Remote Sensing Applications: Society and Environment 41 (2026) 101958

Contents lists available at ScienceDirect

Remote Sensing Applications: Society and Environment

journal homepage: www.elsevier.com/locate/rsase

Patch-based anomaly detection on SAR images to localize hotspots on the North and Baltic Sea coasts

Wadim Koslow [a],*, Kathrin Rack [a], Tobias D. Grabosch [a], Alexander Rüttgers [a], Luca Dell'Amore [b], Paola Rizzoli [b]

[a] *Institute of Software Technology, German Aerospace Center (DLR), High-Performance Computing Department, Linder Höhe, Köln, 51147, Germany*

[b] *Microwaves and Radar Institute, German Aerospace Center (DLR), Satellite-SAR-Systems Department, Muenchener Strasse 20, Wessling, 82234, Germany*

ARTICLE INFO

Keywords:
Synthetic Aperture Radar
Anomaly detection
Hotspot localization
Extreme weather events
Coastal protection
Unsupervised learning

ABSTRACT

In recent years, the vulnerability of coastal regions has increased significantly due to the effects of climate change. Measures must be taken to protect these coastal regions, which are disproportionately affected by extreme weather events and other damaging factors, and to increase their resilience. In this study, we propose a conceptual patch-based extension to the unsupervised Local Outlier Factor (LOF) anomaly detection algorithm to enable hotspot detection in Earth observation data. We validate our approach on Synthetic Aperture Radar (SAR) data using both synthetic and real-world anomalies and demonstrate that these methods outperform an autoencoder and a temporal Reed–Xiaoli (RX) approach, which are widely used for anomaly detection. Additionally, we generate coastal hotspot maps that identify areas requiring greater protection against extreme weather events and other hazards. These maps allow us to provide recommendations to decision-makers and governance bodies.

1. Introduction

The coastal regions of the North Sea and the Baltic Sea are increasingly threatened by rising sea levels and extreme weather events like storms and storm tides, which are linked to climate change and expected to grow in frequency and intensity (Intergovernmental Panel on Climate Change, 2023). Resulting coastal changes, such as erosion, can severely impact landscapes, populations, and infrastructure, making early detection crucial.

This work differentiates between short-term, event-driven changes (*anomalies*) and longer-term, recurring or gradually intensifying changes (*hotspots*). Anomalies reflect sudden impacts, whereas hotspots mark areas of sustained or escalating risk that require ongoing observation.

We employ time series of space-borne synthetic aperture radar (SAR) images to identify anomalies and hotspots. The main advantage of SAR over optical systems is its ability to continuously and regularly monitor the regions of interest without being significantly affected by atmospheric factors (Flores et al., 2019). In fact, the acquisition of SAR data can be performed independently of cloud cover, day-light or weather conditions. Traditional approaches to anomaly detection in SAR imagery rely heavily on statistical and signal based detectors. One of the most wildly adopted of these is the RX (Reed and Yu, 1990) detector and its numerous variants, such as the kernel RX (Kwon and Nasrabadi, 2005). These algorithms remain foundational in hyperspectral and SAR anomaly detection due to their interpretability and computational efficiency.

---

* Corresponding author.
E-mail address: wadim.koslow@dlr.de (W. Koslow).

https://doi.org/10.1016/j.rsase.2026.101958
Received 22 October 2025; Received in revised form 29 January 2026; Accepted 2 March 2026

In recent years deep learning methods have increasingly dominated SAR anomaly detection research. An AE learns low dimensional representations of normal background scenes flagging high reconstruction errors as anomalies (Chen et al., 2018). More advanced variants have been successfully applied to SAR data such as variational autoencoders (Sinha et al., 2021) or a self-supervising adversarial autoencoder (Muzeau et al., 2022). Ideas related to reconstruction have also advanced rapidly in the field of hyperspectral remote sensing. For instance, there are blind-spot architectures that suppress anomalous pixels during background reconstruction (Wang et al., 2024), semisupervised GAN-based methods that improve background-anomaly separability via spectral-domain transformations (Wang et al., 2023), lightweight cross-image VAE-style approaches designed to generalize to unseen scenes (Wu et al., 2025), and masking-based dual-autoencoder designs for robust background reconstruction in target detection settings (Chen et al., 2025). While these works are not SAR-specific, they motivate reconstruction- and patch-based formulations that reduce reliance on dense annotations and aim to improve robustness across scenes.

Supervised deep architectures such as U-Nets, are now standard in specific anomaly contexts including oil spill detection (Zakzouk et al., 2025) and ship detection (Cui et al., 2021). These models achieve state-of-the-art accuracy but require large, annotated datasets that are often unavailable in remote coastal regions. Furthermore, their generalization across sensors, incidence angles and geographic contexts remains a challenge. Closely related “zero-shot” and test-time reconstruction priors have also been explored in remote sensing. For example, there are enhanced deep image prior variants for hyperspectral super-resolution that learn scene-specific structure directly from observations (Li et al., 2025). This highlights a complementary direction when paired training data are scarce. The LOF (Breunig et al., 2000) algorithm offers a compelling alternative to both statistical and deep learning approaches.

It is a density-based method that has been successfully applied to various problems, such as the monitoring of complex multimode industrial processes (Ma et al., 2013), the detection of anomalies in computer networks (Auskalnis et al., 2018), money laundering detection (Gao, 2009) and even video data of hybrid rocket combustion tests (Rüttgers and Petrarolo, 2021). However its use in SAR anomaly detection remains largely unexplored. The LOF computes the density of each point compared to its neighbors to identify regions with significantly lower densities, indicating outliers, while adapting to local variations in density. This is particularly well suited to the spatially heterogeneous backscatter of coastal SAR images. However, when applied directly to image data, LOF operates only at the image level and therefore lacks the ability to localize anomalies within an image. To overcome this limitation, we propose a patch-based LOF framework that enables pixel-level anomaly localization in SAR imagery. Furthermore, we employ the Mahalanobis distance, which is better suited to the statistical characteristics of SAR data than the commonly used Euclidean distance. Additionally, we implement a temporal variant of the RX algorithm as well as a simple AE, which serve as a baseline and comparison to the results of the LOF.

However, detecting and interpreting anomalies and hotspots in coastal regions remains difficult. Tides introduce natural variability, making it hard to distinguish abnormal changes. Additionally, the absence of ground-truth data hinders reliable evaluation. To address this, we incorporate synthetic anomalies for validation which is a common approach when working with unlabeled datasets (Bouman et al., 2024).

The remainder of this article is organized as follows: Section 2 introduces the SAR and weather data, the preprocessing steps, and our approach to generating synthetic anomalies. Then, in Section 3, we present the LOF model, the RX, and the AE including the proposed training procedure. In particular, we introduce our implementation of a patch-based anomaly detection. In Section 4, we apply the different methods to synthetic and real-world anomalies, compare the results, and validate the results by correlating weather data to the anomaly detection. In Section 5, we analyze further implementations of an AE, examine whether the RX or LOF can be favored, analyze hotspot maps, and discuss how promising the methods are for the future. Finally, the conclusions are drawn in Section 6.

2. Data

Here, we introduce the SAR data processing, the selection of coastal region, the considered weather data, and the way we have inserted synthetic anomalies into the data. The considered SAR data have been acquired between 2016 and 2023 for the North Sea and between 2018 and 2020 for the Baltic Sea, both by the Sentinel-1 (S1) mission. The temporal repeat cycle was 6–12 days, depending on whether one or two satellites were operational. The S1-satellites have a polar orbit with an altitude of 693 km and a use a C-band SAR at 5.405 GHz. Although the considered data contains a large part of the Baltic and North Sea, we focus on two German sub-regions, i.e. Fehmarn island and the Cuxhaven region, depicted in Fig. 1. However, our findings can be extended to other regions.

2.1. Processing of SAR data

The time series used consists of SAR images acquired in interferometric wide swath (IW) mode, with a nominal pixel spacing of about $3.7\,\mathrm{m}$ and $14\,\mathrm{m}$ in ground range and azimuth direction, respectively. These S1 images are typically characterized by a wide swath of about $250\,\mathrm{km}$, along the range dimension. As a first step, co-registration was performed by resampling all images on the same spatial grid in terms of radar coordinates. Then, in addition to the generation of basic SAR products, i.e. amplitude and intensity data, we derived all possible interferometric coherence data between image pairs at 6–12-days temporal baselines. Finally, all SAR products were geocoded to a final position of $20\,\mathrm{m}$, in order to associate each pixel with its actual position on the ground as a function of latitude/longitude coordinates. Additionally, as a preprocessing step before applying the algorithms, we take the logarithm of the data, in order to make it more gaussian.

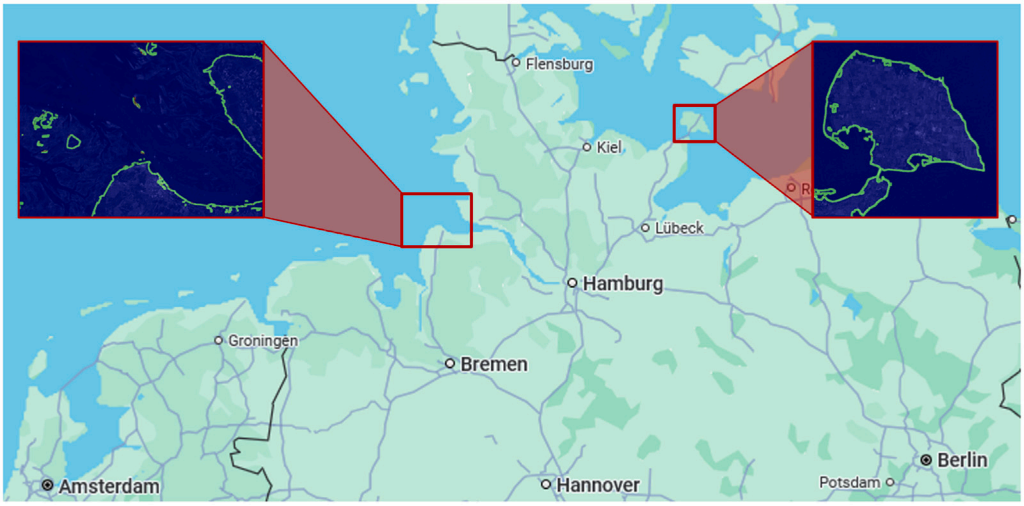

Fig. 1. Illustration of the two regions Fehmarn island (right) and Cuxhaven region (left), in the context of north Germany, monitored in this work.

2.2. Coastal mask

Here, we are interested in the coastal region, thus we apply a coastline mask, that allows to exclude irrelevant areas. To create the masks, a U-Net architecture (Ronneberger et al., 2015) was first trained to segment images into water and non-water bodies. A small architecture has been implemented (only three layers), given the very simple problem that has to be addressed. For this scope, the training dataset was retrieved over the Mediterranean Sea and using coastal regions where tides are assumed to be negligible. In particular, several S1 short time-series (with about 15 acquisitions each, processed at a final posting of $20\,\mathrm{m} \times 20\,\mathrm{m}$) were processed, one for each test site depicted in Fig. 1; also, all possible coast orientations and coastlines were used for training the network in order to generalize the trained models as much as possible. The water masks for each image of the time-series have been averaged to derive a more robust estimation of the coastline. The resulting coastline mask can be seen in Fig. 1 over the island of Fehmarn and the Cuxhaven region. Note that the mask also contains shorelines of lakes and other bodies of water. Although they are not the focus of our study, we keep these areas in our dataset because they are relatively small in number and do not have a negative impact on our methods.

2.3. Weather data

In order to determine which weather prevailed on the considered dates, we consult the data from meteostat (Lamprecht, 2025). The weather station with number 10131 in Cuxhaven, from which we took the data, is located at the coordinates N53.8683° E8.699°. Prominent storm events in the area of Cuxhaven in the considered time period are storm Herwart on 2017-10-27, Storm Eugen on 2021-05-05, and Storm Malik (German name: Nadia) on 2022-02-27, compare Fig. 10.

2.4. Synthetic anomalies

Analyzing real anomalies as expected due to erosion, storm surges, or other natural phenomena is difficult, due to the lack of knowledge about the appearance of these kind of anomalies due to their rareness. In addition, it cannot be ruled out that future anomalies will appear differently from previous ones. A valuable addition is the use of synthetic anomalies, also called artificial anomalies, which allow accurate validation, provided that the generated anomalies are qualitatively equivalent to the real anomalies. Recent approaches to synthetic anomalies in image data can be found in Fung et al. (2025) and in Kim and Lee (2024). Due to the lack of information about the appearance of natural anomalies, it is important to see how accurately the model can find anomalies of various shape from large (i.e. areas) to small (e.g. single pixels) sizes. Fig. 2 depicts the three spatial distribution types of synthetic anomalies that we investigate here: An annulus, single pixels, and areas.

Furthermore, in order to design the anomalies as realistic as possible, we analyzed the distributions of all pixel values along the coastline for each image in our time series by fitting and selecting the best fitting distribution using the “Fitter” Python library. The results show, that most of the pixel values in the dates follow a log-normal distribution, where the values are in floating-point representation ranging from zero (black pixel) to one (white pixel). While single-look SAR amplitude data is usually Rayleigh-distributed according to literature, this does not apply to our case. One explanation might be that we are focusing on the coastline instead of analyzing the full image. The mean parameters of the distribution for all dates are $\mu = 0.19(\pm 0.03)$, $\sigma = 0.52(\pm 0.07)$, and $\mathit{shift} = 0.004(\pm 0.01)$. However, for the above-mentioned storm events the pixel values follow different distributions with a much higher average image brightness. Fig. 3 shows a comparison between a sample of three regular dates and the three storm dates. By randomly sampling the synthetic anomalies from the distributions of these storm events, they represent real-world anomalies.

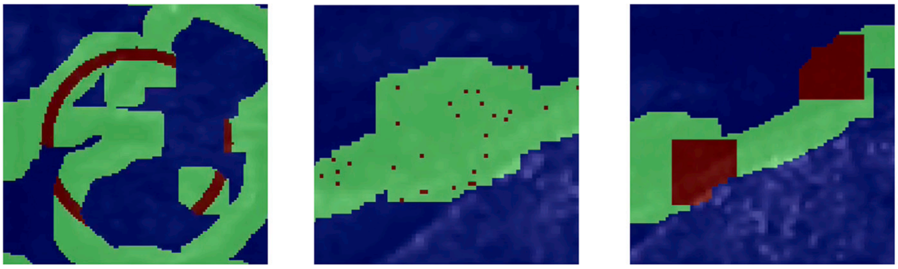

Annulus Single Area

Fig. 2. Three different types of pattern for synthetic anomalies. Pixels that have been altered with synthetic anomalies are depicted in red, while the pixels in green have not been modified. Blue pixels are outside of the coastline mask and are not considered by our algorithms.

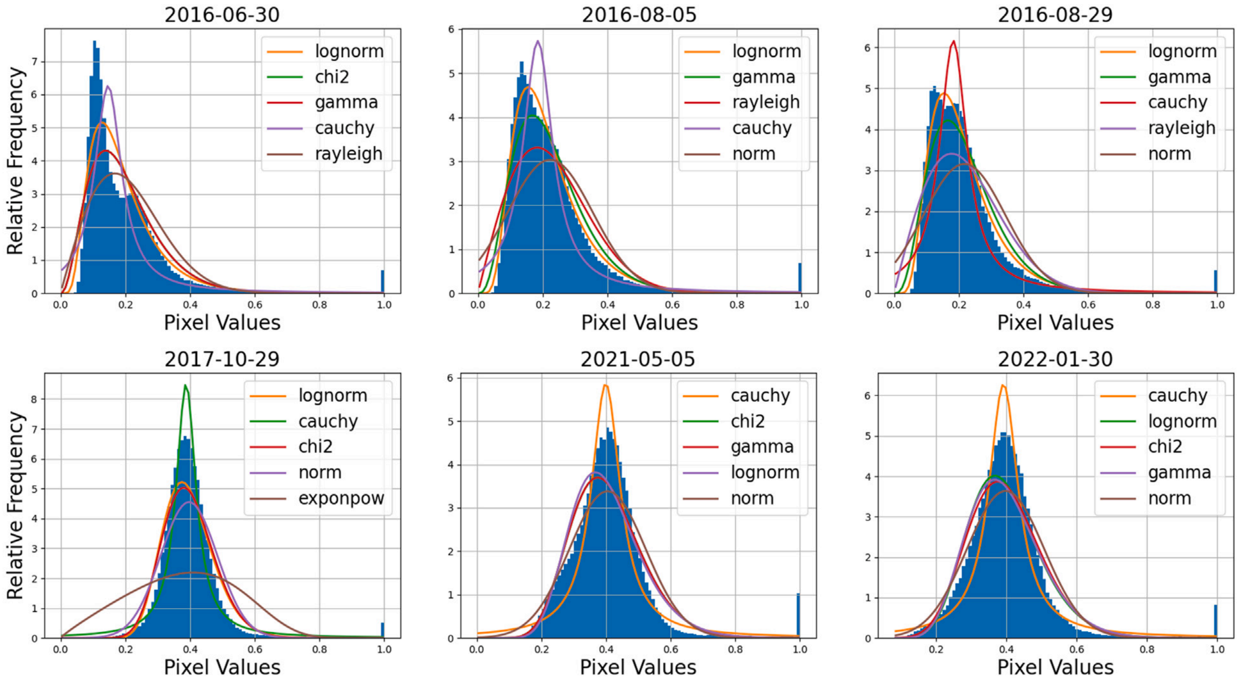

Fig. 3. Comparison of distributions of pixel values between images of dates. Top: distributions of image pixel values during ordinary weather. Bottom: distributions of images taken near major storm events.

3. Methods

In the following, we will briefly introduce our main methods, the autoencoder (AE), local outlier factor (LOF), and Reed-Xiaoli (RX) and describe our proposed application and evaluation procedure. In particular, we will present our novel patch-based extension approach that allows us to use LOF for anomaly and hotspot detection. Additionally we present a temporal variant of the RX algorithm and its adaptation for the patch-based approach. In our work, all of the methods are implemented in PyTorch and run on a workstation with Nvidia CUDA support.

3.1. Autoencoder

An autoencoder (AE) (Vincent et al., 2008), is a convolutional neural network (Krizhevsky et al., 2017) consisting of two parts, the encoder and the decoder. Its architecture is shown in Fig. 4. The encoder (first part) learns to reduce the data to a lower dimensional space. The lowest representation is the bottleneck or latent space of the AE, where the extraction of the most prominent features are expected. Then, the decoder (second part) learns to reconstruct the original data based on the latent space. The reconstruction is compared to the original image by a loss function, which is minimized. As a loss function we use the mean squared error (MSE).

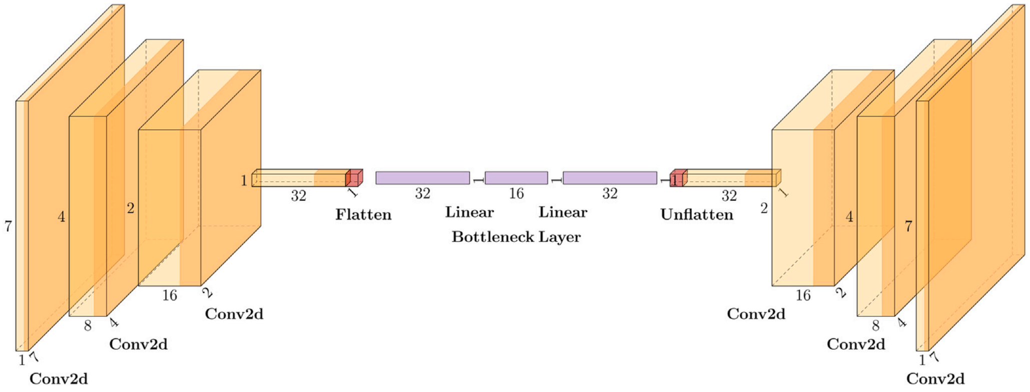

Fig. 4. Diagram of the autoencoder architecture used here. In the autoencoder, three 2D convolutions are followed by flattening and a linear transformation resulting in the latent space. The decoder is structured vice versa.

This is given by

$$\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}(I_i - R_i)^2, \tag{1}$$

with I being the image vector and R the reconstruction vector. The architecture of our AE is shown in Fig. 4. Starting from our patches in the encoder a 2D Convolution (Conv2D) is performed. The convolution is followed by a flattening of the planes and a fully connected affine linear transformation (Linear) resulting in the bottleneck with a target dimension of 16. For the decoder, the described sequence is reversed, resulting in a tensor with a dimension equal to the input. As activation functions we apply sigmoid for the last layer and rectified linear unit (ReLU) function otherwise.

The advantage of using the AE for anomaly detection, is that the trained AE will reconstruct anomalous data worse than normal data, assuming that abnormal data does not or just rarely occur in the training data. Thus, by comparing the original image to its reconstruction, reconstruction errors above a certain threshold are defined as anomalies. For thresholding we use Otsu's method (Otsu, 1979). It estimates an ideal threshold by maximizing the inter-class variance of the pixel value. Here, Otsu's threshold is estimated based on all considered images, i.e. over all dates, in order to make the results comparable.

3.2. Local outlier factor

The local outlier factor (LOF) is a density-based algorithm used for unsupervised outlier detection. Its main advantage is the ability to detect outliers whose distances to other points in the dataset are relatively small and would not be detected by other algorithms. This is achieved by comparing the local densities of the neighboring points rather than their distances.

The basic concept of LOF as described by Breunig et al. (2000) involve the k-distance of an object and the *reachability distance*. Let $d(a,b)$ represent a distance measure, with a and b being images in the dataset. The k-distance(a) is defined as the distance from a to its kth nearest neighbor. The complete set of the $k \in \mathbb{N}$ nearest neighbors includes all images within a distance $\leq$ k-distance(a) and is denoted as $N_k(a)$. It is important to note that the cardinality of $N_k(a)$, indicated by $\operatorname{card}\left(N_k(a)\right)$, can exceed k if multiple images have the same distance from a. Subsequently, the reachability distance from a to b is defined as

$$\mathrm{reach} - \mathrm{dist}(a,b) \;=\; \max\left\{k - \mathrm{distance}(b), d(a,b)\right\}. \tag{2}$$

Eq. (2) does not define a mathematical distance because it is not symmetric. The purpose of Eq. (2) is to establish a dissimilarity measure that is mostly resistant to statistical fluctuations, which occur in very similar images with small distances. In these cases, the actual distance is substituted with k-distance(b). The hyperparameter k determines the smoothing effect; for $k = 1$, the standard (unsmoothed) distance measure $d(\cdot,\cdot)$ is obtained. The *local reachability distance* (lrd) of an image a is defined as

$$\mathrm{lrd}_k(a) = \frac{\operatorname{card}\left(N_k(a)\right)}{\sum_{b\in N_k(a)} \mathrm{reach} - \mathrm{dist}(a,b)}. \tag{3}$$

Eq. (3) specifies the inverse of the average reachability distance of image a. Notably, if image a has at least k duplicates (meaning the reach-dist of all duplicates is zero), Eq. (3) approaches infinity. This can happen in our application when there is no change between adjacent images in a time series. Consequently, these non-informative images have been excluded from the dataset. As a result, the LOF of a is defined as

$$\mathrm{LOF}_k(a) \;=\; \frac{\sum_{b\in N_k(a)} \mathrm{lrd}_k(b)}{\operatorname{card}\left(N_k(a)\right)\cdot \mathrm{lrd}_k(a)}. \tag{4}$$

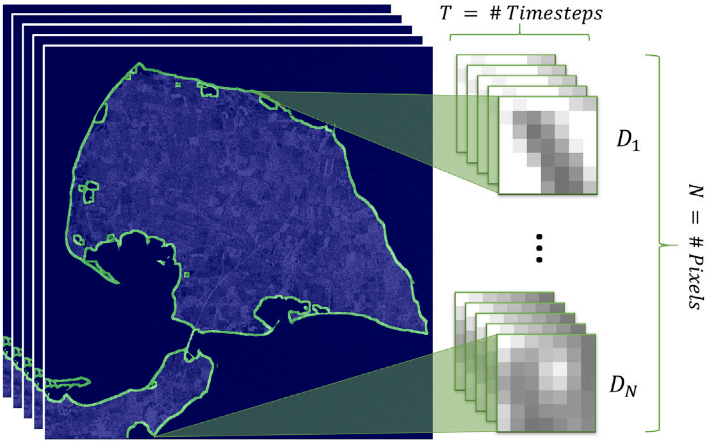

Fig. 5. SAR images of the island of Fehmarn on 2020-01-06. The image is overlayed with our coastline mask. The green line depicts the coastline, whereas the blue region is the irrelevant data. On the right side of the figure, the preprocessing of the patches is shown. We create N separate datasets $D_{i,i\in[1,N]}$, where each dataset corresponds to a single pixel of the coastline mask and contains T patches.

LOF determines the extent to which a is an outlier by comparing the local reachability distances of its neighbors. Three different scenarios can arise in Eq. (4):

$$\mathrm{LOF}_k(a) = \begin{cases} < 1, & \text{if density is higher than} \\ \approx 1, & \text{if density is similar to} \\ > 1, & \text{if density is lower than} \end{cases}$$

the neighbor density. Consequently, if $\mathrm{LOF}_k(a) > 1$, a is identified as a potential outlier. Conversely, $\mathrm{LOF}_k(a) \leq 1$ signifies potential inliers.

3.3. Novel patch-based anomaly and hotspot detection

The classic LOF is a well-known and proven method for finding anomalous images within a time series. However, our goal is to find the exact location of an anomaly within the image. Therefore, we propose to apply an ensemble of LOF algorithms to a series of images, by splitting each image into small subareas (patches) and applying the LOF algorithm to each subarea independently. The result is a patch-based extension of the LOF. To generate the patches, we chose an approach similar to Kalinicheva et al. (2019) and Koslow et al. (2024) for an AE-utilization. For each pixel i that is present in the generated coastline mask that was described in Section 2.2, we extracted a $p \times p$ patch where i is in the center of the patch. For our initial experiments, we chose $p = 3$, so that all patches are small enough to avoid the inclusion of non-relevant data from the sea and mainland, while still preserving spatial information. Thus, our full dataset consists of $T \times N$ image patches (samples) where T is the number of images at different dates and N is the number of pixels within our coastline mask. Each of the N subsets $D_{i,\ i\in\{1,\dots,N\}}$ corresponds to a single pixel within the coastline mask. An illustration of the subsets can be seen in Fig. 5. We assume that for one patch the variability of the pixel values in time are recurring rather than constant; e.g. they might be different for different weather condition (wet and dry) or tide-dependent. We anticipate that the typical pixel values form dense point-clouds, whereas anomalies are located outside these clouds. Therefore, we expect local density methods to be beneficial here. By creating a subset for each pixel in the image the LOF ensemble is able to locate anomalies on a scale corresponding to the resolution of a single pixel. The full flow of the method can be seen in Fig. 6.

For each subset $D_{i,\ i\in\{1,\dots,N\}}$, where N is the number of relevant pixels in each image $I_t, t \in [1,\ T]$, we apply N independent LOF algorithms. A single subset D_i consists of a time series of patches of size $p \times p$ that were extracted around pixel i. Each LOF instance produces anomaly scores S_i for each patch within its subset D_i. If an anomaly score $s_t \in S_i$ exceeds a certain threshold, s_t is considered to be an anomaly within S_i. These scores are then used to create anomaly maps by reconstructing the original images that have been divided into subsets and inserting the anomaly scores into the corresponding pixels.

In order to have a fair comparison between the AE, LOF, and RX we also set up a patch-wise approach for the AE. This means, in detail, that the AE is trained to reproduce the same $p \times p$ patches. During evaluation, the original values and the reconstruction are compared for every patch. However, since the AE generates reconstruction errors for each pixel in every patch, the reconstruction errors of each pixel are averaged over pixel value of each patch that this pixel is present in. The full pipeline can be seen in Algorithm 1 along with the creation of coastal hotspot maps, which will be discussed later.

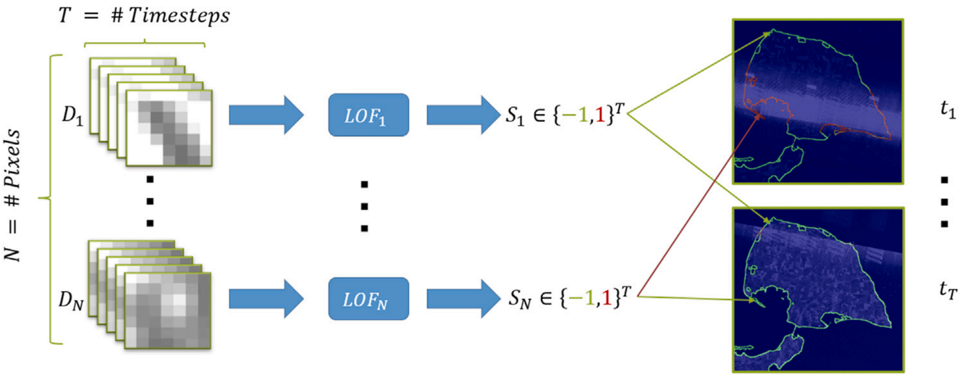

Fig. 6. Illustration of the novel patch-based method: For each subset $D_{i,\ i\in\{1,\dots,N\}}$, the LOF algorithm is applied separately. Each LOF produces a classification array $S_i \in \{-1,1\}^T$, where -1 indicates an inlier (green) and 1 indicates an outlier (red). The results of all LOF algorithms are then combined to produce T anomaly maps. The Fehmarn images on the right show the result of the LOF algorithm for two images that both show radio frequency interference (RFI) that resembles noise in large parts of the images. The LOF is able to precisely locate these anomalies in both images.

Algorithm 1 Full patch-based LOF pipeline

```
Input: Time series of SAR images $I_t$
Output: score_map, anomaly_map, hotspot_map
Generate average water mask
Generate coastline mask with width W by masking a W × W neighborhood around each water pixel bordering a non-water pixel
for each pixel $i$ in coastline mask do
    for each $t$ in $0, ..., T$ do
        Extract $p \times p$ patch $x_{t,i}$ around i
    end for
    Receive timeseries subset $D_i = \{x_{0,i}, ...x_{T,i}\}$
end for
Initialize:
    score_map ← negative ones of same size as $I_t$
    anomaly_map ← negative ones of same size as $I_t$
for each pixel $i$ in coastline mask do
    Apply LOF on $D_i$ and receive scores $S_i \in \mathrm{R}^T$
    for each $t$ in $0, ..., T$ do
        Update the score map: score_map$[i, t] \leftarrow S_i(t)$
        if $S_i(t) >$ threshold then
            Update the anomaly map: anomaly_map$[i, t] \leftarrow 1$
        else
            Update the anomaly map: anomaly_map$[i, t] \leftarrow 0$
        end if
    end for
end for
Initialize:
    hotspot_map ← negative ones of same size as $I_t$
for each pixel $i$ in coastline mask do
    Select top $K$ values of $S_i$
    Update the hotspot map: hotspot_map$[i] \leftarrow$ mean of top $K$ values of $S_i$
end for
End
```

3.4. Mahalanobis distance

An important consideration when using the LOF is the choice of distance measure. The authors of the original paper propose using the Euclidean distance. However, depending on the data that the LOF is applied on, other distance measures might be more suitable, e.g. the structural similarity index measure used on video data of hybrid rocket combustion test in Rüttgers and Petrarolo (2021). Since we are using small patches, this measure is not applicable here. Another distance measure is the Mahalanobis distance (Mahalanobis, 2018), which is also used in the RX algorithm. Let $a, b \in \mathbb{R}^d$ be two feature vectors and let $\boldsymbol{\Sigma} \in \mathbb{R}^{d\times d}$ denote the covariance matrix of the data. The Mahalanobis distance between a and b is then defined as

$$d(a, b) = \sqrt{(a-b)^\top \boldsymbol{\Sigma}^{-1} (a-b)}. \tag{5}$$

In our case, a and b are the flattened patches in $\mathbb{R}^{p^2}$ and $\boldsymbol{\Sigma} \in \mathbb{R}^{p^2\times p^2}$ is the covariance matrix between all patches of a subset $D_{i,\ i\in\{1,\dots,N\}}$. Unlike the Euclidean distance, which assumes that features are independent and have equal variance, the Mahalanobis

Table 1
Table of the fraction of detected artifacts for three anomaly detection (ad-model) models, for three synthetic anomaly shapes. For every case the highest value is bold.

ad-model / synthetic anomaly type	AE			LOF			RX		
	7x7	5x5	3x3	7x7	5x5	3x3	7x7	5x5	3x3
Annulus	0.49	0.56	0.03	**1.0**	**1.0**	0.99	**1.0**	**1.0**	0.99
Area	0.46	0.52	0.05	**1.0**	**1.0**	0.96	**1.0**	**1.0**	0.93
Single	0.83	0.84	0.19	0.47	0.81	**0.94**	0.77	0.81	0.87

distance accounts for correlations and varying variances between features. It measures distance relative to the underlying data distribution, making it particularly suitable for data with correlated feature structures, such as the time series of a neighborhood around a coastal pixel in a SAR image.

3.5. Temporal Reed–Xiaoli

The RX algorithm (Reed and Yu, 1990) is a traditional approach to anomaly detection in SAR imagery. It is a statistical and signal-based detector. Originally, the RX, algorithm computes a score for each pixel in a SAR image, by computing the Mahalanobis distance between the pixel and the mean of the background:

$$S(p) = (p - \mu)^{\top} \boldsymbol{\Sigma}^{-1} (p - \mu), \tag{6}$$

where μ is the mean vector of the background (typically computed from the neighboring pixels or from the entire image), and $\boldsymbol{\Sigma}$ is the covariance matrix of the background, usually estimated from the rest of the pixels in the image. A high score indicates that the pixel is far from the expected background distribution, marking it as an anomaly. Since we are interested in anomalies within a timeseries in this work, we define the background as timeseries of patches that were derived above so that

$$\mu_i = \frac{1}{T} \sum_{j=1}^{T} x_{ij}, \qquad \text{for } x_i \in D_{i,\ i \in \{1,\ldots,N\}} \tag{7}$$

$$\boldsymbol{\Sigma}_i = \frac{1}{T-1} \sum_{j=1}^{T} (x_j - \mu_i)(x_j - \mu_i)^{\top} \qquad \text{for } x_i \in D_{i,\ i \in \{1,\ldots,N\}} \tag{8}$$

where $D_{i,\ i \in \{1,\ldots,N\}}$ are the patch subsets described above. This approach allows us to create comparable results to our patch-based LOF pipeline.

4. Results

For the analysis of the results, we focus on the Cuxhaven region, shown in Fig. 8. The LOF and RX are applied to the full time series (2016 to 2023), while the AE is trained on the year 2016 and the remaining data is used for evaluation. After analyzing the loss values for the validation data for a training with different epochs, we chose 20 epochs for the training of the AE. For the k-neighbors parameter, we chose a value of 10 after some testing, which seemed to perform the best.

4.1. Synthetic data

In order to better quantify the performance of our models, we evaluate our algorithms on datasets with synthetic anomalies that we introduced in Section 2.4. Table 1 compares the accuracies of the detected anomalies, given by

$$\mathrm{acc} = \frac{\#\ \text{detected anomalous pixels}}{\#\ \text{total anomalous pixels}}. \tag{9}$$

where anomalous pixels are the artificially changed pixels. It shows the mean accuracies over all timesteps containing synthetic anomalies. We deliberately refrain from analyses based on false positives as the precision or metrics derived from them. Since we integrated the artificial anomalies into the real images, further detected anomalies might refer to real-world anomalies and are not an indication that the method is detecting something incorrectly. We alter 6 images within the dataset and also vary the size of the patches. The threshold for the anomaly scores were chosen so that the top 5% would be classified as anomalies to make the results comparable between the three algorithms. In all cases, we replace the values of the selected pixels by sampling from a lognormal distribution with shape parameter 0.19, mean -0.03 and standard deviation 0.42, which is the distribution that best fit the values during the storm on the 2017-10-29, compare Fig. 3.

We observe, that the accuracies of the AE remain consistently below 0.2. For the LOF and RX, however, the values are above 0.94 and 0.87 respectively for 3 × 3 patches. This illustrates the importance of choosing the correct patch size. If we increase the patch size to 5 × 5 and 7 × 7, the accuracies for areas and the annulus increases to a detection rate of 100%. This is plausible, since there are more anomalous pixels inside a patch for connected anomaly types if the patch size increases and is thus easier

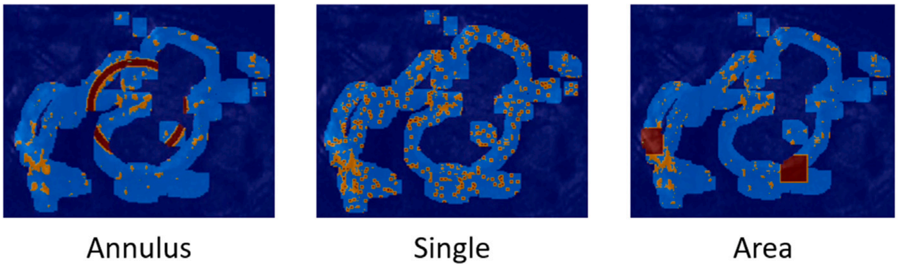

Fig. 7. Anomaly detection maps created with LOF for different types of synthetic anomalies. Red pixels represent correctly detected synthetic anomalies (true positive), green pixels represent missed anomalies (false negative), and orange pixels represent detected anomalies for unaltered pixels.

to identify. However, the single anomaly type shows a significant decrease in accuracy compared to a 3 × 3 patch size. Since the accuracies for a 3 × 3 patch size are already quite high for the area and annulus, this patch size seems to be more appropriate for most applications. Additionally, we can see that the LOF outperforms the RX, especially on single anomaly types. However, the RX itself performs quite well when compared to the AE.

For the sake of completeness, we used the LOF and RX on 1 × 1 patches, which means that each timestep consists of only a single value. The resulting accuracies were approximately constant for all types of anomalies at 0.32 and 0.57 for the LOF and RX respectively. This result is worse than even the performance on the single synthetic anomaly type on 7 × 7 patches and it confirms the importance of using small patches in order to incorporate spatial context into the data. Fig. 7 shows a small area of the anomaly maps generated by the LOF algorithm with 3 × 3 patches. It is important to note that although the synthetic anomalous pixels are detected quite accurately, indicated by the large number of red pixels, there are also many orange areas around them, indicating detected anomalies that are not synthetically modified. This is because the patches that are close enough to a synthetic anomaly pixel are affected by it. Using a size of 3 × 3, all patches within a square radius of 1 pixel will also include the center pixel. This results in the orange spots around the red areas in Fig. 7. Therefore, it is important to note that a found anomaly does not necessarily mean that the exact pixel is anomalous, but rather that there is an anomaly in a $p \times p$ area, which further encourages the use of 3 × 3 patches in order to make the predicted anomalies more precise.

4.2. Real-world anomaly detection

Just because a method is able to successfully detect synthetic anomalies does not imply that this method is also suitable for real-world applications. Therefore, we check here how well our methods detect known severe weather events and their effects on the landscape.

In Fig. 8 the resulting anomaly maps of three different dates are shown. The SAR image of 2019-10-31 (left column) contains an artifact in the top right corner, which is shown in Fig. 9. The artifact can be ascribed to radio frequency interference during the acquisition of the SAR image, similar to the image artifacts that can be seen in Fig. 6. We can observe that the AE (bottom row) localizes less pixels that contain this artifact compared to the other two algorithms. On the 2022-01-30 (right column) the storm front “Malik” caused major flooding in the region, leading to the detection of a large number of anomalies on that date by the LOF and RX but not the AE. Comparing the LOF and RX results, it is worth noting, that the anomalies found by the LOF are more sparse and seem to be more precise then those of the RX when comparing the storm images.

However, the cases analyzed so far were only exemplary. To obtain a more general understanding of the capability of the methods, we compare the number of anomalies of each date with certain weather variables during that time. Specifically we take a look at the average daily wind speed, as this is likely to have the greatest impact on the environment. Fig. 10 shows the number of anomalies as bars versus the average wind speed as points. We have omitted some values in the middle of the time series, where no interesting findings could be observed in order to improve the visibility of the plots. We also include wind speeds that exceed $30\,\mathrm{km/h}$, as lower speeds should have no impact on the environment. Three dates show a large number of anomalies and these dates correspond to the severe storm events mentioned in Section 2.3. Thus, we observe that both the LOF and RX can clearly distinguish these dates as anomalies and, as seen in Fig. 8, are also able to localize the specific areas which were hit by the storm. There is one more date, the 2018-03-04, where a lot of anomalies were detected by the LOF and RX, which does not have a particularly high wind speed. After further investigation, the date in question exhibited the lowest temperatures in the entire time series, which may have led to the formation of icy surfaces and, consequently, anomalous features in the SAR imagery.

Again, the AE performs worse. Fig. 10 shows the number of anomalies per time point and the daily mean reconstruction error. The three dates closest to the known storm events are marked in orange. Neither the number of anomalies nor the height of the values seem to be related to storm events. If we examine the top ten dates for the number of anomalies or the top ten values of the mean reconstruction error, we see that only the date 2021-05-05 belongs to the top ten for the mean reconstruction error, specifically it is placed at position eight, but not for the other dates or for the number of anomalies. Thus, the AE hardly recognize these dates.

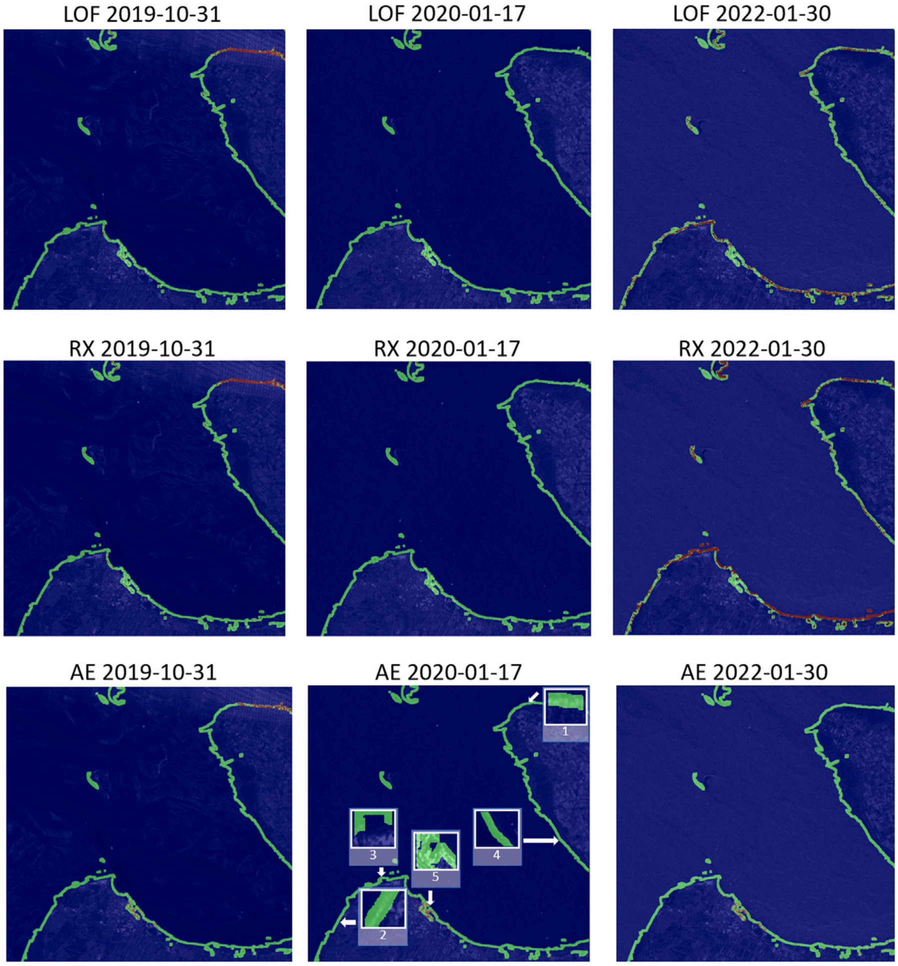

Fig. 8. Anomaly maps for three dates of LOF (top), RX (middle), and AE (bottom). Additionally, the center AE image shows the five snippet regions that are analyzed below.

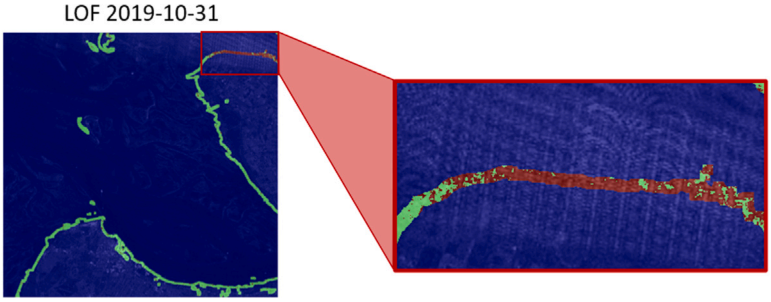

Fig. 9. Anomaly map for LOF on 2019-10-31 with enlarged section of the radio frequency interference artifact.

More generally, we consider all weather specific variables, by calculating the Pearson correlation coefficients (Pearson, 1896) between the number of anomalies and the weather variable values (see Fig. 11). As a rule, the higher the absolute value of the Pearson correlation coefficient, the higher the correlation. In support of our previous findings, we observe a significant correlation between the anomalies and the wind speed parameters for the LOF and RX. However, we observe a negligible correlation between the weather variables and the daily anomaly count produced by our algorithm. This finding is consistent with the fact that SAR images are invariant to weather changes. The only other notable correlation is a slight negative correlation with the temperature, indicating a higher number of anomalies during colder seasons, which is also when most storms are observed. Furthermore, it is

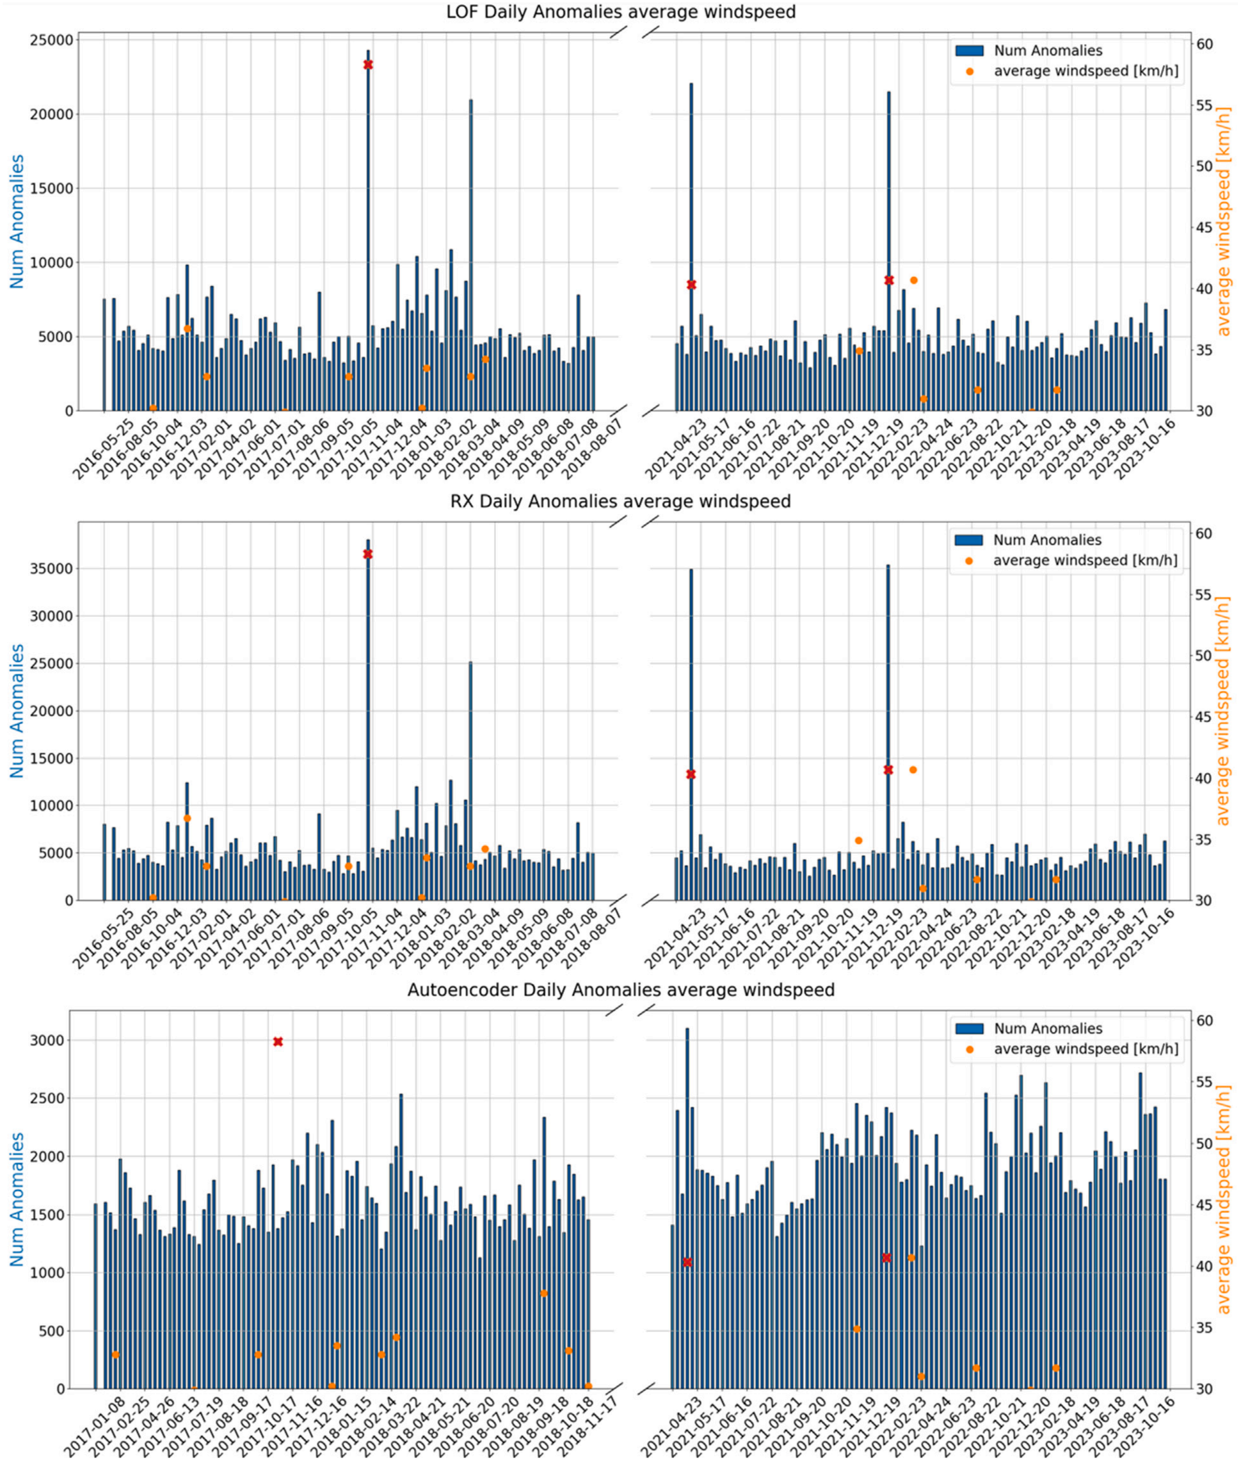

Fig. 10. Number of anomalies found by each algorithm (blue bars) vs. average wind speed (orange dots). The wind speed values are only shown if they are larger then 30 km/h. The red crosses represent the three storm events that were introduced in Section 2.3.

important to note that a higher correlation does not necessarily indicate better performance. A Pearson coefficient of 1 would mean that the model only detects strong wind speeds and ignores everything else. This is not what is expected from a good anomaly detection algorithm. For the AE, no correlation is observed for any of the weather parameters.

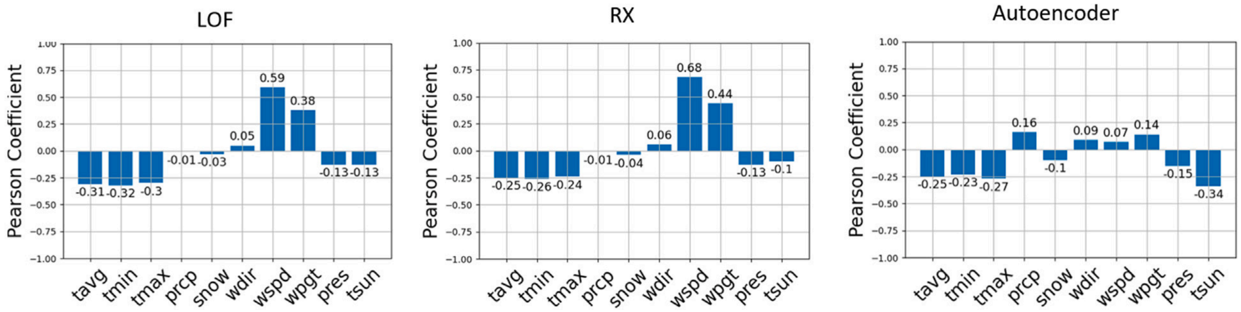

Fig. 11. Pearson coefficient between the daily number of anomalies found by each algorithm and different weather variables. Average wind speed (wdsp) and maximum wind speed (wpgt) were adjusted, so that only wind speed values of $> 30\,\mathrm{km/h}$ or $> 50\,\mathrm{km/h}$, respectively were considered.

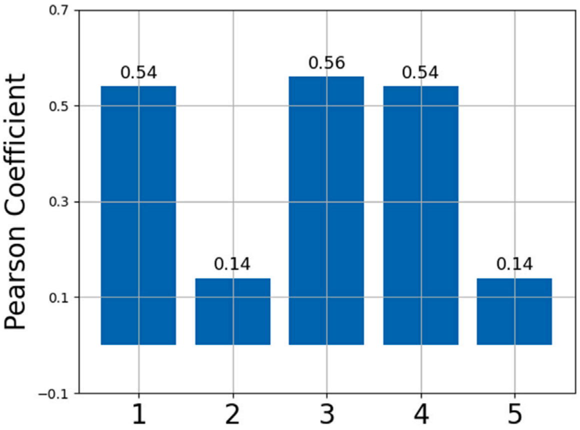

Fig. 12. Pearson coefficients regarding the number of anomalies and the average wind speed (wdsp) for snippets of the Cuxhaven region. Compared to the larger images, the Pearson coefficient has increased, but there is a large variability.

5. Discussion

Our investigations into anomaly and hotspot detection show that for our novel patch-wise extension approach the LOF and RX successfully identify synthetic anomalies, artifacts, and especially coastal changes due to severe weather. In contrast to this, the AE performs worse. The blindness of the AE for the different types of anomalies may be due to the fact that the AE is trained to reconstruct the patches. However, it is only one AE and thus all the patches together define the model parameters. It is possible that the variability of the pixel values is very high on one patch and lower on others. Anomalies will then be missed on the latter. In order to analyze this in more detail, we investigate the performance of the AE for smaller segments, called snippets, cut out of the investigated images. The AE for each snippet was trained with the same hyperparameters as for the larger ones. For the snippets, the storm events are much better recognized. Fig. 12 shows the correlation between the average wind speed to the number of anomalies in the snippets. Thus, for snippets the performance of the AE is better than for the whole image and for three of the five cases even comparable to the LOF and RX. Since we expect a smaller variability of pixel values in the snippets, this supports our assumption.

By decreasing the size of the snippets to the chosen patch size, it becomes equivalent to training an AE for every patch. Although this resulted in better performance compared to the training of only one AE for all patches, it was still worse than the results of the LOF and RX. Additionally, the training time for this case was 72 h and the inference time was also much higher due to the overhead of creating and evaluating a separate model for each pixel in the dataset. Therefore, it is not feasible for future applications, such as for continuous monitoring. Likewise unfeasible is to apply the AE on the full image without patches, because this leads to problems with computer memory. Another approach is to split the images into tiles without overlapping them leading to much larger and fewer patches, however this also did not improve the results. We tested the performance for a relatively small region around Helgoland (125 × 188 pixels) with synthetic anomalies. We found that the detection rate slightly increases, however so does the number of obvious false positives. Hence, we conclude that the AE approach does not suit the complexity and variability of our SAR data to accurately detect meaningful anomalies.

However, we are aware that this is a simple autoencoder architecture and we cannot rule out the possibility that better results can be achieved with modifications to it, as variational or adversarial autoencoders. However, first tests with a variational autoencoder, did lead to worse results. Since the main modification in both cases is a more variable latent space, this is not surprising for a

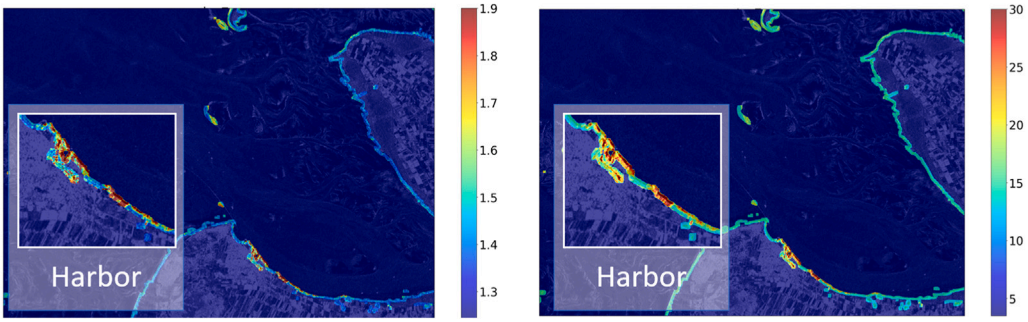

Fig. 13. Coastal hotspot maps obtained as average of the top 40 outlier score maps over time for the LOF (left) and the RX (right). The colorbars represent the average outlier scores for each pixel.

training of all patches simultaneously. Nevertheless, for an autoencoder per patch employing these methods could be advantageous, but also requires large computing resources, as discussed elsewhere for our case. Classical methods on the other hand do not require a large training data set, but only enough neighboring points to have a sufficient number of neighbors for each patch to be able to make a statement about the local density.

We conclude that the LOF and RX are better suited for this application. However, the results with synthetic anomalies in Table 1 show, that the LOF outperforms the RX, especially on single pixel anomalies.

Whereas, for real-world applications, the performance seems to be comparable. One difference becomes evident, by analyzing the coastal hotspot maps shown in Fig. 13. These maps highlight areas that exhibit recurring or gradually intensifying anomalies. By incorporating background information, an informed observer can determine whether the observed features represent anticipated hotspots, such as those seen at the harbor or constitute more unusual activities. The latter could indicate a risk of infrastructure damage due to more frequent flooding or even coastal erosion. The left side of Fig. 13 shows the averaged anomaly score map of the top LOF 40 anomaly scores of each pixel subset, while the right side shows the averaged anomaly scores of the RX. For better contrast, the maximum value of the color bar of the averaged scores of the LOF and RX was limited to 1.9 and 30, respectively. A closer look at the images shows that the harbor area, which has been magnified in the image, clearly shows a higher frequency of anomalies for both algorithms. This could be explained by the higher activity that is present at a harbor and therefore a higher number of anomalies in these areas. Although both algorithms highlight similar areas, the LOF shows a higher contrast between areas of interest, while the RX seems to be more uniform, even though several different values for the upper threshold of the colorbar were tried, of which 30 seemed to show the best contrast. This might lead to some areas of interest being overlooked. However, more data is needed to validate this approach with confidence. One limitation of our work is the low resolution of our data, as well as the lack of ground truth real-world anomalies that we have access to. Although the storm events identified by our algorithms provide a credible validation, more subtle anomalies like coastal erosion are impossible to detect with an image resolution of 20 meters per pixel. Thus, better data could greatly improve the validity of the results and increase the amount of information that can be gained from them. On the one hand, the clear detection of the harbor suggests that we are detecting anomalies that deviate from the regularity, which is exactly what we wanted to show. On the other hand, one could argue that there should be a differentiation between dangerous anomalies and anomalies caused by human activity. However, our aim is to provide a tool for the detection of anomalies of all kinds. Because the appearance of significant coastal changes are unpredictable, the occurrence of obvious hotspots must be accepted. By simply linking to other map material, such as the ESA landcover map,[1] it is easily possible to draw further conclusions. Although this work focuses on specific coastal regions, we expect that the proposed patch-based methods are applicable to all coastal regions because they rely solely on fundamental, location-independent principles of anomaly detection. More precisely, our methods operate on statistical and structural patterns in time for small patches without taking predefined geographic features into account. Therefore, we hypothesize that these methods are not limited to coastal regions, but rather can be applied to all types of Earth observation data, e.g. vegetation, inland waters or urban areas. This highlights the relevance of our work for a wide range of applications.

6. Conclusion

In summary, we demonstrated our novel patch-based extensions of LOF for unsupervised anomaly detection in SAR amplitude data, enabling the spatial identification of anomalies. Additionally we created coastal hotspot maps which highlight areas with frequent anomalous activities, thereby indicating potentially elevated risk in these regions. Although our study focused on the North and Baltic Sea, we presume this method can be applied to other coastal and potentially non-coastal regions, due to the unsupervised and non-specialized nature of our method. Different types of synthetic anomalies, varying in size and shape, as well as image artifacts were clearly recognized.

---

[1] https://esa-worldcover.org

We compared the LOF to a patch-wise AE implementation and a temporal adaptation of the RX algorithm which is foundational in the area of earth observation. We showed, that the patch-wise AE implementation achieved the lowest overall performance and even the RX was outperformed by our LOF approach on all considered types of artificial anomalies. Some regions that experience more frequent activities, as the harbor, and in particular changes due to storm events were identified as real-world anomalies. The latter were proven by a clear correlation between the number of anomalies and wind speed, indicated by a high Pearson correlation coefficient with concurrent low correlation for other weather variables. Despite the encouraging performance of the proposed approach, further investigation is required in order to validate it on higher-resolution data, where more subtle anomalies, such as coastal erosion, can be present. Preliminary testing using artificial anomalies that model coastal erosion yielded promising results: up to 25% higher accuracy with the LOF compared to the RX. However, further investigation is needed to ensure that the artificial anomalies closely follow a real-world erosion process. Furthermore, the methods developed in this work focus on static anomaly detection, and therefore may not fully capture temporal anomalies or gradual changes. To address this limitation, we propose future extensions based on additional feature engineering that explicitly model temporal trends in the data. While such approaches are beyond the scope of this study and would require more extensive validation data, our results suggest that this represents a promising direction for future research.

Overall, our new patch-based LOF implementations show strong potential as tools for monitoring and protecting coastal infrastructure, supporting decision-makers, and mitigating the effects of climate change.

CRediT authorship contribution statement

Wadim Koslow: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation. **Kathrin Rack:** Writing – original draft, Visualization, Validation, Software, Investigation. **Tobias D. Grabosch:** Validation, Software, Investigation. **Alexander Rüttgers:** Writing – original draft, Supervision, Project administration. **Luca Dell'Amore:** Writing – review & editing, Data curation. **Paola Rizzoli:** Supervision, Project administration.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used DeepL in order to improve the readability of this article. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was carried out under the project *Resiliente Versorgungsinfrastruktur und Warenströme im Kontext küstennaher Extremwetterereignisse* (RESIKOAST) by the German Aerospace Center (DLR).

Data availability

The input SAR data is made publicly available at Koslow (2025). The code is not publicly available.

References

Auskalnis, J., Paulauskas, N., Baskys, A., 2018. Application of local outlier factor algorithm to detect anomalies in computer network. Elektron. Elektrotechnika 24 (3), 96–99. http://dx.doi.org/10.5755/j01.eie.24.3.20972.

Bouman, R., Bukhsh, Z., Heskes, T., 2024. Unsupervised anomaly detection algorithms on real-world data: How many do we need? J. Mach. Learn. Res. 25 (105), 1–34. http://dx.doi.org/10.48550/arXiv.2305.00735.

Breunig, M., Kriegel, H.P., Ng, R., Sander, J., 2000. LOF: identifying density-based local outliers. In: Proc. 2000 ACM SIGMOD Int. Conf. Manag. Data. pp. 93–104. http://dx.doi.org/10.1145/335191.335388.

Chen, Z., Gao, H., Lu, Z., Zhang, Y., Ding, Y., Li, X., Zhang, B., 2025. MDA-HTD: Mask-driven dual autoencoders meet hyperspectral target detection. Inf. Proc. Manag. 62 (4), 104106. http://dx.doi.org/10.1016/j.ipm.2025.104106, URL: https://www.sciencedirect.com/science/article/pii/S0306457325000482.

Chen, Z., Yeo, C., Lee, B., Lau, C., 2018. Autoencoder-based network anomaly detection. In: Proc. 2018 Wirel. Telecommun. Symp. WTS, pp. 1–5. http://dx.doi.org/10.1109/WTS.2018.8363930.

Cui, Z., Wang, X., Liu, N., Cao, Z., Yang, J., 2021. Ship detection in large-scale SAR images via spatial shuffle-group enhance attention. IEEE Tran. Geosci. Remote. Sens. 59 (1), 379–391. http://dx.doi.org/10.1109/TGRS.2020.2997200.

Flores, A., Herndon, K., Thapa, R., Cherrington, E., 2019. The SAR Handbook: Comprehensive Methodologies for Forest Monitoring and Biomass Estimation. http://dx.doi.org/10.25966/nr2c-s697.

Fung, C., Qiu, C., Li, A., Rudolph, M., 2025. Model selection of anomaly detectors in the absence of labeled validation data. IEEE Trans. Artif. Intell. 1–10. http://dx.doi.org/10.1109/TAI.2025.3562505.

Gao, Z., 2009. Application of cluster-based local outlier factor algorithm in anti-money laundering. In: 2009 Int. Conf. Manag. Serv. Sci. pp. 1–4. http://dx.doi.org/10.1109/ICMSS.2009.5302396.

Intergovernmental Panel on Climate Change, 2023. Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge University Press, http://dx.doi.org/10.1017/9781009157896.

Kalinicheva, E., Sublime, J., Trocan, M., 2019. Change detection in satellite images using reconstruction errors of joint autoencoders. In: Lect. Notes Comp. Sci. Springer International Publishing, pp. 637–648. http://dx.doi.org/10.1007/978-3-030-30508-6_50.

Kim, H., Lee, C., 2024. Enhancing anomaly detection via generating diversified and hard-to-distinguish synthetic anomalies. In: Proc. 33rd ACM Int. Conf. Inf. Knowl. Manag. CIKM '24, Assoc. for Comput. Mach. New York, NY, USA, pp. 1089–1098. http://dx.doi.org/10.1145/3627673.3679623.

Koslow, W., 2025. SAR images over the baltic and north sea [Dataset]. http://dx.doi.org/10.5281/zenodo.15686791.

Koslow, W., Rack, K., Ruettgers, A., Dell'Amore, L., Rizzoli, P., 2024. Artifact detection in SAR images with AI methods. In: EUSAR 2024; 15th Eur. Conf. Synth. Apert. Radar. VDE, pp. 463–468, URL: https://ieeexplore.ieee.org/abstract/document/10659585.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2017. ImageNet classification with deep convolutional neural networks. Commun. ACM 60 (6), 84–90. http://dx.doi.org/10.1145/3065386.

Kwon, H., Nasrabadi, N., 2005. Kernel RX-algorithm: a nonlinear anomaly detector for hyperspectral imagery. IEEE Trans. Geosci. Remote. Sens. 43 (2), 388–397. http://dx.doi.org/10.1109/TGRS.2004.841487.

Lamprecht, C., 2025. Meteostat. URL: https://meteostat.net/. (Accessed 20 June 2025).

Li, J., Zheng, K., Gao, L., Han, Z., Li, Z., Chanussot, J., 2025. Enhanced deep image prior for unsupervised hyperspectral image super-resolution. IEEE Trans. Geosci. Remote. Sens. http://dx.doi.org/10.1109/TGRS.2025.3531646.

Ma, H., Hu, Y., Shi, H., 2013. Fault detection and identification based on the neighborhood standardized local outlier factor method. Ind. Eng. Chem. Res. 52 (6), 2389–2402. http://dx.doi.org/10.1021/ie302042c.

Mahalanobis, P.C., 2018. Reprint of: Mahalanobis, P.C. (1936) "on the generalised distance in statistics." Sankhya A 80 (1), 1–7. http://dx.doi.org/10.1007/s13171-019-00164-5.

Muzeau, M., Ren, C., Angelliaume, S., Datcu, M., Ovarlez, J.-P., 2022. Self-supervised learning based anomaly detection in synthetic aperture radar imaging. IEEE Open J. Signal Process. 3, 440–449. http://dx.doi.org/10.1109/OJSP.2022.3229618.

Otsu, N., 1979. A threshold selection method from gray-level histograms. IEEE Trans. Syst. Man Cybern. 9 (1), 62–66. http://dx.doi.org/10.1109/TSMC.1979.4310076.

Pearson, K., 1896. VII. Mathematical contributions to the theory of evolution.—III. Regression, heredity, and panmixia. Philo. Trans. R. Soc. Lon. Ser. A 187, 253–318. http://dx.doi.org/10.1098/rsta.1896.0007.

Reed, I., Yu, X., 1990. Adaptive multiple-band CFAR detection of an optical pattern with unknown spectral distribution. IEEE Trans. Acoust. Speech Signal Process. 38 (10), 1760–1770. http://dx.doi.org/10.1109/29.60107.

Ronneberger, O., Fischer, P., Brox, T., 2015. U-Net: Convolutional networks for biomedical image segmentation. In: Med. Image Comp. and Comp.-Assist. Interv. MICCAI 2015, pp. 234–241. http://dx.doi.org/10.1007/978-3-319-24574-4_28.

Rüttgers, A., Petrarolo, A., 2021. Local anomaly detection in hybrid rocket combustion tests. Exp. Fluids 62 (7), 136. http://dx.doi.org/10.1007/s00348-021-03236-1.

Sinha, S., Giffard-Roisin, S., Karbou, F., Deschatres, M., Karas, A., Eckert, N., Coléou, C., Monteleoni, C., 2021. Variational autoencoder anomaly-detection of avalanche deposits in satellite SAR imagery. In: Proc. 10th Int. Conf. Clim. Inform. In: CI2020, Assoc. for Comput. Mach. New York, USA, pp. 113–119. http://dx.doi.org/10.1145/3429309.3429326.

Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A., 2008. Extracting and composing robust features with denoising autoencoders. In: Proc. of the 25th ICML. pp. 1096–1103. http://dx.doi.org/10.1145/1390156.1390294.

Wang, D., Gao, L., Qu, Y., Sun, X., Liao, W., 2023. Frequency-to-spectrum mapping GAN for semisupervised hyperspectral anomaly detection. CAAI Tran. Intell. Technol. 8 (4), 1258–1273. http://dx.doi.org/10.1049/cit2.12154.

Wang, D., Zhuang, L., Gao, L., Sun, X., Zhao, X., 2024. Global feature-injected blind-spot network for hyperspectral anomaly detection. IEEE Trans. Geosci. Remote. Sens. Lett. 21, 1–5. http://dx.doi.org/10.1109/LGRS.2024.3449635.

Wu, S., Zhang, X., Wang, G., Chen, P., Gu, J., Cheng, X., Jiao, L., 2025. VJDNet: A simple variational joint discrimination network for cross-image hyperspectral anomaly detection. Remote. Sens. 17 (14), 2438. http://dx.doi.org/10.3390/rs17142438.

Zakzouk, M., Abdulaziz, A., Abou El-Magd, I., Dahab, A., Ali, E., 2025. Automated oil spill detection using deep learning and SAR satellite data for the northern entrance of the Suez Canal. Sci. Rep. 15, http://dx.doi.org/10.1038/s41598-025-03028-1.