AI-driven multimodal data fusion for hazardous object detection in maritime and coastal environments*

Vitaly Brevus[1,*,†], Halyna Brevus[2,†], Roman Gashynskyi[2,†], Aser Kashosi[2,†] and, Roman Yuzefovych[3]

[1] *Ternopil Ivan Puluj National Technical University, 56 Ruska str., Ternopil 46001, Ukraine*
[2] *Rework-Space LLC, 10 Berezhans'ka str, office 82, Ternopil 46027, Ukraine*
[3] *Karpenko Physico-Mechanical Institute NAS of Ukraine, Department of Methods and Facilities for Acquisition and Processing of Diagnostic Signals, 5 Naukova Str., Lviv 79060, Ukraine*

Abstract

This paper introduces a novel AI-driven platform that integrates satellite imagery, unmanned aerial vehicle data, and Automatic Identification System signals into an analytics pipeline for hazardous object detection in maritime and coastal environments. The system leverages YOLO11 for object detection and a knowledge graph based on Blue Brain Nexus to achieve semantic interoperability. Results demonstrate the ability of the developed information technology to detect maritime debris, oil spills, and vessel activity, while enabling adaptive route planning and decision support. This approach provides a scalable framework for emergency response and environmental monitoring, aligning with current advances in Artificial Intelligence, machine learning, and applied modeling in information technologies.

Keywords

AI, multimodal data fusion, knowledge graph, satellite imagery, emergency response

1. Introduction

The frequency of natural disasters, maritime accidents, and climate-related hazards has increased significantly in recent years [1]. Traditional monitoring systems that rely solely on satellite data or unmanned aerial vehicle (UAV)-based inspections often fail to provide a complete and timely situational picture [2, 3]. There is a growing need for integrated approaches that combine multiple heterogeneous data sources with Artificial Intelligence (AI) to support emergency response [4, 5]. This is particularly relevant in the context of digital sovereignty and cloud infrastructure projects like Gaia-X [6]. For near-real-time hazardous object detection, and actionable insights for safer maritime operations within coastal areas or large rivers one needs the solution, merging satellite data with accurate drone imagery and heterogeneous data sources enriching the precision and accuracy of the system.

This paper presents a multimodal AI platform that merges satellite imagery, UAV data, and Automatic Identification System (AIS) signals into a unified knowledge graph, enabling near-real-time hazardous object detection and decision support. The goal of this paper is to demonstrate the feasibility of multimodal fusion for maritime monitoring, present a knowledge-driven architecture for semantic interoperability, and validate the system.

AdvAIT-2025: 2nd International Workshop on Advanced Applied Information Technologies: AI & DSS, December 05, 2025, Khmelnytskyi, Ukraine, Zilina, Slovakia

1* Corresponding author.
† These authors contributed equally.
v_brevus@tntu.edu.ua (V. Brevus); ceo@rework-space.com (H. Brevus); rgashynskyi@rework-space.com (R. Gashynskyi); akashosi@rework-space.com (A. Kashosi); roman.yuzefovych@ipm.lviv.ua (R. Yuzefovych)
0000-0002-7055-9905 (V. Brevus); 0009-0009-5614-0241 (H. Brevus); 0009-0002-1681-4378 (R. Gashynskyi); 0000-0002-9059-2358 (A. Kashosi); 0000-0001-5546-453X (R. Yuzefovych)

2. Related works

Research on AI for remote sensing has expanded rapidly [3], focusing on applications such as vessel detection, oil spill monitoring, and disaster assessment, based on developments of neural networks for image analysis [7, 8, 9]. Multimodal data fusion has been studied extensively [10], with approaches ranging from simple statistical techniques to advanced semantic integration. Knowledge graphs (e.g., Blue Brain Nexus) have emerged as effective tools for organizing heterogeneous data while ensuring interoperability [11]. Despite progress, existing works rarely integrate near-real-time AI models, multimodal fusion, and knowledge graphs into a single system for maritime safety and emergency response. Integration of cutting-edge object detection algorithms, such as YOLO, becomes critical for enhancing near-real-time situational awareness and rapid response in maritime environments. The YOLO object detection algorithm [9, 12] is used for visual recognition of quadcopter streaming videos. Cargo ships can use drones for monitoring dangerous sea routes. Having such drones in air and sea all the time would be too expensive and ineffective. Our platform combines available satellite and tracking data for preliminary analysis and deploy sea drones or UAVs when needed, to receive local accurate data, transform these and analyze them all in real time.

Relevant data, offered by monitoring services, can be classified by source:

- Copernicus Marine Service (CMEMS): provides comprehensive sea surface temperature, salinity, and currents data, along with marine ecosystem data and sea ice concentration and extent measurements.
- European Marine Observation and Data Network (EMODnet): offers bathymetry, seabed habitats, and human activities information, complemented by oceanographic data including tides and currents, as well as geological, biological, and chemical datasets.
- National Oceanic and Atmospheric Administration (NOAA): contributes ocean temperatures, salinity, and currents data, marine debris monitoring and habitat mapping services, and weather and sea surface data collected from buoys and drones.
- Ocean Observatories Initiative (OOI): delivers ocean temperature, salinity, and chemical properties data, seafloor imaging and topography information, and biological observations.

To collect necessary data, MariNeXt service [13] can be utilized to classify critical maritime hazards like oil spills and debris. Its robust predictions ensure effective hazard monitoring for maritime surveillance, even in challenging environments.

3. Methodology

This study presents the development and validation of a UAV-based object detection system for maritime surveillance applications. While the broader multimodal data fusion framework encompasses satellite imagery (Sentinel-1 SAR, Sentinel-2), Automatic Identification System (AIS) data, and knowledge graph integration, this paper specifically focuses on the UAV component utilizing YOLOv11 object detection algorithms.

3.1. System architecture overview

The proposed information technology implements a five-component modular pipeline architecture (Figure 1). The data ingestion layer aggregates heterogeneous sources including Copernicus Sentinel Hub API, UAV-mounted sensors, and AIS data streams. Data preprocessing is performed through the QueryOptima™ platform for harmonization and normalization. A Blue Brain Nexus-powered knowledge graph manages semantic data storage and interoperability. The analytics layer integrates YOLOv11 object detection with graph-based reasoning, while the visualization component provides real-time dashboards and geospatial interfaces for end-user interaction.

3.2. UAV-based object detection system

A Minimum Viable Product (MVP) of the system was developed and tested during the Case Study: 8th CASSINI Hackathon (November 2024) at the POLE Product Design Center, Ukraine.

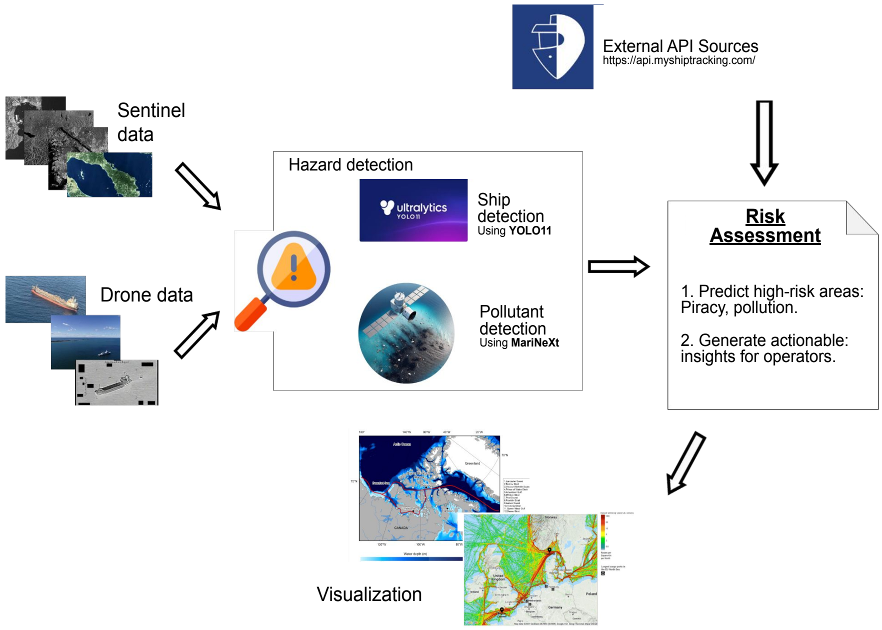

Figure 1: Diagram of AI-driven multimodal data fusion and risk assessment process.

The core contribution of this work centers on the development of a specialized UAV-based maritime object detection system employing the YOLOv11 architecture. The detection model was trained exclusively on the High-Resolution Ship Collections 2016 Multi-Scale (HRSC2016-MS) maritime dataset [14], representing a significant enhancement over the original HRSC2016 dataset.

Dataset Characteristics and Preprocessing. The HRSC2016-MS dataset comprises 1,680 high-resolution optical remote sensing images containing 7,655 annotated ship instances. The dataset exhibits comprehensive environmental diversity, encompassing maritime scenes across multiple operational conditions: sea and coastal environments, diurnal and nocturnal imaging scenarios, and varied meteorological conditions including clear and cloudy weather patterns. The dataset's multi-scale nature provides images with varying resolutions and aspect ratios, essential for training robust detection algorithms capable of identifying vessels across different scales and perspectives.

YOLOv11 Model Training and Fine-tuning. The YOLOv11 object detection architecture was selected for its demonstrated superiority in real-time object detection tasks and computational efficiency suitable for UAV deployment scenarios. The model was trained end-to-end on the HRSC2016-MS dataset, which was split into 60% for training, 20% for testing and 20% for validation, using transfer learning from pre-trained weights, with specific hyperparameter optimization for maritime object detection. Fine-tuning procedures incorporated UAV-specific imagery to enhance detection performance in coastal operational environments. The training process employed data augmentation techniques including geometric transformations, photometric adjustments, and multi-scale training to improve model generalization across diverse maritime conditions. Model convergence was monitored through validation metrics including precision, recall, and mean Average Precision (mAP) at multiple Intersection over Union (IoU) thresholds.

3.3. Integration framework

While this paper focuses on the UAV detection component, the broader system integrates with semantic knowledge representation through Blue Brain Nexus, enabling cross-modal data fusion and reasoning capabilities. The knowledge graph architecture facilitates semantic interoperability between UAV detection results, satellite imagery analysis, and AIS tracking data, supporting comprehensive maritimesituational awareness and risk assessment applications.

4. Results

The system was validated through three primary use cases: (1) detection and classification of oil spills, debris, and environmental anomalies using MariNeXt; (2) ship detection using Sentinel-1 SAR imagery (trained on SAR-Ship-Dataset[2]) split into 60% for training, 20% for testing and 20% for validation that was cross-validated with UAV imagery; and (3) prediction of vessel routes with anomaly detection using AIS data.

The comparison highlights that the YOLO11 model, fine-tuned with UAV imagery, demonstrates reliable ship detection in coastal environments. While some false positives occur, the overall bounding box alignment with labeled data confirms high precision (Figure 2).

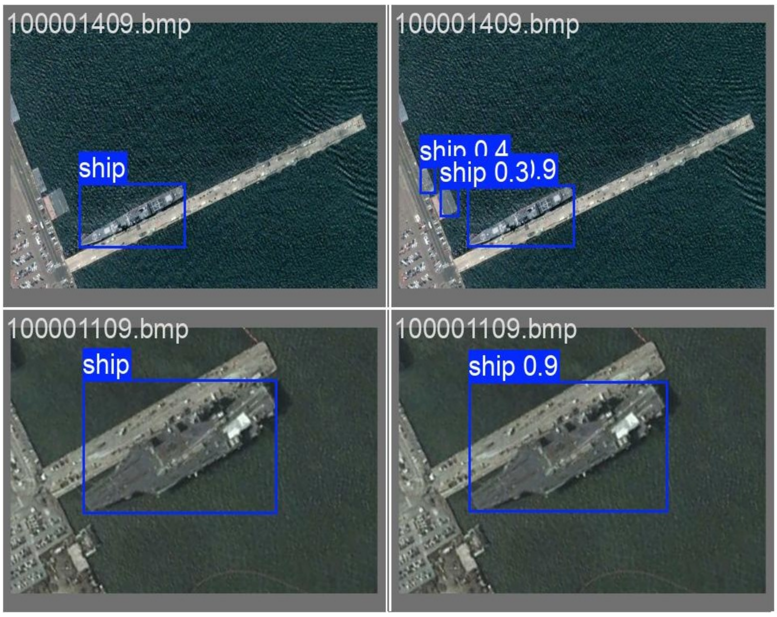

(a) (b)

Figure 2: Comparison of labeled (a) and predicted (b) ships.

The Figure 3 illustrates the evolution of validation metrics during model training. The initial drop in precision is an artifact of the learning rate scheduler and optimizer state at the beginning of the training. Precision rapidly stabilizes around 0.9, while Recall converges near 0.75. The mean Average Precision at Intersection over Union threshold 0.5 reaches approximately 0.85, and the stricter metric mAP@[0.5:0.95] converges to 0.66. These results indicate consistent model improvement across epochs and confirm robust generalization for maritime object detection.

All losses decrease steadily, with the most significant reduction observed during the first 20 epochs, after which the curves gradually converge and stabilize. The absence of divergence between training

[2]https://github.com/CAESAR-Radi/SAR-Ship-Dataset

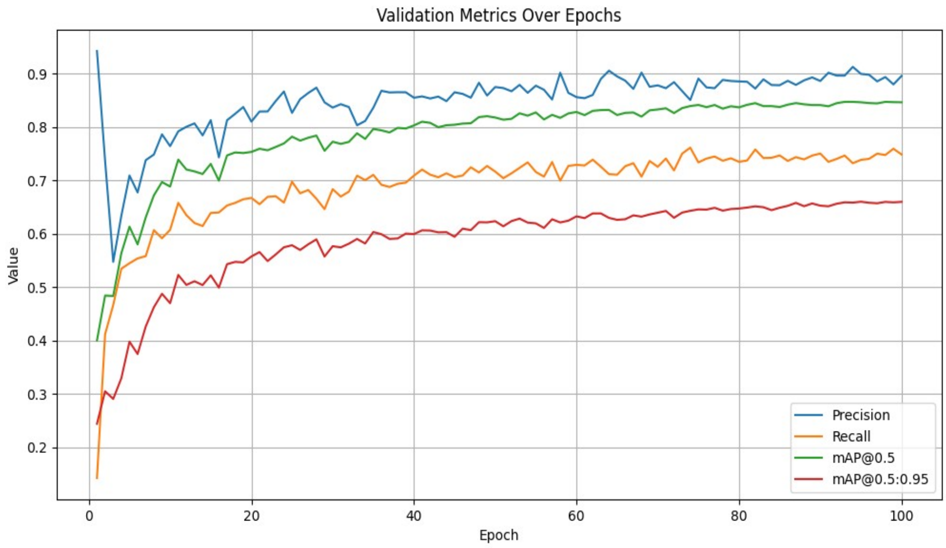

Figure 3: Validation metrics over training epochs.

and validation losses indicates that the model does not suffer from overfitting and generalizes well to unseen data (Figure 4).

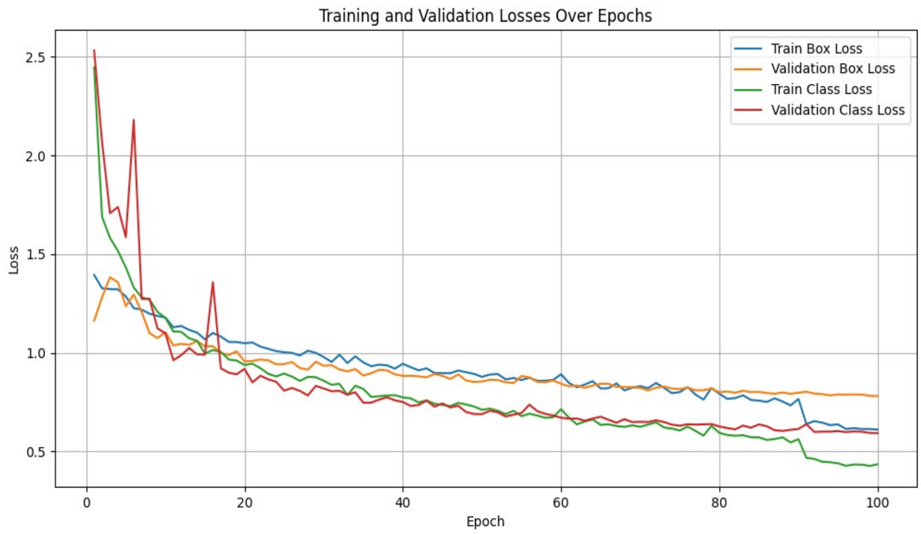

Figure 4: Training and validation losses over training epochs.

Therefore, the main results of the information technology validation can be summarized as follows:

- **Detection Accuracy**: The fine-tuned YOLO11 model achieved a precision of 90% and a recall of 75%, demonstrating robust performance in identifying maritime objects with UAV-enhanced data.
- **Processing Latency**: The analytics pipeline maintained an average processing time of under 5 seconds per frame, ensuring near-real-time performance suitable for time-critical applications.

- **Data Integration**: The platform successfully integrated multimodal data sources (satellite, UAV, and AIS) into a unified knowledge graph using Blue Brain Nexus, enabling effective semantic linking and querying.

The MVP was awarded third place at the national level of the 8th CASSINI Hackathon competition.

5. Discussion

The validation of the proposed information technology confirmed its strengths, including the provision of analytics with low latency, semantic interoperability through a knowledge graph, and a scalable pipeline applicable to multiple domains such as maritime, coastal, and disaster response. However, a few challenges were identified. These include limitations in data access, such as restricted AIS feeds and cloud dependency, the high computational demand of near-real-time pipelines, and the limited availability of labeled datasets for training YOLO11 on maritime hazards. Future work should focus on expanding the platform's capabilities beyond maritime safety to address other disaster types like floods and wildfires. Furthermore, important future developments include integration with European Data Spaces initiatives and ensuring compliance with FAIR data principles[3] for open science.

6. Conclusion

The proposed information technology demonstrates the feasibility of an AI-driven multimodal data fusion system for maritime safety and environmental monitoring. By combining satellite, UAV, and AIS data within a knowledge graph architecture, the system provides near-real-time hazardous object detection and decision support. The fine-tuned YOLO11 model achieved a precision of 90% and a recall of 75%, with a mAP@[0.5:0.95] of 0.66, demonstrating robust performance. Data owners will benefit from the developed information technology by maximizing the value of their data assets to provide high-quality insights for maritime security.

The ship detection model from Sentinel-1 SAR data performs exceptionally well achieving 93% precision and 92.5% recall. This demonstrates high accuracy in identifying ships with minimal false positives, critical for cross validation with AIS third party API data.

The results from the CASSINI Hackathon validate its effectiveness and highlight the potential of such approaches in broader applications, including emergency response, sustainable maritime operations, and environmental protection.

Acknowledgements

This work is based on the DroneSight project, which was awarded third place in the Ukrainian national selection of the 8th CASSINI Hackathon: EU Space for Defence and Security[4]. The authors gratefully acknowledge the CASSINI Hackathon, an initiative of the European Union, for providing the framework and premium access to satellite imagery that were instrumental to this research. Rework-Space LLC acknowledges support from the European Innovation Council through the Seeds of Bravery (UASEEDs) project.

Declaration on Generative AI

During the preparation of this work, the authors used GPT-4 and Grammarly in order to: Grammar and spelling check. After using these tools/services, the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

---

[3]https://www.go-fair.org/fair-principles/
[4]https://taikai.network/cassinihackathons/hackathons/euspace-defence-security/

References

[1] R. Khurana, D. Mugabe, X. L. Etienne, Climate change, natural disasters, and institutional integrity, World Development 157 (2022). doi:10.1016/j.worlddev.2022.105931.

[2] I. Chandran, K. Vipin, Multi-UAV networks for disaster monitoring: challenges and opportunities from a network perspective, Drone Systems and Applications 12 (2024) 1–28. doi:10.1139/dsa-2023-0079.

[3] J. Wang, K. Zhou, W. Xing, H. Li, Z. Yang, Applications, evolutions, and challenges of drones in maritime transport, Journal of Marine Science and Engineering 11 (2023). doi:10.3390/jmse11112056.

[4] S. P. H. Boroujeni, A. Razi, S. Khoshdel, F. Afghah, J. L. Coen, L. O'Neill, P. Fule, A. Watts, N.-M. T. Kokolakis, K. G. Vamvoudakis, A comprehensive survey of research towards AI-enabled unmanned aerial systems in pre-, active-, and post-wildfire management, Information Fusion 108 (2024). doi:10.1016/j.inffus.2024.102369.

[5] A. Bazrafkan, C. Igathinathane, N. Bandillo, P. Flores, Optimizing integration techniques for UAS and satellite image data in precision agriculture — a review, Frontiers in Remote Sensing 6 (2025). doi:10.3389/frsen.2025.1622884.

[6] R. Adler-Nissen, K. A. Eggeling, The discursive struggle for digital sovereignty: Security, economy, rights and the cloud project Gaia-X, JCMS: Journal of Common Market Studies 62 (2024) 993–1011. doi:10.1111/jcms.13594.

[7] I. Konovalenko, P. Maruschak, V. Brevus, Steel surface defect detection using an ensemble of deep residual neural networks, Journal of Computing and Information Science in Engineering 22 (2021). doi:10.1115/1.4051435.

[8] I. Konovalenko, P. Maruschak, V. Brevus, O. Prentkovskis, Recognition of scratches and abrasions on metal surfaces using a classifier based on a convolutional neural network, Metals 11 (2021). doi:10.3390/met11040549.

[9] L.-h. He, Y.-z. Zhou, L. Liu, W. Cao, J.-h. Ma, Research on object detection and recognition in remote sensing images based on YOLOv11, Scientific Reports 15 (2025). doi:10.1038/s41598-025-96314-x.

[10] Y. Li, M. El Habib Daho, P.-H. Conze, R. Zeghlache, H. Le Boité, R. Tadayoni, B. Cochener, M. Lamard, G. Quellec, A review of deep learning-based information fusion techniques for multimodal medical image classification, Computers in Biology and Medicine 177 (2024). doi:10.1016/j.compbiomed.2024.108635.

[11] M. F. Sy, B. Roman, S. Kerrien, D. M. Mendez, H. Genet, W. Wajerowicz, M. Dupont, I. Lavriushev, J. Machon, K. Pirman, D. Neela Mana, N. Stafeeva, A.-K. Kaufmann, H. Lu, J. Lurie, P.-A. Fonta,A. G. R. Martinez, A. D. Ulbrich, C. Lindqvist, S. Jimenez, D. Rotenberg, H. Markram, S. L. Hill, Blue brain nexus: An open, secure, scalable system for knowledge graph management and data-driven science, Semantic Web 14 (2023) 697–727. doi:10.3233/SW-222974.

[12] B. Zhao, J. Zhao, R. Song, L. Yu, X. Zhang, J. Liu, Enhanced yolo11 for lightweight and accurate drone-based maritime search and rescue object detection, PLOS ONE 20 (2025) 1–24. doi:10.1371/journal.pone.0321920.

[13] K. Kikaki, I. Kakogeorgiou, I. Hoteit, K. Karantzalos, Detecting marine pollutants and sea surface features with deep learning in Sentinel-2 imagery, ISPRS Journal of Photogrammetry and Remote Sensing 210 (2024) 39–54. doi:10.1016/j.isprsjprs.2024.02.017.

[14] W. Chen, B. Han, Z. Yang, X. Gao, Mssdet: Multi-scale ship-detection framework in optical remote-sensing images and new benchmark, Remote Sensing 14 (2022). URL: https://www.mdpi.com/2072-4292/14/21/5460. doi:10.3390/rs14215460.