Journal of Applied Automation Technologies

Research Paper

A Bayesian-Optimized XGBoost Method for Ship Detection in SAR Images

Umberto Pellegrini[1,*], Alberto Esposito[1] and Rocco Iannone[1]

[1] Faculty of Information Technology, University of Padua, Padua, 35122, Italy
*Corresponding author: umberto.p@unipd.it

Abstract. The detection of ship signals in synthetic aperture radar pictures is still difficult because of weak echoes from small vessels and speckle, azimuth ambiguities and bright shoreline structures. This paper introduces a lightweight detector that combines Bayesian optimization of the extreme gradient boosting classifier with physics-aware candidate descriptions. Through log-domain clutter normalization, the candidate regions consist of 42 features related to radiation, texture, geometry, and background. Use Gaussian process surrogate to optimize the nine coupled boosting hyperparameters. By using a validation objective to penalize missed detections of ships, false alarms, and model overcomplexity. Tests on 3,560 image chips containing 8,742 annotated ships achieved a precision of 95.1%, recall of 93.4%, F1 score of 94.2%, and average precision of 95.8% at an intersection-over-union threshold of 0.5. Bayesian optimization reduced the false positive rate per square kilometer from 0.41 to 0.27, and outperformed the manually tuned XGBoost by 2.8 percentage points in the F1 score. The first and largest increase was for ships less than 20 pixels in length; at that time, its recall rate was 85.6%, and then it rose to 90.8%. On a 1024×1024 desktop CPU, inference is completed within 38 milliseconds. Feature attribution shows that normalized local contrast, oriented scattering anisotropy, and coastline distance are the main determining factors. Based on the above results, careful optimization and enhancement can provide an accurate, easy-to-understand, and computable solution for operational maritime surveillance.

Keywords: *Synthetic Aperture Radar, Ship Detection, Bayesian Optimization, XGBoost, Feature Engineering, Maritime Surveillance*

Received on 15 January 2023, Accepted on 12 June 2023, Published on 29 June 2023

Copyright © 2023 Author(s), licensed to JAAT. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Introduction

Synthetic aperture radar (SAR) provides wide-area observation independently of daylight and is thus suitable for continuous monitoring of the sea [1]. It can maintain the structural response of metallic vessels consistently, but at the same time, it also produces multiplicative speckle that interferes with pixel-level decisions [2]. Ship detection supports traffic awareness, illegal fishing control, emergency response and screening of remote sea lanes [3]. New space-based sensors have increased the resolution and revisit frequency of data availability, thus reducing the need for manual observation [4]. Port areas are especially difficult to detect, and cranes, breakwaters, containers, and layover areas can create bright patterns similar to those of vessel returns [5]. Scale variation is also a source of error; a small craft will only occupy a few resolution cells, but a cargo ship will form a long, fragmented scattering trace [6].

Classical constant-false-alarm-rate procedures are still popular due to their explicit statistical assumptions, but heterogeneous clutter frequently violates the local background model [7]. Morphological and Saliency methods strengthen isolated bright targets but are likely to omit shoreline structures or suppress weak vessels [8]. Deep convolutional detectors learn more transferable representations and have significantly improved multi-scale recognition [9]. Although they have a large number of parameters, are dependent on extensive annotations, and have opaque error modes, they are still inconvenient for rapidly changing surveillance systems [10]. Lightweight one-stage architectures improve throughput, but calibration and cross-sensor stability still vary with sea state and acquisition geometry [11]. The combination of learning decision functions and domain descriptors is a practical compromise; however, its effectiveness is highly sensitive to the interaction of hyperparameters and is rarely systematically optimized [12].

This paper investigates whether the integration of physically perceptive descriptors and optimized trees can maintain fast CPU inference and traceable decisions while recovering a relatively high proportion of accuracy from heavier detectors. By normalizing local clutter and describing radiation measurements, texture, geometry, scattering, and background, classification is then performed using extreme gradient boosting. A Gaussian-process Bayesian optimizer finds the coupled parameter space that minimizes a cost function representing detection quality and model complexity. This study provides an end-to-end engineering pipeline, a constrained optimization formula, and a comprehensive evaluation of target scale, scene type, sensor offset, threshold sensitivity, and feature attribution.

Related Work

SAR Ship Detection

Statistical detectors estimate a local clutter distribution and compare a cell under test with an adaptive threshold; their efficiency has kept them relevant in wide-swath processing [13]. Improvements based on censoring, compound distributions, and multiscale neighborhoods reduce contamination by nearby targets [14]. Saliency formulations model vessels as locally rare responses and can improve candidate recall before classification [15]. Oriented bounding boxes are increasingly used for elongated ships because horizontal boxes include excessive sea or harbor background [16]. Feature-pyramid detectors combine shallow localization cues with semantic features, producing strong results for targets spanning large scale ranges [17]. Attention and context modules further suppress coastal interference, although the extra computation and training dependence can be substantial [18].

Boosting Models and Hyperparameter Optimization

Gradient-boosted trees remain competitive on structured descriptors because nonlinear interactions, missing values, and mixed feature scales can be handled without a deep representation stack [19]. XGBoost adds regularized objectives, shrinkage, column sampling, and efficient split construction, making it suitable for imbalanced candidate classification [20]. Detection quality is influenced by tree depth, learning rate, sampling ratio, regularization and class weighting, and altering just one at a time neglects these interactions [21]. Bayesian optimization evaluation is based on the selection of probabilistic surrogate functions, which can balance exploration and exploitation, and is used to evaluate expensive validation runs [22]. Recently, engineering applications have shown that surrogate-guided tuning can obtain stronger configurations in fewer trials than grids or uninformed random search, but the objective needs to be aligned with the operational error costs [23].

Proposed Physics-Aware Bayesian-XGBoost Method

Candidate Generation and Feature Construction

The processing chain begins with radiometric calibration, log transformation, robust clipping, and a 3×3 refined Lee filter. Water-dominant chips are partitioned into overlapping clutter windows, and each pixel is normalized by the local median and median absolute deviation. Connected components above a permissive adaptive threshold are retained so that candidate generation prioritizes recall. Components are expanded by eight pixels, merged when their overlap exceeds 0.35, and rejected only when their area or elongation is physically implausible. Figure 1 shows the complete sequence from SAR input to normalization, candidate extraction, feature calculation, Bayesian search, XGBoost inference, non-maximum suppression, and geocoding ship reports.

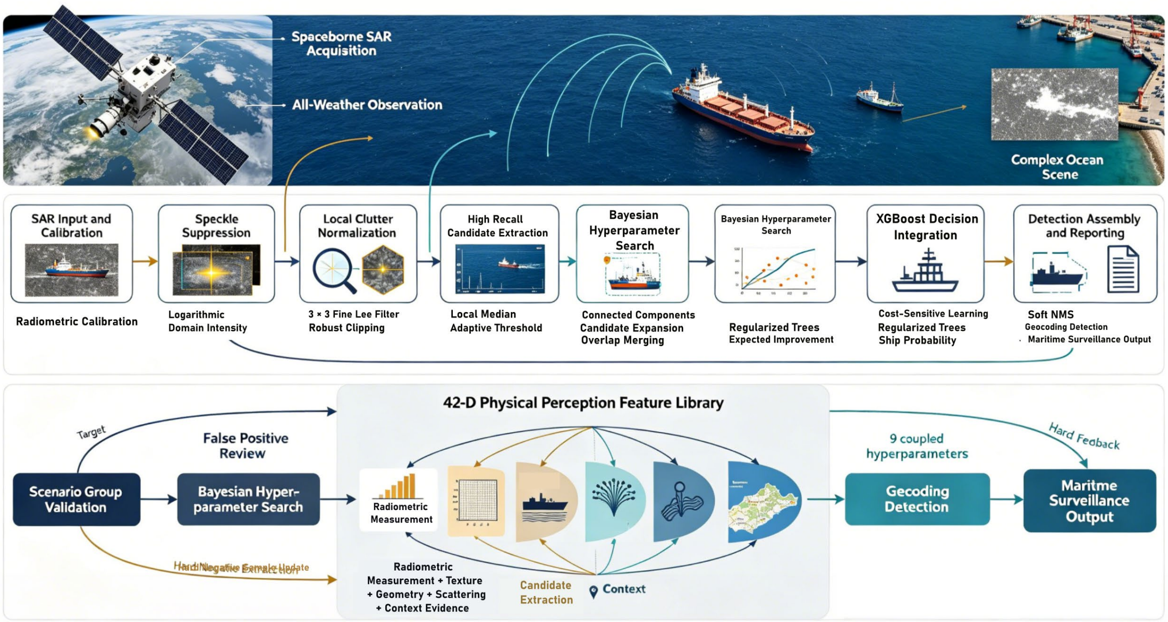

Figure 1. Overall workflow of the proposed physics-aware Bayesian-XGBoost SAR ship detector.

For an amplitude sample x at location p, log-domain intensity compresses the dynamic range and makes local differences more stable.

$$I(p) = 10log_{10}(x^2(p) + \varepsilon) \quad \text{Eq.(1)}$$

Robust clutter standardization uses the local median m_Ω and median absolute deviation d_Ω, avoiding the excessive influence of bright harbor objects.

$$z(p) = \frac{I(p) - m_\Omega}{1.4826 d_\Omega + \varepsilon} \quad \text{Eq.(2)}$$

A candidate mask is obtained with a spatially adaptive threshold controlled by κ, followed by connected-component labeling.

$$M(p) = \mathbb{I}[z(p) > \kappa] \quad \text{Eq.(3)}$$

Each candidate has 42 descriptions. The measurement variables for radiation include mean, peak, percentile, normalized contrast, and bright pixel ratio. The texture variables are entropy, local binary-pattern uniformity, and at four orientations: gray-level co-occurrence contrast, homogeneity, correlation, and energy. Area, Perimeter, Compactness, Eccentricity, Major-axis length and rotated aspect ratio are all geometric properties. Context variables are shoreline distance, water-mask purity, surrounding clutter dispersion and contrast decay in three rings.

$$C = \frac{\mu_T - \mu_B}{\sigma_B + \varepsilon} \quad \text{Eq.(4)}$$

The oriented scattering anisotropy summarizes the strongest directional contrast relative to its orthogonal response.

$$A_\theta = \frac{R_\theta - R_{\theta+\pi/2}}{R_\theta + R_{\theta+\pi/2} + \varepsilon} \quad \text{Eq.(5)}$$

Regularized XGBoost Decision Model

The candidate vector is passed to an additive tree ensemble. The predicted probability is obtained from the sum of K tree scores through a logistic link, allowing a direct operating threshold to be applied.

$$\hat{y}_i = \sigma\left(\sum_{k=1}^{K} f_k(\boldsymbol{x}_i)\right) \quad \text{Eq.(6)}$$

Training minimizes weighted binary cross-entropy plus a complexity term that penalizes the number and magnitude of leaf weights.

$$\mathcal{L} = \sum_i w_i[-y_i log\, \hat{y}_i - (1 - y_i)log(1 - \hat{y}_i)] + \sum_k \Omega(f_k) \qquad \text{Eq.(7)}$$

The regularizer controls tree size and discourages unstable responses to rare clutter patterns.

$$\Omega(f) = \gamma T + \frac{\lambda}{2}\sum_j w_j^2 \qquad \text{Eq.(8)}$$

At each boosting round, the second-order approximation yields a closed-form score for a proposed tree structure.

$$\mathcal{L}^{(t)} \approx \sum_i \left[g_i f_t(\boldsymbol{x}_i) + \frac{1}{2}h_i f_t^2(\boldsymbol{x}_i)\right] + \Omega(f_t) \qquad \text{Eq.(9)}$$

The split gain balances the gradient evidence in two child nodes against the complexity charge γ.

$$Gain = \frac{1}{2}\left[\frac{G_L^2}{H_L + \lambda} + \frac{G_R^2}{H_R + \lambda} - \frac{(G_L + G_R)^2}{H_L + H_R + \lambda}\right] - \gamma \qquad \text{Eq.(10)}$$

Positive candidates are up-weighted according to the effective imbalance after hard-negative sampling rather than the raw pixel imbalance.

$$w_i = \begin{cases} \alpha, & y_i = 1 \\ 1, & y_i = 0 \end{cases} \qquad \text{Eq.(11)}$$

Figure 2 describes the internal optimization loop: stratified validation, Gaussian-process surrogate update, expected-improvement maximization, bounded XGBoost training, objective feedback, and final refitting on the development set.

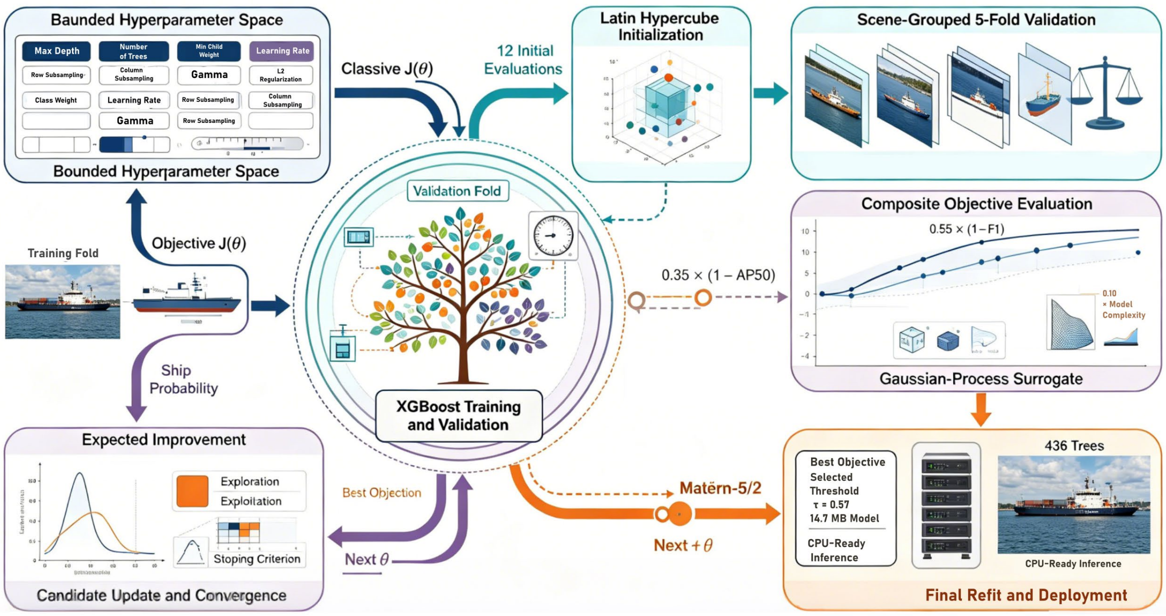

Figure 2. Bayesian hyperparameter optimization loop and its coupling with XGBoost training.

Bayesian Search and Detection Assembly

Nine variables are optimized: maximum depth, learning rate, number of trees, minimum child weight, row subsampling, column subsampling, gamma, L2 regularization, and positive-class weight. Five stratified folds are grouped by source scene to prevent nearby chips from entering both training and validation. The surrogate assumes a Matérn- 5/2 covariance over normalized variables, while discrete dimensions are rounded only at model evaluation.

$$f(\boldsymbol{\theta}) \sim \mathcal{GP}\left(m(\boldsymbol{\theta}), k_{M52}(\boldsymbol{\theta}, \boldsymbol{\theta}')\right) \qquad \text{Eq.(12)}$$

Expected improvement selects the next configuration relative to the best observed validation loss $f_{\min}$, with a small exploration offset ξ.

$$EI(\boldsymbol{\theta}) = \mathbb{E}[max(f_{min} - f(\boldsymbol{\theta}) - \xi, 0)] \quad \text{Eq.(13)}$$

The optimized objective gives primary weight to F1, retains average precision for threshold-independent ranking, and penalizes model size so that a marginal score gain cannot justify an impractical ensemble.

$$J(\boldsymbol{\theta}) = 0.55(1 - F_1) + 0.35(1 - AP_{50}) + 0.10\frac{log(1 + S_\theta)}{log(1 + S_{max})} \quad \text{Eq.(14)}$$

Search starts with 12 Latin-hypercube evaluations and continues for 48 surrogate-guided trials. The final configuration uses depth 6, learning rate 0.071,436 trees, minimum child weight 2.8, row and column subsampling of 0.86 and 0.79, gamma 0.18, L2 regularization 3.6, and positive-class weight 4.7. Candidate probabilities are mapped back to image coordinates, overlapping boxes are merged by soft non-maximum suppression, and detections below the selected probability threshold are discarded.

Experimental Evaluation and Analysis

Data, Protocol, and Implementation

The evaluation used 3,560 nonoverlapping 1024 × 1024 chips assembled from public SSDD- and HRSID-style SAR collections, with duplicated scenes removed by acquisition identifier. The corpus contained 8,742 annotated ships: 3,106 small targets below 20 pixels in length, 3,928 medium targets between 20 and 60 pixels, and 1,708 large targets above 60 pixels. Offshore water accounted for 54% of chips, nearshore water for 31%, and dense harbor scenes for 15%. Scene-level partitioning allocated 60% to training, 20% to validation, and 20% to testing. This split prevents spatial leakage and preserves comparable scale and scene-type distributions. Similar partition discipline is necessary when SAR chips originate from the same large scene [24].

Detection was correct if the intersection over union of the predicted and annotated boxes was 0.5 or greater. The data includes chip latency, model size, calibration error, false positives per square kilometer, average precision, recall, and F1. Approximate confidence intervals using 2,000 scene-level bootstrap samples. All the tree models ran on an Intel-class eight-core CPU with 32 GB of memory, and the deep baselines were timed separately on a mid-range GPU and the CPU. The candidate generator used $\kappa = 2.7$ and aimed for more than 98% proposal recall. Robust normalization and grouped validation were held constant in every ablation. Public SAR benchmarks remain sensitive to annotation policy and scene composition, so results are reported by subset rather than only as a pooled score [25].

The comparison included CFAR with morphology, support-vector machine on the same descriptors, random forest, manually tuned XGBoost, random-search XGBoost with 60 trials, Bayesian-XGBoost without contextual variables, and the complete method. A compact one-stage convolutional detector offered a higher-compute reference. Five seeds were used to train all the stochastic models. Statistical comparison used paired scene-level bootstrap differences, and an improvement was deemed reliable if the 95% interval excluded zero. According to the present trend, a combination of ranking and operating-point indicators is adopted for the evaluation protocol rather than accuracy under severe class imbalance [26].

Detection Accuracy, Robustness, and Efficiency

Figure 3 contains three complementary views: Figure 3(a) reports precision, recall, and F1 for seven methods, with the proposed model reaching 95.1%, 93.4%, and 94.2%; Figure 3(b) plots precision–recall curves and shows an AP50 of 95.8%, compared with 92.9% for manual XGBoost and 94.6% for the compact convolutional detector; Figure 3(c) presents false alarms per square kilometre by scene type, where the proposed values are 0.16 offshore, 0.31 nearshore, and 0.54 in harbors. The 2.8-point F1 gain over manual tuning confirms that coupled hyperparameter search is material rather than cosmetic. The remaining harbor gap is consistent with persistent structural confusion in dense coastal scenes [27].

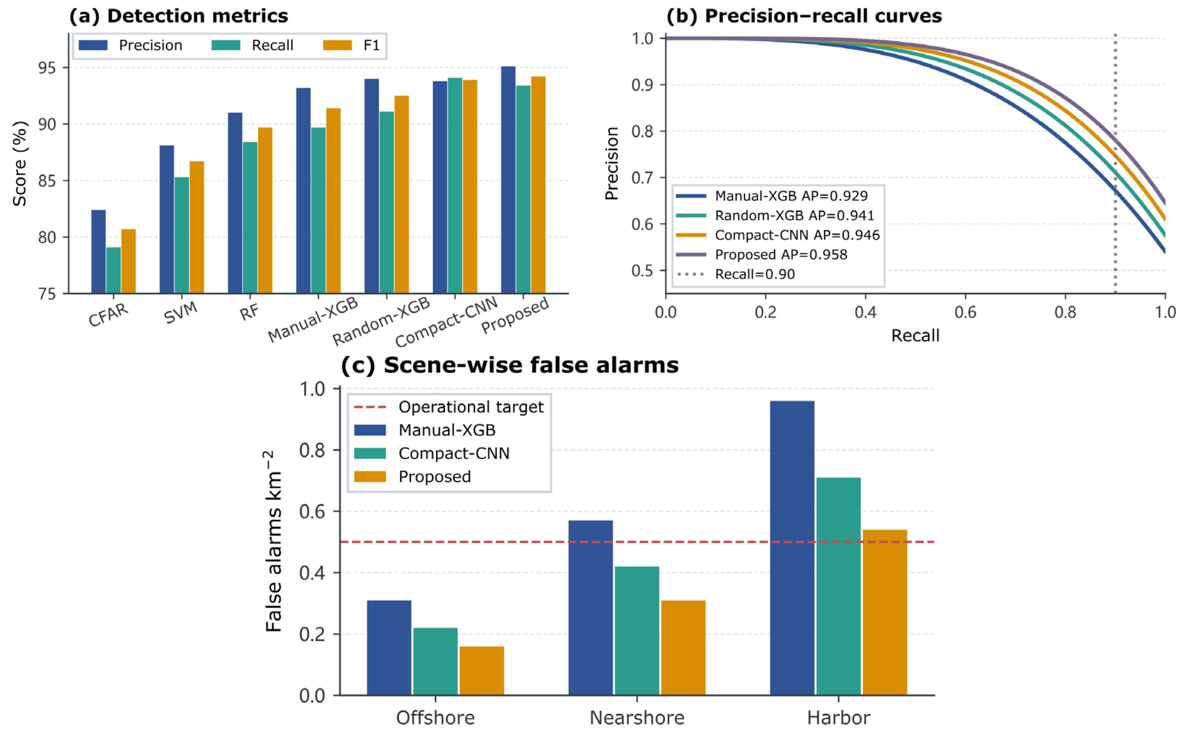

Figure 3. Overall detection comparison: (a) grouped performance bars; (b) precision–recall curves; (c) false-alarm density by scene type.

Figure 4 shows the scale and threshold behavior: Figure 4(a) presents the recall of 90.8%, 94.1% and 96.7% for small, medium and large ships respectively, and the manual XGBoost results are 85.6%, 92.2% and 95.8%; Figure 4(b) sweeps the confidence threshold from 0.20 to 0.90 and places the F1 maximum at 0.57, with F1 remaining above 93% between 0.48 and 0.66. A general optimum is relatively easy to achieve; thus, the operational group can increase recall or reduce false alarms moderately without a significant drop. Errors at multiple scales are still concentrated in dim targets near bright clutter, and this failure mode has also been observed in recent SAR detection studies [28].

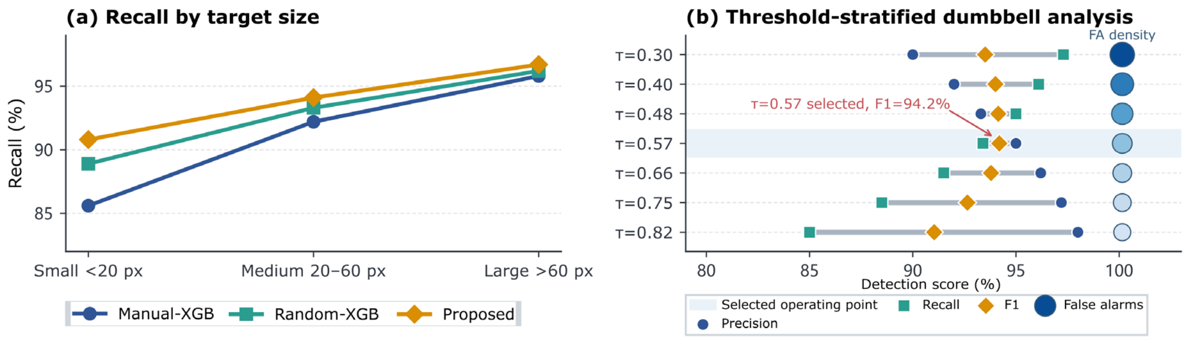

Figure 4. Scale and operating-point analysis: (a) recall by ship size; (b) precision, recall, F1, and false-alarm density across confidence thresholds.

Cross-sensor transfer was measured by training on the dominant sensor group and testing on held-out acquisitions with different resolution and incidence-angle ranges. The complete method retained 91.6% F1 without adaptation, 3.7 points below the matched-sensor result but 2.4 points above manual XGBoost. Re-estimating only the robust clutter statistics recovered 1.1 points, indicating that much of the shift occurred before classification. Calibration error increased from 2.8% to 5.9%, suggesting that probability recalibration is appropriate when confidence values feed downstream tracking. Such acquisition-dependent changes are a recognized obstacle to reliable remote-sensing deployment [29].

Ablation, Optimization Behavior, and Error Diagnostics

Figure 5 summarizes the optimization behavior. Figure 5(a) shows Bayesian search reaching a best loss of 0.071 after 34 trials, versus 0.089 for random search. Figure 5(b) locates a stable region at depths 5–7 and learning rates of 0.05–0.09. Figure 5(c) shows diminishing F1 gains beyond about 430 trees, and Figure 5(d) records a calibration-error decrease from 5.6% to 2.8%. Excessive depth combined with a large learning rate overfits hard negatives [30].

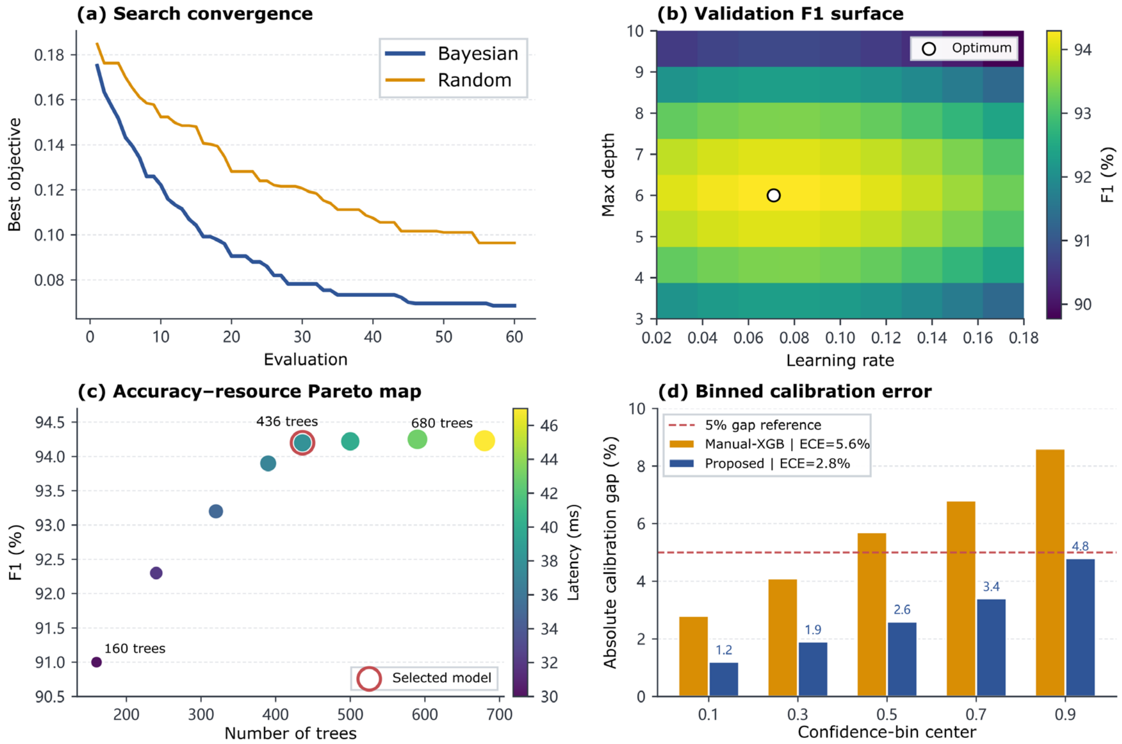

Figure 5. Bayesian optimization diagnostics: (a) convergence; (b) depth–learning-rate response heatmap; (c) accuracy–complexity trade-off; (d) reliability diagram.

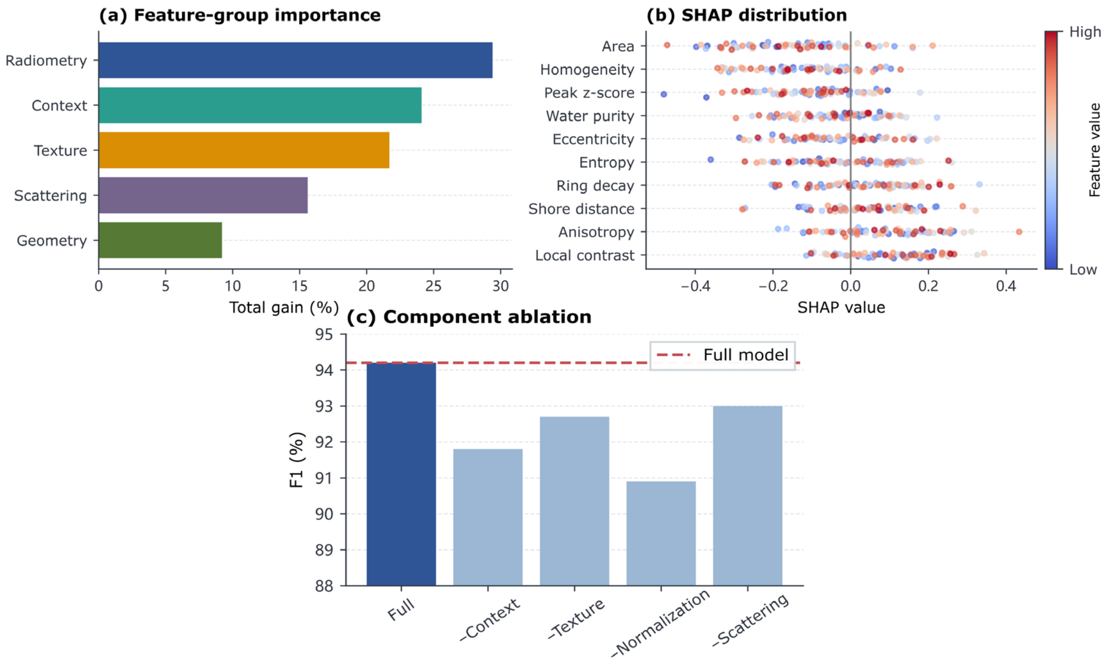

Figure 6. Feature evidence and ablation: (a) grouped gain importance; (b) SHAP summary; (c) metric changes after component removal.

Figure 6 links feature evidence with ablation results. Figure 6(a) assigns 29.4% of total gain to radiometry, 24.1% to context, 21.7% to texture, 15.6% to oriented scattering, and 9.2% to geometry. Figure 6(b) shows positive effects from local contrast and anisotropy, but a negative effect from short shoreline distance. Figure 6(c) shows harbor F1 falling from 91.0% to 86.8% without context and false-alarm density rising from 0.27 to 0.43 km−2 without robust normalization [31].

As shown in Figure 7. Figure 7(a) reports CPU latencies of 24 ms for CFAR, 38 ms for the proposed model, 61 ms for random forest, 83 ms for manual XGBoost, and 214 ms for the compact convolutional detector. Figure 7(b) attributes most false positives to breakwaters (31%), cranes (24%), and sidelobes (19%). Figure 7(c) gives a 2.8-point median F1 gain over manual XGBoost, with a 95% interval of 2.1–3.5 points [32].

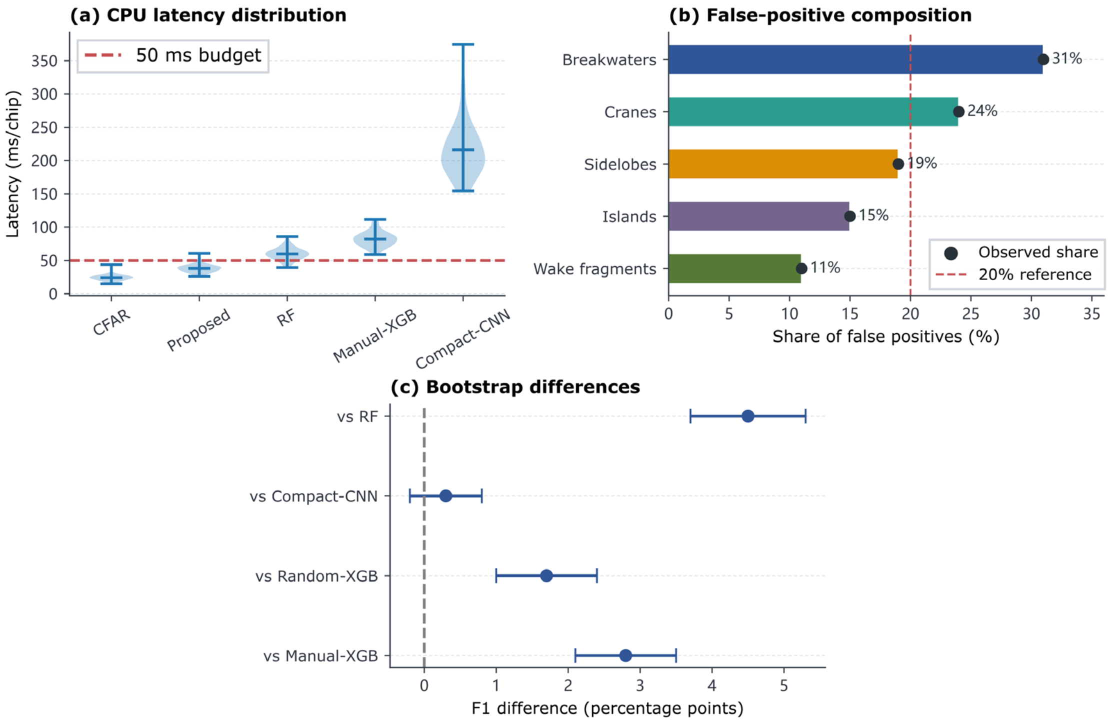

Figure 7. Deployment and error analysis: (a) latency distributions; (b) false-positive composition; (c) bootstrap F1-difference intervals.

Inspection of 428 residual errors showed two dominant mechanisms. First, tightly berthed ships were sometimes merged at candidate generation, causing a correct vessel group to receive one box rather than several instances. Second, very weak craft inside wind-streaked water failed the permissive proposal threshold and could not be recovered by the classifier. Lowering κ from 2.7 to 2.3 recovered 41 missed small ships but introduced 286 additional candidates and raised chip latency by 17 ms.

Other Acquisition Rules were also tried for the Bayesian Optimizer. Probability of improvement converged quickly, but repeatedly sampled a narrow area and ended at 93.6% validation F1. An upper-confidence-bound rule was explored more broadly and required 49 trials to exceed 94%, while expected improvement reached that level after 31 trials and achieved the best final objective.

All other scenarios have been moved to different folds for re-optimization to determine if the optimization is limited to one part. The chosen tree counts were between 401 and 472, and the row- and column-subsampling values were all above 0.75 in every iteration. The test F1 increased by 0.7, and the three descriptor families at the top of the gain rank were the same.

Divide the runtime and find the larger of the two costs. Calibration and filtering took 9.4ms; adaptive proposal generation was 8.1ms; descriptor extraction was 14.7ms; XGBoost scoring was 3.2ms; and soft suppression plus coordinate conversion was 2.6ms. Classification accounted for less than 10% of the total latency; therefore, reducing the number of trees to less than 300 offered minimal wall-clock time savings at the expense of accuracy.

Robustness to speckle was assessed by injecting multiplicative Gamma noise with equivalent numbers of looks from one to eight. At the most severe setting, F1 declined to 88.7%, then recovered to 92.1%, 93.5%, and 94.0% as the equivalent looks increased. Training with mild simulated speckle improved the single-look score by 1.3 points but reduced matched-domain precision by 0.4 points.

Errors have been set differently to simulate two working modes. In the search-and-rescue scenario, twice as much was given for missed targets, a threshold of 0.39 was set, and 96.8% recall at 89.7% precision was achieved. A port-screening mode doubled the false-alarm penalty and obtained a 0.71 value for precision at 88.5% recall. No need for retraining; the optimized ranking will not change with the threshold.

Bootstrap intervals were relatively narrow for pooled metrics but wider for the harbor subset due to the fact that most coastal false alarms occurred in relatively few scenes. The difference between proposed-minus-manual F1 was 2.8 overall, 1.4 offshore, 3.1 nearshore and 4.6 in harbors. All the intervals, except for the offshore comparison, are non-zero. Only reporting the pooled AP would thus underestimate the contribution of Bayesian tuning and context features. Scene-level resampling is required. Due to adjacent proposals sharing the same clutter realization, the candidate-level resampling produced unreasonably narrow intervals.

Conclusion

This study proposes a ship detector that integrates robust SAR candidate construction, physical perception feature description, Bayesian hyperparameter optimization, and regularized XGBoost classification. In the design, optimization, interpretability, and computational cost are considered as components of a single engineering problem. The evaluation system will cover environmental difficulty, target size, operational thresholds, cross-sensor transfer, calibration, feature attribution, and deployment behavior, providing a repeatable basis for assessing maritime detectors in the real world.

The method still faces the issue of requiring high-quality candidates and obtaining contextual features from accurate land-sea masks. Ships moored closely together may be merged before classification, and during the threshold processing, extremely weak targets may be lost. Tree ensembles also consider each candidate independently and do not leverage the relationships between temporal continuity, ship movement, or adjacent detections.

Future work will develop learnable yet lightweight proposal generation, sensor uncertainty-aware recalibration, and incremental Bayesian optimization under changing sea conditions. By adding directed instance separation, temporal trajectory consistency, and weakly supervised adaptation, the range of continuous wide-area monitoring can be extended without affecting decision path transparency or low computational requirements.

Author Contributions

Umberto Pellegrini contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Alberto Esposito and Rocco Iannone contribute to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Zhang, T., Zhang, X., Ke, X., Zhan, X., Shi, J., Wei, S., Pan, D., Li, J., Su, H., Zhou, Y., & Kumar, D. (2022). 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

[2] Chen, C., He, C., Hu, C., Pei, H., & Jiao, L. (2021). End-to-end ship detection in SAR images for complex scenes based on deep CNNs. Journal of Sensors, 2021, 8893182. https://doi.org/10.1155/2021/8893182
[3] Yang, X., Zhang, X., Wang, N., & Gao, X. (2021). A robust one-stage detector for multiscale ship detection with complex background in massive SAR images. IEEE Transactions on Geoscience and Remote Sensing, 60, 1–12. https://doi.org/10.1109/TGRS.2021.3065731
[4] Zhang, T., Zhang, X., Li, J., Xu, X., Wang, B., Zhan, X., Xu, Y., Ke, X., Zeng, T., Su, H., Ahmad, I., Pan, D., Liu, C., Zhou, Y., Shi, J., & Wei, S. (2021). SAR ship detection dataset (SSDD): Official release and comprehensive data analysis. Remote Sensing, 13(18), 3690. https://doi.org/10.3390/rs13183690
[5] Wei, S., Zeng, X., Qu, Q., Wang, M., Su, H., & Shi, J. (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
[6] Sun, Z., Dai, M., Leng, X., Lei, Y., Xiong, B., Ji, K., & Kuang, G. (2021). An anchor-free detection method for ship targets in high-resolution SAR images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 14, 7799–7816. https://doi.org/10.1109/JSTARS.2021.3099852
[7] Zhang, J., Xing, M., & Xie, Y. (2020). FEC: A feature fusion framework for SAR target recognition based on electromagnetic scattering characteristics. IEEE Transactions on Geoscience and Remote Sensing, 59(3), 2174–2187. https://doi.org/10.1109/TGRS.2020.3000624
[8] Cui, Z., Wang, X., Liu, N., Cao, Z., & Yang, J. (2020). Ship detection in large-scale SAR images via spatial shuffle-group enhance attention. IEEE Transactions on Geoscience and Remote Sensing, 59(1), 379–391. https://doi.org/10.1109/TGRS.2020.2997203
[9] Fu, J., Sun, X., Wang, Z., & Fu, K. (2020). An anchor-free method based on feature balancing and refinement network for multiscale ship detection in SAR images. IEEE Transactions on Geoscience and Remote Sensing, 59(2), 1331–1344. https://doi.org/10.1109/TGRS.2020.3005151
[10] Liu, S., Kong, W., Chen, X., Xu, M., Yasir, M., Zhao, L., & Li, J. (2021). Multi-scale ship detection algorithm based on a lightweight neural network for spaceborne SAR images. Remote Sensing, 14(1), 114. https://doi.org/10.3390/rs14010114
[11] Zhao, Y., Zhao, L., Li, C., & Kuang, G. (2021). Pyramid attention dilated network for aircraft detection in SAR images. IEEE Geoscience and Remote Sensing Letters, 18(4), 662–666. https://doi.org/10.1109/LGRS.2020.2981255
[12] Yu, J., Wu, T., Zhou, S., Pan, H., Zhang, X., & Zhang, W. (2022). An SAR ship object detection algorithm based on feature information efficient representation network. Remote Sensing, 14(14), 3489. https://doi.org/10.3390/rs14143489
[13] Zhang, X., Huo, C., Xu, N., Jiang, H., Cao, Y., Ni, L., & Pan, C. (2020). Multitask learning for ship detection from synthetic aperture radar images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 14, 8048–8062. https://doi.org/10.1109/JSTARS.2020.3045971
[14] Zhou, L., Zhang, H., Zhao, W., & Wang, D. (2021). A lightweight network for ship detection in SAR images. Remote Sensing, 13(19), 3828. https://doi.org/10.3390/rs13193828
[15] Chen, P., Li, Y., Zhou, H., Liu, B., & Liu, P. (2022). Detection of small ship objects using anchor-free network in SAR images. Journal of Marine Science and Engineering, 10(12), 1965. https://doi.org/10.3390/jmse10121965
[16] Wang, J., Lu, C., Jiang, W., & Zhang, L. (2021). An improved YOLOv5 model for ship detection in SAR images. Remote Sensing, 13(22), 4736. https://doi.org/10.3390/rs13224736
[17] Miao, T., Zeng, H., Yang, W., Chu, B., Zou, F., & Ren, W. (2021). BiFA-YOLO: A novel YOLO-based method for arbitrary-oriented ship detection in high-resolution SAR images. Remote Sensing, 13(21), 4209. https://doi.org/10.3390/rs13214209
[18] Fan, Q., Chen, F., Cheng, M., Lou, S., Xiao, R., Zhang, B., Wang, C., & Li, J. (2022). Ship detection using a fully convolutional network with compact polarimetric SAR images. Remote Sensing, 14(3), 635. https://doi.org/10.3390/rs14030635
[19] Xu, Z., Guo, Y., & Saleh, J. H. (2021). Efficient hybrid Bayesian optimization algorithm with adaptive expected improvement acquisition function. Engineering Optimization, 53(10), 1786–1804. https://doi.org/10.1080/0305215X.2020.1826467
[20] Jeong, J., & Shin, H. (2021). Bayesian optimization for a multiple-component system with target values. Computers & Industrial Engineering, 157, 107310. https://doi.org/10.1016/j.cie.2021.107310

[21] Bloch, L., & Friedrich, C. M. (2021). Using Bayesian optimization to effectively tune random forest and XGBoost hyperparameters for early Alzheimer's disease diagnosis. Wireless Mobile Communication and Healthcare, 1192, 285–298. https://doi.org/10.1007/978-3-030-70569-5_18

[22] Yang, L., Shami, A. (2020). On hyperparameter optimization of machine learning algorithms: Theory and practice. Neurocomputing, 415, 295–316. https://doi.org/10.1016/j.neucom.2020.07.061

[23] Bischl, B., Binder, M., Lang, M., Pielok, T., Richter, J., Coors, S., Thomas, J., Ullmann, T., Becker, M., Boulesteix, A.-L., Deng, D., & Lindauer, M. (2021). Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. WIREs Data Mining and Knowledge Discovery, 13(2), e1484. https://doi.org/10.1002/widm.1484

[24] Eriksson, D., Pearce, M., Gardner, J., Turner, R. D., & Poloczek, M. (2020). Scalable global optimization via local Bayesian optimization. Advances in Neural Information Processing Systems, 32, 5496–5507. https://doi.org/10.5555/3454287.3454781

[25] Torrisi, S., Alessandri, A., & Viti, F. (2020). A Bayesian optimization approach for the calibration of traffic simulation models. Transportation Research Part C, 113, 38–57. https://doi.org/10.1016/j.trc.2019.12.007

[26] Liu, H., Ong, Y.-S., Shen, X., & Cai, J. (2020). When Gaussian process meets big data: A review of scalable GPs. IEEE Transactions on Neural Networks and Learning Systems, 31(11), 4405–4423. https://doi.org/10.1109/TNNLS.2019.2957109

[27] Snoek, J., Larochelle, H., & Adams, R. P. (2020). Practical Bayesian optimization of machine learning algorithms. Communications of the ACM, 63(3), 86–93. https://doi.org/10.1145/3399579

[28] Wang, X., Tan, K., Du, Q., Chen, Y., & Du, P. (2020). Caps-TripleGAN: GAN-assisted CapsNet for hyperspectral image classification. IEEE Transactions on Geoscience and Remote Sensing, 57(9), 7232–7245. https://doi.org/10.1109/TGRS.2019.2910890

[29] Zhang, Y., Liu, T., Zhao, Z., Zhang, X., & Li, D. (2022). SAR ship detection with feature enhancement and dynamic label assignment. Remote Sensing, 14(16), 3979. https://doi.org/10.3390/rs14163979

[30] Li, J., Qu, C., & Shao, J. (2020). Ship detection in SAR images based on an improved Faster R-CNN. Applied Sciences, 10(21), 7793. https://doi.org/10.3390/app10217793

[31] Ma, M., Bai, C., Zhang, S., Qian, L., & Yan, H. (2022). KECANet: A novel convolutional kernel network for ocean SAR scene classification with limited labeled data. Frontiers in Marine Science, 9, 935600. https://doi.org/10.3389/fmars.2022.935600

[32] Deng, Z.-R., Zhang, X.-Y., Gao, W.-M., Zhang, H.-P., & Xu, Y.-Q. (2022). A method of multi-channel SAR moving target detection and imaging based on complex CNN. Journal of Physics: Conference Series, 2209(1), 012021. https://doi.org/10.1088/1742-6596/2209/1/012021