

# SAR-UNet: Small Attention Residual UNet for Explainable Nowcasting Tasks

Mathieu Renault<sup>†</sup>, Siamak Mehrkanoon<sup>‡\*</sup>

<sup>†</sup> Department of Advanced Computing Sciences, Maastricht University, Maastricht, Netherlands

<sup>‡</sup> Department of Information and Computing Sciences, Utrecht University, Utrecht, Netherlands  
mathieu.renault@student.maastrichtuniversity.nl, s.mehrkanoon@uu.nl

**Abstract**—The accuracy and explainability of data-driven nowcasting models are of great importance in many socio-economic sectors reliant on weather-dependent decision making. This paper proposes a novel architecture called Small Attention Residual UNet (SAR-UNet) for precipitation and cloud cover nowcasting. Here, SmaAt-UNet is used as a core model and is further equipped with residual connections, parallel to the depthwise separable convolutions. The proposed SAR-UNet model is evaluated on two datasets, i.e., Dutch precipitation maps ranging from 2016 to 2019 and French cloud cover binary images from 2017 to 2018. The obtained results show that SAR-UNet outperforms other examined models in precipitation nowcasting from 30 to 180 minutes in the future as well as cloud cover nowcasting in the next 90 minutes. Furthermore, we provide additional insights on the nowcasts made by our proposed model using Grad-CAM, a visual explanation technique, which is employed on different levels of the encoder and decoder paths of the SAR-UNet model and produces heatmaps highlighting the critical regions in the input image as well as intermediate representations to the precipitation. The heatmaps generated by Grad-CAM reveal the interactions between the residual connections and the depthwise separable convolutions inside of the multiple depthwise separable blocks placed throughout the network architecture.

**Index Terms**—UNet, Precipitation Nowcasting, Cloud Cover Nowcasting, Deep Learning

## I. INTRODUCTION

An accurate precipitation nowcast is essential in many domains such as urbanization, agriculture and tourism. Sudden heavy rainfalls can lead to major catastrophes such as floods and landslides and have a tremendous economic impact. Early Warning Systems (EWSs) use the nowcast/forecast to limit damages due to climate hazards. Therefore, quick, accurate and trustworthy forecasts are invaluable, especially for the near future, i.e., nowcasting. For instance, highly competitive mechanical sports such as Formula 1 rely on forecasts as detailed as minute-by-minute predictions on different parts of a circuit to decide the race setup and strategy. The aviation sector depends on various factors such as precipitation, wind speed and direction to ensure a safe journey for the passengers daily. State-of-the-art NWP based models rely on numerical methods simulating the intrinsic physical dynamics of the climate. These simulations are often complex and require vast computational resources. In this context, the development of advanced data-driven models such as Deep Learning (DL) in

the past years has gained a lot of attention even has challenged the performance of NWP models [1]. Deep Learning based models have been previously successfully applied for forecasting tasks in a range of application domains including crop yield [2], solar irradiance [3], traffic [4] and weather [5–12]. The weather datasets are time-series with observations of a weather element at each time step. Recurrent Neural Networks (RNNs) are previously designed to include the order and time-dependency observed in datasets [13]. Over the last decade, literature has observed the development of Convolutional Neural Networks for several computer vision tasks [14]. In particular, LeNet-5 [15], GoogLeNet [16] and U-Net [17] have shown promising results for image segmentation or object detection, especially in the medical field [18]. The literature on traditional computer vision tasks such as image classification is rich. However, image-to-image forecasting is rather a new field, and many architectures have yet to be developed to address the problems as the one studied here.

Despite outperforming performance of deep learning approaches, they are not yet fully transparent in how they reach to a particular decision. Recently, researchers have developed Explainable AI (XAI) techniques to help the user understand the logic that leads the model to the prediction [19]. Standard methods to achieve this are gradient-based saliency maps, Class Activation Maps, or Excitation Backpropagation. These post hoc methods produce heatmaps by computing the layers' activation after the network is trained. Other popular post hoc algorithms are perturbation-based, where one determines the essential features by altering or removing parts of the input and quantifying the impact on the prediction performance. Here, we use post hoc algorithms used previously in image classification and segmentation to give additional insights on model predictions. This paper proposes a novel architecture called Small Attention Residual UNet (SAR-UNet) which uses SmaAt-UNet [20] as core model and equips it with a Residual Connection parallel to each Depthwise Separable Convolution on both encoder and decoder paths. We evaluate the introduced model on precipitation and cloud cover nowcasting over the Netherlands and France respectively. The main contributions of the paper are the following:

- We introduce a novel deep architecture for precipitation as well as cloud cover nowcasting. This network relies on the use of Depthwise Separable Convolution (DSC)

\*Corresponding author.

with Residual connections and Attention mechanisms in its encoder and decoder. The SAR-UNet outperforms its predecessor, the SmaAt-UNet, in precipitation and cloud cover nowcasting tasks.

- We utilize Grad-CAM, a visual explanation technique, in order to provide additional insights on the nowcasts. Given an input image, Grad-CAM produces heatmaps of the activations in different levels of our network. To the best of our knowledge, this is the first adaptation of Grad-CAM for an image-to-image nowcasting task.

This paper is organized as follows. A brief overview of the related research works is given in Section II. Section III, introduces the proposed SAR-UNet model and the explains the used visual explanation technique. The experimental settings and description of the used datasets are given in Section IV. The obtained results are discussed in Section V and the conclusion is drawn in Section VI.

## II. RELATED WORK

Data-driven weather forecasting has recently gained a lot of attention. Among many successful deep learning architectures, Recurrent Neural Networks [21] and Long short-term memory (LSTM) [22] have shown promising results in sequential data analysis. Xu et al. [23] combined a Generative Adversarial Network with LSTM architecture to produce a network suitable for prediction using satellite cloud maps. An attempt to design an LSTM block that would encompass the treatment of the spatial aspect of the given input image led to the Convolutional LSTM Network (ConvLSTM) [24], which has been successfully applied on precipitation nowcasting and outperformed the Fully Connected LSTM network.

UNet [17], the widely used architecture in computer vision, has been previously extended and modified in various ways. For instance, the authors in [20] introduced the SmaAt-UNet, using a UNet core model, by replacing all regular convolutions with Depthwise-Separable Convolution (DSC) [25] as well as adding Convolutional Block Attention Modules (CBAMs) [26] to the encoder part. The explored DSC mechanism in the SmaAt-UNet significantly reduced the number of trainable parameters of the network. Moreover, the utilized CBAM blocks create attention maps both for the channel dimensions and the spatial dimensions, therefore can better learn the inter-channel as well as inter-spatial relationships of the weather variables. Diakogiannisa et al. [27] proposed a UNet backbone enhanced with residual connections parallel to the convolutional blocks. The use of the residual connections avoids the problem of exploding and vanishing gradients, allowing the creation of deeper networks.

Explaining and interpreting deep weather forecasting models is still in its early stage. Abdellaoui et al. [28] used occlusion analysis to infer the importance of weather variables. The great advantage of this process is that it can be applied to any network architecture. In [7], the author quantifies the uncertainty in the prediction through Test Time Dropout, a method that approximate the Bayesian inference of Bayesian Neural Networks. An uncertainty map is then produced to

visualize and interpret the prediction. The authors in [29, 30] use Local Interpretable Model-agnostic Explanations (LIME) in a weather forecast context to interpret the decisions from their model. LIME approximates the deep learning network with a simple model as a linear one to understand the relationships in the weather features. Most of the explainable AI techniques adapted to CNN architectures have been applied to image classification or segmentation tasks. For instance, the authors in [31] introduced Gradient-weighted Class Activation Maps (Grad-CAM) to interpret a CNN's output.

## III. METHOD

### A. Proposed SAR-UNet model

1) *Architecture*: Fig. 1(a) shows a diagram that summarizes the proposed model. The model receives several images of size 288 by 288 stacked over the channels as input. Similar to UNet, its encoder-decoder transforms the data into spatially smaller images with more channels, then back into spatially larger images with fewer channels on each level before the output. Each level of the encoder is made of three transformations: Residual DSC Block (blue arrow), CBAM (yellow arrow), and 2x2 Max Pooling (red arrow). The CBAM's output is used as input for Max Pooling layer as well as it is send to the corresponding level (purple arrow) of the decoder part, where it is concatenated to the images from the lower level. The 2x2 Max Pooling reduces the spatial size of the images by 2, and its output is used as input to the next (lower) level of the encoder. A decoder level begins with 1x1 convolution followed by upsampling, represented by the green arrow in the diagram. The 1x1 convolution aims to halve the number of channels, while the upsampling doubles the spatial size of the images. The images are then concatenated to the CBAM's output over the channels and fed into a Residual DSC block. Note that the Residual DSC blocks double the number of channels in the encoder whilst they halve them in the decoder. A 1x1 output convolution (black arrow) layer is used in the end to obtain a single output image with the prediction.

2) *Residual DSC Block*: Residual connections act as a shortcut in the network and therefore make it easier to optimize the network parameters which itself can result in improving the training of the network and its accuracy. In addition, the residual connections allow for deeper networks to be trained without suffering from the vanishing gradients problem. In the proposed Residual DSC Block, shown in Fig. 1(b), we combine the Residual connection with the Depthwise Separable Convolution mechanism. The input goes through DSC with 3x3 kernel, Batch Normalization and ReLU twice. Moreover, this same input goes, in parallel, through a residual connection with 1x1 convolution to match the output channels of DSC. The output of the Residual DSC block is obtained by summing the output of the residual and the DSC paths.

3) *CBAM*: We place CBAMs after each Residual DSC block in the encoder. As opposed to [20], where the CBAM's output is only used in the skip connection, here it is also used as the input of the next level. The last adaptation of the SmaAt-UNet resides in the number of channels in the bottleneck

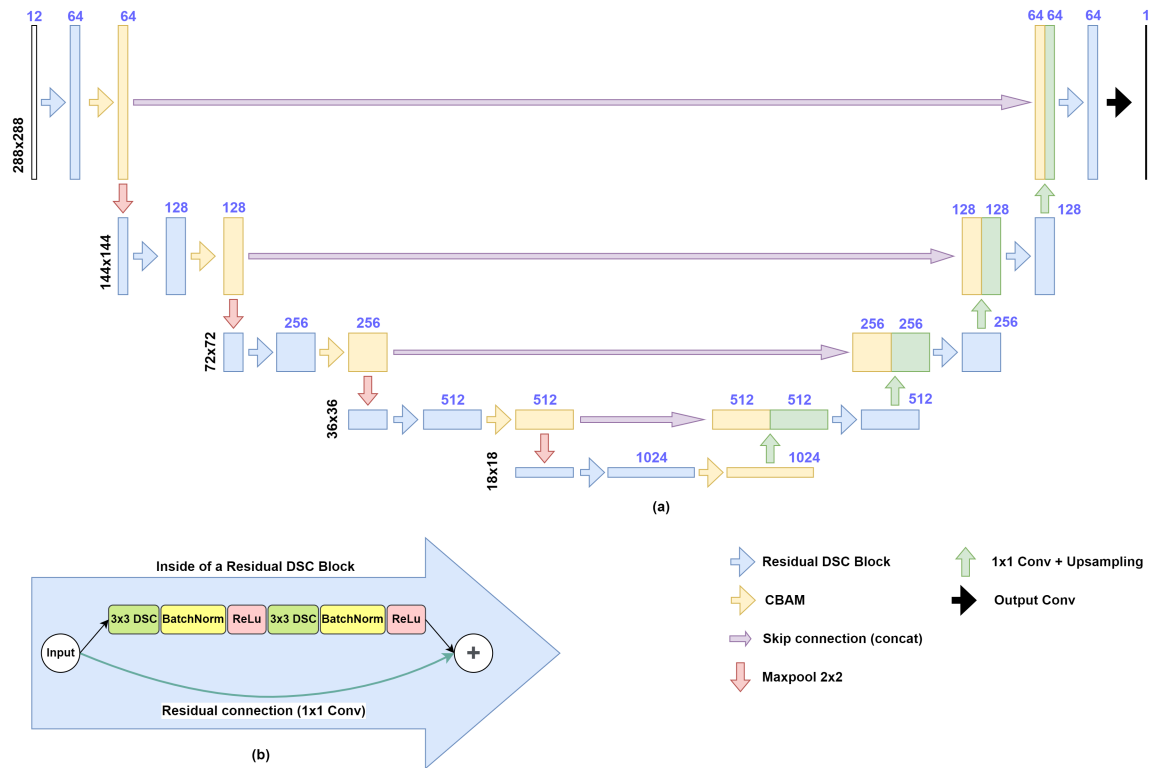


Fig. 1. SAR-UNet architecture. Arrows stand for a given transformation applied to the data, while the rectangles show the shape of the data after each step. The height and width of the images are shown on the left side at each level of the network. The number of filters used in the convolutions can be seen through the number of channels in the resulting images at each step. In the bottom left corner, we show the inside of a Residual DSC Block: the input goes, in parallel, through two DSCs on one side, through 1x1 convolution on the other, and both sides are then summed. The network uses 12 channels as input, corresponding to 60 minutes of image data.

layer. Our network has 1024 channels, twice as many as in the bottleneck of the SmaAt-UNet. For this reason, we use 1x1 Convolutions before each upsampling operation to reduce the channels in the next decoder level. These adaptations induce a slight increase in the network’s total number of trainable parameters. We compare the total number of parameters of UNet, SmaAt-UNet and SAR-UNet in Fig. 2.




Fig. 2. Comparison of the number of trainable parameters between networks

## B. Training

For training we followed the guidelines of [20]. The initial learning rate is set to 0.001 with a learning rate scheduler dividing the learning rate by 10 every time the validation loss has not decreased in four consecutive epochs. The Adam optimizer is used for training the model. The maximum

number of epochs is set to 200, with an early stopping criterion of 15 epochs, effectively stopping training if there was no improvement over the last 15 epochs. The models are trained using a batch size of 6 on the Google Colab Pro platform with a GPU (Nvidia Tesla P100).

## C. Evaluation

To evaluate the performance of the examined models, we use the Mean Squared Error (MSE) which is also used as the loss function during training. We also set a range of metrics after binarizing the output image. Following the lines of [20], using a threshold of  $0.5\text{mm}/h$ , we binarize each pixel of the output image and count the number of True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN). Thanks to this, we get the Accuracy, Precision, Recall, and the F1-score of the model. We emphasize the MSE as an indicator of performance as this is a regression task, and the objective is to predict the exact quantity of precipitation in millimeters. Moreover, we compare the performance of the proposed model with that of SmaAt-UNet and the Persistence method, which uses the last input image of a sequence as the prediction image.

## D. Explainability

Explainability is an essential part of deep learning now-casting. Here, we explain the predictions through activation

heatmaps. These heatmaps show which areas of the input are responsible for the high values in the output, thus making the nowcast more transparent. Grad-CAM was originally developed by [31] to produce a visual explanation of a classification network. It takes the gradient of the score for a class with respect to feature map activations to obtain neuron importance weights. Then it performs a linear combination of the forward activation maps followed by ReLU to obtain the final result. As mentioned before, after binarizing the output radar image for evaluation metrics, we obtain an image that is segmented with rain and no rain classes. We thus obtain explanations for our nowcasting task by transforming it into an image segmentation task. We set the algorithm to produce heatmaps for the rain class on the entire image.

#### IV. EXPERIMENTS

##### A. Precipitation nowcasting

We use the same precipitation map dataset as in [8, 20]. This dataset is collected from 2016 to 2019 by the Royal Netherlands Meteorological Institute (Koninklijk Nederlands Meteorologisch Instituut) with two radars in Den Helder and De Bilt, The Netherlands, capturing precipitation intensities every 5 minutes. Every pixel of the resulting image corresponds to one square kilometer of land, and its numerical value stands for the amount of rainfall, in hundredth of a millimeter. That means a pixel value of 1 translates to 0.01mm of rainfall. We follow the preprocessing steps of [20], by cropping the images to a size of 288x288 pixels, centered on the Netherlands. As our interest is focused on images with rain, similar to [20], we select images with a threshold set to 50% of the total pixels having a value strictly higher than 0.

The sequence of images used as input for the network are stacked over the channel dimension. We have conducted a set of experiments to quantify the nowcasting performance in different situations. In particular, we use three different input sizes with 6, 12 and 18 channels, corresponding to 30, 60 and 90 minutes of data input for the network. For each of the input sizes, we conduct a range of nowcasting tasks: 30, 60, 90, 120 and 180 minutes ahead. This amounts to a total of 15 different nowcasting tasks. We used checkpoints to save the model after each epoch during training and for each setup we selected the checkpoint with lowest loss on the validation set. An example of the nowcastings obtained by our proposed SAR-UNet model is shown in Fig. 3.

##### B. Cloud cover nowcasting

We also examine our model on the French cloud cover dataset used in [20]. The images in this dataset are binary and of size 256x256; each pixel has a value 1 if there is a cloud and 0 if there is none. The images are recorded every 15 minutes. Following the lines of [20], we output six images (corresponding to the nowcasts for the next 1.5 hours ahead) using four images (corresponding to the past 1 hour data) as input. An example of the four input images is shown in Fig. 4.




Fig. 3. Example of the nowcasting tasks, using 30 minutes input data. The first column presents the 6 input images. The second column shows the ground truth for nowcasting from 30 to 180 minutes ahead. The proposed model predictions are shown in the last column.




Fig. 4. Example of the cloud cover data input. The time interval between each image is 15 minutes.



## V. RESULTS AND DISCUSSION

### A. Precipitation and cloud cover nowcasting

This section presents the results obtained for the precipitation and cloud cover nowcasting tasks. The obtained results of the precipitation task are summarized in Table I, which show that SAR-UNet model is better than the other tested models in 13 of the 15 different setups when considering only the MSE, which we emphasize this by plotting the average MSE (average of all minutes ahead predictions) by the amount of input used for each model in Fig. 5. In this

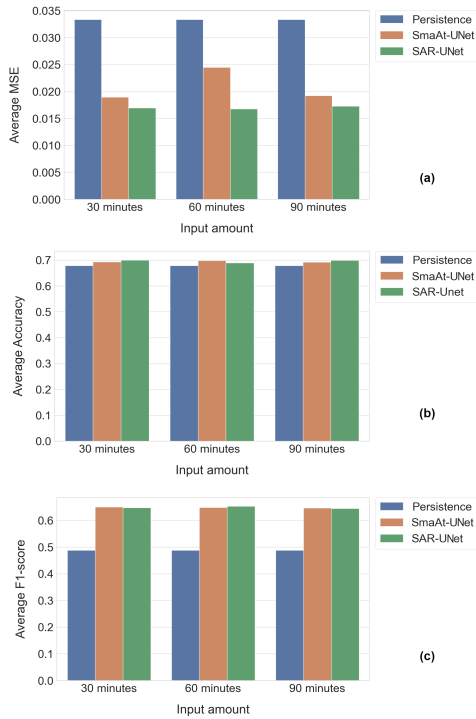


Fig. 5. The comparison of different metrics obtained by averaging over multi-step ahead nowcasts using different amounts of data as input. (a) The average MSE. (b) The average Accuracy. (c) The average F1-score.

Figure, we compare the performance of the examined models and visualize the superiority of SAR-UNet model over the SmaAt-UNet and the Persistence models. In addition, there does not seem to be a substantial increase or decrease in the model performance when varying input data amounts. Table I shows that nowcasting further away in the future is a more challenging task, and in general the losses increase as the minutes ahead increases. For Precision, Recall and F1-score, the SAR-UNet is better than the SmaAt-UNet in about as many setups as when the SmaAt-UNet is better than the SAR-UNet. Therefore, SAR-UNet and SmaAt-UNet are comparable in terms of these tested metrics. This result might be explained by the arbitrary choice of a threshold for binarization of the prediction used in calculating these metrics. The obtained metrics for the cloud cover nowcasting task are tabulated in Table II. Similar to precipitation task, here SAR-UNet in most of the cases outperforms the SmaAt-UNet and Persistence models.

### B. Activation heatmaps

The collection of heatmaps obtained using Grad-CAM bring additional clarity to how SAR-UNet functions and what parts of the images are responsible for the prediction. In Fig. 6 and Fig. 7 we show the heatmaps of network's activations for the precipitation nowcasting. These heatmaps are obtained using the same data shown in Fig. 3. The figures contain heatmaps of the activations of different inner layers of the SAR-UNet to explain how they operate. Therefore, we output a heatmap for the whole Residual DSC blocks, as well as each DSC sequence and Residual connection inside of these blocks, and finally for all CBAMs in the encoder path of the model. We split the plots between the encoder part in Fig. 6 and decoder part in Fig. 7 where the complete path followed by the data, from the very beginning to the very end are shown. This provides additional insights on which part of the input image is most important at each layer and level of the model and that how these information are combined to give a final prediction.

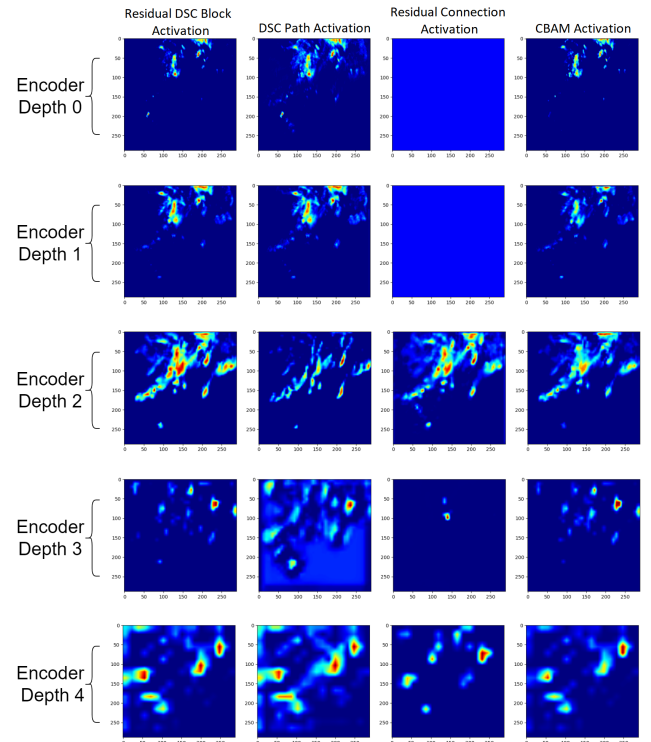


Fig. 6. Activation heatmaps obtained using Grad-CAM on all pixels predicted with rain after binarization. The first column represents the activation from an entire Residual DSC Block, while the second and third columns focus on parts of it, i.e., the activation of a DSC and the activation of the Residual connection respectively. The activations of the CBAM are shown on the last column. Each row stands for a level (depth) of the encoder part of SAR-UNet model. Thus, the first row is the first transformation for the input, while the bottom row is the last transformation of the encoder part.

We start by analyzing the first column of the figures, representing the Residual DSC Blocks. The first three levels of the encoder in Fig. 6 seem to be focused on the areas where the last input image presents the highest precipitation amounts. In particular, first only the areas with very high precipitation

TABLE I  
COMPARISON OF THE MODELS FOR THE **PRECIPITATION NOWCASTING TASKS**.  $\uparrow$  MEANS THE BEST VALUE FOR A METRIC IS THE HIGHEST WHILE  $\downarrow$  MEANS THE BEST VALUE IS THE LOWEST.

Input amount	Prediction time	Models	MSE $\downarrow$	Precision $\uparrow$	Recall $\uparrow$	Accuracy $\uparrow$	F1 score $\uparrow$	
30 minutes	30 minutes ahead	Persistence	0,0249	0,678	0,643	0,756	0,66	
		SmaAt-UNet	0,0151	0,675	<b>0,883</b>	0,8	0,765	
		SAR-UNet	<b>0,0124</b>	<b>0,73</b>	0,835	<b>0,826</b>	<b>0,779</b>	
	60 minutes ahead	Persistence	0,0318	0,603	0,522	0,698	0,56	
		SmaAt-UNet	<b>0,0160</b>	<b>0,634</b>	0,749	<b>0,749</b>	0,687	
		SAR-UNet	0,0161	0,607	<b>0,799</b>	0,736	<b>0,69</b>	
	90 minutes ahead	Persistence	0,0360	0,558	0,43	0,665	0,486	
		SmaAt-UNet	0,0177	<b>0,561</b>	0,682	<b>0,687</b>	0,616	
		SAR-UNet	<b>0,0170</b>	0,532	<b>0,812</b>	0,668	<b>0,643</b>	
	120 minutes ahead	Persistence	0,0375	<b>0,532</b>	0,355	0,648	0,426	
		SmaAt-UNet	0,0257	0,506	<b>0,769</b>	0,639	<b>0,61</b>	
		SAR-UNet	<b>0,0189</b>	0,524	0,666	<b>0,654</b>	0,586	
	180 minutes ahead	Persistence	0,0368	0,478	0,229	<b>0,624</b>	0,31	
		SmaAt-UNet	<b>0,0201</b>	0,463	<b>0,75</b>	0,588	<b>0,573</b>	
		SAR-UNet	0,0203	<b>0,483</b>	0,614	0,616	0,54	
	60 minutes	30 minutes ahead	Persistence	0,0249	0,678	0,643	0,756	0,66
			SmaAt-UNet	0,0248	0,677	<b>0,878</b>	0,801	0,764
			SAR-UNet	<b>0,0120</b>	<b>0,697</b>	0,868	<b>0,813</b>	<b>0,774</b>
60 minutes ahead		Persistence	0,0318	0,603	0,522	0,698	0,56	
		SmaAt-UNet	0,0166	0,582	<b>0,843</b>	0,719	<b>0,688</b>	
		SAR-UNet	<b>0,0163</b>	<b>0,623</b>	0,747	<b>0,74</b>	0,679	
90 minutes ahead		Persistence	0,0360	<b>0,558</b>	0,43	0,665	0,486	
		SmaAt-UNet	0,0314	0,553	0,737	<b>0,684</b>	<b>0,632</b>	
		SAR-UNet	<b>0,0185</b>	0,542	<b>0,747</b>	0,674	0,628	
120 minutes ahead		Persistence	0,0375	0,532	0,355	0,648	0,426	
		SmaAt-UNet	0,0196	<b>0,54</b>	0,653	<b>0,667</b>	0,591	
		SAR-UNet	<b>0,0176</b>	0,485	<b>0,831</b>	0,613	<b>0,612</b>	
180 minutes ahead		Persistence	0,0368	0,478	0,229	<b>0,624</b>	0,31	
		SmaAt-UNet	0,0302	<b>0,488</b>	0,685	0,619	0,57	
		SAR-UNet	<b>0,0196</b>	0,477	<b>0,712</b>	0,607	<b>0,571</b>	
90 minutes		30 minutes ahead	Persistence	0,0249	0,678	0,643	0,756	0,66
			SmaAt-UNet	0,0131	<b>0,720</b>	0,84	<b>0,821</b>	<b>0,776</b>
			SAR-UNet	<b>0,0125</b>	0,703	<b>0,851</b>	0,813	0,77
	60 minutes ahead	Persistence	0,0318	0,603	0,522	0,698	0,56	
		SmaAt-UNet	0,0197	0,576	0,792	0,709	0,667	
		SAR-UNet	<b>0,0171</b>	<b>0,615</b>	<b>0,802</b>	<b>0,742</b>	<b>0,696</b>	
	90 minutes ahead	Persistence	0,0360	<b>0,558</b>	0,43	0,665	0,486	
		SmaAt-UNet	0,0205	0,547	<b>0,752</b>	<b>0,679</b>	<b>0,633</b>	
		SAR-UNet	<b>0,0176</b>	0,539	0,749	0,672	0,627	
	120 minutes ahead	Persistence	0,0375	<b>0,532</b>	0,355	0,648	0,426	
		SmaAt-UNet	0,0211	0,499	<b>0,754</b>	0,631	<b>0,601</b>	
		SAR-UNet	<b>0,0191</b>	0,52	0,659	<b>0,65</b>	0,581	
	180 minutes ahead	Persistence	0,0368	0,478	0,229	<b>0,624</b>	0,31	
		SmaAt-UNet	0,0217	<b>0,487</b>	<b>0,649</b>	0,619	<b>0,556</b>	
		SAR-UNet	<b>0,0202</b>	0,484	0,641	0,616	0,552	

TABLE II  
COMPARISON OF THE MODELS FOR THE **CLOUD COVER NOWCASTING**.  $\uparrow$  MEANS THE BEST VALUE FOR A METRIC IS THE HIGHEST WHILE  $\downarrow$  MEANS THE BEST VALUE IS THE LOWEST.

Models	MSE $\downarrow$	Precision $\uparrow$	Recall $\uparrow$	Accuracy $\uparrow$	F1 score $\uparrow$
Persistence	0.1491	0.872	0.872	0.851	0.872
SmaAt-UNet	0.0794	<b>0.892</b>	0.921	0.889	0.906
SAR-UNet	<b>0.0787</b>	<b>0.892</b>	<b>0.923</b>	<b>0.890</b>	<b>0.907</b>

are activated in the encoder depth 0, progressing to a larger activated area in encoder depth 2 corresponding to areas with rain in the input. The deeper levels of the network, shown in Fig. 6 and 7, namely Encoder depth 3,4 and Decoder depth 3, appear to activate in more abstract areas, with a few dots being significantly activated, leaving the rest of the image

inactivated. The final three levels (depth 2,1 and 0) of the decoder in Fig. 7 are much more similar to the prediction made by the network. These layers smooth out the prediction so that areas vary gradually from little rain to substantial rain, giving a more realistic appearance to the prediction.

The activation from a Residual DSC block is generally similar to the ones from the DSC path and the Residual connection as it is the sum of both. Nevertheless, it appears in Fig. 6 that the DSC path is closer to the output of the Residual DSC block, while the Residual connection is less important with the first two depth levels being uniformly activated. On the other hand, in Fig. 7 it is the Residual connection that is activated similarly to the Residual DSC Block of the same row. Therefore, in the encoder and decoder parts, we can infer that the two paths of a Residual DSC Block switch

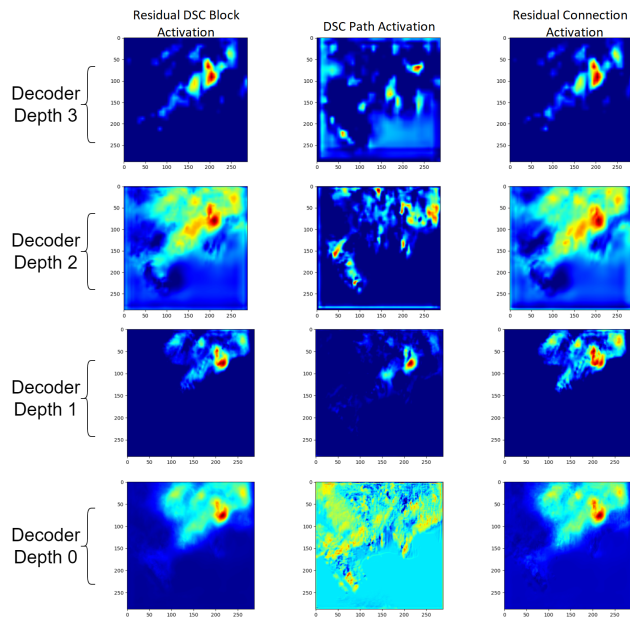


Fig. 7. Activation heatmaps obtained through the use of Grad-CAM on all pixels predicted with rain after binarization. The first column represents the activation from an entire Residual DSC Block, while the second and third columns focus on parts of it, i.e., the activation of a DSC and the activation of the Residual connection respectively. Each row stands for a level of the decoder part of SAR-UNet model. Thus, the first row is the bottom level of the decoder, while the final row is the last transformation before the output convolutional layer.

roles and their importance change gradually. Finally, the fourth column in Fig. 6 presents the activations of the CBAMs in the Encoder part of the network. We observe that the CBAMs activate almost exactly as the Residual DSC Blocks. The visible difference is in the intensity of the activation, especially near the points where the activation is the maximum, i.e., the red and orange areas. The activation of the CBAM is less intense in these areas, leading to fewer red areas and more yellow and green ones. Fig. 8 shows the activation heatmaps of Encoder depth 1 and 4, and Decoder depth 3 and 1 of the SAR-UNet for the cloud cover dataset. We notice the Encoder depth 1, placed at the beginning of the network, is activated at the borders between cloud and non-cloud zones with high precision. Encoder depth 4 is also activated on these borders but in wider patches, leading to larger red zones on the heatmap. In Decoder depth 2, we observe activation zones in the center of the cloudy areas of the image. It is focused on the inside of the cloud area delimited in the previous layers. Decoder depth 1 is one of the final layers of the network. Its activation heatmap is very different from the other shown heatmaps. This layer’s activation covers almost entirely the cloud areas of the image. It is therefore a combination of the previous layers, with the borders and the centre of the clouds all activated areas.

## VI. CONCLUSION

In this paper a novel Small Attention Residual UNet (SAR-UNet) is proposed for weather element nowcasting tasks.

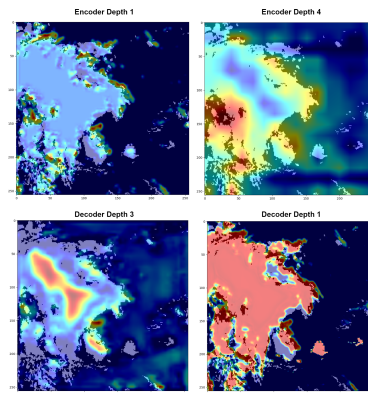


Fig. 8. The obtained activation heatmaps of four different residual DSC blocks within the SAR-UNet, for the cloud cover nowcasting task.

The model is based on a U-shape convolutional network that combines Depthwise Separable Convolutions, Residual Connections and the Convolutional Block Attention Modules to outperform its predecessor, the SmaAt-UNet. The proposed model is evaluated on two nowcasting tasks, i.e., precipitation and cloud cover nowcasting. The experimental results demonstrate that SAR-UNet model outperforms the other tested models for the studied datasets. In order to shed lights on the explainability of the inner workings of the proposed SAR-UNet model, we have visualized the activation heatmaps for the nowcasts using Grad-CAM technique. The implementation of our SAR-UNet model is available at GitHub<sup>1</sup>.

## REFERENCES

- [1] M. G. Schultz, C. Betancourt, B. Gong, F. Kleinert, M. Langguth, L. H. Leufen, A. Mozaffari, and S. Stadler, “Can deep learning beat numerical weather prediction?” *Philosophical Transactions of the Royal Society A*, vol. 379, no. 2194, p. 20200097, 2021.
- [2] A. Wolanin, G. Mateo-García, G. Camps-Valls, L. Gómez-Chova, M. Meroni, G. Duveiller, Y. Liangzhi, and L. Guanter, “Estimating and understanding crop yields with explainable deep learning in the indian wheat belt,” *Environmental research letters*, vol. 15, no. 2, p. 024019, 2020.
- [3] H. Wang, R. Cai, B. Zhou, S. Aziz, B. Qin, N. Voropai, L. Gan, and E. Barakhtenko, “Solar irradiance forecasting based on direct explainable neural network,” *Energy Conversion and Management*, vol. 226, p. 113487, 2020.
- [4] Z. Zhao, W. Chen, X. Wu, P. C. Chen, and J. Liu, “Lstm network: a deep learning approach for short-term traffic forecast,” *IET Intelligent Transport Systems*, vol. 11, no. 2, pp. 68–75, 2017.
- [5] R. L. Cañar, A. Fontaine, P. L. Morillo, and S. El Yacoubi, “Deep learning to implement a statistical weather forecast for the andean city of quito,” in *2020 IEEE ANDESCON*. IEEE, 2020, pp. 1–6.

<sup>1</sup><https://github.com/mathieurenault1/SAR-UNet>

- [6] L. Zhu, G. Zhu, L. Han, N. Wang *et al.*, “The application of deep learning in airport visibility forecast,” *Atmospheric and Climate Sciences*, vol. 7, no. 03, p. 314, 2017.
- [7] Y. Yang and S. Mehrkanoon, “AA-Transunet: Attention augmented transunet for nowcasting tasks,” in *the proc. of IEEE International Joint Conference on Neural Networks (IEEE-IJCNN)*, 2022, pp. 1–8.
- [8] J. G. Fernández and S. Mehrkanoon, “Broad-unet: Multi-scale feature learning for nowcasting tasks,” *Neural Networks*, vol. 144, pp. 419–427, 2021.
- [9] J. G. Fernández, I. A. Abdellaoui, and S. Mehrkanoon, “Deep coastal sea elements forecasting using unet-based models,” *Knowledge-Based Systems*, vol. 252, p. 109445, 2022.
- [10] D. Aykas and S. Mehrkanoon, “Multistream graph attention networks for wind speed forecasting,” in *the proc. of IEEE Symposium Series on Computational Intelligence (IEEE-SSCI)*, 2021, pp. 1–8.
- [11] T. Stańczyk and S. Mehrkanoon, “Deep graph convolutional networks for wind speed prediction,” in *the proc. of European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, 2021, pp. 147–152.
- [12] S. Mehrkanoon, “Deep shared representation learning for weather elements forecasting,” *Knowledge-Based Systems*, vol. 179, pp. 120–128, 2019.
- [13] S. Razavi, “Deep learning, explained: Fundamentals, explainability, and bridgeability to process-based modelling,” *Environmental Modelling & Software*, vol. 144, p. 105159, 2021.
- [14] S. Lawrence, C. L. Giles, A. C. Tsoi, and A. D. Back, “Face recognition: A convolutional neural-network approach,” *IEEE transactions on neural networks*, vol. 8, no. 1, pp. 98–113, 1997.
- [15] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [16] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [17] 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.
- [18] P. Natekar, A. Kori, and G. Krishnamurthi, “Demystifying brain tumor segmentation networks: interpretability and uncertainty analysis,” *Frontiers in computational neuroscience*, vol. 14, p. 6, 2020.
- [19] G. Ras, N. Xie, M. van Gerven, and D. Doran, “Explainable deep learning: A field guide for the uninitiated,” *Journal of Artificial Intelligence Research*, vol. 73, pp. 329–397, 2022.
- [20] K. Trebing, T. Stanczyk, and S. Mehrkanoon, “SmaAt-UNet: Precipitation nowcasting using a small attention-unet architecture,” *Pattern Recognition Letters*, vol. 145, pp. 178–186, 2021.
- [21] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, “Learning internal representations by error propagation,” California Univ San Diego La Jolla Inst for Cognitive Science, Tech. Rep., 1985.
- [22] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [23] Z. Xu, J. Du, J. Wang, C. Jiang, and Y. Ren, “Satellite image prediction relying on gan and lstm neural networks,” in *ICC 2019-2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–6.
- [24] X. Shi, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, “Convolutional lstm network: A machine learning approach for precipitation nowcasting,” *Advances in neural information processing systems*, vol. 28, 2015.
- [25] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, “Mobilenets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint arXiv:1704.04861*, 2017.
- [26] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, “Cbam: Convolutional block attention module,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 3–19.
- [27] F. I. Diakogiannis, F. Waldner, P. Caccetta, and C. Wu, “Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 162, pp. 94–114, 2020.
- [28] I. A. Abdellaoui and S. Mehrkanoon, “Deep multi-stations weather forecasting: explainable recurrent convolutional neural networks,” *arXiv preprint arXiv:2009.11239*, 2020.
- [29] S. Gao and Y. Wang, “Explainable deep learning powered building risk assessment model for proactive hurricane response,” *Risk analysis*, 2022.
- [30] M. Zdravković, I. Ćirić, and M. Ignjatović, “Explainable heat demand forecasting for the novel control strategies of district heating systems,” *Annual Reviews in Control*, 2022.
- [31] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.