Matting operator and building of inputs

The four-channel images of B_p (resp B_c) are noted y'_i (resp. z_i). The power of the back-scattered signal is additive, hence, for two clusters of targets R and U in the same radar gate, one can write:

$$\begin{aligned} Z_h &= Z_h(R) + Z_h(U) \\ Z_v &= Z_v(R) + Z_v(U) \\ Z_{dr} &= Z_v/Z_h \\ &= 1/Z_h[Z_h(R) \times Z_{dr}(R) + Z_h(U) \times Z_{dr}(U)] \\ &= m_h(Z_{dr}(R), Z_{dr}(U)) \end{aligned}$$

Where m_h stands for the mean between the two Z_{dr} weighted by Z_h . For K_{dp} , the same ponderation holds when the phase formula is linearised at the first order : $K_{dp} = m_h(K_{dp}(R), K_{dp}(U))$.

Combining the ρ_{hv} would imply the use of the variance of the signal which was not available. By default, we chose to use $\rho_{hv} = m_h(\rho_{hv}(R), \rho_{hv}(U))$. Consequences of these choices are discussed in VI

¹<https://www.aeris-data.fr/>

²https://github.com/space-latmos/deep_washing_radar_data

ring (gates)	2^{th} (152-280)	3^{th} (256-384)	4^{th} (360-488)	5^{th} (464-592)	Total
B_p	37,000	36,000	35,000	33,000	150,000
B_c	87,000	36,000	20,000	15,000	150,000

TABLE I

COMPOSITION OF B_p AND B_c AFTER SELECTION. ON THE 87,000 IMAGES OF B_c IN THE FIRST COLUMN, 46,000 ARE FROM THE INTERNAL RING (48-176).

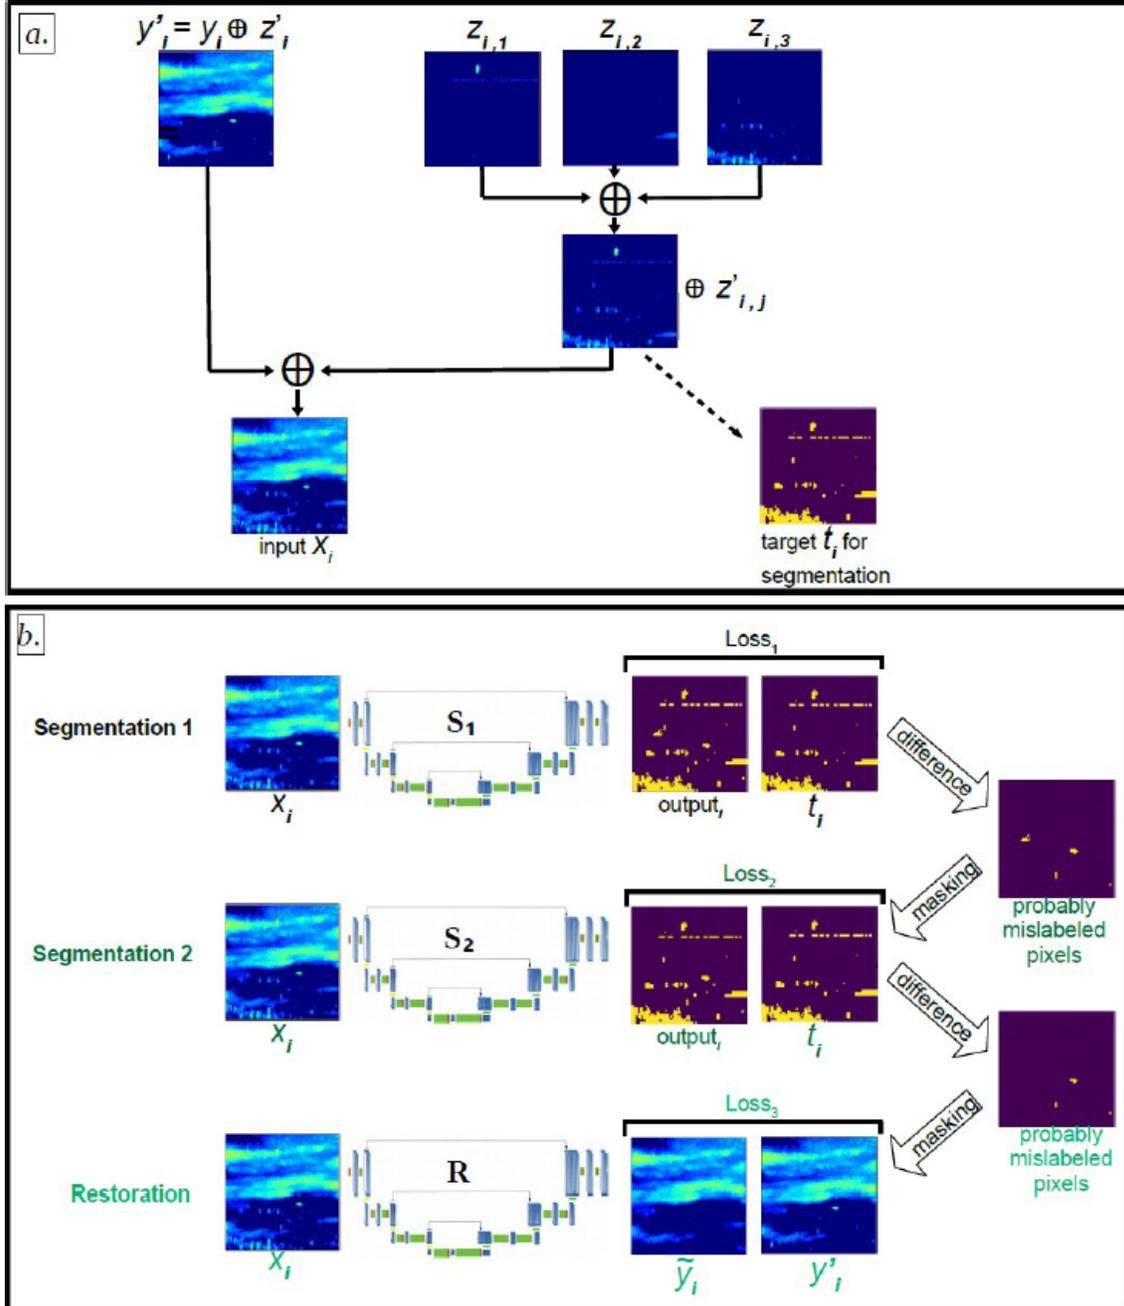


Fig. 4. a. Building of inputs and targets. The yellow regions of the target t_i . The known clutter $z_{i,j}$ from B_c are added to the corrupted rain echoes y'_i taken in B_p , to form the input x_i (cf. eq. 1). The regions corrupted by known clutter are encoded in the target t_i (cf. eq. 3). b. Successive trainings. Firstly, S_1 is trained to segment regions corrupted by known clutter. Then, pixels that are detected by S_1 as probably mislabeled are masked when S_2 is trained. Finally, for the blind restoration, S_2 is used as a mask to train R . The losses are defined in eq. 4, eq. 5 and eq. 6

The matting operator \oplus finally corresponds to a simple sum for the first component (Z_h) and the weighted mean (m_h) for the three others. Similarly, the y'_i may be decomposed in: $y'_i = y_i \oplus z'_i$ where y_i contains the rain echoes, and z'_i , the clutter.

It is assumed that the z'_i are independent from the y_i and that they follow the same distribution as the z_i . Under these assumptions, once inputs are written as $x_i = y'_i \oplus z_i$, it is no more possible to distinguish the z'_i from the z_i . The proportion of clutter can be interpreted as a label noise ratio [19]. Adding several images of B_c reduces this ratio but too much additional signals largely complicates the restoration problem. In the end, the addition of three images $z_{i,1}, z_{i,2}, z_{i,3}$ taken from B_c appeared to be a good trade-off.

For both segmentation and restoration, the inputs are finally given by:

$$x_i = y'_i \oplus \left[\bigoplus_{j=1}^{j=3} z_{i,j} \right] \quad (1)$$

where i indexes the image in the training batches and will further be omitted.

B. Data augmentation and stratified sampling

The matting is implemented as a data augmentation step and is made on the fly. As the physical parameters and the frequencies of clutter vary according to the range, the y'_i and the $z_{i,j}$ are ring-dependent. During the training, the rings are uniformly distributed in our mini-batches.

Other augmentation data functions are applied before matting. Images from B_p are randomly cropped (from 128×128 to 64×64) and an azimuthal flip is applied with a probability of 0.5. Images from B_c are moreover affected by a random cyclic permutation of the azimuthal indices. Thanks to these random operations, the clutter is uniformly distributed in the input images, even though they are caused by unmoving targets.

C. Segmentation and restoration

The U-net architecture [13] is then trained on segmentation and restoration tasks. The initialisation of the weights, the optimisation method and the learning rate (10^{-4}) are set to standard values [29], [30].

A first network, S_1 , is trained on a binary segmentation task “corrupted” vs. “not corrupted”. For a given input $x = y' \oplus \left[\bigoplus_{j=1}^{j=3} z_j \right]$, the target is the matrix $t = (t^k)_{k \in E}$ where E stands for the set of the 64×64 pixels and:

$$t^k := 1 \quad \text{if} \quad \left[\bigoplus_{j=1}^{j=3} z_j \right]_{refl.}^k > -9 \text{ dBZ} \quad (2)$$

$$0 \quad \text{if} \quad \text{not} \quad (3)$$

The output of S_1 is the matrix of probabilities $p_{S_1}^k$ for the pixel k to be corrupted. The loss function is defined with the cross-entropy:

$$Loss_1 = - \sum_{k \in E} t^k \ln(p_{S_1}^k) + (1 - t^k) \ln(1 - p_{S_1}^k) \quad (4)$$

Then, a second U-net, S_2 , is initiated on S_1 and trained with a modified loss function. Following the work of Brodley et al. [24], the weight of the probably mislabeled pixels (those of $F_{noisy \ label}^{S_1}$ – see eq. (5)) is reduced. On synthetic data, a null weight offers better results, meaning mislabeled pixels are likely masked during the training of S_2 . $Loss_2$ is defined by replacing E in $Loss_1$ by E^{S_1} where:

$$E^{S_1} = E \setminus F_{noisy \ label}^{S_1} \quad (5)$$

where:

$$F_{noisy \ label}^{S_1} = \{k \in E \mid p_{S_1}^k > 0.5 \text{ and } t^k = 0\}$$

Finally, a third U-net, represented by R , is trained on a blind restoration task. The target is the first component of y' , i.e. the horizontal reflectivity $y'_{refl.}$. The last layer of the network is modified to generate a scalar field \tilde{y} .

In recent work on images restoration, the distance between the output and the target is generally computed with RMSE [31], [2], [32], [33] or with Mean Absolute Error [3], [4]. Using RMSE in a context of noisy target would favour a compromise between suppression and conservation of the clutter. On the other hand, Mean Absolute Error (MAE), favours the removal of all the clutter as soon as they're mostly removed from the target. This claim is formalised in the work of Lehtinen et al. [18]. To take advantage of the preceding segmentation steps, the probably mislabeled pixels are masked thanks to the S_2 output:

$$Loss_3 = \frac{1}{|E^{S_2}|} \sum_{k \in E^{S_2}} |\tilde{y}^k - y'_{refl.}| \quad (6)$$

where $|E^{S_2}|$ stands for the cardinality of E^{S_2} .

In the testing phase, to preserve regions that have not been corrupted by clutter, modifications are limited in corrupted areas according to S_2 . Output \tilde{y} is hence defined by:

$$\tilde{y}^k := \tilde{y}^k \quad \text{if } p_{S_2}^k(\text{“corrupted”}) > 0.5 \quad (7)$$

$$y'^k \quad \text{if not} \quad (8)$$

D. Refinement with an ensemble method

Applying R on images of B_p provides a first estimate of the clean images y_i . But the addition of significant clutter during the training phase raises a common issue (cf. [34]): the images of B_p are statistically much less affected by clutter than the training inputs. This “domain shift” may alter the performances of the model. Thus, instead of directly applying the algorithm on B_p images, the matting process is repeated

Area	nb. images	Mean ref. (inputs)	Mean ref. (outputs)
Sainte-Assise	8,000	- 5.4 dBZ	-8.9 dBZ
AZ296	10,000	-3.3 dBZ	-8.5 dBZ

TABLE II
MEAN REFLECTIVITIES COMPUTED FROM NON-PRECIPIATING SITUATIONS BEFORE (*inputs*) AND AFTER (*outputs*) PROCESSING.

according to a naive weighted average: from an image y'_{test} , the N inputs are formed as:

$x_{test,l} = y'_{test} \oplus \bigoplus_{j=1}^3 z_{l,j}$, $l = 1, \dots, N$. To compute the final estimate \hat{y} at pixel k , the $x_{test,l}$ which have not been knowingly corrupted are averaged :

$$\hat{y}^k = \frac{1}{|G^k|} \left[\sum_{l \in G^k} \tilde{y}_l^k \right] \quad (9)$$

where $G^k = \{l \mid \forall j, z_{l,j}^k = -9 \text{ dBZ}\}$.

This extra step offers improved performances on a set of synthetic data used as a reference test (not presented here). In practice, it is applied separately for each ring, with $N = 16$. To avoid time-consuming random loading of $z_{l,j}$, these images were preselected before the global cleaning.

This latter operation of an individual picture takes about 30s using low-cost GPU and CPUs (Nvidia GeForce 1080 with four Intel Xeon E5).

V. RESULTS

The developed method gives excellent results on individual images. Both the segmentation and the restoration appear visually nominal in most cases. Nevertheless, quantitative evaluation of the performances have been designed and are presented hereafter.

A. Results on known clutter and conservation of precipitation echoes

The first performance evaluation step was to check that clutter collected in non-precipitating situations were correctly removed. This verification was made on images of B_c extracted in the Sainte-Assise and the AZ296 areas (see III). The averaged reflectivities of these non-precipitating situations are shown on Table II. They are much closer to the no-signal threshold of -9 dBZ after processing, showing that clutter have been properly removed. This result is consistently slightly better on Sainte-Assise which was part of the training database, than on AZ296 which was not.

The second performance evaluation step was to check that reflectivities due to precipitations was not affected by the cleaning process . A first set of about 2000 pixels was selected in the Sainte-Assise area, in a precipitating cell affected by clutter of low intensity (clear air echoes). A second set of about 2000 pixels was selected in the AZ296 sector where contamination by clutter is highly unlikely.

Results are shown on Table III. The Mean Absolute Error on Sainte-Assise is small and could be at least partially attributed to the removal of clear air echoes. The perfect

Area	Mean ref. (inputs)	MAE
Sainte-Assise	25.88 dBZ	0.85 dBZ
AZ296	37.6 dBZ	0 dBZ

TABLE III
INITIAL AVERAGED REFLECTIVITY (*inputs*) AND MEAN ABSOLUTE ERRORS AFTER CORRECTION ON A SET OF PIXELS WITH PRECIPITATION.

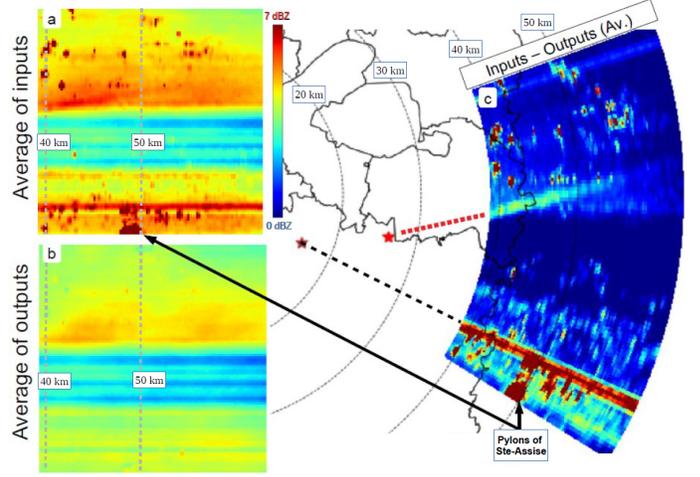


Fig. 5. Averages of 8000 horizontal reflectivity images over the Ste-Assise area, taken from B_p . a. : Mean of input's reflectivity (vertical axis: azimuth, horizontal axis: range). b. : idem a) but after restoration. c. : difference between a) and b) projected in the longitude-latitude coordinate system. The red star and red dashed line give respectively the positions of the runway 26 of Orly airport and its air corridor, the brown star locates a radar transmitter based in Palaiseau, suspected to be the source of the main interference. Administrative borders of the Île-de-France administrative regions also appear as black solid lines.

score on AZ296 shows that S_2 correctly does not detect clutter.

B. Ground clutter and interferences

Figure 5 shows the average reflectivity fields before (*inputs*) and after (*outputs*) processing and the difference between the two. The massive echo due to the pylons of the Sainte-Assise transmitter as well as the ground clutter (red spots) are successfully removed. Furthermore, the difference (fig. 5.c) shows very clearly a non-radial strip which corresponds to a section of two of the Orly western air corridors showing that intermittent echoes (planes) are also detected and properly removed.

On Figure 5.a and b, the effect of partial beam blocking (horizontal blue lines) is appropriately left uncorrected. Notwithstanding that masking effect, the smoothness of the average after processing tends to show that more than being detected, ground clutter is inpainted in a very consistent way, when compared to their surroundings.

Using the difference between *inputs* and *outputs* presented on figure 5.c, we delimit regions of the map that are often corrupted by strong ground clutter. Apart from ground clutter, precipitations are the only acknowledged source of

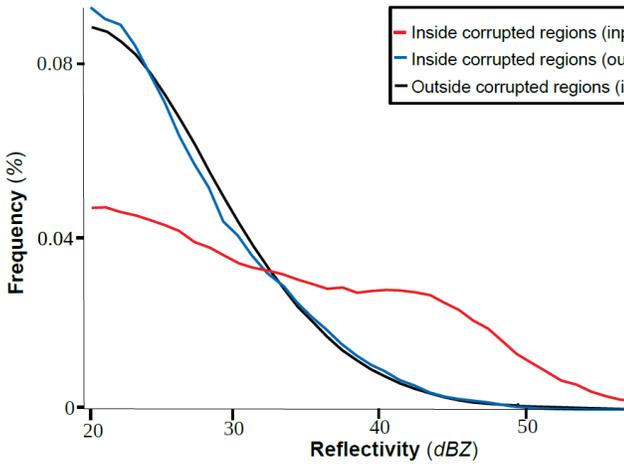


Fig. 6. Estimated distributions of horizontal reflectivities higher than 20 dBZ in the Sainte-Assise area from 8000 images of 128×128 pixels. Frequencies are estimated from histograms of 1 dBZ bin size. The distribution inside the regions corrupted by ground clutter (red curve) is shifted to lower values after restoration (blue curve). These new values match the distribution of the pixel outside the corrupted regions (black curve).

high reflectivity. The high reflectivities outside the corrupted regions could hence be considered as uncorrupted. Their distribution is plotted on Figure 6 (black solid). On the same figure, the pdf of pixels inside the corrupted regions before (red solid) and after restoration (blue solid) are presented. As expected, the black and blue curves match almost perfectly.

To further assess the method's performances on regions that are corrupted by interferences, we focus on the area of Pont-sur-Yonne, located the fourth ring. At this range, as strong ground clutter is rare, the main source of unwanted reflectivity higher than 10 dBZ is the Palaiseau strip (see figure 7.a). This time, pixels with initial $Z_h > 10\text{dBZ}$ and that were corrected are used to plot the red (resp. blue) pdf curve before (resp. after correction) of Figure 7. The black curve on the same figure is the pdf of the initial Z_h of all the other pixels (corrected with initial $Z_h < 10\text{dBZ}$ and uncorrected but with initial $Z_h > 10\text{dBZ}$).

For $Z_h > 10\text{dBZ}$, the restored distribution fits with the distribution of the pixels unaffected by the interference. Between 4 and 12 dBZ, the shaded area on Figure 7.a, is the expected discrepancy between the fraction of pixels whose initial Z_h is below 10 dBZ and were corrected for transient clutters others than interference.

Below 4 dBZ, the effect of the low-SNR threshold imposed on the initial Z_h (see VI.) explains the discrepancy highlighted with the green shade (see section Data). Along the 4th ring, this threshold varies from 1 dBZ at 86 km to 4 dBZ at 117 km. Hence, for input images, the support of the distribution of reflectivity are bounded by $Z_{min} = 1\text{dBZ}$. But this is not the case for the outputs. This point is discussed in VI.

The same results are obtained for the AZ296 region but are not presented here since they do not provide any extra piece

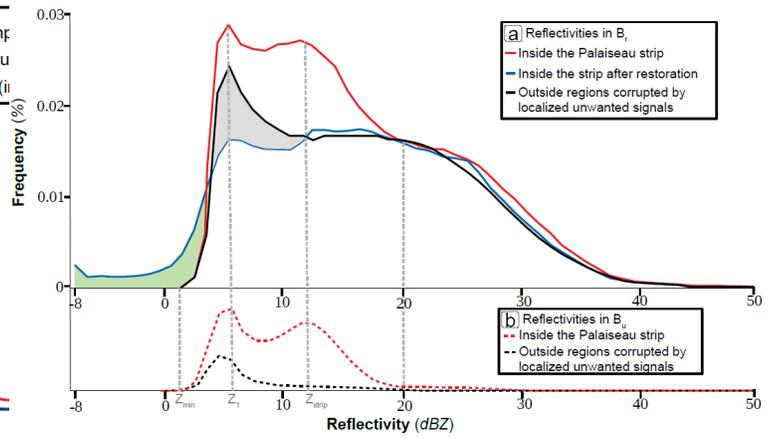


Fig. 7. Distribution of the horizontal reflectivities over the Pont-sur-Yonne area (estimated from 8000 images, $\approx 10^7$ pixels, with a bin's size of 1 dBZ). a. Distributions of horizontal reflectivity from images of corrupted precipitations (images of B_p). The red curve (resp. the blue curve) shows the distribution inside the region frequently filled by the Palaiseau strip before (resp. after) restoration. The red distribution presents two modes Z_0 and Z_{strip} . The reflectivity outside the corrupted regions are represented by the black curve. b. Distributions of horizontal reflectivity in images with known clutter (images of B_c). The red and black dotted curves are associated to the same regions as in a.

of information.

C. Case study

Three case studies are presented in more details and the results are compared to the segmentation of Météo-France's (MF) operational algorithm.

1) Strong ground clutters:

The upper row of Figure 8 shows a situation of light stratiform precipitation over Paris. Echoes of the Paris buildings (noticeably Montparnasse and Eiffel Towers in the white circle) are embedded in the precipitation area. After correction, the building clutter is well removed while restored rain field appears to blend smoothly with the rest of the precipitation area. Similarly, the building's signature in the white rectangle are removed and the restored reflectivity field is consistent with its surroundings. When compared to the mask provided by Météo-France, all contaminated areas appear to have been adequately processed with no apparent discontinuities in the restored Z_h .

In the middle row of Figure 8 the Sainte-Assise pylon return is embedded at the edge of a high reflectivity area (white circle). In Addition, the Palaiseau strip is visible in the white elongated rectangle across the rain patch. Both contamination appear correctly removed and no discontinuities can be detected after restoration. Once again, when compared to the MF product, all clutter seem to have been adequately removed. The main difference seems to be mostly situated along the edge of the innermost part of the domain where MF algorithm classifies

some pixels as rainy while the restored field clearly shows no signal.

2) *Clear air echoes:*

In the broad sense, clear air echoes (CAE) are caused by targets as diverse as aerosols, insects, birds or turbulence. They may occur anywhere, but are very rare beyond 100 km. When misclassified as rain echoes, they could be advected by predictive weather models after assimilation. They also contribute to the radar-based climatology possible bias.

The bottom row of figure 8), shows a situation of a steady south-western flux (along the dashed white line) with suspected clutter presence (white rectangle) along the flux between two active rain patches. The restored image shows that between these two rain patches, the observed clutter is mostly due to CAE. This seems to be confirmed by two vertically-pointing radars set upstream in the south-western corner of the white rectangle (ROXI, X-band and BASTA, W-band) that did not record any precipitation in the stream for a few hours. The Météo-France classification is more ambiguous in the same rectangle area, identifying a mixture of CAE, rain, noise and other unwanted signals. Recent changes in the fuzzy-logic algorithm used for this classification seems to consolidate the idea that most of the pixels labelled as "rain" were effectively CAE (Météo-France/CMR, personal communication).

3) *Anomalous Propagation:*

Ground clutter may be due to Anomalous Propagation (AP) of the electromagnetic wave, when a strong stratification of the atmosphere curves the radar beam downward. This kind of clutter is very transient and more difficult to detect than standard partial beam-blocking due to fixed obstacles. Furthermore, AP may occur far from the radar where ground clutter is usually unlikely to occur. Figure 9 shows such an occurrence on 08/08/2017. On the input reflectivity field, the trace of the high-voltage line between Boisville-la-Saint-Père and Artenay could exceptionally be seen (white dotted line). Other clutter around the power line could also be attributed to AP echoes. After processing, clutter has been removed. Furthermore, the edges of the restored cell seem to match almost exactly those predicted by the Météo-France classification algorithm.

D. *failure cases*

Figure 10 shows two example when the restoration process clearly failed. The top row shows a situation where clutter in the white rectangle is embedded in a large patch of light precipitation. After restoration, if the ground clutter were clearly identified and removed, the restored field exhibits unrealistic holes. Furthermore, in the upper-left corner of the image, the cropping of the precipitation area seems also unrealistic. This latter "excessive cropping" is certainly more often observed than the ground clutter anomaly. It is to be noted that it only concerns reflectivities that are lower than 5 dBZ.

The bottom row of Figure 10 shows a rare occurrence of underestimation affecting a whole precipitating area, as in the white rectangle of the bottom right image.

Ancillary data, namely observations from the aforementioned ROXI and BASTA radars show that in these two particular instances, the precipitation under consideration are actually two cases of snow events. These are relatively rare in the Ile-de-France region (no more than a few days a year) and the consequences are discussed in VI.

VI. DISCUSSION

The two failure cases presented on figure 10 occurred both in snowfall conditions. Once again, these situations being rare in the region where the Trappes radar is located (no more than 3 to 5 days in 2017), they are likely very underrepresented in the database. The snow has very different polarimetric signatures when compared to rain and if the network is capable of correctly identifying the ground clutter, it is unable to properly inpaint the gaps, as in the first row of fig. 10. This is conformed by the bottom example (09/01/2017) when around the same time, ROXI sees a marked transition from rain to snowfall. It is thus possible that the cells that are properly processed are indeed rain while the failing ones are the snowfall. This would require further investigation to be consolidated.

Another question regards the inpainting itself. The network could be filling the regions where clutter is removed with a median value eventually adjusted somewhat randomly. This would lead to a proper averaged reflectivity over the sum of images but would not be satisfactory in terms of the algorithm expected performances on an individual image. Although it is extremely difficult to find a definite answer to this point, two remarks can be made. First, the individual images (except in the snowfall case) do not show any visible discontinuity of the inpainted pixels with respect to their surroundings. Second, the very good pdf match shown on Fig. 6 and 7 would also advocate for a proper inpainting of de-cluttered regions at least for interferences and ground clutter.

A remark, complementary with the aforementioned point, is the apparent decrease in averaged reflectivities presented Fig. 5 particularly in areas where neither ground clutter nor interferences were corrected. This aspect can be addressed through a more detailed analysis of Fig. 7 and more specifically the grey shaded area. The black curve shows the pdf of all the pixels that were not corrected for interference or ground clutter. This does not mean that more transient clutter like clear-air echoes or radar noise went undetected and uncorrected for these pixels. It is our assumption (supported also by pdf comparisons, not shown) that the network correctly segmented and inpainted these, eroding the black curve above the grey shaded area hence lowering the average Z_h . The distribution of these transient clutter appears to be only slightly dependent on the azimuth and the network performs a correction of about -0.5 to -1.5

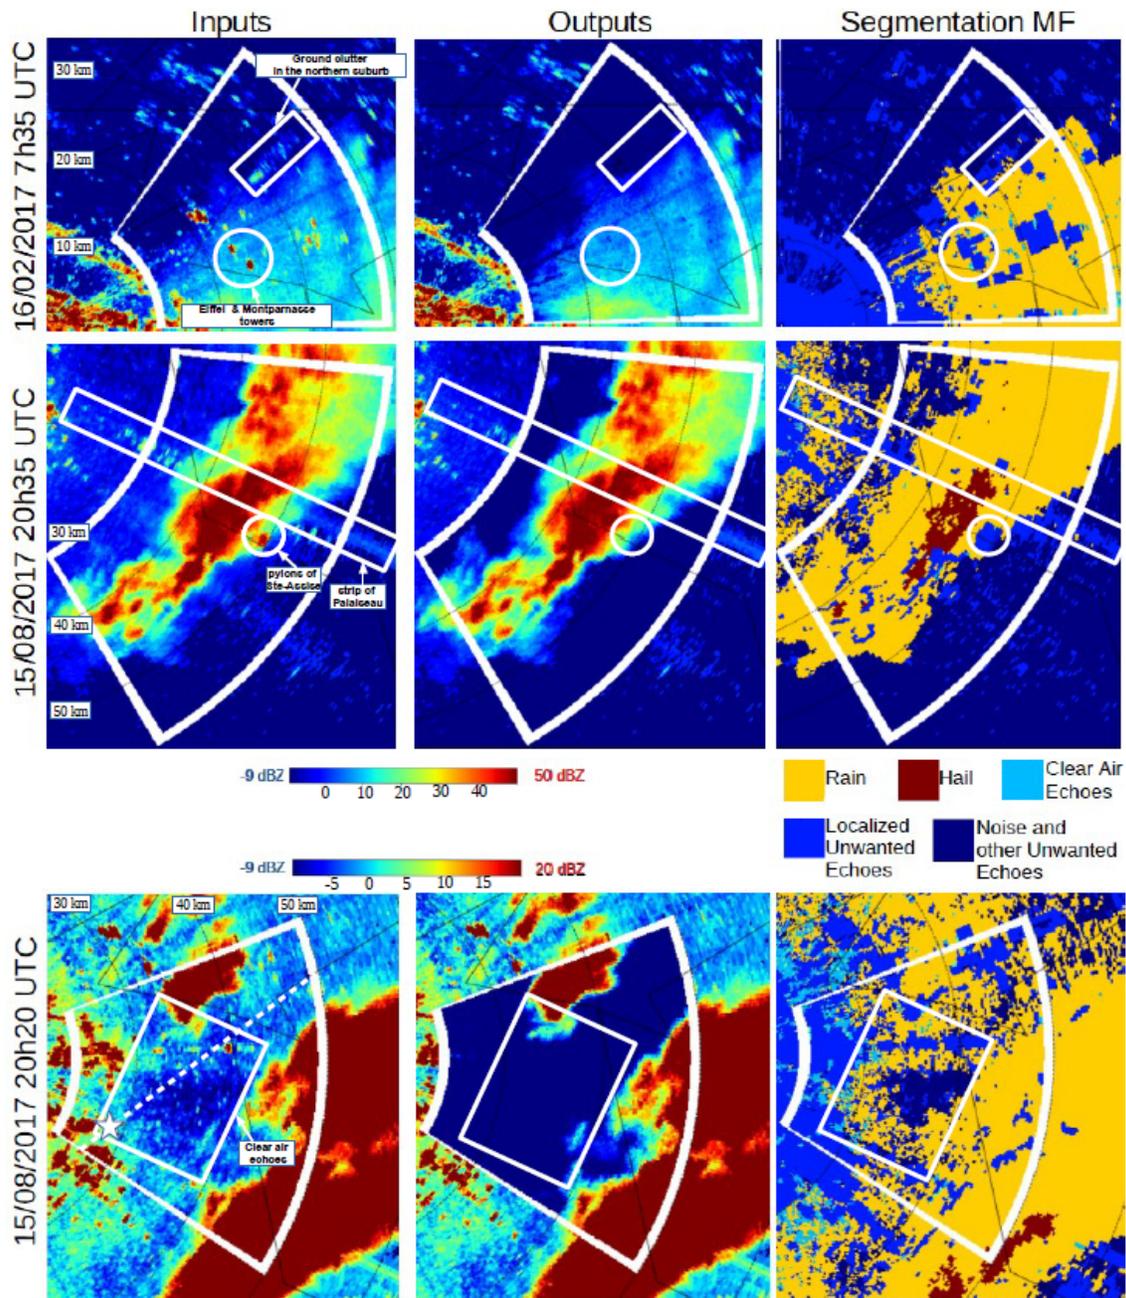


Fig. 8. Removal of ground clutter and clear air echoes. In the output column, the network is only applied into the white section of ring. The third column shows the output of the segmentation algorithm of Météo-France (MF). Top row: echoes of Parisian buildings embedded in (white circle) or at the borders of a weak stratiform precipitation. Middle row: Embedded ground clutter and clear air echoes surrounding a thunderstorm. Bottom row: Clear air echoes around a thunderstorm cell. The white star shows the locations of the two radars, ROXI and BASTA. The white dashed-line represents the atmospheric transect probed by these two radars in the main air flow around 20h20 UTC.

dB on the average Z_h .

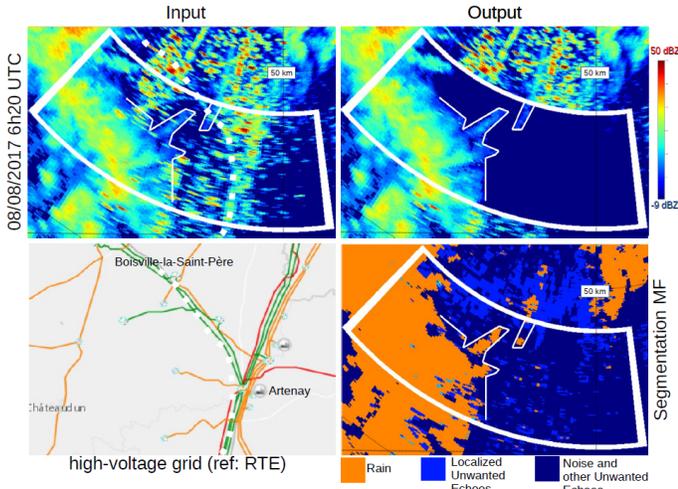


Fig. 9. Example of anomalous propagation echoes. The input reflectivity shows linear echoes that match with the high-voltage power grid around Artenay, located at 75 km from the radar. The thin white contour marks the border of the restored rain cell. This contour corresponds approximately to the borders of the rainy region predicted by the MF algorithm.

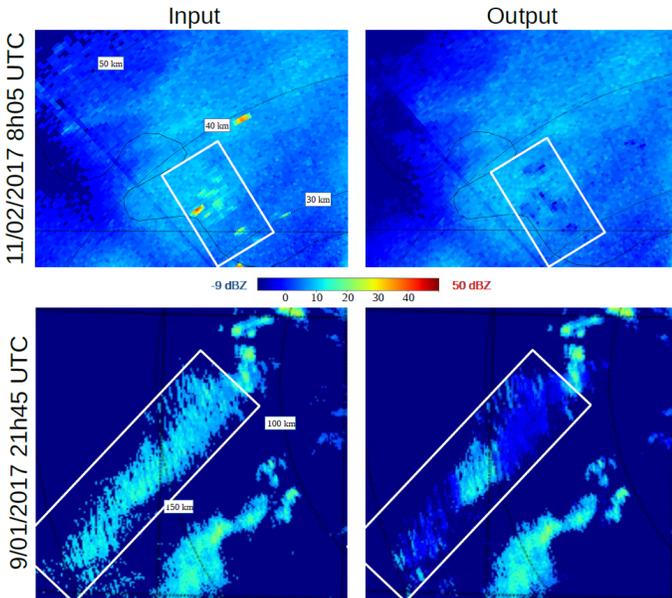


Fig. 10. Failure cases. First row: the restored reflectivity shows unlikely holes after ground clutter removal. Moreover, the westernmost border of the cell have been cropped and lowered. Second row: the structured precipitation cell in the white rectangle has been lowered.

Similarly, the existence of the green-shaded area of Fig. 7 can be traced to a couple of reasons. It might be caused by an imperfect removal of clutter. Figures 5 and table II show that some clutter is not totally removed. However, these residuals are rare and the low-reflectivity values are still observable on distribution outside ground clutter. It could also be related to the processing of the cell edges. When corrupted by clutter, edges of a rain cell are rebuilt as the median border could be. As the low-SNR threshold (see III) grows with the range, the median border is smoother than in the outer rings, so the network tend to smooth the impact of the low-SNR threshold.

Last, some residual misclassification may remain, to the fact that B_c still contains echoes of weak precipitation: for rainrate whose corresponding reflectivities fall under 10 dBZ, the rain is hardly recorded by Météo-France rain gauges. Indeed, a 10 dBZ rain corresponds, through the classical Marshall-Palmer ZR law, to a rain rate of only 0.15 mm/h and is thus extremely hard to filter out with certainty.

VII. CONCLUSION AND PERSPECTIVES

An algorithm to detect and restore clutter in meteorological radar images has been developed. It is based on Full convolutional Networks, which stand among the more efficient architectures on issues as semantic segmentation or restoration, in a supervised or weakly-supervised framework. Based on a weakly supervised approach, the algorithm avoids time-expensive labelling and does not rely on a few events, as approaches based on fuzzy logic generally do. Instead, a large number of clear-air situations with clutter were selected. These images were then matted with the images to be cleaned. The network learned to remove these clear-air clutter with the genuine clutter of the rainy situation.

Despite a strong hypothesis on the way the precipitation echoes and the clutter are combined in the radar picture, results are promising. Ground clutter is removed in most of the case, even if some further work needs to be made on snowfall situation. For echoes harder to collect, such as clear sky echoes and anomalous propagation, the method gives very satisfactory results.

Further work needs to be made to generalise the algorithm and tests its robustness. First, beyond 150 km, the strong effects of the low-SNR threshold and the coarser resolution of the measurement make the generalisation harder. From our attempts, the borders of the rain cells are strongly cropped at far ranges. Among the solution to overcome these effects, using range-specific U-net or adding the radar range in the input will be tested.

Second, the network was also tested on images from the Météo-France radar of Abbeville. It succeeded to remove ground clutter due to wind turbines, although the latter were not in any of the learning databases. However, this

preliminary result also showed that some of the rain echoes were unduly removed also. A specific training on images from other radars will be tested.

Last, the network's robustness will need to be tested on a much larger dataset spanning several years of data to check if extreme events are indeed properly processed.

REMERCIEMENTS

We thanks the CMR (Centre de Météorologie Radar) from Météo-France for the raw data and the results of their segmentation used in our case study. We thanks in particular Clotilde Augros, Valérie Vogt and Nicolas Gaussiat for their constructive comments. This work was supported in part by the CNES/TOSCA and IPSL fundings and conducted in collaboration with Météo-France.

REFERENCES

- [1] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [2] X.-J. Mao, C. Shen, and Y.-B. Yang, "Image restoration using convolutional auto-encoders with symmetric skip connections," *arXiv preprint arXiv:1606.08921*, 2016.
- [3] Y. P. Liu, J. shan Pan, and Z. Su, "Deep blind image inpainting," *CoRR*, vol. abs/1712.09078, 2017.
- [4] A. Hertz, S. Fogel, R. Hanocka, R. Giryas, and D. Cohen-Or, "Blind visual motif removal from a single image," *arXiv preprint arXiv:1904.02756*, 2019.
- [5] A. V. Ryzhkov and D. S. Zrnica, "Polarimetric rainfall estimation in the presence of anomalous propagation," *Journal of Atmospheric and Oceanic Technology*, vol. 15, no. 6, pp. 1320–1330, 1998.
- [6] R. Cornelius, R. Gagon, and F. Pratte, "Optimization of wsr-88d clutter processing and ap clutter mitigation," *Final Report*, 1995.
- [7] V. Lakshmanan, K. Hondl, G. Stumpf, and T. Smith, "Quality control of weather radar data using texture features and a neural network," in *Preprints, 31st Radar Conference*. Citeseer, 2003, pp. 522–525.
- [8] M. Steiner and J. A. Smith, "Use of three-dimensional reflectivity structure for automated detection and removal of nonprecipitating echoes in radar data," *Journal of Atmospheric and Oceanic Technology*, vol. 19, no. 5, pp. 673–686, 2002.
- [9] J. J. Gourley, P. Tabary, and J. Parent du Chatelet, "A fuzzy logic algorithm for the separation of precipitating from nonprecipitating echoes using polarimetric radar observations," *Journal of Atmospheric and Oceanic Technology*, vol. 24, no. 8, pp. 1439–1451, 2007.
- [10] Y.-H. Cho, G. W. Lee, K.-E. Kim, and I. Zawadzki, "Identification and removal of ground echoes and anomalous propagation using the characteristics of radar echoes," *Journal of Atmospheric and Oceanic Technology*, vol. 23, no. 9, pp. 1206–1222, 2006.
- [11] V. Chandrasekar, R. Keränen, S. Lim, and D. Moisseev, "Recent advances in classification of observations from dual polarization weather radars," *Atmospheric Research*, vol. 119, pp. 97–111, 2013.
- [12] J. Vivekanandan, D. Zrnica, S. Ellis, R. Oye, A. Ryzhkov, and J. Straka, "Cloud microphysics retrieval using s-band dual-polarization radar measurements," *Bulletin of the american meteorological society*, vol. 80, no. 3, pp. 381–388, 1999.
- [13] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [14] P. Zhang, W. Liu, H. Wang, Y. Lei, and H. Lu, "Deep gated attention networks for large-scale street-level scene segmentation," *Pattern Recognition*, vol. 88, pp. 702–714, 2019.
- [15] R. Caruana, "Multitask learning," *Machine learning*, vol. 28, no. 1, pp. 41–75, 1997.
- [16] P. Xiao, Y. Guo, and P. Zhuang, "Removing stripe noise from infrared cloud images via deep convolutional networks," *IEEE Photonics Journal*, vol. 10, no. 4, pp. 1–14, 2018.
- [17] T. Halperin, A. Ephrat, and Y. Hoshen, "Neural separation of observed and unobserved distributions," *arXiv preprint arXiv:1811.12739*, 2018.
- [18] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, "Noise2noise: Learning image restoration without clean data," *arXiv preprint arXiv:1803.04189*, 2018.
- [19] B. Fréney and M. Verleysen, "Classification in the presence of label noise: a survey," *IEEE transactions on neural networks and learning systems*, vol. 25, no. 5, pp. 845–869, 2013.
- [20] S. Sukhbaatar, J. Bruna, M. Paluri, L. Bourdev, and R. Fergus, "Training convolutional networks with noisy labels," *arXiv preprint arXiv:1406.2080*, 2014.
- [21] D. Rolnick, A. Veit, S. Belongie, and N. Shavit, "Deep learning is robust to massive label noise," *arXiv preprint arXiv:1705.10694*, 2017.
- [22] N. Heller, J. Dean, and N. Papanikolopoulos, "Imperfect segmentation labels: How much do they matter?" in *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*. Springer, 2018, pp. 112–120.
- [23] O. Petit, N. Thome, A. Charnoz, A. Hostettler, and L. Soler, "Handling missing annotations for semantic segmentation with deep convnets," in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*. Springer, 2018, pp. 20–28.
- [24] C. E. Brodley and M. A. Friedl, "Identifying mislabeled training data," *Journal of artificial intelligence research*, vol. 11, pp. 131–167, 1999.
- [25] J. CHATELET, "Aramis, le réseau français de radars pour la surveillance des précipitations," *La Météorologie [ISSN 0026-1181]*, 2003, *Série 8, N° 40 ; p. 44-52*, vol. 8, 01 2011.
- [26] R. J. Doviak and D. S. Zrnica, *Doppler radar and weather observations / Richard J. Doviak, Dusan S. Zrnica*. Academic Press Orlando, Fla, 1984.
- [27] P. Tabary, C. Augros, J.-L. Champeaux, J.-L. Chèze, D. Faure, D. Idziorek, R. Lorandel, B. Urban, and V. Vogt, "Le réseau et les produits radars de météo-france," *La Météorologie*, vol. 8ème série, pp. 15–27, 11 2013.
- [28] M. D. Dilmli, "Méthode de classification des séries temporelles: application à un réseau de pluviomètres," Ph.D. dissertation, 2019.
- [29] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 2010, pp. 249–256.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] J. Xie, L. Xu, and E. Chen, "Image denoising and inpainting with deep neural networks," in *Advances in neural information processing systems*, 2012, pp. 341–349.
- [32] S. Chaudhury and H. Roy, "Can fully convolutional networks perform well for general image restoration problems?" in *2017 Fifteenth IAPR International Conference on Machine Vision Applications (MVA)*. IEEE, 2017, pp. 254–257.
- [33] N. Cai, Z. Su, Z. Lin, H. Wang, Z. Yang, and B. W.-K. Ling, "Blind inpainting using the fully convolutional neural network," *The Visual Computer*, vol. 33, no. 2, pp. 249–261, 2017.
- [34] J. Donahue, P. Krähenbühl, and T. Darrell, "Adversarial feature learning," *arXiv preprint arXiv:1605.09782*, 2016.