

# Finding Ground-Based Radars in SAR Images: Localizing Radio Frequency Interference Using Unsupervised Deep Learning

Kristian Aalling Sørensen<sup>1</sup>, Graduate Student Member, IEEE, Anders Kusk<sup>1</sup>,  
Peder Heiselberg<sup>1</sup>, and Henning Heiselberg<sup>1</sup>

**Abstract**—Synthetic aperture radar (SAR) satellite images are used increasingly more for Earth observation. While SAR images are useable in most conditions, they occasionally experience image degradation due to interfering signals from external radars, called radio frequency interference (RFI). RFI-affected images are often discarded in further analysis or preprocessed to remove the RFI. However, few on-ground radars can cause RFI in SAR images and such information can thus increase domain awareness greatly over both land and sea, where, e.g., localizing and characterizing RFI signals in the ocean could help classify otherwise overlooked ships. The aim of the current study is to detect and localize RFI signals automatically in Sentinel-1 level-1 images and further characterize the on-ground radar. The spatial structure of RFI signals vary greatly. A convolutional autoencoder (CAE) was therefore developed to reconstruct RFI-free Sentinel-1 images. Conversely, RFI-affected images could not be well reconstructed. Anomalous heatmaps were then developed to automatically detect and localize RFI anomalies in the images under varying environmental and geographical conditions, whereafter the external radar characteristics were extracted manually from Sentinel-1 level-0 data. We could consequently classify and localize RFI signals believed to originate from both stationary radars and ship-borne radars. We further argue that the calculated ship-borne radar characteristics correspond to those of air-surveillance radars. Empirically, the method showed better detection results than those of previous studies. Our study shows that more information can be extracted from certain detected objects, such as ships, from SAR images.

**Index Terms**—Anomaly classification and localization, convolutional autoencoder (CAE), deep learning, radio frequency interference (RFI), synthetic aperture radar (SAR).

## I. INTRODUCTION

**N**EW surveillance methods applicable for maritime domain awareness are essential for monitoring and controlling maritime traffic safety, piracy, smuggling, invasion of

foreign navy, and more. Dark ships are noncooperative vessels with nonfunctioning transponder systems such as the automatic identification system (AIS) for ships. Their transmission may be jammed, spoofed, sometimes experience erroneous returns, or simply turned off deliberately or by accident.

Foreign navy, in particular, rarely broadcast, e.g., identification or location information. Other noncooperative surveillance systems are therefore required, such as satellite or airborne imaging systems; electronic warfare support measures; or coastal radars. Furthermore, knowing that a detected vessel carries, e.g., an air-surveillance radar provides valuable information to decision makers and gives important insight to the vessel characteristics.

Satellite-borne synthetic aperture radar (SAR) imaging sensors have all-weather, all-hours imaging capabilities, and can therefore be used to image, e.g., noncooperative ships [1]. However, emitted pulses from external active radars with similar center frequency as the SAR can cause interference in the SAR images. This is called radio frequency interference (RFI) [2], [3] and is caused by, e.g., wireless communication such as radio and television in the P- or L-band spectrum [4], or by air-surveillance radars in the C-band spectrum. Furthermore, mutual RFI can occur between two SAR satellites [5], [6], [7]. However, the location of the mutual RFI between two satellites can be predetermined by the orbits of the satellites [6]. Alternatively, by analyzing the RFI, it can be determined to originate from mutual RFI, due to, e.g., the periodical structure of the mutual RFI in the entire image caused by the superimposing of the two images [5]. Shao et al. [8] analyzed the impact of RFI signals in SAR ship detection, and concluded that RFI signals pose a serious threat in automatic ship detection algorithms. More importantly, few on-ground radars are capable of causing RFI signals in C-band SAR images mid-sea. Localizing the RFI signals in an SAR image would therefore enhance maritime domain awareness by: increasing the amount of detected noncooperative ships, since detection schemes developed for several polarizations might fail when an image is contaminated by an RFI [8]; providing valuable information on the detected vessel, since only few types of vessels carry air-surveillance radars, with some warships being among them.

The few published studies of RFI signals in SAR images focus on detecting images containing RFI signals [9], [10],

Manuscript received 27 January 2023; revised 12 May 2023 and 9 June 2023; accepted 28 June 2023. Date of publication 13 July 2023; date of current version 11 August 2023. (Corresponding author: Kristian Aalling Sørensen.)

Kristian Aalling Sørensen and Henning Heiselberg are with the Center for Security, National Space Institute of Denmark, Technical University of Denmark, 2800 Kongens Lyngby, Denmark (e-mail: kaaso@space.dtu.dk; hh@dtu.dk).

Anders Kusk is with the Department of Microwaves and Remote Sensing, National Space Institute of Denmark, Technical University of Denmark, 2800 Kongens Lyngby, Denmark.

Peder Heiselberg is with the Department of Geodesy and Earth Observation, National Space Institute of Denmark, Technical University of Denmark, 2800 Kongens Lyngby, Denmark.

Digital Object Identifier 10.1109/TGRS.2023.3295002

or to remove the RFI signal using different mitigation techniques [11], [12], [13], [14], [15]. In particular, Monti-Guarnieri et al. [9] implemented a global Sentinel-1 RFI-contaminated probability map. Instead of using the processed SAR image or full level-0 data, Monti-Guarnieri et al. [9] exploited the noise measurement of the rank echoes in the level-0 Sentinel-1 data. Monti-Guarnieri et al. [9] implemented a Fisher's Z test and the Kullback–Leibler divergence to find RFIs using the statistical measurement of the noise in the rank echoes. Since March 2022, their RFI detection and mitigation schema has been applied to all Sentinel-1 image to remove RFI [16]. However, their method assumes that all external contributions will be present in the rank echoes. Meyer et al. [17] used a conventional incoherent notch-filtering algorithm for the later RFI mitigation applied to L-band SAR data acquired from the PALSAR instrument. They furthermore extracted RFI information such as the RFI type and RFI range bandwidth. In [18], they used the Sentinel-1 quick-look images to find and flag images containing RFI. They transformed all three RGB values of the Sentinel-1 quick-looks to a single gray-scale image. In their preprocessing procedure, they had to empirically find a threshold for each image that removed low-intensity pixels, whereafter they used a convolutional neural network as a reference classifier. Likewise, Artiemjew et al. [10] implemented a convolutional neural network as a supervised binary classifier. Artiemjew et al. [10] and Chojka et al. [18] find images with RFI signals without localizing the signals and are dependent on manually decided thresholds specific for each image. Most importantly, most RFI deep learning algorithms in the literature are supervised, and are thus heavily dependent on the manually labeled training data or simulated data [10], [12], [14], [18], [19]. RFI signals have very different spatial structures and supervised models cannot capture the variability of RFI signals. It is therefore a necessity to use either out-of-distribution deep learning, or unsupervised anomaly detectors.

Lately, research has been conducted on detecting and localizing RFI signals from stationary RFI sources. In [20], they detected and localized RFI signals using location-specific threshold and tolerance values applied in a combination of the Sentinel-1 VV and VH polarizations. While they did find an approximate location using both the ascending and descending orbits, they also assumed that an RFI was present in each analyzed image, and they did not determine any characteristics of the radar. In [21], they detect RFI signals in Sentinel-1 quick-looks using an advanced deep learning model trained on labeled, segmented RFI anomalies. Similar to, e.g., autoencoders, their network compressed the images into a bottleneck that they then up-sampled. They furthermore added atrous convolutions, separable convolutions, and attention layers etc., and showed very good results compared to other models, such as U-Net and MobileNet-V2. They labeled 100 training and 100 validation images and trained their model using the ground truth images. They showed that their model could detect RFI signals from both on-ground radars and from mutual RFI. While their model and results are very impressive, they used positively labeled images for training, meaning it is necessary to label RFI images for each geographical location.

We argue that RFI signals provide valuable information. Likewise, we argue that a mitigation strategy performed on Sentinel-1 level-0 data reduces the information, since RFI affected values oftentimes are replaced by, e.g., the nearest unaffected value.

In this study, we present an unsupervised RFI localization method using deep learning. By training a convolutional autoencoder (CAE) on RFI-free Sentinel-1 interferometric wide-swath (IW) quick-look images, anomalies can be detected in SAR images. The anomalies are then localized and classified as RFI anomalies using domain knowledge on the structures of RFI signals. By analyzing the Sentinel-1 level-0 data corresponding to the location of the RFI signals, we extract information about the ground-based radars.

The novelty of this study lies in a workflow that can be used to find and extract on-ground radar characteristics using data from the Sentinel-1 SAR constellation without explicitly labeling RFI signals in the images. The workflow can be summarized in three steps.

- 1) Detect large-scale anomalies in Sentinel-1 IW quick-look images using deep learning.
- 2) Classify the detected anomalies as RFI signals using domain knowledge of the processed level-1 images.
- 3) Analyze the corresponding level-0 Sentinel-1 image burst and characterize the ground-based radar.

Section II describes the localization methodology and the extraction of RFI information. Section III shows our results followed by a discussion in Section IV.

## II. METHOD AND MATERIALS

Anomalies were detected in Sentinel-1 quick-look images with an unsupervised CAE, whereafter the anomalies were localized and classified as either RFI signals, or other anomalies using knowledge on the characteristics of the anomaly types. The Sentinel-1 level-0 data for the respective RFI signal were then used to extract information of the on-ground radars. The methodology therefore initially describes the RFI formation in Sentinel-1 data. Autoencoders will further be described, whereafter the dataset and reconstruction model is introduced.

### A. Sentinel-1

An RFI signal is an external contribution to the received signal of the SAR, caused by an external radar with a similar center frequency, emitting a signal toward the receiving SAR antenna. The total received signal,  $P_r$ , under the influence of an external RFI signal,  $P_I$ , and other noise contributions,  $P_N$ , is then given by

$$P_r = P_{r_s} + P_N + P_I \quad (1)$$

where  $P_{r_s}$  is the contribution of the received signal derived from the transmitted SAR signal. The external RFI contribution,  $P_I$ , travels a distance dependent on the external source; if the source is ground-based, it travels half the distance of  $P_{r_s}$  and can therefore oversaturate  $P_r$ , resulting in very high-intensity pixels in the SAR data.

The Sentinel-1 SAR instrument transmits microwave pulses and receives the echo from each pulse. In the IW mode, the data are acquired in bursts of approximately 1500 echoes, corresponding to  $\approx 1$  s satellite flight time, during which the elevation beam pattern is kept constant and the azimuth beam pattern is swept from looking slightly aft to looking slightly ahead, thus imaging an azimuth extent longer than the distance traveled by the satellite during the burst. This allows the elevation beam pattern to be cycled from burst to burst between three different swaths (IW1–IW3), increasing the ground coverage at the expense of azimuth resolution and a more complicated processing. The level-0 echoes received by the radar in each burst are downlinked from the satellite, and are provided by the European space agency (ESA) as the level-0 product. This product is not relevant for most users, as the resolution in both range and azimuth is in the order of kilometers and an SAR focusing must be carried out to retrieve useful images.

1) *Sentinel-1 IW Processing*: For the IW mode, the pulses transmitted by Sentinel-1 are linear frequency modulated with a length,  $T_{\text{SAR}}$ , of approximately  $50 \mu\text{s}$  and a bandwidth,  $BW$ , from 42–56 MHz, depending on the swath.

The first step of focusing is the pulse/range compression, which is carried out individually on each echo, and is a cross correlation of the received pulse with a replica of the transmitted pulse [22], [23]. This has the effect of “compressing” the received pulse echoes to a width of  $\approx 1/BW$ , resulting in a range resolution ( $c/2BW$ )  $\approx 3$  m, with  $c$  being the speed of light. RFI signals from an external radar using a different pulse modulation will not be compressed and may even be extended beyond their original length dependent on the RFI pulse modulation, and will therefore often be as long as the swath in the range direction.

SAR focusing in the azimuth direction works in a similar way by exploiting the approximately linearly changing Doppler frequency arising from the relative motion between the satellite and a stationary target on the ground. Focusing is achieved by cross-correlating the range compressed signal matrix in the azimuth direction with an linear frequency-modulated correlation kernel. In the azimuth direction, an RFI signal from a single echo is a delta function, and will thus be smeared over all focused image lines to which the RFI-affected pulse contributes. For IW mode images, the smearing is changed by the beam sweeping and the net effect is that a single RFI signal should, in IW level 1 images, extend over 3–5 km in the azimuth direction.

### B. Convolutional Autoencoders

In deep learning, algorithms are used to automatically find nonlinear functional relationships between an input and an output using a model trained with dataset. In an autoencoder, a model is trained to reconstruct the original input using an encoder and decoder, respectively, by parameterizing the input,  $\mathbf{x}$ , into a latent space,  $\mathbf{z}$ . For a good  $\mathbf{z}$ , the autoencoder learns how to reconstruct a noise-free sample,  $\hat{\mathbf{x}}$ , and can therefore be used to automatically and unsupervised find anomalies by comparing a single input  $\mathbf{x}$  with the reconstructed  $\hat{\mathbf{x}}$ , i.e., a sample is anomalous when  $\mathbf{x} \not\approx \hat{\mathbf{x}}$ . An autoencoder finds

TABLE I  
COMPARISON OF SENTINEL-1 IMAGE PRODUCTS. PRODUCTS COMPARED:  
SLC IW, GRD IW, AND QUICK-LOOK

Parameter	SLC IW	GRD IW	Quick-look
Range resolution [m]	$\approx 3.1$	$\approx 20.4$	$\approx 1020$
Azimuth resolution [m]	$\approx 21.7$	$\approx 21.7$	$\approx 1062$
Bits per pixel	16 I/Q	16	8
Pixel spacing [m]	NA	10x10	500x490
Radiometric Resolution	3	1.6	NA
File size	7.31 GB	1.75 GB	351 kB

all anomalies in the data, but does not classify the type of anomaly as, e.g., an RFI signal.

An encoder maps an input,  $\mathbf{x}$  into a smaller latent space,  $\mathbf{z}$ , as follows:

$$g_\phi : \mathbf{x} \in \mathbb{R}^{d(\mathbf{x})} \rightarrow \mathbf{z} \in \mathbb{R}^{d(\mathbf{z})} \quad (2)$$

where the latent space dimension is smaller than the input dimension, i.e.,  $d(\mathbf{z}) \ll d(\mathbf{x})$ . A decoder then makes a reconstruction,  $\hat{\mathbf{x}}$ , of the original input

$$q_\theta : \mathbf{z} \in \mathbb{R}^{d(\mathbf{z})} \rightarrow \hat{\mathbf{x}} \in \mathbb{R}^{d(\mathbf{x})}. \quad (3)$$

The objective of a CAE is to reconstruct an input array using convolutional operations. Increasingly more complex patterns can be found by placing several layers consecutively. Hence, the encoder in a CAE uses convolutional layers to parameterize  $\mathbf{x}$  using the latent space,  $\mathbf{z}$ . Conversely, the decoder uses inverted convolutions, called deconvolutions, and aims to reconstruct the original input using  $\mathbf{z}$ , i.e.,  $\hat{\mathbf{x}} \approx \mathbf{x}$ . For  $\mathbb{R}^{d(\mathbf{z})} \approx \mathbb{R}^{d(\mathbf{x})}$ , a CAE likely overfit the training data and make a perfect reconstruction of the training images. Conversely, a too small  $\mathbb{R}^{d(\mathbf{z})}$  does not allow for the decoder to reconstruct an image.

### C. Dataset

SAR images from the regions shown in Fig. 1 were obtained from the Sentinel-1 SAR satellites. The first two bands from the quick-look images, corresponding to the VV and VH polarizations, were acquired from the Sentinel-1 IW ground range detected (GRD) VV/VH images.

Since it was predominately the large-scale spatial features of the RFI signals that were used to characterize the RFI anomalies, the quick-look resolution was deemed sufficient for both RFI detection and localization. From Table I, we see a comparison between the different Sentinel-1 image products. The spatial resolution for the quick-looks are  $50 \times 50$  times lower than, e.g., the GRD images. Consequently, the co- and cross polarization file sizes are more than  $2000\times$  smaller, and the quick-looks reduce the computational cost correspondingly.

The dataset for training and validation was constructed such that it contain images without RFI signals. Conversely, a testing dataset was constructed such that it contains images both with and without RFI signals, but was not further annotated. A third dataset was manually labeled with both 45 RFI-free and 40 RFI-affected SAR images. This dataset was not used for neither training, nor testing of the CAE, but instead for determining global thresholds. Table II illustrates the regional data distribution of the different datasets.






Fig. 1. (Left) Geographical locations of the four regions where the Sentinel-1 quick-look images were acquired from. (Right) Zoomed-in view on the region of Dubai, UAE, showing nine different image acquisition areas.

TABLE II

DISTRIBUTION OF TRAINING QUICK-LOOKS AND TESTING QUICK-LOOKS IN THE DIFFERENT REGIONS. TRAINING DATA WERE RFI-FREE, WHEREAS TESTING DATA WERE UNANNOTATED

Region	Training	Testing	Annotated RFI/ No RFI	Combined
Denmark	86	26	10/12	134
Israel	96	422	13/12	543
UAE	171	247	12/12	442
Japan	272	72	5/9	358
Total	625	767	85	1477

As illustrated in Fig. 1, data from four different regions were acquired. The images, and corresponding metadata, were retrieved from the Alaskan Satellite Facility Distributed Active Archive Center and processed by ESA. Sentinel-1 SAR images are often affected by RFI from sources in the regions near the United Arab Emirates (UAE), Israel, and Denmark. Similarly, Sentinel-1 data are occasionally affected by targets mid-sea near Japan. The areas thus illustrated the capability of the method to detect and localize RFI signals in different regions.

All images were subsequently normalized with the min-max normalization using the values from the training data, i.e., all training images were within the range of  $\{0, 1\}$  and were center-cropped to the same size of  $340 \times 500 \times 2$  with the last two channels corresponding to the VV and VH polarization, respectively. Fig. 2 illustrates a subset of the testing data with the VV polarization (top) and VH polarization (bottom). Both images had RFI signals in their VH polarizations, while the image from UAE also had RFI in its VV polarization. We see that the images in the range extent of the RFI signals is the length of a swath with the RFI near UAE covering many swaths, and that the azimuth extent is approximately 3–5 km, as explained in Section II-A.

#### D. Image Reconstruction and RFI Localization

A CAE was made to reconstruct all input SAR images. The model was trained on a dataset consisting of RFI-free SAR images and the CAE consequently reconstructed normal SAR images. Conversely, RFI-affected SAR images were poorly reconstructed and the differences between the original and reconstructed images were therefore used to find anomalies. The CAE found all types of anomalies, not only RFI signals, and a secondary classification schema was therefore made

to determine if the anomalies were RFI signals. In short as follows.

- 1) A single CAE model reconstructed SAR images for all regions, excluding the RFI signals and other large-scale anomalies.
- 2) Similarity heatmaps,  $\mathbf{H}$ , were made between  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  to find anomalous regions.
- 3) The reconstruction errors, i.e., the anomalies,  $\mathbf{A}$ , were calculated only in the anomalous regions in  $\mathbf{H}$ .
- 4) If the anomaly for the individual Sentinel-1 RFI region in  $\mathbf{A}$  was large enough, it was defined as an RFI signal.

1) *CAE Image Reconstruction*: The CAE model is described in Fig. 3 and was implemented using Tensorflow 2.10. Each layer was followed by a batch normalization with a momentum of 0.99 and  $\epsilon = 0.001$  [24] and the first two layers of the encoder were likewise followed by a  $2 \times 2$  kernel max-pooling layer. The hyperparameters (kernel size, number of filters, the latent space size, if dropout was used, L1 and L2 regularization, gradient clipping value, initial learning rate) were found with a discrete hyperband parameter search [25] using the KerasTuner [26]. For each hyperparameter, several different discrete steps were allowed, with a total of 147 different trails in the first iteration and a termination after 150 epochs. For instance, the latent space was allowed a minimum value of 25, a maximum of 350, and a step size of 25. Similarly, the initial learning rate was logarithmic sampled between  $10^{-2}$  and  $10^{-6}$ . The model was trained anew with the found hyperparameters using a fivefold cross-validation with a reducing learning rate and the Adam optimizer. The weights corresponding to the lowest validation loss was saved and used.

The cost function for the reconstruction error was made with the mean structural similarity measure (MSSIM) as the loss function,  $\mathcal{L}$ . The MSSIM is maximizing the perceptual similarity between  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  applied on each polarization and averaged [27]. The MSSIM use local mean, variance, and covariance calculated on sliding boxes and compares the differences between the two images. The MSSIM for a single reconstructed image is given by the mean of the similarity between all local regions,  $r = 1, \dots, R$

$$\text{MSSIM}(\mathbf{x}, \hat{\mathbf{x}}) = \frac{1}{R} \sum_{r=1}^R \left( \frac{(2\mu_{x_r}\mu_{\hat{x}_r} + C_1)(2\sigma_{x_r\hat{x}_r} + C_2)}{(\mu_{x_r}^2 + \mu_{\hat{x}_r}^2 + C_1)(\sigma_{x_r}^2 + \sigma_{\hat{x}_r}^2 + C_2)} \right). \quad (4)$$




Fig. 2. Subset of the test data with (top) VV polarization and (bottom) VH polarization from (left) UAE, United Arab Emirates with strong RFI stripes, and (right) Island of Shikoku, Japan with a single RFI stripe in its VH polarization in the upper right burst.

The mean and variance of a local region  $x_r$  in  $\mathbf{x}$  are given by

$$\begin{aligned} \mu_{x_r} &= \sum_{i=1}^{N=11^2} w_i x_{r,i} \\ \sigma_{x_r}^2 &= \sum_{i=1}^{N=11^2} w_i (x_{r,i} - \mu_{x_r})^2, \text{ and likewise for } \hat{\mathbf{x}} \end{aligned} \quad (5)$$

and the local covariance

$$\sigma_{x_r, \hat{x}_r} = \sum_{i=1}^{N=11^2} w_i (x_{r,i} - \mu_{x_r})(\hat{x}_{r,i} - \mu_{\hat{x}_r})$$

they are calculated using a sliding  $11 \times 11$  weighted Gaussian matrix,  $w$ , to reduce blocking patterns in the similarity map [28]. Wang et al. [28] found that  $11 \times 11$  gave good results visually while experimenting with smaller and larger sizes. Each local region is then defined by a  $11 \times 11$  moving kernel, whereafter the similarity measure is averaged over the total number of local statistics calculated.  $C_1$  and  $C_2$  are small numbers added to avoid numerical instability as in [27].

The cost function is thus given by

$$\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) = \sum_{q=1}^Q 1 - \text{MSSIM}(\mathbf{x}, \hat{\mathbf{x}}) \quad (6)$$

where  $Q$  is the number of SAR images used to update the weights of the model, formerly called the batch size. Consequently,  $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) \rightarrow 0$  for  $\hat{\mathbf{x}} \approx \mathbf{x}$  and  $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) \rightarrow 1$  for  $\hat{\mathbf{x}} \not\approx \mathbf{x}$ . The MSSIM metric is often used in perceptual image processing [29] and outperforms, e.g., the mean-squared-error loss function for image comparisons by minimizing the luminance, contrast, and structure of two images [30]. We likewise experimented with, e.g., the multiscale SSIM (MS-SSIM) [27] and root-mean-square-error. MS-SSIM showed comparable results, and both MSSIM and MS-SSIM are easily implemented through the Tensorflow framework. MS-SSIM is more complex and the MSSIM was therefore chosen [31].

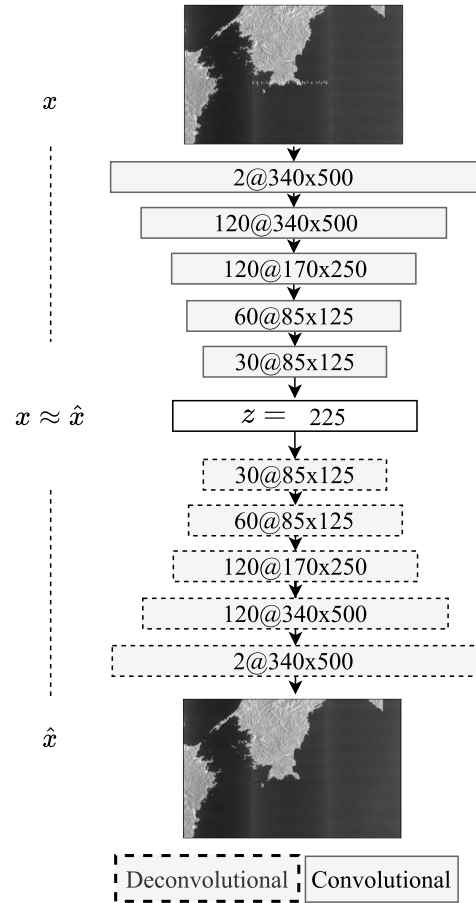


Fig. 3. Diagram of the CAE reconstruction model.

A batch size of 1 was too slow to train, and a batch gradient descent used too much memory. For very large batch sizes, we saw that the model had difficulties reconstructing parts of the input image and empirically, a batch size of 6 showed good results and  $Q = 6$  was therefore used.

2) *Find Anomalous Regions*: The CAE image reconstruction model was implemented in Tensorflow with

Hyperparameters determined using a hyperband optimization scheme, and the model could consequently reconstruct normal scenes. However, anomalous patterns in the input image,  $\mathbf{x}$  could not be well reconstructed.

Only large-scale anomalies were found where  $\hat{\mathbf{x}} \not\approx \mathbf{x}$ . Due to the spatial resolution of the SAR quick-look images, we expected only anomalies such as, e.g., RFI signals, varying ocean states, or poorly reconstructed regions. Poorly reconstructed regions were reduced by having a sufficiently large training set.

Reconstruction error heatmaps,  $\mathbf{H}$ , were made to determine if regions were anomalous using a MSSIM running kernel of size  $10 \times 30$ , thus showing regions with large reconstruction errors

$$\mathbf{H}^{j,k} = \sum_{j'=j}^{j+10} \sum_{k'=k}^{k+30} \text{MSSIM}(\mathbf{x}^{j',k'}, \hat{\mathbf{x}}^{j',k'}) \quad (7)$$

with  $j$  and  $k$  being the row and column pixel indices, respectively. The MSSIM kernel size should be large enough to capture some spatial structure, but a too large kernel would increase the computation time correspondingly, i.e., for each pixel pair, two  $J \times K$  subsets are compared with an arithmetic time complexity of  $O(JKN)$  pr. band with  $J$  and  $K$  being the rows and columns of the MSSIM running kernel and  $N$  the size of the weighted Gaussian matrix. It thus increases quadratically for the size. RFI signals were empirically shown to be more visible in elongated running kernels and the size of  $10 \times 30$  was selected empirically. A symmetrical padding was thereafter used on  $\mathbf{H}$  to get the original size of the images. The heatmaps,  $\mathbf{H}$ , were only used to determine anomalous regions, and not to extract the anomalies.

3) *Extract Anomalies in Anomalous Regions:* Large pixel-wise anomalies,  $\mathbf{A}^{j,k}$  were then extracted in regions with large anomalous errors where  $\mathbf{H}^{j,k} > 0.15$  thereby removing small, insignificant local reconstruction errors, such as the blocking structure explained in [28], i.e.,

$$\mathbf{A}^{j,k} = \begin{cases} |\mathbf{x}^{j,k} - \hat{\mathbf{x}}^{j,k}|, & \text{if } \mathbf{H}^{j,k} > 0.15 \\ 0, & \text{otherwise.} \end{cases} \quad (8)$$

The conservative, global noise threshold was made by comparing the values in  $\mathbf{H}$  from the annotated RFI-affected and RFI-free images (see Fig. 4). The RFI-free images have low  $\mathbf{H}$  values, whereas the RFI-affected images have much higher  $\mathbf{H}$  values. By setting the threshold at 0.15, RFI signals were preserved while most regions with background noise were removed, some background noise not originating from RFI signals was still expected.

4) *Classify Localized Anomalies as RFI Signals:* Elongated anomalies were grouped in  $\mathbf{A}$  by calculating the mean of errors in image subsets divided column-wise into approximately three swaths, 500 columns /3(swaths/image)  $\approx$  166 columns, similar to the swaths of the Sentinel-1 satellites. The quick-look images have a resolution of 490 m/row, and the theoretical height of the RFI signals are 5 km (see Section II-A), and the anomalies were therefore found row-wise according to the approximate height of an RFI signal,  $\lceil 5 \text{ km}/0.49(\text{km}/\text{row}) \rceil = 11$  rows, which in turn can be used to determine approximately

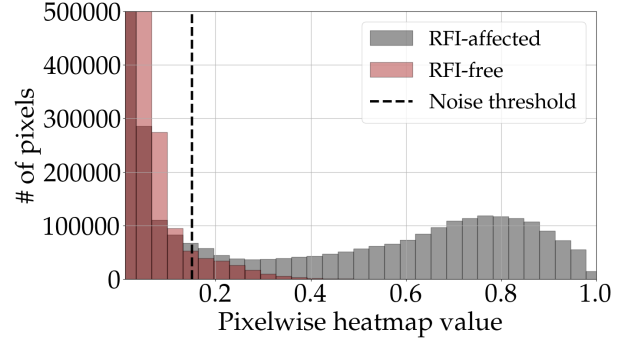


Fig. 4. Pixel-wise values in  $\mathbf{H}$  of (red) RFI-free images and (gray) RFI-affected images. The dotted black line illustrates the global noise threshold of 0.15.

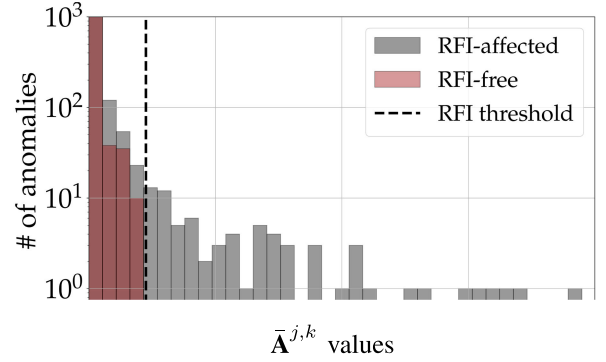


Fig. 5.  $\bar{\mathbf{A}}^{j,k}$  of (red) RFI-free images and (gray) RFI-affected images. The dotted black line illustrates the RFI threshold of 0.09.

which burst the RFI anomaly was located in. We then analyzed the anomaly in each “RFI signal subset” of size  $11 \times 166$ . RFI anomalies were defined where the mean of the anomalies in a subset was larger than 0.09, i.e., for each swath, for each consecutive subset of 11 rows, we now average the boxes of size  $11 \times 166$  in  $\mathbf{A}^{j,k}$ .

This gives the “average” anomaly,  $\bar{\mathbf{A}}^{j,k}$ , of size  $31 \times 3$ , i.e.,  $\lceil 340/11 \rceil \times 500/166 = 31 \times 3$ . We then define an RFI anomaly where  $\bar{\mathbf{A}}^{j,k} > 0.09$ .

Hence, if the mean reconstruction error of a single “RFI height” in a single “swath,” a subset of  $11 \times 166$ , was larger than 0.09, it was defined as an RFI anomaly independent on both geographical region or location, given that  $\mathbf{H}$  defined the region as being anomalous. Otherwise, it was classified as another type of anomaly. Likewise to the threshold for  $\mathbf{H}$ , the threshold for  $\bar{\mathbf{A}}$  was found by comparing  $\bar{\mathbf{A}}^{j,k}$  for the RFI-free and RFI-affected annotated images, as seen in Fig. 5.

#### E. RFI Radar Characterization in Level-0 Data

In few examples, the spatial structure of RFI signals in level-1 data could be used to determine radar characteristics. However, it was observed that in cases where the ground-based radars use, e.g., interleaved pulse repetition frequencies (PRFs) as to the left in Fig. 2, it was not possible. Furthermore, while the PRF could be estimated in some level-1 quick-looks, other characteristics like the pulse modulation could not. It is

therefore preferred to use the level-0 data when analyzing the interfering ground-based radar’s characteristics.

The level-0 data provided by ESA are not affected by the RFI signal smearing due to the SAR focusing, and thus the individual RFI signals from an external radar could be resolved. This allowed us to estimate basic pulse parameters like pulselength, PRF, and, in some cases, pulse modulation.

To identify pulse groups, we extracted a burst affected by RFI from the level-0 Sentinel-1 data, based on the RFI analysis of the quick-look images in  $\bar{A}^{j,k}$ . We then manually picked a prominent individual RFI signal, a reference pulse, extracted it from the level-0 data and cross-correlated each echo line with the extracted pulse in order to enhance pulses with a similar modulation. We then identified the maximum in each enhanced echo line and identified the presence of a pulse in an enhanced line if the maximum value exceeded 10% of the peak value of the autocorrelation of the reference pulse.

See Appendix A for a detailed explanation of the radar characterization in level-0 data.

### III. RESULTS

The RFI signals were classified and localized by a two-stage workflow. First, a CAE, trained on a Nvidia Tesla A100 80 GB GPU, was used to reconstruct the original SAR images. Regional reconstruction errors were used to make reconstruction heatmaps to find anomalous regions. The anomalies were then defined as either RFI anomalies or other using the elongation of the signal. By analyzing the corresponding Sentinel-1 level-0 data, the characteristics of the RFI sources were found.

#### A. CAE Reconstruction

The CAE was trained such that the same model can be used to find anomalies in all scenes illustrated in Fig. 1. Fig. 6 illustrates the fivefold cross-validation loss for the CAE reconstruction model. The fourfold training and onefold validation images did not have any RFI signals, since this would have allowed the model to reconstruct RFI signals. Less represented images in the dataset, i.e., images acquired from scenes with relatively few images, had higher MSSIM loss. The folds were created randomly, and so for each fold, both the training and validation sets differed geographically. The model learned the underlying distributions of each training fold similarly while the validation data had several spikes. These spikes were due to, e.g., less represented acquisition scenes. However, after 100 epochs the CAE learned to reconstruct them as well.

To compare the reconstruction capabilities of CAEs, we conducted experiments with a masked autoencoder (MAE) that used a 60% masking [32] of the input image. In contrast to CAEs, MAEs use transformer blocks [33], which replace convolutional layers with fully connected and multihead self-attention layers [34]. This results in a significant increase in computational complexity compared to CAEs, especially for large images but generally need less training data. We constructed a simple CAE and MAE reaching validation MMSIMs of 0.043 and 0.044, respectively, for the validation data, thus giving similar qualitative results. The




Fig. 6. Training and validation loss for the fivefold cross-validation.

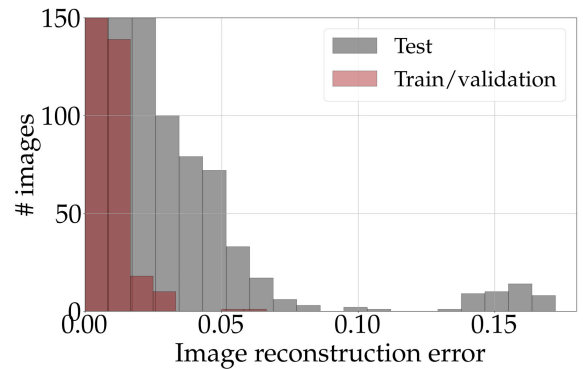


Fig. 7. Training, validation, and testing mean pixelwise reconstruction errors. A high error corresponds to images not well reconstructed, possibly stemming from an RFI signal.

MAE was better at reconstructing images with a high spatial correlation across the images but not as good at reconstructing images with less spatial information, i.e., near shore or islands. The MAE could therefore be slightly better for detecting RFI on land, while the CAE is better for RFI on the ocean. However, the MAE is more computationally expensive and requires a quadratic increase in the number of parameters relative to the input size. With an emphasis on RFI signals caused by ships, the CAE is argued to be better.

Fig. 7 illustrates the mean pixelwise reconstruction error of all the training and testing images inferred from the model with the lowest validation loss from the cross-validation, i.e.,  $\sum_j^{J=340} \sum_k^{K=500} (\mathbf{x}^{j,k} - \hat{\mathbf{x}}^{j,k}) / (J \cdot K)$ . The testing errors were higher than the training and validation errors due to the RFI-affected SAR images, and were not caused by an overfitting of the training images. The images mostly had low reconstruction errors due to the combined score for the entire image pair. For instance, while a geographical region of image  $\mathbf{x}_m$  was affected by an RFI signal, other regions of image  $\mathbf{x}_m$  were not. The resulting reconstruction error for the entire image  $\mathbf{x}_m$  was thus comparatively low. However, since the CAE was unsupervised it could not, by the reconstruction error alone, be verified if the high reconstruction errors were due to RFI signals or from other anomalies.



In the top of Fig. 8 we see three VH images, in the middle their reshaped latent space, and below their reconstructions. Two of the images were acquired near Dubai with one having a clear RFI signal and the other without. The images have similar latent space with minor differences resulting in both images being reconstructed without RFI signals. The bottom figure was acquired mid sea and has an RFI signal. The latent space has very little activation due to the spatially homogeneous sea. The RFI signal mid sea was likewise not reconstructed. The CAE could consequently reconstruct SAR images without RFI signals under varying conditions, even under strong winds causing much higher backscatter than the normality.

Other studies have argued that autoencoders can generalize well enough to reconstruct anomalies. Our CAE did not learn to reconstruct the RFI signals, and it was consequently not necessary to implement, e.g., a memory addressing unit as was done in [35]. But, we saw that some RFI-free images had relatively high reconstruction errors. This could possibly be mitigated by using the radiometric calibrated images, similar to [18], or generally by having more representative training data to better reconstruct the varying sea state.

The CAE model was therefore capable of reconstructing the general, normal structural, and spectral patterns in the input images including both stationary land features and some high intensity patterns on the oceans. Conversely, the model had high image reconstruction errors for the RFI-affected images.

### B. RFI Localization

The CAE reconstructed the SAR quick-look images without large-scale anomalies. The difference between the reconstructed and original images were thereafter used to classify the anomalies as either RFI signals or other anomalies. Fig. 9 illustrates the results from six different acquisition from the test set. The figure shows the VV and VH polarizations, along with the heatmap after the 0.15 noise threshold with the resulting detected RFI anomaly highlighted to the right.

We see that the images, and their RFI signals, have very different spatial characteristics. The RFI signals near UAE and Israel are all very visible in both the VH and VV polarizations corresponding to what is expected from different stationary radars transmitting in both polarizations. The RFI signals near Denmark and Japan are only visible in the VH polarization which is expected if the on-ground radar only transmits in the horizontal polarization. Generally, we see that the RFI signals have high MSSIM heatmap errors while the neighboring non-RFI affected regions have lower MSSIM errors. Especially the non-RFI affected regions near Dubai and Japan had low reconstruction errors. Conversely, we see that the acquisition near Denmark had high reconstruction errors on land. These findings correspond with what was expected from the training data distribution seen in Table II. More training data reduce the reconstruction error and vice versa. Furthermore, the region near Israel is greatly affected by RFI signals in most bursts in the third subswath. The method found RFI in all regions, independent of environmental conditions and was therefore generalizable for many scenarios, given enough training data.

TABLE III  
ACCURACY AND F1-SCORE OF THE 85 MANUALLY  
LABELED SAR IMAGES

Region	Proposed Method Accuracy% /F1%	Ref. [21] Method Accuracy% / F1%
Denmark	69/77	52/70
Israel	100/100	50/67
UAE	100/100	50/67
Japan	93/89	38/53
Overall	91/92	48/64

Nevertheless, in complex regions with little training data, other anomalies are classified as RFI signals, such as near Denmark. By changing the global threshold the number of false positives could be reduced. However, a high false positive rate is preferred over a low true positive rate when assuming that the RFI signals are caused by targets of interest. Instead, more training data should be added to the training.

Our RFI classification results were compared to that in [20] with the 85 annotated images using classification accuracy,  $= (TP + TN)/(TP + FP + TN + FN)$ , and second the F1-Score,  $= (TP)/(TP + 1/2(FP + FN))$ , where TN and FN are the true and false negatives, respectively, and FP and TP the false and true positives, respectively. The method described in [20] defined RFI anomalies pixel-wise by the condition

$$\frac{\sigma_{\text{Cross}}^0 - \sigma_{C_o}^0}{\sigma_{C_o}^0} H(\sigma_{\text{Cross}}^0 - \sigma_{C_o}^0) H(\sigma_{\text{Cross}}^0 - T) > 0 \quad (9)$$

where  $\sigma_{\text{Cross}}^0$  and  $\sigma_{C_o}^0$  are the radiometric calibrated cross and copolarizations of the multilooked SAR image, respectively, and  $H(\cdot)$  the Heaviside function. The threshold,  $T$  has to be fit for each region manually contrarily to our method in which global thresholds were used. The threshold  $T = 0.005$  was in [20] found empirically for a region around UAE and was used in this study (see Fig. 1). The results are shown in Table III. The rather low accuracy near Denmark stems from the high false positive rate due to the high reconstruction errors, as seen in Fig. 9.

The method in [20] classified all annotated images as containing RFI signals, independent of locations and true labels. Leng et al. [20] used a pixel-wise score, meaning all pixels with a certain value were labeled as RFI signals resulting in a very high false alarm rate, illustrating the problem of detecting RFI signals in SAR images by setting local thresholds manually for each scene. Conversely, the method in this study correctly classified anomalies with an accuracy of 91% and an F1 score of 92%, and further localized them automatically (see Table III). The method in the current study used the CAE to exploit the spatial signature of RFI signals to find a region in the SAR image containing an RFI signal, whereafter the RFI signal was localized in that region alone, thereby reducing the false alarm rate greatly while having the same high true positive rate.

However, our method has few shortcomings. The method struggled to locate RFI signals in very high-intensity regions such as cities, as seen to the left in Fig. 10. This was expected since cities themselves had intensities close to 1, and there were therefore no differences between the RFI signal and the






Fig. 8. (Top) Original VH quick-look images,  $\mathbf{x}$  (middle) their latent space illustrated as  $25 \times 5$  images,  $\mathbf{z}$  and (bottom) their reconstructions,  $\hat{\mathbf{x}}$ . (Left middle) UAE. (Right) Open sea near Denmark. The images were reconstructed without the RFI signals as intended.

cities. This could possibly be remedied by using the multilook images with larger dynamic range. Furthermore, empirically, some RFI signals were not extending throughout the entire swath in range or 3–5 km in azimuth. This caused few false negatives. Fig. 10 (right) shows a VH quick-look with an RFI signal that was not found. However, it was detected as an anomaly from the heatmap.

Stationary RFI signals could be localized in several images. By geolocating the localized RFI signals it was consequently possible to determine the approximate location of external, stationary radars as in [20]. This is illustrated in Fig. 11, where the RFI signals from two descending orbits and one ascending orbit were localized and illustrated. The corresponding VH quick-looks are shown below. We see the two RFI signals from the descending orbits are very similar and have the same intersection with the RFI signal from the ascending orbit.

It was possible to somewhat localize the RFI signals under varying geographical and environmental conditions. The heatmaps compared  $\mathbf{x}$  with  $\hat{\mathbf{x}}$ , and by removing the background noise, large anomalies were found. The anomalies were caused by varying phenomena, and the RFI signals were found by calculating the average anomaly in a region the size of an expected RFI signal. Our method could localize RFI signals mid ocean and on land. Localized RFI signals on land could further be used to determine the location of stationary radars. Visual inspection of the original image showed a correlation between the expected anomalies and the detected anomalies.

### C. RFI Characterization

Once RFI signals were detected and localized in the SAR quick-look images, it was possible to characterize the on-ground radar characteristics using Sentinel-1 level-0 data.

TABLE IV

RADAR CHARACTERIZATION CORRESPONDING TO FIG. 12

Parameter	Left image	Right image
Pulse width [ $\mu\text{s}$ ]	1	1
Phase modulation	None	None
PRF [Hz]	1000	400 / 500
Approx. blind speed [m/s]	$679 \cdot k$	$272 \cdot k / 340 \cdot k$

Fig. 12 illustrates two VH quick-look images, classified with RFI anomalies with our method. The RFI signals in the two images had similar spatial structures and were both located mid ocean.

The corresponding radar characteristics are shown in Table IV. The radars causing the RFI signals in Fig. 12 had very similar characteristics, and it was therefore hypothesized that they originate from the same type of radar.

It was found that the interfering radars near, e.g., Japan were often not employing a single pulse and PRF, but instead transmitting groups of, e.g., 16 pulses during which the pulselength, modulation, and PRF were kept constant, after which these parameters were changed. This is consistent with radars employing pulse-Doppler/Moving target indicator techniques to measure both target position and speed, i.e., to accommodate the blind-speed problem in which targets moving at the doppler frequency cannot be observed by the radar [36]. The approximated blind speeds for the radars in Table IV,  $v_b$ , were calculated as follows [36]:

$$v_b = \frac{k \cdot c \cdot \text{PRF}}{2f_c} \quad (10)$$

where  $c$  is the speed of light,  $f_c$  the center frequency, here approximated as the carrier frequency of the SAR and  $k = 1, 2, \dots$  the blind speed order. Using only a PRF

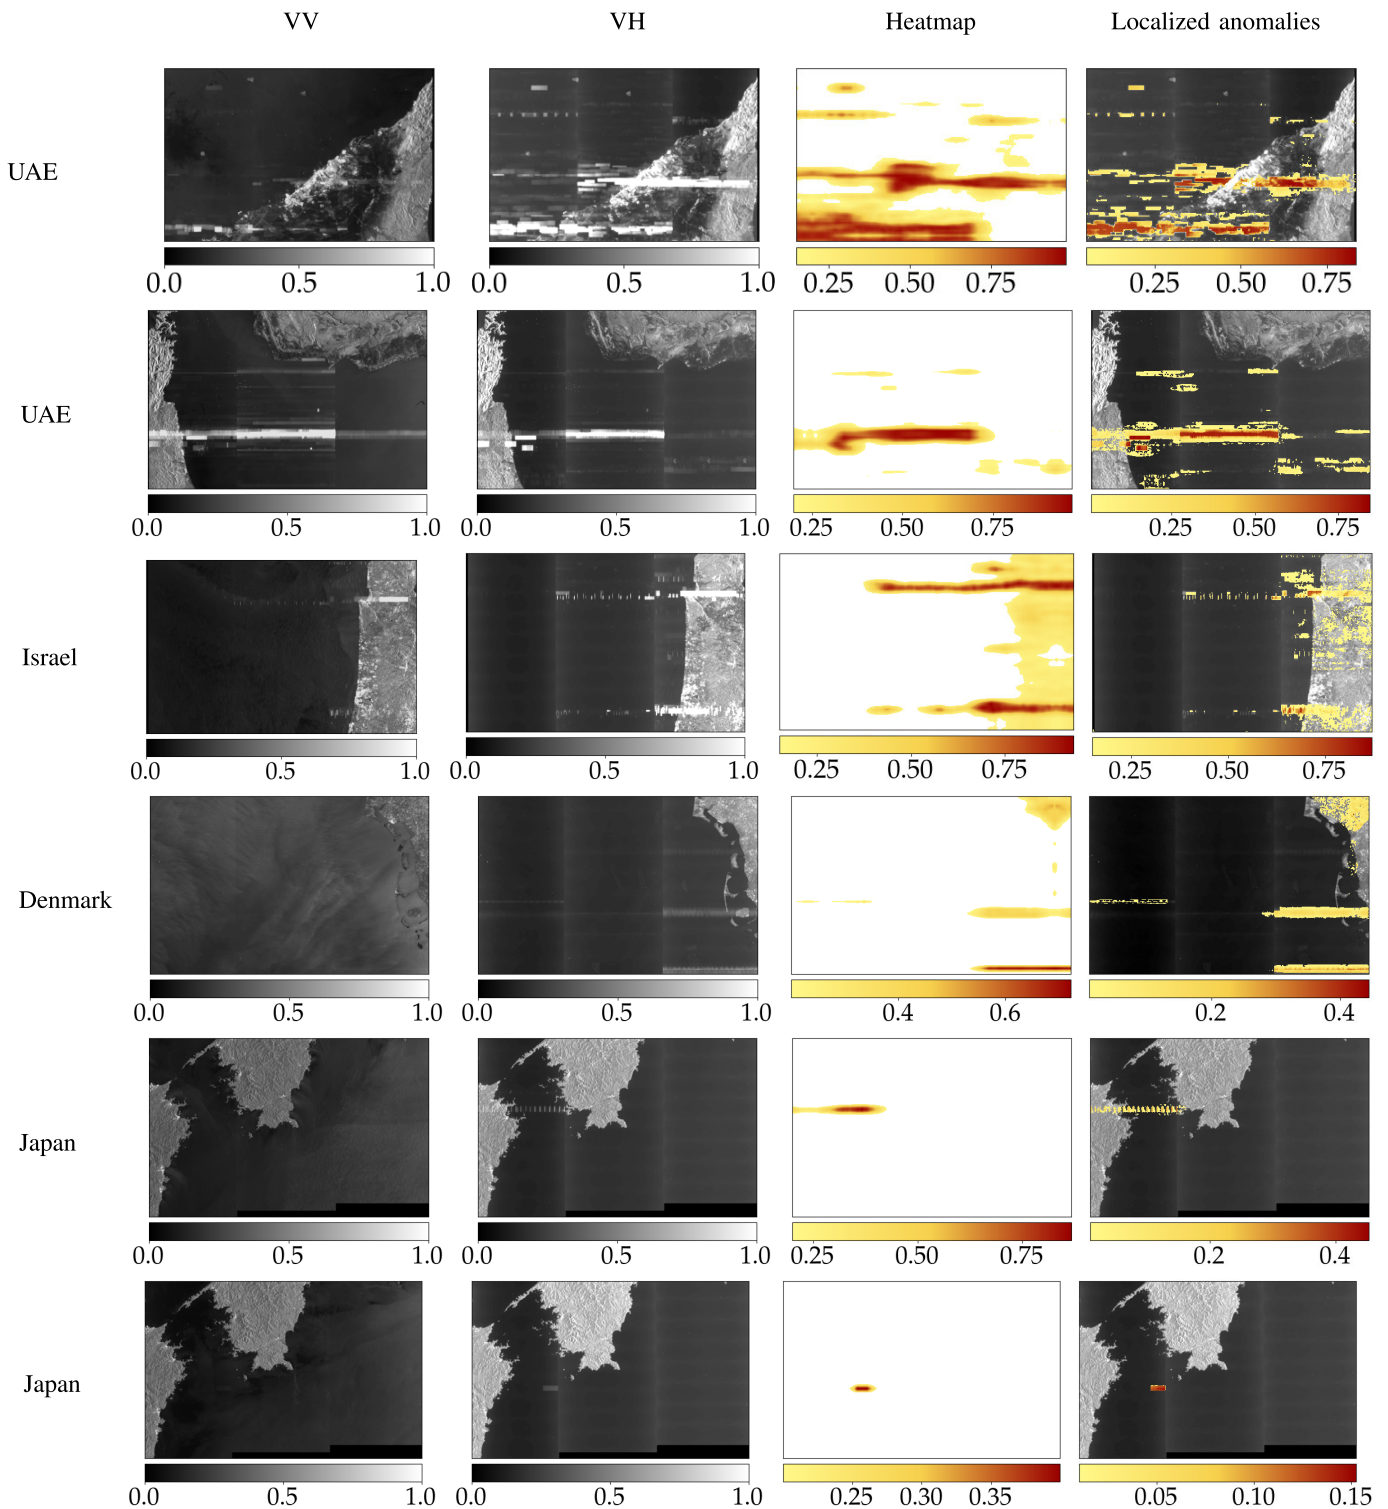


Fig. 9. Qualitative results of the RFI localization. (Left) Original VV quick-look image and (middle left) Original VH quick-look image, (middle right) Heatmap following a global, static, noise threshold of 0.15 with the color-bar illustrating the heatmap values and (right). The classified and localized RFI signals, highlighted on the VH quick-look images with the color-bar corresponding to the pixel-wise anomalies. The localized RFI signals had a mean reconstruction error of 0.15 within its burst.

of 500 Hz, it would not be possible to detect moving targets with a speed of approximately 340 m/s corresponding to Mach 1, or 679 m/s and thus motivating the use of staggered PRFs.

The RFI-affected level-0 data near Israel and UAE had similar structures. However, each level-0 image had very complicated radar characteristics with several different pulse lengths and several interleaved PRFs within each image,




Fig. 10. (Left) Anomaly classified as an RFI signal. However, the RFI signal on the city was not localized. (Right) Small RFI signal that was found and classified as a non-RFI anomaly.

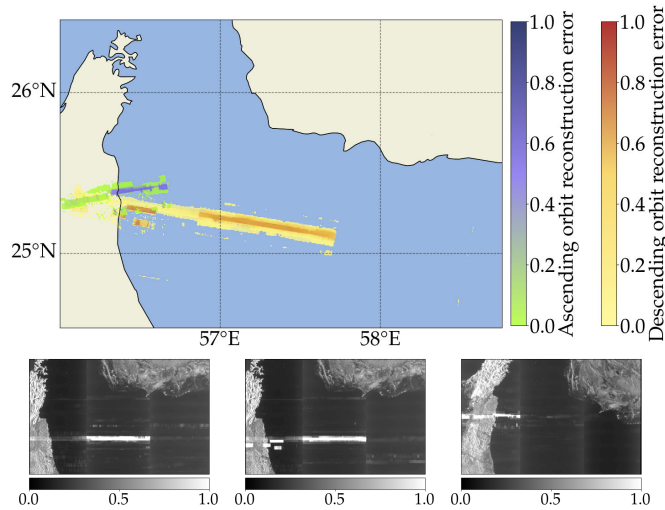


Fig. 11. (Top) Location of ground-based radar from (bottom) three images acquired near Dubai.




Fig. 12. Two quick-look images near the Island of Shikoku, Japan, both containing a single RFI signal originating from mid-sea ground-based radars. Images were acquired on (left) June 28, 2018 and (right) March 25, 2020.

complicating the characterization greatly. This might be caused by, e.g., a varying PRF.

For some RFI signals, we could likewise estimate the PRF using only the level-1 quick-looks. In Fig. 13, we see five VH quick-looks all acquired from the ocean near Japan. The signals all have periodical structures where two adjacent pulses are displaced horizontally after each other. This could be used to characterize both the PRF and the pulselength. However, for RFI signals like the ones near Dubai or Israel in Fig. 9 it was not possible to do so, since there was no clear periodicity in the signal. All RFI signal sources in Fig. 13 are believed to be ship-borne radars.

We detected the RFI signals in the level-1 quick-look images and extracted the information from the corresponding bursts in the level-0 data. From the level-0 data, we could determine,




Fig. 13. Five RFI signals located near Japan. All signals have comparable structures.

e.g., the pulselength, bandwidth, and PRF as described in Appendix A. For “simple” RFI signals as the ones observed in the ocean, this method worked well. For complex RFI signals as the ones near Israel it was very difficult.

#### IV. DISCUSSION

In this study, we demonstrate the possibility of localizing RFI signals in Sentinel-1 SAR quick-looks images automatically using a CAE. We first find anomalies in the images whereafter we classify them as RFI anomalies. The method consequently classifies and localizes RFI signal from several different geographical regions under varying environmental conditions with a mean accuracy of 91%. We furthermore show that the Sentinel-1 level-0 data can be used to determine specific characteristics of the RFI source.

Prior work has demonstrated the possibility of finding RFI signals in both level-0 and level-1 Sentinel-1 SAR data. Finding RFI signals in level-0 data is highly computational although it provides detailed information on the RFI sources. Conversely, finding RFI in, e.g., quick-look images reduces the needed computational power. Earlier studies classified both quick-look and Sentinel-1 IW GRD images into either a no-RFI or RFI class. But, such methods cannot be used to localize the RFI signals, and are used mainly to discard RFI affected images and is dependent on a large, representative manually created dataset. For instance, Chojka et al. [18] implemented a CNN to classify images with RFI signals, whereafter they could remove them in further analysis.



Nonetheless, a binary CNN classifier cannot handle out-of-distribution RFI signals, nor can it localize the RFI signals in the images. Chojka et al. [18] was further dependent on a highly limited labeled dataset, limiting the geographical regions. Hence, motivating our implemented methodology that handles both out-of-distribution RFI signal detection and localization in previously seen regions. Leng et al. [20] localized RFI signals in multilook images under the assumption that an RFI signal was present, by implementing a manual threshold-based method. They localized the RFI signals in few images. However, static threshold methods are not generalizable and thresholds must be found empirically for every region and cannot be used for images under varying geographical or environmental conditions. The technique described in [9] uses the rank echoes of the level-0 data. It can therefore not detect RFI signals if the external contribution was not present in the rank echoes, while our method has no such restriction.

We found that, given enough training images, our CAE can reconstruct SAR images in all geographical regions while excluding large anomalies, such as RFI signals. The RFI signals from complex regions are adequately localized. Likewise, RFI signals from spatially simple regions are easily localized owing to the large structural and spectral discrepancy between the RFI signatures and the image backgrounds. Our model thus localizes and highlights RFI over land and water bodies alike under different conditions.

As our model is implemented on SAR quick-look images, information on the RFI source is lost. Even so, test time for the CAE is of the order of seconds for an image, which means the CAE can be used efficiently on many images. By exploiting the power of deep learning methods, our model is generalizable to very different RFI structures on a global scale given a sufficiently large RFI-free training dataset, as opposed to [20]. However, if the model were to be used on larger scale, changes must be made to accommodate the varying latent space.

In this study, we furthermore determine the characteristics of the RFI source radar by analyzing the interference in the Sentinel-1 level-0 data. From the level-0 data, it is possible to determine radar characteristics such as the PRF and pulse length. We therefore exploit the advantages of both the quick-look images and the level-0 data.

While our model performs well in regions with enough training data, it struggles in regions with little training data. Some regions have little training data, such as near Denmark or Israel, as a result of the high frequency of RFI-affected images. This increases the total reconstruction error in complex regions, even when no RFI is present. The customized localization method tries to handle both situations by considering the known structures of RFI signals.

Satellite-borne SAR sensors have center frequencies in the microwave part of the electromagnetic spectrum, whereas few on-ground radars have similar frequencies, which means they do not have the ability to cause RFI signals in SAR images. However, few examples capable of causing RFI signals include stationary antimissile systems or anti-aircraft warfare radars placed on-board warships and some weather radars. With our model, we can localize the different RFI signals. For instance, monitoring RFI activity in, e.g., the ocean might increase

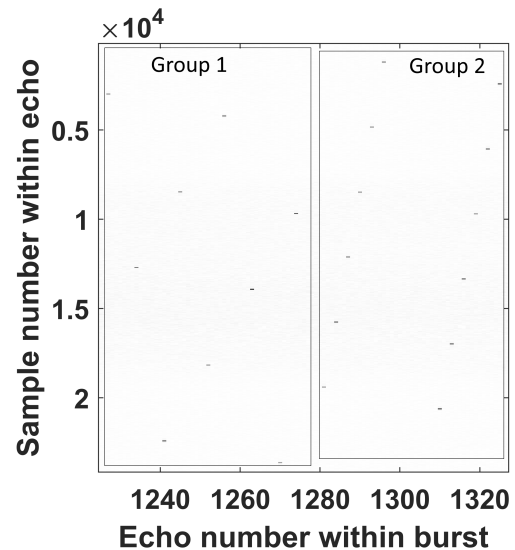


Fig. 14. Two RFI pulse groups of differing PRF seen in the raw data.

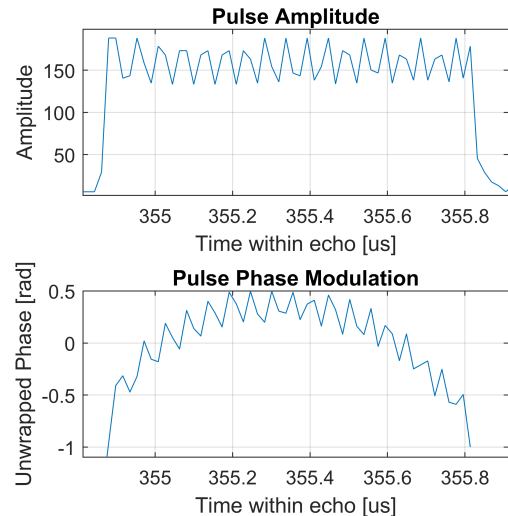


Fig. 15. Amplitude and phase modulation of a single RFI pulse.

maritime domain awareness by giving decision makers information on foreign navy vessel activity, while also increasing vessel detection [8]. Navy vessels rarely transmit location and voyage information, and such RFI activity will hence add information on the vessel type and potential threat.

Radars have different frequencies depending on the objective of their missions. The Sentinel-1 satellite has a C-band SAR sensor and can, therefore, only monitor RFI originating from on-ground C-band radars. Several different frequencies must subsequently be observed to monitor RFI activity fully, although the method developed in this study can be used for all frequencies given a representative dataset. The method developed in this study exploits the quick-look images for RFI localization and the level-0 data for RFI radar source characterization. Our method can detect stationary RFI signals and by exploiting the ascending/descending orbits of the Sentinel-1 satellite we can further estimate the corresponding geographi-



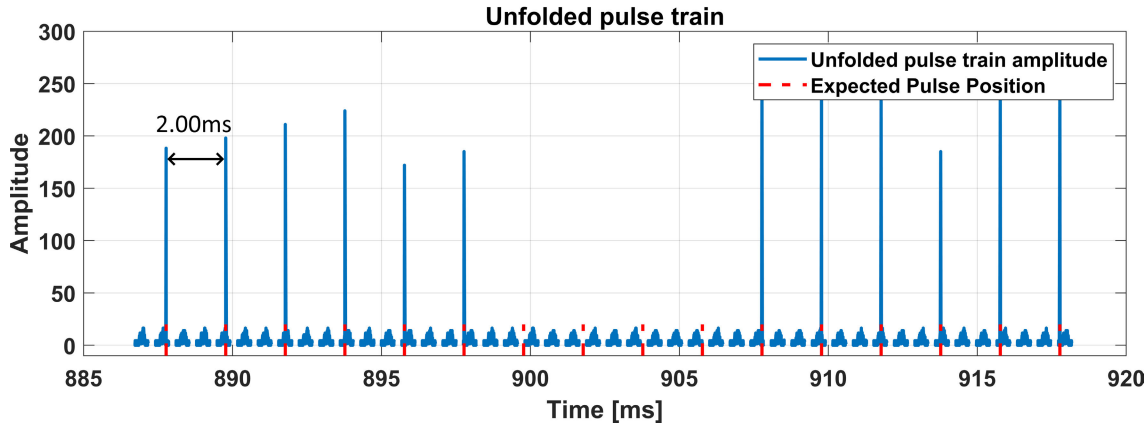


Fig. 16. Unfolded pulse train for pulse group 2 (see Fig. 14). A PRF of 500 Hz is estimated from the measured minimum pulse separation of 2 ms. Expected pulse positions with this PRF are indicated by red lines, showing that some pulses fall outside the Sentinel-1 sampling window. The time axis is referenced to the first sample of the first echo in the full burst.

cal locations. Likewise, our method detects nonstationary RFI signals, e.g., mid-sea, hypothesized to originate from shipborne air-surveillance radars. While the method cannot be used to detect the actual ships, it can be combined with other vessel detection methods to improve maritime domain awareness greatly (see [1]).

Future work could combine and compare RFI detection and ship detection automatically. Either such that after an RFI signal has been detected mid ocean, the corresponding Sentinel-1 single look complex (SLC) image could be analyzed for ships, or vice versa. Furthermore, more work could be done on RFI classification. By analyzing more structures in the level-0 data and correlating the structures to the level-1 images, it could be possible to further classify the radar type and characteristics using only the level-1 images followed by an analysis of, e.g., where the different radar types are observed. On land and or ocean.

#### APPENDIX A DERIVING RFI PULSE CHARACTERISTICS FROM LEVEL 0 DATA

The Sentinel-1 level-0 data provided by ESA contain the raw data packets downlinked from the SAR instrument on the satellite. The packets contain the raw echoes received by the radar, as well as annotation information required to unpack and accurately time the individual echoes. We used an open source software package [37] to unpack the level-0 echoes. As described in Section II-A, data are acquired in burst of approximately 1 s, switching every second between the three swaths IW1-IW3, so the natural unit of analysis is one burst. In the following, we will explain how the RFI in Fig. 12 (left) was characterized from the raw data. First, the swath and burst containing the RFI was inferred. Then, the raw echo matrix for the identified burst was extracted from the corresponding level-0 dataset. This matrix contains the complex I/Q-demodulated echoes (a superposition of echoes of the Sentinel-1 radar pulses and RFI from interfering sources). By plotting the raw data matrix (see Fig. 14), we can identify groups of RFI pulses with the same PRF, as they occur along

straight lines in the raw data image. Note that Sentinel-1 does not receive continuously, which means some RFI pulses of a constant PRF pulse group may be missed.

In the following section, we proceed by analyzing pulse group 2 from Fig. 14. First, a single pulse is extracted manually from the raw data by visual inspection of an echo line. In Fig. 15, the amplitude and phase of the complex samples covering the pulse is shown. The phase modulation is inferred by unwrapping the phase across the pulsewidth and removing a linear trend, that could result from the Sentinel-1 receiver and the RFI radar employing different center frequencies. In this case, we see a negligible phase modulation, suggesting that the RFI transmitter employs simple constant frequency pulses. The pulsewidth can be directly determined from the amplitude plot, and is measured to be  $1 \mu\text{s}$ . Note that in general, RFI pulses are not so well-defined if, e.g., they are linear frequency modulated, as parts of the pulse frequency sweep may fall outside the analog bandwidth of the Sentinel-1 receiver, and also they may be further distorted by the decimation filter employed online in the SAR instrument to reduce the sampling rate of the raw data from 300 MHz to a swath-dependent value between 47 and 64 MHz before downlink.

#### APPENDIX B DERIVING THE PRF OF THE RFI RADAR FROM LEVEL-0 DATA

To derive the PRF of the pulse group under investigation, we first do a matched filtering (cross correlation) of each echo in the pulse group with a copy of the (complex) extracted pulse that was found in Appendix A (see Fig. 15). This results in a more well-defined peak at the pulse leading edge, and optimizes the signal-to-noise ratio for detecting replicas of that pulse. In case of multiple interfering radars using different pulse modulations, this also makes it easier to distinguish between them. For each of the matched filtered echo lines, we then find the peak and median amplitude, and if the peak value is at least  $5\times$  the median, we interpret this as the presence of an RFI pulse at the sample corresponding to the peak position. Note that a more advanced method is required

in case the RFI radar PRF is so high that more than one pulse is received in an echo line, specifically if  $T_{\text{echo}} > 1/\text{PRF}_{\text{RFI}}$ , where  $T_{\text{echo}}$  is the duration of a received echo line. For each peak, we determine the exact time of reception,  $T_{pk}$ , referenced to the first sample of the first echo line in the examined pulse group, by

$$T_{pk} = \frac{M}{\text{PRF}_{S1}} + \frac{N}{f_s} \quad (11)$$

where  $M$  is the number of the echo lines containing the peak,  $\text{PRF}_{S1}$  is the PRF employed by the Sentinel-1 SAR,  $N$  is the sample number of the peak within the echo line,  $f_s$  is the sampling frequency employed by Sentinel-1 to sample each echo, and it is assumed that the Sentinel-1 sampling window start time (SWST) does not change during the acquisition of the pulse group. We then calculate the difference between successive peaks, and select the minimum as the RFI pulse repetition interval,  $\text{PRI}_{\text{RFI}}$ , related to the PRF by  $\text{PRF}_{\text{RFI}} = 1/\text{PRI}_{\text{RFI}}$ . The reason for using the minimum is that Sentinel-1 does not sample continuously, so all RFI pulses may not arrive within the sampling windows. To illustrate this, we plot in Fig. 16 the amplitude of all the SAR echoes covering pulse group 2 on a continuous time scale. The estimated  $\text{PRI}_{\text{RFI}}$  of 2.00 ms (corresponding to a  $\text{PRF}_{\text{RFI}}$  of 500 Hz) is illustrated, and we indicate also the expected pulse positions using the estimated PRF. From the illustration, it is clear that 4 out of 16 expected pulses fall outside the sampling windows, and we hypothesize that the radar did send groups of 16 pulses, even though we only observe 12 of them. An analysis of pulse group 1 using the same approach led to an estimated PRF of 400 Hz, and a pulse group length of 16.

#### ACKNOWLEDGMENT

The authors would like to thank the Danish Defense for support, in particular, the joint arctic command (JAC) and the Danish Defense Acquisition and Logistics Organization (DALO).

#### REFERENCES

- [1] P. Heiselberg, K. A. Sørensen, H. Heiselberg, and O. B. Andersen, "SAR ship-iceberg discrimination in Arctic conditions using deep learning," *Remote Sens.*, vol. 14, no. 9, p. 2236, 2022.
- [2] P. de Mattheaïs, R. Oliva, Y. Soldo, and S. Cruz-Pol, "Spectrum management and its importance for microwave remote sensing [technical committees]," *IEEE Geosci. Remote Sens. Mag.*, vol. 6, no. 2, pp. 17–25, Jun. 2018.
- [3] M. Tao, J. Su, Y. Huang, and L. Wang, "Mitigation of radio frequency interference in synthetic aperture radar data: Current status and future trends," *Remote Sens.*, vol. 11, no. 20, p. 2438, Oct. 2019.
- [4] R. Natsuaki, T. Motohka, and T. T. S. Suzuki, "Polarimetric characteristics of temporarily coherent RFI in ALOS-2 PALSAR-2," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2017, pp. 3155–3158.
- [5] N. Li, Z. Lv, and Z. Guo, "Observation and mitigation of mutual RFI between SAR satellites: A case study between Chinese GaoFen-3 and European Sentinel-1A," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5112819.
- [6] P. Meadows, G. Hajduch, and N. Miranda, "Sentinel-1 long duration mutual interference," Sentinel-1 Mission Perform. Center, Saint-Agne, France, Tech. Rep. MPC-0432, 2018.
- [7] G. Hajduch, P. Meadows, and N. Miranda, "Sentinel-1 RadarSat-2 mutual interference," Sentinel-1 Mission Perform. Center, Saint-Agne, France, Tech. Rep. MPC-0353, 2017.
- [8] P. Shao, X. Lu, P. Huang, W. Xu, and Y. Dong, "Impact analysis of radio frequency interference on SAR image ship detection based on deep learning," in *Proc. IEEE Int. Geosci. Remote Sensing Symp.*, Sep. 2020, pp. 2447–2450.
- [9] A. Monti-Guarnieri, D. Giudici, and A. Recchia, "Identification of C-band radio frequency interferences from Sentinel-1 data," *Remote Sens.*, vol. 9, no. 11, p. 1183, Nov. 2017.
- [10] P. Artiemjew, A. Chojka, and J. Rapinski, "Deep learning for RFI artifact recognition in Sentinel-1 data," *Remote Sens.*, vol. 13, no. 1, p. 7, Dec. 2020.
- [11] A. Recchia, D. Giudici, R. Piantanida, N. Franceschi, A. Monti-Guarnieri, and N. Miranda, "On the effective usage of Sentinel-1 noise pulses for denoising and RFI identification," in *Proc. 12th Eur. Conf. Synth. Aperture Radar*, Jun. 2018, pp. 1–5.
- [12] W. Fan et al., "Interference mitigation for synthetic aperture radar based on deep residual network," *Remote Sens.*, vol. 11, no. 14, p. 1654, Jul. 2019.
- [13] P. Parasher, K. M. Aggarwal, and V. M. Ramanujam, "RFI detection and mitigation in SAR data," in *Proc. URSI Asia-Pacific Radio Sci. Conf. (AP-RASC)*, Mar. 2019, pp. 1–4.
- [14] J. Yu, J. Li, B. Sun, J. Chen, and C. Li, "Multiclass radio frequency interference detection and suppression for SAR based on the single shot MultiBox detector," *Sensors*, vol. 18, no. 11, p. 4034, Nov. 2018.
- [15] J. Su, M. Xi, Y. Gong, M. Tao, Y. Fan, and L. Wang, "Time-varying wideband interference mitigation for SAR via time-frequency-pulse joint decomposition algorithm," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5118816.
- [16] G. Hajduch, N. Franceschi, M. Pinheiro, and A. Valentino, "SAR-MPC Sentinel-1: Using the RFI annotations," Sentinel-1 Mission Perform. Center, Saint-Agne, France, Tech. Rep. MPC-0540, 022.
- [17] F. J. Meyer, J. B. Nicoll, and A. P. Doulgeris, "Correction and characterization of radio frequency interference signatures in L-band synthetic aperture radar data," *IEEE Trans. Geosci. Remote Sens.*, vol. 51, no. 10, pp. 4961–4972, Oct. 2013.
- [18] A. Chojka, P. Artiemjew, and J. Rapinski, "RFI artefacts detection in Sentinel-1 Level-1 SLC data based on image processing techniques," *Sensors*, vol. 20, no. 10, p. 2919, May 2020.
- [19] Y. Junfei, L. Jingwen, S. Bing, and J. Yuming, "Barrage jamming detection and classification based on convolutional neural network for synthetic aperture radar," in *Proc. IEEE Int. Geosci. Remote Sens. Symp.*, Jul. 2018, pp. 4583–4586.
- [20] X. Leng, K. Ji, and G. Kuang, "Radio frequency interference detection and localization in Sentinel-1 images," *IEEE Trans. Geosci. Remote Sens.*, vol. 59, no. 11, pp. 9270–9281, Nov. 2021.
- [21] M. Tao et al., "Radio frequency interference signature detection in radar remote sensing image using semantic cognition enhancement network," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 5231714.
- [22] R. Piantanida, G. Hajduch, and J. Poullaouec, "Sentinel-1 Level 1 detailed algorithm definition," Sentinel-1 Mission Perform. Center, Saint-Agne, France, Tech. Rep. MPC-0307, 2022.
- [23] *Sentinel-1 Level-0 Data Decoding Package*, European Space Agency, Paris, France, 2020.
- [24] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 448–456.
- [25] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh, and A. Talwalkar, "Hyperband: A novel bandit-based approach to hyperparameter optimization," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 6765–6816, 2018.
- [26] T. O'Malley, E. Bursztein, J. Long, F. Chollet, H. Jin, and L. Invernizzi. (2019). *Keras-tuner*. [Online]. Available: <https://github.com/keras-team/keras-tuner>
- [27] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *Proc. 37th Asilomar Conf. Signals, Syst. Comput.*, 2003, pp. 1398–1402.
- [28] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [29] K. Ding, K. Ma, S. Wang, and E. P. Simoncelli, "Comparison of full-reference image quality models for optimization of image processing systems," *Int. J. Comput. Vis.*, vol. 129, no. 4, pp. 1258–1281, Apr. 2021.
- [30] U. Sara, M. Akter, and M. S. Uddin, "Image quality assessment through FSIM, SSIM, MSE and PSNR—A comparative study," *J. Comput. Commun.*, vol. 7, no. 3, pp. 8–18, 2019.

- [31] A. K. Venkataramanan, C. Wu, A. C. Bovik, I. Katsavounidis, and Z. Shahid, "A hitchhiker's guide to structural similarity," *IEEE Access*, vol. 9, pp. 28872–28896, 2021.
- [32] K. He, X. Chen, S. Xie, Y. Li, P. Dollár, and R. Girshick, "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2021, pp. 16000–16009.
- [33] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," 2021, *arXiv:2010.11929*.
- [34] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 1–11.
- [35] D. Gong et al., "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 1705–1714.
- [36] P. Sedivy, "Radar PRF staggering and agility control maximizing overall blind speed," in *Proc. Conf. Microw. Techn. (COMITE)*, Apr. 2013, pp. 197–200.
- [37] J. Friedt. (2021). *Sentinel-1 Level 0 Decoding Package*. [Online]. Available: [https://github.com/jmfriedt/sentinel1\\_level0](https://github.com/jmfriedt/sentinel1_level0)



**Kristian Aalling Sørensen** (Graduate Student Member, IEEE) received the M.Sc.Eng. degree in Earth and space physics and engineering from the Technical University of Denmark (DTU), Kongens Lyngby, Denmark, in 2020, with specialization in Earth observation, where he is currently pursuing the Ph.D. degree with the Center for Security, National Space Institute of Denmark.

He was a Sergeant with the Engineering Regiment in Denmark. He is an Alumnus from the University of British Columbia, Vancouver, BC, Canada, and the Pohang University of Science and Technology, Pohang, South Korea. His research interest is on increasing maritime domain awareness using satellite data and artificial intelligence and include; ship detection in satellite data, data fusion and aggregation or various data sources, and anomaly detection in both temporal and spatial data.



**Anders Kusk** received the M.Sc. degree in electrical engineering and the Ph.D. degree in synthetic aperture radar (SAR) processing techniques from the Technical University of Denmark (DTU), Kongens Lyngby, Denmark, in 2002 and 2006, respectively.

He was a Post-Doctoral Researcher with DTU. Since 2009, he has been with the Microwaves and Remote Sensing Division, DTU, as a Research Engineer. His research interests include synthetic aperture radar (SAR) processing, ice flow mapping from SAR data, and participation in system development, processor development, and campaign activities for ESA's P-band polarimetric airborne radar ice sounder (POLARIS).



**Peder Heiselberg** received the bachelor's degree in physics and the master's degree in geophysics from the University of Copenhagen, Copenhagen, Denmark, in 2014 and 2017, respectively. He is currently pursuing the Ph.D. degree with the Technical University of Denmark, Kongens Lyngby, Denmark.

His research interests include machine learning, ship detection in synthetic aperture radar (SAR) images, and AIS-SAR data fusion.



**Henning Heiselberg** received the B.Sc. degree in mathematics and the Ph.D. degree in physics from the University of Aarhus, Aarhus, Denmark, in 1982 and 1987, respectively.

He was a Post-Doctoral Fellow with the Niels Bohr Institute, Copenhagen, Denmark, and the University of Illinois at Urbana–Champaign, Champaign, IL, USA. He was a Staff Scientist with the Lawrence Berkeley Laboratory, Berkeley, CA, USA. He was an Associate Professor with the Nordic Institute of Theoretical Physics, Stockholm, Sweden. He was an Adjoint Professor with the University of South Denmark, Odense, Denmark. He was a Senior Scientist with Danish Defense Research for 15 years, where he also participated in the MoD's New Fighter Program, and writing the Arctic and National Space Strategies. He is currently a Senior Scientist with the Technical University Denmark (DTU) Space, Kongens Lyngby, Denmark, and the Head of the Center for Security. His research interests cover number of topics within physics, astrophysics, sensor and space technology in general, remote sensing and Earth observation, and Arctic surveillance.