

Snow Avalanche Segmentation in SAR Images With Fully Convolutional Neural Networks

Filippo Maria Bianchi , Jakob Grahn, Markus Eckerstorfer, Eirik Malnes, and Hannah Vickers 

Abstract—Knowledge about frequency and location of snow avalanche activity is essential for forecasting and mapping of snow avalanche hazard. Traditional field monitoring of avalanche activity has limitations, especially when surveying large and remote areas. In recent years, avalanche detection in Sentinel-1 radar satellite imagery has been developed to improve monitoring. However, the current state-of-the-art detection algorithms, based on radar signal processing techniques, are still much less accurate than human experts. To reduce this gap, we propose a deep learning architecture for detecting avalanches in Sentinel-1 radar images. We trained a neural network on 6345 manually labeled avalanches from 117 Sentinel-1 images, each one consisting of six channels that include backscatter and topographical information. Then, we tested our trained model on a new synthetic aperture radar image. Comparing to the manual labeling (the gold standard), we achieved an *F1* score above 66%, whereas the state-of-the-art detection algorithm sits at an *F1* score of only 38%. A visual inspection of the results generated by our deep learning model shows that only small avalanches are undetected, whereas some avalanches that were originally not labeled by the human expert are discovered.

Index Terms—Convolutional neural networks (CNNs), deep learning, saliency segmentation, Sentinel-1 (S1), snow avalanches, synthetic aperture radar (SAR).

I. INTRODUCTION

KNOWLEDGE about the spatio-temporal distribution of snow avalanche (hereafter referred to as avalanche) activity in a given region is critical for avalanche forecasting and hazard mapping. An increase in avalanche activity or magnitude of releasing avalanches leads to an increase in avalanche risk. Conventionally, avalanche activity is primarily monitored through field measurements, which is time-consuming, expensive, and can only be done for very few accessible areas. Monitoring avalanche activity using satellite-borne synthetic aperture radar (SAR) has, therefore, gained considerable interest in recent years. SAR products enable continuous covering of very large areas, regardless of light and weather conditions [1].

Manuscript received October 6, 2020; accepted November 2, 2020. Date of publication November 10, 2020; date of current version January 6, 2021. This work was supported by in part the ‘Satskred’ Project (NIT.01.20.5), financed by the Norwegian Space Centre, the Norwegian Water and Energy Resource Directorate, and the Norwegian Public Road Administration. (Corresponding author: Filippo Maria Bianchi.)

Filippo Maria Bianchi is with the Department of Mathematics and Statistics, UiT The Arctic University of Norway, 9019 Tromsø, Norway, and also with the NORCE Norwegian Research Centre AS, 5008 Bergen, Norway (e-mail: filippo.m.bianchi@uit.no).

Jakob Grahn, Markus Eckerstorfer, Eirik Malnes, and Hannah Vickers are with the NORCE Norwegian Research Centre AS, 5008 Bergen, Norway (e-mail: jgra@norceresearch.no; maec@norceresearch.no; eima@norceresearch.no; havi@norceresearch.no).

Digital Object Identifier 10.1109/JSTARS.2020.3036914

An experienced operator can identify avalanche debris (the depositional part of an avalanche) in SAR change detection composites (showing temporal radar backscatter change) with high accuracy. On the other hand, automatic signal processing methods based on radar backscatter thresholding and segmentation often fail and produce a large number of false alarms due to the highly dynamic nature of snow in the SAR images [2]. A key limitation of classical segmentation methods is that they mainly focus only on pixelwise information in radar backscatter, without accounting for the contextual information around the pixel and high-level features, such as the shape and the texture of avalanche debris. Also, local topography in which the avalanches occur is largely disregarded since it has only been used to mask out areas where avalanches are unlikely to occur. However, the occurrence of avalanches is strongly correlated to topographical conditions and avalanche debris exhibits characteristic shapes, which should both be taken into account when performing the detection.

Convolutional neural networks (CNNs) have attracted considerable interest for their ability to model complex contextual information in images [3]. Prominent examples in remote sensing are terrain surface classification [4], [5], categorization of aerial scenes [6], detection of changes in the terrain over time from SAR and optical satellite sensors [7], [8], and segmentation of objects from airborne images [9], [10]. Nevertheless, few research efforts have been devoted so far toward detecting avalanche activity from SAR data, which remains an open and challenging endeavor. In our previous work [11], we proposed a deep learning architecture to perform binary classification of avalanches in Northern Norway. In particular, we used a CNN to classify fixed-size patches of SAR images in two classes: 1 if the patch contains at least one avalanche, or 0 otherwise. Our approach was successively adopted later on for SAR-borne avalanche detection in the Alps [12] and in other locations in Norway [13]. As a major limitation, patchwise classification cannot determine the presence of multiple avalanches within the same patch. Also, the results are heavily influenced by the patch size, which makes it difficult to evaluate the detection performance. In particular, for large windows is easier to correctly predict the presence of at least one avalanche, but the resolution of the detection is too coarse and not very useful.

In this work, we approach avalanche detection as a saliency segmentation task, where the classification is not done at the patch level, but rather at the individual pixel level. We adopt a fully convolutional network (FCN) architecture, which generates for each input image a segmentation mask. This solves the

drawback of the dependence from the window size and makes it possible to determine the exact location of the avalanches. Our work provides important contributions to the fields of earth science, remote sensing, and avalanche risk assessment.

- 1) We explore, for the first time, the capability of deep learning models in detecting the presence of avalanches in SAR products at a pixel granularity and surpass the current state-of-the-art avalanche detection algorithm [2]. Our work was possible thanks to a large dataset of SAR products manually annotated by an avalanche expert.
- 2) We advance our knowledge on topographical features to identify areas where avalanches are highly likely to occur.

Notably, we introduce a new topographical feature, called *potential angle of reach* (PAR), which indicates how likely it is for an avalanche to reach a specific location. We do not use the PAR to filter input images or detection results, but we rather provide the PAR as an exogenous input feature to the FCN. We first estimate how informative is the PAR in the discrimination of avalanche and not-avalanche pixels. Then, in the experimental section, we evaluate how much the detection performance of the deep learning model improves when providing the FCN with the PAR feature map.

II. SAR DATASET

The dataset consists of data from the Sentinel-1 (S1) satellites. In particular, data acquired in the interferometric wide-swath mode was considered in terms of the ground range detected product. In total, 118 SAR scenes covering two mountainous regions in Northern Norway in the period October 2014–April 2017.

A. Preprocessing

Each SAR product was 1) radiometrically calibrated to radar backscatter (sigma naught) values, 2) spatially downsampled from 10 to 20 m resolution, 3) geocoded onto a 20 m resolution UTM-grid (EPSG:32633) using a 10 m resolution digital elevation model (DEM) [14], 4) radiometrically transformed to decibel (dB) values and clipped to range values from -25 to -5 dB, to remove noise and restrict the range of the backscatter to intervals where avalanches are visible. The preprocessed products were then grouped by their satellite geometry, such that the scenes within a group have the same viewing geometry. For each group, scenes were paired chronologically into *reference* and *activity* image pairs. For the two S1 satellites, the reference image is acquired six days before the activity image (12 days before the launch of S1B in 2015). The resulting products have an approximate size of 11.500×5.500 pixels, and each pixel covers 20×20 m.

B. Generation of SAR Features

We considered three SAR features to generate the images to be processed by the deep learning model. The first two are the difference of the horizontal and vertical polarization between the reference and the activity image: $VV = VV^{\text{activ}} - VV^{\text{ref}}$, $VH = VH^{\text{activ}} - VH^{\text{ref}}$. The difference values are rescaled to $[0, 1]$ [see Fig. 1(a) and (b)]. The third feature is the pointwise

product of the difference images squared: $VVVH = VV^2 * VH^2$ [see Fig. 1(c)]. We did not consider radar shadow, layover masks, or land masks depicting avalanche runout zones, which are not available for all areas.

C. Labeling

For each product, a human expert generated a binary segmentation mask that indicates whether a pixel in the product is an avalanche or not. To create the segmentation mask, the human expert looked for changes in a difference image obtained from the following three channels: $R[VV^{\text{reference}}]$, $G[VV^{\text{activity}}]$, $B[VV^{\text{reference}}]$. We considered visual detection as the golden standard and used it as ground truth to train and evaluate our deep learning model. The whole dataset contains a total of 6345 avalanches; 3 667 355 474 pixels are classified as “non-avalanche” and 712 945 (0.000194% of the total) as “avalanche.”

III. TOPOGRAPHICAL FEATURES

Since avalanches are caused by steep terrain, the topography is an important factor to determine where avalanches can appear. In particular, the local slope needs to be steep enough for an avalanche to release and the slope typically needs to flatten out for the avalanche to stop. Therefore, it is reasonable to consider such information when performing the detection task and we generated two feature maps from the DEM, which is available for the entire Norway in 10 m pixel resolution. The first is the local slope angle of the terrain; the second is a new topographical feature introduced in this work, which is called PAR.

A. Slope Angle

The slope angle feature map is directly computed by taking the gradient of the DEM [see Fig. 1(d)]. The terrain slope is often considered when detecting avalanches, as they typically start in terrain between 35 and 45 degrees steepness and deposit on less steep slope angles. In previous work, the slope was used to derive a runout mask that indicated where avalanches are most likely to deposit [2]. Since the mask is applied to filter out areas in a preprocessing operation, the slope feature did not contribute to the actual detection. Most importantly, since run-out masks are obtained by manual thresholding the slope, if a wrong threshold is chosen some avalanches will not be detected. To address this issue, we provide the slope as an additional layer of the input image and let our neural network learn how to optimally exploit it to solve the segmentation task, without applying manually chosen thresholds.

Fig. 2 shows that the distribution of the slope angle is different for the avalanche and nonavalanche pixels in our dataset. In particular, avalanche pixels are mostly concentrated around $[20, 35]$ degrees. The difference in the two distributions indicates that the slope angle can be exploited to discriminate between “avalanche” and “nonavalanche” classes.

B. Potential Angle of Reach

The angle of reach of an avalanche, sometimes denoted α and referred to as the alpha-angle, indicates how far an avalanche travels from its triggering point in relation to the descent it




Fig. 1. (a), (b) SAR features obtained from the difference in the VV and VH channels. (c) Product VVVH of the squared differences. (d), (e) Slope and PAR feature maps. Only a small area ($1k \times 1k$ pixels) of the actual scene is depicted here.




Fig. 2. Distribution of the slope angle for avalanche and nonavalanche pixels.

makes. Specifically, it is defined as the elevation angle of the line between the point of furthest avalanche runout and the point of highest release. For most avalanches, this angle ranges between 20 and 40 degrees [15]–[17].

While the angle of reach is defined only for an existing avalanche, we here introduce the *PAR* (denoted as $\tilde{\alpha}$), which is defined for a hypothetical avalanche located at any given point in the DEM. Ideally, this feature will range values between 20 and 40 degrees in terrain where avalanches can accumulate. Assuming that avalanches normally releases in steep terrain,




Fig. 3. Definition of the PAR $\tilde{\alpha}$, where $\theta(x)$ denotes the angle between the horizontal and the line drawn from a point in a release zone, denoted as x , to the point of interest.

e.g., in slopes of 30–50 degrees, the PAR angle is obtained by 1) computing the elevation angle to all neighboring release points x (within a 4 km radius), and 2) by taking the maximum of all such angles, as illustrated in Fig. 3. By computing the PAR for each point in the DEM, a PAR feature map can be obtained and used as an additional channel of input images.

Fig. 4 depicts the distribution of the PAR angles for avalanche and nonavalanche pixels using the training data. It is possible to see that for avalanche pixels the distribution is more regular and has a single peak centered around 40 degrees. While the true angle of reach is expected to range 20–40 degrees, the PAR is




Fig. 4. Distribution of the PAR for avalanche and nonavalanche pixels.

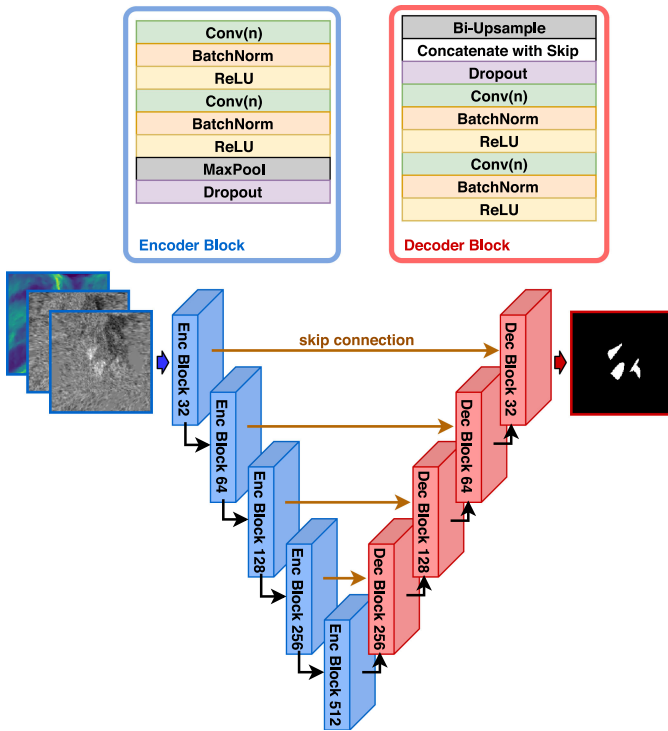


Fig. 5. FCN architecture used for segmentation. Conv(n) stands for a convolutional layer with n neurons. For example, $n = 32$ in the first encoder block, 64 in the second, and so on.

consequently biased toward higher values. We concluded that the PAR is informative since the two distributions are different for the two classes. Contrarily to the slope, the PAR is not simply concatenated to the other layers of the input image but is rather used to encourage the deep learning model to focus on specific areas (see Section IV-C).

IV. DEEP LEARNING MODEL

The FCN network used for segmentation is based on the U-Net architecture [18], which consists of an *encoder* and a *decoder*, respectively, depicted in blue and red in Fig. 5. The encoder hierarchically extracts feature maps that indicate the presence of the patterns of interest in the image. By reducing the spatial dimensions and increasing the number of filters, the deeper layers in the encoder capture patterns of increasing complexity and with a larger spatial extent in the input image. The decoder

gradually transforms the high-level features and, in the end, maps them into the output. The output is a binary segmentation mask, which has the same height/width of the input and indicates which are the pixels that belong to the avalanche class. The skip connections link the feature maps from the encoding to the decoding layers, such that some information can bypass the bottleneck located at the bottom of the “U” shape. In this way, the network still learns to generalize from the high-level latent representation but also recovers the spatial information through a pixelwise semantic alignment with the intermediate representations.

Fig. 5 shows the architecture details: the number n in each Enc/Dec Block indicates the quantity of 3×3 filters in the Conv(n) layers. The encoder reduces the spatial dimension with max pooling, whereas the decoder restores it through bilinear up-sampling. Each block contains two batch normalization [19] and one dropout layer [20], which are, respectively, used to facilitate the training convergence and improve the model generalization capability. We note that batch norm layers are not present in the original U-Net architectures but, as also verified in preliminary experiments, their presence improves the segmentation performance. The last encoder block (Enc Block 512 in Fig. 5) does not have dropout, whereas the last decoder block (Dec Block 32) is followed by a Conv layer with one 1×1 filter and a sigmoid activation. Since the network is fully convolutional (there are no dense layers), it can process inputs of variable size.

We note that it would be possible to use more powerful FCN architectures such as DeepLabV3+ [21], which achieves state-of-the-art results in segmenting natural images. However, models with a larger capacity, such as DeepLabV3+, require very large datasets to be trained on. In remote sensing applications, a smaller network such as U-Net is often preferred, given the limited amount of training data. Moreover, the U-Net outperforms other architectures in detecting small objects [22], such as the snow avalanches in our work.

A. Class Balance

Avalanches are small objects and the avalanche class is highly under-represented in the dataset (avalanche pixels are only 0.019% of the total). Therefore, a trivial model that classifies each pixel as “nonavalanche” would reach a classification accuracy of 99.98%. A solution to handle class unbalance is to differently weight the loss relative to the pixels of the different classes so that the model is more penalized when it misclassifies the underrepresented class [9]. Specifically, we configured the loss to give twice more importance to the classification errors on the avalanche pixels. We also experimented with loss functions specifically designed to handle class unbalance, such as the Jaccard-distance loss [23] and the Lovász-Softmax loss [24], but we obtained worse results than optimizing the FCN using binary cross-entropy loss with class balancing.

B. Data Augmentation

To avoid overfitting during training and to enhance the model generalization to new data, we perform data augmentation by

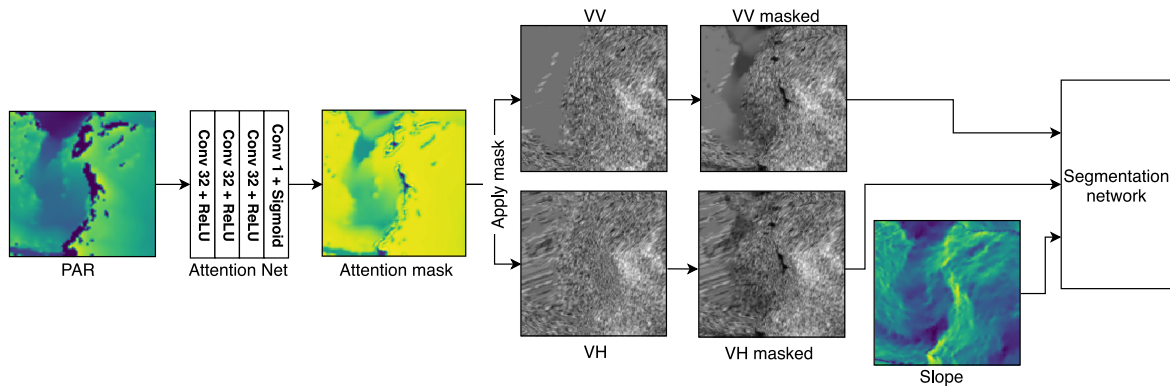


Fig. 6. For each patch, the Attention Net generates an attention mask from the PAR features and applies it to the VV, VH, and VVVH SAR features. The masked SAR features and the slope (not masked) are then fed into the U-Net. Attention Net and U-Net are jointly trained by minimizing the segmentation error. Note that the VVVH feature is not shown in the figure for conciseness.

randomly applying (on the fly) horizontal and vertical flips, horizontal and vertical shifts, rotations, zooming, and shearing to the training images. To ensure consistency, the same transformations on the input images are also applied to their labels (avalanche masks).

To compute the prediction of a whole SAR product at inference time, we could *slide* the FCN on the large image and compute predictions for one window at a time. However, this approach usually generates checkerboard artifacts and border effects close to the window edges. To obtain smoother and more accurate predictions, we consider overlapping windows by sliding the FCN with a stride equal to half the window size. Furthermore, we apply to each window all the possible 90° rotations and flips; then, we compute the predictions and, finally, revert the transformations on the predicted outputs. To obtain the final segmentation, we first merge the multiple predictions available at each pixel location (stemming from the geometric transformations and the overlapping windows) and then we join them by using a second-order spline interpolation.

C. Attention Mask

Following our hypothesis that the PAR feature map can highlight areas where it is more likely to find an avalanche, we propose a neural attention mechanism [25] that generates an *attention mask* conditioned on the PAR. The intention is to learn an attention mask that encourages the segmentation procedure to put more focus on specific regions of the input image. Specifically, we use a small network that takes as input the PAR and generates the attention mask that is, subsequently, applied pixelwise to the SAR channels (VV, VH, and VVVH) before they are fed into the segmentation network (see Fig. 6). We note that the attention mask is not applied to the input channel containing the slope feature map.

The attention network consists of three stacked Conv layers with $32 \times 3 \times 3$ filters and ReLU activations and a Conv layer with one 3×3 filter and sigmoid activation. The attention network has a small receptive field (seven pixels), meaning that each attention value only depends on the local PAR. This is acceptable since the PAR already yields highly nonlocalized features

from the DEM and captures long-range relationships in the scene.

The attention network is also fully convolutional and is jointly trained with the segmentation network. Our solution allows learning end-to-end on how to generate and apply the attention mask in a way that is optimal for the downstream segmentation task. This is a more flexible approach than masking out parts of the input (e.g., by applying precomputed runout masks), or directly premultiplying the SAR channels with the PAR feature map.

D. Model Training and Evaluation

We trained the FCN by feeding it with small square patches, rather than processing entire scenes at once, which would also be unfeasible due to the memory limitations of the GPU.¹ By using small patches it is also possible to inject stochasticity in the learning phase by randomly shuffling and augmenting the data at each epoch. This limits overfitting and decreases the chances of getting stuck in local minima. We experimented with patches of 160×160 or 256×256 pixels, which is a size compatible with the receptive field of the filters in the innermost network layer (Enc Block 512), which is 140. After preliminary experimentation, we obtained the best performance with the 160×160 patches. The training and validation sets are generated by randomly partitioning these patches in order to prevent biasing either the training or validation sets toward any particular imaging parameters, such as the incidence angle. It should, moreover, be noted that image pairs are only constructed from the same satellite orbit number, such that the viewing geometries of the activity and reference images are nearly identical. To build the training/validation set, we considered only the patches containing at least one pixel classified as “avalanche” by the human expert. We ended up with $\approx 35,000$ patches, of which 10% were used as a validation set for model selection and early stopping. Finally, out of the 118 available S1 scenes, one scene with date April 17, 2018, which contains 99 avalanches, was isolated from the rest and used as the test set.

¹Two Nvidia GTX2080 were used to train and evaluate the model.

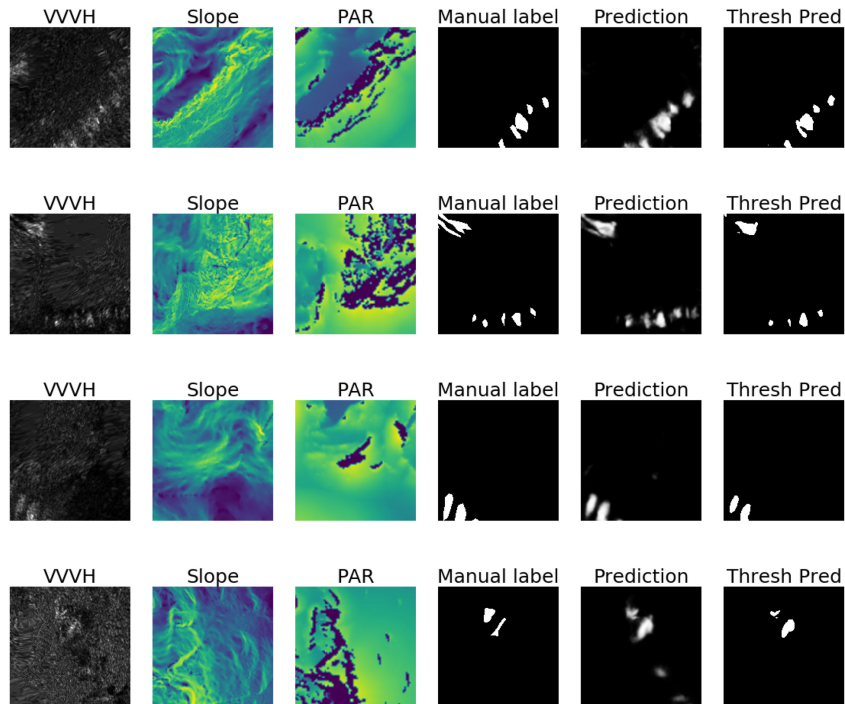


Fig. 7. Examples of prediction on individual patches of the validation set. From the left: (i) VVH input channel fed to FCN; (ii) slope feature fed to FCN; (iii) PAR feature fed to Attention Net; (iv) ground truth labels manually annotated by the expert; (v) raw output of the FCN; and (vi) FCN output thresholded at 0.5.

V. RESULTS AND DISCUSSION

The network is trained with Adam optimizer [26] with default parameters; we used minibatches of size 16 and dropout rate 0.4. Examples of FCN predictions are depicted in Fig. 7. Since the networks predict real values in $[0, 1]$, a binary segmentation mask (last column) is obtained by thresholding the soft output (third column) at 0.5.

Since the avalanche class is highly underrepresented, accuracy is not a good measure to quantify the performance and, therefore, we evaluated the quality of the segmentation result by using different metrics. The first is the $F1$ score, which is computed at the pixel level and is defined as

$$F1 = 2 \frac{\text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

where precision is defined as $\frac{TP}{TP+FP}$ and recall is $\frac{TP}{TP+FN}$ (TP = True Positives, FP = False Positives, FN = False Negatives). The $F1$ score is also evaluated during training on the validation set and used for early stopping and for saving the best model.

To evaluate the segmentation results at a coarser resolution level, we considered the bounding boxes containing the avalanches in the ground truth and in the predicted mask. To quantify how much the bounding boxes overlap in the ground truth and the predicted segmentation mask, we computed the intersection over union (IoU)

$$\text{IoU} = \frac{\text{Area of bounding boxes intersection}}{\text{Area of bounding boxes union}}.$$

We compared the proposed deep learning method with the state-of-the-art algorithm for automatic avalanche detection,

TABLE I
SEGMENTATION RESULTS FROM THE TEST IMAGE WITH 99 AVALANCHES. WE REPORT THE $F1$ SCORE (IN PERCENT), IOU OF THE BOUNDING BOXES (IN PERCENT), TRUE POSITIVE (CORRECT HITS), FALSE NEGATIVE (MISSED AVALANCHES DETECTION), AND FALSE POSITIVE (FALSE AVALANCHES DETECTION)

Method	$F1$ (%)	IoU (%)	TP (#)	FN (#)	FP (#)
Baseline	38.13	33.11	44	45	11
FCN	66.6	54.3	72	17	32

which is currently used in production pipelines [2]. Such a segmentation algorithm is primarily driven by change detection and filtering methods to enhance potential avalanche features; dynamic thresholding based on the statistics of image pairs controls the final delineated features. The baseline algorithm is, to a large extent, dependent on additional input layers such as slope, vegetation maps and runout zone information that restrict the areas where features are allowed to be detected, thereby reducing the number of false alarms as much as possible.

Table I reports the results obtained on the test image. Compared to the baseline, the FCN achieved a much higher agreement with the manual labels, as indicated by the higher $F1$ and IoU values. Out of the 99 avalanches in the test image, FCN correctly identified 72 of them and missed 17. However, most of the FN are small avalanches that are difficult to detect. FCN also identified 32 FP: most of them are due to particular terrain structures, which cause high backscatter that resembles avalanches (see Fig. 8). Interestingly, some of those FPs are




Fig. 8. Comparison between manual labeling and FCN output overlay onto an RGB change detection image. From the left: (i) agreement between FCN detection and manual annotations; (ii) avalanches missed by the FCN; (iii) false detection from the FCN algorithm; and (iv) avalanches correctly detected by the FCN but overlooked during the manual annotation.

TABLE II
ABLATION EXPERIMENT RESULTS

VV	VH	VVVH	Slope	PAR	PAR (attn.)	F1
✓						55.4
✓	✓					63.0
✓	✓					64.9
✓	✓	✓	✓			65.2
✓	✓	✓	✓	✓		65.4
✓	✓	✓	✓		✓	66.6

actual avalanches that have been overlooked during the manual annotation.

A. Ablation Study

The ablation study consists of removing some features from the model or from the input data to evaluate how these affect the performance. In particular, we study how much each SAR channel and the topographical feature maps contribute to the segmentation results. We also evaluate the difference in concatenating the PAR to the other input channels (VV, VH, VVVH, and slope) or using it to compute the attention mask that is applied pixelwise to the SAR channels (see the details in Section IV-C).

The results reported in Table II indicate that the most important improvement comes from including the difference image obtained by the VH channels, compared to using the VV channel alone. By adding the slope and PAR features it is possible to further increase the segmentation performance. Finally, the results show that the proposed attention mechanism allows to better exploit the information yield by the PAR, compared to just concatenating the PAR feature map to the other input channels.

VI. CONCLUSION

In this work, we proposed the first deep learning approach for saliency segmentation of avalanches in S1 SAR images. As channels of the images provided as input to the segmentation network, we used the time difference of the radar backscatter information, as well as topographical information. The latter

consists of the terrain slope and the newly introduced PAR, which indicates the likelihood of finding avalanches at different locations. The topographical feature maps were provided along with the SAR features to an FCN, which was trained to perform avalanche segmentation. The ground truth segmentation masks used to train the deep learning model came from the manual labeling of avalanche pixels performed by a human expert. A total of 118 S1 SAR products were labeled, of which 117 were used for training and one single product was used for testing the segmentation performance on unseen data.

The FCN was extended with an additional attention block, jointly trained with the rest of the segmentation network, which computes an attention mask conditioned on the PAR. The mask was applied to the input SAR features to let the segmentation network focusing more on the critical areas.

The results show the effectiveness of the proposed method, improving the $F1$ score of 38.1% achieved by a baseline signal processing algorithm to 66.6%. The $F1$ score was computed based on the manual labeling of the human expert. The proposed deep learning model only fails to detect some of the smaller avalanches, whereas detects additional avalanches that have been missed by the expert.

By being the first of its kind, we believe that our work will pave the way for pixel-level classification of snow avalanches in SAR data with deep learning and will serve as a future reference in the field of earth science and remote sensing. Our analysis and the obtained results suggest that the PAR is well correlated with the presence of avalanches. Therefore, we believe that the proposed PAR feature will be useful for future work in this field. In the next step, we aim to extend our dataset to evaluate the FCN's performance on SAR images with different snow conditions (wet or dry).

REFERENCES

- [1] M. Eckerstorfer, E. Malnes, and K. Müller, "A complete snow avalanche activity record from a Norwegian forecasting region using Sentinel-1 satellite-radar data," *Cold Regions Sci. Technol.*, vol. 144, pp. 39–51, 2017.
- [2] H. Vickers, M. Eckerstorfer, E. Malnes, and A. Douleris, "Synthetic aperture radar (SAR) monitoring of avalanche activity: An automated detection scheme," in *Proc. Scand. Conf. Image Anal.*, 2017, pp. 136–146.

- [3] X. X. Zhu *et al.*, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geosci. Remote Sens. Mag.*, vol. 5, no. 4, pp. 8–36, Dec. 2017.
- [4] Y. Zhou, H. Wang, F. Xu, and Y.-Q. Jin, "Polarimetric SAR image classification using deep convolutional neural networks," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 12, pp. 1935–1939, Dec. 2016.
- [5] M. Kampffmeyer, A.-B. Salberg, and R. Jenssen, "Urban land cover classification with missing data modalities using deep convolutional neural networks," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 11, no. 6, pp. 1758–1768, Jun. 2018.
- [6] O. A. Penatti, K. Nogueira, and J. A. Dos Santos, "Do deep features generalize from everyday objects to remote sensing and aerial scenes domains?," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. Workshops*, 2015, pp. 44–51.
- [7] L. T. Luppino, F. M. Bianchi, G. Moser, and S. N. Anfinsen, "Unsupervised image regression for heterogeneous change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 57, no. 12, pp. 9960–9975, Dec. 2019.
- [8] L. T. Luppino, F. M. Bianchi, G. Moser, and S. N. Anfinsen, "Remote sensing image regression for heterogeneous change detection," in *Proc. IEEE 28th Int. Workshop Mach. Learn. Signal Process.*, Sep. 2018, pp. 1–6.
- [9] M. Kampffmeyer, A.-B. Salberg, and R. Jenssen, "Semantic segmentation of small objects and modeling of uncertainty in urban remote sensing images using deep convolutional neural networks," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit. Workshops*, 2016, pp. 680–688.
- [10] F. M. Bianchi, M. M. Espeseth, and N. Borch, "Large-scale detection and categorization of oil spills from SAR images with deep learning," *Remote Sens.*, vol. 12, no. 14, 2020, Art. no. 2260.
- [11] P. E. Kummervold, E. Malnes, M. Eckerstorfer, I. Arntzen, and F. M. Bianchi, "Avalanche detection in Sentinel-1 radar images using convolutional neural networks," in *Proc. Int. Snow Sci. Workshop*, 2018, pp. 377–381.
- [12] S. Sinha *et al.*, "Can avalanche deposits be effectively detected by deep learning on Sentinel-1 satellite SAR images?," in *Proc. Climate Informat.*, Paris, France, Oct. 2019.
- [13] A. U. Waldeland, J. H. Reksten, and A.-B. Salberg, "Avalanche detection in SAR images using deep learning," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, 2018, pp. 2386–2389.
- [14] T. Grydeland and Y. Larsen, "Beyond plane sailing: Solving the range-Doppler equations in a reduced geometry," in *Proc. 12th Eur. Con. Synthetic Aperture Radar*, 2018, pp. 1–4.
- [15] S. Bakkehoi, U. Domaas, and K. Lied, "Calculation of snow avalanche runout distance," *Ann. Glaciol.*, vol. 4, pp. 24–29, 1983.
- [16] D. Delparte, B. Jamieson, and N. Waters, "Statistical runout modeling of snow avalanches using GIS in Glacier National Park, Canada," *Cold Regions Sci. Technol.*, vol. 54, no. 3, pp. 183–192, 2008.
- [17] K. Johnston, B. Jamieson, and A. Jones, "Estimating extreme snow avalanche runout for the Columbia Mountains, British Columbia, Canada," in *Proc. 5th Can. Conf. Geotechn. Nat. Hazards*, Kelowna, BC, Canada, vol. 49, no. 11, 2011, pp. 1309–1318.
- [18] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assisted Intervention*, 2015, pp. 234–241.
- [19] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. 32nd Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [20] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [21] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vision*, 2018, pp. 801–818.
- [22] M. Krestenitis, G. Orfanidis, K. Ioannidis, K. Avgerinakis, S. Vrochidis, and I. Kompatsiaris, "Oil spill identification from satellite images using deep neural networks," *Remote Sens.*, vol. 11, no. 15, 2019, Art. no. 1762. [Online]. Available: <https://www.mdpi.com/2072-4292/11/15/1762>
- [23] G. Csurka, D. Larlus, F. Perronin, and F. Meylan, "What is a good evaluation measure for semantic segmentation?," in *Proc. Brit. Mach. Vision Conf.*, 2013, vol. 27, p. 2013.
- [24] M. Berman, A. Rannen Triki, and M. B. Blaschko, "The Lovász-Softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks," in *Proc. IEEE Conf. Comput. Vision Pattern Recognit.*, 2018, pp. 4413–4421.
- [25] K. Xu *et al.*, "Show, attend and tell: Neural image caption generation with visual attention," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2048–2057.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Represent.*, 2015.



Scientist with the NORCE Norwegian Research Centre, Bergen, Norway. His research interests include machine learning, complex networks, and dynamical systems. He applied his work in several areas, including remote sensing.



Jakob Grahn received the M.Sc. degree in radio and space science from Chalmers University of Technology, Gothenburg, Sweden, in 2012, and the Ph.D. degree in physics from the University of Tromsø, Tromsø, Norway, in 2018.

Since 2018, he has been a Researcher with the NORCE Norwegian Research Centre, Bergen, Norway (previously NORUT). His current research focuses on radar remote sensing applications, such as mapping of snow properties, detection of extreme weather at sea, and 3-D deformation mapping by means of InSAR.



Markus Eckerstorfer received the Ph.D. degree in physical geography from the University Centre in Svalbard, Longyearbyen, Norway, in 2013.

Since 2013, he has been with the NORCE Norwegian Research Centre AS (former Norut), Bergen, Norway, where his focus lies on snow avalanche detection using radar satellite data.



Eirik Malnes received the M.Sc. and Dr.Sci. degrees in space plasma physics from the University of Tromsø, Tromsø, Norway, in 1991 and 1994, respectively.

From 1995 to 1997, he held a postdoctoral position with Tromsø and Grenoble, France. From 1997 to 2001, he was a Senior Scientist with The Norwegian Defense Research Establishment, Oslo, Norway, working on radar applications. From 2001 to 2019, he was a Senior Scientist with the Northern Research Institute (Norut), Tromsø. Since 2019, he has been a Senior Scientist with NORCE Norwegian Research Centre, Oslo. His field of interest is various applications of SAR, including physical modeling of scattering mechanisms and signal processing in connection with earth observation of the cryosphere surface water. He has led two EU projects and several ESA and national funded projects related to cryosphere monitoring. He is the author of more than 120 scientific papers and technical reports in the field of remote sensing.



Hannah Vickers received the Ph.D. degree in ionospheric physics from the Radio and Space Plasma Physics Group, University of Leicester, Leicester, U.K., in 2011.

She is currently a Research Scientist with NORCE Norwegian Research Centre AS, Bergen, Norway, where her interests lie in remote sensing of the cryosphere and its applications in climate change research.