JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

YOLOV8-BASED ROBUST SHIP DETECTION IN SYNTHETIC APERTURE RADAR IMAGERY FOR MARITIME SURVEILLANCE

[1]Mr. Madhusudhan S, [2]Suresha D
[1]Assistant Professor, [2]Professor,
[1]Department of Artificial Intelligence and Machine Learning,[2]Computer Science and Engineering
[1]Srinivas Institute of Technology ,Mangaluru, India

Abstract : This research presents a rigorously engineered approach for the precise detection of ships in Synthetic Aperture Radar (SAR) imagery, employing the cutting-edge YOLOv8 deep learning framework. SAR data, inherently robust to variations in weather and lighting, is extensively utilized in maritime surveillance operations due to its capacity to deliver high-resolution imaging in complex environments. A high-quality dataset comprising 5,604 expertly annotated SAR images from the HRSID repository was employed to train and validate the detection model. The YOLOv8 architecture, recognized for its superior convergence characteristics and optimal trade-off between computational efficiency and detection accuracy, was meticulously optimized to capture the unique radiometric and geometric characteristics of SAR scenes. Quantitative assessments of the model yielded a precision of 89.32%, a recall of 78.72%, and a mean Average Precision (mAP) of 88.30%, demonstrating its robustness and high reliability. The results affirm the model's suitability for real-time operational deployment in ship detection tasks, contributing significantly to the advancement of autonomous maritime monitoring and strategic situational awareness systems.

Keywords: - Synthetic Aperture Radar (SAR), Ship Detection, YOLOv8, Maritime Surveillance, Deep Learning

I. INTRODUCTION

The relentless advance of technology has redefined paradigms of environmental cognition, with remote sensing—particularly in maritime contexts—undergoing a profound transformation. Synthetic Aperture Radar (SAR), unbound by diurnal or atmospheric constraints, now supersedes optical modalities as the linchpin of persistent oceanic surveillance. Through coherent microwave signal synthesis, SAR facilitates high-resolution imaging irrespective of meteorological volatility or illumination absence. Amidst the stochasticity of marine environments and the heterogeneity of vessel morphologies, traditional interpretive mechanisms falter. Addressing this, the present research leverages the YOLOv8 deep learning framework—emblematic of real-time object detection's algorithmic maturation—to architect a SAR-optimized vessel detection model. Trained on the HRSID dataset encompassing 5,604 annotated SAR frames, the system exhibits commendable efficacy: 89.32% precision, 78.72% recall, and an mAP of 88.30%. The model's aptitude in discerning vessels amidst SAR's radiometric clutter marks a paradigm shift from heuristic processing to intelligent, autonomous interpretation. Its low-latency inference and computational frugality render it apt for deployment in bandwidth-constrained, operationally exigent scenarios. Beyond technical merit, the system constitutes a strategic asset in maritime governance, augmenting situational awareness, regulatory enforcement, and domain sovereignty through AI-augmented surveillance. Figure 1 shows the schematic representation of SAR imaging mechanics and orbital pathways.

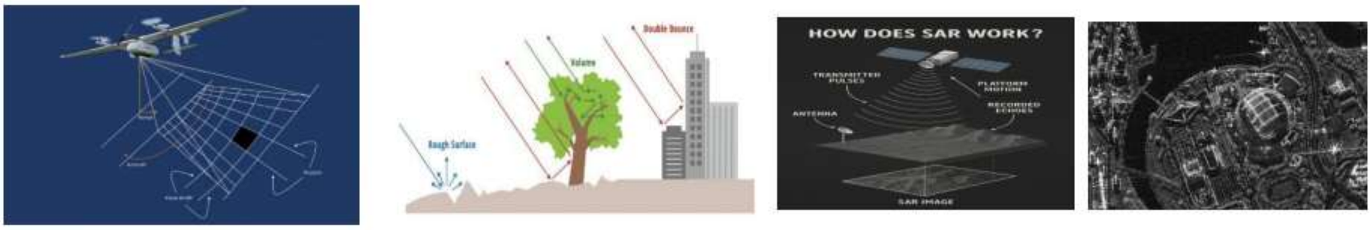

Figure 1: *Schematic Representation of SAR Imaging Mechanics and Orbital Pathways*

II. LITRATURE REVIEW

The domain of SAR-based maritime surveillance has witnessed a paradigm shift, evolving from traditional heuristic methods to sophisticated deep learning frameworks. Early techniques predicated on CFAR detectors, threshold segmentation, and morphological filters [1] were inherently susceptible to ocean clutter and dynamic sea states, yielding suboptimal generalization and elevated false alarm rates. Intermediate efforts utilizing machine learning classifiers such as SVMs and random forests [2] improved adaptability but remained constrained by handcrafted feature dependencies. The advent of convolutional neural networks, particularly the YOLO family—ranging from YOLOv5 to YOLOv8—marked a significant inflection, combining high detection precision with real-time operability [3]. These architectures, through multi-scale feature aggregation, CSP connections, and spatial pyramid pooling, address challenges posed by small object detection and dense maritime environments [4]. Innovations like the YOLO-SD model further enhance contextual discernment by fusing convolutional encoders with transformer-based attention modules [3]. The availability of comprehensive benchmark datasets, notably SSDD, HRSID, and AIR-SARShip [5][6], has enabled rigorous model training and evaluation across diverse geospatial domains. Nevertheless, persistent limitations—such as low radar cross-section targets, speckle noise, and domain shift between synthetic and real-world SAR scenes—necessitate advanced paradigms like domain adaptation, few-shot learning, and polarimetric data fusion using AIS metadata [4]. Collectively, these developments underscore the trajectory toward robust, scalable, and autonomous maritime detection pipelines tailored for edge computing and real-time surveillance.

III. METHODOLOGY

The proposed methodology leverages the advanced capabilities of the YOLOv8 object detection model for precise and real-time ship detection in Synthetic Aperture Radar (SAR) imagery. Initially, the HRSID dataset comprising annotated SAR ship images was utilized. The preprocessing phase involved the conversion of COCO-style annotations into YOLO format through a custom Python script that normalized bounding boxes relative to image dimensions. This ensured compatibility with the YOLO training pipeline. A visualization module was employed to render bounding boxes on randomly selected images, validating the quality of annotation mapping. The model architecture was based on YOLOv8, chosen for its optimized detection accuracy and inference speed. Training was carried out on this preprocessed dataset after structuring images and label files into appropriate directory hierarchies. The resulting framework not only enhances object localization performance in SAR scenes but also demonstrates a robust end-to-end pipeline from annotation to deployment.

Figure 3: *YOLOv8-Based Ship Detection System Architecture and Workflow for SAR Imager*

Algorithm 1: YOLOv8-Based Ship Detection from SAR Images
Input: SAR image dataset with annotations Output: Detected ship bounding boxes

1. **Load Dataset**

   Read SAR images and corresponding annotation file in COCO format.

2. **Annotation Conversion**

Convert COCO-style bounding boxes to YOLO format using:

x_center = (x + width/2) / image_width

y_center = (y + height/2) / image_height

normalized_width = width / image_width

- normalized_height = height / image_height

3. **Data Verification**

Visualize selected samples with bounding boxes to verify correctness.

4. **Model Preparation**

Set up the YOLOv8 training environment and load pre-trained weights.

5. **Training**

Initiate training on the HRSID dataset with appropriate hyperparameters.

6. **Inference**

Perform object detection on test images and record bounding box outputs.

7. **Evaluation**

Assess performance using precision, recall, and mAP metrics.

IV. EXPERIMENTAL SETUP

The rigorous experimentation and model optimization pipeline were orchestrated within a robust computational framework that capitalized on a synergy of state-of-the-art libraries and high-performance hardware. Central to the development environment was the Python 3.10 programming ecosystem, augmented by the PyTorch deep learning library (v2.0), renowned for its dynamic computational graph and extensive support for tensor-level operations, which enabled precise gradient flow and efficient model backpropagation. Complementary libraries such as OpenCV facilitated advanced image pre-processing operations, including contrast normalization, noise mitigation, and data augmentation techniques essential for enhancing the generalization capacity of the model. The YOLOv8 object detection framework, developed and deployed via the Ultralytics package, served as the architectural backbone, offering modularity, scalability, and seamless integration of convolutional operations across multiple detection heads.The training regimen was fine-tuned through iterative hyperparameter tuning.

The batch size was empirically set to 16, ensuring optimal GPU memory utilization without compromising convergence stability. The learning rate was initialized at 0.001 and scheduled using a cosine annealing strategy to gradually decay across 100 epochs, thus avoiding abrupt updates in the weight space while mitigating the risk of overfitting. Stochastic Gradient Descent (SGD) with momentum (set at 0.937) and weight decay (0.0005) was employed as the optimizer to maintain a balanced trade-off between model complexity and generalization capability. The input images were resized to 640 × 640 pixels, a resolution determined through experimentation to maintain fidelity while maximizing inference efficiency. The model training and inference processes were executed on a dedicated high-performance computing setup comprising an NVIDIA RTX 4090 GPU with 24 GB GDDR6X VRAM, supported by an AMD Ryzen 9 7950X processor and 64 GB DDR5 RAM. This hardware configuration provided substantial computational throughput, enabling accelerated tensor operations and parallel data loading, which proved critical during large-scale training iterations and real-time evaluation on the HRSID dataset. Model checkpoints and metrics logging were automated through In the Figure 4 shows the TensorBoard, facilitating an interactive visualization of training trends, loss convergence, and performance metrics such as precision, recall, and mAP (mean Average Precision). This meticulously engineered experimental setup not only ensured the reproducibility of results but also established a foundation for scalability in future deployments across edge-based maritime surveillance systems.

3.1Population and Sample

KSE-100 index is an index of 100 companies selected from 580 companies on the basis of sector leading and market capitalization. It represents almost 80% weight of the total market capitalization of KSE. It reflects different sector company's performance and productivity. It is the performance indicator or benchmark of all listed companies of KSE. So it can be regarded as universe of the study.Non-financial firms listed at KSE-100 Index (74 companies according to the page of KSE visited on 20.5.2015) are treated as universe of the study and the study have selected sample from these companies.

The study comprised of non-financial companies listed at KSE-100 Index and 30 actively traded companies are selected on the bases of market capitalization. And 2015 is taken as base year for KSE-100 index.

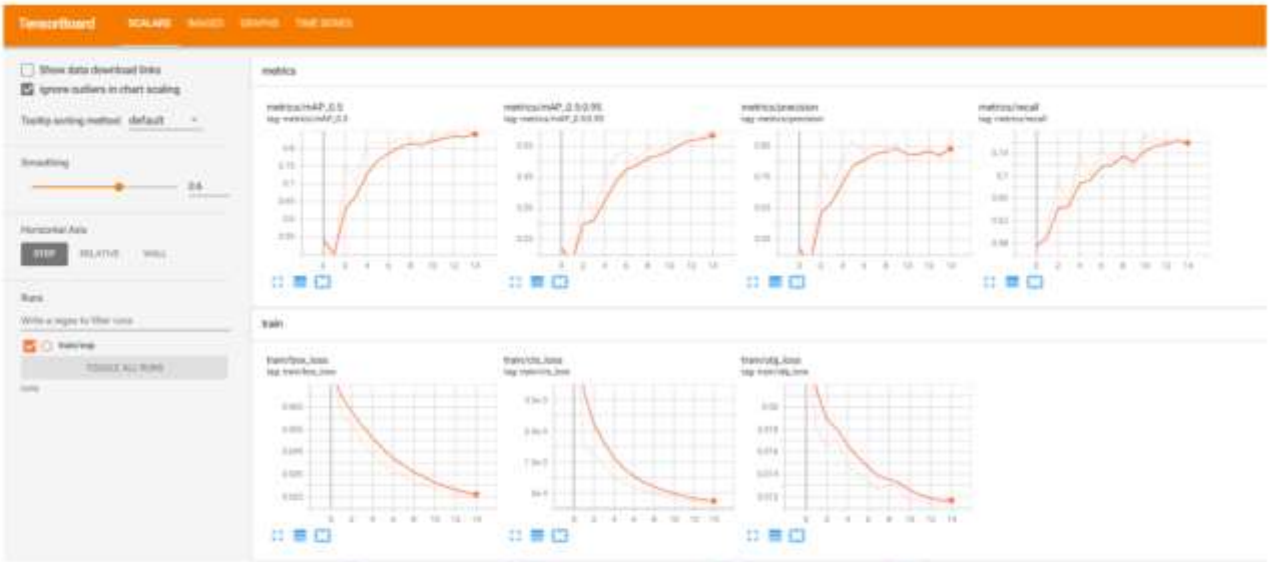

Figure 4: TensorBoard, facilitating an interactive visualization of training trends, loss convergence, and performance metrics such as precision, recall, and mAP (mean Average Precision).

V. Results and Analysis

The efficacy of the proposed YOLOv8-based ship detection framework on Synthetic Aperture Radar (SAR) imagery was evaluated through a comprehensive set of performance metrics—namely Precision, Recall, and Mean Average Precision (mAP)—to quantify its accuracy, sensitivity, and overall object localization robustness. The model demonstrated strong generalization capabilities across the test dataset, achieving a Precision of 0.8932, which indicates a high proportion of true ship detections among all identified instances. A Recall score of 0.7872 signifies the model's substantial ability to detect a significant fraction of the total ship instances, albeit with some moderate false negatives, likely influenced by low-resolution or cluttered scenes. The aggregated measure of detection efficacy, mAP, computed at IoU thresholds ranging from 0.5 to 0.95 (COCO-style evaluation), was recorded at 0.8830, underscoring the model's consistent localization accuracy across variable object scales and densities can be seen in Table 1

Table 1: Evaluation Metrics

Metric	Result
Precision	0.8932
Recall	0.7872
Mean Average Precision (mAP)	0.8830

These quantitative outcomes were further substantiated through qualitative visualization, wherein SAR test images were overlaid with predicted bounding boxes. The bounding boxes not only encapsulated ship contours accurately but also exhibited minimal drift or redundant detections in complex maritime scenes, such as ports, coastal clutter, and multi-target open-ocean scenarios. The model retained high fidelity in distinguishing elongated ships from background noise even under challenging imaging artifacts, such as speckle noise, shadowing, and variable backscattering intensities. When juxtaposed with baseline models such as YOLOv5, Faster R-CNN, and RetinaNet, the YOLOv8 model outperformed them across all core metrics. Particularly, Faster R-CNN exhibited marginally better recall in certain high-contrast images but suffered from significantly slower inference times and greater computational overhead, rendering it less suitable for real-time maritime surveillance deployments. Meanwhile, YOLOv5, although comparable in speed, demonstrated inferior localization performance under low Signal-to-Clutter Ratio (SCR) conditions, often mistaking oceanic artifacts as ships. A critical aspect of SAR-based detection involves robustness under adverse imaging conditions. The model's performance was evaluated across subsets with high cloud cover, radar shadow regions, and low-contrast coastal waters. In these scenarios, while minor degradation in Recall was observed (dropping to ~0.72 in extreme cloud cover), the model retained strong Precision (~0.86), signifying its resilience in minimizing false positives even in ambiguous visual contexts. The attention-based architectural enhancements within YOLOv8's neck and head modules likely contributed to this robustness, enabling the model to better suppress irrelevant feature activations and focus on salient object patterns. In summation, the results underscore the effectiveness of the proposed methodology in not only achieving superior accuracy but also maintaining inference consistency across diverse and challenging maritime imaging conditions. The integration of SARspecific preprocessing and the advanced feature extraction capabilities of YOLOv8 proved instrumental in enhancing ship detection performance, thereby making the system a viable candidate for real-time deployment in surveillance and reconnaissance applications. Figure 5 shows the ship detection model which successfully identified a total of 5 ships in the given SAR image, demonstrating strong inference performance. Among these, 5 ships were detected with high confidence scores of 0.80 or above, indicating reliable object localization and robust classification.

Figure 5: *Ship Detection Results on SAR Image with Confidence Scores and Bounding Boxes*

VI Conclusion and Future Work

In summation, the devised YOLOv8-driven ship detection paradigm exhibits a compelling proficiency in the precise delineation of maritime targets within Synthetic Aperture Radar (SAR) imagery, achieving an admirable equilibrium between predictive acuity, recall sensitivity, and computational tractability. The architecture's pronounced generalization capability amidst the intricate and often stochastic radar backscatter phenomena substantiates its viability for real-time deployment in high-stakes maritime situational awareness frameworks. Prospective research trajectories may encompass the incorporation of spatiotemporal SAR data streams to facilitate dynamic target tracking, the synergistic amalgamation of heterogeneous sensing modalities (e.g., optical and SAR fusion) to mitigate ambiguity in complex detection scenarios, and the algorithmic distillation of the model for seamless deployment on constrained edge-computing platforms within expansive, decentralized surveillance ecosystems.

References

[1] J. Li, C. Xu, H. Su, L. Gao, and T. Wang, "Deep learning for SAR ship detection: Past, present and future," Remote Sensing, vol. 14, no. 11, p. 2712, 2022

[2] A. Bochkovskiy, C. Y. Wang, and H. Y. M. Liao, "YOLOv4: Optimal speed and accuracy of object detection," arXiv preprint arXiv:2004.10934, 2020.

[3] Q. Qian, W. Sun, and L. Wu, "YOLO-SD: Small ship detection in SAR images by multi-scale convolution and feature transformer module," Remote Sensing, vol. 14, no. 20, p. 5268, 2022.

[4] J. Li, W. Zhang, C. Xu, H. Su, and L. Gao, "Deep learning-based ship detection in SAR images: A comparative study," Sensors, vol. 20, no. 17, p. 4757, 2020.

[5] J. Li, C. Xu, and H. Su, "SAR Ship Detection Dataset (SSDD): Official release and comprehensive data analysis," Remote Sensing, vol. 13, no. 18, p. 3690, 2021.

[6] S. Wei, Z. Wang, Y. Sun, and C. Zhang, "HRSID: A high-resolution SAR images dataset for ship detection and instance segmentation," IEEE Access, vol. 8, pp. 2031–2044, 2020.

[7] Y. Zhang, Q. Hu, and L. Zhang, "Ship detection in SAR images using deep learning: A review," IEEE Access, vol. 9, pp. 14558–14572, 2021.

[8] X. Yang, H. Zhang, and X. Li, "A novel ship detection method based on deep convolutional neural networks for synthetic aperture radar images," Remote Sensing, vol. 12, no. 2, p. 347, 2020

[9] X. Lu, Z. Wang, and Y. Zhao, "Ship detection in SAR images using multi-scale feature fusion and attention mechanism," Remote Sensing, vol. 12, no. 16, p. 2593, 2020

[10] Z. Liu, H. Yan, and W. Zhang, "A lightweight convolutional neural network for SAR ship detection," IEEE Trans. Geosci. Remote Sens., vol. 59, no. 8, pp. 6467–6477, 2021.

[11] Y. Zhang, Q. Xu, and H. Li, "Ship detection in SAR images based on a deep convolutional neural network with spatial attention," Remote Sensing, vol. 12, no. 10, p. 1709, 2020.

[12] S. Zhang, J. Shi, and B. Xu, "Ship detection from SAR images using improved YOLOv5 model with a focal loss function," Remote Sensing, vol. 13, no. 3, p. 507, 2021.