Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Deep Learning Object Detection for Maritime Security: A Critical Review of Task Evidence, Benchmark Comparability and Deployment Constraints

Syed Muhammad Waqas
muhammad.waqas01@iqra.edu.pk
Iqra University
Muhammad Ahsan
Iqra University
Fahad Ur Rehman
Iqra University
Muhammad Bilal Jeeva
Iqra University
Mazliham Mohd Su'ud
Multimedia University
Romeis Riaz
Iqra University
Muhammad Usman Khan
Iqra University
Muhammad Mansoor Alam
Riphah International University

Research Article

Keywords: maritime security, object detection, deep learning, remote sensing, edge computing, multimodal fusion

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11592068/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Deep Learning Object Detection for Maritime Security: A Critical Review of Task Evidence, Benchmark Comparability and Deployment Constraints

Syed Muhammad Waqas[1*], Muhammad Ahsan[1], Fahad Ur Rehman[1], Muhammad Bilal Jeeva[1], Mazliham Mohd Su'ud[2], Romeis Riaz[1], Muhammad Usman Khan[1], Muhammad Mansoor Alam[3]

[1]Faculty of Engineering, Science and Technology, Iqra University, Karachi, Pakistan.
[2]Faculty of Computing and Informatics, Multimedia University, Cyberjaya, Malaysia.
[3]Faculty of Computing, Riphah International University, Islamabad, Pakistan.

*Corresponding author(s). E-mail(s): muhammad.waqas01@iqra.edu.pk;

Abstract

Deep learning has been used successfully for object detection in optical, infrared, SAR, and multispectral images of maritime environments. Nevertheless, its deployment in maritime security systems remains challenging. This review analyses major object detection architectures, sensing methods, application domains, and deployment constraints in this domain. The most frequently represented detection architectures in the reviewed papers include You Only Look Once (YOLO) detectors (199 papers) and attention-based and transformer models (123 papers). Tracking, multimodal fusion, threat reasoning, and detector security are much less represented. There are also major deficiencies in evaluation practices. For example, 121 out of 299 papers reporting detection metrics do not provide an intersection-over-union threshold value. Inference latency, energy consumption, and target hardware are also insufficiently reported in many studies. Besides detection accuracy, we discuss how security-task evidence affects benchmark selection and detector deployment in communication-aware settings. The review provides a reporting checklist and a roadmap for future research.

Future work should concentrate on multimodal benchmarks, small-object detection, cross-sensor generalisation, trustworthy sensor fusion, edge intelligence, and realistic evaluation of detectors.

Keywords: maritime security, object detection, deep learning, remote sensing, edge computing, multimodal fusion

1 Introduction

Maritime borders are vast, open, and impossible to patrol continuously, so states rely on sensing and the automated processing of the resulting data to build maritime domain awareness (Mitra et al. 2017; Müller and Brook 2010). The security tasks involved—countering piracy, smuggling and trafficking, illegal, unreported and unregulated (IUU) fishing, and unauthorised intrusion—share a common perception problem: detecting and identifying vessels and other objects at sea (Eski 2019; Okafor-Yarwood et al. 2024). Traditional surveillance combines cooperative reporting through the Automatic Identification System (AIS), coastal radar, and human operators, but AIS can be switched off or spoofed, radar struggles with small non-metallic craft in clutter, and sustained visual monitoring exceeds human endurance (Eaton et al. 2018). Deep learning (DL) for computer vision has changed this picture: convolutional neural networks (CNNs), one-stage detectors such as You Only Look Once (YOLO), and transformer-based detectors now localise and classify ships, small boats, swimmers, buoys and floating objects in optical, infrared, radar, synthetic aperture radar (SAR) and multispectral imagery (Hassan et al. 2026; Soloviev et al. 2020; Gamage et al. 2023). Maritime scenes, however, are harder than the general-purpose benchmarks on which most detectors are developed: targets can be a few pixels wide, the background of waves, glare and wakes is non-stationary, the sensor is often on a moving platform, and adverse weather is common (Zhang et al. 2017). The field is also young: 423 of the 655 primary studies reviewed here were published in 2024–2026.

1.1 Scope

This review covers detection and the closely related classification, segmentation, tracking and sensor-fusion methods for objects on or near the sea surface, using electro-optical, infrared, radar, SAR and multispectral sensing on shore, ship, unmanned aerial vehicle (UAV), unmanned surface vehicle (USV) and satellite platforms. AIS data are considered when they are fused with, or checked against, sensor detections; AIS is not treated as a substitute for non-cooperative sensing. Underwater acoustics and vessel-trajectory prediction without a detection or security element are out of scope. The corpus comprises 655 primary studies (2010–2026): 111 studies from exploratory and targeted searches, which anchor the narrative synthesis, and 544 studies from documented database searches, which enter the corpus-level counts and the technology–application analysis (Section 2).

The framing is security-oriented, but the evidence base is broader. Most studies name no specific task; they are retained because they supply the detection capability

on which security systems build. Claims about specific security applications therefore rest on far fewer studies than claims about general detection, and the two are kept apart throughout.

1.2 Contribution and relation to existing surveys

Fourteen related surveys were identified. They address selected sensors, platforms, tasks or technologies, including electro-optical detection and tracking (Prasad et al. 2017), spaceborne optical vessel detection (Kanjir et al. 2018), electro-optical sea-surface detection (Lyu et al. 2023), maritime vision datasets (Su et al. 2023), anomaly and threat detection (Gamage et al. 2023), small-object detection (Rekavandi et al. 2025), occlusion handling (Banu and Sidram 2019), CNN accelerators on field-programmable gate arrays (FPGAs) (Haijoub et al. 2025), the security of object detectors (AlZaabi et al. 2020), cyber threats to Internet of Things (IoT)-enabled maritime systems (Ashraf et al. 2023), unmanned surface vehicles (Boretti 2025) and underwater detection (Song et al. 2026). The closest recent overlaps are the review of civilian maritime detection and tracking (Alsaaq 2026) and the review of AIS–SAR–optical fusion (Hanmante et al. 2026). Table 1 compares the seven closest reviews with the present one. None of them examines, for the same corpus, whether detection studies provide evidence about security tasks, whether their reported results are comparable, and whether the detectors can operate within maritime communication limits; this joint analysis is the focus here.

The review makes five contributions: (i) a coding of 655 primary studies by contribution, technology, sensor, platform and task; (ii) an analysis of reported benchmark results and evaluation protocols that shows why many headline scores cannot be compared; (iii) a task-level synthesis that separates direct maritime-security evidence from adjacent rescue, navigation and environmental applications; (iv) an analysis of the link and compute constraints that determine where detectors run and what a system can transmit; and (v) a reporting checklist and a prioritized research agenda. All five are syntheses of published work; no new experiments are reported.

The review addresses five questions: (RQ1) how is the published corpus distributed between direct maritime-security tasks and adjacent perception applications? (RQ2) which detector families and enabling techniques dominate the literature, and what trade-offs can be inferred without treating incompatible benchmarks as a common leaderboard? (RQ3) which datasets and evaluation protocols support reported results, and where are the main gaps? (RQ4) how do maritime communication capacity and latency constrain detection deployment? (RQ5) what methodological, technical and operational priorities are supported by the evidence? Section 2 describes the review method. RQ1 is answered by the corpus profile (Sections 2 and 3) and the application analysis (Section 5); RQ2 by Section 4; RQ3 by Sections 4.3 and 5.3; RQ4 by Section 6; and RQ5 by Section 7. Section 8 states the limitations.

Table 1 The seven related surveys closest in scope to the present review, and the boundary between each and this review.

Survey	Primary emphasis	Boundary relative to the present review
(Prasad et al. 2017)	Electro-optical video processing for maritime object detection and tracking; includes the Singapore Maritime Dataset.	Predates much of the recent deep-learning, transformer and multi-sensor literature; its centre of gravity is EO video rather than security-specific evidence and deployment constraints.
(Kanjir et al. 2018)	Vessel detection and classification from spaceborne optical imagery.	Sensor-specific; does not provide a cross-modality comparison of detector evidence or a security-task taxonomy.
(Su et al. 2023)	Maritime vision datasets and benchmarks.	Dataset-centred; the present review connects dataset and protocol differences to reported detector results and application claims.
(Gamage et al. 2023)	Machine learning for maritime surveillance, anomaly and threat detection.	Behaviour- and threat-oriented; the present review additionally analyses object-detection architectures, dataset protocols and the gap between object detection and a verified security outcome.
(Rekavandi et al. 2025)	Small-object detection across domains, with maritime surveillance as a case study.	A cross-domain methodological guide rather than a corpus-level synthesis of maritime tasks, sensing modalities, benchmark comparability and communication constraints.
(Alsaaq 2026)	Civilian maritime object detection and tracking across optical, infrared, SAR and other sensing settings.	Closely related in sensing and detection coverage; the present review places more emphasis on security-specific versus adjacent-task evidence, protocol comparability and communications-aware deployment.
(Hanmante et al. 2026)	Multimodal maritime surveillance using AIS, SAR and optical fusion, including explainability.	Fusion-focused; the present review also compares detector families, reported benchmark results, task categories and deployment constraints beyond fusion alone.

2 Review methodology

The transparency and reproducibility of the literature synthesis in this survey require that the process used to identify, screen, classify and select the studies discussed in Sections 3–7 be reported. The review therefore followed a structured, PRISMA-informed identification and screening process with four principal stages: preliminary investigation, database and citation searching, publication screening, and data extraction and synthesis.

2.1 Preliminary investigation

The literature synthesis began with a preliminary investigation to clarify the scope of AI-based maritime security perception and to identify the terminology used across the computer-vision, remote-sensing, ocean-engineering and security literatures. Google Scholar was used as the primary discovery platform because of its broad coverage of journal articles, conference papers, book chapters and dataset papers regardless of publisher. Initial searches used the key phrases “maritime security”, “maritime domain awareness”, “ship detection”, “vessel detection”, “sea-surface object detection”, “maritime surveillance” and “multi-object tracking”.

After an initial set of relevant publications had been identified, further terminology was extracted from their titles, abstracts, keywords and reference lists, including “small-object detection”, “adverse-weather detection”, “dark vessel”, “illegal fishing”, “search and rescue”, “synthetic aperture radar”, “Automatic Identification

System", "unmanned aerial vehicle" and "unmanned surface vehicle". Google Scholar was additionally used for backward and forward citation chaining, and a targeted search of OpenAlex, Crossref and publisher websites was used to locate public maritime benchmark datasets, SAR-based dark-vessel and illegal-fishing studies, and the maritime-communication literature required for Section 6. The review period covers 2010 to October 2026, spanning the transition from hand-crafted maritime detectors to deep-learning, transformer and foundation-model approaches. Three International Maritime Organization (IMO) conventions were retained as regulatory background (International Maritime Organization 1974, 2003, 1972).

The search vocabulary was organized around four concept groups: (i) the maritime setting, (ii) the perception task, (iii) the learning approach and sensing modality, and (iv) the security application. Representative exploratory queries included (maritime OR sea OR ship OR vessel OR coastal) AND ("object detection" OR "ship detection" OR tracking); ("maritime surveillance" OR "maritime security") AND ("deep learning" OR CNN OR YOLO OR transformer); (SAR OR radar OR infrared OR satellite OR UAV OR USV) AND ("vessel detection" OR "small-object detection"); and ("dark vessel" OR "illegal fishing" OR "search and rescue" OR "port security") AND ("computer vision" OR "object detection"). These searches produced an exploratory corpus of 102 primary studies, to which the targeted search added nine.

2.2 Conducting research

A documented database search was then used as the principal structured search strategy, because the maritime-detection literature is large, growing rapidly and distributed across computer-vision, remote-sensing, ocean-engineering and security venues. The search was run on 20 October 2026 in IEEE Xplore and Semantic Scholar, with publication years restricted to 2010–2026, and repeated on 21 October 2026 in OpenAlex as an additional multidisciplinary index. In IEEE Xplore, the maritime and task terms were restricted to the title and the learning and application terms to the abstract:

```
(maritime OR ship(s) OR vessel(s) OR boat OR "sea surface" OR coastal OR
harbo(u)r) [title]
AND (detection OR detector OR tracking) [title]
AND ("deep learning" OR "neural network" OR convolutional OR YOLO OR
transformer) [abstract]
AND (security OR surveillance OR "search and rescue" OR "illegal fishing" OR
"dark vessel" OR "situational awareness" OR "domain awareness") [abstract]
```

Semantic Scholar applies these terms to title and abstract combined, so its results were afterwards restricted to records whose title contains both a maritime term and a detection or tracking term; OpenAlex applied the query to title and abstract. The exploratory and documented searches complement each other: the former covers classical detectors, system descriptions and studies without explicit security wording, whereas the latter covers deep-learning studies that describe a security or surveillance application. The following inclusion and exclusion criteria were applied during screening.

Inclusion: Records were included when they were peer-reviewed journal articles, conference papers or book chapters written in English and addressed at least one of the following: (i) a method for maritime object detection, classification, segmentation or tracking; (ii) a maritime dataset, benchmark or evaluation protocol; (iii) an application relevant to maritime safety or security, such as coastal surveillance, dark-vessel detection, illegal fishing, search and rescue, port monitoring or autonomous navigation; or (iv) an enabling technology directly relevant to maritime perception, such as sensor fusion, image enhancement, edge inference, or explainable or privacy-aware learning.

Exclusion: Records were excluded when they (i) did not concern maritime objects; (ii) did not involve detection or tracking; (iii) addressed only underwater or acoustic sensing, or vessel-trajectory prediction without a detection element; (iv) were not primary studies; (v) could not be verified as peer reviewed, such as preprints, theses and conference abstracts; (vi) were another version of an included record; or (vii) had been retracted. Duplicates across sources were counted once.

2.3 Examining publications

Publication examination followed a staged process (Fig. 1). The documented searches identified 2,004 records: 217 from IEEE Xplore, 1,138 from Semantic Scholar and 649 from OpenAlex. Records outside the title restriction, duplicates, and records already in the exploratory corpus or already retrieved were removed, leaving 678 records for title-and-abstract screening. Screening was carried out independently by two authors, with disagreements resolved by discussion or by a third author (Cohen's $\kappa = 0.96$ for IEEE Xplore and Semantic Scholar, and $\kappa = 0.82$ for OpenAlex), and 595 records were retained for eligibility assessment.

Eligibility was assessed against the criteria above from the complete bibliographic record and abstract of each retained record, following the authors' eligibility criteria and coding rules. After 51 records were excluded, 544 studies were added to the 111 exploratory and targeted studies, giving a review corpus of 655 primary studies. The narrative synthesis draws mainly on the 111 exploratory and targeted studies and cites documented-search studies where they add evidence.

From this corpus, the studies discussed individually in the synthesis were selected for their methodological or analytical contribution to the review questions. A study was cited when it (i) proposed or evaluated a detector, tracker or fusion method that illustrates a technique or trend; (ii) introduced or evaluated a maritime dataset or benchmark; (iii) reported results or evaluation protocols used in the comparability analysis; (iv) described a security application, deployment or system; or (v) addressed an enabling technology, limitation or trust issue discussed in the challenges. This yielded 156 primary studies that are cited and synthesized in the manuscript. Together with 14 related surveys used for positioning and 30 background sources on policy, regulation and maritime communications, they form the 200 works in the reference list.

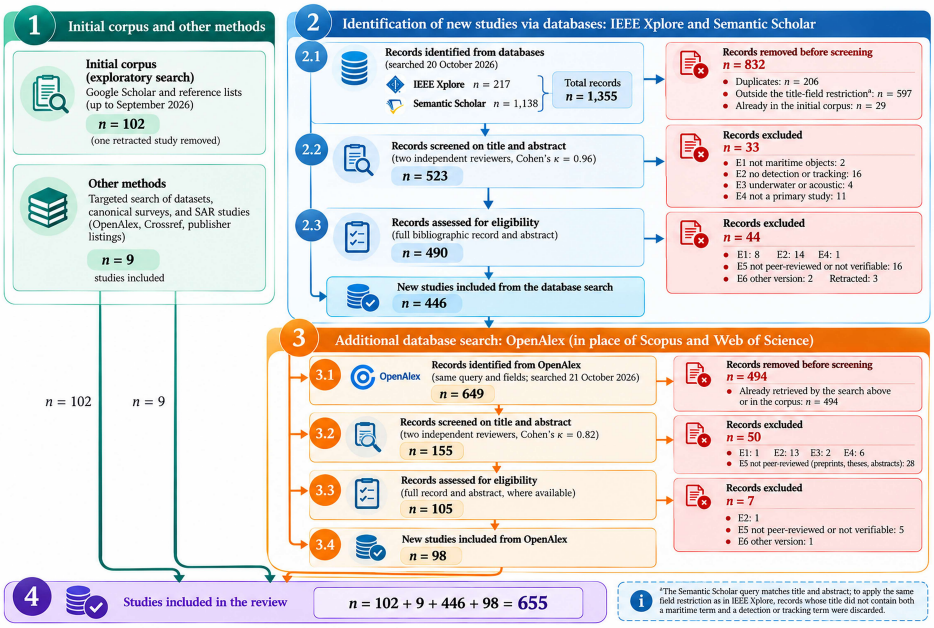


Fig. 1 PRISMA-style flow diagram for the documented and supplementary search stages. The 102-study exploratory corpus and nine targeted studies were identified separately; these were combined with 446 newly included IEEE Xplore/Semantic Scholar studies and 98 OpenAlex studies.

2.4 Extracting data

To address the research questions, a structured coding sheet documented the principal characteristics of each primary study. The recorded attributes were publication year, venue and publication type; contribution type (method, dataset or benchmark, comparative evaluation, or system, platform or sensor description); technology code; perception task; sensor modality; observation platform; and security applications. In addition, the complete abstracts of 557 studies were retrieved from IEEE Xplore, Semantic Scholar and Crossref, and the public benchmarks named, the metrics and intersection-over-union (IoU) thresholds used, the numerical results, and the reported speed, hardware and model size were extracted from them. Benchmark names were matched against 40 maritime and remote-sensing datasets and verified against the sentence in which each occurs. For the 19 studies with openly accessible full texts among those reporting results on common benchmarks, the evaluation protocol (dataset split, input size, hardware and repeated runs) was extracted from the full text (Supplementary Tables S1 and S2).

By contribution, 589 primary studies present methods, 33 datasets or benchmarks, 24 comparative evaluations and 24 system descriptions; because these labels are non-exclusive, their sum (670) exceeds 655 (Supplementary Table S5). Of the 655 studies, 383 are journal articles and 272 conference papers or book chapters; 166 do not specify the sensor in the title or abstract, and 59 name more than one sensor. By application, 504 studies name no specific task and are coded as general maritime surveillance, and

151 are coded as search and rescue (49), border and coastal surveillance including dark vessels (27), autonomous navigation (25), threat recognition (14), vessel identification (10), illegal fishing (9), port security (8), environmental monitoring (5) and detector security (4). These nine labels are not equally security-specific: border and coastal surveillance, illegal fishing, port security, threat recognition and detector security are direct security tasks, whereas search and rescue, autonomous navigation and environmental monitoring are adjacent operational domains, and vessel identification can support security without establishing a security outcome by itself.

The literature was organized according to the central distinction of this survey between object-level perception evidence and evidence about a security outcome. The synthesis therefore groups studies by application and by enabling technology, and contrasts them with respect to small and dense targets, sea-surface dynamics, adverse weather, occlusion, sensor diversity, domain shift, real-time processing, edge deployment and operational trustworthiness. Particular attention was given to how results are reported: whether a study names its dataset version and split, states the IoU threshold, reports repeated runs, and measures speed and energy on named hardware. This provides the basis for the comparability analysis in Section 4.3.

The extracted evidence was synthesized qualitatively through narrative comparison, methodological categorization, application mapping and comparative tables. No formal meta-analysis or universal quality score was applied, because the studies differ substantially in datasets, splits, metrics, hardware and objectives; reported scores are interpreted relative to the dataset and protocol on which they were obtained. Bibliographic records were checked against OpenAlex and Crossref, including for retractions, and the limitations of the review process are discussed in Section 8.

3 Technical scope and taxonomy

Terminology. The literature uses *maritime security* (protection against intentional unlawful acts) and *maritime safety* (unintentional hazards) with different operational requirements, yet often interchangeably (Eski 2019); *maritime domain awareness* is the common operational picture built by fusing sensors and databases (Mitra et al. 2017); *AIS* is a cooperative report that complements, but cannot replace, sensor-based detection (Eaton et al. 2018); *maritime object detection* covers detection, classification, segmentation and tracking as one perception chain (Sarajlić et al. 2026; Bousetouane and Morris 2016; Sharma et al. 2024; d'Afflisio et al. 2024; Niu et al. 2024); and a *dark vessel*—one not transmitting AIS—can be found only by non-cooperative sensing and is the key border-security target (Eaton et al. 2018). The key distinction is operational: a detector identifies an object or class, whereas a security decision additionally requires context, corroboration and often temporal evidence. Large ships may be easy to localize yet difficult to classify from a distant view; small boats and people can occupy very few pixels; wakes can indicate movement but are indirect evidence rather than an object class. AIS reports provide identity and position claims that should be cross-checked against sensor observations, especially when a vessel is absent, delayed or inconsistent in AIS (Eaton et al. 2018; Fischer and Bauer 2010; Bloisi et al. 2017).

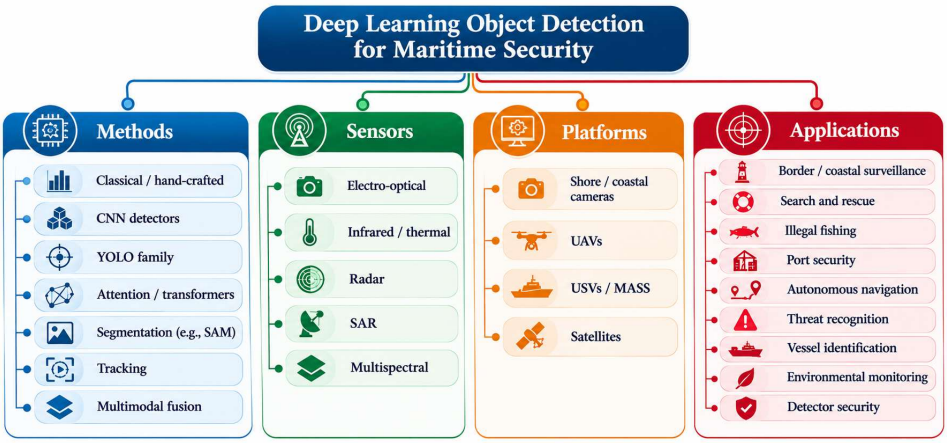


Fig. 2 Taxonomy of deep learning-based object detection for maritime security (methods, sensors, platforms, and applications).

The literature has progressed from hand-crafted features and classical vision—fast and interpretable but sensitive to waves, illumination and clutter (Prasad et al. 2017; Chan and Weed 2012; Shi et al. 2018)—to early neural and CNN-based ship detectors (de Nicolas et al. 2013; Cruz and Bernardino 2016) and the first public deep-learning benchmarks, such as SeaShips and MaSTr1325 (Shao et al. 2018; Bovcon et al. 2019), then to one-stage YOLO detectors, attention and transformer models (Akash et al. 2026; Messaoud et al. 2024; Feng et al. 2025), and most recently segmentation foundation models and vision-language perception (Meenakshi et al. 2025; Shi et al. 2026). These stages overlap rather than replace one another: classical methods remain useful for constrained sub-tasks, while learned detectors dominate complex benchmarks. Fig. 2 organizes the corpus along its method, sensor, platform and task dimensions and shows how many studies fall into each category.

Why maritime detection is distinctive. General-purpose detectors transfer imperfectly to maritime imagery for five recurring reasons. (1) *Small and dense targets*: ships viewed from satellites or high UAVs span a few pixels, and harbours are crowded, motivating multi-scale fusion, feature pyramids and small-object-aware transformers (Rekavandi et al. 2025; Chen et al. 2023a; Feng et al. 2025; Zhang et al. 2026). (2) *Adverse weather and illumination*: fog, haze, rain, low light and glare cut contrast, addressed by weather-aware detection, dehazing, low-light enhancement and joint denoising–detection (Chen et al. 2024; Tang et al. 2026; Zheng et al. 2026; Kaviyarasu et al. 2026). (3) *Dynamic background and non-stationary platforms* break the static-background assumption (Zhang et al. 2017; Chan and Weed 2012). (4) *Occlusion and large scale variation* (Banu and Sidram 2019; Chen et al. 2026b). (5) *Real-time and resource constraints* on edge platforms and satellites; FPGA-specific acceleration is a separate hardware-design question (Silva et al. 2024; Haijoub et al. 2025; Madabhushi et al. 2024).

Metrics and comparability. Detection is reported with precision, recall, F1, and average precision (AP)/mean AP at one or more intersection-over-union (IoU) thresholds, often broken down by object size; efficiency with frames per second, latency, parameters and floating-point operations; tracking with identity-preserving multi-object measures (Niu et al. 2024; Guo et al. 2024); and operational tasks with measures such as false alarms per hour and detection range, which the reviewed studies rarely report. Efficiency figures are meaningful only for named hardware (Silva et al. 2024; Haijoub et al. 2025). Because datasets, thresholds and hardware differ, raw figures are not comparable across papers (Su et al. 2023; Alfaro et al. 2025); Section 4.3 quantifies this problem and Section 7 proposes a reporting checklist to address it. Representative datasets are discussed with the applications in Section 5.

4 Critical synthesis of detection methods and reported evidence

Every application rests on a processing chain: sensing, pre-processing, detection or segmentation, optional tracking, fusion with cooperative data, and decision support. Errors propagate along that chain. Low-contrast frames can lower recall before tracking begins; inconsistent timestamps or spatial registration can degrade sensor fusion; and accurate detection alone does not establish intent or unlawful activity. This section therefore compares the model families and evaluation evidence while distinguishing local object recognition from the higher-level security interpretation that may follow.

4.1 Detector families

Four broad detector paradigms recur in the corpus. *CNN detectors* were the first deep models used at sea and remain the mature baseline for benchmarking, often with transfer learning for scarce labels and multi-scale designs to preserve small objects (Alghazo et al. 2021; Farahnakian et al. 2021; Soloviev et al. 2020). *One-stage YOLO detectors* are by far the most common choice—mentioned in 199 of 655 studies—valued for the speed–accuracy trade-off and adapted across all sensors by attention modules, richer pyramids, tiny-box losses and pruning; the adaptations are strikingly similar across papers (Akash et al. 2026; Ma et al. 2026; Yang et al. 2025), and a data-centric comparison suggests that training data matter as much as architecture (Kapoor and Sinha 2024). *Attention and transformer models* (123 studies), as attention modules on CNNs or DETR-style detectors, target small and dense objects in SAR and infrared at higher computational cost (Messaoud et al. 2024; Feng et al. 2025; Zhang et al. 2026; Venkata et al. 2025). *Segmentation and foundation models* add per-pixel outlines and open-vocabulary flexibility (Sharma et al. 2024; Shi et al. 2026); a CPU implementation of the Segment Anything Model shows that foundation models can run without a GPU, but their accuracy on tiny, low-contrast objects and their cost remain open (Meenakshi et al. 2025). Table 2 compares the families qualitatively; as Section 4.3 shows, the reported numbers do not support a ranking.

Table 2 Qualitative comparison of detector families used in maritime security. Strengths and limitations summarise the cited examples and the results in Supplementary Table S1; they are qualitative judgments, not measured scores, and do not rank the families.

Family	Strengths	Limitations	Typical maritime use	Examples
Classical / hand-crafted	Light, interpretable, no training data	Sensitive to waves and glare, manual tuning	Early systems, simple static scenes	(Shi et al. 2018; Chan and Weed 2012; Frost and Tapamo 2013; Banu and Sidram 2022)
Two-stage CNN (Faster R-CNN style)	Mature, good localization	Slower, loses fine detail through down-sampling	Radar and satellite images, offline analysis	(Sharma et al. 2024; Alghazo et al. 2021; Gonzales-Martínez et al. 2021)
Single-stage (YOLO family)	Fast, easy to deploy, large ecosystem	Small and dense targets need tuning	Real-time UAV, USV, shore and edge systems	(Nandal et al. 2025; Yang et al. 2025; Bi et al. 2026; Ruan et al. 2026)
Attention and transformer models	Global context; gains reported on small and dense targets	Higher computation, data hungry	SAR, infrared and dense harbor scenes	(Messaoud et al. 2024, 2023; Feng et al. 2025; Zhang et al. 2026)
Segmentation and foundation models	Pixel-level outlines, flexible categories, labeling aid	Heavy computation, weak on tiny targets	Outlining, trajectory analysis, annotation	(d'Afflisio et al. 2024; Cane and Ferryman 2018; Meenakshi et al. 2025)
Vision-language and multimodal models	Open vocabulary, sensor complementarity	Alignment and evaluation immature	Emerging security systems	(Shi et al. 2026; Priyadharshini and Vadivazhagan 2024; Liu et al. 2025)

4.1.1 CNN-based and two-stage detectors

CNN-based detectors remain a useful reference point because their performance can be understood against established object-detection pipelines and fine-tuned from pre-trained weights when maritime annotations are scarce. Region-based or two-stage approaches generate candidate regions before classification and refinement, which can be advantageous when localization matters more than maximum throughput. Their main constraints at sea are computational cost and the loss of fine spatial detail through repeated down-sampling: a vessel that occupies only a small number of pixels may disappear in deeper feature maps. Multi-scale feature pyramids, higher-resolution inputs and small-object-aware losses address parts of this problem, but higher input resolution also increases memory and latency. Reported gains should therefore be interpreted together with input dimensions, parameter count and evaluation protocol, not as an architecture-only effect (Soloviev et al. 2020; Alghazo et al. 2021).

4.1.2 Single-stage detectors and the YOLO family

YOLO is the most frequently identified family in the coded corpus (199 studies), which is consistent with its practical appeal for real-time cameras and embedded platforms. A single forward pass produces bounding boxes and classes, and a

large deployment ecosystem supports quantization, pruning and hardware-specific implementations. Maritime adaptations commonly add attention, multi-scale feature aggregation, smaller-object losses or feature-fusion blocks. These changes respond to genuine maritime difficulties, but the similarity of proposed modifications makes attribution difficult: if two papers use different image sources, split strategies, augmentation, resolution and training schedules, a small accuracy gain cannot be attributed confidently to the architectural component alone. A data-centric comparison of YOLO and transformer detectors is an important corrective to architecture-driven narratives: improvements in annotation quality, coverage of sea conditions and split design can matter as much as a new module (Kapoor and Sinha 2024). For an operational choice, latency on the intended platform and recall at the target range should be assessed alongside AP, not assumed from the model family name (Akash et al. 2026; Yang et al. 2025).

4.1.3 Attention, transformers and global context

Attention modules inserted into CNNs and transformer-based detectors should not be treated as a single interchangeable class. Local or channel attention can reweight salient feature maps while preserving much of the existing detector pipeline; transformer-style self-attention can model longer-range relationships among targets and background regions, which may help in dense harbors or scenes with small, partially occluded vessels. The cost is not only inference complexity: training may require more labelled data, and the benefit can depend on image resolution, object density and the quality of the backbone features. In SAR imagery, the speckle and clutter structure differs from optical images (Venkata et al. 2025; Ma et al. 2024), so a reported benefit in one modality should not be generalized to another without cross-dataset evidence. The 123 studies coded as attention or transformer approaches include heterogeneous designs and are not a count of pure transformer detectors. Future comparisons should separate attention-as-a-module from end-to-end transformer detectors and use matched baselines, repeated runs and size-stratified metrics (Messaoud et al. 2024; Ma et al. 2024).

4.1.4 Segmentation, foundation models and multimodal perception

Segmentation provides object masks rather than only bounding boxes and is useful when the boundary, extent or surrounding wake matters. However, mask-level output is more expensive to annotate and evaluate, and segmentation quality can be fragile for tiny targets with weak boundaries. Foundation and vision-language models introduce flexible prompting or category descriptions that may reduce the need for a fixed class inventory, but the evidence in the reviewed maritime literature is still emerging. A model that can segment common objects in general imagery is not automatically reliable for small ships, lifeboats, buoys or low-contrast targets at sea; latency and prompt sensitivity also need maritime evaluation. The recent vision-language line is promising for open-set recognition and operator assistance, but it should be treated as complementary evidence rather than as a replacement for specialized detectors until systematic comparisons are available (Sharma et al. 2024; Meenakshi et al. 2025; Shi et al. 2026).

4.2 Choosing a detector under maritime operating conditions

The evidence does not support a single preferred family across all maritime settings. For coastal cameras and UAV video, throughput and temporal continuity make single-stage detectors attractive (Yang et al. 2025; Geng et al. 2025b), but camera motion, glare, occlusion and the need to follow small objects mean that frame-level AP alone is insufficient (Zhang et al. 2017; Zhu et al. 2025; Banu and Sidram 2019). Evaluation should report recall at operationally relevant object sizes, false alarms over time, and the effects of tracking or temporal confirmation. For satellite optical imagery, fine spatial detail can support visual classification but depends on illumination and cloud conditions (Kanjir et al. 2018); for SAR, day/night and cloud-penetrating capability is offset by speckle, unfamiliar target appearance and modality-specific clutter (Venkata et al. 2025; Paolo et al. 2022). A method tuned on optical data should not be described as robust to SAR or radar unless that transfer has been evaluated.

Benchmark evidence is concentrated on SAR ship detection (Section 4.3). A high score on a familiar SAR split establishes performance only on that benchmark and protocol; it cannot establish broad maritime-security reliability. Likewise, a detector trained for search and rescue may be valuable for finding small people or craft (Varga et al. 2022; Zhao et al. 2025a) but does not automatically meet a border-surveillance requirement, where false alerts, identity corroboration and evidence persistence may carry greater operational weight (Eaton et al. 2018). Model choice should therefore begin with the task, sensor, range, platform and error costs, then compare models under a common protocol.

Complementary techniques. Tracking associates detections over time for USVs and UAVs, with identity switches, glare and camera motion as the main failure modes, and long-term maritime tracking data still lacking (Zhu et al. 2025; Zhang et al. 2022; Niu et al. 2024; Alsaaq 2026). *Multi-sensor fusion* (coded as a technology in 31 studies) combines the complementary strengths of optical, infrared, radar, SAR, multispectral and AIS inputs at data, feature or decision level, but must handle misalignment, differing resolutions and missing modalities (Fischer and Bauer 2010; Markchom et al. 2026; Priyadharshini and Vadivazhagan 2024; Liu et al. 2025). Modality selection is conditional rather than a universal ranking (Supplementary Table S4): optical imagery preserves visual detail in suitable illumination (Lyu et al. 2023; Kanjir et al. 2018); thermal infrared helps in low light but lacks colour and texture (Teutsch and Krüger 2010; Ma et al. 2026); SAR supports day/night, cloud-penetrating observation but creates a different clutter and speckle problem (Venkata et al. 2025; Paolo et al. 2022); and marine radar supplies range and motion information but gives weak returns from small wooden or fibreglass craft and no visual class detail (Gonzales-Martínez et al. 2021; Külcü 2026). These are physical and operational trade-offs, not a benchmark-derived score. Fusion is useful only when registration, timing, missing sensors and inconsistent AIS are handled explicitly. *Label-efficient learning* (transfer learning, synthetic/game-engine data, rebalancing, foundation-model masks, open-set detection) addresses scarce, imbalanced and domain-shifted data, yet no reviewed study explicitly adapts across changing sea states or sensors (Farahnakian et al. 2021; Vaidya et al. 2026; Zhao et al. 2025a; Yoon et al. 2024). *Image enhancement* before or

jointly with detection is common, though an image that looks better is not necessarily better for detection, and benefits are rarely compared across studies (Chen et al. 2024; Tang et al. 2026; Zheng et al. 2026; Kaviyarasu et al. 2026).

4.3 What the reported results do and do not show

Abstracts were available for 557 of the 655 studies and were parsed for benchmarks, metrics, results, speed and hardware; for the most frequently used benchmarks, the reported values are listed in Supplementary Table S1, and the evaluation protocols of the 19 openly accessible full texts in Supplementary Table S2. Five findings follow (Fig. 3).

(1) Evaluation concentrates on a few SAR benchmarks. Of 557 abstracts, 182 name a public maritime or remote-sensing benchmark and 92 of these name a SAR one; SSDD (66) and HRSID (55) dominate, ahead of SeaDronesSee (25), SeaShips (18), HRSC2016 (15) and LS-SSDD (13). Only four recent studies use the infrared ship detection dataset ISDD (Tang and Ji 2025; Sun and Lian 2025; Gao et al. 2025; Niu et al. 2026a), and no maritime-radar benchmark is shared across studies: the two radar results in Supplementary Table S1 come from 350 simulated frames, with an external check on one real radar dataset (Külcü 2026), and from 4,237 images collected by the authors (Cen et al. 2025).

(2) Metrics are often under-specified. Of 299 abstracts reporting a detection metric, 121 give AP with no IoU threshold, 113 give AP_{50} and only 39 the stricter $AP_{50:95}$. The difference is large: on SSDD the same studies report ≈98% at IoU 0.5 but ≈72–73% averaged over 0.5–0.95 (Hoang-Ngoc et al. 2025; Niu et al. 2026b), and a threshold-free AP (Ma et al. 2024; Zhao et al. 2023) or a value on a modified split such as SSDD+ (Cheng et al. 2026) is not comparable to either.

(3) SSDD is saturated. Ten studies report AP_{50} of 96.8–99.0% on SSDD (Fan et al. 2024; Fang et al. 2025; Bu et al. 2026; Huang et al. 2026b; Chen et al. 2025; Niu et al. 2026b; Wang et al. 2026; Zhang and Huang 2025; Hoang-Ngoc et al. 2025; Huang and Shang 2024), a range of about two points. Because these studies differ in split, input size and training schedule (Supplementary Table S2) and none reports run-to-run variation, differences of this size cannot be attributed to the architectures; they do not show what causes the differences either. Where $AP_{50:95}$ is reported, the values are even closer (71.9–73.2%) (Chen et al. 2025; Hoang-Ngoc et al. 2025; Wei and Cai 2026; Bu et al. 2026; Niu et al. 2026b; John and Guan 2025). HRSID separates methods more clearly (AP_{50} from 85.6% to 95.9%) (Zhang et al. 2024; He and Huang 2024), and LS-SSDD, with small ships in large scenes, yields values of about 77% (Garaiman and Radoi 2025; Guan et al. 2025). New SAR detectors should therefore be judged on $AP_{50:95}$, HRSID or LS-SSDD and small-object metrics rather than on AP_{50} for SSDD.

(4) The same benchmark name is not the same test set. On SeaShips, AP_{50} ranges from 79.7% on the SeaShips 7000 release (Cai and Zhou 2026) to 98.9% on an extended version (Geng et al. 2025a) and 99.0% on an unspecified split (Liang and Han 2025); on SeaDronesSee, from 56.9% on version 2 (Yang et al. 2025) to 78.1–80.5% in studies that do not state the version (Wang et al. 2025a; Lu et al. 2026; Feng et al. 2025). These spreads exceed the typical reported gain over a baseline, which is often below two points (e.g. +0.7 AP_{50} and +0.3 $AP_{50:95}$ on HRSID (He and Huang 2024), +0.5 and

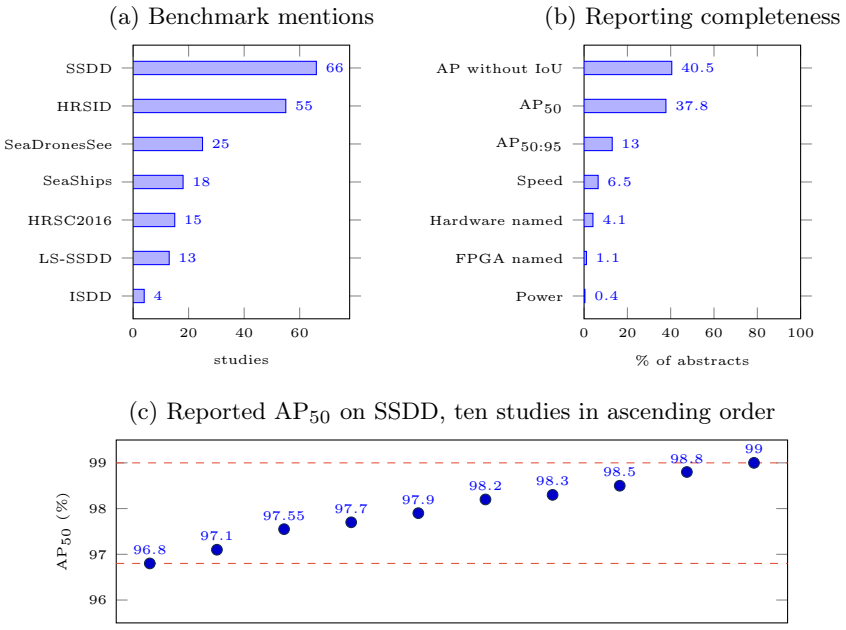


Fig. 3 Evaluation evidence in the corpus. (a) Public benchmarks named in the 557 available abstracts; no maritime-radar benchmark is used by more than one study, so none is shown. (b) Share of abstracts reporting each item: metric and IoU threshold among the 299 abstracts that report a detection metric; speed, hardware and power among all 557. (c) The ten AP_{50} values reported for SSDD, spanning 2.2 points; the six $AP_{50:95}$ values reported for SSDD lie between 71.9% and 73.2% (sources in Section 4.3). Counts are from the abstract-level extraction described in Section 2; the SSDD values were checked against the studies' abstracts.

+0.2 on SeaDronesSee (Yang et al. 2025)). Among the 19 openly accessible protocols, only one reports multiple runs with standard deviations (Supplementary Table S2), so gains of this size cannot in general be separated from training variance. Model size is reported more consistently: several detectors exceed 97% AP_{50} on SSDD with fewer than 3 million parameters (Huang and Shang 2024; Chen et al. 2025; Wang et al. 2026).

(5) Deployment evidence is thin. Only 36 of the 557 abstracts report speed, 23 name the hardware (six an FPGA), and two report power consumption, both for FPGA implementations (e.g. Lv et al. 2026). These few observations use different models, inputs and measurement conditions, so we do not use them to rank embedded platforms or to infer an accuracy penalty for edge deployment. This gap is itself consequential: workstation accuracy is not a reliable proxy for performance on a buoy, USV or UAV. Deployment studies should report model, quantisation, input resolution, hardware, batch size, end-to-end latency, power measurement method and accuracy under the same test conditions; these requirements motivate the reporting checklist in Section 7.

4.4 Technology–application coverage

Coding all 655 studies by technology and application yields the matrix in Table 3 (counts). The YOLO family dominates (199 studies) ahead of attention/transformers (123), with tracking (41), fusion/AIS (31), enhancement (17), federated/explainable/security learning (11) and reasoning (3) much rarer. Coverage is uneven: lightweight and edge approaches appear in 83 studies but only 14 target a specific application; search and rescue has 49 studies but only two on tracking, one on fusion and three on enhancement; illegal fishing and port security have no tracking study and one edge study each; threat recognition holds all three reasoning studies but none on attention, enhancement or edge. Empty and sparse cells indicate research gaps, though abstract-only coding may also reflect how studies describe themselves. The coded records indicate that most studies address one or two stages of the perception chain; end-to-end evaluation from sensing through security-relevant decision support remains uncommon.

Table 3 The studies are allocated to the rows of technology codes and the columns of application codes; the columns correspond to the nine task-specific labels defined in the text. Each cell gives the number of studies with that technology–application code; "–" means that no study has this combination. Technology codes are non-exclusive, so a study can appear in several rows; each study has one application column.

Technology	General surv.	Border / coastal	Search and rescue	Illegal fishing	Port security	Auton. navigation	Threat recognition	Vessel identification	Environ. monitoring	Detector security
Classical (hand-crafted)	10	1	0	–	–	1	–	1	–	–
CNN / deep detectors	238	16	13	8	4	9	8	6	2	–
YOLO family	157	2	25	–	3	6	2	2	2	–
Attention, transformers, vision-language models	101	5	12	–	2	1	–	2	–	–
Segmentation / foundation models	17	0	0	–	–	4	1	–	–	1
Tracking	24	3	2	–	–	9	1	1	1	–
Multimodal fusion and AIS	18	4	1	3	–	2	2	1	–	–
Enhancement / pre-processing	13	0	3	–	–	1	–	–	–	–
Edge / embedded / lightweight	69	0	8	1	1	2	–	–	1	1
Federated learning, explainable AI, detector security	4	1	1	–	–	1	–	–	–	4
Reasoning and anomaly detection	0	0	0	–	–	0	3	–	–	–
Systems and sensing hardware	5	1	0	–	1	0	–	–	–	–
Studies per application	504	27	49	9	8	25	14	10	5	4

5 Maritime-security applications and datasets

Detection acquires value only through the task it enables, but the corpus indicates an important distinction between where maritime object detectors are studied and where they are demonstrated to improve security outcomes. The nine task labels are not equally security-specific: border and coastal surveillance, illegal-fishing detection, port security, threat recognition and detector security concern security directly; vessel identification can support security but does not alone establish illegality; search and rescue, autonomous navigation and environmental monitoring are adjacent operational areas. The latter provide valuable transfer evidence, but not direct proof of security effectiveness, particularly because many studies report object-level metrics without evaluating the downstream decision.

5.1 Direct security tasks

Border and coastal surveillance. This is the central application for non-cooperative detection, especially when vessels do not transmit AIS or transmit inconsistent identity and position information. SAR is useful for wide-area observation and for acquisition through cloud or darkness, whereas coastal electro-optical cameras and UAVs offer different advantages in revisit rate, resolution and local tracking. Hybrid CNN–transformer and attention methods appear in the reviewed corpus (Venkata et al. 2025; Ma et al. 2024); public datasets such as SSDD, HRSID and LS-SSDD support SAR vessel detection (Zhang et al. 2021b; Wei et al. 2020; Zhang et al. 2020), and xView3-SAR was built specifically for dark fishing activity (Paolo et al. 2022). Yet a ship-shaped target in SAR is not automatically a verified dark vessel: that inference depends on temporal and spatial alignment with AIS, coverage of the AIS feed, vessel identity matching and the treatment of missing or delayed messages (Eaton et al. 2018; Paolo et al. 2024). The literature is strongest on image-level detection and less consistent in demonstrating end-to-end, operational detection of non-cooperative vessels. Satellite coverage and revisit also differ from shore, ship and UAV systems, so evidence cannot be transferred between these settings without qualification.

Illegal fishing and port security. Illegal-fishing detection overlaps with dark-vessel surveillance and satellite observation, but only nine studies in the coded corpus are assigned this label. Detecting a vessel does not establish that illegal fishing is occurring. An operational inference requires location and time context, evidence of fishing activity and, where available, cross-checking against cooperative reports; the nine studies assigned this task label are too few to establish general operational performance (Lalasa et al. 2024; Paolo et al. 2022, 2024). Port-security work similarly needs more than vessel localization: the relevant outcomes may include approach to a restricted zone, anomalous activity or an alert that is correctly reviewed by an operator. Dense traffic, occlusions and local camera geometry make harbor deployment different from open-sea satellite detection. The corpus contains eight port-security records, and the limited number makes it difficult to infer general performance or operational readiness from the reported methods (Hassan et al. 2026; Zhang et al. 2026; Tu-Salaam and Alamaniotis 2026).

Threat recognition and detector security. Threat recognition adds a reasoning layer to perception. Detection supplies observations; tracking links them over time; behavior analysis or anomaly detection then considers whether a sequence is unusual in context. The activity–behavior–intent distinction is useful because an unusual trajectory can be a cue for further review but is not, by itself, evidence of hostile intent (Gamage et al. 2023; Shahir et al. 2014; Chen et al. 2026a). Fourteen records are coded as threat-recognition studies, whereas only four are coded as detector-security studies. The latter category concerns the reliability of the AI system itself, including adversarial inputs and cyber interfaces, rather than the detection of threats by the system. The imbalance suggests that detector robustness and secure update/fusion pathways deserve explicit testing, while the small sample means that broad conclusions about maritime-specific adversarial resistance remain premature (Ashraf et al. 2023; AlZaabi et al. 2020).

5.2 Adjacent operational applications and transfer limits

Search and rescue. Search and rescue is the largest task-specific group (49 records), centred on detecting small people, rafts and distressed craft, often from UAV imagery. SeaDronesSee and related collections support tiny-target research, while low contrast, class imbalance, flight height, weather and night operation remain practical constraints (Varga et al. 2022; Dong et al. 2025; Geng et al. 2025b). This work can transfer useful methods to border monitoring because both tasks require detecting small objects against dynamic water backgrounds. Nevertheless, the operational error costs differ: a rescue detector may prioritize missed people and recall, whereas a security system may also need strict false-alarm control and corroboration before escalation. A detector trained for one objective should not be assumed suitable for the other without revalidation.

Autonomous navigation. The 25 studies coded for autonomous navigation focus on onboard perception of vessels, obstacles and buoys and are often tied to collision avoidance. MaSTr1325, MODS and LaRS support parts of this problem, but navigation scenes differ from long-range surveillance in viewing geometry, object scale, motion, latency tolerance and the consequences of a false negative (Mehla et al. 2023; Bovcon et al. 2019; Žust et al. 2023). Tracking and temporal consistency are especially relevant here, though the presence of tracking in a navigation study does not itself resolve identity, unlawful behavior or security intent.

Environmental and vessel-identification tasks. Environmental-monitoring studies reuse maritime perception for tasks such as identifying vessels or observing wakes and related scene cues. Ten records are coded for vessel identification and five for environmental monitoring. These are useful when they supply transferable recognition methods or sensor evidence, but their metrics need to be linked explicitly to a security objective before they are treated as security findings. The same caution applies to general surveillance: it is the largest group, but it often reports detection performance without identifying a particular security scenario.

From detection to a security outcome. The reviewed work often stops at an intermediate perception measure. A detector can locate a ship-shaped object without establishing the vessel's identity, whether it is transmitting valid AIS, whether it is inside a restricted area, or whether its behaviour is unlawful. Evidence for a security

system therefore spans at least three layers: (i) perception quality under the stated sensor and conditions; (ii) fusion and temporal association, including matching to cooperative reports without treating those reports as ground truth; and (iii) the operational decision, such as a correctly reviewed alert or confirmed non-cooperative vessel. Most reported results in this corpus concern the first layer. We retain those studies because they develop useful enabling capabilities, but we do not treat their AP or recall values as direct evidence that the full security task has been solved.

5.3 Dataset coverage and what it can establish

Public benchmarks exist for shore-based optical detection (Prasad et al. 2017; Shao et al. 2018; Moosbauer et al. 2019), USV obstacle detection and segmentation (Bovcon et al. 2019, 2022; Žust et al. 2023), UAV-based search and rescue (Varga et al. 2022; Dong et al. 2025), and SAR ship and dark-vessel detection (Zhang et al. 2021b; Wei et al. 2020; Zhang et al. 2020; Paolo et al. 2022); their use is concentrated on SAR (Fig. 3a). Benchmark mentions in abstracts measure use, not benchmark quality or protocol reproducibility. Infrared evidence rests on a single dataset used by four studies, and no maritime-radar benchmark is shared across studies, so comparative evaluation remains limited for the two modalities that operate best at night.

Dataset name alone is insufficient to establish comparability. Versions, geographic regions, target-size distributions, train/test splits, data leakage controls, image resolution, augmentation, and metrics at different intersection-over-union (IoU) thresholds can produce materially different scores. The reported SeaShips and SeaDronesSee ranges illustrate that changed versions and splits can shift results more than the gains claimed for some new architectures. We therefore treat the benchmark values as a record of what individual studies report, not as a harmonized leaderboard. The reported results and protocols are listed in Supplementary Tables S1–S2 and representative datasets in Supplementary Table S6.

Evidence from the documented-search additions. Of the 544 studies added by the documented searches, 106 carry a task label and 17 introduce datasets or benchmarks. Search and rescue is the largest added group (44 studies), dominated by YOLO and real-time detection transformer (RT-DETR) variants for tiny objects on UAV or embedded hardware (Zhao et al. 2025b; Chen and Chen 2025; Lu et al. 2026; Li et al. 2025); fewer address bad weather, low visibility, sun glint or night operation (Wang et al. 2025b; Huang et al. 2026a; Fei et al. 2024; Rajesh et al. 2026). Several of the 19 added border and dark-vessel studies fuse SAR or optical imagery with AIS (Milios et al. 2019; Kim et al. 2025; Monani et al. 2025; Mazzeo et al. 2025), but few report how disagreements between AIS and sensor data are handled. Only one added study presents operational results; its authors report a six-month deployment on patrol boats (Ebenezar et al. 2026), and one addresses the security of the detector itself by monitoring for adversarial inputs on an embedded board (Bouhlal and Brust 2026). The 17 added datasets broaden optical, SAR and tracking coverage and include one registered visible–infrared dataset (Xu et al. 2026), but none is a marine-radar benchmark. Overall, the evidence supports progress in object-level perception but is insufficient for claims about verified security outcomes, long-duration reliability or field readiness.

6 Deployment: communication, efficiency and trustworthiness

Where a maritime detector can run, what it can transmit, and how it can be updated are decided by the communication link as much as by the model—yet no primary study in the corpus evaluates detection accuracy as a function of available data rate. This section quantifies the link budget from the maritime-communication literature and reviews communication-aware inference (Alqurashi et al. 2023; Wei et al. 2021).

Links in numbers. Capacity falls with distance from shore (Table 4). AIS uses two 9.6 kbit/s VHF channels for standardized short maritime reports, not as a generic transport for arbitrary detector payloads (International Telecommunication Union 2014); the VHF Data Exchange System (VDES) adds ≈307 kbit/s terrestrial and ≈2–101 kbit/s satellite channels (Lázaro et al. 2019; International Telecommunication Union 2022; Luglio et al. 2018); coastal cellular gives tens of Mbit/s; and offshore only satellites remain—geostationary links add ≈480 ms round-trip latency, while measured low-Earth-orbit (LEO) service reaches 50–100 Mbit/s down and 4–12 Mbit/s up at 40–50 ms for fixed users (Mohan et al. 2024).

Table 4: Indicative maritime communication links and their relevance to AI-based detection. Rates and latencies are source-reported; low-Earth-orbit (LEO) values refer to measurements for fixed terrestrial users and may be lower at sea. For the illustrative comparison, a compressed 1080p30 stream is assumed to use about 5 Mbit/s; actual H.265 rate depends on scene, encoder and quality. A minimal payload of 50 detections/s at 12 bytes per detection equals 4.8 kbit/s before timestamps, identifiers, metadata, framing, error correction or retransmission. AIS is a standardized cooperative-reporting system, not a generic detector-payload channel. LTE: Long Term Evolution.

Link class	Coverage and indicative capacity / latency	Relevance to detection and sources
AIS (VHF)	Coastal line of sight; satellite reception may extend reception offshore. Two shared 9.6 kbit/s channels; standardized reports arrive every few seconds to minutes.	Provides cooperative identity and position claims and can help flag missing or inconsistent reports. Fuse with sensor detections; do not use as a generic image or detector-payload transport. (International Telecommunication Union 2014; Eaton et al. 2018; Fischer and Bauer 2010)
VDES terrestrial (VDE-TER)	Coastal. Up to 307.2 kbit/s nominal per 100 kHz channel; low latency.	May carry alerts and compact detection summaries from vessels or buoys; not a practical video channel at the stated nominal rate. (Lázaro et al. 2019; Luglio et al. 2018)
VDES satellite (VDE-SAT)	Open sea, subject to service and access. Rates vary; roughly 2–29 kbit/s downlink and 2–101 kbit/s uplink.	Compact alerts or detection summaries may fit depending on uplink rate and overhead; a 4.8 kbit/s bare payload will not fit every service condition. (Luglio et al. 2018)
Cellular (LTE/5G)	Near shore and in ports; tens of Mbit/s or more with low latency under adequate coverage.	Supports video, edge-server offloading and detections in coastal areas; actual capacity depends on coverage and congestion. (Alqurashi et al. 2023; Wei et al. 2021)
Satellite, LEO	Broad coverage, service dependent. Fixed-user measurements: about 50–100 Mbit/s down and 4–12 Mbit/s up, with 40–50 ms median round-trip latency.	Can support detections, image chips and selected compressed streams; cited rates are not assumed to represent all maritime conditions. (Mohan et al. 2024)

Link class	Coverage and indicative capacity / latency	Relevance to detection and sources
Satellite, GEO	Broad coverage except polar regions; service dependent. Capacity varies; propagation round-trip latency is at least about 480 ms by geometry.	Suitable for alerts, detections, image chips and delay-tolerant traffic; interactive control and streaming are constrained by latency and service rate. (Alqurashi et al. 2023; Wei et al. 2021)
Hybrid / integrated networks	Cooperating satellites, UAVs, USVs, ships and shore stations; capacity varies with links and routing.	Potentially supports sharing of detections, tracks and model updates, but requires attention to interoperability, authentication and intermittent links. (Wei et al. 2021; Pang et al. 2020; Zhao et al. 2024)

What to transmit. Sensor data far exceed these links. Uncompressed 1920 × 1080 video at 24 bit and 30 fps is ≈1.49 Gbit/s. Streaming guidance recommends about 4.5–5.8 Mbit/s for such a stream in H.265 (Apple Inc. 2026; Sullivan et al. 2012); taking 5 Mbit/s as an illustrative value, compression reduces the rate roughly 300-fold. A minimal raw payload of 50 detections per second at 12 bytes per detection is ≈4.8 kbit/s before timestamps, identifiers, metadata and protocol overhead. Even compressed video thus needs roughly three orders of magnitude more bandwidth than a detection list, fitting only the coastal cellular and higher-rate LEO links, whereas compact alerts or lower-rate detection summaries may fit VDES only when the available uplink capacity exceeds the payload plus protocol overhead; richer detection streams and imagery require more capable links. Running the detector on the platform—as in the corpus edge studies and onboard satellite and SAR processing (Silva et al. 2024; Haijoub et al. 2025; Prete et al. 2024; Madabhushi et al. 2024)—converts the bandwidth bottleneck into a compute one, but keeps the raw evidence that an operator (to verify) and a developer (to retrain) both need on the platform; sending detections continuously, image chips on demand and full frames only when bandwidth allows is a plausible policy, although no corpus study evaluates it. The available evidence does not establish a matched-condition comparison of workstation and edge accuracy, so on-platform performance should not be assumed equal to workstation performance.

Communication-aware inference and federated cost. Three strands from the communication literature address inference over constrained links—split computing and edge offloading, semantic/task-oriented communication, and federated learning over satellite and integrated networks (Supplementary Table S3) (Madabhushi et al. 2024; Matsubara and Levorato 2021; Zhang et al. 2021a; Dou et al. 2026; Bourtsoulatze et al. 2019; Shao et al. 2022; Chen et al. 2023b; Razmi et al. 2022; Chen et al. 2022)—but none has been tested on a maritime detection benchmark, and no corpus detector uses them, making their combination a clear opportunity for cross-disciplinary work. Federated learning is attractive for privacy, but each round must transmit model parameters or an update. As an illustrative upper-bound scenario in which a full model-sized payload is sent in 32-bit precision, YOLOv11n (about 2.6 million parameters), Marine-YOLO (4.2 million) and RT-DETR-X (67.3 million) correspond to approximately 10.4, 16.8 and 269 MB, respectively (Ruan et al. 2026). At a 4–12 Mbit/s LEO uplink, these payloads require roughly 7–21 s to 3–9 min; at 100 kbit/s, the largest would require about 6 h. These are idealized transfer times, excluding protocol overhead, retransmissions, quantization and compression; actual federated updates may be smaller than a full model. Maritime federated-learning studies report model

performance and privacy but seldom quantify per-round traffic or intermittent-link behaviour (Nandal et al. 2025; Huang et al. 2025), leaving update compression, asynchronous rounds and link-aware participant selection as open directions (Razmi et al. 2022; Chen et al. 2022).

Communication-layer and detector security. Connectivity enlarges the attack surface: AIS can be disabled or spoofed, IoT maritime systems face broad cyber threats, integrated space–air–ground–sea networks are exposed at every interface, and adversarial perturbations transfer between detectors (Eaton et al. 2018; Ashraf et al. 2023; Guo et al. 2022; Manasut et al. 2025). Detections, alerts and model updates sent over the air should be authenticated, and a fusion system should treat AIS as a claim to verify rather than ground truth. The implication for research is to evaluate detectors jointly with their data-rate requirement and, where possible, under link traces (reflected in RQ4 and the roadmap of Section 7).

7 Open challenges and prioritised research directions

The preceding sections support five lessons: the field is young and dominated by general surveillance; evidence concentrates on satellite SAR and optical ship detection with YOLO and attention networks; benchmarks exist for optical, USV, UAV and SAR data, but infrared evidence is thin and no radar benchmark is shared across studies; reported gains over baselines are usually smaller than the differences between dataset versions and splits, and are rarely accompanied by repeated runs; and robustness, adversarial security and privacy are evaluated separately from detection accuracy rather than as part of it. Fourteen challenges follow in three groups; Table 5 lists, for each, the current approaches and future directions.

Perception. (1) *Small, dense and occluded targets*: resolution-adaptive detectors, temporal confirmation from video, and detection coupled with tracking, building on size-aware designs and size-diverse data (Rekavandi et al. 2025; Feng et al. 2025; Zhang

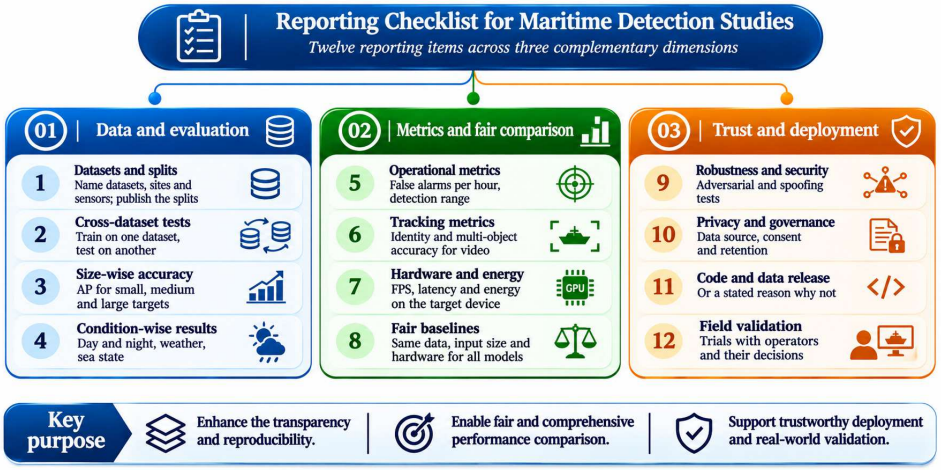


Fig. 4 Proposed reporting checklist for future studies on deep learning-based maritime security detection.

et al. 2026; Chen et al. 2026b). (2) *Adverse weather and sensor degradation*: standard robustness evaluation, out-of-distribution detection, and loss-free joint enhancement–detection across sensors (Chen et al. 2024; Tang et al. 2026; Zheng et al. 2026; Sun et al. 2021; Kaviyarasu et al. 2026). (3) *Scarce, imbalanced and domain-shifted data*: multimodal, climate- and platform-diverse datasets, transfer learning, synthetic data, rebalancing and foundation-model labelling, and—still absent in the corpus—explicit cross-sea-state/cross-sensor domain adaptation (Soloviev et al. 2020; Su et al. 2023; Farahnakian et al. 2021; Vaidya et al. 2026; Zhao et al. 2025a; Alsaaq 2026). (4) *Open-world and continual learning* for unpredictable threats and open vocabularies, including on-platform adaptation without forgetting (Yoon et al. 2024; Shi et al. 2026).

System. (5) *Real-time operation on constrained platforms*: preserving small-object accuracy after pruning/quantisation/distillation, edge–cloud scheduling, and benchmarks that report accuracy and energy together; hardware-aware architecture design for maritime tasks is absent from the corpus (Silva et al. 2024; Haijoub et al. 2025; Prete et al. 2024; Liu et al. 2026). (6) *Multimodal fusion and cooperative data* robust to missing/misaligned inputs, cross-checking rather than trusting AIS, and evaluated on public multimodal benchmarks that mostly do not yet exist (Eaton et al. 2018; Bloisi et al. 2017; Priyadharshini and Vadivazhagan 2024; Liu et al. 2025; Hanmante et al. 2026). (7) *Foundation and vision-language models*, whose accuracy on tiny objects, real-time behaviour and label efficiency are open, and which could mislead operators with fluent but wrong output (Meenakshi et al. 2025; Shi et al. 2026; Asim et al. 2026; Huang et al. 2026c; Zhang and Li 2026; Wu et al. 2026). (8) *Networked, communication-aware sensing* across satellites, UAVs, USVs and shore stations, co-designed with the links of Section 6 (Alqurashi et al. 2023; Wei et al. 2021; Pang et al. 2020; Zhao et al. 2024).

Trust and deployment. (9) *Adversarial robustness and cyber security*, including transferable attacks, tampered or spoofed inputs and physical attacks, the last absent from the corpus (Ashraf et al. 2023; Manasut et al. 2025; AlZaabi et al. 2020). (10) *Privacy, ethics and regulation*: formal privacy guarantees, data governance and a legal basis, recognising that technology alone cannot secure the seas (Okafor-Yarwood et al. 2024; Hall 2026; Rupa et al. 2026; Nandal et al. 2025; Huang et al. 2025). (11) *Explainability, trust and human interaction*: only two reviewed works address explainability (Rupa et al. 2026; Kumar et al. 2026); operators need uncertainty, fast verification and feedback, and evaluation of the human–algorithm team, not the algorithm alone (Brüggemann et al. 2016; Shahir et al. 2014; Chen et al. 2026a). (12) *Benchmarking and fair comparison* under a unified protocol (size-wise and condition-wise accuracy, cross-dataset results, speed/energy on named hardware, operational metrics) with released code and splits (Su et al. 2023; Alfaro et al. 2025; Kapoor and Sinha 2024). (13) *Operational deployment and validation* over long field trials, which are rarely reported (Alsaaq 2026); purpose-built surveillance systems and patrol platforms are the main exceptions (Brauchle et al. 2019; Markchom et al. 2026; Panico et al. 2020). (14) *Integration of detection with higher-level reasoning* into complete pixels-to-decision chains, of which the corpus has few examples (d'Afflisio et al. 2024; Niu et al. 2024; Brüggemann et al. 2016; Shahir et al. 2014).

Table 5: Fourteen open challenges for maritime-security object detection: why each matters, current approaches in the reviewed literature, and future directions.

Challenge	Importance	Current approaches	Future directions
Small, dense, occluded targets	Few pixels per target, overlapping detections	Multi-scale fusion, feature pyramids, attention, tiling (Feng et al. 2025; Si et al. 2023; Chen et al. 2023a; Zhang et al. 2026)	Resolution-adaptive detectors, temporal confirmation, and tracking-assisted detection
Weather and illumination	Fog, night, and glare lower contrast	Weather-aware detection, dehazing, low-light enhancement, joint denoising (Chen et al. 2024; Tang et al. 2026; Zheng et al. 2026; Sun et al. 2021; Kaviyarasu et al. 2026)	Robustness benchmarking, joint enhancement and detection, out-of-distribution alerts
Scarce and shifted data	Bad transfer between sites and sensors	Transfer learning, rebalancing, size-diverse datasets, synthetic data (Zhao et al. 2025a; Farahnakian et al. 2021; Chen et al. 2026b; Vaidya et al. 2026)	Multi-modal multi-climate datasets, synthetic data, standard metrics
Open-world and continual learning	New vessels and novel objects are encountered over time	Unknown-object detection; vision-language perception (Yoon et al. 2024; Shi et al. 2026)	Open-set detection with operator feedback, incremental learning without forgetting
Real-time edge operation	Low power, bandwidth and latency	Lightweight models, distillation, FPGA and onboard processing (Silva et al. 2024; Haijoub et al. 2025; Prete et al. 2024; Liu et al. 2026)	Hardware-aware design, compression maintaining small-object accuracy
Multimodal fusion and cooperative data	Misalignment and lack of sensors; AIS can be spoofed	Feature- and decision-level fusion; AIS cross-checking (Eaton et al. 2018; Fischer and Bauer 2010; Hanmante et al. 2026; Priyadharshini and Vadivazhagan 2024; Liu et al. 2025)	Fusion with tolerance for missing input, AIS validated by sensors, open multimodal benchmark
Foundation and vision-language models	Fewer labels, but not tested on small targets and in real time	CPU segment-anything; vision-language detection (Meenakshi et al. 2025; Shi et al. 2026)	Maritime adaptation and benchmarking against specialized detectors
Networked sensing and communication-aware detection	Wider coverage requires many cooperating platforms; links vary from 9.6 kbit/s to tens of Mbit/s	Multi-altitude systems, adaptive UAV patrols, onboard and semantic transmission (Boretti 2025; Madabhushi et al. 2024; Chen et al. 2023b; Markchom et al. 2026; Panico et al. 2020)	Accuracy reported across several data rates, split computing and task-oriented communication, cooperation over limited links, target hand-over
Adversarial and cyber security	Attacks on machine learning models and IoT networks	Transfer-attack analysis; maritime IoT threat analysis (Ashraf et al. 2023; Manasut et al. 2025; AlZaabi et al. 2020)	Adversarial training, spoofing detection, secure software updates
Privacy and regulation	Personal information, cross-border sharing	Federated learning; explainability-based privacy; governance studies (Okafor-Yarwood et al. 2024; Rupa et al. 2026; Nandal et al. 2025; Huang et al. 2025)	Formal privacy guarantees, governance, and audit procedures

Table 5 (continued)

Challenge	Importance	Current approaches	Future directions
Explainability and trust	Operators need explanations for alerts	Explainable detection; ontologies and behaviour models (Rupa et al. 2026; Brüggemann et al. 2016; Shahir et al. 2014; Chen et al. 2026a)	Explainable detectors, uncertainty visualization, human-in-the-loop evaluation
Benchmarking and reproducibility	Results not comparable across papers	Dataset surveys; comparative and data-centric evaluations (Su et al. 2023; Alfaro et al. 2025; Moosbauer et al. 2019; Kapoor and Sinha 2024)	Protocol, reporting checklist, code, model and split releases
Operational deployment and validation	Benchmarking results do not apply in the field	Purpose-built surveillance systems and patrol platforms (Brauchle et al. 2019; Markchom et al. 2026; Panico et al. 2020)	Long-duration field trials, operational metrics, human-in-the-loop evaluation
Link to higher-level reasoning	Detection alone is not threat identification	Fusion, trajectory classification, situation analysis (Fischer and Bauer 2010; d'Afflisio et al. 2024; Shahir et al. 2014; Chen et al. 2026a)	End-to-end systems from perception to decision support

Reporting checklist and roadmap. Many lessons concern reporting rather than methods; Fig. 4 gives a checklist (data and splits; multiple datasets and conditions; size-wise and operational metrics; speed and energy on named hardware; appropriate baselines; robustness and privacy; released code and data; field validation) whose adoption would make results comparable and show where detectors are deployment-ready. The challenges are then sequenced into a three-stage roadmap (Fig. 5) by a dependency rule—short term for shared resources others depend on, medium term for methods already begun but needing those resources, long term for deployment- and institution-dependent work. *Short term*: agree shared evaluation protocols and operational metrics (Su et al. 2023; Alfaro et al. 2025), release multimodal multi-site datasets, especially for infrared and radar (Varga et al. 2022; Chen et al. 2026b; Žust et al. 2023; Paolo et al. 2022; Patino et al. 2021), test robustness across sensors and weather (Ma et al. 2026), and report the required data rate alongside accuracy. *Medium term*: small-object, low-power detectors (Haijoub et al. 2025; Feng et al. 2025; Liu et al. 2026), joint enhancement–detection (Zheng et al. 2026; Kaviyarasu et al. 2026), fusion architectures that tolerate missing inputs and verify AIS (Eaton et al. 2018; Markchom et al. 2026; Paolo et al. 2024), foundation models benchmarked against specialized detectors (Meenakshi et al. 2025; Shi et al. 2026), federated learning with maritime-appropriate communication cost (McMahan et al. 2017; Razmi et al. 2022; Nandal et al. 2025; Huang et al. 2025), and split, onboard or task-oriented inference evaluated with realistic link traces (Matsubara et al. 2022; Madabhushi et al. 2024; Shao et al. 2022; Chen et al. 2023b). *Long term*: trustworthy systems that explain alerts and resist adversarial and cyber attack (Ashraf et al. 2023; Manasut et al. 2025; Rupa et al. 2026), open-world and continual learning (Yoon et al. 2024), integrated space–air–ground–sea coordination (Boretti 2025; Zhao et al. 2024; Markchom et al. 2026), perception that feeds behaviour and intent reasoning (Brüggemann et al. 2016;

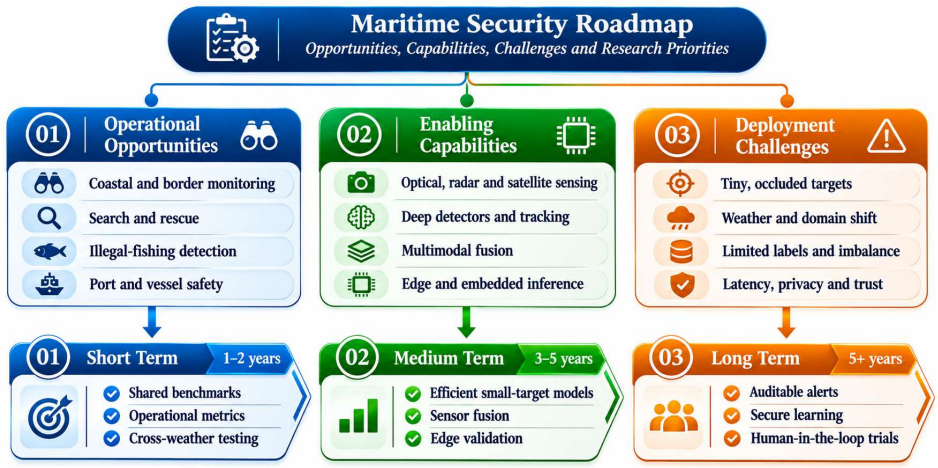


Fig. 5 Maritime security opportunities, enabling capabilities, deployment challenges, and a staged research roadmap.

Shahir et al. 2014; Chen et al. 2026a), and validation in prolonged operational trials with operators in the loop.

8 Limitations of this review

Several limitations should be considered when interpreting the findings. First, this is a structured review based on a documented search and double-screened selection, but its protocol was not registered in advance. Scopus and Web of Science were not searched; although OpenAlex provides additional coverage, it does not index all publishers equally, and some relevant studies may therefore have been missed. The review is also restricted to English-language peer-reviewed literature and excludes defence reports, patents and industrial documentation.

Second, the 544 studies added through the documented searches were assessed and coded primarily using abstracts and bibliographic records rather than full texts. In particular, 24 OpenAlex additions were classified using bibliographic records alone (Section 2). Consequently, some methodological details and application characteristics may not have been fully captured. The technology–application matrix (Table 3) and corpus-level counts are based largely on this classification, and the possibility of coding errors cannot be quantified. The numerical results in Supplementary Table S1 were extracted from abstracts and checked against the available records.

Third, although the supplementary tables report the study-level evidence used in the quantitative synthesis, the complete per-study coding inventory of all 655 studies is not published with this article, limiting the independent reproduction of the aggregate counts. Finally, no meta-analysis was attempted because the datasets, evaluation metrics and hardware configurations differ substantially across studies. Accordingly, the comparison in Table 2 represents a reasoned synthesis rather than a measured ranking. Furthermore, security-specific conclusions are based on the 151 studies assigned a specific task, many of which concern adjacent operational applications rather than

direct maritime security outcomes. Future work could strengthen the evidence base through broader database coverage, full-text assessment, more detailed study-level documentation, periodic updates and re-evaluation of representative detectors using the checklist in Fig. 4.

9 Conclusion

This survey synthesizes 655 primary studies on deep-learning-based maritime object detection and related perception tasks. Its principal result is not that one architecture is superior, but that the evidence base is uneven in both task coverage and evaluation quality. Fewer than a quarter of the studies name a specific task, and those that do include rescue, navigation and environmental aims as well as direct security ones. The distinction matters: strong object-detection performance can support security operations, but it does not by itself establish correct vessel identity, unlawful activity or hostile intent.

YOLO-family models are the most frequently coded approach (199 studies), followed by attention and transformer methods (123). Tracking, multimodal fusion, explainability, federated learning and detector security are considerably less represented. Research is concentrated on optical and SAR ship detection; infrared evidence is limited and no maritime-radar benchmark is shared across the reviewed studies. These counts describe the corpus as coded (Section 8), not the field as a whole.

Reported benchmark values cannot be converted into a fair global ranking. Dataset versions, splits, input resolution, class composition, IoU threshold, training schedule and hardware differ, while many papers omit repeated-run variation or operational measures. The SSDD results illustrate how near-saturated AP_{50} values can obscure more meaningful differences in small-object performance, cross-dataset generalization and deployment cost. A credible comparison should therefore use common splits, more than one dataset, size- and condition-stratified metrics, and latency/energy measurements on named target hardware.

The review also identifies a deployment gap. Maritime links vary greatly in capacity and latency, so transmitting video, transmitting compact detections, and running inference onboard have different consequences for bandwidth, verification and retraining. AIS should be treated as cooperative evidence to cross-check, not a generic communication channel or unquestioned ground truth. The proposed checklist and staged roadmap consequently prioritize reproducible benchmarks, infrared/radar data, robust small-object detection, trustworthy sensor fusion, communication-aware evaluation and long-duration field studies with operator involvement.

Within the limits set out in Section 8, the most actionable conclusion is that maritime AI research should shift from marginal gains on isolated benchmark splits toward comparable evidence of generalization, operational reliability and security-relevant outcomes.

Supplementary information. The online supplementary document contains seven evidence tables: reported benchmark results (Table S1), accessible evaluation protocols (Table S2), communication-aware methods (Table S3), sensing modalities (Table S4),

the corpus-level evidence profile (Table S5), representative datasets (Table S6) and application–platform mappings (Table S7).

Statements and Declarations

Author Contributions. Conceptualization, S.M.W. and M.M.S.; methodology, S.M.W. and M.A.; software, S.M.W. and F.U.R.; validation, S.M.W., M.A. and F.U.R.; formal analysis, S.M.W. and M.A.; investigation, S.M.W., M.A., F.U.R., M.B.J. and R.R.; data curation, S.M.W. and R.R.; writing—original draft preparation, S.M.W.; writing—review and editing, M.M.S., M.M.A., M.A., F.U.R., M.B.J. and M.U.K.; visualization, S.M.W.; supervision, M.M.S. and M.M.A.; project administration, S.M.W.; funding acquisition, M.M.S. All authors have read and agreed to the published version of the manuscript.

Funding. This research received no specific grant or financial support from any funding agency in the public, commercial, or not-for-profit sectors. The research and preparation of this manuscript were conducted without dedicated research funding. The publication charges, if applicable, will be covered by Mazliham Mohd Su'ud upon acceptance of the manuscript.

Competing interests. The authors declare no conflict of interest.

Ethics approval and consent to participate. Not applicable. This article reviews published literature and did not involve human participants, human data or animals.

Consent for publication. Not applicable.

Data availability. The study-level data supporting the quantitative synthesis (reported results, evaluation protocols, sensing modalities, corpus-level counts, datasets and application mappings) are provided in Supplementary Tables S1–S7. The per-study coding records of the 655 studies are not published with this article.

Use of generative AI. The authors used ChatGPT (OpenAI; model: GPT-6; source: https://chatgpt.com; accessed October 2026) as a generative AI tool to make minor improvements to specific sentences and enhance the presentation quality of selected figures. All AI-assisted changes were reviewed and verified by the authors, who take full responsibility for the final content of the manuscript.

References

Akash K, Maria HH, Kayalvizhi R (2026) YOLOv8 based intelligent maritime surveillance using satellite imagery. In: International Conference on AI for the Oceans, pp 97–113, https://doi.org/10.1007/978-981-95-3978-9_7

Alfaro MM, Wisniewska E, Baumgaertner NS, et al (2025) Performance evaluation of object detection models in maritime applications. In: International Conference on Science, Engineering Management and Information Technology, pp 51–66, https://doi.org/10.1007/978-3-032-01948-6_4

Alghazo J, Bashar A, Latif G, et al (2021) Maritime ship detection using convolutional neural networks from satellite images. In: 2021 10th IEEE International Conference on Communication Systems and Network Technologies (CSNT), pp 432–437, https://doi.org/10.1109/csnt51715.2021.9509628

Alqurashi FS, Trichili A, Saeed N, et al (2023) Maritime communications: A survey on enabling technologies, opportunities, and challenges. IEEE Internet of Things Journal 10(4):3525–3547. https://doi.org/10.1109/JIOT.2022.3219674

Alsaaq F (2026) Ocean-aware deep learning for civilian maritime object detection and tracking in complex ocean environments: A comprehensive review. Frontiers in Marine Science 13:1929278. https://doi.org/10.3389/fmars.2026.1929278

AlZaabi A, Talib MA, Nassif AB, et al (2020) A systematic literature review on machine learning in object detection security. In: 2020 IEEE 5th International Conference on Computing Communication and Automation (ICCCA), pp 136–139, https://doi.org/10.1109/ICCCA49541.2020.9250836

Apple Inc. (2026) HTTP live streaming (HLS) authoring specification for Apple devices. URL https://developer.apple.com/documentation/http-live-streaming/hls-authoring-specification-for-apple-devices, Apple Developer Documentation, accessed 29 October 2026

Ashraf I, et al (2023) A survey on cyber security threats in IoT-enabled maritime industry. IEEE Transactions on Intelligent Transportation Systems 24(2):2677–2690. https://doi.org/10.1109/TITS.2022.3164678

Asim MM, Smallwood CB, Tariq A, et al (2026) RampWatch: an In-the-Wild dataset and Text-Guided detection framework for recreational vessels. In: 2026 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp 7576–7585, https://doi.org/10.1109/WACV61042.2026.00731

Banu R, Sidram MH (2019) Object detection and tracking with occlusion handling in maritime Surveillance-A review. In: Recent Trends in Image Processing and Pattern Recognition (RTIP2R 2018), vol 1035. Springer, Singapore, p 88–103, https://doi.org/10.1007/978-981-13-9181-1_8

Banu R, Sidram MH (2022) Object detection with occlusion handling in maritime surveillance using histogram of orientated gradient. In: Computational Intelligence in Machine Learning: Select Proceedings of ICCIML 2021. Springer Nature Singapore, Singapore, p 381–389, https://doi.org/10.1007/978-981-16-8484-5_36

Bi Z, Yang H, Fu Y, et al (2026) Enhanced YOLO-Security deep learning model for ship detection in maritime security. Results in Engineering 29:108815. https://doi.org/10.1016/j.rineng.2025.108815

Bloisi DD, Previtali F, Pennisi A, et al (2017) Enhancing automatic maritime surveillance systems with visual information. IEEE Transactions on Intelligent Transportation Systems 18(4):824–833. https://doi.org/10.1109/tits.2016.2591321

Boretti A (2025) Navigating the future: the expanding role of unmanned surface vehicles in maritime security. Journal of Transportation Security 18(1):24. https://doi.org/10.1007/s12198-025-00314-x

Bouhlal BE, Brust CA (2026) Multi-layer adversarial detection for ship segmentation on NVIDIA Jetson Orin. In: Deutscher Luft- und Raumfahrtkongress (DLRK), https://doi.org/10.25967/650199

Bourtsoulatze E, Kurka DB, Gündüz D (2019) Deep joint source-channel coding for wireless image transmission. IEEE Transactions on Cognitive Communications and Networking 5(3):567–579. https://doi.org/10.1109/TCCN.2019.2919300

Bousetouane F, Morris B (2016) Fast CNN surveillance pipeline for fine-grained vessel classification and detection in maritime scenarios. In: 2016 13th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp 242–248, https://doi.org/10.1109/avss.2016.7738076

Bovcon B, Muhovič J, Perš J, et al (2019) The MaSTr1325 dataset for training deep USV obstacle detection models. In: 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp 3431–3438, https://doi.org/10.1109/IROS40897.2019.8967909

Bovcon B, Muhovič J, Vranac D, et al (2022) MODS—A USV-oriented object detection and obstacle segmentation benchmark. IEEE Transactions on Intelligent Transportation Systems 23(8):13403–13418. https://doi.org/10.1109/TITS.2021.3124192

Brauchle J, Bayer S, Hein D, et al (2019) MACS-Mar: a real-time remote sensing system for maritime security applications. CEAS Space Journal 11(1):35–44. https://doi.org/10.1007/s12567-018-0207-7

Brüggemann S, Bereta K, Xiao G, et al (2016) Ontology-based data access for maritime security. In: European Semantic Web Conference, pp 741–757, https://doi.org/10.1007/978-3-319-34129-3_45

Bu L, Yang S, Li B, et al (2026) DSSA-RT-DETR: dynamic Scattering-Aware SAR ship detection for complex backgrounds and small targets. Symmetry 18(10):1731. https://doi.org/10.3390/sym18101731

Cai Z, Zhou W (2026) Research on ship target detection in complex sea surface scenarios based on improved YOLOv7. Applied Sciences 16(4):1769. https://doi.org/10.3390/app16041769

Cane T, Ferryman J (2018) Evaluating deep semantic segmentation networks for object detection in maritime surveillance. In: 2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp 1–6, https://doi.org/10.1109/avss.2018.8639077

Cen Z, Chen H, Huang X, et al (2025) Enhanced YOLO with radar attention and Multi-Scale fusion for maritime target detection. In: 2025 2nd International Conference on Image Processing, Multimedia Technology and Machine Learning (IPMML), pp 142–146, https://doi.org/10.1109/IPMML68499.2025.11406982

Chan MT, Weed C (2012) Vessel detection in video with dynamic maritime background. In: 2012 IEEE Applied Imagery Pattern Recognition Workshop (AIPR), pp 1–5, https://doi.org/10.1109/aipr.2012.6528222

Chen C, Zeng W, Zhang X (2023a) HFPNet: super feature aggregation pyramid network for maritime remote sensing small-object detection. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 16:5973–5989. https://doi.org/10.1109/jstars.2023.3286483

Chen H, Xiao M, Pang Z (2022) Satellite-based computing networks with federated learning. IEEE Wireless Communications 29(1):78–84. https://doi.org/10.1109/MWC.008.00353

Chen H, He J, Meng S, et al (2023b) Semantic segmentation aided image transmission for maritime UAV surveillance network. In: 2023 International Conference on Wireless Communications and Signal Processing (WCSP), pp 689–694, https://doi.org/10.1109/WCSP58612.2023.10404483

Chen J, Huang J, Tan Y, et al (2025) LGNet: a lightweight Ghost-Enhanced network for efficient SAR ship detection. Remote Sensing 17(23):3800. https://doi.org/10.3390/rs17233800

Chen L, Patino L, Boyle J, et al (2026a) Semantic behaviour modelling for threat detection in maritime security. Applied Soft Computing 203:115968. https://doi.org/10.1016/j.asoc.2026.115968

Chen M, Sun J, Aida K, et al (2024) Weather-aware object detection method for maritime surveillance systems. Future Generation Computer Systems 151:111–123. https://doi.org/10.1016/j.future.2023.09.030

Chen W, Wu X, Yang Q, et al (2026b) MARS and ETMS: a Large-Scale dataset and Ship-Size sensitive detection method in maritime applications. IEEE Transactions on Intelligent Transportation Systems https://doi.org/10.1109/TITS.2026.3727609, early access

Chen Z, Chen Q (2025) Study on the application of improved target detection algorithm for UAV in maritime search and rescue. In: International Conference on

Mechatronics and Intelligent Control (ICMIC 2024), https://doi.org/10.1117/12.3045499

Cheng C, Jin F, Chang M (2026) SSANet: a scale and shape awareness network for ship detection in remote sensing scenes. In: Lecture Notes in Electrical Engineering, pp 207–217, https://doi.org/10.1007/978-981-95-8232-7_20

Cruz G, Bernardino A (2016) Aerial detection in maritime scenarios using convolutional neural networks. In: Lecture Notes in Computer Science, pp 373–384, https://doi.org/10.1007/978-3-319-48680-2_33

Dong H, Liu G, Liu W (2025) UAV-Based maritime search and rescue: Dataset construction and object detection evaluation. In: 2025 7th International Conference on Frontier Technologies of Information and Computer (ICFTIC), pp 570–574, https://doi.org/10.1109/icftic68075.2025.11324804

Dou J, Liu M, Ding Y, et al (2026) Heterogeneous USV-UAV cooperative edge service with energy-efficiency-aware joint optimization for maritime mobile edge computing. Journal of Systems Architecture 104051. https://doi.org/10.1016/j.sysarc.2026.104051

d'Afflisio E, Millefiori LM, Braca P, et al (2024) Maritrac: Maritime trajectory classification using object instance segmentation with model-based generated data augmentation. In: 2024 27th International Conference on Information Fusion (FUSION), pp 1–8, https://doi.org/10.23919/fusion59988.2024.10706510

Eaton RS, German S, Balasuriya A (2018) Maritime border security using sensors, processing, and platforms to detect dark vessels. In: 2018 IEEE International Symposium on Technologies for Homeland Security (HST), pp 1–5, https://doi.org/10.1109/ths.2018.8574147

Ebenezar US, Karthikeyan L, Velmurugan J, et al (2026) Multimodal Physics-Informed graph neural diffusion for edge detection of AIS-Dark illegal fishing vessels. Procedia Computer Science 283:604–612. https://doi.org/10.1016/j.procs.2026.06.143

Eski Y (2019) Maritime security: Problems of security versus safety practices. In: Shapiro LR, Maras MH (eds) Encyclopedia of Security and Emergency Management. Springer, Cham, p 1–5, https://doi.org/10.1007/978-3-319-69891-5_90-3

Fan X, Hu Z, Zhao Y, et al (2024) A small-ship object detection method for satellite remote sensing data. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 17:11886–11898. https://doi.org/10.1109/jstars.2024.3419786

Fang G, Han F, He C, et al (2025) LightSAR-Net: a lightweight multiscale dynamic SAR ship detection network. IEEE Geoscience and Remote Sensing Letters 22:1–5.

https://doi.org/10.1109/LGRS.2025.3628605

Farahnakian F, Zelioli L, Heikkonen J (2021) Transfer learning for maritime vessel detection using deep neural networks. In: 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), pp 1–6, https://doi.org/10.1109/itsc48978.2021.9565077

Fei Z, Xie Y, Deng D, et al (2024) HDetect-VS: tiny human object enhancement and detection based on visual saliency for maritime search and rescue. Applied Sciences 14(12):5260. https://doi.org/10.3390/app14125260

Feng X, Ran Y, Zhang Y, et al (2025) FX-DETR: a lightweight and efficient model for maritime small object detection. In: 2025 10th International Conference on Automation, Control and Robotics Engineering (CACRE), pp 165–169, https://doi.org/10.1109/cacre66141.2025.11119563

Fischer Y, Bauer A (2010) Object-oriented sensor data fusion for wide maritime surveillance. In: 2010 International WaterSide Security Conference, pp 1–6, https://doi.org/10.1109/wssc.2010.5730244

Frost D, Tapamo JR (2013) Detection and tracking of moving objects in a maritime environment using level set with shape priors. EURASIP Journal on Image and Video Processing 2013(1):42. https://doi.org/10.1186/1687-5281-2013-42

Gamage C, Dinalankara R, Samarabandu J, et al (2023) A comprehensive survey on the applications of machine learning techniques on maritime surveillance to detect abnormal maritime vessel behaviors. WMU Journal of Maritime Affairs 22(4):447–477. https://doi.org/10.1007/s13437-023-00312-7

Gao Y, Wu C, Ren M, et al (2025) Refined Anchor-Free model with feature enhancement mechanism for ship detection in infrared images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 18:12946–12960. https://doi.org/10.1109/JSTARS.2025.3569294

Garaiman E, Radoi A (2025) Real-time ship detection in SAR images using a transformer-based architecture under Schedule-Free optimization. In: 2025 International Conference on Machine Intelligence for GeoAnalytics and Remote Sensing (MIGARS), pp 1–4, https://doi.org/10.1109/MIGARS67156.2025.11232155

Geng H, Liu X, Ge X, et al (2025a) SKAN-YOLO: a precision detection algorithm for small maritime targets. In: 2025 3rd International Conference on Algorithm, Image Processing and Machine Vision (AIPMV), pp 84–89, https://doi.org/10.1109/AIPMV67185.2025.11290047

Geng W, Yi J, Cheng L (2025b) An efficient detector for maritime search and rescue object based on unmanned aerial vehicle images. Displays 87:102994. https://doi.org/10.1016/j.displa.2025.102994

Gonzales-Martínez R, Machacuay J, Rotta P, et al (2021) Faster R-CNN with a cross-validation approach to object detection in radar images. In: 2021 IEEE International Conference on Aerospace and Signal Processing (INCAS), pp 1–4, https://doi.org/10.1109/incas53599.2021.9666929

Guan T, Chang S, Wang C, et al (2025) SAR small ship detection based on enhanced YOLO network. Remote Sensing 17(5):839. https://doi.org/10.3390/rs17050839

Guo H, Li J, Liu J, et al (2022) A survey on space-air-ground-sea integrated network security in 6g. IEEE Communications Surveys & Tutorials 24(1):53–87. https://doi.org/10.1109/COMST.2021.3131332

Guo Y, Shen Q, Ai D, et al (2024) Sea-IoUTracker: a more stable and reliable maritime target tracking scheme for unmanned vessel platforms. Ocean Engineering 299:117243. https://doi.org/10.1016/j.oceaneng.2024.117243

Haijoub A, Hatim A, Guerrero-Gonzalez A, et al (2025) CNN architectures on FPGA for maritime object detection: A review of methods and challenges. In: 2025 7th International Symposium on Advanced Electrical and Communication Technologies (ISAECT), pp 1–6, https://doi.org/10.1109/ISAECT68904.2025.11318704

Hall AE (2026) Maritime security technologies and coastal neo-fortification. Political Geography 126:103483. https://doi.org/10.1016/j.polgeo.2025.103483

Hanmante S, Kolhar S, Patil S (2026) A comprehensive review of multimodal artificial intelligence techniques in maritime surveillance using AIS, SAR, and optical data fusion. Array 30:100931. https://doi.org/10.1016/j.array.2026.100931

Hassan SI, Karnani C, Zia MYI (2026) YOLO-Based object detection for intelligent maritime surveillance and risk identification. In: 2026 Global Conference on Wireless and Optical Technologies (GCWOT), pp 1–5, https://doi.org/10.1109/gcwot69191.2026.11499730

He J, Huang C (2024) YOLOv8-SARship: an efficient and accurate detection method for SAR ship images. In: 2024 5th International Conference on Computer Vision, Image and Deep Learning (CVIDL), pp 280–284, https://doi.org/10.1109/CVIDL62147.2024.10604104

Hoang-Ngoc D, Nguyen DMT, Huynh-The T (2025) SSC-YOLO: Spectral-Spatial fusion convolution for ship detection in SAR imagery. In: 2025 International Conference on Advanced Technologies for Communications (ATC), pp 1–6, https://doi.org/10.1109/ATC67618.2025.11268806

Huang L, Liao M, Li S, et al (2026a) MSA-YOLO: an optimized UAV object detection algorithm for Low-Visibility maritime. Remote Sensing 18(13):2065. https://doi.org/10.3390/rs18132065

Huang Y, Liu W, Lin Y, et al (2025) FLCSDet: federated learning-driven cross-spatial vessel detection for maritime surveillance with privacy preservation. IEEE Transactions on Intelligent Transportation Systems 26(1):1177–1192. https://doi.org/10.1109/TITS.2024.3488497

Huang Z, Shang W (2024) DSP-YOLO: a SAR ship detection algorithm for multiscale sequence fusion based on fusion attention. In: Proceedings of the 2024 7th International Conference on Computational Intelligence and Intelligent Systems, pp 44–52, https://doi.org/10.1145/3708778.3708785

Huang Z, Guan W, Zhou J, et al (2026b) Shape-Prior dynamic refinement network with diffusion wavelets for Fine-Grained ship detection in remote sensing images. Remote Sensing 18(16):2649. https://doi.org/10.3390/rs18162649

Huang Z, et al (2026c) Robust maritime object detection via a hybrid DINOv2 and YOLOv8n architecture. Electronics 15(16):3708. https://doi.org/10.3390/electronics15163708

International Maritime Organization (1972) Convention on the International Regulations for Preventing Collisions at Sea (COLREGs), 1972. IMO, London, U.K.

International Maritime Organization (1974) International Convention for the Safety of Life at Sea (SOLAS), 1974. IMO, London, U.K.

International Maritime Organization (2003) International Ship and Port Facility Security (ISPS) Code. IMO, London, U.K.

International Telecommunication Union (2014) Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile frequency band. Tech. Rep. Recommendation ITU-R M.1371-5, International Telecommunication Union, URL https://www.itu.int/rec/R-REC-M.1371

International Telecommunication Union (2022) Technical characteristics for a VHF data exchange system in the VHF maritime mobile band. Tech. Rep. Recommendation ITU-R M.2092-1, International Telecommunication Union, URL https://www.itu.int/rec/R-REC-M.2092

John A, Guan D (2025) YOLOBiEMA: SAR ship detection model based on spatial learning and feature fusion. In: International Conference on Measurement, Communication, and Virtual Reality (MCVR 2024), https://doi.org/10.1117/12.3067597

Kanjir U, Greidanus H, Oštir K (2018) Vessel detection and classification from spaceborne optical images: A literature survey. Remote Sensing of Environment 207:1–26. https://doi.org/10.1016/j.rse.2017.12.033

Kapoor A, Sinha BB (2024) Object detection in maritime surveillance: A Data-Centric comparison of YOLO and Transformer-Based models. In: 2024 IEEE 8th International Conference on Information and Communication Technology (CICT), pp 1–6, https://doi.org/10.1109/cict64037.2024.10899718

Kaviyarasu A, Nivashini V, Rajesh G, et al (2026) Noise-aware maritime object detection in UAV imagery using joint Denoising-Detection learning. In: 2026 International Conference on Wireless Communications Signal Processing and Networking (WiSPNET), pp 1–6, https://doi.org/10.1109/wispnet69615.2026.11489438

Kim DJ, Song J, Han S, et al (2025) Leveraging SAR satellites and AI for Real-Time dark vessel detection and maritime surveillance. In: OCEANS 2025 Brest, pp 1–5, https://doi.org/10.1109/OCEANS58557.2025.11104622

Kumar KP, Stavroulakis PJ, Thiruthuvanathan MM, et al (2026) Intelligent maritime vessel detection and classification toward offshore wind energy zone surveillance using YOLOv11 and explainable hybrid machine learning for autonomous maritime monitoring. Ocean Engineering 365:127443. https://doi.org/10.1016/j.oceaneng.2026.127443

Külcü S (2026) Enhancing radar PPI-Based maritime object detection using hybrid preprocessing and deep learning approach. IEEE Access 14:111715–111730. https://doi.org/10.1109/ACCESS.2026.3715472

Lalasa KL, Srija RJV, Kumar KP (2024) Maritime security-illegal fishing detection using deep learning. In: 2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS), pp 1–5, https://doi.org/10.1109/ickecs61492.2024.10617407

Lázaro F, Raulefs R, Wang W, et al (2019) VHF data exchange system (VDES): an enabling technology for maritime communications. CEAS Space Journal 11(1):55–63. https://doi.org/10.1007/s12567-018-0214-8

Li J, Hua Y, Xue M (2025) MSO-DETR: a lightweight detection transformer model for small object detection in maritime search and rescue. Electronics 14(12):2327. https://doi.org/10.3390/electronics14122327

Liang L, Han T (2025) YOLOv8-IPAN: enhancing ship detection in maritime surveillance with integrated pyramid attention and progressive feature fusion. In: 2025 5th Asia-Pacific Conference on Communications Technology and Computer Science (ACCTCS), pp 485–493, https://doi.org/10.1109/ACCTCS66275.2025.00090

Liu Q, Chen H, Zhao F (2025) Maritime target detection algorithm based on fusion of visible and infrared images. The Journal of Supercomputing 81(1):22. https://doi.org/10.1007/s11227-024-06509-8

Liu R, Zhu Z, Ge H, et al (2026) Towards scale-adaptive and lightweight maritime ship object detection via dual-cross multi-scale knowledge distillation. Ocean Engineering 343:123206. https://doi.org/10.1016/j.oceaneng.2025.123206

Lu K, Wu Z, Qi H, et al (2026) TCLS-YOLOv11n: a multi-module collaborative optimization network for tiny-target detection in UAV-based maritime search and rescue. Measurement Science and Technology 37(28):285408. https://doi.org/10.1088/1361-6501/ae8612

Luglio M, Roseti C, Zampognaro F (2018) VDES performance evaluation for future e-navigation services. In: Proceedings of the 15th International Joint Conference on e-Business and Telecommunications (ICETE), pp 67–75, https://doi.org/10.5220/0006850200670075

Lv J, Wang Z, Dang C, et al (2026) D 2 -MS: a ship detection accelerator via reparameterized Channel-Splitting network and Dual-Parallel Non-Blocking pipelined architecture. IEEE Geoscience and Remote Sensing Letters pp 1–1. https://doi.org/10.1109/LGRS.2026.3743053

Lyu H, Shao Z, Cheng T, et al (2023) Sea-surface object detection based on electro-optical sensors: A review. IEEE Intelligent Transportation Systems Magazine 15(2):190–216. https://doi.org/10.1109/MITS.2022.3198334

Ma J, Zhou Y, He L, et al (2026) Infrared ship detection in challenging maritime environments: A YOLOv12-Based robustness and benchmark evaluation. Journal of Ocean Engineering and Science 11(6):2103–2124. https://doi.org/10.1016/j.joes.2026.06.004

Ma X, Cheng J, Li A, et al (2024) AMANet: advancing SAR ship detection with adaptive multi-hierarchical attention network. IEEE Access 12:105952–105967. https://doi.org/10.1109/access.2024.3436591

Madabhushi B, Zhang Z, Alves TAO, et al (2024) Efficient onboard SAR image processing for real-time ship detection and data reduction. In: OCEANS 2024 - Halifax, pp 1–10, https://doi.org/10.1109/OCEANS55160.2024.10753824

Manasut P, Ibtasham MS, Yaradanakul Z, et al (2025) Investigating the transferability of TOG adversarial attacks in YOLO models in the maritime domain. IEEE Access 13:126444–126464. https://doi.org/10.1109/access.2025.3589512

Markchom T, et al (2026) Multi-altitude, multimodal maritime surveillance system. IEEE Sensors Journal 26(5):7101–7119. https://doi.org/10.1109/JSEN.2025.3649549

Matsubara Y, Levorato M (2021) Neural compression and filtering for edge-assisted real-time object detection in challenged networks. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp 2272–2279, https://doi.org/10.1109/

ICPR48806.2021.9412388

Matsubara Y, Levorato M, Restuccia F (2022) Split computing and early exiting for deep learning applications: Survey and research challenges. ACM Computing Surveys 55(5):90. https://doi.org/10.1145/3527155

Mazzeo A, et al (2025) Quasi-Tracking of dark ships with an AI-Based Multi-Mission Multi-Frequency approach. In: IGARSS 2025 - 2025 IEEE International Geoscience and Remote Sensing Symposium, pp 6912–6916, https://doi.org/10.1109/IGARSS55030.2025.11242273

McMahan B, Moore E, Ramage D, et al (2017) Communication-efficient learning of deep networks from decentralized data. In: Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), pp 1273–1282

Meenakshi K, Joshi S, Dhipu TM, et al (2025) CPU implementation of SAM based maritime object segmentation and detection. In: 2025 IEEE Space, Aerospace and Defence Conference (SPACE), pp 1–6, https://doi.org/10.1109/space65882.2025.11170617

Mehla N, Ishita, Talukdar R, et al (2023) Object detection in autonomous maritime vehicles: comparison between YOLO v8 and EfficientDet. In: International Conference on Data Science and Network Engineering, pp 125–141, https://doi.org/10.1007/978-981-99-6755-1_10

Messaoud W, Trabelsi R, Cabani A, et al (2023) Multi-head self attention for enhanced object detection in the maritime domain. In: 2023 International Conference on Cyberworlds (CW), pp 179–184, https://doi.org/10.1109/cw58918.2023.00034

Messaoud W, Trabelsi R, Cabani A, et al (2024) Maritime object detection using attention mechanism. Signal, Image and Video Processing 18(2):1833–1845. https://doi.org/10.1007/s11760-023-02897-1

Milios A, Bereta K, Chatzikokolakis K, et al (2019) Automatic fusion of satellite imagery and AIS data for vessel detection. In: 2019 22th International Conference on Information Fusion (FUSION), pp 1–5, https://doi.org/10.23919/FUSION43075.2019.9011339

Mitra N, Maji A, Agarwalla P (2017) Integrated maritime security & surveillance. In: 2017 Devices for Integrated Circuit (DevIC), pp 766–768, https://doi.org/10.1109/devic.2017.8074055

Mohan N, et al (2024) A multifaceted look at starlink performance. In: Proceedings of the ACM Web Conference 2024, pp 2723–2734, https://doi.org/10.1145/3589334.3645328

Monani UJ, Ramudu B, Panigrahi N (2025) Enhanced SAR-Optical image fusion for maritime dark ship detection using principal component analysis and wavelet transform. In: 2025 IEEE India Geoscience and Remote Sensing Symposium (InGARSS), pp 956–961, https://doi.org/10.1109/InGARSS67683.2025.11584037

Moosbauer S, König D, Jäkel J, et al (2019) A benchmark for deep learning based object detection in maritime environments. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp 916–925, https://doi.org/10.1109/cvprw.2019.00121

Müller R, Brook C (2010) Net based waterside security applications: From small solutions to maritime security networks. In: 2010 International WaterSide Security Conference, pp 1–6, https://doi.org/10.1109/wssc.2010.5730266

Nandal P, Bohra N, Mann P, et al (2025) YOLOv11 with transformer attention for real-time monitoring of ships: A federated learning approach for maritime surveillance. Results in Engineering 27:106297. https://doi.org/10.1016/j.rineng.2025.106297

de Nicolas JM, Mata-Moya D, Jarabo-Amores M, et al (2013) Neural network based solutions for ship detection in SAR images. In: 2013 18th International Conference on Digital Signal Processing (DSP), pp 1–6, https://doi.org/10.1109/ICDSP.2013.6622836

Niu J, et al (2026a) DGG-YOLO: a robust framework for infrared ship detection with graph reasoning and enhanced Multi-Scale fusion. IEEE Access 14:37321–37335. https://doi.org/10.1109/ACCESS.2026.3672240

Niu J, et al (2026b) LD-YOLO: a lightweight dynamic Convolution-Based YOLOv8n framework for robust ship detection in SAR imagery. IEEE Geoscience and Remote Sensing Letters 23:1–5. https://doi.org/10.1109/LGRS.2025.3630098

Niu L, Fan Y, Liu T, et al (2024) MOMT: a maritime real-time visual multiobject tracking algorithm based on unmanned surface vehicles. IEEE Sensors Journal 24(21):35429–35447. https://doi.org/10.1109/jsen.2024.3455572

Okafor-Yarwood I, Eastwood O, Chikowore N, et al (2024) Technology and maritime security in Africa: Opportunities and challenges in Gulf of Guinea. Marine Policy 160:105976. https://doi.org/10.1016/j.marpol.2023.105976

Pang Y, Wang D, Wang D, et al (2020) A space-air-ground integrated network assisted maritime communication network based on mobile edge computing. In: 2020 IEEE World Congress on Services (SERVICES), pp 269–274, https://doi.org/10.1109/SERVICES48979.2020.00060

Panico A, Fragonara LZ, Al-Rubaye S (2020) Adaptive detection tracking system for autonomous UAV maritime patrolling. In: 2020 IEEE 7th International Workshop on Metrology for AeroSpace (MetroAeroSpace), pp 539–544, https://doi.org/

10.1109/metroaerospace48742.2020.9160214

Paolo F, Lin TT, Gupta R, et al (2022) xView3-SAR: detecting dark fishing activity using synthetic aperture radar imagery. In: Advances in Neural Information Processing Systems 35 (NeurIPS 2022), Datasets and Benchmarks Track, pp 37604–37616, https://doi.org/10.52202/068431-2726

Paolo FS, et al (2024) Satellite mapping reveals extensive industrial activity at sea. Nature 625(7993):85–91. https://doi.org/10.1038/s41586-023-06825-8

Patino L, Cane T, Ferryman J (2021) A comprehensive maritime benchmark dataset for detection, tracking and threat recognition. In: 2021 17th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS), pp 1–8, https://doi.org/10.1109/avss52988.2021.9663739

Prasad DK, Rajan D, Rachmawati L, et al (2017) Video processing from electro-optical sensors for object detection and tracking in a maritime environment: A survey. IEEE Transactions on Intelligent Transportation Systems 18(8):1993–2016. https://doi.org/10.1109/TITS.2016.2634580

Prete RD, et al (2024) Enhanced maritime monitoring via onboard processing of raw multi-spectral imagery by deep learning. In: IGARSS 2024-2024 IEEE International Geoscience and Remote Sensing Symposium, pp 1713–1717, https://doi.org/10.1109/igarss53475.2024.10641068

Priyadharshini SD, Vadivazhagan K (2024) Enhanced vessel detection in maritime surveillance using multi-modal data integration and deep learning. In: 2024 8th International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)(I-SMAC), pp 1090–1099, https://doi.org/10.1109/i-smac61858.2024.10714663

Rajesh G, Nivashini V, Akther MR, et al (2026) RGB-to-Infrared translation and object detection for Drone-Based maritime surveillance using CycleGAN and YOLOv8. In: 2026 International Conference on Machine Learning and Autonomous Systems (ICMLAS), pp 40–46, https://doi.org/10.1109/ICMLAS67792.2026.11483947

Razmi N, Matthiesen B, Dekorsy A, et al (2022) Ground-assisted federated learning in LEO satellite constellations. IEEE Wireless Communications Letters 11(4):717–721. https://doi.org/10.1109/LWC.2022.3141120

Rekavandi AM, Xu L, Boussaid F, et al (2025) A guide to image-and video-based small object detection using deep learning: Case study of maritime surveillance. IEEE Transactions on Intelligent Transportation Systems 26(3):2851–2879. https://doi.org/10.1109/TITS.2025.3530678

Ruan C, Hong L, Gao J (2026) Marine-YOLO: a high-precision object detection algorithm for complex maritime environments. Applied Ocean Research 169:104994.

https://doi.org/10.1016/j.apor.2026.104994

Rupa C, Nandini CD, Chandana BHN (2026) An XAI based approach for privacy preservation of maritime vessel detection. In: International Conference on Information and Communication Technology for Competitive Strategies, pp 183–193, https://doi.org/10.1007/978-3-032-20609-1_18

Sarajlić E, Karna H, Braović M, et al (2026) AI-Based object detection for improved maritime vessel type identification. In: 2026 11th International Conference on Smart and Sustainable Technologies (SpliTech), pp 1–6, https://doi.org/10.23919/SpliTech70133.2026.11685499

Shahir HY, Glässer U, Nalbandyan N, et al (2014) Maritime situation analysis: A multi-vessel interaction and anomaly detection framework. In: 2014 IEEE Joint Intelligence and Security Informatics Conference, pp 192–199, https://doi.org/10.1109/jisic.2014.36

Shao J, Mao Y, Zhang J (2022) Learning task-oriented communication for edge inference: An information bottleneck approach. IEEE Journal on Selected Areas in Communications 40(1):197–211. https://doi.org/10.1109/JSAC.2021.3126087

Shao Z, Wu W, Wang Z, et al (2018) SeaShips: a large-scale precisely annotated dataset for ship detection. IEEE Transactions on Multimedia 20(10):2593–2604. https://doi.org/10.1109/TMM.2018.2865686

Sharma R, Saqib M, Lin CT, et al (2024) Maritime surveillance using instance segmentation techniques. In: International Conference on Data Science and Communication, pp 31–47, https://doi.org/10.1007/978-981-99-5435-3_3

Shi B, Wang Z, Cao T, et al (2026) MMD-Net: a Vision-Language perception model for maritime object detection and application. Journal of Ocean University of China 25(2):505–518. https://doi.org/10.1007/s11802-026-6153-7

Shi J, Jin J, Zhang J (2018) Object detection based on saliency and sea-sky line for USV vision. In: 2018 IEEE 4th Information Technology and Mechatronics Engineering Conference (ITOEC), pp 1581–1586, https://doi.org/10.1109/itoec.2018.8740763

Si J, Song B, Wu J, et al (2023) Maritime ship detection method for satellite images based on multiscale feature fusion. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 16:6642–6655. https://doi.org/10.1109/jstars.2023.3296898

Silva HV, de Figueiredo FA, Mafra SB (2024) Performance evaluation of edge computing object detection models for maritime surveillance on a Raspberry Pi. In: 2024 IEEE Latin-American Conference on Communications (LATINCOM), pp 1–6, https://doi.org/10.1109/latincom62985.2024.10770682

Soloviev V, Farahnakian F, Zelioli L, et al (2020) Comparing CNN-based object detectors on two novel maritime datasets. In: 2020 IEEE International Conference on Multimedia & Expo Workshops (ICMEW), pp 1–6, https://doi.org/10.1109/icmew46912.2020.9106019

Song W, et al (2026) A comprehensive review and case study of underwater object detection from traditional techniques to large vision language models. Discover Artificial Intelligence 6:247. https://doi.org/10.1007/s44163-026-00895-z

Su L, Chen Y, Song H, et al (2023) A survey of maritime vision datasets. Multimedia Tools and Applications 82(19):28873–28893. https://doi.org/10.1007/s11042-023-14756-9

Sullivan GJ, Ohm JR, Han WJ, et al (2012) Overview of the high efficiency video coding (HEVC) standard. IEEE Transactions on Circuits and Systems for Video Technology 22(12):1649–1668. https://doi.org/10.1109/TCSVT.2012.2221191

Sun X, Liu T, Yu X, et al (2021) Unmanned surface vessel visual object detection under all-weather conditions with optimized feature fusion network in YOLOv4. Journal of Intelligent & Robotic Systems 103(3):55. https://doi.org/10.1007/s10846-021-01499-8

Sun Y, Lian J (2025) IRSD-Net: an adaptive infrared ship detection network for small targets in complex maritime environments. Remote Sensing 17(15):2643. https://doi.org/10.3390/rs17152643

Tang C, Ji Z (2025) IRS-DETR: efficient infrared ship detection transformer for complex SeaAir backgrounds. IEEE Geoscience and Remote Sensing Letters 22:1–5. https://doi.org/10.1109/LGRS.2025.3601741

Tang Q, Liu X, Lin Q, et al (2026) Remote sensing image dehazing for maritime object detection: A physics-driven and frequency-aware network. Ocean Engineering 363:126775. https://doi.org/10.1016/j.oceaneng.2026.126775

Teutsch M, Krüger W (2010) Classification of small boats in infrared images for maritime surveillance. In: 2010 International WaterSide Security Conference, pp 1–7, https://doi.org/10.1109/wssc.2010.5730289

Tu-Salaam AI, Alamaniotis M (2026) Benchmarking deep learning architectures for Image-Based surveillance to improve maritime port security. In: International Conference on Emerging Internet, Data & Web Technologies, pp 156–167, https://doi.org/10.1007/978-3-032-22798-0_15

Vaidya SK, et al (2026) A virtual testbed for maritime search and rescue: Customizing YOLO for small object detection in Unreal Engine. Journal of Ocean Engineering and Science 11(2):597–613. https://doi.org/10.1016/j.joes.2025.12.018

Varga LA, Kiefer B, Messmer M, et al (2022) SeaDronesSee: a maritime benchmark for detecting humans in open water. In: 2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp 3686–3696, https://doi.org/10.1109/wacv51458.2022.00374

Venkata SB, Ashirova A, Karwande VS, et al (2025) Towards accurate maritime surveillance: A hybrid CNN-Transformer architecture for ship detection in SAR imagery. In: 2025 International Conference on Innovations in Intelligent Systems: Advancements in Computing, Communication, and Cybersecurity (ISAC3), pp 1–6, https://doi.org/10.1109/isac364032.2025.11156808

Žust L, Perš J, Kristan M (2023) LaRS: a diverse panoptic maritime obstacle detection dataset and benchmark. In: 2023 IEEE/CVF International Conference on Computer Vision (ICCV), pp 20247–20257, https://doi.org/10.1109/ICCV51070.2023.01857

Wang C, Li L, Lu K, et al (2025a) TPS-YOLOv8n: an advanced algorithm for small object detection in complex maritime search and rescue scenarios using unmanned aerial vehicle imagery. Measurement Science and Technology 36(7):075404. https://doi.org/10.1088/1361-6501/ade326

Wang C, Lu K, Ma B, et al (2025b) A multi-task joint learning-based algorithm for small object detection in maritime UAV imagery under adverse weather conditions. Measurement Science and Technology 36(9):095414. https://doi.org/10.1088/1361-6501/ae080c

Wang H, Li D, Wang C, et al (2026) FSP-YOLO: a lightweight algorithm for ship detection in SAR images under complex backgrounds. IEEE Access 14:118353–118366. https://doi.org/10.1109/ACCESS.2026.3719648

Wei D, Cai C (2026) A lightweight SAR ship detection method with multi-scale feature fusion and task-aligned mechanism. Measurement Science and Technology 37(29):295105. https://doi.org/10.1088/1361-6501/ae8926

Wei S, Zeng X, Qu Q, et al (2020) HRSID: a high-resolution SAR images dataset for ship detection and instance segmentation. IEEE Access 8:120234–120254. https://doi.org/10.1109/ACCESS.2020.3005861

Wei T, Feng W, Chen Y, et al (2021) Hybrid satellite-terrestrial communication networks for the maritime internet of things: Key technologies, opportunities, and challenges. IEEE Internet of Things Journal 8(11):8910–8934. https://doi.org/10.1109/JIOT.2021.3056091

Wu G, Zhang X, Xu X, et al (2026) Ship target detection in SAR images with vision foundation models. In: 2026 7th International Conference on Pattern Recognition and Machine Learning (PRML), pp 98–101, https://doi.org/10.1109/PRML71097.2026.11687309

Xu C, Gao L, Zhang J, et al (2026) Maritime multimodal data resource system—Infrared-visible dual-modal dataset for ship detection. Journal of Radars https://doi.org/10.12000/JR25144

Yang C, Wang J, Li K, et al (2025) An improved YOLOv8 with AFPN for maritime UAV object detection. In: 2025 6th International Conference on Computer Communication and Network Security (CCNS), pp 158–162, https://doi.org/10.1109/ccns66849.2025.11338038

Yoon H, Lee Y, Lee H, et al (2024) MarUCOD: unknown but concerned object detection in maritime environments. In: International Conference on Pattern Recognition, pp 253–268, https://doi.org/10.1007/978-3-031-78447-7_17

Zhang J, Huang X (2025) FFC-YOLO: a SAR ship detection network with enhanced feature representation and dynamic cross-scale fusion. Engineering Research Express 7(2):025233. https://doi.org/10.1088/2631-8695/adcdca

Zhang J, Li Y (2026) SPFDet: CLIP Text-Prior-Guided Structure-Enhanced detector for Fine-Grained ship detection in remote sensing images. Sensors 26(14):4476. https://doi.org/10.3390/s26144476

Zhang K, Gui X, Ren D, et al (2021a) Energy–latency tradeoff for computation offloading in UAV-assisted multiaccess edge computing system. IEEE Internet of Things Journal 8(8):6709–6719. https://doi.org/10.1109/JIOT.2020.2999063

Zhang T, et al (2020) LS-SSDD-v1.0: a deep learning dataset dedicated to small ship detection from large-scale Sentinel-1 SAR images. Remote Sensing 12(18):2997. https://doi.org/10.3390/rs12182997

Zhang T, et al (2021b) SAR ship detection dataset (SSDD): official release and comprehensive data analysis. Remote Sensing 13(18):3690. https://doi.org/10.3390/rs13183690

Zhang W, Gao XZ, Yang CF, et al (2022) A object detection and tracking method for security in intelligence of unmanned surface vehicles. Journal of Ambient Intelligence and Humanized Computing 13(3):1279–1291. https://doi.org/10.1007/s12652-020-02573-z

Zhang W, Shang Y, Chen S, et al (2026) Edge-Topology transformer for dense infrared maritime object detection. In: 2026 International Annual Conference on Complex Systems and Intelligent Science (CSIS-IAC), pp 185–190, https://doi.org/10.1109/csis-iac70275.2026.11585167

Zhang Y, Li QZ, Zang FN (2017) Ship detection for visual maritime surveillance from non-stationary platforms. Ocean Engineering 141:53–63. https://doi.org/10.1016/j.oceaneng.2017.06.022

Zhang Y, Jia Y, Tang Y (2024) Accurate detection of arbitrary ship directions using SAR based on RTMDet. Remote Sensing Letters 16(2):156–169. https://doi.org/10.1080/2150704X.2024.2440666

Zhao B, Li J, Zhao J, et al (2025a) Optimizing drone-captured maritime rescue image object detection through dataset rebalancing under sample constraints. The Visual Computer 41(13):11221–11237. https://doi.org/10.1007/s00371-025-04098-y

Zhao B, Zhao J, Song R, et al (2025b) Enhanced YOLO11 for lightweight and accurate drone-based maritime search and rescue object detection. PLOS One 20(7):e0321920. https://doi.org/10.1371/journal.pone.0321920

Zhao C, Fu X, Dong J, et al (2023) LPDNet: a lightweight network for SAR ship detection based on Multi-Level laplacian denoising. Sensors 23(13):6084. https://doi.org/10.3390/s23136084

Zhao H, Ji F, Wang Y, et al (2024) Space–air–ground–sea integrated network with federated learning. Remote Sensing 16(9):1640. https://doi.org/10.3390/rs16091640

Zheng Y, Ao X, Xie W, et al (2026) An enhanced end-to-end low-light ship detection framework for UAV-based maritime surveillance. Ocean Engineering 366:127916. https://doi.org/10.1016/j.oceaneng.2026.127916

Zhu F, Pan B, Im N (2025) DMODT: dynamic maritime environment object detection and tracking model for MASS. Ocean Engineering 340:122270. https://doi.org/10.1016/j.oceaneng.2025.122270

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- MaritimeSecuritySurveyPaper.pdf