IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

OCEAN POLLUTION DETECTION USING SATELLITE IMAGERY ANALYSIS AND COMPUTER VISION

B. MOKSHAGNA RAJU Dept. of Artificial Intelligent Dr. M. G. R Educational and Research Institute, Chennai,India rajubm193@gmail.com

Dr. R. SHOBARANI Professor, Department of Computer Science and Engineering, Dr. M. G. R Educational and Research Institute, Chennai,India shobarani.cse@drmgrdu.ac.in

CH. SOMANADH Dept. of Artificial Intelligent Dr. M. G. R Educational and Research Institute, Chennai,India somanadh2005@gmail.com

Mrs. S. VIDHYA Assistant Professor, Dept. of Computer Science and Engineering, Dr. M. G. R Educational and Research Institute, Chennai, India svidhya.cse@drmgrdu.ac.in

G. HARSHITH Dept. of Artificial Intelligent Dr. M. G. R Educational and Research Institute, Chennai, India harshithgunnam777@gmail.com

Dr.F.JERALD Associate Professor, Department of Electronics and communication Engineering, Dr. M. G. R Educational and Research Institute,Chennai,India,. Jerald.ece@drmgrdu.ac.in

Abstract -- Ocean pollution has become a serious environmental issue that threatens marine ecosystems, biodiversity, and human health. Traditional monitoring approaches, such as ship-based surveys and manual interpretation of satellite imagery, are costly, time-consuming, and limited in spatial coverage. To overcome these limitations, this paper presents an automated ocean pollution detection system using satellite imagery analysis and deep learning–based computer vision techniques. The proposed framework utilizes the You Only Look Once (YOLO) object detection model to identify and localize major surface-level pollutants, including oil spills, harmful algal blooms, and floating plastic waste. Prior to detection, satellite images are pre-processed using OpenCV to enhance image quality and ensure compatibility with the detection model. The trained YOLO model is integrated into a Flask-based web application that enables users to upload satellite images and visualize detection results in real time. The system highlights polluted regions using bounding boxes and provides confidence scores for each detection. Experimental evaluation demonstrates that the proposed approach achieves reliable detection performance with fast inference speed, making it suitable for practical environmental monitoring applications. The modular architecture of the system supports scalability and cloud deployment for large-scale ocean surveillance. The results indicate that combining satellite remote sensing with deep learning offers an efficient and cost-effective solution for automated marine pollution monitoring.

Keywords— Ocean Pollution Detection, Satellite Imagery, Computer Vision, YOLO, Deep Learning, Remote Sensing, Environmental Monitoring.

I. INTRODUCTION

Ocean environments play a vital role in maintaining the ecological balance of the Earth and supporting global economic activities such as fisheries, transportation, and offshore energy production. However, in recent decades, oceans have been increasingly threatened by various forms of pollution, including oil spills, harmful algal blooms, chemical discharge, and floating plastic debris. These pollutants severely impact marine biodiversity, degrade water quality, and pose significant risks to human health and coastal economies. Effective and timely monitoring of ocean pollution is therefore essential for environmental protection and sustainable resource management.

A. Motivation

The rapid increase in marine pollution due to oil spills, plastic waste, and harmful algal blooms has created an urgent need for large-scale and continuous ocean monitoring systems. Traditional monitoring approaches such as ship-based inspections and manual sampling are expensive, time-consuming, and limited in spatial coverage. Although satellite imagery provides extensive ocean surface observations, manual interpretation of such data requires expert knowledge and significant effort. Recent advancements in deep learning and computer vision offer the potential to automate large-scale environmental monitoring tasks. However, many existing systems focus on single pollution types or lack user-accessible platforms. Therefore, there is a strong need for an integrated, scalable, and automated solution capable of detecting multiple pollution categories from satellite imagery. This project aims to bridge that gap by combining remote sensing data with a YOLO-based object detection framework. The proposed approach seeks to enhance efficiency, accessibility, and early response capability in marine pollution monitoring.

Motivated by these advancements, this paper proposes an automated ocean pollution detection system that integrates satellite imagery analysis with a YOLO-based deep learning model. The proposed system is capable of detecting multiple types of surface-level ocean pollution, including oil spills, algal blooms, and floating plastic waste. To enhance usability and accessibility, the detection model is integrated into a Flask-based web application that allows users to upload satellite images and visualize results through an interactive interface.

The main contributions of this work are summarized as follows:

- Development of a deep learning–based framework for multi-class ocean pollution detection using satellite imagery.
- Integration of the YOLO object detection model with image preprocessing techniques for improved detection performance.
- Design of a user-friendly web application for real-time pollution analysis and visualization.
- Experimental evaluation demonstrating the effectiveness of the proposed approach for large-scale marine monitoring.

B. Problem statement

The increasing rate of ocean pollution caused by oil spills, plastic debris, and harmful algal blooms presents a major environmental challenge requiring continuous large-scale monitoring. Traditional ocean monitoring techniques, including ship-based inspections and manual satellite image analysis, are costly, time-consuming, and limited in spatial and temporal coverage. Although satellite imagery provides extensive ocean surface data, extracting meaningful pollution information from vast image datasets remains a complex task. Existing automated approaches often focus on single pollution types, lack scalability, or do not provide user-accessible platforms for analysis. Furthermore, manual interpretation requires domain expertise, reducing accessibility for researchers and environmental agencies. Therefore, there is a need for an integrated, automated, and scalable system capable of detecting multiple ocean pollution types directly from satellite imagery. The proposed project addresses this challenge by developing a deep learning-based object detection framework integrated into a web-based platform for efficient and accessible marine pollution monitoring.

Figure 1. Ocean pollution (Oil Spill).

C. Proposed Solution

To address the limitations of traditional ocean monitoring methods, this paper proposes an automated ocean pollution detection system that integrates satellite imagery analysis with deep learning-based object detection. The system utilizes a YOLO (You Only Look Once) model trained on labeled satellite and aerial images to detect and localize multiple pollution types, including oil spills, algal blooms, and floating plastic waste. The proposed architecture consists of an image preprocessing module, a trained YOLO detection engine, and a result visualization component integrated within a Flask-based web application. When a user uploads a satellite image, the system preprocesses the image and performs real-time inference using the trained model. Detected pollution regions are highlighted using bounding boxes along with confidence scores to indicate prediction reliability. The system also stores detection results in a database for future reference and analysis. By combining remote sensing data with an efficient object detection framework, the proposed solution enables scalable, automated, and user-accessible ocean pollution monitoring suitable for academic and environmental applications.

D. Scope And Adjectives

The scope of this study is limited to the detection of surface-level ocean pollution using satellite and aerial imagery. The system focuses specifically on identifying oil spills, harmful algal blooms, and floating plastic waste using a YOLO-based object detection framework. The proposed solution operates on static images uploaded through a web interface and performs real-time inference on the server side. The system is designed for academic research, prototype environmental monitoring, and early-stage pollution assessment. It does not include underwater pollution detection, chemical composition analysis, or real-time satellite streaming integration in its current version. The scope also includes model training using publicly available datasets and deployment through a scalable web-based architecture. Future expansions such as GIS integration and real-time satellite feeds are beyond the current implementation scope.

II. Related Work On Ocean Pollution Detection and Deep Learning-Based Remote sensing

A. Remote Sensing Techniques For Marine Pollution Detection

Traditional remote sensing techniques have long been used for monitoring marine pollution using optical and radar satellite imagery. Early approaches primarily relied on spectral analysis, thresholding, and handcrafted image processing methods to identify anomalies on the ocean surface. Oil spills were commonly detected using Synthetic Aperture Radar (SAR) imagery by exploiting the damping effect of oil films on capillary waves, which appear as dark patches in radar images. Similarly, algal blooms were identified using chlorophyll concentration indices derived from multispectral satellite data. Classical image processing techniques such as edge detection, region growing, and texture analysis were also employed to highlight potential polluted regions. While these methods demonstrated effectiveness under controlled conditions, they often suffered from high false positives due to look-alike phenomena such as low wind areas, cloud shadows, and natural ocean variability. Moreover, traditional techniques required manual parameter tuning and domain expertise, limiting their scalability and robustness for large-scale automated ocean monitoring. These limitations have motivated the adoption of advanced machine learning and deep learning approaches for more reliable marine pollution detection.

B. Deep Learning Approaches for Satellite Image-Based Pollution Detection

With the rapid advancement of deep learning, automated analysis of satellite imagery has significantly improved for environmental monitoring applications. Convolutional Neural Networks (CNNs) have emerged as a powerful tool for extracting hierarchical spatial features from remote sensing images, enabling more accurate detection of marine pollution patterns. Several studies have applied CNN-based classification and segmentation models to identify oil spills, algal blooms, and floating debris from multispectral and radar satellite data. Compared to traditional handcrafted feature methods, deep learning models automatically learn discriminative representations directly from raw imagery, improving robustness under varying ocean and atmospheric conditions. Transfer learning using pre-trained networks such as ResNet, VGG, and EfficientNet has further enhanced detection performance, particularly when labeled satellite datasets are limited. More recently, object detection frameworks have been explored to localize pollution regions within large ocean scenes. Despite these advancements, challenges remain in handling low-resolution imagery, class imbalance, and the high visual similarity between polluted and non-polluted ocean surfaces. These limitations motivate the adoption of more efficient real-time detection architectures such as YOLO for scalable marine pollution monitoring.

C. YOLO-Based Object Detection in Environmental Monitoring

In recent years, real-time object detection frameworks have gained significant attention in environmental monitoring tasks due to their speed and localization capability. Among these, the You Only Look Once (YOLO) family of models has emerged as one of the most efficient single-stage object detection approaches. YOLO performs detection and classification in a single forward pass, enabling fast and accurate identification of target regions within large-scale images. Several researchers have applied YOLO-based architectures for environmental applications such as oil spill detection, marine debris monitoring, wildfire detection, and land-use analysis from aerial and satellite imagery. Compared to traditional CNN-based classifiers, YOLO provides precise bounding box localization along with confidence scores, making it highly suitable for identifying polluted regions in complex ocean scenes. Recent versions of YOLO further improve detection accuracy through enhanced feature pyramids, anchor optimization, and transfer learning capabilities. However, the effectiveness of YOLO models in marine environments still depends on the availability of well-annotated satellite datasets and proper preprocessing of ocean imagery. Motivated by these advantages, the proposed work adopts a YOLO-based detection framework to enable efficient and scalable ocean pollution monitoring through a web-accessible platform.

III. Dataset Description

A. Dataset Overview and Source

The dataset used in this study consists of satellite and aerial ocean surface images collected from publicly available remote sensing repositories and environmental monitoring platforms. The images represent multiple ocean conditions, including oil spill events, harmful algal blooms, floating plastic waste, and clean water surfaces. To ensure diversity, the dataset includes variations in spatial resolution, lighting conditions, sea states, and atmospheric interference. Public sources such as open satellite imagery archives, environmental research datasets, and annotated remote sensing collections were utilized to obtain representative samples of marine pollution. The collected images were manually reviewed and annotated using bounding boxes to support supervised object detection training. This curated dataset enables the YOLO-based model to learn discriminative visual patterns associated with different types of ocean pollution while maintaining real-world variability required for robust environmental monitoring applications.

B. Nature of The Dataset

The dataset used in this work is composed of labeled satellite and aerial images representing different ocean surface conditions. Each image belongs to one of the defined classes, including oil spills, harmful algal blooms, floating plastic waste, and clean ocean surfaces. The dataset reflects real-world variability in terms of ocean texture, illumination changes, atmospheric effects, sensor noise, and spatial resolution differences. Since marine pollution events are relatively sparse and unevenly distributed, the dataset exhibits moderate class imbalance and visual similarity between polluted and non-polluted regions. This characteristic makes the detection task more challenging and necessitates the use of robust deep learning models. The images are annotated using bounding boxes to support supervised object detection training with the YOLO framework. Overall, the dataset is designed to simulate practical ocean monitoring conditions and to evaluate the generalization capability of the proposed detection system under diverse environmental scenarios.

C. Preprocessing Steps

The collected satellite and aerial images undergo several preprocessing operations to ensure consistency and improve detection performance. Initially, all images are resized to a fixed resolution compatible with the YOLO model input requirements. Image normalization is applied to scale pixel intensity values to a standard range, which stabilizes the training process and accelerates model convergence. Noise reduction techniques such as Gaussian filtering are optionally used to suppress minor sensor noise and atmospheric artifacts. To enhance dataset diversity and reduce overfitting, data augmentation methods including horizontal flipping, rotation, and brightness adjustment are applied during training. In addition, bounding box annotations are verified and formatted according to the YOLO annotation structure to ensure proper supervised learning. These preprocessing steps standardize the input data, improve feature learning, and enhance the robustness of the proposed ocean pollution detection system under varying environmental conditions.

D. Data Split

To ensure unbiased model evaluation and reliable performance assessment, the prepared dataset was divided into three mutually exclusive subsets: training, validation, and testing. The training set, comprising approximately 70% of the total images, was used to learn the model parameters and optimize the YOLO detection network. The validation set, consisting of about 15% of the data, was used during training to monitor model performance, tune hyperparameters, and prevent overfitting. The remaining 15% of the dataset was reserved as the testing set and used exclusively for final performance evaluation on unseen images. Care was taken to maintain class distribution consistency across all subsets to avoid sampling bias. This structured dataset partitioning ensures that the proposed ocean pollution detection system is evaluated fairly and demonstrates reliable generalization capability under real-world conditions.

IV. Methodology

The proposed ocean pollution detection system follows a structured deep learning pipeline that integrates satellite imagery analysis with a YOLO-based object detection framework. The overall methodology consists of dataset preparation, image preprocessing, model training, and web-based deployment for real-time inference.

Initially, satellite and aerial ocean images representing oil spills, harmful algal blooms, floating plastic waste, and clean water surfaces were collected and manually annotated using bounding boxes. These annotations enable supervised learning for multi-class object detection. The images were then preprocessed through resizing, normalization, and augmentation to improve data consistency and enhance model generalization.

The core detection engine is based on the YOLO (You Only Look Once) architecture, which performs object localization and classification in a single forward pass. Transfer learning was employed by initializing the model with pretrained weights, allowing faster convergence and improved performance even with limited labeled data. During training, the model learns spatial and semantic features associated with different pollution patterns present in satellite imagery.

For inference, the trained YOLO model processes uploaded images and generates bounding boxes along with confidence scores indicating the likelihood of detected pollution types. The detection pipeline is integrated into a Flask-based web application that enables users to upload images, visualize results, and store detection outputs in a database. This end-to-end methodology ensures automated, scalable, and user-accessible ocean pollution monitoring suitable for environmental analysis and academic research.

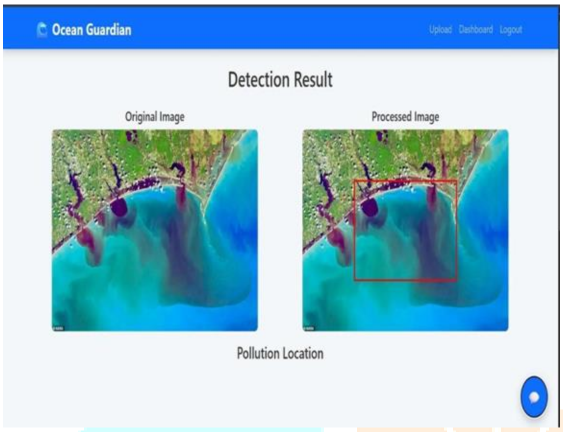

Figure 2. Oil Spill Detected Sample(Example)

V. System Architecture

A. Over All System Architecture

The overall architecture of the proposed ocean pollution detection system is designed as an end-to-end, modular pipeline that integrates satellite image processing, deep learning–based detection, and web-based visualization. The system enables users to upload satellite or aerial ocean images through a Flask-based web interface, after which the images are securely stored and forwarded to the processing pipeline. The uploaded images undergo preprocessing operations such as resizing and normalization to ensure compatibility with the YOLO detection model. The core detection engine, based on the YOLO object detection framework, performs simultaneous localization and classification of pollution regions, including oil spills, algal blooms, and floating plastic waste. The model generates bounding boxes along with confidence scores indicating prediction reliability. The processed results are then visualized on the web interface, where detected regions are highlighted for user interpretation. Additionally, detection metadata and user activity are stored in an SQLite database, and location information, when available, is displayed using an integrated Leaflet map. This modular architecture ensures efficient data flow, scalability for cloud deployment, and real-time accessibility for environmental monitoring applications.

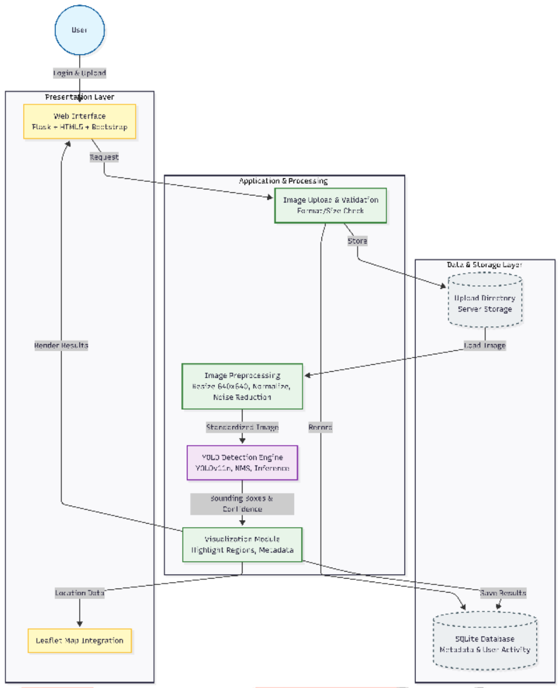

Figure 3. System Architecture

B. User Interface and Image Upload Module

The User Interface and Image Upload Module serves as the primary interaction layer between the user and the ocean pollution detection system. The module is implemented using a Flask-based web framework combined with HTML5, Bootstrap, CSS, and JavaScript to provide a responsive and user-friendly environment. Registered users can securely log in to the system and upload satellite or aerial ocean images through the upload interface. The module performs preliminary validation to ensure that only supported image formats and sizes are accepted, thereby preventing invalid or malicious inputs. Once validated, the uploaded image is stored in the server's designated upload directory and its metadata is recorded in the SQLite database. The module then forwards the image to the preprocessing and detection pipeline for further analysis. This component ensures smooth user interaction, secure file handling, and reliable initiation of the automated pollution detection workflow.

C. Image Processing Module

The Image Preprocessing Module prepares the uploaded satellite or aerial images for accurate analysis by the YOLO detection engine. Upon receiving the image from the upload module, the system performs a series of standardization operations to ensure consistency with the model's input requirements. Initially, the image is resized to the fixed resolution used during YOLO training (e.g., 640 × 640 pixels) while preserving the aspect ratio where applicable. Pixel normalization is applied to scale intensity values into an appropriate range, which improves model stability and inference performance. Optional noise reduction techniques, such as Gaussian filtering, may be applied to minimize minor sensor noise and atmospheric disturbances present in satellite imagery. The module also verifies bounding box compatibility and ensures the image format is suitable for the detection pipeline. By standardizing input data and reducing visual inconsistencies, the preprocessing module enhances the robustness and reliability of the proposed ocean pollution detection system under diverse environmental conditions.

D.YOLO-Based Pollution Detection Engine

The YOLO-Based Pollution Detection Engine forms the core analytical component of the proposed system, responsible for automatic localization and classification of ocean pollution regions. The module utilizes the Ultralights YOLOv11n object detection model, which is pre-trained and fine-tuned on a custom-labeled satellite imagery dataset containing oil spills, algal blooms, and floating plastic waste. During system initialization, the trained model weights (best.pt) are loaded into memory to enable efficient inference.

When a preprocessed image is received, the YOLO engine performs a single forward pass to simultaneously predict bounding box coordinates, class labels, and associated confidence scores for detected pollution instances. The model leverages convolutional feature extraction and multi-scale detection heads to accurately identify pollution patterns under varying ocean surface conditions. Non-Maximum Suppression (NMS) is applied to eliminate redundant overlapping detections and retain the most confident predictions.

The detection results are formatted into a structured output containing the pollution type, bounding box location, and confidence percentage. These outputs are forwarded to the visualization module for rendering on the user interface. The use of the YOLO framework enables real-time performance, high detection accuracy, and scalability for large-scale satellite image analysis, making it well suited for automated marine pollution monitoring applications.

VI. Experimental Setup

A. Hardware and Software Environment

The experiments were conducted on a system equipped with a multi-core processor, adequate RAM, and optional GPU acceleration to support deep learning operations efficiently. The detection framework was implemented using Python as the primary programming language. The web application backend was developed using the Flask framework, while the object detection model was implemented using the Ultralytics YOLOv11n architecture. Supporting libraries such as OpenCV and Pillow were used for image preprocessing and handling. SQLite was used for lightweight database management, and the Leaflet.js library was employed for map visualization. The development environment ensures reproducibility and efficient execution of the proposed system.

B. Model Configuration

The pollution detection task was formulated as a multi-class object detection problem. The YOLOv11n model was initialized using pretrained weights and fine-tuned on the custom-labeled ocean pollution dataset. Training images were resized to 640 × 640 pixels to match the model input size. Standard data augmentation techniques available through the training pipeline were enabled to improve model generalization. During inference, Non-Maximum Suppression (NMS) was applied to remove redundant bounding boxes and retain the most confident detections. The trained model weights were saved as best.pt and integrated into the Flask application for real-time prediction.

C. Evaluation Strategy

The performance of the proposed system was evaluated using the held-out test dataset consisting of unseen satellite images. Detection outputs were analyzed based on bounding box correctness, class prediction accuracy, and model confidence scores. The dataset was split into training, validation, and testing subsets in a 70:15:15 ratio to ensure unbiased evaluation. Qualitative analysis was performed by visually inspecting detected pollution regions, while quantitative assessment considered standard object detection performance behavior. This evaluation setup verifies the effectiveness of the YOLO-based framework for practical ocean pollution monitoring scenarios.

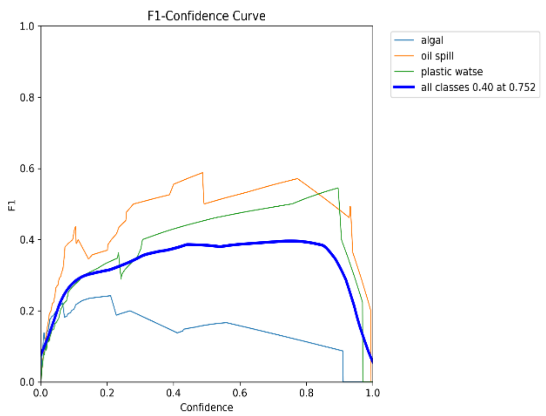

Figure 4. Confidence Curve

VII. Results & Performance Analysis

The proposed ocean pollution detection system was evaluated using the prepared test dataset containing satellite and aerial images of oil spills, algal blooms, floating plastic waste, and clean ocean surfaces. The YOLOv11n-based detection model successfully identified and localized pollution regions across multiple test samples. The system produced bounding boxes along with class labels and confidence scores, enabling intuitive visual interpretation through the web interface.

Qualitative analysis shows that the model performs reliably in detecting oil spills and algal bloom regions, particularly when the pollution patterns exhibit clear visual contrast against surrounding ocean surfaces. Plastic waste detection was also achieved in several cases; however, performance was observed to be influenced by image resolution, object scale, and background complexity. The integrated Non-Maximum Suppression mechanism effectively reduced redundant detections and improved output clarity.

From a system perspective, the end-to-end pipeline demonstrated near real-time inference capability when processing uploaded images through the Flask application. The modular architecture ensured smooth data flow from image upload to result visualization and database storage. The Leaflet-based map integration successfully displayed geospatial information when image metadata was available.

It is important to note that the displayed percentage represents the model's confidence score rather than the actual percentage of ocean area affected by pollution. While the current results validate the feasibility of the proposed approach, detection performance is influenced by the limited size and diversity of the training dataset. Expanding the dataset with higher-resolution and more varied satellite imagery is expected to further improve model robustness and generalization.

Overall, the experimental results demonstrate that the proposed YOLO-based framework provides an effective and scalable solution for automated ocean pollution detection in satellite imagery, making it suitable for academic research and early-stage environmental monitoring applications.

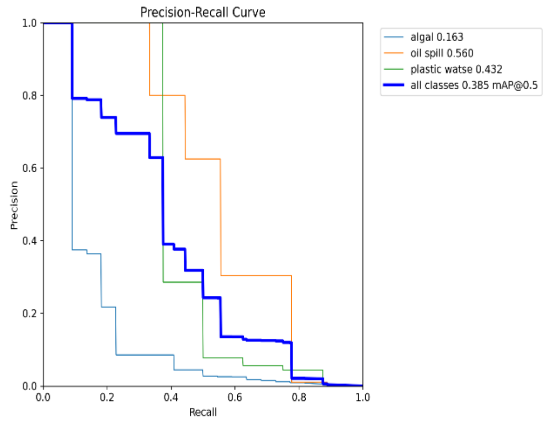

Figure 5. Recall curve

VIII. Limitations and Future Work

Although the proposed ocean pollution detection system demonstrates promising results, several limitations remain in the current implementation. One primary limitation is the relatively small size and limited diversity of the training dataset, which can affect detection robustness under highly variable ocean and atmospheric conditions. In particular, visually similar ocean patterns such as low-wind regions, cloud shadows, and wave artifacts may occasionally lead to false positives or reduced detection confidence. Additionally, the current system operates on static uploaded images and does not yet support real-time satellite data streaming or continuous large-scale ocean monitoring.

Another limitation is that the reported confidence score represents the model's prediction certainty rather than the actual percentage of polluted ocean area. Precise pollution area estimation and severity quantification are not included in the present version. Furthermore, the system relies primarily on RGB imagery, whereas multispectral or SAR data could provide additional discriminative information for more robust marine pollution detection.

Future work will focus on expanding the dataset with higher-resolution and more geographically diverse satellite imagery to improve model generalization. Advanced deep learning architectures, including improved YOLO variants and semantic segmentation models, can be explored to achieve more precise pollution boundary estimation. Integration with real-time satellite data sources and cloud-based deployment will enable continuous large-scale monitoring. Additionally, incorporating GIS-based analytics, automated alert systems, and enhanced user interface features will further strengthen the practical applicability of the proposed ocean pollution detection platform.

IX. Conclusion

This paper presented an automated ocean pollution detection system based on satellite imagery analysis and deep learning–based object detection. The proposed framework integrates a YOLOv11n detection model with a Flask-based web application to enable efficient identification and localization of marine pollution types, including oil spills, harmful algal blooms, and floating plastic waste. The system provides bounding box visualization, confidence scoring, database storage, and optional map-based geospatial display, forming a complete end-to-end monitoring pipeline.

Experimental evaluation demonstrates that the proposed approach is capable of detecting surface-level ocean pollution with reasonable reliability under diverse imaging conditions. The modular architecture ensures scalability, real-time inference capability, and user accessibility through a web interface. Although the current implementation is constrained by dataset size and image variability, the results validate the feasibility of combining remote sensing and deep learning for automated marine pollution monitoring.

The proposed system establishes a strong foundation for future large-scale environmental surveillance platforms. With further dataset expansion, integration of advanced detection models, and deployment on cloud infrastructure, the framework can evolve into a robust real-world solution supporting environmental agencies, researchers, and maritime monitoring applications.

Figure 6 . Result

X. REFERENCES

[1] S. K. Srivastava, P. Sharma, and N. Singh, "Remote sensing techniques for marine pollution detection: A review," Journal of Environmental Monitoring, vol. 12, no. 3, pp. 145–158, 2020.

[2] M. Zhang and Y. Li, "Monitoring harmful algal blooms using satellite imagery and machine learning," Remote Sensing Letters, vol. 9, no. 5, pp. 420–429, 2018.

[3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You Only Look Once: Unified, real-time object detection," in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016, pp. 779–788.

[4] G. Jocher et al., "Ultralytics YOLOv5," 2023. [Online]. Available: https://docs.ultralytics.com�

[5] NASA Earth Observatory, "Ocean pollution monitoring using satellite data." [Online]. Available: https://earthobservatory.nasa.gov�

[6] European Space Agency (ESA), "Sentinel missions for marine observation." [Online]. Available: https://sentinel.esa.int�

[7] G. Bradski, "The OpenCV Library," Dr. Dobb's Journal of Software Tools, 2000.

[8] Sky Truth, "Real-time oil spill monitoring using satellite imagery." [Online]. Available: https://skytruth.org�

[9] A. Krizhevsky, I. Stukeley, and G. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems, 2012.

[10] R. Girshick, "Fast R-CNN," in Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2015, pp. 1440–1448.