



Deep learning-based object detection of offshore platforms on Sentinel-1 imagery and the impact of synthetic training data

Robin Spanier ^a, Thorsten Hoeser ^a and Claudia Kuenzer ^{a,b}

^aGerman Remote Sensing Data Center, Earth Observation Center, EOC of the German Aerospace Center, DLR, Wessling, Germany; ^bInstitute for Geography and Geology, Department of Remote Sensing, University of Würzburg, Würzburg, Germany

ABSTRACT

The recent and ongoing expansion of marine infrastructure, including offshore wind farms, oil and gas platforms, artificial islands, and aquaculture facilities, highlights the need for effective monitoring systems. The development of robust models for offshore infrastructure detection relies on comprehensive, balanced datasets, but falls short when samples are scarce, particularly for underrepresented object classes, shapes, and sizes. By training deep learning-based YOLOv10 object detection models with a combination of synthetic and real Sentinel-1 satellite imagery acquired in the fourth quarter of 2023 from four regions (Caspian Sea, South China Sea, Gulf of Guinea, and Coast of Brazil), this study investigates the use of synthetic training data to enhance model performance. We evaluated this approach by applying the model to detect offshore platforms in three unseen regions (Gulf of Mexico, North Sea, Persian Gulf) and thereby assess geographic transferability. This region-holdout evaluation demonstrated that the model generalizes beyond the training areas. In total, 3529 offshore platforms were detected, including 411 in the North Sea, 1519 in the Gulf of Mexico, and 1593 in the Persian Gulf. The model achieved an F1 score of 0.85, which improved to 0.90 upon incorporating synthetic data. We analysed how synthetic data enhances the representation of unbalanced classes and overall model performance, taking a first step towards globally transferable detection of offshore infrastructure. This study underscores the importance of balanced datasets and highlights synthetic data generation as an effective strategy to address common challenges in remote sensing, demonstrating the potential of deep learning for scalable, global offshore infrastructure monitoring.

ARTICLE HISTORY

Received 6 November 2025
Accepted 23 December 2025

KEYWORDS

Earth observation; object detection; oil rigs; offshore; offshore platforms; remote sensing; sentinel-1; YOLOv10

CONTACT Robin Spanier robin.spanier@dlr.de German Remote Sensing Data Center, Earth Observation Center, EOC of the German Aerospace Center, DLR, Münchener Str. 20, Weßling 82234, Germany

Supplemental data for this article can be accessed online at <https://doi.org/10.1080/01431161.2026.2612908>

© 2026 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

1. Introduction

In the past few decades, the expansion of maritime persistent infrastructure has accelerated, particularly in relation to energy infrastructure such as oil and gas platforms and wind farms (Bugnot et al. 2021; Liu et al. 2019). Among these, offshore platforms play a central role in supporting a wide range of marine operations. Platforms serve as a base for offshore activities, provide work and living facilities, and are essential for the exploitation of marine resources, research, production, rescue operations, maritime logistics, storage, and transport (X. Zhang et al. 2025). They occur in various structural configurations adapted to water depth and environmental conditions, designed to operate for over 25 years in harsh environments (Figure 1) (Sadeghi 2007). This maritime infrastructure contributes significantly to economic development and drive social and technological progress in regions that have previously been poor in resources. In 2020, offshore oil and gas extraction reached a peak at USD 988 billion in gross output value (OECD 2025). Stakeholders such as governments, regulatory authorities, and operators seek to harmonize international installation and decommissioning practices of this infrastructure (OECD, Organisation for Economic Co-operation and Development (2025); Potter and Pearson 2023). The aim is to effectively manage the dynamic development of the maritime space while taking into account the associated ecological, economic, and social interests, for example, in the areas of ecosystem services, fisheries, logistics, shipping, the recreation industry, and nature conservation (Ma et al. 2024; Posen et al. 2020; J. Wang et al. 2014; Williamson et al. 2016). Reliable data on offshore infrastructure is essential for maritime domain awareness and facilitates environmental management, decision making, and policy to address challenges across their life cycle (OECD 2025; Potter and Pearson 2023; OECD 2021). Some installations transmit positional data via AIS to enhance monitoring (March et al. 2021) and regulations such as the EU's Environmental Impact Assessment (EIA) mandate reporting of new facilities (European Commission and Directorate-General for Environment 2022). However, such regulations are regionally

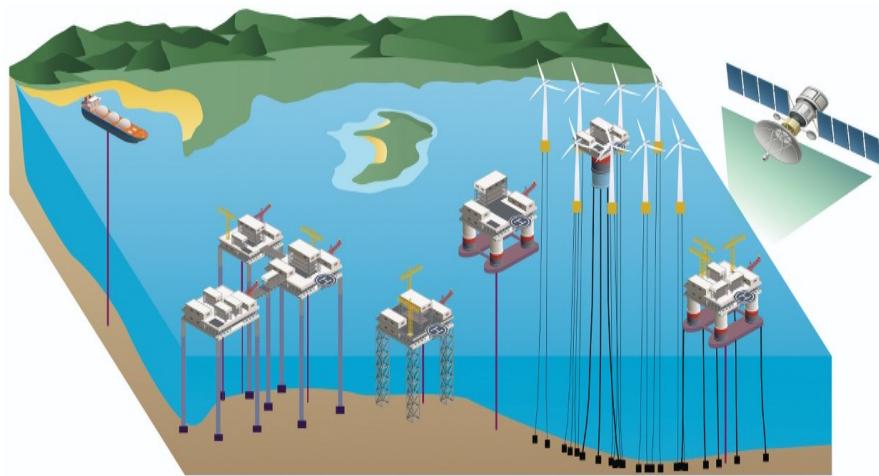


Figure 1. Schematic overview of offshore energy infrastructure, in particular, different platform types. Fixed platforms are anchored to the seabed using piles or jackets and are common in shallow waters, while semi-permanent structures are partially mobile and used in deeper settings. Figure adjusted from (Spanier and Kuenzer 2024).

limited and inconsistently enforced, and existing data are often incomplete or restricted due to confidentiality and security concerns (Liu et al. 2016; Sun et al. 2020; Virdin et al. 2021; Wong, Thomas, and Halpin 2019). Given these challenges, the development of reliable, scalable methods for accurately and quickly detecting the location and characteristics, such as size, type, or design, of offshore infrastructure is a key concern.

Over the past 50 years, tens of millions of satellite images have been collected and archived, opening up new possibilities for environmental and infrastructure monitoring. Advances in this field provide an overview of the condition and dynamics of large-scale Earth surfaces (Baumhoer et al. 2023; Esch et al. 2022; Hoeser and Kuenzer 2022a). Satellite remote sensing offers regular and effective monitoring without the need for direct access, particularly for regions that are difficult to access, such as offshore areas (UNCTAD 2021; Yang et al. 2013). With a short repeat cycle, global coverage, and the continuous availability of long-term measurements, a wide range of data is collected by various sensors and made publicly available in a constantly growing, updated, and nearly global database. In addition, earth observation (EO) has benefited significantly from advances in cloud computing and data science. Platforms such as Google Earth Engine (GEE) make it possible to efficiently analyse large amounts of remote sensing data without having to store or download it locally (Amani et al. 2020; Gorelick et al. 2017).

The use of remote sensing data to investigate offshore infrastructure condition and development via automatic detection has become increasingly prevalent in recent years. While early work primarily made use of the high thermal and optical contrast of gas flares in infrared data (Casadio, Arino, and Minchella 2012), synthetic aperture radar (SAR)-based methods have now become the predominant approach, due to their ability to operate under all weather conditions and strong sensitivity to metallic sea-surface targets such as ships, platforms, and wind turbines.

Conventional approaches often employ constant false alarm rate (CFAR) detectors or threshold-based methods combined with geometric or multiscale filters to extract potential offshore targets (Cheng et al. 2013; Marino, Velotto, and Nunziata 2017; J. Zhang, Wang, and Su 2019). Multitemporal composites are used to effectively isolate static backscatter signatures of man-made structures and suppress moving targets, both using purely SAR-based methods (Wong, Thomas, and Halpin 2019) and in combination with optical data (W. Xu et al. 2020). Furthermore, both optical-based (Liu et al. 2019) and SAR-based studies (T. Zhang et al. 2021) demonstrate that offshore infrastructure can be automatically extracted on a global scale.

The development of robust deep learning models for SAR data continues to remain challenging due to small object sizes and the limited availability of high-quality training data. Several studies demonstrate that targeted architectural or preprocessing adjustments can counteract these challenges. These include hybrid or multi-stage pipelines that combine convolutional feature extraction, morphological or sparse coding-based preprocessing, and context and shape models to map small or finely structured targets more reliably (Aghaei, Akbarizadeh, and Kosarian 2022a; Ghara, Fatemeh, and Akbarizadeh 2022; Samadi, Akbarizadeh, and Kaabi 2019; Sharifzadeh, Akbarizadeh, and Seifi Kavian 2019; Taibi, Akbarizadeh, and Farshidi 2019; Tirandaz, Akbarizadeh, and Kaabi 2020). In addition, hybrid classification chains demonstrate the potential of sequentially linking empirical and deep learning methods to reduce false alarms in complex maritime scenes

(Blondeau-Patissier et al. 2023). Furthermore, lightweight, speckle-robust backbones and cross-domain architecture designs improve transferability with limited ground truth (Aghaei, Akbarizadeh, and Kosarian 2022b; Davari, Akbarizadeh, and Mashhour 2022).

A related approach to overcoming limited training data is SyntEO, a framework for the generation of synthetic training samples (Hoeser and Kuenzer 2022b). Using this, DeepOWT successfully extracts offshore windfarms and turbines from Sentinel-1 (S1) on a global scale (Hoeser and Kuenzer 2022a; Hoeser, Feuerstein, and Kuenzer 2022). Recent work is pursuing multimodal or architecturally enhanced approaches, for example, through the fusion of SAR and multispectral data or the use of pseudo-Siamese models, which better leverage complementary modalities (Ma et al. 2024; Paolo et al. 2024; Qiu et al. 2024; X. Zhang et al. 2025). Beyond this, machine learning-based SAR applications, such as the derivation of hub-height-related wind fields from S1 (Montera et al. 2022), underscore the broad potential of data-driven methods for maritime analysis.

Existing research often addresses detecting offshore infrastructure in regionally limited study areas, using multi-temporal composites to remove moving targets and extract static signatures of offshore infrastructures. Nevertheless, key challenges remain, including the limited availability of annotated training data, the limited generalizability across regions, and the high computational effort required to process large remote sensing datasets. Because of this, deep learning methods are frequently used in this field (Spanier and Kuenzer 2024).

The motivation for this study is based on the importance of an optimized training strategy for deep learning-based object detection in EO data. The study evaluates the effects of these optimizations using established YOLO base models, which are known for their efficiency and high detection performance. The central contributions of this work can be divided into four main aspects:

- The incorporation of synthetic data which significantly improves detection performance, especially for underrepresented targets.
- A region holdout strategy enabling model generalization and geographical transferability of the models beyond the training areas.
- The introduction of a fully trainable offshore platform detection framework based on the GEE and the Google Cloud Platform (GCP), which delivers robust results for hotspot regions of maritime infrastructure.
- A scalable approach that can be applied to the entire S1 archive, enabling global detection and spatiotemporal analysis.

2. Materials and methods

2.1. Study regions

The study focused on seven regions known for their abundance of offshore energy infrastructure, particularly oil, gas, and wind: the North Sea (NS), Persian Gulf (PG), Gulf of Mexico (GoM), South China Sea (SCS), Caspian Sea (CS), Gulf of Guinea (GoG), and the Coast of Brazil (CoB) (Figure 2). The latter four regions were used to train the object detection (OD) model, while NS, PG, and GoM served for inference. This region-holdout setup enabled evaluation of model generalization and geographic transferability beyond the training areas. The boundaries of the inference regions (NS, PG, GoM) followed the

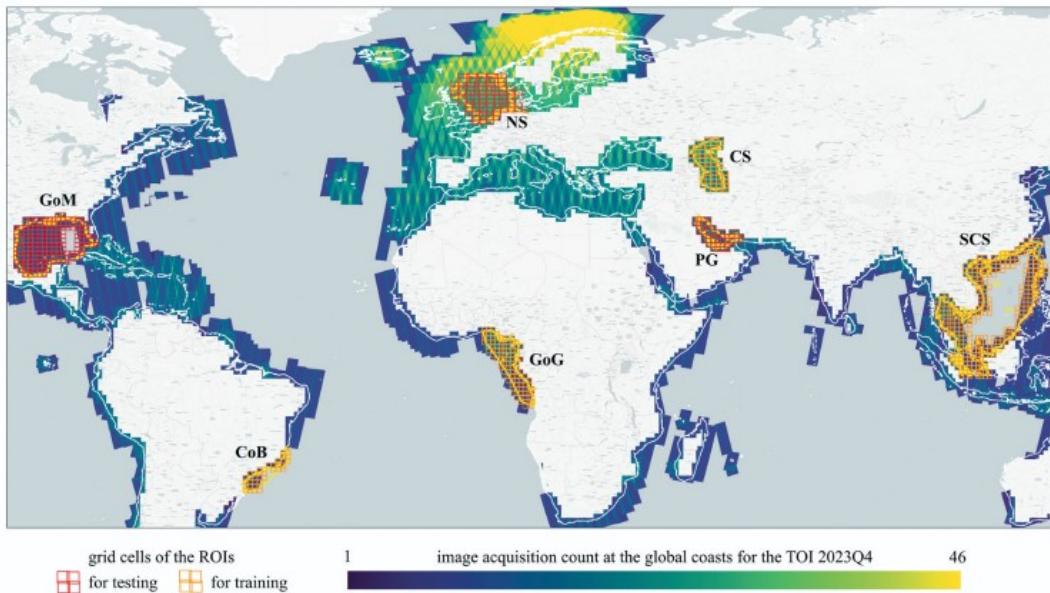


Figure 2. Overview of the seven study regions: South China Sea (SCS), the Caspian Sea (CS), the Gulf of Guinea (GoG), and the coast of Brazil (CoB) used for training (orange) and North Sea (NS), the Gulf of Mexico (GoM), and the Persian Gulf (PG) used for testing (red). The boundaries of the inference regions are based on the boundaries defined by the International hydrographic Organisation (IHO 2018), while the training regions were defined manually (yellow). In addition, the map displays the spatial distribution and image acquisition count of Sentinel-1 SAR scenes that overlap with the global coasts in 2023Q4. It demonstrates the significant regional variation in acquisition density (higher density in Europe, lower density in America and Asia) and the geographical separation between training and test data, which is used to evaluate the geographical transferability of the model.

official definitions of the International Hydrographic Organisation (Flanders Marine Institute 2018). While the training regions were informed by this data, they were ultimately manually delineated to focus primarily on the targeted coverage of areas of high platform density rather than formal regional delineation.

2.2. Image acquisition and preprocessing

SAR imagery from the European Space Agency (ESA) Copernicus S1 mission (ESA 2025a) was used due to its proven reliability for detecting metallic sea-surface targets, independence of weather, global coverage, high spatial and temporal resolution, and free access (El-Darymli et al. 2013; Spanier and Kuenzer 2024). The S1 constellation comprises two C-band SAR satellites, S1A (launched 2014) and S1B (2016), in sun-synchronous orbits with a 180° phase offset, providing a 6-day combined revisit. Due to the failure of S1B in 2021 and the selected time of interest (TOI), the fourth quarter of 2023 (2023Q4), this study used S1A only. The resulting delayed revisit rate of approximately 12 days is a temporary limitation which reduces the number of available SAR acquisitions for our 3-month median composite. This was taken into account when selecting the composite length, so that the effect remains minor for the 2023Q4 study period. However, a change in acquisition frequency may become more relevant for longer study periods and when combining different revisit cycles. Additionally, acquisition frequency varies regionally

with orbital overlap and mission priorities, resulting in denser coverage over Europe (ESA 2025a) (Figure 2).

We used the Ground Range Detected (GRD) Level-1 product provided by GEE (GEE, Google Earth Engine and ESA 2025), which processes the backscatter coefficient (σ^0) in decibels (dB). Each scene was preprocessed in the following steps: orbit data correction, removal of edge artefacts and thermal noise, radiometric calibration, and terrain correction. For our analysis, we used VH polarization data in Interferometric Wide (IW) Swath mode, with a spatial resolution of approximately 10 m. The S1 archive was queried for all scenes acquired during the TOI 2023Q4, and for the regions of interest (ROI), the SCS, CS, GoG, and CoB (for model training) and the NS, GoM, and PG (for testing). To enable efficient data handling and export on GEE, we applied a grid over the ROIs (Figure 2). The grid was based on the tiling system of S2 that subdivides Earth into a predefined set of tiles (UTM/WGS84), using a 100 km step. However, each tile has a surface area of 110×110 km to provide a large overlap with the neighbouring tiles (ESA 2021).

These image tiles were preprocessed, using the GEE API, to create deep-learning-ready inputs and to remove transient or mobile objects. First, all image tiles were stacked and reduced to a median composite. This temporal aggregation effectively removed moving or temporary objects and thus reduced potential false positives even before model inference (Wong, Thomas, and Halpin 2019). A total of 928 median-composite tiles across all study sites were generated and exported (Table 1).

To minimize data volume, the median composites were converted from 16-bit floating point (backscatter amplitude in dB) to 8-bit integers by clipping and linearly mapping the backscatter signal range of -40 dB to 0 dB onto values 0 to 255. The resulting tiles were then further tiled into 640×640 pixel chips with a 20% overlap, input sizes compatible with the DL model, and to prevent boundary artefacts.

On SAR images, offshore platforms appear as bright backscatter clusters in front of the darker sea with a low backscatter coefficient (Figure 3). The backscatter signal of a platform is larger than its actual size. For example, the signal of a 120×70 m platform can be approximately 290×230 m on a radar image. The signal is caused by several backscatter effects, including direct reflections from the highest structures of the platform (layover), double reflections between the platform's vertical structures and the sea surface, and triple reflections (or even higher order) between the platform and the surrounding sea surface (Jackson and Apel 2004; Marino, Velotto, and Nunziata 2017). Figure 3 illustrates the complexity of platform types and their backscatter signals on S1 SAR images. As mentioned above, they differ in their functionality, shape, and size, and

Table 1. Summary of Sentinel-1 satellite imagery downloaded and used for each study region in 2023Q4. The table demonstrates the significant regional variation in data availability, particularly with regard to acquisition density.

Regions of interest	approx. area (km^2)	S1 scenes	median comp. tiles
North Sea	570,000	852	94
Persian Gulf	250,000	291	51
Gulf of Mexico	1,340,000	554	199
South China Sea	2,050,000	946	363
Caspian Sea	400,000	457	69
Gulf of Guinea	820,000	354	109
Coast of Brazil	150,000	57	44
Sum Σ		3,511	928

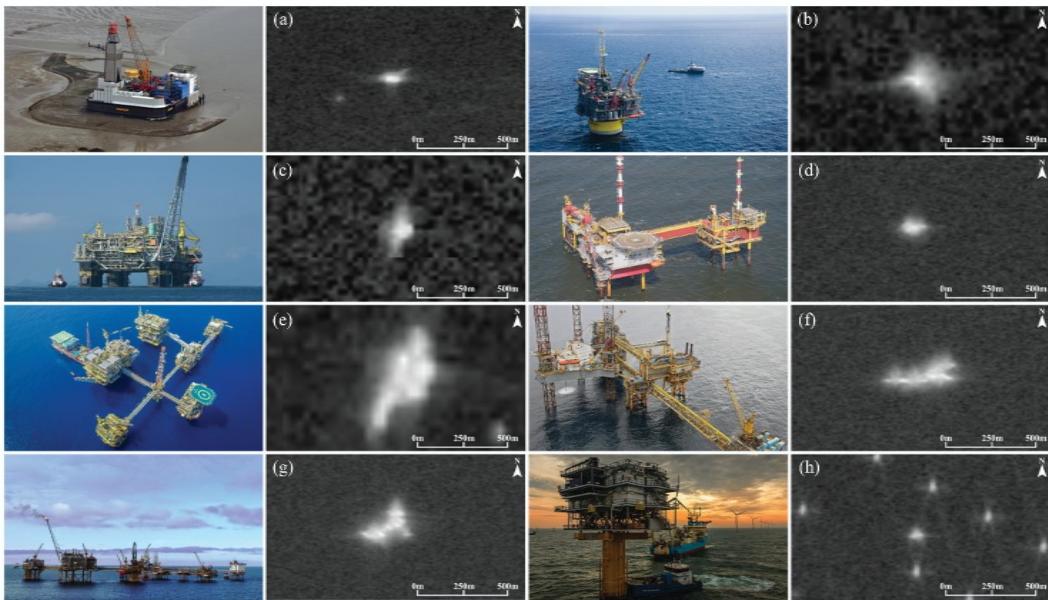


Figure 3. Examples of the structural and radar image diversity of offshore platforms in Sentinel-1 SAR data. This figure shows optical aerial images (left) and the corresponding S1 images (right) for (a–c) individual offshore platforms, (d–g) complex platform clusters, and (h) an offshore wind farm with turbines and a substation. The figure demonstrates the high variability in shape, extent, and backscatter signature, especially in platform clusters. Aerial imagery retrieved from commons.Wikimedia.org, credit to (Komen 2008; Petrobras 2009; Raudøy 2016; Roletschek 2011).

whether they are single platforms or multiple interconnected platforms, as well as in their resulting backscatter signal. We have assigned a dedicated class for wind turbines as off-targets in our dataset. Detecting them is an important step, as their backscatter signals are similar in strength to those of platforms (Figure 3h) and can lead to misclassifications. Due to their distinct radar signature, their detection and separation from platforms is possible (Hoeser, Feuerstein, and Kuenzer 2022; Wong, Thomas, and Halpin 2019; W. Xu et al. 2020; T. Zhang et al. 2021).

2.3. *Ground truth data*

For model training and testing, a comprehensive and consistent ground truth (GT) dataset was established. As explained in the previous chapter, some of the platforms are complex structures consisting of several individual platforms (referring hereinafter as platform clusters) (Figure 3d–g). Platform clusters comprise all occurrences consisting of more than one individual platform; in other words, contiguous platform complexes (Figure 3d–g). Multiple closely spaced backscatter signals from individual platforms result in stronger and blended noise in the S1 radar image, which manifests itself as a coherent complex backscatter signal. Individual platforms can hardly be derived from this, which is why we have labelled each complex as a single object. The goal of this study was to identify platforms in general. We did not label these objects as platforms, but rather as platform clusters to test whether we could optimize our model performance by treating single platforms and platform clusters separately during training and then combining

them into a unified platform class for evaluation. We therefore defined three classes: *single platforms*, *platform clusters*, and *wind turbines* (off-targets). The classes were assigned based on their spatial footprint as seen in the S1 images, and not on the basis of their structural typology (as in Figure 1).

Label creation was performed manually by a trained annotator based on a defined protocol. The primary reference, S1 GRD imagery from 2023Q4, was supported by Sentinel-2 MSI imagery, high-resolution optical data from DigitalGlobe (Google Earth), and auxiliary datasets such as vector layers and web-based infrastructure databases (Table 2). On this basis, separate GeoJSON files were created for each ROI and class, where each GeoJSON represented precise segmentation masks for each class's respective backscatter signature. Single platforms were only assigned if there was a clearly isolated signature. Two or more spatially connected platforms (for example, via walkways or bridge structures) were classified as platform clusters. Typical borderline cases involved individual platforms located very close to each other, whose signatures partially merge visually in the SAR image. In such cases, high-resolution optical data were used to assess whether multiple signatures actually represent separate objects or functionally belong to a cluster (see Figure S1 in the supplementary materials).

To ensure label quality, all polygons were systematically compared with the additional data listed in Table 2. Cross-validation among multiple data sources ensured positional accuracy and minimized labelling errors, providing a reliable foundation for consistent ground truth generation. Bounding boxes (bboxes) were then derived from the final vector masks. The final GT dataset for all seven areas comprised more than 14,000 polygons (41% *single platform*, 4% *platform cluster*, 55% *wind turbine*) (Table 3).

To avoid data leakage and ensure independent evaluation, training, and test datasets of pairs of image chips and labels (referring hereinafter as image-label pairs) were kept

Table 2. Overview of auxiliary data used for the labelling process of the ground truth dataset.

Dataset	Year	Coverage	Source
S1 GRD SAR imagery	2023	globally	ESA (2025a)
S2 MSI imagery	2023	globally	ESA (2025b)
DigitalGlobe imagery	–	globally	
Global Energy Monitor	2025	globally	Global Energy Monitor (2025)
OpenSeaMap	2025	globally	OSM (2025b)
4C Offshore	2025	globally	(4C Offshore 2025)
North Sea Energy	2022	North Sea	North Sea Energy (2025)
OSPAR	2023	North Sea	(OSPAR 2025)
BOEM	2024	Gulf of Mexico	BOEM (2024)

Table 3. Overview of all ground truth labels collected per study region, by class type. The table shows regional differences in object distribution and the occurrence of classes. Compared to *single platforms* and *wind turbines*, the *platform cluster* class is notably underrepresented.

Study regions	<i>single platform</i>	<i>platform cluster</i>	<i>wind turbines</i>	sum Σ
North Sea	418	80	4,949	5,447
Persian Gulf	1,550	163		1,713
Gulf of Mexico	1,562	187		1,749
South China Sea	1,054	98	2,920	4,072
Caspian Sea	311	81		392
Gulf of Guinea	932	89		1,021
Coast of Brazil	33	3		36

strictly separate. All GT image-label pairs from the SCS, CS, GoG, and CoB were used for model training, while pairs from the NS, GoM, and PG served exclusively for testing.

The image-label pairs used for model training were randomly divided into 90% training and 10% validation subsets to track model performance during training. We ensured that image chips originating from the same image tile were not split between sets. The entire set of 8909 labels from NS, GoM, and PG constituted the fixed test dataset, which remained constant across all experiments to ensure comparability.

The entire training set comprised 5521 labels, with only 271 labels belonging to the *platform cluster* class (Table 3). This highlights the significant underrepresentation of this class, which, if not balanced, would disappear in the training signal due to the dominance of single platforms and wind turbines. To improve training stability and model generalization, we aimed for an approximately balanced distribution among the three classes. Since obtaining additional labels was not feasible due to the limited availability of real-world samples, two potential strategies remained: expanding the dataset with *platform clusters* from new regions or generating synthetic data to increase the number of image-label pairs.

2.4. Synthetic data generation

We wanted to generate synthetic data and integrate it into training to optimize and improve our training data through better generalization. The effects of this are a key focus of this study. With the synthetic data, we wanted to increase the number of image-label pairs, especially for the underrepresented class of *platform clusters*. Instead of oversampling the class by duplication of samples together with data augmentation (such as clipping, rotation, scaling, and mosaicking), we expected significantly better detection performance by generating images that are completely new to the model. Therefore, the goal was to create synthetic data that closely resembled real-world, hand-labelled S1 samples in appearance and diversity. For that task, we used the SyntEO framework (Hoeser and Kuenzer 2022b), a comprehensive toolkit for creating custom synthetic EO data.

SyntEO composes a remote sensing scene from 2D discrete objects and background entities. Sensor-specific textures are then applied to simulate radar signatures, while the geometric composition provides bbox annotations. We extended the original SyntEO toolset to include routines for generating platform cluster geometries, in addition to single platforms and wind turbines. Two aspects were crucial in the entire process of generating synthetic data: full automation and randomization to produce thousands of diverse image-label pairs simultaneously, and minimizing the risk of the model learning artificial or repetitive patterns. To ensure natural variability, the appearance and configuration of each synthetic object differed across scenes. The overall workflow is summarized in Figure 4, which illustrates the main processing stages adopted in this study. For a detailed technical description of SyntEO, refer to Hoeser and Kuenzer (Hoeser and Kuenzer 2022b).

- (1) Input data consisted of randomly selected S1 scenes in WGS84 projection (EPSG:4326) from around the world. We defined the image extent in pixels and ensured a representative mix of open-sea, coastal, and land scenes. Images containing interfering objects, such as existing platforms, wind turbines, aquaculture,

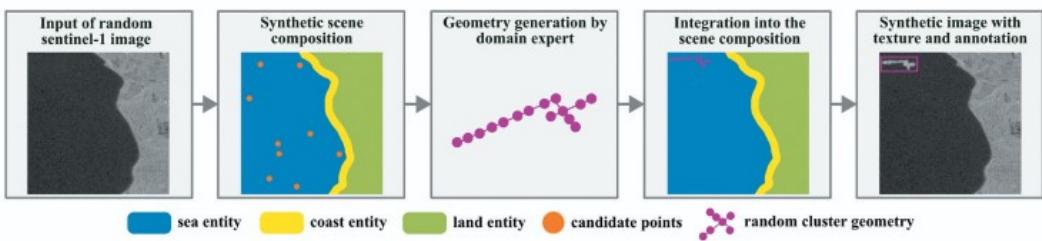


Figure 4. Simplified schematic overview of using the SyntEO framework to generate synthetic training data of platform clusters. Based on real Sentinel-1 scenes, suitable maritime regions are identified, procedural cluster geometries are generated, and clusters are integrated into the scene using texture-based backscatter kernels. The figure illustrates the modular structure of the approach and the separation of geometry and texture modelling.

or lighthouses, were removed using a thresholding process to obtain 'empty' marine area entities suitable for synthetic insertion.

- (2) Images were processed with SyntEO to derive coherent 2D entities for sea, coast, and land. The land entity was determined from the OSM Land Polygon Vector dataset (OSM 2025a). A 1 km buffer from land defined the coast entity, and the remaining area defined the sea entity. Within the sea entity, we created a grid of potential object anchor points (candidate entities). These points were randomized per scene, creating a 'canvas' for placing synthetic objects.
- (3) To realistically reproduce platform clusters, we first analysed real S1 radar signatures to identify common geometric features. Each cluster could be effectively described as a set of lines and points arranged along those lines. We generated random cluster geometries defined by meta parameter values such as the number and length of lines, connection angles, and point spacing (see Table S1 for the complete set of parameter ranges used).
- (4) Each geometry was randomly rotated and placed in the scene at one of the candidate points of our grid.
- (5) Radar backscatter was simulated by applying KDE2D kernels, which create an overlay of Gaussian-like point scatter functions, to the points of the previously generated cluster geometries. The kernel parameters were randomly sampled from predefined ranges, such as size and intensity, for each point to ensure high variance in the resulting textures (see Table S1 for the complete set of parameter ranges used). A separate additive or multiplicative speckle model was not implemented, meaning that the generated textures are based solely on the overlaid kernels. The textured geometries were then inserted into the prepared S1 scenes, resulting in complete synthetic images. Bbox annotations were automatically derived from object extent and position. Representative examples comparing synthetic and real-world scenes are shown in Figure 5.

2.4.1. Datasets

Five training datasets were compiled based on combinations of real and synthetic image-label pairs. The composition of these datasets is summarized in Table 4. The 'base' training dataset consisted of GT from SCS, CS, GoG, and CoB, with *single platforms* and *platform*

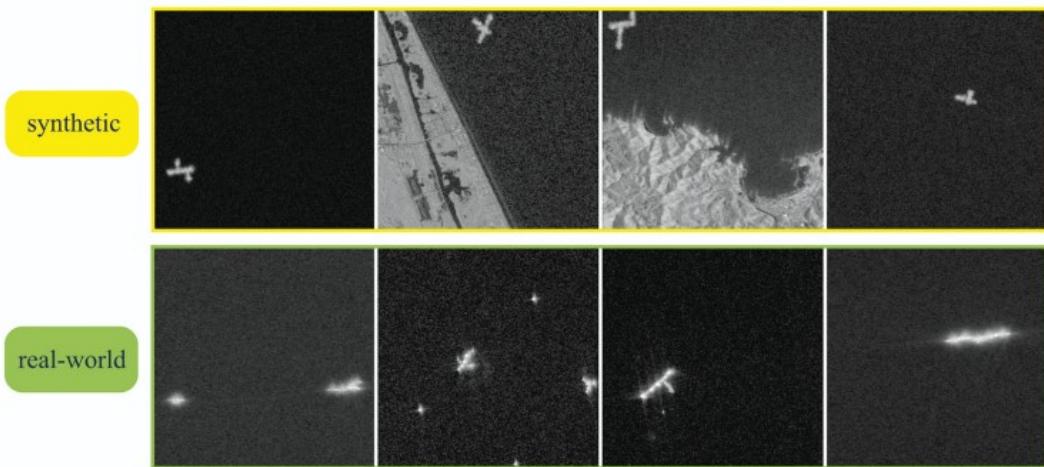


Figure 5. Comparison between synthetic scenes of platform cluster occurrences made by the SyntEO workflow (top) and scenes from the real world (bottom).

Table 4. Composition of training and evaluation datasets with the number of bounding boxes per object class. The different training strategies and the fixed evaluation dataset are shown. An asterisk (*) indicates synthetically generated training samples. The table illustrates the addition of synthetic samples in different proportional steps to address class imbalances and investigate the scalability of training.

Dataset	single platform	platform cluster	wind turbines	Sum Σ
Base (SCS, CS, GoG, CoB)	2,601		2,920	5,521
Split (SCS, CS, GoG, CoB)	2,330	271	2,920	5,521
Cluster-enriched	2,330	271 + 2206*	2,920	7,727
Fully balanced	2,333 + 2677*	271 + 4729*	2,920 + 2080*	15,000
Synthetic-only	5,000*	5,000*	5,000*	15,000
Evaluation (NS, PG, GoM)	3,530	430	4,949	8,909

clusters as a combined class, while the ‘split’ dataset split these two classes. To systematically assess the impact of synthetic training data, three additional datasets were created. In the cluster-enriched dataset, synthetic image-label-pairs were added to the underrepresented *platform cluster* class, resulting in a more balanced distribution across all three classes (2,330–2,920 pairs each). The fully balanced dataset was generated by augmenting all three classes with synthetic examples to reach 5000 pairs per class. Finally, a synthetic-only dataset was produced to test model performance when trained exclusively on synthetic data.

For evaluation, a fixed test dataset containing real-world GT labels from the NS, PG, and GoM was used across all experiments to ensure consistency and comparability.

2.5. Offshore platform detection pipeline

We employed version 10 of the YOLO OD algorithm family, released in May 2024 (A. Wang et al. 2024), to evaluate the effects of optimized training strategies for deep learning-based object detection in EO data and detecting offshore platforms in our test regions. YOLO is established and widely used in EO (Hoeser and Kuenzer 2020; Hoeser, Bachofer, and Kuenzer 2020). We chose YOLOv10, a one-stage detector, because it offers a more

efficient ratio of accuracy to latency and computational cost compared to other established detectors, especially two-stage detectors such as Faster R-CNN. For resource-efficient processing of large-scale EO data in GCP, such models are crucial for reducing runtime and hardware requirements. Compared to its predecessor, YOLOv8, version 10 uses a slimmer, NMS-free training architecture that supports stable detection of small, variable objects. Furthermore, while YOLOv8 offers broader multi-task functionality, this is not required for our purely object detection-based application.

The YOLOv10 model takes the S1 scenes as input, processes them through a convolutional neural network (CNN), and outputs the probabilities of the predefined classes: *single platform*, *platform cluster*, and *wind turbine*.

An overview of the detection workflow is shown in Figure 6, and the fundamentals are described in the following four stages.

- (1) Model training: Various training scenarios were developed and tested under three distinct training scenarios: training solely with real data, training solely with synthetic data, and training with real and synthetic data combined (Table 4). The central research question was how synthetic data would influence model performance and geographic generalization. We hypothesized that the targeted inclusion of synthetic samples would lead to a measurable improvement in accuracy and robustness. To prepare the trained models for Inference, the trained YOLOv10 artefacts were exported as a TorchScript model, packaged using the Torch-Model-Archiver (Meta AI 2025), and uploaded as a deployable artefact.
- (2) Inference: Model management and inference were carried out on the GCP. We utilized Vertex AI, GCP's machine-learning service, for model deployment and inference. Model endpoints were created on Vertex AI, and cloud CPUs were

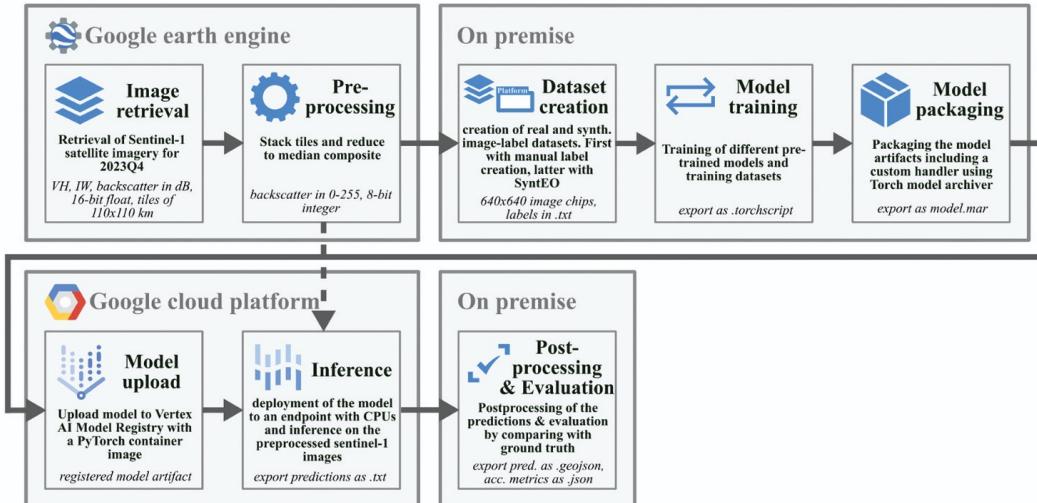


Figure 6. Workflow of the deep learning-based offshore platform detection pipeline. The workflow includes the key processing steps from S1 data acquisition and preprocessing, through the creation of real and synthetic training datasets, model training and packaging, to inference, post-processing, and evaluation. The workflow illustrates the separation between cloud-based and on-premise components as well as the scalability of the pipeline for large-scale applications.

provisioned to perform inference over the preprocessed image chips. Importantly, all satellite data remained on the GCP; only prediction outputs were exported as text files containing detection metadata and geolocation. This design minimized data transfer and reduced cost by limiting billing to compute time only.

- (3) Postprocessing and evaluation: After inference, the predictions from all image chips were merged into a single dataset. Bboxes were reprojected to EPSG:4326 (WGS84) to enable spatial operations, and each detection was assigned a unique identifier. Predictions below the defined confidence threshold of 0.5 were discarded. To suppress noise, detections where all pixels were below 150 (≈ -16.5 dB) were removed. This threshold was derived from the analysis of radar backscatter characteristics of verified platforms in the 8-bit S1 composites. Figure S2 in the supplementary materials illustrates the effect of this threshold on predictions on a sample S1 scene. To reduce duplicates caused by tile overlap, overlapping detections were grouped using an intersection-over-union (IoU) threshold of 0.2. For each group, the most reliable prediction was selected based on class consensus and confidence score. The threshold values used for confidence, pixel value, and IoU grouping were informed by an ablation analysis in which we evaluated the effects of different threshold combinations on precision, recall, and F1 score (Supplementary Material Table S2). The cleaned detection set was exported as GeoJSON for further spatial analysis.

3. Results

To verify the effectiveness of the methods developed in this study, various training experiments were conducted and tested, resulting in a model that efficiently and accurately detects offshore platforms on S1 radar images and generalizes reliably to previously unseen regions, thus being geographically transferable and scalable. First, the results of the training experiments and the influence of synthetic training data are described before finally the results of the best model are presented.

3.1. Training experiment results and synthetic data influence

For the experiments, we used pretrained YOLOv10n and s, and v11n model variants (as provided by A. Wang et al. 2024) that were finetuned with the respective training datasets. After 50 epochs, training was finished by an early stopping mechanism, tracking the validation performance. For testing, we compared model predictions against the GT of our fixed test dataset using a confidence threshold of 0.5, ensuring that only predictions with a probability of over 50% were considered. Further, predictions overlapping a GT object with a higher IoU value of 0.3 were counted as TPs, while lower overlaps were counted as FPs, and undetected GT objects as FNs. Due to the very small spatial extent of offshore platforms in Sentinel-1 data, an IoU threshold of 0.3 was used. In our tests, this resulted in an F1 score that was 0.03 higher than a threshold of 0.5. In the base experiment, IoU matching was performed with separate GT classes for *single platform* and *platform cluster*. In all other experiments, both the model predictions and the corresponding GT labels were merged into a uniform *platform* class before evaluation. The IoU

distributions for single platforms and platform clusters (before merging) and the corresponding mean values are shown in Figure S3 in the supplement materials. The results of all experiments are summarized in [Table 5](#).

In the first experiment, we employed the smallest YOLOv10 variant (10n, 2.3 million parameters), trained on the base dataset. Despite its compact size, the model achieved strong performance, correctly detecting 3537 of 3960 platforms (F1 = 0.82). Using the split dataset next allowed training of *single platforms* and *platform clusters* instead of them being combined. The intermediate results (before merging the *single platforms* and *platform clusters* into a uniform class) showed that the model achieved an F1 score of 0.85 for *single platforms* and 0.42 for *platform clusters*. This highlights the model's difficulties in recognizing *platform clusters*. However, these intermediate results still contain misclassifications between *single platforms* and *platform clusters*, which are eliminated in the final result by merging the classes. With an F1 score of 0.84, which is higher than in the first experiment, our assumption that the best recognition results for platforms can be achieved if the classes are first separated and only merged again after training and inference was confirmed. We therefore retained this approach for the following experiments.

Using YOLOv11n (2.6 million parameters) in the next experiment did not improve the results, while switching to larger weights (v10s, 7.2 million parameters) slightly increased the F1 score to 0.85. Although even larger models (m, l, x) might have yielded marginal gains, we decided against them to ensure scalability when performing inference on GCP or other compute environments.

The subsequent experiments evaluated the influence of synthetic data on model performance. Since the *platform cluster* class was severely underrepresented and achieved F1 scores below 0.48 in the previous models, we enriched the training data with over 2000

Table 5. Results of the training experiments on the fixed evaluation dataset. The table shows the impact of different training strategies and pre-trained YOLO models used on detection performance. It highlights the advantages of combining real and synthetic data, but also reveals the limitations of an exclusively synthetic training strategy.

Dataset	YOLO model	Class	GT	TP	FP	FN	Pr	Rc	F1
Base	10n (2.3M params)	single platform + platform cluster	3,960	3,537	1,115	423	0.76	0.89	0.82
Split	10n (2.3M params)	single platform	3,530	2,983	509	547	0.85	0.85	0.85
		platform cluster	430	173	223	257	0.44	0.40	0.42
		platform	3,960	3,319	569	641	0.85	0.84	0.84
Split	11n (2.6M params)	single platform	3,530	3,011	552	519	0.85	0.85	0.85
		platform cluster	430	274	479	156	0.36	0.64	0.46
		platform	3,960	3,475	841	485	0.81	0.88	0.84
Split	10s (7.2M params)	single platform	3,530	3,025	567	505	0.84	0.86	0.85
		platform cluster	430	210	242	220	0.46	0.49	0.48
		platform	3,960	3,387	657	573	0.84	0.86	0.85
Cluster-enriched	10s (7.2M params)	single platform	3,530	3,036	419	494	0.88	0.86	0.87
		platform cluster	430	240	63	190	0.79	0.56	0.65
		platform	3,960	3,412	346	548	0.91	0.86	0.88
Fully balanced	10s (7.2M params)	single platform	3,530	3,083	438	447	0.88	0.87	0.87
		platform cluster	430	255	95	175	0.73	0.59	0.65
		platform	3,960	3,523	348	437	0.91	0.89	0.90
Synthetic-only	10s (7.2M params)	single platform	3,530	11	54	3,519	0.17	0.00	0.01
		platform cluster	430	0	0	430	0.00	0.00	0.00
		platform	3,960	13	52	3,947	0.20	0.00	0.01

synthetic cluster samples to balance the class distribution. This adjustment is particularly evident in the intermediate results: for the *platform clusters*, precision improved significantly from 0.46 to 0.79 and the F1 score from 0.48 to 0.65, with 30 more TPs and 179 fewer FPs. This improvement carried over to the overall result for the unified *platform* class ($Pr = 0.91$, $Rc = 0.86$, and $F1 = 0.88$), the best result at this stage. To test whether a wider balance across all classes would further improve performance, the training data was expanded to 5000 samples per class with synthetic samples. This approach improved recall in the intermediate results of the *single platforms* and *platform clusters* while maintaining high precision, leading to our best overall performance of the unified *platform* class ($Pr = 0.91$, $Rc = 0.89$, $F1 = 0.90$). Finally, to determine the upper limit of the usability of the synthetic data we generated, the model was trained with 5000 synthetic samples per class exclusively. With this approach, we were unable to generalize the model to real images. The model achieved precision and recall close to zero. The model was apparently unable to learn or derive the features of the classes from synthetic data alone, which it needed in order to subsequently recognize the objects in the real data. We can assume that the model only learned synthetic specific features that could not be transferred to real scenes. Accordingly, we suspect a potential bottleneck in the data generated with SyntEO up to this point. Improving the designed textures could lead to significant detection performance (Hoeser and Kuenzer 2022b; Hoeser, Feuerstein, and Kuenzer 2022) and further increase performance when integrated with real-world data.

The experiments demonstrate both the effectiveness and limitations of synthetic data in training deep learning models for EO applications. Synthetic data significantly improved performance when combined with real-world data, particularly through class balancing. However, real-world examples remained indispensable.

3.2. *Detection results and geographic transferability*

The ‘fully balanced’ model achieved the best overall performance and was used to extract offshore platforms in the three unseen test regions of the NS, GoM, and PG. Figure 7 provides an overview of the detections of offshore platforms in these regions for 2023Q4 and shows the spatial distribution of this offshore infrastructure for these high-energy-producing regions. In total, the model detected 3529 offshore platforms, corresponding to a recall of 0.89 and an F1 score of 0.90. The detection performance varied slightly from region to region: 411 platforms were detected in the NS ($Rc = 0.83$), 1519 in the GoM ($Rc = 0.87$), and 1593 in the PG ($Rc = 0.93$). Furthermore, with an expected calibration error (ECE) of 0.12, the model demonstrates satisfactory calibration of the confidence intervals output, with a slight overconfidence occurring as expected, since offshore platforms comprise only a few effective pixels and speckle-related variability reduces signature stability.

Geographically, the detection maps show that the model was able to reliably generalize to previously unseen regions, proving its robust transferability beyond the training areas (CS, SCS, GoG, and CoB). The dense arrangement of platforms in the central NS, the northern GoM, and the agglomerations in the oil fields in the PG is captured with high spatial accuracy, and even smaller installations near the coast are reliably detected.

The model showed high sensitivity to different platform types, from single installations to complex platform clusters. These results demonstrate the model’s ability to adapt to varying regions with different environmental conditions, detection geometries, and

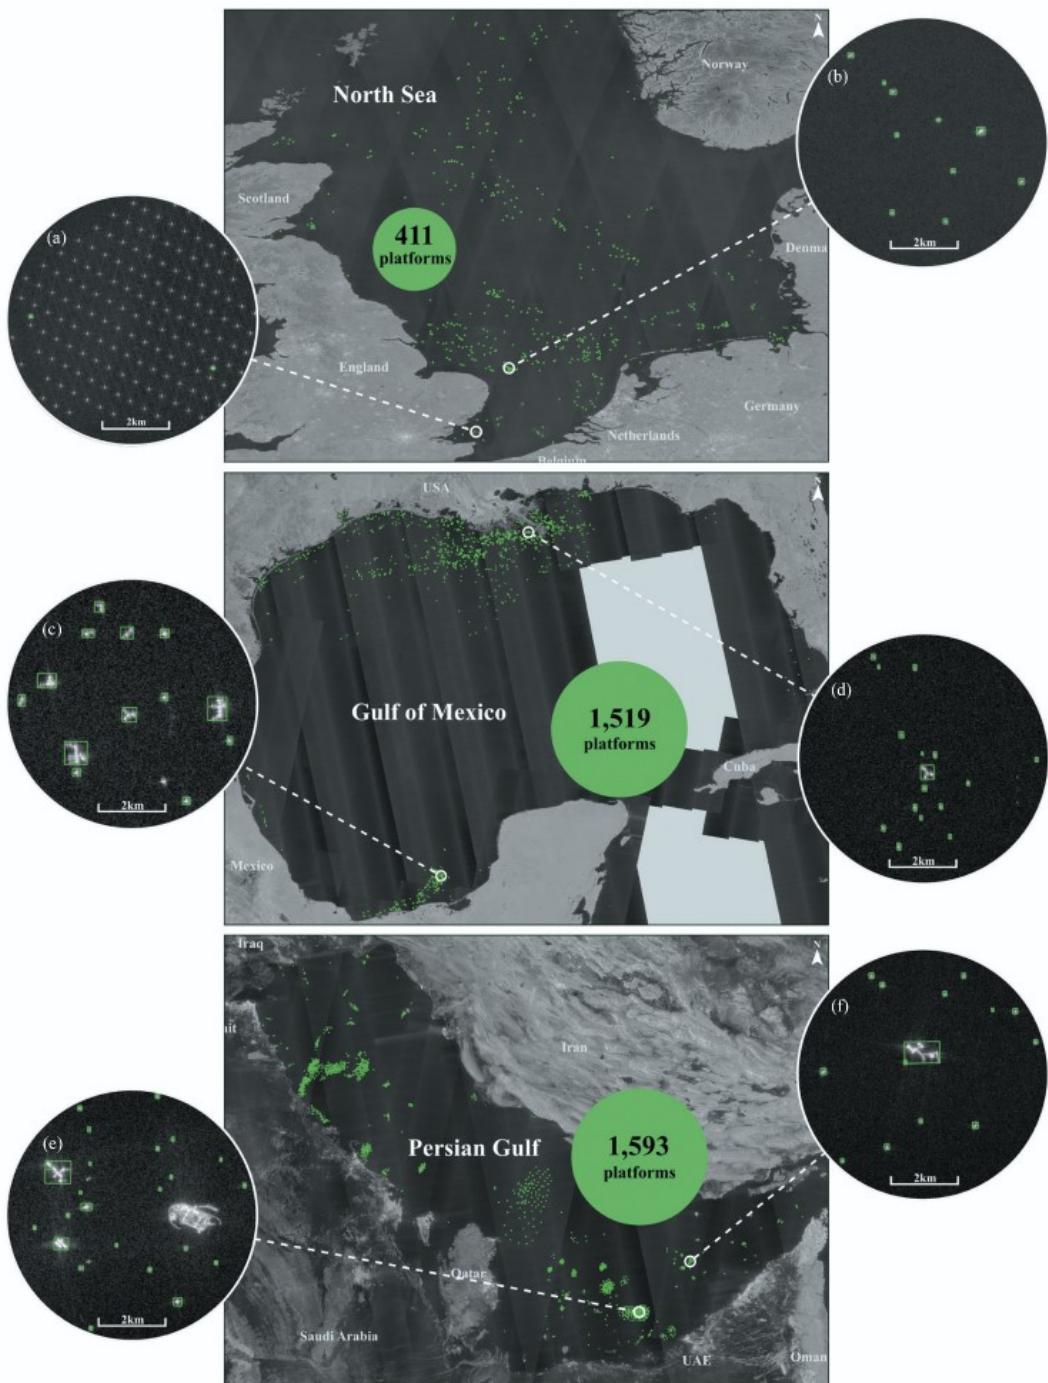


Figure 7. Offshore platform detection results in the three test regions North Sea, the Gulf of Mexico, and the Persian Gulf, for 2023Q4, based on the preprocessed S1 median composites and by the fully balanced model, which outperformed all other models tested in this study. The enlarged images show various challenging areas that are often crowded with different types of platforms (b-f). Wind turbines could be effectively separated from offshore platforms such as substations in a wind farm (a).

		predictions		
		platform	windturbine	background FN
ground truth	platform	3,523 / 3,960 89.0%	0.4% 15	10.7% 422
	windturbine	0.7% 36	4,783 / 4,949 96.6%	2.6% 130
background FP		312	80	

Figure 8. Confusion matrix of the postprocessed detection results of the final model on the fixed evaluation dataset. Correct classifications and misclassifications between offshore platforms, wind turbines, and background are shown. The matrix illustrates the high level of discrimination between platforms and wind turbines, as well as the remaining false positive and false negative detections.

structural complexities, confirming its potential for scalable global mapping of offshore infrastructure. Notably, it was possible to detect the off-target class of *wind turbines* with an F1 score of 0.97, which allowed them to be separated from platforms effectively. Distinguishing substations in offshore wind farms from turbines is a crucial step of particular importance for refining the detection of offshore wind farms (W. Xu et al. 2020; Z. Xu et al. 2022; T. Zhang et al. 2021).

Figure 8 shows the corresponding confusion matrix of the final model. It shows high accuracy for both the *platform* class and the *wind turbine* off-target class (89% for *platforms*, 96.6% for *wind turbines*). Misclassifications were low and are discussed in Section 4.

Overall, the results confirm that YOLO-based object detection on S1 radar images provides a reliable and transferable framework for the automatic identification of offshore platforms in various marine environments.

4. Discussion

In all three test regions, the North Sea, Gulf of Mexico, and Persian Gulf, 3523 of 3960 platforms (89%) were correctly identified, with 422 unidentified (FN) and 312 platforms predicted where none exist (FP). The model performed better on the off-target class of *wind turbines*: 4783 of 4949 turbines (96.6%) were correctly identified, with only 51 cross-class misclassifications. This result demonstrates that the model is capable of reliably distinguishing structurally similar objects on SAR data and underscores the robustness of YOLOv10-based feature extraction, despite this object class not being the focus of the

training. It further underscores that while the model generalizes reliably across all classes, the detection of platforms remains the more complex task.

FNs were associated with the detection of very small platforms, particularly in the Gulf of Mexico (Figure 9d), where numerous structures exist that are smaller than 100 m^2 . At a spatial resolution of 10 m, such targets approach or fall below the effective pixel scale of S1, making their reliable detection difficult. An important factor affecting FNs was the trade-offs made with the chosen confidence and IoU thresholds. Strict thresholds reduced the number of FPs, but unavoidably excluded some TPs with lower confidence. These are mainly very small individual platforms in coastal areas where calm sea conditions or shadowing effects further weaken the backscatter signal. Through iterative tuning, a balance was found that minimized overall error. The fact that higher FN rates occur predominantly where platforms are very small (especially in the Gulf of Mexico) is also reflected in Table S3 (supplementary materials), which provides a comprehensive comparison of regional environmental conditions. It also shows that low coverage density can contribute to higher FN values. At the same time, some environmental parameters, such as the significantly rougher sea conditions in the North Sea, deviate more strongly from the training regions and correlate with regionally distinct recall rates.

Misclassifications were mainly associated with objects that generate radar backscatter similar to that of platforms. These include, for example, buoys, lighthouses, and coastal reefs (Figure 9a-b). Including additional off-target classes in the training dataset could help to avoid such confusion for the model. However, the limited availability of real-world reference data would likely require supplementation with synthetic samples.

In a few cases in the NS, the model also incorrectly classified wind turbines under construction as platforms (Figure 9c). These unfinished turbines, which consist only of the foundation, are very similar to offshore platforms in both geometry and radar response.

Despite these challenges, the model performed particularly well in areas with dense and complex offshore infrastructure, such as the AKAL oil field in the southern GoM or the Upper Zakum field off the coast of the United Arab Emirates, where there are several platform clusters and artificial islands

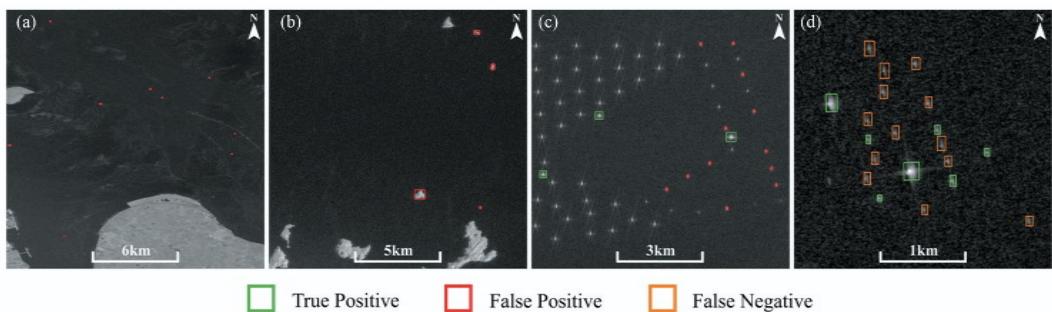


Figure 9. Demonstration of exemplary misclassifications of the model in the test regions. Misinterpretations of buoys and a lighthouse near Bremerhaven, Germany in the North Sea (a), reefs at the northern coast of Oman in the Persian Gulf (b), wind turbine foundations at the gode wind farm north of the Norderney island, Germany in the North Sea (c), and undetected small platforms to the southeast of New Orleans in the northern Gulf of Mexico (d). The figure demonstrates that misclassifications primarily occur with very small targets or with structurally similar off-target objects with similar backscatter patterns.

(Figure 7c, e). The accurate detection of such clustered targets demonstrates the model's ability to represent structural diversity and spatial context, capturing both individual installations and large platform complexes.

This study also demonstrates the critical importance of incorporating synthetic training data for the overall performance of the model. However, while synthetic samples significantly improved performance, especially for underrepresented classes, experiments with purely synthetic training data revealed distinct limitations. As stated in the previous chapter, we suspect a potential bottleneck in the data generated with SyntEO.

Further refinement of the synthetic training data will therefore be a key focus of future work. As a first step, we plan to systematically quantify the domain gap between real and synthetic S1 scenes. Comparative analyses of the distributions of backscatter values, texture-based descriptors, and feature embeddings are planned to more clearly determine whether the performance loss of the purely synthetically trained model is due to texture differences, geometric simplifications, or label structures. Further experiments will also be conducted to reduce FPs. These will investigate whether the addition of specific off-target classes (e.g. buoys or lighthouses) to the synthetic data pool can further reduce the number of FPs. To this end, we would first analyse these off-target objects on S1 median images and derive characteristic patterns in terms of size, shape, and backscatter intensity. On this basis, suitable geometries, textures and parameter values would be developed and integrated into an initial test dataset per class to evaluate the effect on FP rates.

In addition, future work will also consider alternative strategies for generating synthetic SAR data in order to further improve texture and noise modelling. These include physics-based SAR simulation approaches (e.g. Dong, Meng, and Guo 2023; He et al. 2024), which explicitly model electromagnetic scattering processes and potentially generate realistic backscatter patterns, but are currently not suitable for large-scale training datasets due to their high computational complexity and the need for detailed 3D geometries. Generative or domain-adaptive methods (e.g. Wu et al. 2024; Zeng et al. 2024) are also promising. These methods derive SAR textures from real examples using machine learning and are thus better able to replicate speckle-typical patterns, but they require extensive training data and careful model stabilization. Photogrammetric or 3D image-based methods (e.g. Jia, Yang, and Wu 2023; Luo et al. 2025) also enable the realistic reconstruction of complex offshore structures but are only suitable for global scaling to a limited extent due to the considerable effort required for data acquisition and processing. SyntEO was deliberately chosen for this study because it represents a lightweight, fully automatable, and highly scalable procedural approach that does not require complex EM simulation, extensive 3D models, or additional training processes, thus enabling the efficient generation of large amounts of synthetic training data.

Overall, the results show that YOLO-based object detection on S1 radar images is both robust and transferable to different marine environments, but sensitive to very small structures and class ambiguity. Addressing these remaining challenges through targeted dataset refinement and advanced synthetic data generation will be key to achieving consistent mapping of offshore infrastructure on a global scale.

5. Conclusion

In the fourth quarter of 2023, we detected a total of 3529 offshore platforms in three of the most important regions of the global offshore energy infrastructure: 411 in the North Sea, 1519 in the Gulf of Mexico, and 1593 in the Persian Gulf. The objects were detected by applying deep learning to ESA's S1 radar data. In our study, we demonstrated an approach that accomplishes the optimization of the training process of the deep learning-based object detector by integrating synthetic images with real EO data. Although synthetic data significantly improves training performance, especially for underrepresented classes, the purely synthetically trained model showed almost no detection capability. Our results suggest that the bottleneck lies less in the concept and use of synthetic data itself and more in the currently designed SyntEO textures, whose backscatter and noise modelling do not yet sufficiently reflect key aspects of real SAR signatures. Further development of synthetic textures is crucial if synthetic data is to be used not only as a supplement but also as a replacement in the future. Despite this limitation, combined training with synthetic and real data shows that robust and highly generalizable models can be developed that reliably recognize both individual and complex groups of platforms and clearly distinguish substations from turbines in offshore wind farms. Overall, the results underscore the high robustness, generalizability, and transferability of the approach to previously unknown areas.

The presented method contributes to addressing the current lack of up-to-date and complete data on offshore infrastructure, which is crucial for resource management, environmental protection, and marine surveillance. It provides a first step towards an objective and scalable tool that can monitor changes in the inventory and development of offshore platforms across broad spatial and temporal scales. Overall, this approach marks an important step forward and offers new insights into the transition to global, fully automated identification and monitoring of maritime energy infrastructure. It creates a solid foundation for improving government and scientific databases, supports more sustainable resource management, and contributes to transparency and traceability of changes in the world's oceans.

Future work will focus on further optimizing the proposed method and transferring it into an operational context. Building on its successfully demonstrated regional transferability, the developed GCP framework provides the basis for global, near real-time monitoring of the entire Sentinel-1 archive. This would enable continuous, automated updating of the global inventory of offshore platforms, representing an important step towards operational monitoring. In addition, the model will be applied to different temporal scales to systematically record, analyse, and evaluate changes in maritime energy infrastructure.

Acknowledgment

The authors gratefully acknowledge ESA's Copernicus program for providing free access to the Sentinel-1 data and the Google Earth Engine platform for preprocessing and making the data accessible.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Robin Spanier  <http://orcid.org/0009-0005-5959-6210>
 Thorsten Hoeser  <http://orcid.org/0000-0002-7179-3664>
 Claudia Kuenzer  <http://orcid.org/0009-0007-4933-5898>

Author contributions

Robin Spanier: Conceptualization, methodology, code development, data curation, validation, visualization, original manuscript writing. Thorsten Hoeser: Methodology, code development, manuscript reviewing. Claudia Kuenzer: Supervision, manuscript reviewing.

Data availability statement

Data will be made available upon reasonable request.

References

4C Offshore. 2025. *Open Global Offshore Map*. <https://map.4coffshore.com/offshorewind/>.

Aghaei, N., G. Akbarizadeh, and A. Kosarian. 2022a. "Greywolfslsm: An Accurate Oil Spill Detection Method Based on Level Set Method from Synthetic Aperture Radar Imagery." *European Journal of Remote Sensing* 55 (1): 181–198. <https://doi.org/10.1080/22797254.2022.2037468>.

Aghaei, N., G. Akbarizadeh, and A. Kosarian. 2022b. "Osdes_Net: Oil Spill Detection Based on Efficient_Shuffle Network Using Synthetic Aperture Radar Imagery." *Geocarto International* 37 (26): 13539–13560. <https://doi.org/10.1080/10106049.2022.2082545>.

Amani, M., A. Ghorbanian, S. A. Ahmadi, M. Kakooei, S. M. M. Armin Moghimi, S. H. A. Moghaddam, S. H. A. Moghaddam, et al. 2020. "Google Earth Engine Cloud Computing Platform for Remote Sensing Big Data Applications: A Comprehensive Review." *IEEE Journal of Selected Topics Applied Earth Observations Remote Sensing* 13:5326–5350. <https://doi.org/10.1109/JSTARS.2020.3021052>.

Baumhoer, C. A., A. J. Dietz, K. Heidler, and C. Kuenzer. 2023. "Icelines - A New Data Set of Antarctic Ice Shelf Front Positions." *Scientific Data* 10 (1): 138. <https://doi.org/10.1038/s41597-023-02045-x>.

Blondeau-Patissier, D., T. Schroeder, G. Suresh, Z. Li, F. I. Diakogiannis, P. Irving, C. Witte, and A. D. L. Steven. 2023. "Detection of Marine Oil-Like Features in Sentinel-1 SAR Images by Supplementary Use of Deep Learning and Empirical Methods: Performance Assessment for the Great Barrier Reef Marine Park." *Marine Pollution Bulletin* 188:114598. <https://doi.org/10.1016/j.marpolbul.2023.114598>.

BOEM, Bureau of Ocean Energy Management. 2024. *Oil & Gas Mapping and Data: Oil & Gas Infrastructure*. <https://www.boem.gov/oil-gas-energy/mapping-and-data>.

Bugnot, A. B., M. Mayer-Pinto, L. Airoldi, E. C. Heery, E. L. Johnston, L. P. Critchley, E. M. A. Strain, et al. 2021. "Current and Projected Global Extent of Marine Built Structures." *Nature Sustainability* 4 (1): 33–41. <https://doi.org/10.1038/s41893-020-00595-1>.

Casadio, S., O. Arino, and A. Minchella. 2012. "Use of ATSR and SAR Measurements for the Monitoring and Characterisation of Night-Time Gas Flaring from Off-Shore Platforms: The North Sea Test Case." *Remote Sensing of Environment* 123:175–186. <https://doi.org/10.1016/j.rse.2012.03.021>.

Cheng, L., K. Yang, L. Tong, Y. Liu, and M. Li. 2013. "Invariant Triangle-based Stationary Oil Platform Detection from Multitemporal Synthetic Aperture Radar Data." *Journal of Applied Remote Sensing* 7 (1): 73537. <https://doi.org/10.11117/1.JRS.7.073537>.

Davari, N., G. Akbarizadeh, and E. Mashhour. 2022. "Corona Detection and Power Equipment Classification Based on GoogleNet-AlexNet: An Accurate and Intelligent Defect Detection Model Based on Deep Learning for Power Distribution Lines." *IEEE Transactions on Power Delivery* 37 (4): 2766–2774. <https://doi.org/10.1109/TPWRD.2021.3116489>.

Dong, C.-L., X. Meng, and L.-X. Guo. 2023. "Research on SAR Imaging Simulation Based on Time-Domain Shooting and Bouncing Ray Algorithm." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 16:1519–1530. <https://doi.org/10.1109/JSTARS.2023.3237628>.

El-Darymli, K., P. McGuire, D. Power, and C. Moloney. 2013. "Target Detection in Synthetic Aperture Radar Imagery: A State-of-the-Art Survey." *Journal of Applied Remote Sensing* 7 (1): 71598. <https://doi.org/10.1117/1.JRS.7.071598>.

ESA, European Space Agency. 2021. "Sentinel-2 Products Specification Document." S2-PDGS-TAS-DI-PSD. <https://sentinels.copernicus.eu/documents/247904/0/Sentinel-2-product-specifications-document-V14-9.pdf>.

ESA, European Space Agency. 2025a. Sentinel-1. <https://sentiwiki.copernicus.eu/web/sentinel-1>.

ESA, European Space Agency. 2025b. Sentinel-2. <https://sentiwiki.copernicus.eu/web/sentinel-2>.

Esch, T., E. Brzoska, S. Dech, B. Leutner, D. Palacios-Lopez, A. Metz-Marconcini, M. Marconcini, A. Roth, and J. Zeidler. 2022. "World Settlement Footprint 3D - A First Three-Dimensional Survey of the Global Building Stock." *Remote Sensing of Environment* 270:112877. <https://doi.org/10.1016/j.rse.2021.112877>.

European Comission, and Directorate-General for Environment. 2022. *Environmental Impact Assessment of Projects – Rulings of the Court of Justice of the European Union*.

GEE, Google Earth Engine, and ESA, European Space Agency. 2025. SAR GRD Collection (COPERNICUS/S1\GRD). Sentinel-1. https://developers.google.com/earth-engine/datasets/catalog/COPERNICUS_S1_GRD.

Ghara, M., S. B. S. Fatemeh, and G. Akbarizadeh. 2022. "A New Technique for Segmentation of the Oil Spills from Synthetic-Aperture Radar Images Using Convolutional Neural Network." *IEEE Journal of Selected Topics in Applied Earth Observations Remote Sensing* 15:8834–8844. <https://doi.org/10.1109/JSTARS.2022.3213768>.

Global Energy Monitor. 2025. *Global Oil and Gas Extraction Tracker*. <https://globalenergymonitor.org/projects/global-oil-gas-extraction-tracker/tracker-map/>.

Gorelick, N., M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore. 2017. "Google Earth Engine: Planetary-Scale Geospatial Analysis for Everyone." *Remote Sensing of Environment* 202:18–27. <https://doi.org/10.1016/j.rse.2017.06.031>.

He, Y., L. Xu, J. Huo, H. Zhou, and X. Shi. 2024. "A Synthetic Aperture Radar Imaging Simulation Method for Sea Surface Scenes Combined with Electromagnetic Scattering Characteristics." *Remote Sensing* 16 (17): 3335. <https://doi.org/10.3390/rs16173335>.

Hoeser, T., F. Bachofer, and C. Kuenzer. 2020. "Object Detection and Image Segmentation with Deep Learning on Earth Observation Data: A Review—Part II: Applications." *Remote Sensing* 12 (18): 3053. <https://doi.org/10.3390/rs12183053>.

Hoeser, T., S. Feuerstein, and C. Kuenzer. 2022. "DeepOWT: A Global Offshore Wind Turbine Data Set Derived with Deep Learning from Sentinel-1 Data." *Earth System Science Data* 14 (9): 4251–4270. <https://doi.org/10.5194/essd-14-4251-2022>.

Hoeser, T., and C. Kuenzer. 2020. "Object Detection and Image Segmentation with Deep Learning on Earth Observation Data: A Review-Part I: Evolution and Recent Trends." *Remote Sensing* 12 (10): 1667. <https://doi.org/10.3390/rs12101667>.

Hoeser, T., and C. Kuenzer. 2022a. "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. <https://doi.org/10.1016/j.jag.2022.102957>.

Hoeser, T., and C. Kuenzer. 2022b. "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. <https://doi.org/10.1016/j.isprsjprs.2022.04.029>.

IHO Sea Areas, Version 3. 2018.

Jackson, C. R., and J. R. Apel. 2004. Synthetic Aperture Radar: Marine User's Manual. <https://repository.library.noaa.gov/view/noaa/1118>.

Jia, H., P. Yang, and R. Wu. 2023. "Comparative Study of Sentinel-1-Focused and Simulated SAR Images Using Lidar Point Cloud Modeling for Coastal Areas." *Electronics* 12 (20): 4312. <https://doi.org/10.3390/electronics12204312>.

Komen, J. 2008. Oil Platform North of Ameland, Holland (9429668332).jpg. [https://commons.wikimedia.org/wiki/File:Oil_platform_north_of_Ameland,_Holland_\(9429668332\).jpg](https://commons.wikimedia.org/wiki/File:Oil_platform_north_of_Ameland,_Holland_(9429668332).jpg).

Liu, Y., C. Hu, Y. Dong, B. Xu, W. Zhan, and C. Sun. 2019. "Geometric Accuracy of Remote Sensing Images Over Oceans: The Use of Global Offshore Platforms." *Remote Sensing of Environment* 222:244–266. <https://doi.org/10.1016/j.rse.2019.01.002>.

Liu, Y., C. Sun, Y. Yang, M. Zhou, W. Zhan, and W. Cheng. 2016. "Automatic Extraction of Offshore Platforms Using Time-Series Landsat-8 Operational Land Imager Data." *Remote Sensing of Environment* 175:73–91. <https://doi.org/10.1016/j.rse.2015.12.047>.

Luo, B., H. Cao, J. Cui, X. Lv, J. He, H. Li, and C. Peng. 2025. "SAR-PATT: A Physical Adversarial Attack for SAR Image Automatic Target Recognition." *Remote Sensing* 17 (1): 21. <https://doi.org/10.3390/rs17010021>.

Ma, P., M. Macdonald, S. Rouse, and J. Ren. 2024. "Automatic Geolocation and Measuring of Offshore Energy Infrastructure with Multimodal Satellite Data." *IEEE Journal of Oceanic Engineering* 49 (1): 66–79. <https://doi.org/10.1109/JOE.2023.3319741>.

March, D., K. Metcalfe, J. Tintoré, and B. J. Godley. 2021. "Tracking the Global Reduction of Marine Traffic During the COVID-19 Pandemic." *Nature Communications* 12 (1): 2415. <https://doi.org/10.1038/s41467-021-22423-6>.

Marino, A., D. Vellootto, and F. Nunziata. 2017. "Offshore Metallic Platforms Observation Using Dual-Polarimetric TS-X/TD-X Satellite Imagery: A Case Study in the Gulf of Mexico." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 10 (10): 4376–4386. <https://doi.org/10.1109/JSTARS.2017.2718584>.

Meta AI. 2025. "Torch-Model-Archiver: PyTorch Foundation." <https://github.com/pytorch/serve/blob/master/model-archiver/README.md>.

Montera, L. D., H. Berger, R. Husson, P. Appelghem, L. Guerlou, and M. Fragoso. 2022. "High-Resolution Offshore Wind Resource Assessment at Turbine Hub Height with Sentinel-1 Synthetic Aperture Radar (SAR) Data and Machine Learning." *Wind Energ. Sci* 7 (4): 1441–1453. <https://doi.org/10.5194/wes-7-1441-2022>.

North Sea Energy. 2025. *Project Atlas - An Interactive Atlas of the North Sea*. <https://northseaenergy-projectatlas.app/atlas/introduction-energy/visibility>.

OECD, Organisation for Economic Co-operation and Development. 2021. "OECD Science, Technology and Industry Working Papers." *Blueprint for Improved Measurement of the International Ocean Economy: An Exploration of Satellite Accounting for Ocean Economic Activity*. OECD Publishing.

OECD, Organisation for Economic Co-operation and Development. 2025. "The Ocean Economy to 2050." OECD Publishing.

OSM, OpenStreetMap. 2025a. *Land Polygons*. <https://osmdata.openstreetmap.de/data/land-polygons.html>.

OSM, OpenStreetMap. 2025b. *OpenseaMap*. <http://openseaMap.org/>.

OSPAR, Oslo and Paris Conventions. 2025. *Inventory of Offshore Installations - 2023*. <https://odims.ospar.org/en/maps/map-inventory-of-offshore-installations-2023/>.

Paolo, F. S., D. Kroodsma, J. Raynor, T. Hochberg, P. Davis, J. Cleary, L. Marsaglia, S. Orofino, C. Thomas, and P. Halpin. 2024. "Satellite Mapping Reveals Extensive Industrial Activity at Sea." *Nature* 625 (7993): 85–91. <https://doi.org/10.1038/s41586-023-06825-8>.

Petrobras, D. 2009. Oil Platform P-51 (Brazil).jpg. <https://commons.wikimedia.org/w/index.php?curid=5621984>.

Posen, P. E., K. Hyder, M. Teixeira Alves, N. G. Taylor, and C. P. Lynam. 2020. "Evaluating Differences in Marine Spatial Data Resolution and Robustness: A North Sea Case Study." *Ocean & Coastal Management* 192:105206. <https://doi.org/10.1016/j.ocecoaman.2020.105206>.

Potter, R. W. K., and B. C. Pearson. 2023. "Assessing the Global Ocean Science Community: Understanding International Collaboration, Concerns and the Current State of Ocean Basin Research." *NPJ Ocean Sustain* 2 (1). <https://doi.org/10.1038/s44183-023-00020-y>.

Qiu, Y., X.-M. Li, L. Yan, and Z. Chen. 2024. "Synergic Sensing of Light and Heat Emitted by Offshore Oil and Gas Platforms in the South China Sea." *International Journal of Digital Earth* 17 (1). <https://doi.org/10.1080/17538947.2024.2441932>.

Raudøy, A. N. 2016. Bligh Bank 2 Offshore Substation BE 2016.jpg. https://commons.wikimedia.org/wiki/File:Bligh_Bank_2_offshore_substation_BE_2016.jpg.

Roletschek, R. 2011. 11-09-04-mittelplate-by-RalfR-18.jpg. <https://commons.wikimedia.org/wiki/File:11-09-04-mittelplate-by-RalfR-18.jpg>.

Sadeghi, K. 2007. "An Overview of Design, Analysis, Construction and Installation of Offshore Petroleum Platforms Suitable for Cyprus Oil/Gas Fields." *GAU Journal of Social Applied Science* 2 (4): 1–16.

Samadi, F., G. Akbarizadeh, and H. Kaabi. 2019. "Change Detection in SAR Images Using Deep Belief Network: A New Training Approach Based on Morphological Images." *IET Image Processing* 13 (12): 2255–2264. <https://doi.org/10.1049/iet-ipr.2018.6248>.

Sharifzadeh, F., G. Akbarizadeh, and Y. Seifi Kavian. 2019. "Ship Classification in SAR Images Using a New Hybrid CNN–MLP Classifier." *Journal of Indian Society of Remote Sensing* 47 (4): 551–562. <https://doi.org/10.1007/s12524-018-0891-y>.

Spanier, R., and C. Kuenzer. 2024. "Marine Infrastructure Detection with Satellite Data—A Review." *Remote Sensing* 16 (10): 1675. <https://doi.org/10.3390/rs16101675>.

Sun, C., Y. Liu, S. Zhao, and S. Jin. 2020. "Estimating Offshore Oil Production Using DMSP-OLS Annual Composites." *ISPRS Journal of Photogrammetry and Remote Sensing* 165:152–171. <https://doi.org/10.1016/j.isprsjprs.2020.05.019>.

Taibi, F., G. Akbarizadeh, and E. Farshidi. 2019. "Robust Reservoir Rock Fracture Recognition Based on a New Sparse Feature Learning and Data Training Method." *Multidimensional Systems and Signal Processing* 30 (4): 2113–2146. <https://doi.org/10.1007/s11045-019-00645-8>.

Tirandaz, Z., G. Akbarizadeh, and H. Kaabi. 2020. "POLSAR Image Segmentation Based on Feature Extraction and Data Compression Using Weighted Neighborhood Filter Bank and Hidden Markov Random Field-Expectation Maximization." *Measurement* 153:107432. <https://doi.org/10.1016/j.measurement.2019.107432>.

UNCTAD, United Nations Conference on Trade and Development. 2021. "Exploring Space Technologies for Sustainable Development." <https://unctad.org/publication/exploring-space-technologies-sustainable-development>.

Virdin, J., T. Vegh, J.-B. Jouffray, R. Blasiak, S. Mason, H. Österblom, D. Vermeer, H. Wachtmeister, and N. Werner. 2021. "The Ocean 100: Transnational Corporations in the Ocean Economy." *Science Advances* 7 (3). <https://doi.org/10.1126/sciadv.abc8041>.

Wang, A., H. Chen, L. Liu, K. Chen, Z. Lin, J. Han, and G. Ding. 2024. "YOLOv10: Real-Time End-to-End Object Detection." <https://arxiv.org/abs/2405.14458>.

Wang, J., M. Li, Y. Liu, H. Zhang, W. Zou, and L. Cheng. 2014. "Safety Assessment of Shipping Routes in the South China Sea Based on the Fuzzy Analytic Hierarchy Process." *Safety Science* 62:46–57. <https://doi.org/10.1016/j.ssci.2013.08.002>.

Williamson, B. J., P. Blondel, E. Armstrong, P. S. Bell, C. Hall, and J. J. Waggett. 2016. "A Self-Contained Subsea Platform for Acoustic Monitoring of the Environment Around Marine Renewable Energy Devices—Field Deployments at Wave and Tidal Energy Sites in Orkney, Scotland." *IEEE Journal of Oceanic Engineering* 41 (1): 67–81. <https://doi.org/10.1109/JOE.2015.2410851>.

Wong, B. A., C. Thomas, and P. Halpin. 2019. "Automating Offshore Infrastructure Extractions Using Synthetic Aperture Radar & Google Earth Engine." *Remote Sensing of Environment* 233:111412. <https://doi.org/10.1016/j.rse.2019.111412>.

Wu, B., H. Wang, C. Zhang, and J. Chen. 2024. "Optical-to-SAR Translation Based on CDA-GAN for High-Quality Training Sample Generation for Ship Detection in SAR Amplitude Images." *Remote Sensing* 16 (16): 3001. <https://doi.org/10.3390/rs16163001>.

Xu, W., Y. Liu, W. Wu, Y. Dong, W. Lu, Y. Liu, B. Zhao, H. Li, and R. Yang. 2020. "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. <https://doi.org/10.1016/j.rser.2020.110167>.

Xu, Z., H. Zhang, Y. Wang, X. Wang, S. Xue, and W. Liu. 2022. "Dynamic Detection of Offshore Wind Turbines by Spatial Machine Learning from Spaceborne Synthetic Aperture Radar Imagery." *Journal of King Saud University - Computer and Information Sciences* 34 (5): 1674–1686. <https://doi.org/10.1016/j.jksuci.2022.02.027>.

Yang, J., P. Gong, R. Fu, M. Zhang, J. Chen, S. Liang, B. Xu, J. Shi, and R. Dickinson. 2013. "The Role of Satellite Remote Sensing in Climate Change Studies." *Nature Climate Change* 3 (10): 875–883. <https://doi.org/10.1038/nclimate1908>.

Zeng, Z., X. Tan, X. Zhang, Y. Huang, J. Wan, and Z. Chen. 2024. "ATGAN: A SAR Target Image Generation Method for Automatic Target Recognition." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17:6290–6307. <https://doi.org/10.1109/JSTARS.2024.3370185>.

Zhang, J., Q. Wang, and F. Su. 2019. "Automatic Extraction of Offshore Platforms in Single SAR Images Based on a Dual-Step-Modified Model." *Sensors (Basel, Switzerland)* 19 (2): 231. <https://doi.org/10.3390/s19020231>.

Zhang, T., B. Tian, D. Sengupta, L. Zhang, and Y. Si. 2021. "Global Offshore Wind Turbine Dataset." *Scientific Data* 8 (1): 191. <https://doi.org/10.1038/s41597-021-00982-z>.

Zhang, X., M. Qiu, S. Tao, X. Ge, and M. Wang. 2025. "OPDNet: An Offshore Platform Detection Network Based on Bitemporal Bimodal Remote-Sensing Images and a Pseudo-Siamese Structure." *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 18:6409–6421. <https://doi.org/10.1109/JSTARS.2025.3539674>.