

---

# GLOBAL OFFSHORE WIND INFRASTRUCTURE: DEPLOYMENT AND OPERATIONAL DYNAMICS FROM DENSE SENTINEL-1 TIME SERIES

---

PREPRINT

 **Thorsten Hoerer**<sup>\*1</sup>,  **Felix Bachofer**<sup>1</sup>, and **Claudia Kuenzer**<sup>1,2</sup>

<sup>1</sup>Earth Observation Center (EOC), German Aerospace Center (DLR), Oberpfaffenhofen, 82234 Wessling, Germany

<sup>2</sup>Institute for Geography and Geology, University of Wuerzburg, 97074 Wuerzburg, Germany

April 23, 2026

## ABSTRACT

The offshore wind energy sector is expanding rapidly, increasing the need for independent, high-temporal-resolution monitoring of infrastructure deployment and operation at global scale. While Earth Observation based offshore wind infrastructure mapping has matured for spatial localization, existing open datasets lack temporally dense and semantically fine-grained information on construction and operational dynamics. We introduce a global Sentinel-1 synthetic aperture radar (SAR) time series data corpus that resolves deployment and operational phases of offshore wind infrastructure from 2016Q1 to 2025Q1. Building on an updated object detection workflow, we compile 15,606 time series at detected infrastructure locations, with overall 14,840,637 events as analysis-ready 1D SAR backscatter profiles, one profile per Sentinel-1 acquisition and location. To enable direct use and benchmarking, we release (i) the analysis ready 1D SAR profiles, (ii) event-level baseline semantic labels generated by a rule-based classifier, and (iii) an expert-annotated benchmark dataset of 553 time series with 328,657 event labels. The baseline classifier achieves a macro F1 score of 0.84 in event-wise evaluation and an area under the collapsed edit similarity-quality threshold curve (AUC) of 0.785, indicating temporal coherence. We demonstrate that the resulting corpus supports global-scale analyses of deployment dynamics, the identification of differences in regional deployment patterns, vessel interactions, and operational events, and provides a reference for developing and comparing time series classification methods for offshore wind infrastructure monitoring.

**Keywords** Earth Observation · Offshore Wind Energy · Offshore Infrastructure · Sentinel-1 · Time Series

## 1 Introduction

Offshore wind energy has experienced its most substantial growth phase within the past five years, see Figure 1. The year 2021 stands out as a record year, during which 3,418 offshore wind turbines (OWTs) were deployed, the highest annual number to date. For comparison, the median annual deployment between 2017 and 2024 was 1,100 units, with the second highest year being 2023 (1,269 units). Regional differences show that the surge in 2021 was driven by China, which accounted for 77.3% (2,642 units) of global deployments. The United Kingdom contributed 10.5% (359 units), and the European Union 4.8% (163 units).

As of March 2025, the offshore wind market remains concentrated in these three regions: China accounts for 50.8% (7,676 units) of all globally installed OWTs (15,100 total), followed by the European Union with 26% (3,919 units) and the United Kingdom with 19.3% (2,907 units).

This trend is expected to continue. Germany and the United Kingdom have set targets of 30 GW [Bundesrepublik Deutschland, 2016] and 40 GW [UK Government, 2021], respectively, by 2030. Looking further ahead to 2050, the

---

\*Corresponding author [thorsten.hoeser@dlr.de](mailto:thorsten.hoeser@dlr.de)

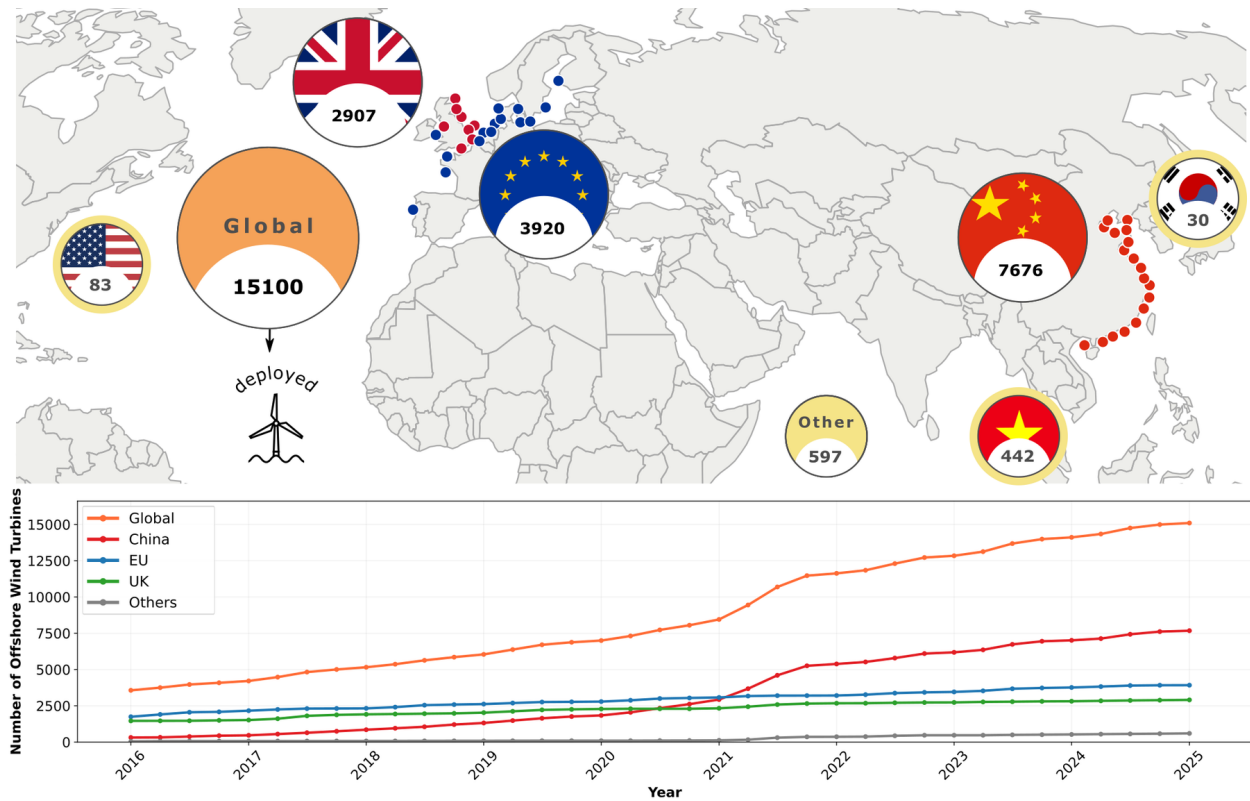


Figure 1: Major offshore wind energy markets with the number of readily deployed turbines in the first quarter of 2025 and their corresponding historic evolution aggregated to quarterly sums.

European Union aims to reach 300 GW of offshore wind energy installed capacity [European Commission, 2020]. China plans to expand its combined installed solar and wind energy capacity to six times its 2020 level, reaching approximately 3,600 GW [United Nations, Department of Global Communications, 2025], with offshore wind representing a rapidly growing component in this strategy. Considering these targets, the recent surge in offshore wind infrastructure deployment is only the beginning of a much larger expansion over the coming decades with a regional focus in the North Sea Basin and along the coast of China.

Beyond established markets, additional regions are expected to contribute to future expansion. Emerging technologies, particularly floating offshore wind turbines successfully demonstrated in pilot projects over the past decade, enable deployment in deeper waters and open unused marine areas for development [Roddier et al., 2010, Bento and Fontes, 2019, Hoenser and Kuenzer, 2022a]. Today the offshore wind energy sector is characterized by continued densification of established hotspot regions and the expansion into new offshore areas. These parallel developments require systematic monitoring of infrastructure deployments and operations, supported by freely available, high-quality data for stakeholders across industry, policy, and science.

From a lifecycle perspective, the construction phase of offshore wind infrastructure is the most resource-intensive, logistically complex, and environmentally impactful stage. Improved understanding and monitoring of this phase are essential for optimizing deployment processes and planning future decommissioning activities.

Although the availability of global location data for offshore wind infrastructure has improved in recent years, particularly through Earth Observation derived products [Wong et al., 2019, Zhang et al., 2021, Hoenser et al., 2022, Paolo et al., 2024, Wang et al., 2024a, Zhang et al., 2024, Ding et al., 2024, He et al., 2025, Fei et al., 2025, Liu et al., 2026], the temporal resolution and semantic detail of existing open datasets remain limited. Specifically, they often lack the capability to resolve discrete events and temporally precise transitions between construction phases, despite the high temporal resolution of freely available Earth Observation data.

This publication advances automated offshore wind infrastructure mapping by extending current automated mapping toward high-density global time series analysis using Sentinel-1 radar Earth Observation data. The manuscript presents a comprehensive approach to monitoring offshore wind infrastructure using Sentinel-1 Synthetic Aperture Radar (SAR)

data. The primary objective is to create a consistent dataset that allows for the analysis of dense time series and events at these infrastructure locations, for regional and global, but also local and highly detailed investigations.

We present global Sentinel-1 SAR time series with median revisit intervals of up to one day within and 12 days outside the European Union, resolving deployment and operational phases of offshore wind infrastructure from 2016 to 2025. The contributions of this paper are as follows:

- A global dataset of 15,606 offshore wind infrastructure locations with global spatial coverage.
- 14,840,637 SAR-based event signals compiled into 15,606 analysis ready time series corresponding to individual infrastructure locations spanning 2016Q1 to 2025Q1.
- A rule-based baseline classification that assigns semantic labels to each event within all time series, distinguishing deployment and operational phases.
- A benchmark dataset consisting of 553 time series with 328,657 expert annotations at the event level, enabling evaluation and methodological comparison.
- Collectively, these contributions establish an open, Earth Observation data corpus that combines global spatial coverage, high temporal resolution, analysis ready time series, baseline predictions, and an annotated benchmark dataset [Hoeser et al., 2026].

Together, these resources enable in-depth analysis of recent global offshore wind infrastructure dynamics and support the development, evaluation, and comparison of time series analysis and classification methods for offshore wind infrastructure applications.

## 2 Related Research

Earth Observation based mapping of offshore wind infrastructure has received increasing attention in recent years. Numerous methods and data products have been developed for large scale automated detection of offshore wind turbines (OWTs) in remote sensing imagery.

Early approaches relied primarily on multispectral optical data. Liu et al. [2016] used Landsat-8 imagery and identified offshore platforms in general by exploiting temporal and positional invariance of detected objects. Xu et al. [2020] analyzed Sentinel-2 and Landsat data, applying statistical filtering and thresholding followed by spatial selection within reported wind farm areas in the North Sea Basin to derive OWT locations.

Several studies demonstrated the suitability of SAR data for OWT detection. Zhang et al. [2019], Wang et al. [2019], and Wong et al. [2019] applied constant false alarm rate (CFAR) to detect OWT locations in SAR imagery, leveraging the strong radar backscatter characteristics of turbine structures. These works laid an important foundation for SAR based OWT detection.

A number of studies have addressed global scale mapping. Zhang et al. [2021] applied morphological filtering to Sentinel-1 SAR imagery within a rule-based approach and produced a first freely available, EO based, global scale OWT dataset called GOWF. A similar rule-based approach was presented by Wang et al. [2024a] for global OWT location mapping. Paolo et al. [2024] motivated by the work of Wong et al. [2019] used a combination of CFAR on Sentinel-1 SAR data and convolutional neural network (CNN) on Sentinel-1 SAR and Sentinel-2 multispectral data to detect and classify marine infrastructure on a global scale with a dedicated "wind" category. Hoeser et al. [2022] introduced DeepOWT, a global scale dataset of offshore wind infrastructure locations that distinguishes OWTs, OWT foundations, and support platforms as separate classes, using a CNN based object detection approach on Sentinel-1 SAR data. Zhang et al. [2024] employed Sentinel-2 imagery and an ensemble of CNN based detection models to reduce false positives in global OWT detection.

Regionally focused studies have also been conducted, particularly for Chinese waters. Liu et al. [2026] and Ding et al. [2024] proposed SAR based OWT mapping using CNN object detection methods, while Song and Li [2025] investigated multispectral Sentinel-2 imagery also using CNN based object detection. Wang et al. [2024b] combined multispectral Sentinel-2 and Sentinel-1 SAR data, applying handcrafted spatial feature detectors for OWT mapping.

For this study, the temporal dimension of offshore wind infrastructure deployment is of particular importance. Existing approaches that incorporate temporal information can broadly be grouped into two categories.

The first category investigates binary changes in SAR amplitude or multispectral reflectance time series to detect persistent changes indicating turbine installation. The studies of Xu et al. [2020], Zhang et al. [2021], Ding et al. [2024], Paolo et al. [2024], and Liu et al. [2026] fall into this category. These approaches provide binary signals (non-turbine/turbine) with annual to monthly temporal resolution. While effective for estimating the onset of deployment,

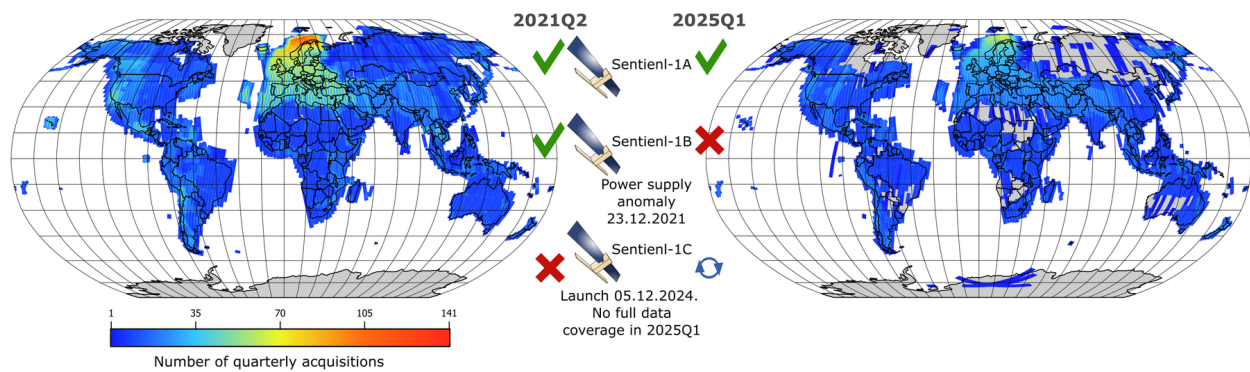


Figure 2: Global distribution of the number of quarterly acquisitions by the Sentinel-1 mission with different platforms in orbit and actively operating.

they do not differentiate fine-grained construction stages such as foundation installation, tower and nacelle assembly, or periods of active construction.

The second category analyzes spatial SAR backscatter patterns to distinguish multiple deployment stages. Hoenser et al. [2022] introduced this category by using three-month aggregated Sentinel-1 SAR data to differentiate between no turbine, turbine foundation, and deployed turbine stages. However, the quarterly aggregation reduces temporal precision and limits the ability to capture short-term events such as vessel interactions or precise phase transitions (e.g., from open water to completed foundation, or from foundation to fully installed turbine).

Overall, current research demonstrates strong progress in spatial localization of offshore wind turbines, particularly at regional and global scales. However, there remains a clear gap in resolving fine-grained temporal dynamics of the deployment process. Accurate spatial detection combined with temporally detailed and semantically differentiated information is essential for exploiting the full potential of Earth Observation derived offshore wind infrastructure datasets.

Such enriched datasets support market analyses and strategic outlooks [Jung and Schindler, 2023, Jung et al., 2024], enable investigation of regional deployment and decommissioning dynamics [Gourvenec et al., 2022], provide detailed inputs necessary for spatial planning and modelling [Schwartz-Belkin and Portman, 2023], and support ecosystem impact studies through precise temporal correlation of construction phases with animal behaviour [Lai et al., 2025].

To address these limitations, we aim to establish a foundation for high-density time series analysis of offshore wind infrastructure, enabling temporally precise and semantically rich monitoring at a global scale.

### 3 Data and Materials

This study uses data from the Sentinel-1 satellite constellation, part of the European Space Agency’s (ESA) Copernicus Programme. Sentinel-1 provides C-band (5.6 cm wavelength) SAR imagery which is largely independent of cloud cover and solar illumination [Torres et al., 2012]. Combined with revisit intervals ranging from 1 to 12 days, depending on number of platforms, latitude and geographic region, Sentinel-1 is well suited for generating high-density time series at global scale, particularly for coastal and offshore regions.

Spatial and temporal coverage highly depend on the number of operational platforms within the constellation, see figure 2. During the second quarter of 2021, both Sentinel-1A and Sentinel-1B were fully operational, with the highest data take frequency over the European Union. Following the power supply anomaly of Sentinel-1B on 2021-12-23 [European Space Agency (ESA), n.d.], and prior to the operational phase of Sentinel-1C (launched but not acquiring data in the first quarter of 2025), Sentinel-1A was the only operating platform during this time. Even when this resulted in reduced spatial coverage and lower revisit frequency, coastal and offshore regions continued to receive consistent acquisitions. The constellation’s fallback design demonstrates that even under an unplanned platform failure, Sentinel-1 remains capable of delivering continuous time series suitable for coastal and offshore monitoring applications.

For persistent offshore infrastructure detection, SAR backscatter characteristics are particularly suitable, see figure 3. This study investigates Sentinel-1 Ground Range Detected (GRD) products acquired in Interferometric Wide (IW)

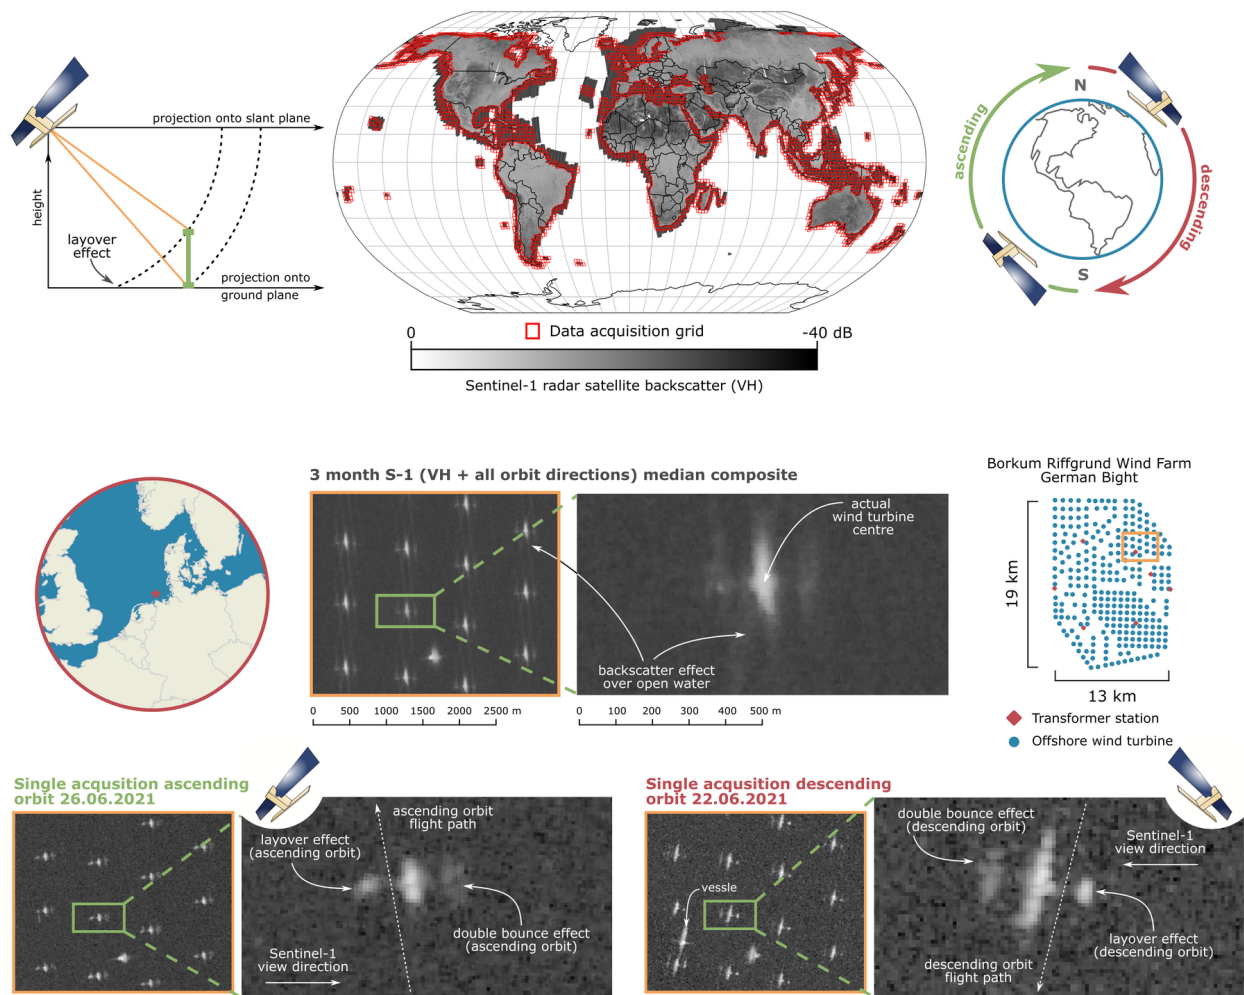


Figure 3: Global processing grid overlaid on a three-month Sentinel-1 median composite, together with schematics of the satellite's polar orbit and radar acquisition geometry and representative Sentinel-1 scenes demonstrating viewing-direction-dependent imaging effects.

swath mode with VH (vertical transmit, horizontal receive) polarization, preprocessed and made available via Google Earth Engine [Gorelick et al., 2017, Google Earth Engine Team, 2026].

Due to the near-polar orbit and right-looking acquisition geometry of Sentinel-1, images are acquired in ascending and descending passes, meaning targets are observed either from west-to-east or east-to-west viewing directions. For vertical metallic structures such as offshore wind turbines, this geometry produces characteristic radar signatures [Hoeser and Kuenzer, 2022a].

A layover effect occurs when e.g. signals reflected from elevated structures return to the sensor before signals reflected from the ground or sea surface at the structure's base. For OWTs, strong reflections from the nacelle generate high backscatter values that, in the final image data, appear displaced toward the sensor relative to the true turbine location. Visually, the turbine appears to lean toward the sensor. Depending on orbit direction, this layover signature appears to the left for ascending passes and right for descending passes relative to the OWT centre. This horizontal signature is a key element of the dataset design and subsequent analysis presented in this study.

When combining acquisitions from different orbit directions into temporal composites like a median stack, the turbine centre and layover-related backscatter peaks are sharpened, while speckle over open water is reduced. Transient objects such as vessels are largely suppressed in median composites, if they do not remain stationary during the compositing period [Hoeser et al., 2022, Paolo et al., 2024].

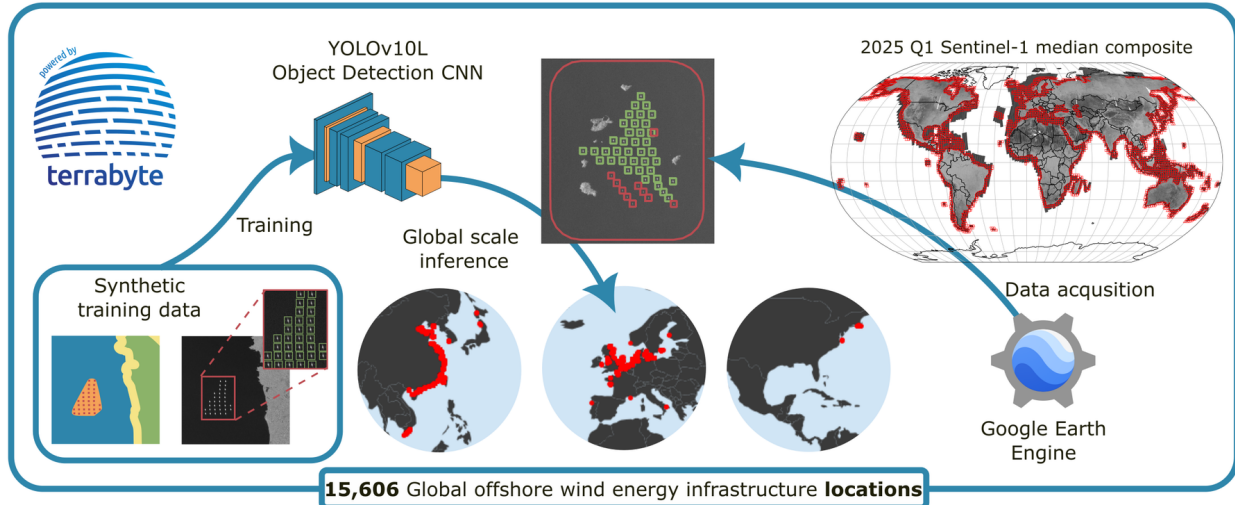


Figure 4: Deep learning based, global, offshore wind infrastructure detection workflow based on Hoeser et al. [2022].

For these reasons, Sentinel-1 SAR imagery is highly suitable for global offshore wind infrastructure detection and for generating high-density time series capable of resolving structural changes during deployment and operation.

## 4 Method

### 4.1 Updated global offshore wind infrastructure detection

To derive offshore wind infrastructure locations at global scale, we updated the detection workflow proposed by Hoeser et al. [2022]. The workflow consists of a global offshore wind farm proposal network, and a second-stage detector that identifies individual infrastructure units within the proposed regions.

As illustrated in figure 4, we trained YOLOv10-L object detection models provided by Ultralytics [Wang et al., 2024c] to improve inference throughput in the two-stage global detection process. This replaces the Faster R-CNN [Ren et al., 2017] architecture used in the preceding study. Following Hoeser and Kuenzer [2022b], Hoeser et al. [2022], we generated 90,000 synthetic training samples with corresponding annotations for both offshore wind farms and individual offshore wind infrastructure units to train the models for both stages.

Training was conducted on four NVIDIA A100 GPUs hosted on DLR’s Terrabyte HPC [DLR terrabyte, 2026]. The models were trained using an input size of  $1024 \times 1024$  pixels, a batch size of 32, and 20 epochs. A linearly decaying learning rate schedule was applied, decreasing from 0.01 to 0.0001 following a three-epoch warm-up phase.

Inference was performed on median composites of the Sentinel-1 GRD IW product. All images which are within the global Exclusive Economic Zones (EEZs) [Flanders Marine Institute, 2024] with VH polarization for both orbit directions are queried and processed on GEE for 2025Q1. The resulting tiles were exported to Terrabyte HPC and organized as a global mosaic using the GDAL tile index (GTI). A virtual chip grid of  $2048 \times 2048$  pixels with 50% overlap is used to enable highly parallelized data access via the GTI during a CPU based inference process running up to 24,000 concurrent model instances.

Detection bounding boxes of the second object detection stage were reduced to centroid coordinates to get point representations of offshore wind infrastructures for validation. Validation was performed using manually annotated offshore wind infrastructure locations derived from the Sentinel-1 2025Q1 median composite across three validation sites.

We updated the validation datasets from Hoeser et al. [2022] to include infrastructure constructed up to 2025Q1 in the North Sea Basin and the East China Sea. Additionally, a new validation site was introduced along the southeast coast of Vietnam to include challenging to detect offshore wind farm layouts characteristic for this region. Overall, we provide 9,770 offshore wind energy infrastructure validation locations. Because classification of deployment state and object type is addressed in the subsequent time series analysis, detection was evaluated in a class-agnostic manner, treating turbine foundations, deployed turbines, and support platforms as a single class.

Performance was quantified using Precision, Recall, and F1 score at site level, along with macro- and micro-averaged scores across sites to account for class imbalance. A detection was considered a true positive (TP) if its centroid lay within a radius of 100 m of a ground truth point. Detections outside this radius were counted as false positives (FP), and ground truth points without a corresponding detection within 100 m were counted as false negatives (FN).

We handled edge cases, so that if multiple detections matched the same ground truth point, only the closest detection was counted as TP, additional detections were counted as FP. Furthermore, if a detection fell within 100 m of multiple ground truth points, it was assigned to the nearest ground truth location as TP and unmatched ground truth points were counted as FN.

Evaluation metrics were computed as:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad (1)$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad (2)$$

$$F_1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (3)$$

With this modernized workflow building upon Hoeser et al. [2022], we provide an updated global dataset of offshore wind infrastructure locations, which serves as the foundation for compiling and analysing high-density time series of offshore wind infrastructure dynamics.

## 4.2 High density Sentinel-1 time series compilation

The primary contribution of this publication is the extension of global offshore wind infrastructure detections with high-density Sentinel-1 based time series, providing fine-grained details of temporally resolved, Earth Observation based information on deployment and operational phases of offshore wind infrastructure. To this end, we query and preprocess every available Sentinel-1 acquisition at each detected offshore wind infrastructure location individually.

Figure 5 illustrates a final example time series for a single offshore wind turbine covering the period from 2016 to 2025. The 1D backscatter profiles are generated by applying a column-wise maximum reduction along the horizontal axis of each 2D image. These profiles capture characteristic SAR signatures whose spatial patterns vary depending on the infrastructure state. The resulting time series consist of one 1D backscatter profile per Sentinel-1 acquisition per turbine location.

Data acquisition was performed on GEE, while client-side postprocessing was used to compile the final analysis ready time series dataset, see figure 6. GEE's strongly uses parallelizable map-reduce operations, where in our case, the applied operation consists of a column-wise maximum reduction to derive 1D profiles from 2D Sentinel-1 imagery crops.

Three main challenges had to be addressed:

1. Defining a spatial feature representation compatible with GEE's region-based reduction.
2. Designing a scalable and efficient batch processing pipeline which matches GEE's image collection filtering.
3. Implementing a data model compliant with GEE constraints while minimizing export size and data egress costs.

A naïve approach would query an image, crop the region of interest, and apply a column-wise reduction to obtain a 1D profile. However, GEE's region reduction assigns a single aggregated value per geometry, which would reduce each detection box into a single scalar rather than a profile.

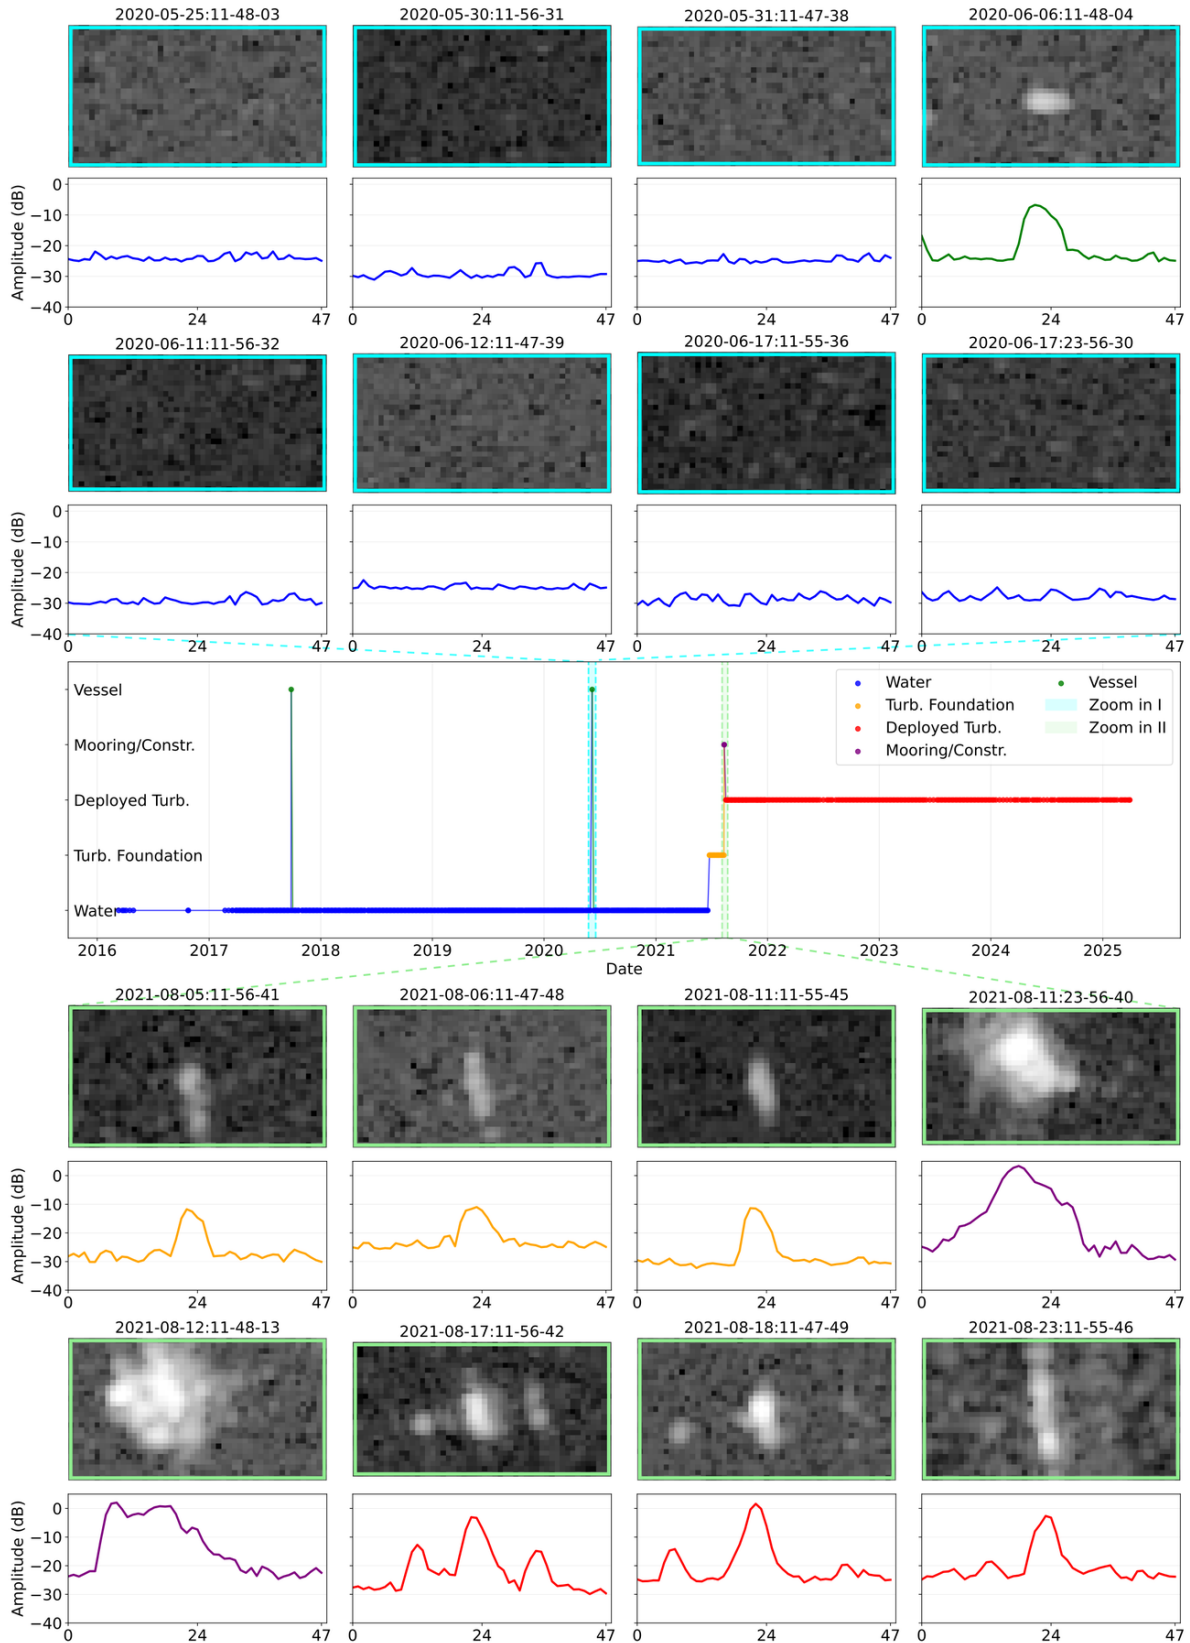


Figure 5: Example time series of a single offshore wind turbine. The central panel shows predicted event labels, the upper and lower panels display corresponding 2D Sentinel-1 backscatter images and derived 1D backscatter profiles.



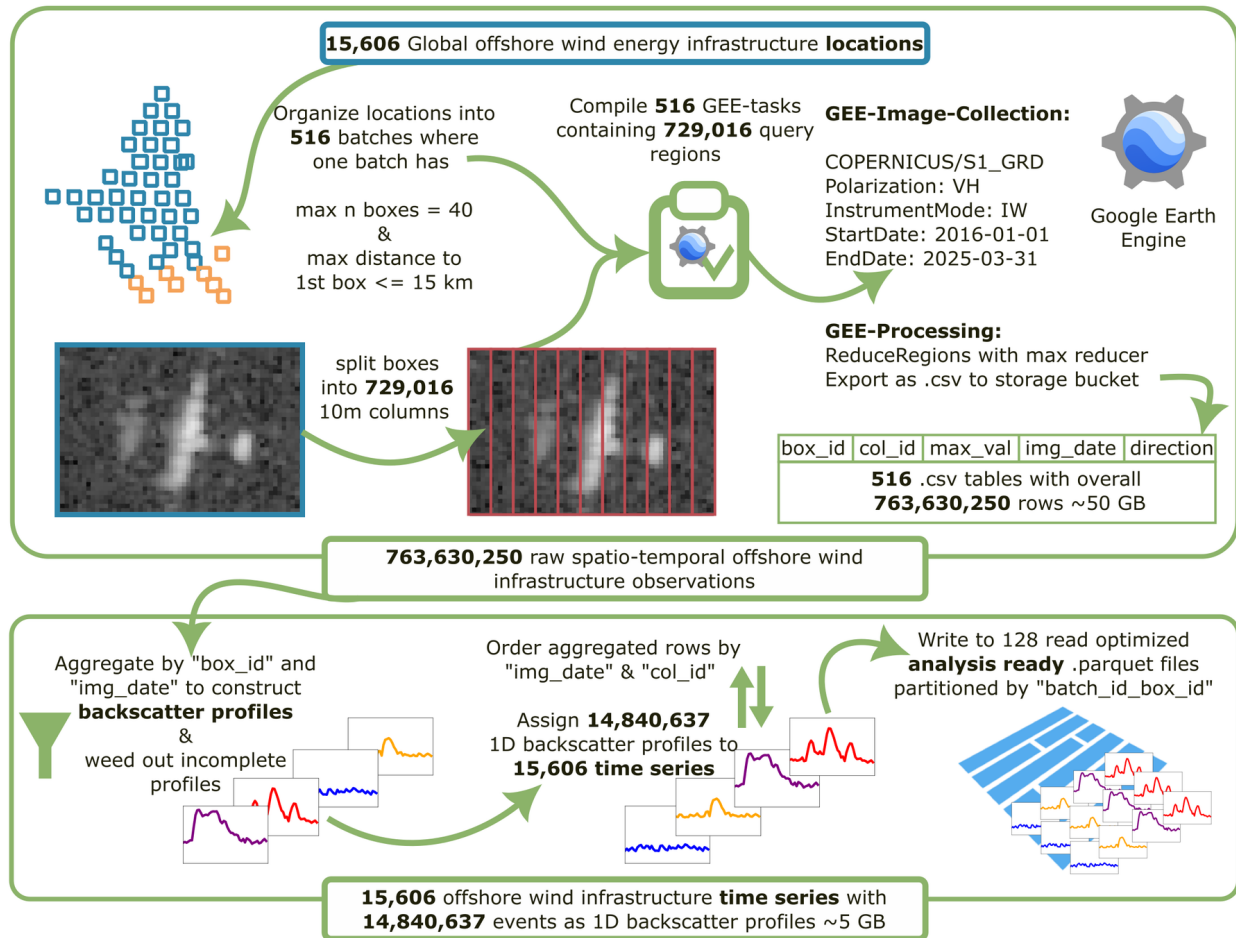


Figure 6: Time series compilation workflow, based on 15,606 offshore wind infrastructure spatial detections to efficiently acquire 14.8 million Sentinel-1, 1D backscatter profiles with unique time and location relations.

To overcome this limitation, each of the 15,606 detection boxes was split client side into 10 m wide vertical columns along the horizontal axis, resulting in 729,016 narrow spatial processing units. Instead of launching one GEE task per column, boxes (not columns) were grouped into spatial batches under the following constraints:

- A batch includes all boxes within 15 km of a randomly selected seed box.
- A batch contains no more than 40 boxes.

The 15 km spatial threshold increases the likelihood that all boxes in a batch share common image footprints, minimizing unnecessary NaN values due to partial coverage. The limit of 40 boxes per batch was determined empirically and maintains high parallelism while preventing overly large export objects that may cause GEE task failures. Applying this batching strategy resulted in 516 processing batches for the 15,606 detected locations.

Each GEE batch exported a CSV file. The tabular format was selected to allow inclusion of metadata alongside reduced backscatter values. The data model contains:

- `box_id`: identifier of the parent detection box
- `column_id`: identifier of the 10 m column within the box
- `max_value`: maximum backscatter value for the column
- `img_date`: Sentinel-1 acquisition date
- `orbit_direction`: ascending or descending pass

Including orbit direction is essential for later interpretation of directional SAR signatures. The 516 exported CSV files collectively contain 763,630,250 rows. GEE task preparation and submission was managed client side, maintaining up to 25 concurrent or queued GEE tasks with automated retry logic after failure. Data were exported to a Google Cloud Storage bucket, resulting in about 50 GB of data and total egress cost of about 6 USD. This cloud-based workflow demonstrates both computational and resource efficiency for large-scale SAR time series acquisition.

After download, batch level CSV files were processed individually. Rows were aggregated by `box_id` and `img_date` to group all backscatter profile values belonging to a single location and acquisition. Entries containing NaN values were removed. To reconstruct coherent 1D backscatter profiles, column values were sorted by `column_id`. This resulted in 14,840,637 valid 1D backscatter profiles associated with offshore wind infrastructure locations and unique acquisition dates.

For efficient storage and data access, the dataset was saved into 128 partitioned Parquet files. Partitioning was implemented by computing a hash of `batch_id + box_id` and mapping the resulting integer to the range 0-127 using a modulo operation. This scheme ensures that all profiles belonging to the same location are stored within the same partition while maintaining near-uniform partition sizes. As a result, loading individual location-specific time series is highly efficient.

### 4.3 Rule-based time series classification

The goal is to provide a rule-based classifier as baseline approach, which is able to assign the labels water, turbine foundation, deployed turbine, vessel, mooring / active construction, or as fallback unclear to each 1D backscatter profile of the prepared time series.

In a preprocessing step, before classification, for each backscatter profile a peak detection is performed [Virtanen et al., 2020] after a Gaussian smoothing of the original backscatter profile to focus on major peaks. During peak detection, candidate peaks must satisfy a minimum inter-peak distance of 5 profile bins and a minimum prominence of 2 dB. The detected peaks are mapped back to their positions to the unsmoothed profile indices. To distinguish peaks resulting from the actual target position and layover effects, the centre peak is defined as the peak closest to the profile centre. Layover peaks are defined based on their orbit direction. For descending passes, the right adjacent peak serves as the valid side peak, and for ascending passes, the left adjacent peak is used. Furthermore, for each profile we calculate the mean, standard deviation and range. This set of features is used during profile classification.

The rule-based classifier is designed in two stages. In the first stage, each profile is classified independently using thresholds on profile statistics (mean, standard deviation, range) and peak characteristics (presence, prominence, width, and amplitude). Figure 5 should be revisited here to get an intuition of the characteristics of the 1D backscatter profiles and different profile types. A water label is assigned if the profile has low overall backscatter or lacks distinct peaks, including cases of low mean and variability, weak or absent centre peaks, or uniformly low peak prominence. A mooring / active construction label would indicate the presence of a larger construction vessel or maintenance vessel, for which we assume strong and broad backscatter profiles. This label is assigned when broad peaks or high dynamic range indicate strong but spatially diffuse responses, or when low variance appears together with high mean backscatter. A deployed turbine label is assigned when a pronounced centre peak and a valid layover peak are present with sufficient amplitude and prominence. An unclear label is assigned to profiles with elevated side peaks that do not satisfy any of the preceding criteria, while a turbine foundation label is assigned when a sufficiently prominent centre peak is present and no earlier rule applies. All remaining profiles default to unclear after this first stage.

In the second stage, labels are refined using sequential context. First, a smoothing operation is performed by applying a moving window with an adaptive kernel that ignores transient labels (unclear, vessel, mooring) when determining effective neighbours. Targeted passes correct isolated non-water profiles flanked by water and enforce local consistency by relabelling profiles that disagree with two agreeing neighbours, except for vessel labels, which are preserved. Next, profiles still labelled unclear are reevaluated using relaxed peak-based criteria to resolve ambiguities between mooring / active construction and turbine foundation classes. The smoothing step is then repeated to propagate updated labels.

Subsequent refinements correct isolated water labels between turbine-related classes using a sliding window. Larger temporal contexts are then considered through segment-based refinement, where adjacent deployed turbine and turbine foundation segments are compared chronologically. Treating mooring and vessel labels as transparent, the shorter of two conflicting adjacent segments is relabelled to match the longer one. This process is then iterated until convergence.

Finally, a timeseries refinement resolves platform related ambiguities by examining the global label distribution. Depending on the dominance and ordering of vessel, mooring, turbine foundation, and deployed turbine labels, subsets of these labels are reassigned to the platform class, including cases where turbine foundation labels persist long after the

last deployed turbine occurrence. The final labels are written to the time series parquet files holding the corresponding 14,840,637 swath profiles.

To evaluate the performance of the rule-based classification, we examined and annotated 328,657 swath profiles of 553 time series randomly drawn from all 15,606 available time series. This benchmark data set is part of the data corpus provided with this publication. We used two validation strategies, single event classification evaluation and time series evaluation.

For single event evaluation, the ground-truth event labels in the benchmark dataset are compared against the labels predicted by the rule-based classifier using a confusion matrix, from which class-wise Precision, Recall, and F1-scores are derived, as well as overall and grouped macro averaged scores.

To complement these point wise scores with a sequence sensitive metric, we compute the edit similarity score. Edit similarity is defined as the normalized Levenshtein distance [Levenshtein, 1966] between the ground-truth and predicted label sequences, which we chose due to different sequence lengths of the time series:

$$\text{EditSim}(S, \hat{S}) = 1 - \frac{d_{\text{Lev}}(S, \hat{S})}{\max(|S|, |\hat{S}|)}, \quad (4)$$

where  $d_{\text{Lev}}$  denotes the Levenshtein distance and  $|S|$  and  $|\hat{S}|$  are the sequence lengths. The Levenshtein distance measures how many edit operations like insertions, deletions, or substitutions of labels are necessary to align the predicted sequence with the ground-truth sequence, where fewer required edits express a higher similarity. Edit similarity is computed in two variants. The first variant compares the full label sequences performing a point-wise evaluation, thereby penalizing incorrect labels at any time step. Since this variant is conceptually close to a micro-averaged F1 score obtained from single-event evaluation, we additionally compute edit similarity on collapsed label sequences. During collapsing, segments of identical labels are reduced to a single label. Collapsed edit similarity abstracts away absolute event durations and focuses on the temporal ordering and presence of event phases. This variant penalizes unnecessary transitions, incorrect event ordering, and missing or extra phases in the predicted sequence, making it a suitable complementary metric to single-event evaluation. For benchmarking purposes, edit similarity is further evaluated at uniformly spaced quality thresholds  $q_s \in [0, 1]$  with a step size of 0.05. For each threshold, we compute the fraction of time series with an edit similarity score greater than or equal to  $q_s$ . The overall performance is summarized as the area under the edit similarity-quality threshold curve (AUC).

## 5 Results

### 5.1 Dense SAR backscatter profile time series

As a primary outcome of this study, we provide a globally consistent, high-temporal-resolution time series dataset of offshore wind infrastructure derived from Sentinel-1 SAR imagery available at Hoerer et al. [2026]. While an updated set of offshore wind infrastructure locations for 2025Q1 is included as a structural foundation, the central contribution lies in the analysis-ready time series data and the benchmark annotations.

The 14,840,637 analysis-ready 1D Sentinel-1 backscatter profiles spanning the period from 2016Q1 to 2025Q1 are a starting point for detailed investigation and research in Earth Observation based offshore wind infrastructure monitoring by complementing spatial detection approaches with high density, temporal infrastructure dynamics. Each backscatter profile corresponds to a single acquisition at a detected offshore wind infrastructure location and preserves structurally informative SAR signatures.

The temporal resolution and thus number of events of each time series directly reflects Sentinel-1 acquisition frequency. During dataset engineering, we focused on extracting the densest possible time series from the Sentinel-1 archive. As shown in figure 7, temporal resolution varies geographically. Regions within the European Union show dense coverage corresponding to a higher number of events per time series due to higher acquisition frequencies and overlapping Sentinel-1 orbits, especially at higher latitudes, resulting in revisit intervals of up to one day for certain locations. In contrast, regions such as the Chinese EEZ generally follow the mission’s default 12-day revisit interval of a single Sentinel-1 platform.

The provided 1D backscatter profiles are balancing feature richness, global completeness, and technical usability. The objective was to construct a globally consistent dataset at maximal temporal resolution while preserving SAR features necessary for detailed deployment-phase analysis for a broad user audience, with different computational resources.




Figure 7: Global and regional distributions for the number of events per time series and the median number of days between time series events.

The 14.8 million 1D backscatter profiles need about 5 GB of storage, enabling broad accessibility without requiring large-scale computational infrastructure. In contrast, storing the corresponding 2D Sentinel-1 image patches would increase data volume by an estimated factor of 25, significantly limiting usability to users with advanced computing resources.

## 5.2 Global Offshore Wind Infrastructure Detection

With the updated two stage object detection workflow based on Hoerer et al. [2022], we identified 15,606 offshore wind infrastructure across the entire global EEZ based on Sentinel-1 VH-polarized GRD imagery reduced to a 2025Q1 median composite. Of these, 15,100 were subsequently classified as offshore wind turbines in the time series analysis, see figure 1.

At global scale, two dominant offshore wind hotspots exist: the Chinese EEZ and the North Sea Basin (NSB). Figure 8 illustrates the bounding boxes of offshore wind farm regions detected by the first-stage, which scans the global coastline and offshore areas. Zoomed examples further demonstrate detections of individual offshore wind infrastructure units identified by the second-stage within these candidate regions.

Ongoing expansion of the offshore wind energy sector introduce new spatial layouts that challenge detection models. Figure 9 highlights complex environments such as harbours and near-coastal environments, where offshore wind turbines may be located close to piers or connecting infrastructure. Such sites pose recall challenges, increasing the risk of false omissions. Especially the wind farm layouts along the southeast coast of Vietnam represent challenging patterns that differ from typical patterns in the NSB or Chinese EEZ. At the same time, precision is challenged by coastal and harbour infrastructure with radar signatures similar to these layouts.

Detection performance across the three validation sites, the East China Sea, the NSB, and Southeast Vietnam is summarized in figure 10 using Precision, Recall, and F1 score. Across all validation sites, performance ranges from 95.2% to 100% for individual metrics. The overall macro-averaged F1 score is 98.1%, indicating robust global detection performance. The lowest metric observed is a Recall of 95.2% in the NSB. The false omissions belong mainly to a large cluster of turbine foundations under construction at the Dogger Bank, which were not identified as candidate wind farm regions in the first-stage detection of the applied workflow. Overall, the updated detection workflow provides a well performing and globally consistent spatial foundation for the subsequent high-temporal-resolution time series analysis.

## 5.3 Time series classification

For the 15,606 detected offshore wind infrastructure locations, we classified 15,606 corresponding time series with a total of 14,840,637 event labels. Six semantic classes are distinguished: water, turbine foundation, deployed turbine, support platform, mooring / construction vessel, and vessel.

Validation was performed on a benchmark subset of 553 expert-annotated time series containing 328,657 event labels. On the two important summary metrics, the area under the collapsed edit similarity-quality threshold curve (AUC) and the macro F1 score of event-based comparison, the proposed classifier performed with AUC of 0.785 and an F1 of 0.84.

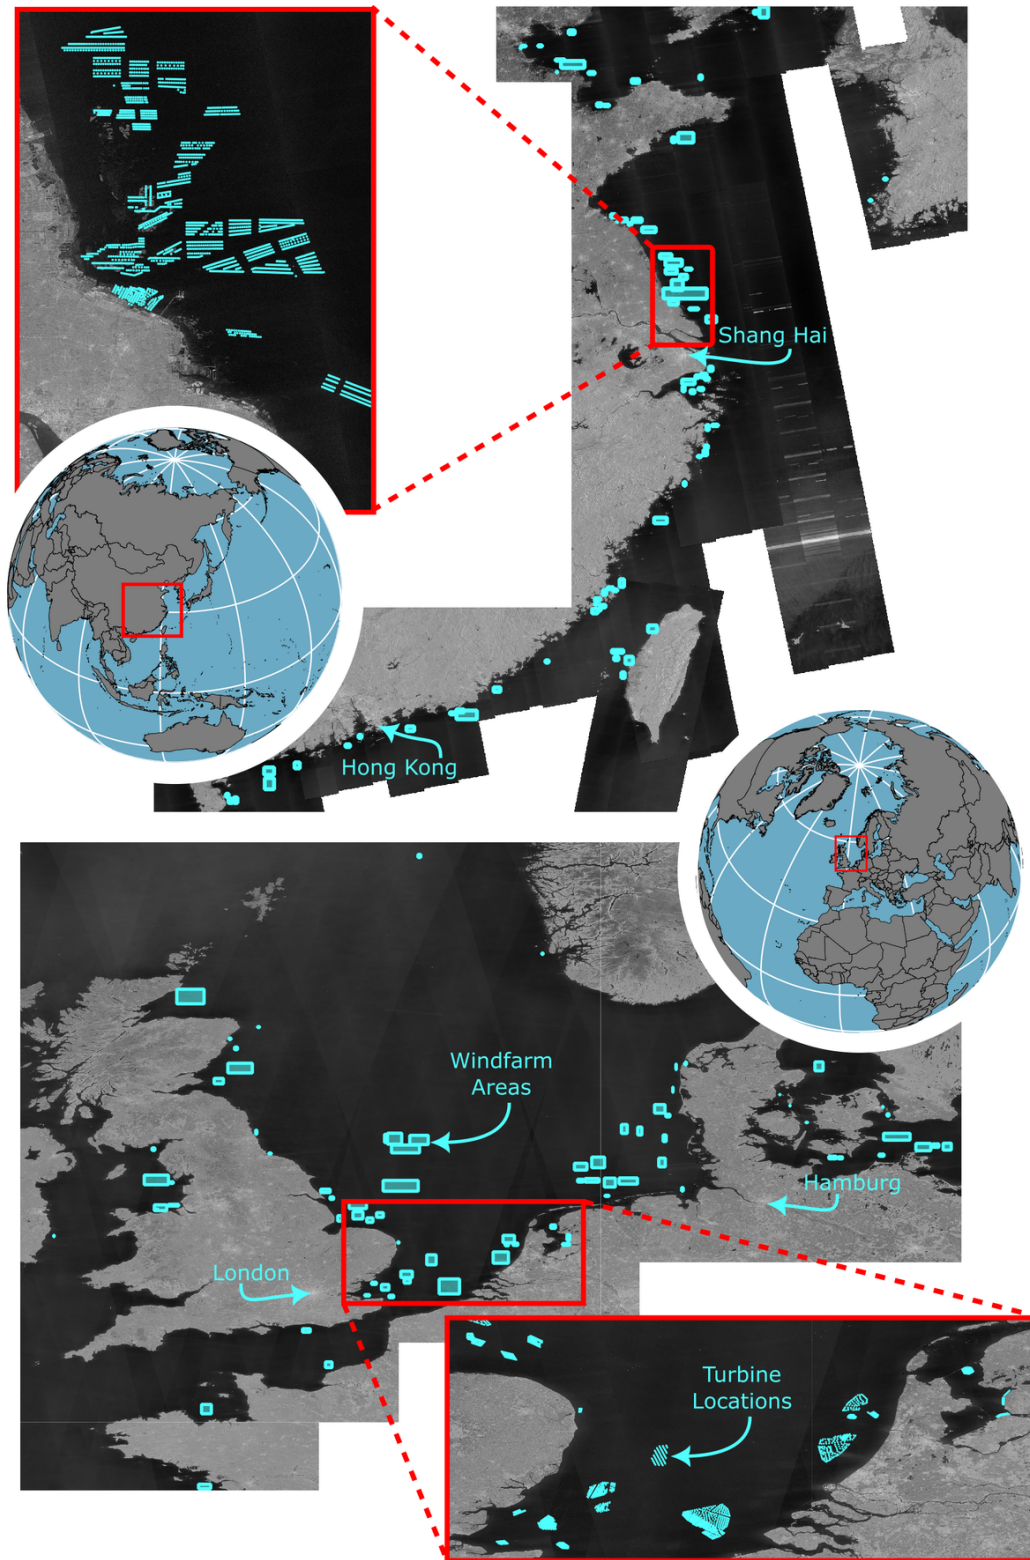


Figure 8: Results of the deep learning based offshore wind infrastructure detection workflow, showing bounding boxes for offshore wind farm areas proposed by a first CNN, and single offshore wind infrastructure locations within, provided by second CNN. The upper part shows detections in Chinese waters, the lower part in the North Sea Basin, the two global wind energy hot spots.

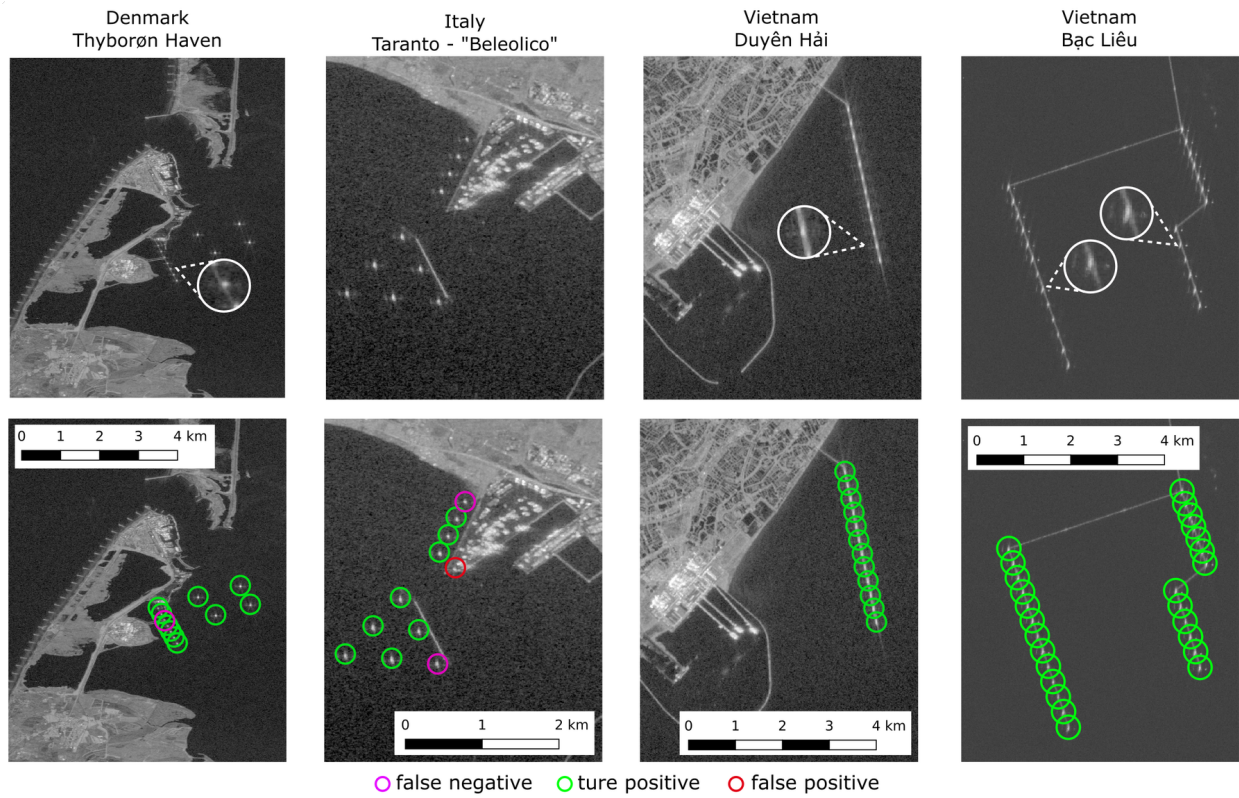


Figure 9: Detection results at challenging near coast and harbor environments.




Figure 10: Detection performance at the three validation sites, the East China Sea, North Sea Basin and southeast Vietnam.

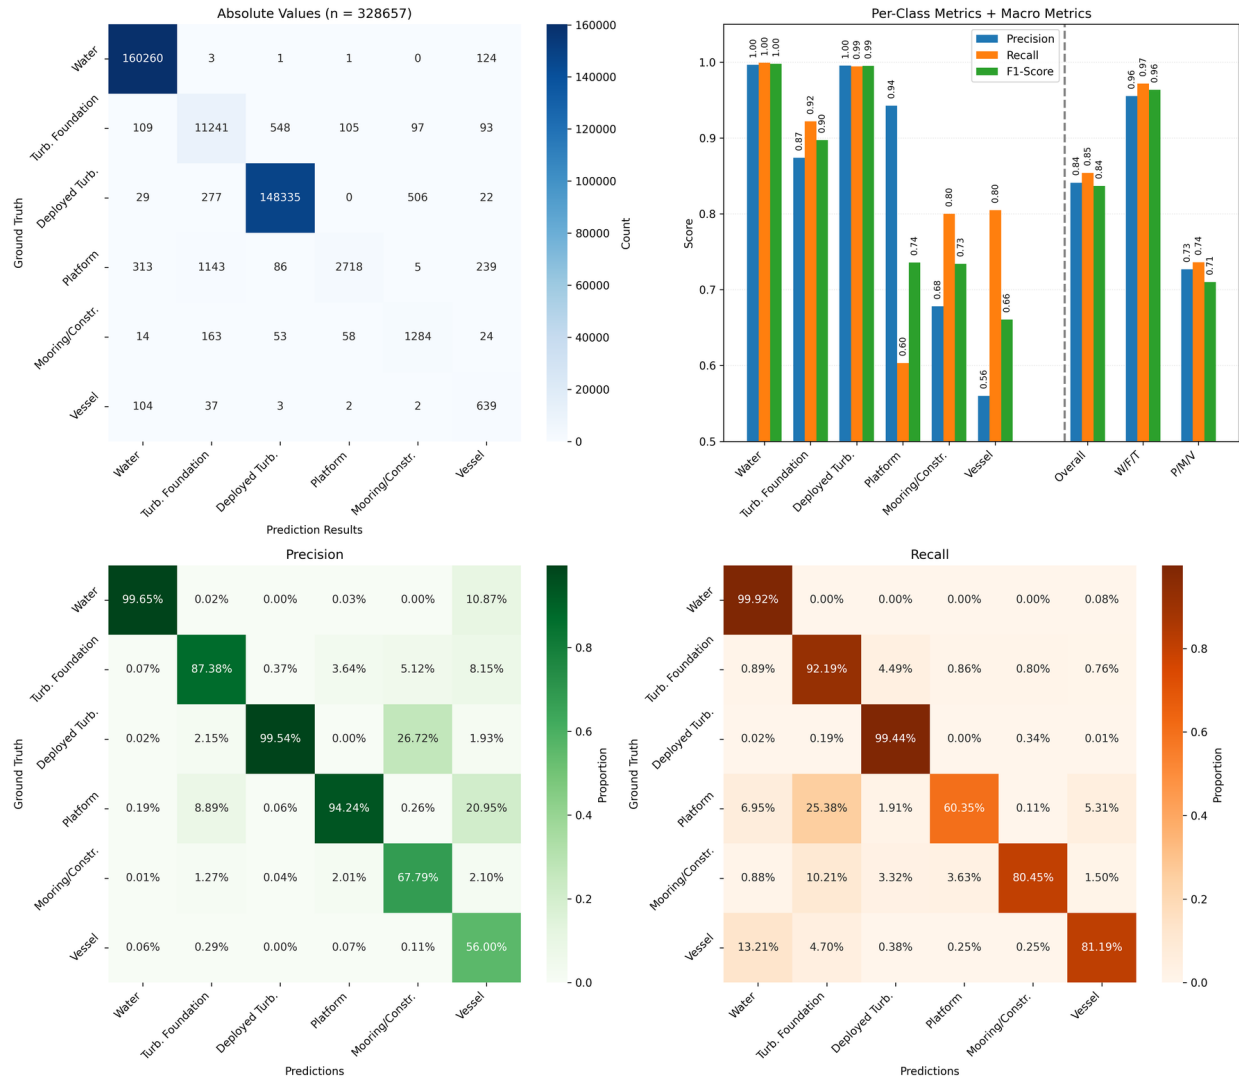


Figure 11: Event level, point-wise classification results as on predictions from the rule-based classifier. In the upper left metric overview of Precision, Recall and F1, Water, Turbine Foundation and Deployed Turbine are aggregated to W/F/T, Platform, Mooring/Construction vessel and Vessel are aggregated to P/M/V.

Looking closer into the point wise evaluation of event classification, figure 11 reveals a clear separation in classification performance between deployment-related turbine classes (water, turbine foundation, deployed turbine; short W/F/T) and non-turbine or transient classes (support platform, mooring / construction vessel, vessel; short P/M/V). The turbine-related classes achieve a higher macro F1 score (0.96), whereas the platform and transient vessel-related classes show lower performance (macro F1 = 0.71).

To complement single-event classification metrics, we evaluate sequential consistency using edit similarity. This metric quantifies how well predicted label sequences match the temporal structure of ground-truth sequences. Particularly informative is the collapsed edit similarity, see figure 12. Its summary statistic AUC provides a measure of sequence and segment agreement and is therefore the most challenging and relevant benchmarking score for future studies. The rule-based classifier scores with an AUC of 0.785, with the steepest decline in successful edit similarity at 0.6 to 0.8 quality levels, stabilizing after the 0.9 quality level at 35% of the classified time series. Meaning, 35% of the classified time series have perfect sequential alignment with the ground truth data.

In figures 13 and 14 we illustrate representative classified time series in spatial and temporal context.

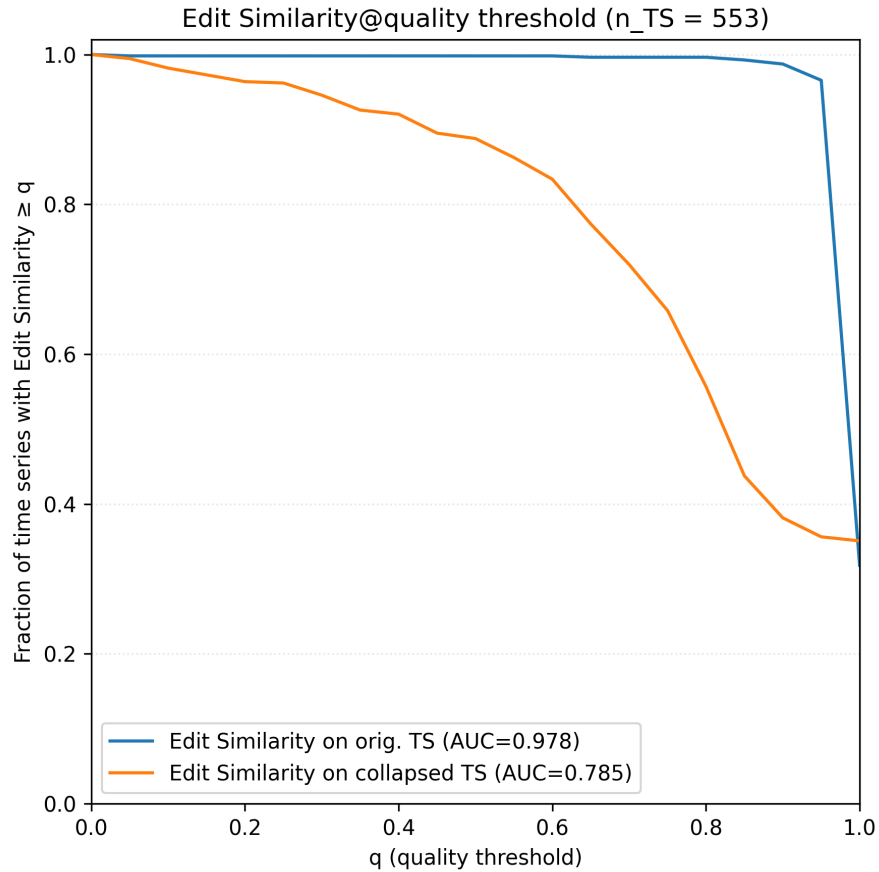


Figure 12: Fraction of Time Series with Edit Similarity scores higher than a given threshold at different threshold levels with corresponding AUC values on step-wise and collapsed classified time series.

In the NSB, deployment phases are clearly separated into distinct foundation construction and turbine installation stages, with a high vessel interaction during peak construction periods. In contrast, examples from the Chinese EEZ show a less distinct separation between foundation and turbine deployment stages. The lag between foundation construction and turbine installation, commonly observed in NSB sites, is often absent, resulting in shorter overall construction durations per turbine. However, vessel interaction patterns remain consistent across regions, with higher event density during construction compared to operational phases.

Using the classified time series, we derive regional deployment duration estimates by identifying the dates of first turbine foundation detection and first deployed turbine detection for each site and computing the elapsed time between these phases. The comparison of China, the European Union, and the United Kingdom shows significantly shorter deployment durations within the Chinese EEZ, see figure 15, which supports the observed differences in the example time series. These results demonstrate how high density, temporal classification enables large-scale analysis of regional deployment patterns.

To illustrate the sensitivity of the classified time series to operational events at a local scale, we present the case study of maintenance events at Hywind Scotland. Figure 16 presents results for the Hywind Scotland floating wind farm. In 2024, all five turbines were temporarily towed back to Norway for maintenance. According to official reports from Global Maritime [Global Maritime, 2024], maintenance began on 2024-05-17, with the first turbine towed back to Norway reaching the port on 2024-05-24. The first turbine returned on 2024-07-05, and the final turbine on 2024-09-22.

These documented events align closely with the observed and classified time series. The classification captures the absence of deployed turbine signatures during the maintenance periods. This case study provides an example, that the time series data of this publication provides temporally detailed insights exceeding those available from summarized reports, demonstrating the capability of the dataset to detect and contextualize operational dynamics at individual sites.



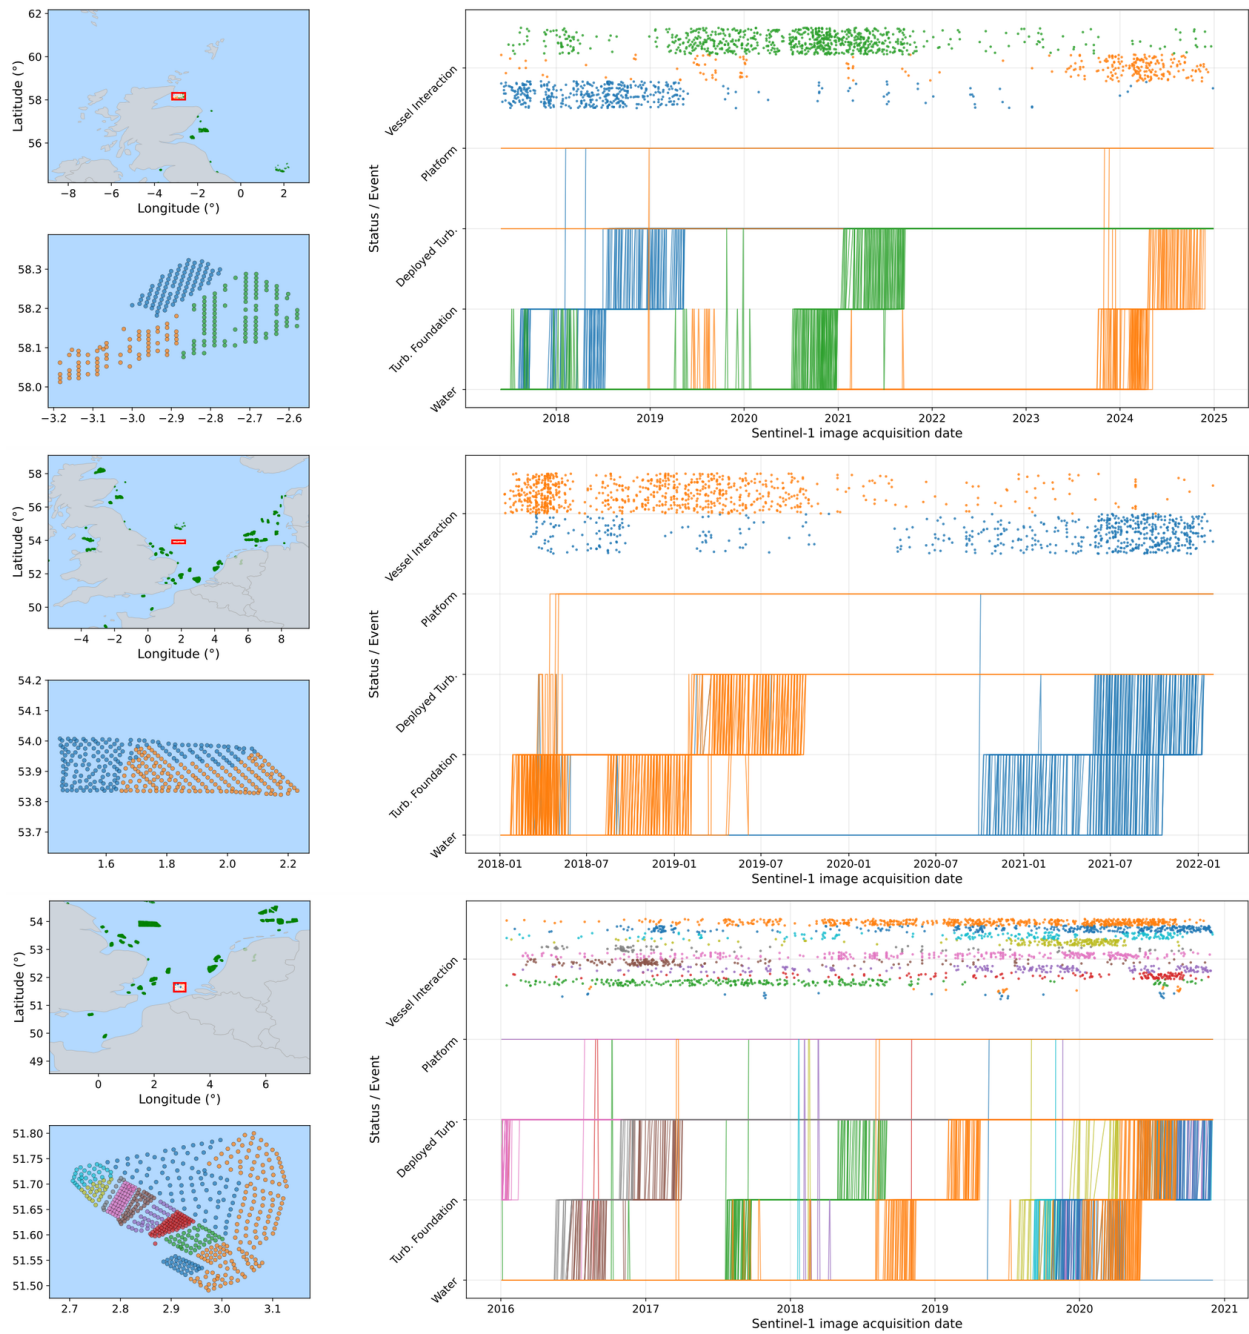


Figure 13: Example time series of three offshore wind farm clusters in the North Sea Basin. Each time series is a coloured line corresponding to the location markers in the close-up overview map. Vessel related labels are markers in the time series with a random jitter for better visualization.

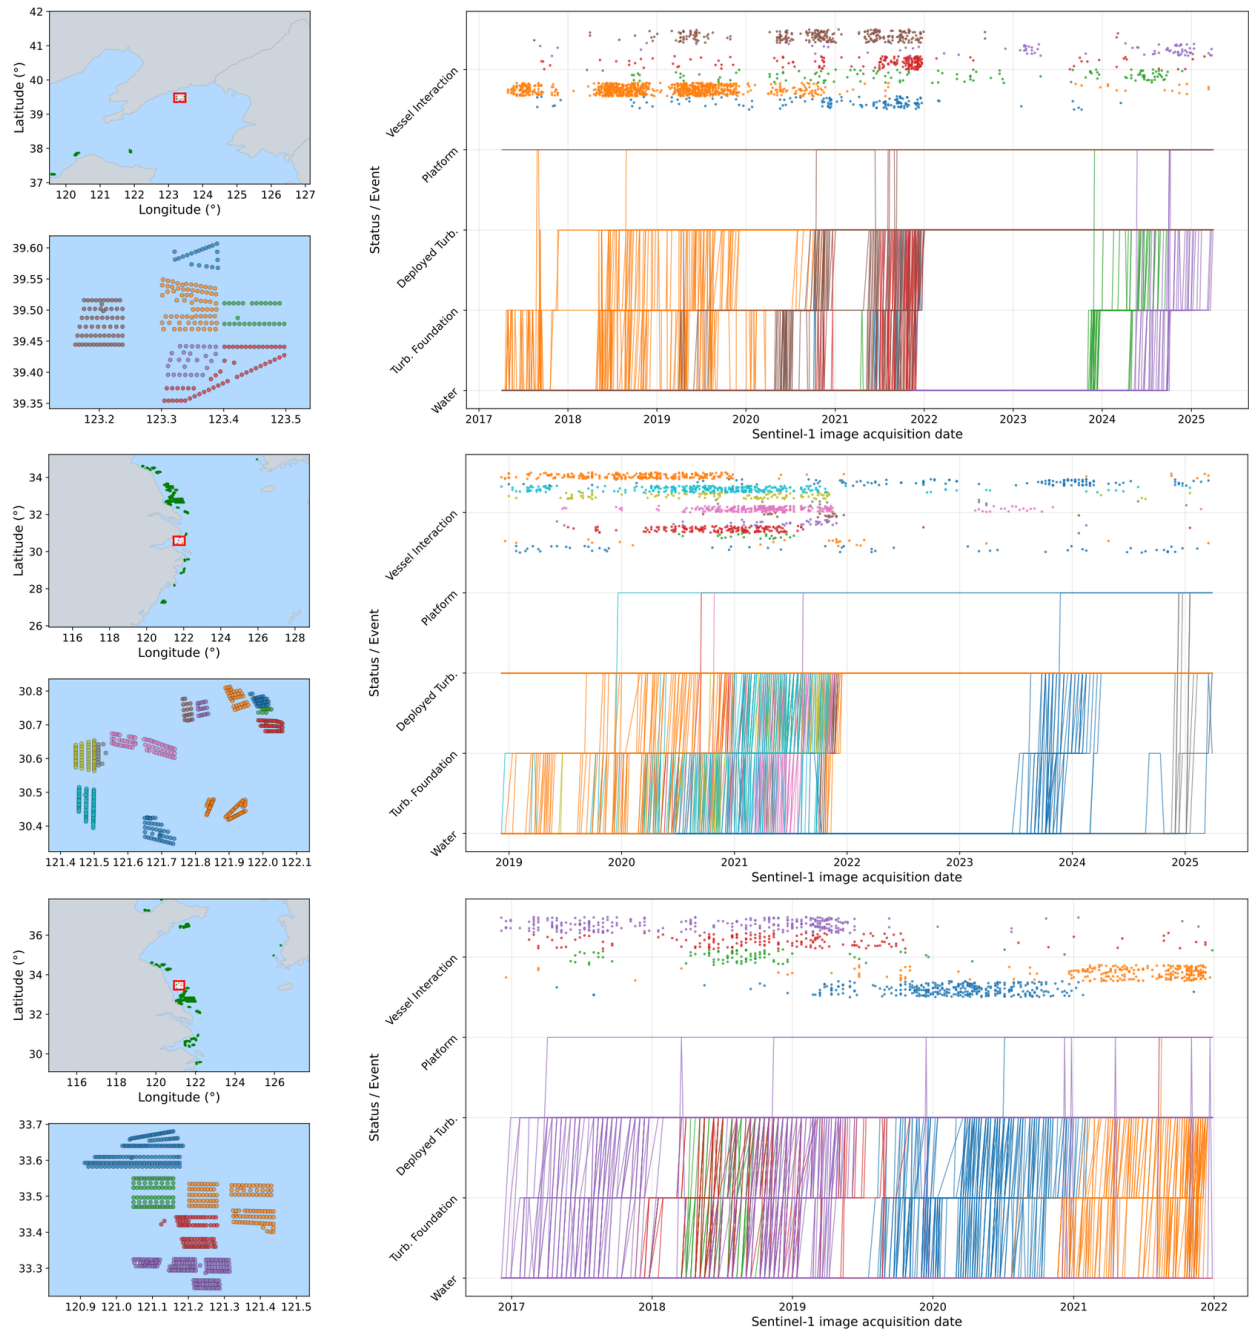


Figure 14: Example time series of three offshore wind farm clusters in the Chinese EEZ. Each time series is a coloured line corresponding to the location markers in the close-up overview map. Vessel related labels are markers in the time series with a random jitter for better visualization.




Figure 15: Distribution of offshore wind turbine deployment duration, derived as elapsed time measured from the first occurrence of a "turbine foundation" label to the first occurrence of a "deployed turbine" label, separated by regions.

Overall, the classified time series reveal deployment patterns, regional construction dynamics, and operational events, demonstrating the analytical potential of high-density SAR-based time series for monitoring offshore wind infrastructure.

## 6 Discussion

Over the past decade, Earth Observation based mapping of offshore wind infrastructure has matured substantially, with robust methods now available for global-scale spatial mapping. The contribution of this work lies in introducing high-density temporal characterization of offshore wind infrastructure at global scale. In this sense, the focus shifts from asking "Where are turbines?" to asking "What is happening at each turbine location, and when?". By providing high density time series and baseline semantic labels, this study moves offshore wind infrastructure monitoring toward dynamic infrastructure intelligence rather than inventory mapping. However, the results also reveal that temporally dense monitoring remains fundamentally dependent on spatial detection. Bridging spatial localization and temporal analysis into a unified spatiotemporal framework represents a central challenge for upcoming research in offshore wind infrastructure monitoring from space.

A central challenge pointed out by this work is the dependency of dense time series analysis on prior spatial detection. The proposed 1D profile-based time series approach is highly effective once a location is known, but it requires initial spatial localization. CNN-based object detectors perform well in complex near-coastal and harbour environments, where heterogeneous background patterns demand high precision [Hoeser et al., 2022, Ding et al., 2024, Zhang et al., 2024, Liu et al., 2024, Song and Li, 2025, Liu et al., 2026]. In contrast, CFAR-based approaches are often more sensitive to generic offshore metallic structures [Wang et al., 2019, Wong et al., 2019, Paolo et al., 2024] and may therefore be better suited to detecting early-stage turbine foundations. The two approaches reflect a precision-recall trade-off, where CNN-based models tend to optimize precision in heterogeneous settings, whereas CFAR-based methods may increase recall for early construction signals but tend to be more prone to false alarms.

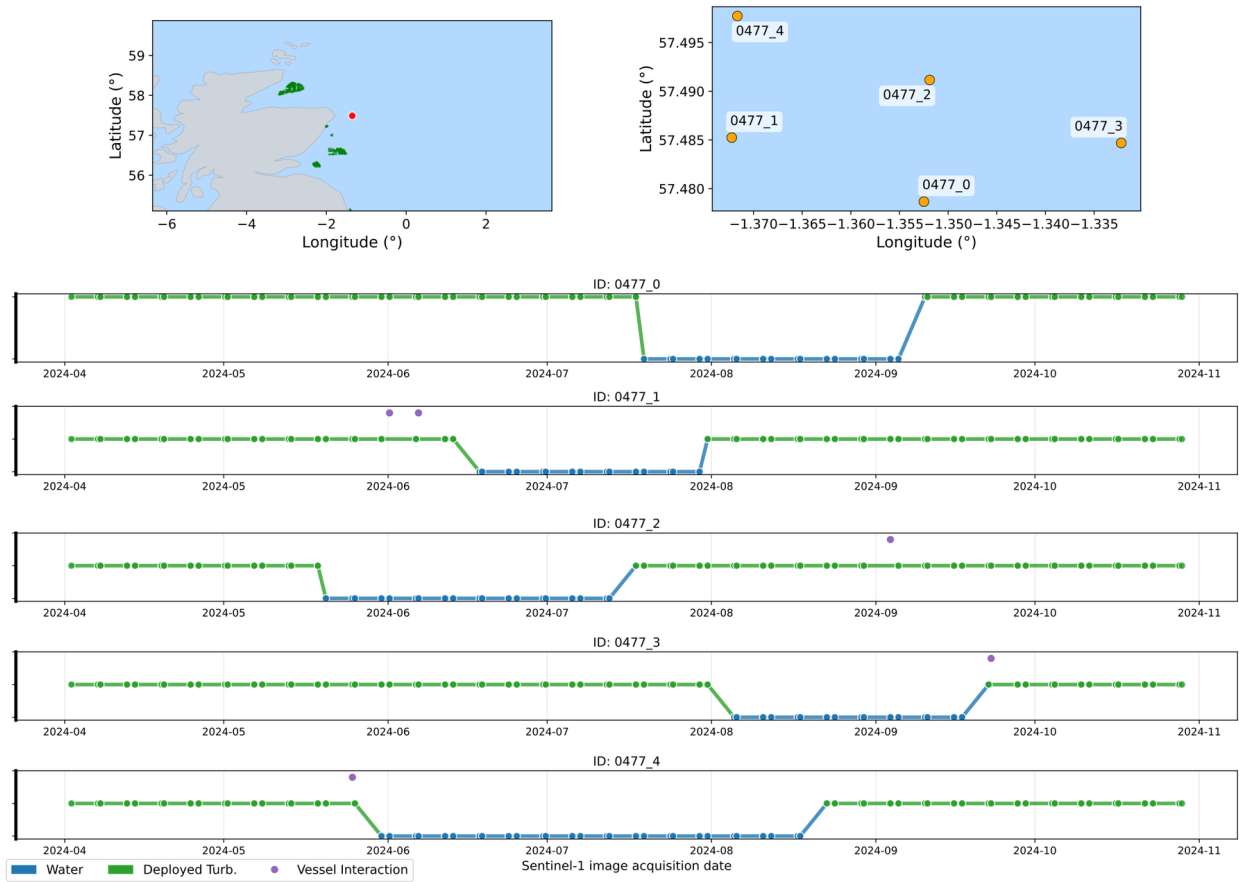


Figure 16: Visualization of the Hywind Scotland (HS) maintenance case study. HS is a floating turbine pilot project, in 2024 all 5 turbines have been towed away from their operational location due to planned maintenance. Reported start 2024-05-17 and end 2024-09-22 of the maintenance align well with classified time series patterns for all turbines.

This trade-off becomes critical when aiming for early onset detection of turbine foundations. Early detection without prior location knowledge remains a conceptually unresolved problem. Future research must address how spatial detection methods can be designed to capture the very first construction signals without overwhelming downstream time series analysis with false positives. Hybrid CNN-CFAR cascades optimized specifically for early-stage foundations, or the integration of temporal anomaly detection directly at the spatial detection stage are promising directions. Furthermore, moving away from quarterly median composites toward single-acquisition-based detection will be essential for approaching near-real-time (NRT) monitoring. While median compositing increases robustness, it suppresses transient signals and delays early detection. Recent work, such as Paolo et al. [2024], provides a starting point for such single-image detection frameworks, particularly when integrating SAR with AIS or optical data sources and cascading CFAR and CNN target classification.

With respect to time series classification, the rule-based classifier presented here was intentionally designed as a relatable, training-data-free baseline. However, the implemented rule-based classifier has already reached considerable complexity, including multi-stage decision logic, sequential smoothing, segment-level corrections, and global label distribution refinements. This complexity suggests that the decision boundaries between infrastructure states are not trivial and may exceed what threshold-based logic can robustly encode. While the baseline achieves strong event-level performance, future improvements may require supervised temporal deep learning to further improve sequence-wise performance.

The evaluation of time series classification highlights the performance gap between event-level accuracy and sequential consistency. Even when point-wise classification performance is strong, especially for turbine-related classes, the collapsed edit similarity score clearly shows that sequence-level agreement leaves room for improvement. This reflects the methodological limitations of hand-crafted rule-based classifiers, where sequential smoothing and refinements improves coherence but tends to risks over-regularization by enforcing deterministic transitions, or opposed to this miss

to refine overly volatile sequences to avoid over-regularization. Thus, future approaches may benefit from learning temporal transitions directly rather than relying on heuristic post-processing.

Performance differences between turbine-related classes and transient or platform-related classes reveal ambiguities in the 1D SAR backscatter profiles. Vessel activity, maintenance operations, and support platforms can produce similar scattering patterns, and backscatter alone may not always be sufficient to clearly separate these classes. Integration of auxiliary data sources such as AIS vessel tracking or optical imagery may therefore enhance classification robustness. While prior studies have demonstrated the benefits of such integration in 2D imagery analysis [Paolo et al., 2024], an open research question is whether similar gains can be achieved within a reduced 1D representation without sacrificing computational efficiency.

The representation of 2D SAR imagery as 1D backscatter profiles is a design choice balancing efficiency and feature depth. This compression enables global time series analysis, with 14.8 million events stored in approximately 5 GB, while preserving horizontally directed scattering signatures. However, this representation assumes that most discriminative information lies in the range direction. Rather than reverting directly to full 2D representations, a promising intermediate step would be to incorporate an additional vertical 1D profile derived along the azimuth direction. Such an extension would approximately double data set size while remaining computationally manageable, and may improve discrimination of mooring vessels and subtle differences between turbine foundation and deployed states, where azimuthal Doppler effects become relevant. This would preserve the initial motivation of compact, analysis-ready data while enriching feature representation.

Finally, although the dataset achieves global spatial completeness, it contains uneven temporal density due to variations in Sentinel-1 revisit frequency. European regions benefit from overlapping orbits and higher data take frequency, whereas parts of the Chinese EEZ typically follow a 12-day revisit interval. Due to this fact, deployment duration estimates and early detection timestamps are influenced by sampling frequency. Future analyses should therefore include acquisition density at a local scale, in order to provide event time ranges rather than assigning single timestamps. Defining milestone events as temporal windows bounded by neighbouring acquisitions may provide more robust interregional comparisons and uncertainty-aware deployment metrics.

Taken together, this work demonstrates that spatial mapping of offshore wind infrastructure has reached a high degree of maturity, while fine grained temporal characterization is emerging as the new frontier. The presented dataset and baseline classification shift offshore wind infrastructure monitoring toward time series based event analysis. At the same time, spatial detection remains the bottleneck for early monitoring. The next generation of Earth Observation based offshore wind infrastructure monitoring will likely require detection approaches optimized for early onset signals and trained spatiotemporal models integrating detection and temporal event classification. Bridging spatial detection and temporal characterization into a unified spatiotemporal monitoring framework will be essential for achieving NRT, globally scalable monitoring of offshore wind infrastructure.

## 7 Conclusion

In this study we present an approach and dataset for Earth Observation based offshore wind infrastructure event classification from high-temporal-resolution time series. The dataset contains open, analysis-ready, and globally distributed Sentinel-1 synthetic aperture radar (SAR) time series spanning 2016Q1–2025Q1, centred on 14,840,637 1D backscatter profiles compiled into 15,606 location-specific time series. These profiles provide a compact representation of global scale offshore wind infrastructure temporal dynamics, that preserves key SAR signatures while remaining accessible in storage and compute requirements.

For a spatial foundation, we update a global two-stage detection workflow and provide 15,606 offshore wind infrastructure locations for 2025Q1, with high detection performance across three geographically diverse validation sites. To classify the 1D SAR backscatter profile time series at each of the detected locations, we introduce a rule-based classifier that assigns event-level semantic labels describing deployment and operational states, and we release a benchmark dataset of 553 expert-annotated time series containing 328,657 event labels to support reproducible evaluation and method comparison. The baseline classifier achieves strong performance, particularly for turbine-related events (transition from water to turbine foundations, and transition from turbine foundations to deployed turbines). Example analyses reveal differences in regional patterns concerning deployment duration with China showing significantly lower durations compared to the European Union and the United Kingdom. On a local scale we also show sensitivity to operational dynamics such as maintenance-related towing events for floating offshore wind turbines.

Beyond the dataset contribution, the results highlight key directions for future work. Early-onset detection of wind infrastructure construction remains challenging, motivating research toward single-acquisition and near-real-time detection strategies and hybrid approaches that balance recall for early foundations with precision in complex coastal

environments. For time series interpretation, the complexity required for a handcrafted rule-based classifier indicates that learned sequential models and multi-source data fusion (e.g., AIS and optical) are promising future research areas.

Overall, the released corpus establishes a foundation for future advancements in offshore wind infrastructure monitoring from space by combining global coverage, high temporal resolution, analysis-ready time series, baseline predictions, and expert annotations as benchmark data set. It enables in-depth investigation of offshore wind deployment and operation at scale and supports the development, evaluation, and benchmarking of temporal methods for independent monitoring of this rapidly expanding critical infrastructure.

## Data Availability

The dataset is publicly available on Zenodo at <https://doi.org/10.5281/zenodo.5933966> [Hoeser et al., 2026].

## Abbreviations

AIS	Automatic Identification System
AUC	Area Under the Curve
CFAR	Constant False Alarm Rate
CNN	Convolutional Neural Network
CSV	Comma Separated Value
DLR	German Aerospace Center (Deutsches Zentrum für Luft- und Raumfahrt)
EEZ	Exclusive Economic Zone
ESA	European Space Agency
EO	Earth Observation
FN	False Negative
FP	False Positive
GDAL	Geospatial Data Abstraction Library
GEE	Google Earth Engine
GRD	Ground Range Detected
GTI	GDAL Tile Index
HPC	High Performance Computing
HS	Hywind Scotland
IW	Interferometric Wide
NRT	Near Real-Time
NSB	North Sea Basin
OWT	Offshore Wind Turbine
SAR	Synthetic Aperture Radar
TP	True Positive
TS	Time Series
VH	Vertical transmit, Horizontal receive
YOLO	You Only Look Once

## Author contributions

TH and CK designed the study. TH labelled the ground truth data; developed and implemented the code for data processing, visualization, and evaluation; and prepared the original manuscript, including figures. CK and FB supervised the study, gave suggestions for figures, and repeatedly commented on and discussed the manuscript.

## Competing interests

The contact author has declared that none of the authors has any competing interests.

## Acknowledgements

The authors gratefully acknowledge the Copernicus program of the European Space Agency (ESA) for providing free access to Sentinel-1 data. We also thank DLR's Terrabyte team for their dedicated efforts in platform operations, which enabled the deep learning experiments, synthetic training data generation, and large-scale inference.

## Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

## Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

We acknowledge the use of Grammarly v1.2.124.1571, which includes a generative AI assistant for English language editing. All AI-generated text suggestions have undergone rigorous revision by the authors.

## References

- Bundesrepublik Deutschland. Gesetz zur entwicklung und förderung der windenergie auf see (windseeg). BGBl. I 2016, S. 2258, 2016. URL <https://www.gesetze-im-internet.de/windseeg/>. Zuletzt geändert durch Art. 2 G v. 27.03.2024 (BGBl. I Nr. 107 v. 27.03.2024), rückwirkend zum 01.01.2023; § 1 Abs. 1. last access: 12 February 2026.
- UK Government. Net zero strategy: Build back greener. [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/1033990/net-zero-strategy-beis.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1033990/net-zero-strategy-beis.pdf), 2021. last access: 12 February 2026.
- European Commission. An eu strategy to harness the potential of offshore renewable energy for a climate neutral future. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:52020DC0741>, 2020. last access: 12 February 2026.
- United Nations, Department of Global Communications. Un climate summit 2025, 24 september 2025. Video, UN Web TV, 2025. URL <https://webtv.un.org/en/asset/k12/k12ebey0bu>. Broadcast by the Department of Global Communications; Subjects: Climate Change; Sustainable Development Goals; Environment. last access: 12 February 2026.
- Dominique Roddier, Christian Cermelli, Alexia Aubault, and Alla Weinstein. Windfloat: A floating foundation for offshore wind turbines. *Journal of Renewable and Sustainable Energy*, 2(3):033104, 06 2010. ISSN 1941-7012. doi:10.1063/1.3435339. URL <https://doi.org/10.1063/1.3435339>.
- Nuno Bento and Margarida Fontes. Emergence of floating offshore wind energy: Technology and industry. *Renewable and Sustainable Energy Reviews*, 99:66–82, 2019. ISSN 1364-0321. doi:<https://doi.org/10.1016/j.rser.2018.09.035>. URL <https://www.sciencedirect.com/science/article/pii/S1364032118306841>.
- Thorsten Hoeser and Claudia Kuenzer. Global dynamics of the offshore wind energy sector monitored with sentinel-1: Turbine count, installed capacity and site specifications. *International Journal of Applied Earth Observation and Geoinformation*, 112:102957, 2022a. ISSN 1569-8432. doi:10.1016/j.jag.2022.102957. URL <https://www.sciencedirect.com/science/article/pii/S1569843222001522>.
- Brian A. Wong, Christian Thomas, and Patrick Halpin. Automating offshore infrastructure extractions using synthetic aperture radar & google earth engine. *Remote Sensing of Environment*, 233:111412, 2019. ISSN 0034-4257. doi:<https://doi.org/10.1016/j.rse.2019.111412>. URL <https://www.sciencedirect.com/science/article/pii/S0034425719304316>.
- Ting Zhang, Bo Tian, Dhritiraj Sengupta, Lei Zhang, and Yali Si. Global offshore wind turbine dataset. *Scientific Data*, 8(1):191, 2021. ISSN 2052-4463. doi:10.1038/s41597-021-00982-z. URL [10.1038/s41597-021-00982-z](https://www.nature.com/scientificdata/).
- Thorsten Hoeser, Stefanie Feuerstein, and Claudia Kuenzer. Deepowt: a global offshore wind turbine data set derived with deep learning from sentinel-1 data. *Earth System Science Data*, 14(9):4251–4270, 2022. doi:10.5194/essd-14-4251-2022. URL <https://essd.copernicus.org/articles/14/4251/2022/>.
- Fernando S. Paolo, David Kroodsma, Jennifer Raynor, Tim Hochberg, Pete Davis, Jesse Cleary, Luca Marsaglia, Sara Orofino, Christian Thomas, and Patrick Halpin. Satellite mapping reveals extensive industrial activity at sea. *Nature*, 625(7993):85–91, 2024. ISSN 1476-4687. doi:10.1038/s41586-023-06825-8. URL <https://doi.org/10.1038/s41586-023-06825-8>.
- Kechao Wang, Wu Xiao, Tingting He, and Maoxin Zhang. Remote sensing unveils the explosive growth of global offshore wind turbines. *Renewable and Sustainable Energy Reviews*, 191:114186, 2024a. ISSN 1364-0321. doi:10.1016/j.rser.2023.114186. URL <https://www.sciencedirect.com/science/article/pii/S1364032123010444>.

- Shuai Zhang, Fangxiong Wang, Yingzi Hou, Junfu Wang, and Jianke Guo. Global offshore wind turbine detection: a combined application of deep learning and google earth engine. *International Journal of Remote Sensing*, 45(18): 6601–6623, 2024. doi:10.1080/01431161.2024.2391587.
- Qiannan Ding, Bo Tian, Chunpeng Chen, Yuekai Hu, and Xue Li. Identifying the spatio-temporal distribution characteristics of offshore wind turbines in china from sentinel-1 imagery using deep learning. *GIScience & Remote Sensing*, 61(1):2407389, 2024. doi:10.1080/15481603.2024.2407389.
- Tingting He, Yihua Hu, Fashuai Li, Yuwei Chen, Maoxin Zhang, Qiming Zheng, Yukan Jin, and He Ren. Mapping land- and offshore-based wind turbines in china in 2023 with sentinel-2 satellite data. *Renewable and Sustainable Energy Reviews*, 214:115566, 2025. ISSN 1364-0321. doi:10.1016/j.rser.2025.115566. URL <https://www.sciencedirect.com/science/article/pii/S1364032125002394>.
- Yasen Fei, Yongnian Gao, Hongyuan Gu, Yongqi Sun, and Yanjun Tian. Yolov5\_cdb: A global wind turbine detection framework integrating cbam and dbscan. *Remote Sensing*, 17(8), 2025. ISSN 2072-4292. doi:10.3390/rs17081322. URL <https://www.mdpi.com/2072-4292/17/8/1322>.
- Longxing Liu, Mengquan Wu, Gang Chen, Yijie Tang, Xiaodong Song, Min Zou, Yandong Xu, Xin Zhang, Shikuan Wang, Jinyi Lv, and Longxiao Zheng. Satellite observations and deep learning unveil the rapid expansion of offshore wind turbines in china. *Resources, Conservation and Recycling*, 227:108706, 2026. ISSN 0921-3449. doi:<https://doi.org/10.1016/j.resconrec.2025.108706>. URL <https://www.sciencedirect.com/science/article/pii/S092134492500583X>.
- Thorsten Hoerer, Felix Bachofer, and Claudia Kuenzer. Deepowt v3.25.1: Dense sentinel-1 time series for deployment and operational dynamics, February 2026. URL <https://doi.org/10.5281/zenodo.18735421>.
- Yongxue Liu, Chao Sun, Yuhao Yang, Minxi Zhou, Wenfeng Zhan, and Wangyu Cheng. Automatic extraction of offshore platforms using time-series landsat-8 operational land imager data. *Remote Sensing of Environment*, 175: 73–91, 2016. ISSN 0034-4257. doi:<https://doi.org/10.1016/j.rse.2015.12.047>. URL <https://www.sciencedirect.com/science/article/pii/S0034425715302716>.
- Wenxuan Xu, Yongxue Liu, Wei Wu, Yanzhu Dong, Wanyun Lu, Yongchao Liu, Bingxue Zhao, Huiting Li, and Renfei Yang. Proliferation of offshore wind farms in the north sea and surrounding waters revealed by satellite image time series. *Renewable and Sustainable Energy Reviews*, 133:110167, 2020. ISSN 1364-0321. doi:<https://doi.org/10.1016/j.rser.2020.110167>. URL <https://www.sciencedirect.com/science/article/pii/S1364032120304585>.
- Jing Zhang, Qi Wang, and Fenzhen Su. Automatic extraction of offshore platforms in single sar images based on a dual-step-modified model. *Sensors*, 19(2), 2019. ISSN 1424-8220. doi:10.3390/s19020231. URL <https://www.mdpi.com/1424-8220/19/2/231>.
- Qi Wang, Jing Zhang, and Fenzhen Su. Offshore platform extraction using radarsat-2 sar imagery: A two-parameter cfar method based on maximum entropy. *Entropy*, 21(6), 2019. ISSN 1099-4300. doi:10.3390/e21060556. URL <https://www.mdpi.com/1099-4300/21/6/556>.
- Xike Song and Ziyang Li. Seasonally robust offshore wind turbine detection in sentinel-2 imagery using imaging geometry-aware deep learning. *Remote Sensing*, 17(14), 2025. ISSN 2072-4292. doi:10.3390/rs17142482. URL <https://www.mdpi.com/2072-4292/17/14/2482>.
- Fangxiong Wang, Shuai Zhang, Yingzi Hou, and Junfu Wang. Extraction of offshore wind turbines in china by combining multispectral and sar image data. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 17:9266–9281, 2024b. doi:10.1109/JSTARS.2024.3392695.
- Christopher Jung and Dirk Schindler. The properties of the global offshore wind turbine fleet. *Renewable and Sustainable Energy Reviews*, 186:113667, 2023. ISSN 1364-0321. doi:10.1016/j.rser.2023.113667. URL <https://www.sciencedirect.com/science/article/pii/S1364032123005245>.
- Christopher Jung, Leon Sander, and Dirk Schindler. Future global offshore wind energy under climate change and advanced wind turbine technology. *Energy Conversion and Management*, 321:119075, 2024. ISSN 0196-8904. doi:10.1016/j.enconman.2024.119075. URL <https://www.sciencedirect.com/science/article/pii/S0196890424010161>.
- Susan Gourvenec, Fraser Sturt, Emily Reid, and Federico Trigos. Global assessment of historical, current and forecast ocean energy infrastructure: Implications for marine space planning, sustainable design and end-of-engineered-life management. *Renewable and Sustainable Energy Reviews*, 154:111794, 2022. ISSN 1364-0321. doi:<https://doi.org/10.1016/j.rser.2021.111794>. URL <https://www.sciencedirect.com/science/article/pii/S1364032121010637>.



- Inbar Schwartz-Belkin and Michelle E. Portman. A review of geospatial technologies for improving marine spatial planning: Challenges and opportunities. *Ocean & Coastal Management*, 231:106280, 2023. ISSN 0964-5691. doi:<https://doi.org/10.1016/j.ocecoaman.2022.106280>. URL <https://www.sciencedirect.com/science/article/pii/S0964569122002563>.
- Yi-Chien Lai, Chi-Yeung Choi, Kisup Lee, In-Ki Kwon, Chia-Hsiang Lin, Luke Gibson, and Wei-Yea Chen. Endangered black-faced spoonbills alter migration across the yellow sea due to offshore wind farms. *Ecology*, 106(1):e4485, 2025. doi:[10.1002/ecy.4485](https://doi.org/10.1002/ecy.4485). URL <https://esajournals.onlinelibrary.wiley.com/doi/abs/10.1002/ecy.4485>.
- Ramon Torres, Paul Snoeij, Dirk Geudtner, David Bibby, Malcolm Davidson, Evert Attema, Pierre Potin, Björn Rommen, Nicolas Floury, Mike Brown, Ignacio Navas Traver, Patrick Deghaye, Berthyl Duesmann, Betlem Rosich, Nuno Miranda, Claudio Bruno, Michelangelo L'Abbate, Renato Croci, Andrea Pietropaolo, Markus Huchler, and Friedhelm Rostan. Gmes sentinel-1 mission. *Remote Sensing of Environment*, 120:9–24, 2012. ISSN 0034-4257. doi:<https://doi.org/10.1016/j.rse.2011.05.028>. URL <https://www.sciencedirect.com/science/article/pii/S0034425712000600>. The Sentinel Missions - New Opportunities for Science.
- European Space Agency (ESA). Sentinel-1b in-flight anomaly summary report. <https://sentinels.copernicus.eu/documents/247904/4819394/Sentinel-1B+In-Flight+Anomaly+Summary+Report.pdf>, n.d. last access: 12 February 2026.
- Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 202:18–27, 2017. ISSN 0034-4257. doi:[10.1016/j.rse.2017.06.031](https://doi.org/10.1016/j.rse.2017.06.031). URL <https://www.sciencedirect.com/science/article/pii/S0034425717302900>. Big Remotely Sensed Data: tools, applications and experiences.
- Google Earth Engine Team. Sentinel-1 algorithms: Sentinel-1 preprocessing. <https://developers.google.com/earth-engine/guides/sentinel1?hl=en#sentinel-1-preprocessing>, 2026. Google Earth Engine Developer Guide. last access: 12 February 2026.
- Ao Wang, Hui Chen, and Lihao Liu. Yolov10: Real-time end-to-end object detection. *arXiv preprint arXiv:2405.14458*, 2024c.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6):1137–1149, 2017. doi:[10.1109/TPAMI.2016.2577031](https://doi.org/10.1109/TPAMI.2016.2577031).
- Thorsten Hoeser and Claudia Kuenzer. Synteo: Synthetic dataset generation for earth observation and deep learning – demonstrated for offshore wind farm detection. *ISPRS Journal of Photogrammetry and Remote Sensing*, 189:163–184, 2022b. ISSN 0924-2716. doi:[10.1016/j.isprsjprs.2022.04.029](https://doi.org/10.1016/j.isprsjprs.2022.04.029). URL <https://www.sciencedirect.com/science/article/pii/S0924271622001319>.
- DLR terrabyte. terrabyte. an innovative platform for exploring the earth. <https://docs.terrabyte.lrz.de/>, 2026. last access: 12 February 2026.
- Flanders Marine Institute. Union of the esri country shapefile and the exclusive economic zones (version 4). <https://www.marineregions.org/>, 2024. last access: 12 February 2026.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi:[10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2). URL <https://doi.org/10.1038/s41592-019-0686-2>.
- Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions and reversals. In *Soviet Physics Doklady*, volume 10, pages 707–710, 1966.
- Global Maritime. Heavy maintenance campaign completed on hywind scotland floating offshore wind farm. <https://www.globalmaritime.com/news/heavy-maintenance-campaign-completed-on-hywind-scotland>, October 2024. News article. last access: 12 February 2026.
- Longxing Liu, Mengquan Wu, Jie Zhao, Lei Bing, Longxiao Zheng, Shaopeng Luan, Yunfei Mao, Mingyue Xue, Jiayan Liu, and Bowen Liu. Deep learning-based monitoring of offshore wind turbines in shandong sea of china and their location analysis. *Journal of Cleaner Production*, 434:140415, 2024. ISSN 0959-6526. doi:[10.1016/j.jclepro.2023.140415](https://doi.org/10.1016/j.jclepro.2023.140415). URL <https://www.sciencedirect.com/science/article/pii/S0959652623045730>.