

Oil Spill Thickness Estimation in Polarimetric SAR Imagery Using an Improved CoTNet-SE Model

Taner Başar^{1,*} and Salih Övünç¹

¹ Department of Computer Engineering, Erciyes University, Kayseri 38039, Turkey

*Corresponding author: taner.basar@erciyes.edu.tr

Abstract. Reliable estimation of offshore oil spill thickness from polarimetric SAR imagery is required for emergency response planning, yet the backscatter response of thin, moderate, and thick slicks is often mixed with wind, incidence angle, look direction, and local sea-surface roughness effects. This paper proposes an improved Contextual Transformer Network with Squeeze-and-Excitation recalibration, called CoTNet-SE, for continuous thickness estimation from calibrated PolSAR patches. The method introduces polarimetric damping descriptors, contextual attention, channel-wise excitation, auxiliary grade supervision, and uncertainty-aware regression to jointly model weak slick texture, polarization-specific suppression, and spatially heterogeneous slick morphology. A controlled multi-scene benchmark of 6,240 labelled patches, four thickness grades, six wind-speed intervals, and three incidence-angle groups was selected for the experiment. Compared with the ResNet-50 baseline, the proposed model reduced the mean absolute error from 0.071 mm to 0.038 mm and the root mean square error from 0.104 mm to 0.056 mm, and increased the coefficient of determination from 0.781 to 0.913. Ablation experiments show that contextual attention reduced RMSE by 19.8%, while SE recalibration further reduced it by 11.3% in mixed sea-state conditions. The uncertainty head achieved 91.6% coverage within the predicted 90% interval. These results indicate that polarimetric channel recalibration and contextual dependency modelling provide an effective route for SAR-based oil spill thickness retrieval in engineering monitoring workflows.

Keywords: *Polarimetric SAR, Oil Spill Thickness, CoTNet-SE, Contextual Attention, Marine Remote Sensing*

Received on 16 September 2023, Accepted on 29 January 2024, Published on 11 February 2024

Copyright © 2024 Author(s), licensed to DEA. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Introduction

Monitoring of marine oil spills has shifted from identifying a slick qualitatively to quantitatively assessing the severity of the slick; thus, data now need to be closer to volume and thickness for response vessel deployment, dispersant allocation and environmental risk assessment [1]. Synthetic aperture radar is now one of the most frequently used remote sensing sources for marine oil spill monitoring because it can operate at night and in bad weather, and because oil films suppress short gravity-capillary waves that would normally produce radar backscatter from the sea surface [2]. Polarimetric SAR has a richer response than single-channel SAR; co-polarized and cross-polarized returns can record different information about surface damping, azimuthal anisotropy and scattering stability [3]. Thick emulsified oil, thin sheen, biogenic film, low-wind look-alikes, and ship wakes may have similar dark signatures in intensity imagery, but polarimetric descriptors can reveal more subtle differences in scattering randomness and channel coupling [4]. Field studies have shown that the same nominal thickness can produce different normalized radar cross sections under different conditions of wind speed, incidence angle and oil composition [5]. Recently, emergency-response systems have gradually required pixel- or patch-level thickness classes within hours of acquisition, and retrieval models now need to be accurate, lightweight and explainable enough for engineering deployment [6]. Therefore, the problem is not to find darker pixels that indicate thicker oil; instead, it is how to learn a conditional relationship between polarimetric damping and the actual condition of the surface. A thickness map with reasonable uncertainty can guide skimming priority,

controlled burning assessment, shoreline protection and post-event loss estimation more directly than a binary slick mask. Therefore, the remote sensing problem is now a quantitative retrieval problem under conditions of incomplete observation.

The existing oil spill retrieval pipelines often use handcrafted polarimetric features, empirical thresholds or shallow classifiers to distinguish slicks from clean water and then assign a relatively coarse severity level. Such pipelines are attractive because they are transparent, but their accuracy decreases when the training conditions differ from those of the observed sea state [7]. Deep convolutional networks have enhanced the recognition of smooth areas and changes in texture by learning spatial patterns automatically [8]. However, convolutional receptive fields may still be inadequate if the differentiation of thickness grades requires contextual damping over extended slick structures rather than local intensity alone [9]. Transformer-based models address this issue by modeling long-range dependencies, but the computational cost of standard attention mechanisms is high, which may lead to neglecting polarization-specific channel information while emphasizing scene-level contrast [10]. Channel attention has been introduced in remote sensing recognition to emphasize information-rich bands and suppress unstable channels, but its application in continuous oil thickness estimation under polarimetric synthetic aperture radar conditions is still insufficient [11]. The model for this task needs to address three interrelated sources of ambiguity: the nonlinear relationship between attenuation and thickness, the spatial mixing of oil and clear water at the oil film boundary, and the conditional reliability of each polarization channel under different wind and incidence geometry conditions. The deficiencies prompted us to construct a concise model with context-related relationships, polarization channel recalibration, and regression-based uncertainty quantification.

This paper introduces an improved Contextual Transformer Network with Squeeze-and-Excitation recalibration, named CoTNet-SE, for marine oil spill thickness estimation from PolSAR image patches. In general, it is to maintain the local structural sensitivity of convolution and add contextual attention blocks that adapt to slick morphology, as well as channel-excitation gates that reweight polarimetric evidence according to the observed sea state. The calibrated HH, HV, and VV channels, along with the derived damping and decomposition features, are supplied to the model, and both a continuous thickness estimate and an auxiliary operating grade are outputted. The auxiliary branch stabilizes the training by employing categorical thickness intervals, and the regression branch keeps the quantitative output needed for the following volume approximation. The controlled train-test split of the proposed framework is divided into wind-speed, incidence-angle and slick-thickness regions. The experiments also consider bias, robustness, calibration and computational cost; otherwise, it would only be the average error, and such a method would not be practical in practice. Therefore, the contributions of this paper are methodological and engineering-oriented: adapting contextual transformer representation to PolSAR oil slick structure, introducing explicit channel recalibration for thickness-sensitive evidence, and evaluating the model with metrics that reflect deployment risk rather than only pooled accuracy.

To make the argument traceable from the remote sensing problem to the engineering evidence, the rest of this paper is organized as follows. Section 2 introduces polarimetric SAR oil spill analysis and deep learning models for thickness-sensitive marine remote sensing, as well as the transition from detection to quantitative retrieval. Section 3 introduces the proposed CoTNet-SE framework, which includes feature construction, contextual attention, SE recalibration, dual prediction heads and the learning objective. This section will also show how overlapping patch inference converts patch-level estimates into scene-level thickness products. Section 4 shows the experimental data, comparison models, ablation studies, robustness analysis, uncertainty evaluation and computational assessment in a single table and six data figures. The experimental part will present the results in an engineering context and, instead of listing individual indicators, will also consider the stability of operation and the mean error in determining the thickness. Section 5 is the conclusion of this paper, and it will briefly present the research results and shortcomings. This organization follows the logic of a retrieval paper: first, it presents the physical ambiguity; then, it identifies the modelling gap; next, it introduces the proposed network; finally, it provides evidence through controlled experiments. The Structure also separates the design of the method from the experimental results to show how each architectural choice has been verified.

Related Work

Polarimetric SAR Analysis for Oil Spill Characterization

Polarimetric SAR oil spill analysis has traditionally been based on the physical observation that oil films reduce the Bragg-scale roughness of the sea surface and thus alter both co-polarized and cross-polarized returns. Early studies used normalized radar cross-section, damping ratio, co-polarization ratio and local contrast to separate slick regions from the surrounding water [12]. Later work added polarimetric decomposition variables, such as entropy, anisotropy and the alpha angle, to differentiate between mineral oil and low-wind areas and natural surfactants [13]. The advantage of these descriptors is that they connect image measurements to understandable scattering behavior, and thus are necessary for emergency users who require a justifiable reason for a mapped slick. Another issue is that the relationship between the descriptor values and thickness is not uniform. The same oil layer may appear differently under different sea conditions, and different oils may have overlapping polarization responses at the same thickness [14]. Therefore, some operating systems use a combination of thresholding, region growing and post-event expert interpretation instead of direct quantitative retrieval. This way can quickly identify the problem, but it does not meet the engineering requirement for consistent thickness estimation under various observation conditions [15].

Next is how to find the thickness after detection. A thin film may be close to the noise floor of the sensor, and a thick emulsion will saturate the damping response and cause internal texture due to wave-current interaction. Manual feature selection may fail to capture this nonlinearity because the most informative features are themselves functions of incidence angle, wind speed and local slick age. Patch-level learning is a convenient way; it maintains the local context near the oil-water interface and can use both local contrast and regional texture in the model. A model that only uses intensity features may still mistake thick oil for wind-shadowed water. Therefore, a thickness-oriented PolSAR method should retain polarization-specific information, include physically meaningful derived descriptors, and remain sensitive to spatial context without being dominated by scene-level darkness in the regression.

The second is the problem of label construction. Direct thickness measurements are relatively rare in many marine accidents because sampling vessels cannot cover the entire slick, and the slick may have changed by the time of satellite acquisition. Therefore, remote sensing studies often use response reports, aerial inspections, controlled-spill records or expert-defined grade intervals. These are some resources, but their accuracy in labelling varies. A way to reduce the average error may still be unsuitable if it results in excessively high confidence for poorly observed areas. Therefore, recent engineering-oriented retrieval work has begun to incorporate uncertainty, robustness and class consistency in addition to mean error. The label problem in oil spill thickness estimation also advocates for auxiliary grade supervision. Grade labels are relatively stable compared to the exact thickness values, but continuous regression is still required for volume-sensitive decisions. Both supervision signals can be used to help the model learn a stable ordering of slick severity and maintain a quantitative output.

Deep Learning Models for SAR Retrieval and Attention Mechanisms

Deep learning has changed SAR interpretation by replacing fixed feature thresholds with trainable representations that can learn speckle statistics, texture variations and sensor-specific artifacts. Convolutional Neural Networks have been used to deal with oil spill detection, sea ice classification, ship detection and flood mapping due to their good local-spatial feature extraction abilities for noisy radar images [16]. Encoder-decoder architectures have further improved dense prediction by fusing multiscale features to recover boundaries and small targets [17]. Oil spill thickness estimation does not output a segmentation mask; it is different. The model should learn a continuous response that can change smoothly in magnitude near slick and sharp emulsified patches. Purely convolutional networks need many layers to learn long-range correlations in the slick morphology, and their channel mixing is often implicit rather than physically motivated.

Attention mechanisms offer a way to achieve more adaptive representation learning. Self-attention can relate distant image regions and has been adapted for remote sensing scene classification and dense prediction [18]. Contextual Transformer modules reduce the cost of full attention by combining local convolutional priors with attention-like contextual encoding, making them suitable for high-resolution SAR patches [19]. Channel Attention Mechanisms, such as Squeeze-and-Excitation (SE) modules, increase the representational selectivity

by determining which channels are more informative for the current sample [20]. Despite these advances, few studies have connected contextual attention and channel recalibration specifically to oil spill thickness regression in PolSAR imagery. The current work addresses this problem by embedding SE recalibration in a contextual transformer backbone and using a dual regression-grade objective designed for thickness-sensitive engineering interpretation.

The existing literature has not yet addressed how to apply high-capacity attention models to operational SAR processing. Full-transformer backbones are accurate, but they generally require a large amount of labelled data, considerable memory, and careful regularization. Marine oil spill datasets are generally smaller and more imbalanced than the typical optical remote sensing benchmarks, and for thick oil samples, they are even rarer and more incident-specific. Lightweight convolutional networks are relatively simple to deploy, but they may not separate scene-wide damping from actual thickness variations. A Contextual Transformer structure is a compromise that retains the local convolutional inductive bias and can make features sensitive to large-scale variations in shape. In conjunction with channel excitation, this structure can determine whether HH, HV, VV or the derived descriptors are more suitable for the observed patch. The Design Direction needs to meet the engineering requirements of having a retrieval algorithm accurate enough for response planning and efficient enough for repeated scene-level inference.

Proposed CoTNet-SE Framework

The framework proposed here considers oil spill thickness estimation to be a supervised patch-level regression problem and adds an auxiliary grade prediction task. Each sample is a co-registered PolSAR patch with calibrated HH, HV and VV backscatter channels and derived polarization descriptors. First, perform radiometric calibration, then fine Lee filtering, normalize incidence-angle, and finally extract the ocean mask. The feature cube is then passed to the improved CoTNet-SE backbone, and contextual transformer blocks model long-range dependencies in slick texture, while squeeze-and-excitation units adaptively recalibrate polarization-sensitive channels. Figure 1 is used here to show the entire process of SAR acquisition to thickness and confidence products.




Figure 1. Polarimetric SAR oil spill thickness estimation workflow. Placeholder for a white-background academic flowchart generated externally.

The input tensor is denoted by $X \in \mathbb{R}^{H \times W \times C}$, where H and W are the patch size and C is the number of polarimetric and derived channels. For each patch, the calibrated feature vector includes logarithmic backscatter, co-polarization ratio, cross-polarization damping, local contrast, entropy-like texture, and incidence-angle normalized descriptors. The normalized channel x_c is expressed as in implementation, the channel set includes

calibrated HH, HV, and VV backscatter together with damping ratio, co-polarization ratio, local texture, and incidence-angle descriptors, so the notation is directly tied to observable PolSAR evidence.

$$x_c = \frac{\log(\sigma_c^0) - \mu_c}{\sigma_c + \varepsilon}, c \in \{HH, HV, VV, D_1, \dots, D_K\} \quad \text{Eq.(1)}$$

The normalization reduces scene-to-scene amplitude drift while retaining local slick-water contrast. It also prevents high-backscatter clean-water areas from dominating early convolutional kernels when a patch contains both slick and background pixels.

The Descriptor Set is relatively small. Not all possible polarimetric decomposition terms are included; otherwise, due to high correlation among the descriptors, training instability would occur with a small number of verified thick-oil samples. Therefore, the feature construction stage keeps variables with direct engineering meanings for slick damping and channel contrast. The three calibrated backscatter channels retain the original sensor data. The co-polarization ratio shows how much the *HH* and *VV* responses are suppressed. The cross-polarization damping term accounts for roughness-sensitive scattering that responds differently to wind and film thickness. A local coefficient-of-variation channel records texture strength in the patch and is suitable for regions close to mixed-slick boundaries. Incidence angle is added as a normalized auxiliary layer so that the network does not need to infer viewing geometry from intensity patterns alone.

The contextual transformer block is designed to preserve local edge sensitivity and expand the effective receptive field. A convolutional key generator first encodes local texture, while a contextual attention kernel aggregates spatial dependency inside the slick region. For a feature map F , the contextual response at pixel i is modeled as

$$y_i = \sum_{j \in N(i)} \text{softmax}\left(\frac{q_i^T k_j}{\sqrt{d}}\right) v_j + W_i f_i \quad \text{Eq.(2)}$$

where $N(i)$ denotes the contextual neighborhood, q_i is the query vector, k_j is the local key vector, and v_j is the value vector. In contrast to full global attention, this form keeps the attention computation bounded and makes it suitable for medium-resolution PolSAR patches. Figure 2 is cited here to describe the internal architecture, including stem convolution, contextual transformer blocks, SE recalibration, residual aggregation, and the dual prediction heads.




Figure 2. Improved CoTNet-SE architecture for PolSAR thickness regression. Placeholder for a white-background academic model diagram generated externally.

The SE module reweights each channel according to its global response and the learned reliability of the polarimetric evidence. Given feature map U , global squeezing produces a channel descriptor z_c , and the excitation gate s_c is calculated as

$$z_c = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W U_c(i,j), s = \text{sigmoid}(W_2 \delta(W_1 z)) \quad \text{Eq.(3)}$$

The recalibrated feature $U'_c = s_c U_c$ allows the network to emphasize channels that remain informative under the current incidence angle and sea state. This mechanism is particularly important for thickness retrieval because cross-polarized returns can be weak in thin slicks, while co-polarized damping can saturate in thicker emulsions. The architecture places SE blocks after contextual modules rather than before them so that channel gating is informed by both local texture and contextual slick morphology.

The back-end has four links. The first stage has a relatively fine spatial grid and is concerned with speckle-tolerant edge and contrast features. The second and third stages add contextual transformer blocks with increasing channel width to help the network learn extended slick structures, wind-aligned streaks and internal texture variations. The fourth is a residual aggregation at multiple scales that covers both thin-edge areas and thick-interior regions in the same latent space. Residual shortcuts are kept around each contextual block because the thickness labels contain noise, and the model should not overfit to a single attention pattern. Dropout is used only in the prediction head, and stochastic depth is applied lightly at the deeper contextual stage. The above structure is relatively deterministic and thus suitable for stable scene-level mosaicking.

The SE recalibration is not considered a general classification improvement. The purpose of the above task is to learn conditional trust in polarimetric channels. For example, when the wind speed is low, a uniformly dark area may be caused by weak surface roughness rather than by thick oil, so the excitation mechanism should not rely solely on co-polarized darkness. At a moderate wind speed, damping contrast is more indicative, and the weight of co-polarized ratio can be increased. In a strong wind, the sea surface is rougher; therefore, cross-polarized texture and local contrast descriptors are more discriminative. Therefore, the excitation vector serves as a learned diagnosis of which physical evidence the model is employing in a particular patch. A larger excitation weight is interpreted as stronger conditional reliance on that polarimetric channel under the current sea-state and incidence-angle conditions.

The output layer contains a continuous regression head and an auxiliary grade head. The continuous head predicts thickness $\hat{t}$ in millimeters, while the grade head predicts one of four operational intervals: sheen, thin film, moderate oil, and thick emulsion. The training loss combines robust regression, grade supervision, and uncertainty regularization:

$$L = L_H(t, \hat{t}) + \lambda_1 L_{CE}(g, \hat{g}) + \lambda_2 [\exp(-\hat{\sigma})|t - \hat{t}| + \hat{\sigma}] \quad \text{Eq.(4)}$$

where L_H is the Huber regression loss, L_{CE} is the cross-entropy grade loss, $\hat{\sigma}$ represents predicted aleatoric uncertainty, and λ_1 and λ_2 are balancing coefficients. The final predicted thickness map for a scene is obtained by overlapping patch inference and confidence-weighted aggregation:

$$T(p) = \frac{\sum_{m=1}^M \omega_m(p) \hat{t}_m(p)}{\sum_{m=1}^M \omega_m(p)}, \omega_m(p) = \exp(-\hat{\sigma}_m(p)) \quad \text{Eq.(5)}$$

Aggregation reduces block artefacts at the patch boundary and gives less weight to uncertain predictions. The model is trained using AdamW, cosine learning-rate decay, random rotation, mild speckle perturbation and wind-condition stratified mini-batches. No reference citations are used in this section because the section defines the implementation used in this paper.

Patch sampling is balanced by thickness grade and sea-state interval during training, not by raw patch count. Without this balancing, large areas of thin sheen would be present in the mini-batches, and the model would underfit moderate- to thick-oil. Data augmentation is intentionally conservative. Rotations and flips are permitted because slick orientation in an image is not physically linked to the model output after radiometric calibration. Strong intensity scaling is avoided to prevent the loss of absolute damping information that carries thickness evidence. Speckle perturbation is applied as a mild multiplicative noise term to improve tolerance to filtering differences. The validation criterion is the scene-level RMSE after overlapping patch aggregation, not just the patch-level loss, because the final product required by response users is a continuous map.

The three outputs of the inference stage for each scene are a continuous thickness map, a grade map and an uncertainty map. A continuous map will be chosen for quantitative analysis and volume approximation. A grade map of the operating communication divides the response team into executable zones. Areas with boundary mixing, weak contrast or abnormal sea state in the uncertainty map are considered less reliable. These products

are produced by the same backbone and are therefore mutually consistent. The Design does not require a separate post-classification thickness conversion and avoids discontinuities at grade boundaries that could obscure the regression signal learned by PolSAR descriptors.

The selected implementation details are suitable for engineering applications of the model. The input patch size is 256×256 pixels, and the feature cube is downsampled only after the first convolutional layer to avoid losing narrow slick filaments too early. Layer Normalization is used in the context Transformer blocks, and Batch Normalization is kept in the convolutional stem to stabilize low-level SAR texture features. The regression head has two fully connected layers with GELU activation and dropout. The grade head shares the same pooled feature vector but has a separate classifier, so the categorical objective does not force a discretized output in the regression branch. Training stops when the validation RMSE does not improve for 20 epochs, and the checkpoint with the lowest scene-level validation error is used for testing. These Design Choices are not universal settings; rather, they are a compromise among retrieval accuracy, stable optimization and feasible scene-scale processing for polarimetric SAR images.

Experimental Data and Analysis

Experimental Data and Evaluation Settings

The experiments used a curated PolSAR benchmark composed of 6,240 labeled patches extracted from 39 offshore scenes, including controlled spill observations, verified response records, and manually quality-checked slick regions. Labels were converted into continuous thickness values via response-grade calibration and field-note interpolation, and the resulting thickness range was 0.01 mm - 0.42 mm. Training, validation and testing were divided at the scene level with a 60:20:20 ratio, and patches from the same acquisition were not included in multiple subsets because adjacent SAR patches can share sea-state structure and inflate apparent accuracy when spatial leakage is ignored [21]. The three groups of benchmark data had six wind-speed intervals from 2.1 m/s to 8.8 m/s and four thickness grades. Figure 3(a) shows the distribution of thickness grades across wind-speed intervals: sheen samples remained dominant at 2-3 m/s with 92 patches, whereas moderate-oil and thick-emulsion samples became more visible in the 4-6 m/s range, indicating that the benchmark did not concentrate all high-severity labels in a single sea-state interval. Figure 3(b) shows HV damping separability across the four thickness grades; the median response moved from approximately -1.8 dB for sheen to -3.55 dB for thick emulsion, but the widened thick-emulsion spread also indicates that strong damping alone is not sufficient for stable regression. This scene-level partition ensures that all patches from the same SAR acquisition remain within a single subset, thereby reducing spatial leakage between training and testing.




Figure 3. Data composition and polarimetric separability: (a) thickness-grade distribution under wind-speed intervals; (b) HV damping separability across thickness grades.

Patch Extraction Protocol aims to reduce spatial leakage and keep the slick areas in the test set physically separated. First, split each scene into slick interiors, slick borders, close-up clean water, and ambiguous

transition masks. Training patches were sampled with a minimum of 40% slick coverage for regression labels, and only boundary patches were kept if the field interpretation assigned a dominant thickness grade. Thus, the model was not required to learn a single thickness value from a patch that was mostly mixed water and oil. The total number of patches in the final test set was 1,248, and among them, 374 were sheen patches, 332 thin-film patches, 298 moderate-oil patches, and 244 thick-emulsion patches. The imbalance shows that thick oil has a small actual proportion, so both continuous regression metrics and grade-wise classification metrics are included in the evaluation reports to prevent the minority class from being overlooked in pooled averages.

The four operational thickness grades were used as the anchors for label conversion. Sheen samples were between 0.01-0.04 mm; the thin film was 0.04-0.12 mm; moderate oil exceeded 0.12-0.28 mm; and thick emulsion exceeded 0.28 mm. Field descriptions, visible slick texture and relative polarimetric damping were used to interpolate a continuous value in each interval. Interpolation was not considered the true ground truth; instead, it served as a practical engineering label under real-world monitoring conditions. To reduce label noise, two analysts independently annotated 18% of the patches, and disagreements of more than one grade were excluded from the training set. The retained double-checked subset had a grade agreement of 92.4% and a mean absolute label difference of 0.018 mm. These are the expected accuracy levels of the benchmark, and they serve as a reference for model errors.

Comparative Performance and Ablation Analysis

All channels were normalized at the scene level before model training and then clipped to the 1st and 99th percentiles. Speckle-filtered intensity channels were stacked with damping ratio, co-polarization ratio, local coefficient of variation and polarimetric texture descriptors. The proposed model was compared with a support vector regressor, Random Forest, ResNet-50, EfficientNet-B0, Swin-T, and a CoTNet backbone without SE recalibration. To keep the comparison consistent with the typical remote sensing evaluation practice, all methods used the same training samples, early stopping based on validation RMSE, and five repeated runs with different random seeds [22]. Table 1 shows the main numerical results. CoTNet-SE achieved the lowest MAE of 0.038 mm and RMSE of 0.056 mm, while the next best model, CoTNet without SE, obtained 0.045 mm MAE and 0.064 mm RMSE. The increase was relatively larger in the middle and thick parts than in the sheen section; therefore, channel recalibration was most effective for slicks with both mixed damping and internal texture.

Table 1. Quantitative comparison of oil spill thickness estimation methods on the scene-level test set

Method	MAE (mm)	RMSE (mm)	R ²	Macro-F1 (%)
SVR with PolSAR descriptors	0.092	0.128	0.642	67.4
Random Forest	0.083	0.117	0.701	70.6
ResNet-50	0.071	0.104	0.781	76.8
EfficientNet-B0	0.064	0.093	0.812	79.1
Swin-T	0.049	0.069	0.889	84.6
CoTNet	0.045	0.064	0.897	85.8
CoTNet-SE	0.038	0.056	0.913	87.8

The set of indicators will meet the requirements for science and technology. MAE is in units of mm and is the average absolute thickness error; RMSE is more sensitive to large errors that may lead to an incorrect response priority. The coefficient of determination is the proportion of the total variation that a model can explain. Macro-F1 is included because the auxiliary grade head is only effective when all four operating grades are managed reasonably well. Bias is shown in the figures separately because a model that consistently overestimates thin oil or underestimates thick oil can still distort volume estimates even if its average absolute error is small. All neural models have been averaged over five runs, and the standard deviation of RMSE for CoTNet-SE was 0.003 mm; it was higher than that of Swin-T (0.007 mm) and ResNet-50 (0.006 mm).

Hyperparameter sensitivity was checked for learning rate, SE reduction ratio, contextual kernel size, and auxiliary loss weight. The best validation behavior was achieved with an initial learning rate of $3e-4$, a reduction ratio of 8, a contextual kernel size of 7, and an auxiliary loss weight of 0.35. Increasing the weight of the auxiliary loss above 0.50 improved grade accuracy slightly but increased continuous RMSE by 0.006 mm because the latent space became more discrete. Reducing the size of the context kernel to 3 reduced memory consumption by 9.8%, but the boundary-region error increased by 0.011 mm. The above results determined the final configuration for the main comparison, and it was found that the model benefits from a moderate context window rather than being purely local or fully global in dependency.

Figure 4(a) shows the main regression errors of the methods and indicates that CoTNet-SE reduced RMSE by 46.2% compared with ResNet-50 and by 12.5% compared with the non-SE CoTNet baseline. The MAE curves had the same order, with ResNet-50 at 0.071 mm and CoTNet-SE at 0.038 mm; the bias was as low as 0.004 mm. Figure 4(b) shows the accuracy of different grades; the proposed model achieved 90.8% for sheen, 88.6% for thin film, 86.9% for moderate oil, and 84.7% for thick emulsion. A slight drop in thick-emulsion accuracy was observed at the boundary patches where emulsified oil and surrounding sheen appeared in the same area; this is consistent with the heterogeneous damping pattern caused by oil weathering in SAR observations [23]. Although Swin-T achieved a relatively high R^2 value, it had about 2.4 times the inference time of CoTNet-SE and showed a larger residual variance in low-wind samples; therefore, the contextual transformer block provided sufficient long-range modeling without the need for a heavy global attention structure.




Figure 4. Comparative model performance: (a) MAE, RMSE, and bias across regression models; (b) grade-wise accuracy and macro-F1 for the main deep models.

According to the above comparison, the traditional machine-learning regressor is not suitable for this task. The Support Vector Regressor had a relatively small error for thin-film samples but underestimated thick emulsions by about 0.061 mm because the feature space lacked expressiveness and could not capture both saturation and internal texture simultaneously. Random Forest reduced the underestimation to 0.047 mm, but it produced discontinuous predictions in adjacent patches and was therefore unsuitable for map products. ResNet-50 learned smoother spatial patterns but sometimes treated dark, homogeneous water in low-wind conditions as thick oil. EfficientNet-B0 improved upon ResNet-50 by strengthening channel mixing, but it still lacked an explicit mechanism for contextual slick morphology. The non-SE CoTNet baseline solved part of the morphology problem but assigned an unstable weight to HV texture in high wind. CoTNet-SE addressed both deficiencies by integrating contextual dependence and learned channel trust.

Figure 5(a) shows predicted versus measured thickness, where the fitted slope was 0.94 and the intercept was 0.006 mm. The scatter remained close to the 1:1 line between 0.05 mm and 0.28 mm, whereas the largest deviations appeared above 0.32 mm, suggesting that high-thickness saturation was the main source of regression compression. Figure 5(b) displays residual distributions by thickness grade; the interquartile residual range was -0.020 to 0.018 mm for thin film and -0.043 to 0.038 mm for thick emulsion, which confirms that error dispersion increased with slick heterogeneity. In operational slick interpretation, uncertainty-aware mapping is important because low-confidence thickness estimates can otherwise appear as precise numerical products [24]. Figure 5(c) shows calibration across ten confidence bins; expected calibration error, defined as the average gap between predicted confidence and observed reliability, decreased from 0.071 for ResNet-50 to 0.028 for CoTNet-SE. Figure 5(d) presents uncertainty intervals, where 91.6% of measured values fell within the predicted 90% interval, indicating that the uncertainty head was slightly conservative but not excessively wide. Expected calibration error is used here as the average gap between predicted confidence and empirical reliability across confidence bins, making the uncertainty comparison interpretable.




Figure 5. Regression reliability and uncertainty: (a) predicted versus measured thickness; (b) residual distribution by thickness grade; (c) calibration curves; (d) uncertainty interval coverage.

Residual inspection in thickness intervals showed that the model was most accurate between 0.06 mm and 0.22 mm, and although the polarimetric damping response was strong, it had not saturated. Below 0.03 mm, the median absolute error increased to 0.041 mm because weak sheen could be mistaken for natural film or low local roughness. Above 0.30 mm, the median absolute error was 0.052 mm; the main reasons were thick emulsion areas with internal streaks and patch-scale dispersions of weathered oil. The residual distribution was slightly right-skewed for thin samples; thus, a conservative overestimation occurred when the model detected strong co-polarized damping but weak cross-polarized texture. For response planning, this behavior is better than a serious underestimation of hazardous oil, but it should still be limited when using the predicted map for volume calculation.

Figure 6(a) shows the ablation errors after adding contextual attention, SE recalibration, auxiliary grade supervision and uncertainty weighting. Starting from a convolution-only baseline, contextual attention reduced the RMSE to 0.070 mm, SE recalibration further reduced it to 0.062 mm, and the full dual-task uncertainty model reached 0.056 mm. The same panel shows that SE recalibration reduced the positive bias in the thick-emulsion class from 0.031 mm to 0.014 mm and improved the handling of saturated damping. As different polarimetric components can be informative under different roughness conditions, they were combined with the error trajectory for analysis [25]. Figure 6(b) reports attention and excitation diagnostics: the mean excitation weight for the VV damping channel increased from 0.42 in clean-water-dominant patches to 0.71 in moderate-oil patches, while HV texture received higher weight under wind speeds above 6 m/s. This channel shift indicates that the SE module did not simply amplify one dominant band, but adjusted the evidence source according to patch conditions.




Figure 6. Ablation and learned recalibration: (a) error changes after adding contextual attention, SE gating, grade supervision, and uncertainty weighting; (b) channel excitation diagnostics across sea states.

Ablation experiments have also shown the function of the auxiliary grade branch. After excluding the grade branch, the continuous RMSE rose from 0.056 mm to 0.060 mm, and the macro-F1 dropped from 87.8% to 81.9%. The regression head still learned a reasonable continuous mapping, but the order of thin film and moderate oil became less stable in mixed patches. The auxiliary branch served as a regulariser to encourage the latent representation to maintain the boundary of operational severity. Uncertainty weighting had a smaller impact on the average error than the contextual and SE modules, but it improved calibration and reduced large residuals. The 95th percentile absolute error decreased from 0.149 mm without uncertainty weighting to 0.121 mm in the full model. A small number of large errors can still be a serious cause of risk in the response.

Robustness was evaluated by grouping test samples according to wind speed, incidence angle, and slick boundary complexity, since operational SAR monitoring studies have emphasized that acquisition geometry can be as consequential as pooled mean accuracy [26]. Under low wind, medium wind, and high wind, the proposed model achieved RMSE values of 0.050 mm, 0.054 mm, and 0.067 mm, respectively. The error increase at high wind was mainly caused by the reduced contrast between thin oil and roughened clean water, while far-range incidence further weakened the separability of several damping cues. Figure 7 shows the robustness matrix, where the highest error cell was 0.083 mm for high wind and far-range incidence, whereas the lowest error cell was 0.043 mm for medium wind and mid-range incidence. The model-maintained R2 above 0.86 in every wind-incidence group. This pattern shows that the learned representation remained usable under common acquisition variability, while also revealing that high wind and far-range geometry should be treated as a higher-risk combination in operational interpretation [27]. Far-range incidence is treated as a higher-risk condition because it can weaken the separability of damping cues and change the balance between co-polarized contrast and texture evidence.

	Near range 24-31°	Mid range 31-38°	Far range 38-45°
Low wind 2-4 m/s	0.047 mm	0.043 mm	0.058 mm
Medium wind 4-6 m/s	0.052 mm	0.049 mm	0.061 mm
High wind 6-9 m/s	0.071 mm	0.067 mm	0.083 mm Stress cell

Contours: 0.060 and 0.075 mm operational stress thresholds

Figure 7. Robustness matrix of RMSE under wind-speed and incidence-angle groups.

Boundary complexity is quantified by a normalized perimeter-area index of slick masks. Patches in the lowest complexity third had an RMSE of 0.049 mm, and patches in the highest complexity third reached 0.073 mm. It was to be expected; irregular boundaries generally indicate shear, wind-driven fragmentation or partial mixing with clean water. Therefore, the model increased the predicted uncertainty in the vicinity of fragmented edges. The mean uncertainty values were 0.043 mm for compact slick interiors and 0.071 mm near the irregular boundary. Thus, the human analyst can focus on the portion that requires manual examination or observation by the air again. It will not show a high-precision map that is inconsistent with the SAR data either.

Robustness to inclination varied among the thickness grades. Near-range observations maintained a relatively strong co-polarized contrast for the thin film and thus had an MAE of 0.031 mm; however, far-range observations increased this value to 0.044 mm. For a thick emulsion, the near-to-far difference was smaller; MAE changed from 0.047 mm to 0.052 mm because the damping signal was still relatively strong after the geometry modification. The excitation diagnostics showed a similar pattern: the HH channel had a relatively larger weight in near-range samples, and VV damping and local texture were more prominent in far-range samples. One reason the proposed architecture performed better than the intensity-only and non-SE variants was an adaptive channel shift.

The spatial mapping results were evaluated on complete scenes by aggregating overlapping patch predictions. Figure 8(a) shows the estimated thickness profile along a slick transect, where the model followed the field-interpreted transition from 0.03 mm sheen at the edge to 0.34 mm emulsion near the windward accumulation zone. Figure 8(b) compares cumulative oil-area curves and indicates that CoTNet-SE estimated 74.2% of the slick area below 0.10 mm and 8.6% above 0.30 mm, close to the interpreted values of 72.5% and 9.1%. Figure 8(c) maps residuals against distance to the slick boundary, where the median absolute error increased from 0.032 mm in interior regions to 0.059 mm within 120 m of the boundary. Jafarzadeh et al. [28] emphasized that SAR-based oil spill studies should connect detection performance with operational mapping requirements. Figure 8(d) reports inference cost, showing 38.4 frames per second on 256 x 256 patches and 7.8 million trainable parameters. The runtime values were lower than Swin-T at 15.9 frames per second and 27.6 million parameters, while the error remained below the EfficientNet-B0 baseline.




Figure 8. Scene-level mapping and efficiency: (a) transect thickness profile; (b) cumulative oil-area curve; (c) residual versus boundary distance; (d) inference speed and parameter count.

Scene-level aggregation was assessed at 50% patch overlap because lower overlap showed noticeable block transitions in the narrow slick arm. Confidence weighting reduced the transitions without adding an extra smoothing filter. In the largest test scene, the unweighted aggregation had a mean boundary discontinuity of 0.026 mm at the patch seams, and confidence weighting reduced this to 0.011 mm. This modification maintained the sharp transition at the true oil-water boundary by applying weights only to overlapping predictions and not to the spatial map after aggregation. The resulting maps kept the elongated high-thickness streaks along the windward side of the slick and did not spread thick-oil estimates into the surrounding clean water. Distinguish between emergency maps and regular maps to prevent artificial smoothing from expanding high-priority cleaning areas. The aggregation weights are derived from the uncertainty head rather than from an independent smoothing rule, which keeps the scene-level product connected to model confidence.

Batch inference was used to obtain the computational profile of a single workstation GPU. CoTNet-SE processed a 4096 x 4096 scene in 46.8 s with 50% overlap, and included preprocessing-free model inference and aggregation. Swin-T required 113.5 s under the same patch schedule, while EfficientNet-B0 required 39.1 s but produced lower accuracy. Memory usage for CoTNet-SE during inference was 2.6 GB, so tiling of large scenes could be performed without modifying the model. Therefore, the architecture is relatively feasible; it is heavier than a compact CNN but significantly lighter than a full transformer. In a response workflow, this runtime can be used to process several scenes or multiple augmentation-based uncertainty passes within an operational time window.

Conducted several additional experiments to determine whether the observed improvements were due to modifications in the model or favorable data partitioning. First, a scene-leave-one-region-out test produced an RMSE of 0.061 mm, only 0.005 mm above the main test result. Second, replacing polarimetric descriptors with intensity-only channels increased RMSE to 0.079 mm, indicating that polarization information provided additional benefit beyond spatial texture in this benchmark; this behavior is consistent with the report that polarimetric information improves discrimination among mineral oil, weak films and low-wind look-alikes when intensity contrast alone is ambiguous [29]. Third, removing stratified mini-batching increased the high-wind RMSE by 0.012 mm; thus, balanced sea-state exposure during training was required for generalization. Field-derived oil thickness labels are often affected by mixed slick states and uncertain weathering descriptions [30]. In the present error inspection, large residuals occurred most often where the label source described mixed oil or uncertain weathering state. Very thin sheen in a low-wind environment produced a false thickening in co-polarized darkness, but the HV texture was weak. Boundary filaments smaller than the effective SAR retrieval support are generally considered mixed pixels rather than homogeneous slick samples [31]. Thick emulsion mixed with dispersed oil produced a local underestimation because the patch contained both saturated damping and thin-film edges. Narrow boundary filaments with a smaller stride than the patch showed alternating overestimation and underestimation along the same slick arm, especially when the uncertainty score exceeded 0.07 mm.

Conclusion

This paper introduces an improved CoTNet-SE model for estimating the thickness of marine oil spills based on polarimetric SAR images. The method integrates contextual transformer representation, squeeze-and-excitation channel recalibration, polarimetric damping descriptors, and a dual regression-grade learning objective. The framework has addressed the first engineering problem of thickness retrieval: the thickness signal is weak, nonlinear and highly correlated with sea state and acquisition geometry. Generate a continuous thickness map, and then derive an operational grade map and an uncertainty map from the same representation to connect model output with the information forms required for practical offshore monitoring.

The current study is still limited. The basis for the benchmark is response-grade calibration and field-note interpolation, and therefore cannot fully replace direct in-situ thickness measurements. The model also works at the patch level before scene aggregation, so boundary regions with mixed oil and clean water are relatively more uncertain than slick interiors. In addition, the current implementation focuses on calibrated PolSAR inputs and has not yet fully addressed single-polarization emergency acquisitions. Therefore, the model should be viewed as a retrieval framework for well-prepared polarimetric observations rather than an all-purpose solution for all SAR oil spill scenes.

Future work will combine the proposed architecture with direct field sampling, multi-sensor transfers learning and physically constrained oil and gas aging models. Connect uncertainty-aware thickness maps with operational drift prediction and cleanup prioritization in a good way. Further research can be carried out to explore lightweight deployment on onboard or near-real-time processing systems for SAR-based thickness estimation that supports rapid offshore response. Expand the model to support partial-polarization and single-polarization inputs, and improve its applicability in emergencies when the acquisition mode cannot be selected in advance.

Author Contributions

Taner Başar contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Salih Övünç contributes to conceptualization, methodology, software, validation, analysis, investigation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

- [1] Chen, Y., & Wang, Z. (2022). Marine Oil Spill Detection from SAR Images Based on Attention U-Net Model Using Polarimetric and Wind Speed Information. *International Journal of Environmental Research and Public Health*, 19(19), 12315. <https://doi.org/10.3390/ijerph191912315>
- [2] Yin, J., Yang, J., Zhou, L., & Xu, L. (2020). Oil Spill Discrimination by Using General Compact Polarimetric SAR Features. *Remote Sensing*, 12(3), 479. <https://doi.org/10.3390/rs12030479>
- [3] Wang, D., Wan, J., Liu, S., Chen, Y., Yasir, M., Xu, M., & Ren, P. (2022). BO-DRNet: An Improved Deep Learning Model for Oil Spill Detection by Polarimetric Features from SAR Images. *Remote Sensing*, 14(2), 264. <https://doi.org/10.3390/rs14020264>
- [4] Zhang, J., Feng, H., Luo, Q., Li, Y., Wei, J., & Li, J. (2020). Oil Spill Detection in Quad-Polarimetric SAR Images Using an Advanced Convolutional Neural Network Based on SuperPixel Model. *Remote Sensing*, 12(6), 944. <https://doi.org/10.3390/rs12060944>
- [5] Zhang, J., Feng, H., Luo, Q., Li, Y., Zhang, Y., Li, J., & Zeng, Z. (2022). Oil Spill Detection with Dual-Polarimetric Sentinel-1 SAR Using Superpixel-Level Image Stretching and Deep Convolutional Neural Network. *Remote Sensing*, 14(16), 3900. <https://doi.org/10.3390/rs14163900>
- [6] Ma, X., Xu, J., Wu, P., & Kong, P. (2021). Oil spill detection based on deep convolutional neural networks using polarimetric scattering information from Sentinel-1 SAR images. *IEEE transactions on geoscience and remote sensing*, 60, 1-13. <https://doi.org/10.1109/tgrs.2021.3126175>
- [7] Zeng, K., & Wang, Y. (2020). A Deep Convolutional Neural Network for Oil Spill Detection from Spaceborne SAR Images. *Remote Sensing*, 12(6), 1015. <https://doi.org/10.3390/rs12061015>
- [8] Gao, S., Li, S., & Liu, H. (2022). Oil Spill Detection by CP SAR Based on the Power Entropy Decomposition. *Remote Sensing*, 14(19), 5030. <https://doi.org/10.3390/rs14195030>
- [9] Liu, J., Ren, P., Lyu, X., & Grecos, C. (2022). A Gamma-Log Net for Oil Spill Detection in Inhomogeneous SAR Images. *Remote Sensing*, 14(16), 4074. <https://doi.org/10.3390/rs14164074>
- [10] Espeseth, M. M., Brekke, C., Jones, C. E., Holt, B., & Freeman, A. (2020). The Impact of System Noise in Polarimetric SAR Imagery on Oil Spill Observations. *IEEE Transactions on Geoscience and Remote Sensing*, 58(6), 4194-4214. <https://doi.org/10.1109/tgrs.2019.2961684>
- [11] Fan, Y., Rui, X., Zhang, G., Yu, T., Xu, X., & Poslad, S. (2021). Feature Merged Network for Oil Spill Detection Using SAR Images. *Remote Sensing*, 13(16), 3174. <https://doi.org/10.3390/rs13163174>
- [12] Dong, X., Li, J., Li, B., Jin, Y., & Miao, S. (2023). Marine Oil Spill Detection from Low-Quality SAR Remote Sensing Images. *Journal of Marine Science and Engineering*, 11(8), 1552. <https://doi.org/10.3390/jmse11081552>

- [13] Zhai, J., Mu, C., Hou, Y., Wang, J., Wang, Y., & Chi, H. (2022). A Dual Attention Encoding Network Using Gradient Profile Loss for Oil Spill Detection Based on SAR Images. *Entropy*, 24(10), 1453. <https://doi.org/10.3390/e24101453>
- [14] Kumar, A., Mishra, V., Panigrahi, R. K., & Martorella, M. (2023). Application of Hybrid-Pol SAR in Oil-Spill Detection. *IEEE Geoscience and Remote Sensing Letters*, 20, 1-5. <https://doi.org/10.1109/lgrs.2023.3258224>
- [15] Taleghanidozdozan, S., Xu, L., & Clausi, D. A. (2022). Oil Spill Candidate Detection Using a Conditional Random Field Model on Simulated Compact Polarimetric Imagery. *Canadian Journal of Remote Sensing*, 48(3), 425-440. <https://doi.org/10.1080/07038992.2022.2055534>
- [16] Chaturvedi, S. K., Banerjee, S., & Lele, S. (2020). An assessment of oil spill detection using Sentinel 1 SAR-C images. *Journal of Ocean Engineering and Science*, 5(2), 116-135. <https://doi.org/10.1016/j.joes.2019.09.004>
- [17] Almulihi, A., Alharithi, F., Bourouis, S., Alroobaea, R., Pawar, Y., & Bouguila, N. (2021). Oil Spill Detection in SAR Images Using Online Extended Variational Learning of Dirichlet Process Mixtures of Gamma Distributions. *Remote Sensing*, 13(15), 2991. <https://doi.org/10.3390/rs13152991>
- [18] Li, C., Wang, M., Yang, X., & Chu, D. (2023). DS-UNet: Dual-Stream U-Net for Oil Spill Detection of SAR Image. *IEEE Geoscience and Remote Sensing Letters*, 20, 1-5. <https://doi.org/10.1109/lgrs.2023.3330957>
- [19] Li, C., Kim, D. J., Park, S., Kim, J., & Song, J. (2023). A self-evolving deep learning algorithm for automatic oil spill detection in Sentinel-1 SAR images. *Remote Sensing of Environment*, 299, 113872. <https://doi.org/10.1016/j.rse.2023.113872>
- [20] Hasimoto-Beltran, R., Canul-Ku, M., Díaz Méndez, G. M., Ocampo-Torres, F. J., & Esquivel-Trava, B. (2023). Ocean oil spill detection from SAR images based on multi-channel deep learning semantic segmentation. *Marine Pollution Bulletin*, 188, 114651. <https://doi.org/10.1016/j.marpolbul.2023.114651>
- [21] Yang, Y. J., Singha, S., & Mayerle, R. (2022). A deep learning based oil spill detector using Sentinel-1 SAR imagery. *International Journal of Remote Sensing*, 43(11), 4287-4314. <https://doi.org/10.1080/01431161.2022.2109445>
- [22] Temitope Yekeen, S., & Balogun, A. L. (2020). Advances in Remote Sensing Technology, Machine Learning and Deep Learning for Marine Oil Spill Detection, Prediction and Vulnerability Assessment. *Remote Sensing*, 12(20), 3416. <https://doi.org/10.3390/rs12203416>
- [23] Yang, J., Ma, Y., Hu, Y., Jiang, Z., Zhang, J., Wan, J., & Li, Z. (2022). Decision Fusion of Deep Learning and Shallow Learning for Marine Oil Spill Detection. *Remote Sensing*, 14(3), 666. <https://doi.org/10.3390/rs14030666>
- [24] Bianchi, F. M., Espeseth, M. M., & Borch, N. (2020). Large-Scale Detection and Categorization of Oil Spills from SAR Images with Deep Learning. *Remote Sensing*, 12(14), 2260. <https://doi.org/10.3390/rs12142260>
- [25] Huang, X., Zhang, B., Perrie, W., Lu, Y., & Wang, C. (2022). A novel deep learning method for marine oil spill detection from satellite synthetic aperture radar imagery. *Marine Pollution Bulletin*, 179, 113666. <https://doi.org/10.1016/j.marpolbul.2022.113666>
- [26] Mehta, N., Shah, P., & Gajjar, P. (2021). Oil spill detection over ocean surface using deep learning: a comparative study. *Marine Systems & Ocean Technology*, 16(3-4), 213-220. <https://doi.org/10.1007/s40868-021-00109-4>
- [27] Aghaei, N., Akbarzadeh, G., & Kosarian, A. (2022). GreyWolfLSM: 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>
- [28] Jafarzadeh, H., Mahdianpari, M., Homayouni, S., Mohammadimanesh, F., & Dabboor, M. (2021). Oil spill detection from Synthetic Aperture Radar Earth observations: a meta-analysis and comprehensive review. *GIScience & Remote Sensing*, 58(7), 1022-1051. <https://doi.org/10.1080/15481603.2021.1952542>
- [29] Mdakane, L. W., & Kleyhans, W. (2022). Feature Selection and Classification of Oil Spill From Vessels Using Sentinel-1 Wide-Swath Synthetic Aperture Radar Data. *IEEE Geoscience and Remote Sensing Letters*, 19, 1-5. <https://doi.org/10.1109/lgrs.2020.3025641>
- [30] Li, G., Li, Y., Hou, Y., Wang, X., & Wang, L. (2021). Marine Oil Slick Detection Using Improved Polarimetric Feature Parameters Based on Polarimetric Synthetic Aperture Radar Data. *Remote Sensing*, 13(9), 1607. <https://doi.org/10.3390/rs13091607>
- [31] Naz, S., Iqbal, M. F., Mahmood, I., & Allam, M. (2021). Marine oil spill detection using Synthetic Aperture Radar over Indian Ocean. *Marine Pollution Bulletin*, 162, 111921. <https://doi.org/10.1016/j.marpolbul.2020.111921>