

Binary Oil Spill Segmentation in Synthetic Aperture Radar Imagery Using U-Net

Ryan Alharbi and Hatem Keshk*

Cite <https://doi.org/10.64589/juri/231581>

Submitted: April 15, 2026 Revised: May 04, 2026 Accepted: May 05, 2026

ABSTRACT

Oil spills remain among the major environmental threats to marine ecosystems and coastal operations, making rapid and accurate detection essential for effective monitoring and emergency response. Synthetic aperture radar (SAR) imagery is widely used for this purpose because it enables daytime, nighttime, and all-weather conditions. Nevertheless, reliable oil spill segmentation remains challenging because SAR images are affected by speckle noise, low contrast, and visual similarities between oil-contaminated areas and surrounding sea surfaces. This paper presents a deep learning-based semantic segmentation framework using U-Net as the primary architecture for binary oil spill detection in SAR imagery with potential integration into UAV-assisted monitoring systems. The dataset comprised 1,112 Sentinel-1 SAR images acquired under diverse environmental and oceanic conditions. The workflow included SAR image preprocessing through logarithmic transformation, wavelet-based denoising, contrast enhancement using contrast-limited adaptive histogram equalization, and resizing to 256×256 px before training. The U-Net model was trained using the Adam optimizer, weighted sparse categorical cross-entropy loss, a batch size of 16, and 30 training epochs to improve segmentation performance under class imbalance. The results showed stable convergence, with the U-Net model achieving an accuracy of 93.33%, a mean Dice score of approximately 0.59, and a mean intersection over union above 0.52. Qualitative predictions also demonstrated accurate localization of major oil spill regions with clear boundaries and limited background noise. Benchmarking against a previously developed DeepLabV3+-based model confirmed the suitability of U-Net as a practical and computationally efficient solution for SAR-based oil spill monitoring and intelligent UAV deployment.

Keywords: Oil spill detection, SAR imagery, Semantic segmentation, U-Net, Deep learning, UAV monitoring

1. INTRODUCTION

Rapid and reliable oil spill detection is essential for supporting cleanup operations and reducing ecological and economic damage¹⁻³. Synthetic aperture radar (SAR) has become one of the most widely used remote sensing modalities for marine oil spill monitoring because it can acquire images over large ocean areas under day and night conditions, and is less affected by cloud cover than optical sensors⁴. In SAR imagery, oil-contaminated regions commonly appear as dark formations because oil films suppress short-gravity capillary waves and reduce radar backscatter from the sea surface⁵. However, this imaging mechanism also creates a major interpretation challenge, as several natural and environmental phenomena, including low-wind areas, wave shadows, biogenic films, and other sea surface effects, may produce similar dark signatures. Consequently, accurate oil spill detection in SAR images requires methods that can distinguish true oil spill regions from visually similar backgrounds and patterns. Earlier studies on SAR-based oil-spill detection relied mainly on manual interpretation, thresholding, feature extraction, and conventional classification algorithms^{5,6}. These methods provide an important foundation for automated oil-spill monitoring by identifying

dark formations and extracting descriptive features for classification. However, their performance is often constrained by speckle noise, low image contrast, changing sea states, and the irregular morphology of oil spill regions. These limitations are especially important in pixel-level segmentation tasks, where the objective is to detect the presence of oil and accurately delineate its spatial extent.

Recent advances in deep learning have shifted SAR oil spill analyses from handcrafted feature-based detection to automatic feature learning and semantic segmentation⁷. Krestenitis et al. demonstrated the use of deep neural networks for oil spill identification from satellite imagery and contributed to the development of publicly available SAR datasets, supporting the transition from image-level recognition to pixel-level oil spill delineation⁷. Zeng and Wang proposed OSCNet, a deep convolutional neural network for spaceborne SAR oil spill detection, and showed that learned hierarchical features can improve the automation and robustness of SAR-based oil spill analysis⁸. Shaban et al. introduced a two-stage deep learning framework wherein candidate regions were first identified and then segmented, emphasizing the importance of segmentation-oriented learning when oil spill

pixels occupy a relatively small portion of SAR scenes⁹. In addition, large-scale deep learning approaches highlighted the value of automated detection and categorization of oil spills in complex SAR imagery¹⁰. In this research direction, semantic segmentation architectures such as U-Net and DeepLabV3+ have become particularly relevant because they are designed to produce dense pixel-level predictions. U-Net is effective for segmentation problems requiring both contextual understanding and spatial localization, whereas DeepLabV3+ uses a more complex encoder-decoder structure with an atrous convolution to capture multiscale semantic information and refine the boundaries. Recent studies have explored improved U-Net- and DeepLabV3+-based models for SAR oil spill segmentation, highlighting the importance of both boundary preservation and robust foreground-background separation in complex marine scenes¹¹. However, the accurate binary delineation of oil spill regions remains challenging because the oil spill class is often highly imbalanced relative to the background, and SAR images may contain noise, weak boundaries, and visually similar non-oil regions. Accordingly, this study investigates a U-Net-based semantic segmentation framework for binary oil spill detection in SAR imagery. The proposed approach combines SAR-oriented preprocessing, including logarithmic transformation, wavelet-based denoising, contrast enhancement using contrast-limited adaptive histogram equalization (CLAHE), and class-weighted learning, to improve the separation between oil spill and non-oil regions under challenging imaging conditions. The objective is to evaluate whether the current U-Net configuration can provide reliable and computationally efficient binary segmentation performance using accuracy, Dice score, and intersection over union (IoU) as evaluation metrics. A previously developed DeepLabV3+-based model was also used as a benchmark to discuss the tradeoff between a more complex multiclass segmentation framework and simpler task-focused U-Net model. The main contributions of this study are as follows: first, it develops an SAR-oriented preprocessing pipeline tailored to binary oil spill segmentation; second, it evaluates a U-Net configuration designed to balance segmentation quality and computational efficiency; and third, it provides quantitative and qualitative benchmarking against a DeepLabV3+ framework to clarify the practical value of the proposed model for SAR-based marine monitoring and potential UAV-assisted deployment.

2. METHODOLOGY

2.1. Overall U-Net Framework. The proposed methodology was designed to identify oil-contaminated regions at the pixel level by transforming each input image into a corresponding binary segmentation mask. Sentinel-1 SAR images were used as the main research input, and the model was developed to distinguish oil spill pixels from non-oil background regions. The overall research design was based on an encoder-decoder segmentation architecture. U-Net was selected because it combines feature extraction and spatial localization within a single framework, making it suitable for image segmentation tasks where object boundaries must be preserved. The encoder progressively extracts high-level features from the input image, whereas the decoder reconstructs these features into a segmentation map with the same spatial dimensions as the original input. Skip connections between corresponding encoder and decoder layers preserve important spatial information and improve boundary reconstruction. This encoder-decoder structure is well-suited to binary semantic segmentation tasks that require both global contextual information and accurate spatial localization. In the SAR oil spill imagery, the foreground oil spill regions may be small, irregular, fragmented, and visually similar to the surrounding sea surface. Therefore, preserving the spatial details during reconstruction is important for maintaining spill boundaries. The skip connections in U-Net transfer low-level spatial information from the encoder to the decoder, supporting more accurate boundary recovery, whereas the deeper layers extract contextual features. In addition, U-Net provides a simple and computationally efficient segmentation framework compared with more complex architectures, making it appropriate for the present binary oil spill detection task and for future deployment-oriented monitoring applications where lightweight models are desirable. The implementation was performed using Python with the TensorFlow and Keras libraries. The proposed framework processes preprocessed SAR images and generates predicted binary masks representing the detected oil spill regions. The model output was used to evaluate the segmentation quality and effectiveness of the overall detection pipeline. Figure 1 presents the overall architecture of the proposed U-Net-based oil spill segmentation framework.

The standard U-Net architecture relies on local convolutional operations, limiting its ability to capture long-range dependencies in SAR imagery. This limitation can affect the segmentation




Figure 1. Overall architecture of the proposed U-Net-based oil spill segmentation framework




Figure 2. Image preprocessing and U-Net-based segmentation workflow for oil spill detection in Sentinel satellite imagery

performance when oil spill regions are spatially fragmented or distributed over large areas. Future work may incorporate attention mechanisms or transformer-based modules to enhance the global context modeling.

2.2. Image Preprocessing and Segmentation Workflow. The preprocessing and segmentation workflow was designed to improve the visibility of oil spill features in the SAR imagery before model training and prediction. In the proposed procedure, each Sentinel image was first resized to 256×256 px to ensure a uniform input size for the network and reduce computational complexity. As SAR data are affected by speckle noise and intensity variations, several preprocessing operations were applied before the images were passed to the U-Net model. First, a logarithmic transformation was applied to stabilize the intensity distribution of the SAR images. Wavelet-based denoising was then used to reduce speckle noise while preserving relevant structural information⁶. After denoising, contrast enhancement was performed using CLAHE to improve the visual separation between oil spill regions and surrounding background¹³. Finally, the processed images were normalized before being used as model inputs. These steps improved input data quality and supported more accurate segmentation performance. The segmentation workflow began by feeding the preprocessed SAR image into the U-Net model, which generated a binary mask representing the predicted oil spill regions. In the output mask, the foreground class corresponds to the detected oil spill pixels, whereas the background class represents non-oil regions. This workflow enables pixel-level identification of contaminated areas and provides a clear visual representation of the segmented spill patterns. Figure 2 illustrates the complete preprocessing and segmentation workflow applied in this study using Sentinel satellite imagery and the corresponding predicted output masks.

3. EXPERIMENTAL WORK

3.1. Dataset Description and Preparation. Our experimental dataset comprised 1,112 SAR images collected from the Sentinel-1 satellite under diverse environmental and oceanic conditions. The use of satellite data improved the dataset variability and supported the robustness of the proposed model under different observation conditions. Each image was associated with

the corresponding segmentation mask used for supervised learning. Since the present work focused on binary semantic segmentation, the mask preparation procedure was designed to separate oil spill regions from non-oil background areas¹⁴. The original mask images were resized to a uniform spatial resolution of 256×256 px to match the network input dimensions and reduce computational cost. During mask preparation, the oil spill class was retained as the foreground class, whereas all remaining pixels were assigned to the background class. This produces binary masks suitable for training the U-Net model. To prepare the image data for model development, all SAR images were first resized to 256×256 px. As shown in Table 1, the dataset was divided into training and testing subsets using an 80/20 split, corresponding to 889 training and 223 testing images. This division allows the model to be trained on a larger subset while preserving an independent set for performance evaluation. The overall dataset preparation procedure ensured consistency between the input images and output masks, and provided a suitable foundation for the subsequent training and evaluation stages.

3.2. Training Configuration and Implementation.

The proposed oil spill segmentation model was implemented in Python using the TensorFlow and Keras libraries. A U-Net-based encoder-decoder architecture was used to perform binary semantic segmentation, wherein each pixel was classified as either oil spill or non-oil background. After dataset preparation, the model was trained using the processed SAR images and their corresponding binary masks. The training process employed a weighted sparse categorical cross-entropy loss function to address class imbalance between oil spill and background pixels.

Table 1. Dataset summary and preparation details

Parameter	Description
Total images	1112
Data source	Sentinel-1
Image type	SAR images
Segmentation task	Binary oil spill segmentation
Classes	Oil spill (foreground), non-oil region (background)
Data split	80% training (889 images), 20% testing (223 images)

This weighting strategy improves learning by assigning greater importance to the minority class during optimization. As shown in Table 2, the model was trained using the Adam optimizer with a batch size of 16 for 30 epochs. Several callback functions were used to improve training stability. Early stopping was applied to reduce overfitting, model checkpoint was used to save the best-performing model, and ReduceLROnPlateau was employed to lower the learning rate when the validation performance stopped improving. This configuration provided a suitable balance among convergence stability, computational efficiency, and segmentation performance.

3.3. Evaluation Protocol. The performance of the proposed U-Net model was evaluated using quantitative metrics and qualitative visual inspection. Since the present study focused on binary oil spill segmentation, the evaluation procedure was designed to measure how accurately the predicted masks matched the corresponding ground-truth masks at the pixel level.

Three main performance metrics were used in this study: accuracy, dice score, and IoU^{15,16}. The accuracy was used to measure the overall proportion of correctly classified pixels in the segmentation output. However, because binary oil spill segmentation is affected by class imbalance, accuracy alone is insufficient to fully reflect model performance¹⁷. Therefore, the Dice score was used to evaluate the overlap between the predicted oil spill and ground-truth regions, whereas IoU was used to quantify the intersection area relative to the union of the predicted and reference masks. These two metrics provide a more reliable assessment of segmentation quality, particularly for foreground regions with irregular boundaries. In addition, a qualitative assessment was performed by visually comparing the selected predicted masks with their corresponding input SAR images. This step was used to examine the ability of the model to localize major oil spill regions, preserve boundary shape, and reduce false detections in the background. Furthermore, the obtained U-Net results were compared with a previously developed DeepLabV3+ model to provide a benchmark reference for the segmentation performance. Table 3 summarizes the evaluation metrics used in the study.

4. RESULTS AND DISCUSSION

4.1. Training Behavior and Quantitative Performance. The quantitative results of the proposed U-Net model demonstrated stable learning behavior and effective segmentation performance for binary oil spill detection in SAR imagery. The training history, summarized in Figures 3 and 4, shows the evolution of loss, accuracy, Dice score, and IoU during model

Table 2. Training configuration

Parameter	Value
Optimizer	Adam
Batch size	16
Number of epochs	30
Loss function	Weighted sparse categorical cross-entropy
Callbacks	Early Stopping, Model Checkpoint, Reduce LROn Plateau

Table 3. Evaluation metrics

Metric	Purpose
Accuracy	Measures the proportion of correctly classified pixels
Dice score	Measures the overlap between predicted and ground-truth spill regions
IoU	Measures the intersection area relative to the union of predicted and ground-truth regions

optimization. Although the training configuration was set to a maximum of 30 epochs, the recorded history available for analysis spanned 14 epochs, indicating that the callback-based strategy allowed the model to reach a stable learning stage without unnecessary additional iterations. As shown in Figure 3, the training loss gradually decreased from approximately 0.20 to 0.18, indicating continuous improvement in fitting the training data. The validation loss remained higher than the training loss, fluctuating between approximately 0.285 and 0.323. This behavior suggests that the model encountered a challenging validation set. However, the fluctuations remained limited and did not indicate severe instability or divergence. From a critical perspective, the gap between training and validation loss implies that the segmentation problem remains difficult due to noise, irregular spill boundaries, and strong visual similarity between the oil spill and non-oil regions in the SAR images. Nevertheless, the absence of a sharp increase in the validation loss suggests that the regularization strategy and callback settings were effective in controlling overfitting. The training accuracy gradually increased from approximately 0.925 to 0.932, whereas the validation accuracy remained high, fluctuating around 0.95–0.96. Although validation accuracy is often expected to be lower than training accuracy, the observed behavior can occur in imbalanced segmentation problems and should be interpreted carefully. In the present binary SAR segmentation task, background pixels dominate many images, and pixel-wise accuracy can be strongly influenced by the proportion of easily classified background pixels in each split. Therefore, a validation subset containing a larger proportion of background or fewer ambiguous pixels




Figure 3. Training and validation loss curves of the proposed U-Net model

may produce a higher accuracy, even when the validation loss remains higher than the training loss. In addition, the use of class-weighted sparse categorical cross-entropy, callback-based training, and regularization-oriented training control may render the training objective more sensitive to difficult minority-class oil spill pixels than the accuracy metric alone. Therefore, a higher validation accuracy does not necessarily indicate superior validation learning or an error in the experiment; rather, it highlights the limitation of using accuracy as the main performance indicator for imbalanced segmentation. The Dice score and IoU curves in Figure 4 provide a more informative assessment of foreground overlap, and should, therefore, be considered together with the accuracy and loss curves. The final model accuracy was approximately 93.33%, confirming strong pixel-level classification performance. However, accuracy alone does not fully reflect the segmentation quality in imbalanced datasets where background pixels dominate the image. Thus, the overlap-based metrics shown in Figure 4a and 4b provide a more meaningful interpretation of the model behavior. The Dice score results in Figure 4a indicate that the training Dice value increased progressively and reached approximately 0.594, while the validation Dice score remained above 0.60 throughout training and peaked near 0.629. Similarly, Figure 4b shows that the training IoU increased to approximately 0.529, whereas the validation IoU remained within the range of 0.543–0.562. These results demonstrate that the proposed U-Net model could learn meaningful spatial overlaps between the predicted and reference spill regions. The relatively consistent validation Dice and IoU values also suggest that the model generalizes reasonably well to the unseen data. Overall, the training curves indicated that the proposed framework achieved a satisfactory balance between convergence stability and segmentation performance. These results are important because they show that a comparatively simple U-Net architecture, when combined with suitable SAR pre-processing and class-weighted learning, can provide reliable oil-spill segmentation performance without requiring an excessively complex model design. This finding is significant in the context of current advances in deep-learning-based remote sensing, where practical accuracy, robustness, and computational efficiency are critical for real-world environmental monitoring applications.

4.2. Qualitative Segmentation Results. Figure 5 and Table 4 present the representative qualitative results of the proposed U-Net model, showing the SAR image, corresponding predicted oil-spill mask, and resulting segmentation behavior. The visual output confirms that the model can identify and segment oil-contaminated regions as distinct foreground areas even when the spill patterns appear irregular and spatially separated. In the illustrated example, the predicted mask highlights several candidate spill regions with clear contrast against the background, demonstrating the ability of the model to perform pixel-level delineation of oil spill features in SAR imagery. A key observation from the qualitative results is that the model can capture nonuniform spill structures rather than only simple compact regions.

This is important because real oil spills often exhibit elongated, fragmented, or discontinuous shapes due to sea surface motion, weather conditions, and image noise. The predicted mask in Figure 5 shows that the U-Net model successfully preserved the main foreground regions and separated them from the surrounding background. This indicated that the encoder–decoder structure and skip connections were effective in maintaining spatial information during segmentation. Meanwhile, the visual results also reveal several limitations. Some segmented regions appeared as isolated patches, and certain boundary shapes were not fully smooth or continuous. This suggests that the model may still produce minor false positives or imperfect edge reconstructions in challenging areas. Such behavior is expected in SAR-based oil spill detection, where weak spill signatures and background clutter can exhibit similar intensity characteristics. Therefore, although the model performs well in localizing the principal spill regions, additional refinement may further improve the boundary accuracy and reduce scattered detections. Overall, the qualitative analysis supported the quantitative findings presented in the previous subsection. The predicted masks were visually meaningful and demonstrated that the proposed U-Net framework can provide practical segmentation outputs for oil spill monitoring. These results are significant because they show that the model is not only numerically effective but capable of producing interpretable outputs that may assist marine surveillance and environmental assessment tasks. In addition, the qualitative evidence reinforces the practical value of the proposed framework




Figure 4. a) Training and validation Dice score curves of the proposed U-Net model; b) Training and validation IoU curves of the proposed U-Net model




Figure 5. Representative qualitative segmentation results of the proposed U-Net model showing the SAR image input and the corresponding predicted oil spill mask

Table 4. Comparison of the proposed model with related methods

Method	Reported Results
Proposed model	Training: Loss = 0.182, Accuracy = 93.33%, Dice = 0.594, IoU = 0.529; Validation: Loss = 0.301, Accuracy = 95.40%, Dice = 0.610, IoU = 0.546
U-Net (Ronneberger et al.) ¹⁸	Oil-spill IoU = 53.79%
LinkNet (Krestenitis et al.) ⁷	Oil-spill IoU = 51.53%
PSPNet (Krestenitis et al.) ⁷	Oil-spill IoU = 40.10%
DeepLabv2 (Krestenitis et al.) ¹²	Oil-spill IoU = 25.57%
DeepLabv3+ (Chen et al.) ¹⁹	Oil-spill IoU = 53.38%; mean IoU = 65.06%
Two-stage framework (Shaban et al.) ⁹	Accuracy = 92%, Precision = 84%, Dice = 80%
OSCNet (Zeng and Wang) ⁸	Accuracy = 94.01%, Recall = 83.51%, Precision = 85.70%

for downstream monitoring scenarios, in which consistent spatial delineation is more informative than image-level classification alone.

4.3. Comparison with DeepLabV3+. To provide a broader perspective on the performance of the proposed U-Net framework, the obtained results were compared with those of a previously developed DeepLabV3+-based model. As illustrated in Figure 6, the DeepLabV3+ implementation used in this study was formulated as a multiclass semantic segmentation problem with five classes: sea surface, oil spill, lookalike, ship, and land, whereas the current U-Net framework was designed for binary oil spill segmentation. Therefore, the comparison should be interpreted as a benchmark-based qualitative and performance-oriented discussion, rather than a perfect one-to-one numerical comparison. The DeepLabV3+ model employs an ImageNet-pretrained Xception backbone and is trained using 256×256 input images, an Adam optimizer, and sparse categorical cross-entropy loss. According to Table 5, the model achieved a test accuracy of approximately 95.37%, with precision, recall, and IoU values of approximately 0.9537 under the selected evaluation settings. Thus, the DeepLabV3+ framework provides a strong overall classification performance for multiclass segmentation tasks. In addition, qualitative examples show that the model can identify major labeled regions and generate masks that are visually consistent with the corresponding reference masks. From an analytical perspective, the stronger numerical performance of DeepLabV3+ can be associated with its more complex architecture and its ability to exploit richer semantic features through the Xception backbone and encoder-decoder design. This makes it well-suited for multiclass scene understanding, especially when multiple surface categories must be separated simultaneously.

However, this advantage is accompanied by greater architectural complexity and a heavier computational structure compared to the proposed U-Net model. In contrast, the proposed U-Net model was specifically optimized for the binary delineation of oil spill regions using SAR-oriented preprocessing and class-weighted learning. Although the reported overall accuracy is lower than that of DeepLabV3+, the U-Net framework demonstrates stable convergence, meaningful overlap scores, and visually interpretable segmentation outputs for binary tasks. This suggests that the proposed U-Net configuration offers a practical balance between segmentation accuracy, model simplicity, and computational efficiency. This balance is particularly important for real-time or resource-constrained environmental monitoring applications, where a lighter segmentation pipeline may be more desirable than a more complex multiclass architecture. Overall, the comparison indicated that both models were effective within their respective formulations. DeepLabV3+ exhibited strong multiclass segmentation capability and high global accuracy, whereas the U-Net framework provided reliable binary oil spill segmentation with a simpler, more task-focused design. These findings highlight that model selection should depend not only on numerical performance but also on the intended segmentation objective, class structure, and deployment requirements.

From a practical standpoint, the present findings should be interpreted in light of several constraints. The current dataset size, binary formulation, and challenging SAR background conditions may limit generalization across unseen acquisition scenarios. In addition, the comparison with DeepLabV3+ is informative but indirectly equivalent because the two models were developed for different labeling formulations. These considerations do not weaken the value of the present results, rather they define a clear path for future work involving harmonized

**Figure 6.** Representative qualitative results of the DeepLabV3+ model showing the input image, predicted mask, true mask, predicted RGB mask, and true RGB mask

Table 5. Performance comparison between the proposed DeepLab model and existing DeepLab-based methods

Model	Reported Results
Proposed DeepLab model	Test Accuracy = 95.37%; Precision = 0.9537; Recall = 0.9537; IoU = 0.9537
DeepLabv2 (Krestenitis et al.) ⁷	Oil-spill IoU = 25.57%
DeepLabv3+ (Krestenitis et al.) ⁷	Oil-spill IoU = 53.38%; mean IoU = 65.06%

benchmarking, larger datasets, and further robustness evaluations. Although modern architectures, such as attention-based U-Net, vision transformers, and AMBA-based models, have shown strong performance, this study focused on evaluating a computationally efficient baseline, U-Net, for deployment-oriented applications. Future work will include systematic benchmarking of these architectures.

5. CONCLUSIONS

This paper presents a U-Net-based framework for binary oil spill segmentation using Sentinel-1 SAR imagery. The proposed pipeline combines logarithmic transformation, wavelet-based denoising, CLAHE, normalization, and class-weighted learning to improve foreground-background separation under challenging SAR imaging conditions.

- The U-Net model showed stable training behavior and achieved an accuracy of approximately 93.33%, a mean Dice score of approximately 0.59, and a mean IoU above 0.52, indicating reliable pixel-level segmentation performance. The model localized major oil spill regions and capture irregular spill patterns, although scattered detections and imperfect boundary recovery remain challenging cases.
- A comparison with the previously developed DeepLabV3+ model showed that DeepLabV3+ provides a strong performance in a multiclass setting, whereas the proposed U-Net offers a simpler and more task-focused solution for binary oil spill delineation. Overall, U-Net is a practical and computationally efficient model for SAR-based oil-spill monitoring. Future work should consider larger and more diverse SAR datasets, direct benchmarking under identical experimental conditions, postprocessing to reduce scattered false detections, ablation analysis of preprocessing components, and deployment-oriented optimization for real-time UAV-assisted monitoring²⁰.

AFFILIATIONS AND AUTHOR DETAILS

Undergraduate Author

Ryan Alharbi – Department of Aerospace Engineering, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia;  0009-0008-5373-7029
Email: s202158430@kfupm.edu.sa

Corresponding Author

Hatem Keshk – Research Mentor, Post Doctoral Fellow, Interdisciplinary Research Center for Aviation and Space Exploration, King Fahd University of Petroleum & Minerals, Dhahran 31261, Saudi Arabia;  0000-0002-8043-1227
Email: hatem.keshk@kfupm.edu.sa

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support provided by the Department of Aerospace Engineering and Interdisciplinary Research Center for Aviation and Space Exploration at the King Fahd University of Petroleum and Minerals. The authors appreciate the technical and research environments that supported the completion of this study.

REFERENCES

- (1) Fingas, M. & Brown, C. E. A review of oil spill remote sensing. *Sensors (Switzerland)* vol. 18 Preprint at <https://doi.org/10.3390/s18010091> (2018).
- (2) Torres, R. et al. GMES Sentinel-1 mission. *Remote Sens. Environ.* 120, 9–24 (2012).
- (3) Dong, S., Feng, J., Gu, Z., Yin, K. & Long, Y. A Review of Artificial Intelligence and Remote Sensing for Marine Oil Spill Detection, Classification, and Thickness Estimation. *Remote Sensing* vol. 17 Preprint at <https://doi.org/10.3390/rs17223681> (2025).
- (4) Brekke, C. & Solberg, A. H. S. Oil spill detection by satellite remote sensing. *Remote Sensing of Environment* vol. 95 1–13 Preprint at <https://doi.org/10.1016/j.rse.2004.11.015> (2005).
- (5) Topouzelis, K. N. Oil spill detection by SAR images: Dark formation detection, feature extraction and classification algorithms. *Sensors* vol. 8 6642–6659 Preprint at <https://doi.org/10.3390/s8106642> (2008).
- (6) Donoho, D. L. *De-Noising by Soft-Thresholding*. *IEEE TRANSACTIONS ON INFORMATION THEORY* vol. 41 (1995).
- (7) Krestenitis, M. et al. Oil spill identification from satellite images using deep neural networks. *Remote Sens. (Basel)*. 11, (2019).
- (8) Zeng, K. & Wang, Y. A deep convolutional neural network for oil spill detection from spaceborne SAR images. *Remote Sens. (Basel)*. 12, (2020).
- (9) Shaban, M. et al. A deep-learning framework for the detection of oil spills from SAR data. *Sensors* 21, (2021).
- (10) Bianchi, F. M., Espeseth, M. M. & Borch, N. Large-scale detection and categorization of oil spills from sar images with deep learning. *Remote Sens. (Basel)*. 12, (2020).
- (11) Zhang, J., Yang, P. & Ren, X. Detection of Oil Spill in SAR Image Using an Improved DeepLabV3+. *Sensors* 24, (2024).
- (12) Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. <http://arxiv.org/abs/1412.6980> (2017).
- (13) Mishra, A. Contrast Limited Adaptive Histogram Equalization (CLAHE) Approach for Enhancement of the Microstructures of Friction Stir Welded Joints. <https://doi.org/10.48550/arXiv.1411.4038>.
- (14) Akshansh Mishra, Contrast Limited Adaptive Histogram Equalization (CLAHE) Approach for Enhancement of the Microstructures of Friction Stir Welded Joints, arXiv, <https://arxiv.org/abs/2109.00886>, (2021).
- (15) Costa, L. da F. Further Generalizations of the Jaccard Index. <http://arxiv.org/abs/2110.09619> (2021).
- (16) Shamir, R. R., Duchin, Y., Kim, J., Sapiro, G. & Harel, N. Continuous Dice Coefficient: a Method for Evaluating Probabilistic Segmentations. Preprint at <https://doi.org/10.1101/306977> (2018).

- (17) Taha, A. A. & Hanbury, A. Metrics for evaluating 3D medical image segmentation: Analysis, selection, and tool. *BMC Med. Imaging* 15, (2015).
- (18) Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol. 9351 234–241 (Springer Verlag, 2015).
- (19) Chen, Y. T., Chang, L. & Wang, J. H. Full-Scale Aggregated MobileUNet: An Improved U-Net Architecture for SAR Oil Spill Detection. *Sensors* 24, (2024).
- (20) Keshk, H. & Abdallah, A. Real-Time UAV-Based Oil Pipeline and Visual Anomaly Detection Using YOLOv26n: A Dataset and Edge-Deployment Study. *Drones* 10, 255 (2026).