

# XPRESS: X-band Radar Place Recognition via Elliptical Scan Shaping

Hyesu Jang<sup>1</sup>, Wooseong Yang<sup>2</sup>, Ayoung Kim<sup>2</sup>, Dongje Lee<sup>1</sup>, and Hanguen Kim<sup>1\*</sup>

**Abstract**—X-band radar serves as the primary sensor on maritime vessels, however, its application in autonomous navigation has been limited due to low sensor resolution and insufficient information content. To enable X-band radar-only autonomous navigation in maritime environments, this paper proposes a place recognition algorithm specifically tailored for X-band radar, incorporating an object density-based rule for efficient candidate selection and intentional degradation of radar detections to achieve robust retrieval performance. The proposed algorithm was evaluated on both public maritime radar datasets and our own collected dataset, and its performance was compared against state-of-the-art radar place recognition methods. An ablation study was conducted to assess the algorithm’s performance sensitivity with respect to key parameters.

**Index Terms**—Marine Robotics, Range Sensing, Localization, SLAM

## I. INTRODUCTION

EARLY maritime autopilot systems were primarily designed for open-sea navigation, where the sparse and relatively unstructured environment allowed for sufficient autonomy despite intermittent sensor noise and signal fluctuations. In contrast, near-shore, coastal, and port environments present more complex operational challenges. As demonstrated by Han et al. [1] and Jang et al. [2], global positioning system (GPS) signals in maritime environments are frequently subject to degradation and interference, complicating real-time decision-making in safety-critical scenarios. Additionally, these environments are characterized by high traffic density and dynamic obstacles, which complicate situational awareness and hinder robust localization due to the frequent occlusion and unpredictability of surrounding agents. Furthermore, geographic features shift over time under the influence of tidal effects and constructions. These challenges render the estimation of vessel location based solely on fixed Electronic Navigational Chart (ENC) or satellite images unreliable and necessitate the incorporation of real-time place recognition (PR) with perception to account for dynamic environmental changes. Previous studies [3] have utilized camera and Light Detection and Ranging (LiDAR) sensors to perceive complex near-shore environments and enable autonomous sailing. However, their applicability in broader maritime scenarios is

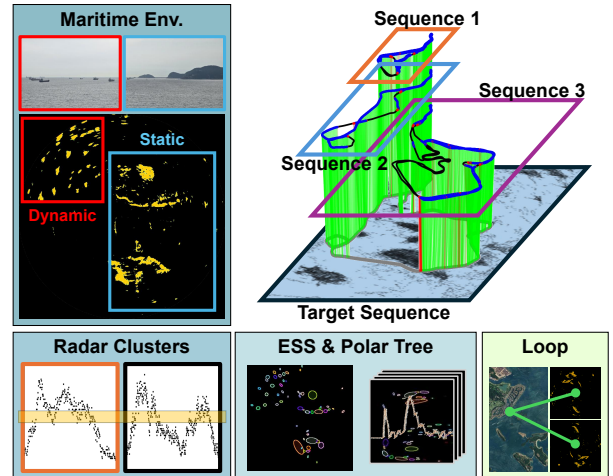


Fig. 1: XPRESS proposes a fast and robust PR method for maritime environments using X-band radar. By leveraging radar cluster information and applying intentional elliptical degradation, the proposed approach minimizes the influence of dynamic elements and fluctuations of static objects.

limited due to their short sensing range, which constrains path planning and vessel control. Moreover, these sensors are vulnerable to adverse weather conditions such as sea fog and squalls, and maintaining hardware functionality over the long term is challenging due to exposure to harsh environments and salt-induced corrosion. Due to these complex factors, radar sensors are predominantly utilized in maritime environments, offering long-range detection capabilities and robust performance under adverse weather conditions. Notably, most existing vessels and yacht platforms are already equipped with X-band radar systems, the importance of developing algorithms tailored to the X-band radar is emphasized for the near-term deployment of autonomous vessel navigation.

The primary challenge lies in that X-band radar offers limited information while generating long-range image data, necessitating the extraction of salient details from low-resolution, large-scale inputs. Furthermore, X-band radar possesses distinctive operational characteristics. Due to its low range resolution, lateral perturbations of the vessel have a negligible effect on the radar returns. However, the system is particularly susceptible to rotational drift, a consequence of its low update rate and long-range detection configuration. Vessel rotation severely distorts the radar imagery, complicating the robust perception. Therefore, algorithmic frameworks designed for X-band radar must prioritize and consider core data compression, rotational robustness, and computational efficiency.

We propose a PR algorithm tailored for X-band radar, which

Manuscript received: July 1, 2025; Revised: September 10, 2025; Accepted: October 20, 2025. This paper was recommended for publication by Editor Giuseppe Loianno upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by the KIMST (07582969), Korea.

<sup>1</sup>H. Jang, D. Lee, and H. Kim are with Seadronix Corp., Seoul, S. Korea [hsjang, dj.lee, hank05]@seadronix.com

<sup>2</sup>W. Yang and A. Kim are with the Department of Mechanical Engineering, SNU, Seoul, S. Korea [yellowish, ayoungk]@snu.ac.kr




Fig. 2: The concept of the proposed algorithm lies in the intentional degradation of radar data. Although static objects are detected, their appearance fluctuates across frames. To address this, radar data is first clustered using connected component labeling. Each labeled component is then approximated by an ellipse. The elliptical data is transformed into polar coordinates, from which a polar histogram is generated to enable KD-Tree-based retrieval. During descriptor matching, candidates are initially filtered based on the number of clusters, followed by estimation of revisitiness among the selected candidates.

resolves challenges due to the inherent characteristics of the sensor data. The main difficulties addressed by the algorithm include efficient extraction of salient information, reduction of retrieval latency, and achievement of rotational invariance. The specific contributions are summarized as follows:

- The first X-band radar PR for maritime application**  
 X-band radar captures abstract and sparse information, limiting to directly extract relevant features. To address this, we reinterpret the limitation as an opportunity to simplify the input representation, enabling the identification of common structural patterns across different observations. Based on the simplification, we design a descriptor specifically suited for PR using X-band radar.
- Fast, rotational-invariant, and fully-metric algorithm design for the huge and ambiguous x-band radar data**  
 To extract candidates from the descriptor dictionary, we propose a two-step retrieval process. Since all radar frames are represented by clustered structures, we first estimate the number of clusters corresponding to each frame. Rather than directly comparing descriptor vectors, we initially apply an integer-based thresholding using the cluster count, reducing the search space. Subsequently, vector-wise matching is performed to identify the most appropriate candidate.
- Robustness evaluation by structuring the ablation studies and own data acquisition**  
 As the proposed work presents the first X-band radar-based PR algorithm, thorough verification is essential. By benchmarking against existing state-of-the-art PR algorithms originally developed for other scanning sensors, we demonstrate that our approach is specifically tailored for X-band radar and achieves robust performance in maritime environments.

## II. RELATED WORK

### A. Place Recognition with Imaging Radar

As noted in the radar survey by Harlow et al. [4], radars can be broadly categorized into imaging radars and automotive radars. X-band radar belongs to the imaging radar category, similar to W-band radar. The primary distinction lies in their operating bandwidths, with one utilizing the X-band and the other the W-band. Due to the lack of PR method dedicated to X-band radar, we briefly summarize the previous PR work on W-band radar. The early works on W-band radar have directly adopted the existing LiDAR-based descriptor into radar imagery. Kim et al. [5] exploited the Scan Context [6]

representation by Radar Cross Section (RCS) values for place retrieval. Building on this, Adolfsson et al. [7] integrated a Scan Context variant with the loop-verification module into TBV Radar SLAM. Hong et al. [8] converted the radar image into point clouds and employed the M2DP descriptor for robust localization across diverse conditions. Jang et al. [9] utilized the Radon transform and the fast Fourier transform (FFT) to create roto-translationally invariant descriptors, RaPlace. Recent approaches have started to consider radar-specific characteristics. Gadd and Newman [10] combined feature aggregation with FFT on radar imagery, Kim et al. [11] leveraged the distribution of unoccupied pixels in radar images as an informative descriptor, named Referee. Learning-based methods have also been proposed for PR with W-band radar. Barnes and Posner [12] proposed self-supervised keypoint detection for global descriptors, and Săftescu et al. [13] embedded radar imaginaries into a metric feature space for similarity comparisons. Yuan et al. [14] introduced uncertainty-aware metric learning to enhance robustness against environmental variations. However, these learning-based approaches exhibit limited generalizability across varying radar noise characteristics in diverse environments. To tackle this issue, some approaches exploited consistent quality LiDAR-derived maps to supervise radar embeddings [15, 16].

As noted by Jang et al. [2], X-band radar has distinct data characteristics against W-band radar, prompting the necessity for specialized PR methodology. We incorporate the unique signal and data properties of the X-band radar to build novel descriptors, enabling reliable open-sea localization.

### B. Maritime SLAM and Perception

The advancements in maritime navigation have largely focused on improving localization accuracy using marine radar, particularly X-band radar, for odometry purposes. For example, Bibby and Reid [17] proposed a hybrid simultaneous localization and mapping (SLAM) framework that integrates both spatial and temporal data, enabling robust localization in dynamic marine environments. Han et al. [1] introduced Coastal SLAM, which combines marine radar with coastal area modeling for GPS-restricted conditions. Schiller et al. [18] improved marine radar odometry by modeling radar resolution and incorporating temporal information. Jang et al. [19] introduced LodeStar, a novel method specifically designed for X-band radar odometry. However, there is a lack of PR methods using marine radar that can support the operation of a full SLAM system. Recent maritime SLAM attempted

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 3: In coastal areas, while the density of anchored vessels remains relatively stable, their individual positions vary consistently. Considering the vessel density within specific regions, rather than relying on exact vessel positions, is advantageous for handling dynamic elements in X-band radar images.

to exploit multi-sensor fusion. Some efforts have integrated LiDAR to enhance berthing state estimation and assistance systems [20, 21], thereby improving situational awareness during complex maneuvers. Furthermore, Sawada and Hirata [3] demonstrated LiDAR-based SLAM for coastal mapping and localization in GNSS-denied environments. Lu *et al.* [22] have further improved berthing navigation by fusing inertial measurement unit (IMU) and RTK data. However, these solutions are primarily applicable to coastal and near-shore environments, where LiDAR can be effective. In contrast, using LiDAR in open-sea navigation proves to be impractical, as highlighted by Jang *et al.* [2]. Therefore, we focus on exploring PR methods that leverage the distinctive features of X-band radar to address the gaps in maritime SLAM systems in wide oceanic areas.

### III. X-BAND RADAR PLACE RECOGNITION

X-band radar provides abstract representations of the surrounding environment, demonstrating challenges in navigation algorithms. The most challenging phenomenon is the inconsistency of detection details, even when the vessel remains in the same position. Furthermore, it is necessary to account for large and clustered dynamic objects anchored in designated zones, potentially appear as dynamic islands in radar imagery. The overall framework of the proposed algorithm consists of two stages, where the first involves generating an X-band radar descriptor that suppresses the effects of fluctuations and clustered dynamic objects, and the second performs efficient descriptor matching with predefined descriptor data. The overview of our system is depicted in Fig. 2

#### A. Marine Descriptor

1) *Connected Component Labeling (CCL)*: LiDAR and W-band radar can capture detailed structural information from the surrounding environment, enabling reliable estimation of surface planes or linear features. In contrast, X-band radar lacks the resolution necessary to visualize precise shapes of detected objects. Therefore, we treat radar returns as aggregated clusters of reflections. To efficiently segment these clusters, we employ CCL [23]. The nearby pixels  $N(p)$  for the individual pixel  $p$  are defined as

$$N(p) = \{(x \pm 1, y), (x, y \pm 1), (x \pm 1, y \pm 1)\} \quad (1)$$

$$N_{half}(p) = \{(x \pm 1, y - 1), (x, y - 1), (x - 1, y)\}.$$

CCL is performed in two phases. In the first phase, the algorithm scans the radar image in a row-wise manner and

assigns a label  $l$  to each non-zero pixel. For each pixel  $p$ , if all neighboring pixels in the half-neighborhood  $N_{half}(p)$  are zero, a new label is assigned to  $p$ . If there is exactly one non-zero neighbor,  $p$  inherits the same label as that neighbor. If multiple non-zero neighbors exist with different labels,  $p$  is assigned one of those labels arbitrarily, and the differing label pairs are stored as equivalence pairs, indicating that they belong to the same connected region.

In the second phase, once the initial labeling is complete, the equivalence pair information is used to rearrange and unify the labels. Each group of equivalent labels is replaced with a unique representative label, ensuring that all pixels within the same connected component are assigned the same label. As a result, distinct clusters in the radar image are identified with unique label values.

2) *Elliptical Scan Shaping*: Exploiting the segmented clusters obtained from CCL, we derive elliptical scan shapes to simplify the representation of detected regions. Given the frame-to-frame fluctuations in radar detection, approximating each cluster with a parametric elliptical model improves the robustness of the descriptor by providing a consistent and compact geometric abstraction. The ellipse is defined by five parameters: the centroid  $(x_c, y_c)$ , the lengths of the major and minor axes  $a$  and  $b$ , and the orientation  $\theta$  of the major axis with respect to the reference coordinate frame. We compute the centroid  $(x_c^l, y_c^l)$  of cluster  $l$  as

$$(x_c^l, y_c^l) = \left( \frac{1}{N_l} \sum_{i=1}^{N_l} x_i^l, \frac{1}{N_l} \sum_{i=1}^{N_l} y_i^l \right) \quad (2)$$

where  $(x_i^l, y_i^l)$  are the coordinates of the  $i$ -th pixel within a cluster  $l$ , and  $N_l$  is the total number of pixels in cluster  $l$ . To determine the best-fit ellipse for each cluster, we compute the covariance matrix of the pixel coordinates within the cluster as

$$C^l = \begin{bmatrix} C_{11}^l & C_{12}^l \\ C_{21}^l & C_{22}^l \end{bmatrix}$$

$$= \frac{1}{N_l} \begin{bmatrix} \sum (x_i^l - x_c^l)^2 & \sum (x_i^l - x_c^l)(y_i^l - y_c^l) \\ \sum (x_i^l - x_c^l)(y_i^l - y_c^l) & \sum (y_i^l - y_c^l)^2 \end{bmatrix} \quad (3)$$

By performing eigen decomposition on the covariance matrix  $C^l$ , we obtain the eigenvalues  $\lambda_{1,2}$  with

$$\lambda_{1,2} = \frac{1}{2} \left( (C_{11}^l + C_{22}^l) \pm \sqrt{(C_{11}^l - C_{22}^l)^2 + 4(C_{12}^l)^2} \right). \quad (4)$$

Assuming a Gaussian pixel distribution, the lengths of the major and minor axes of the ellipse are defined as  $a = 2\sqrt{\lambda_1}$  and  $b = 2\sqrt{\lambda_2}$ , respectively. To estimate the orientation of the ellipse, we use the components of the covariance matrix,

$$\theta^l = \frac{1}{2} \arctan \left( \frac{2C_{12}^l}{C_{11}^l - C_{22}^l} \right). \quad (5)$$

We utilize the elliptical scan-shaped image  $I_{cart}$  as initial x-band radar descriptor, which suppresses irregular fluctuations and enhances the representation of static regions. For all clusters, elliptical scan-shaped images are defined as

$$x^l(t) = x_c^l + a^l \cos(t) \cos(\theta^l) - b^l \sin(t) \sin(\theta^l)$$

$$y^l(t) = y_c^l + a^l \cos(t) \sin(\theta^l) + b^l \sin(t) \cos(\theta^l), \quad (6)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

$$I_{cart} = \bigcup_{l=1}^{l_F} \left\{ (x^l(t), y^l(t)) \mid t \in [0, 2\pi] \right\}. \quad (7)$$

### B. Maritime Place Recognition

The minimum detection range of X-band radar in the MOANA dataset is about  $75\text{ m}$  and the range resolution is about  $3.25\text{ m}/px$ , making the algorithm unsuitable for focusing on lateral movement and invariance. Effective PR algorithms require both fast retrieval and sensor-specific descriptor. Therefore, our algorithm emphasizes the descriptor including rotational invariance while maintaining high retrieval speed.

1) *Cluster Searching*: To minimize the time complexity of the search algorithm, we propose a fast and reliable initial candidate extraction. Maritime scenes captured by X-band radar consist of both static and dynamic objects. Although dynamic objects are fewer in number, their physical size is often substantial which can be misinterpreted as small static islands. This characteristic limits the applicability of conventional feature-based or direct matching techniques for robust place retrieval. Furthermore, as illustrated in Fig. 3, numerous vessels are anchored in anchorage, which subsequently act as dynamic objects. For long-time-interval PR, it is essential to recognize the presence of anchoring zones, while accounting for the characteristic that the positions of individual vessels within the anchoring zones differ from the previous data acquisition. To address this, we treat dynamic objects as individual clusters and consider only the number of clusters, irrespective of their size. This resolves the challenges from both large-scale dynamic objects and anchoring zones, where the configuration of vessels varies over time, yet the overall vessel density remains relatively consistent.

A frame  $F$  is included in the candidate set  $F_{\text{cand}}$  if and only if the difference in cluster count between  $F$  and the query frame  $Q$  satisfies the threshold condition,

$$F \in F_{\text{cand}} \iff |l_F - l_Q| \leq p \quad (8)$$

$l_F$  and  $l_Q$  denote the number of clusters in the candidate frame and the query frame, respectively, and  $p$  is the pre-defined threshold. Compared to vector-to-vector comparisons, thresholding based on a single integer accelerates the candidate selection process.

2) *Polar Ellipse Matching*: After extracting initial candidates, we perform a rotationally invariant search by converting the radar data into polar coordinates. With the central point  $(x_0, y_0)$  of the Cartesian image, the polar coordinate transformation is performed by first translating the points as  $\tilde{x}(t) = x(t) - x_0$ ,  $\tilde{y}(t) = y(t) - y_0$ , and then rearranging them into range-azimuth coordinates, where

$$r(t), \phi(t) = \sqrt{\tilde{x}(t)^2 + \tilde{y}(t)^2}, \arctan\left(\frac{\tilde{y}(t)}{\tilde{x}(t)}\right) \quad (9)$$

$$I_{\text{polar}} = \bigcup_{l=1}^{l_F} \left\{ (r^l(t), \phi^l(t)) \mid t \in [0, 2\pi] \right\} \quad (10)$$

We integrate pixel intensities along each angular direction to generate a range-wise histogram as

$$V_{\text{ESS}} = \sum_{\phi} I_{\text{polar}} \quad (11)$$

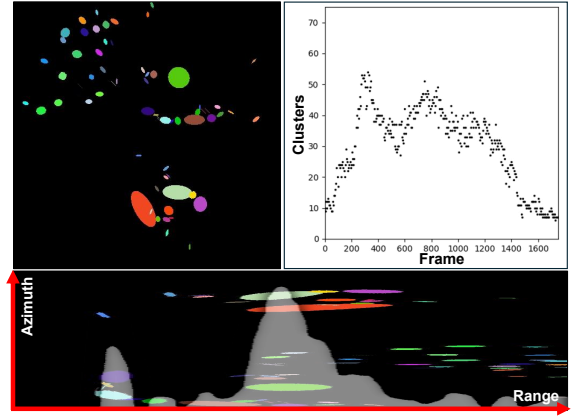


Fig. 4: X-band radar data exhibits fluctuations, the same objects across consecutive frames appear with varying cluster boundaries. To mitigate the impact of this variability, the radar data is degraded into elliptical representations (top-left). The number of resulting clusters (top-right) serves as a criterion for initial candidate selection, and the histogram derived from the polar-transformed data is used as a descriptor for location retrieval (bottom).

that incorporates rotational invariance. This integration also mitigates uncertainty arising from anchoring zones, where dynamic vessels are randomly distributed. Exploiting a KD-tree structure, we efficiently retrieve the nearest neighbor vector among the candidates, and verify the final match by evaluating the cosine similarity distance.

## IV. EXPERIMENT

### A. Maritime Dataset Configuration

To validate the robustness of our algorithm, we conducted experiments on two major public maritime datasets as well as our own collected dataset. The MOANA dataset [2] contains diverse routes with multiple loops, suitable for evaluating both intra-session loop closure and inter-session PR performance. In contrast, the Pohang Canal Dataset [24] does not include intra-session loops but provides repeated traversals of the same routes across different sequences, which is appropriate for assessing inter-session PR capabilities. Additionally, to demonstrate the generalized performance across different datasets, we collected a new dataset covering routes adjacent to the MOANA dataset, enabling the robustness evaluation of inter-dataset PR. The maximum detection range of the radar in our dataset is  $3328\text{ m}$ , with a identical pixel resolution to the MOANA. The detailed experimental environment settings are summarized in Fig. 5.

1) *MOANA dataset*: The MOANA dataset consists of seven sequences designed to evaluate maritime navigation algorithms. Among these, the Island sequences are particularly suited for PR evaluation. We use three Island sequences for intra-session PR and the Island Expedition sequence for inter-session PR experiments. The Single Island sequence provides clear radar data, enabling the evaluation of short-term revisit-ness detection. The Twins Island sequence offers noisier radar data, challenging the robustness of PR. The Complex Island sequence combines both clear and noisy radar data, with a large number of loop candidates due to its longer navigation route. The Island Expedition sequence encompasses all locations

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 5: The left image illustrates our sensor set and the rights are the routes for the evaluation. To evaluate the proposed algorithm, we utilized two public datasets and additionally acquired our own dataset in the same geographic region as the MOANA dataset. Our dataset is divided into two sequences, Coastal and Complex, which are used to assess intra-session and long-term inter-session PR, respectively.

from the Single, Twins, and Complex Island sequences. Thus, we present inter-session PR results across three pairs: Single-Expedition, Twins-Expedition, and Complex-Expedition.

2) *Pohang Canal dataset*: Since the Pohang Canal dataset consists of a single route without loops, we evaluated inter-session recognition performance. The Narrow Canal area represents the most challenging environment for X-band radar, characterized by significant multipath noise and frequent detection failures. The Inner Port area, being wider than the Narrow Canal, is suitable for evaluating tunnel-like PR scenarios where similar features appear repeatedly. The Outer Port area, although relatively short, connects the narrow region to the broader coastal area. The Near-Coastal area provides stable radar operation conditions. We validated our algorithm across these four distinct radar environments at once.

3) *Own dataset*: Our dataset was collected in the same maritime environment as the MOANA dataset, sharing common areas at several locations. Using this dataset, we evaluate both intra-session and long-time-interval inter-session PR by matching against the Island Expedition sequence from the MOANA dataset. This experiment verifies the robustness of the proposed algorithm when applied to independently collected X-band radar data.

## B. Evaluation Criteria

PR in the context of SLAM differs fundamentally from general image retrieval tasks. In SLAM, the objective is to identify the single most probable frame from the previous data to enable loop closure or designate a keyframe. Additionally, the system performance of processing large-scale data must be considered to ensure real-time operation. Taking these requirements into account, we adopt the following criteria for evaluating the main results.

1) *PR Methods*: As introduced in the previous chapter, four recent W-band radar-based PR algorithms are available for comparison[5][9][10][11]. However, we excluded RaPlace from our performance evaluation due to its high computational complexity, which does not align with our objective of assessing the efficiency and lightness of the algorithm. All algo-

gorithms were reproduced using the W-band radar data to verify their integrity, and the results confirmed that performance was not degraded on the W-band radar dataset.

2) *Precision-recall curve and F1 curve*: The precision-recall curve and the area under the curve (AUC) are standard metrics for evaluating PR accuracy. These indices measure the model’s positive predictive value ( $=TP/(TP+FP)$ ) and its ability to correctly identify true matches ( $=TP/(TP+FN)$ ). The F1 score is the harmonic mean of precision and recall, effectively balancing the two metrics to minimize bias toward either. The F1 curve is plotted against recall values, sharing the same x-axis as the precision-recall curve.

3) *False Alarm Rate and Recall@1*: The false alarm rate ( $=FP/(TN+FP)$ ) serves as a critical metric to assess whether the algorithm reliably identifies only true positive matches. In maritime environments where navigation occurs over expansive areas without fixed lanes and loop closures are rare, false alarms can severely impact system reliability, even if precision is high. Recall@1 ( $=TP/GT$ ) is not a fully appropriate metric for SLAM applications, as it assumes the presence of at least one correct match among the candidates, whereas SLAM systems may encounter cases where no true correspondence exists. The false alarm rate necessitates a reasonably performing classifier that guarantees the true positive rate, and recall@1 is an unsafe index since the false alarms are not considered. Both the false alarm rate and Recall@1 require method-specific optimal thresholds to be effectively utilized as evaluation metrics. Therefore, these two factors were jointly considered to ensure a fair and reliable assessment of model performance. For intra-session PR, we employ the area under the Receiver Operator Characteristic (ROC) curve to jointly evaluate the false alarm rate and recall. For inter-session PR, where correct matches are known to exist among the candidates, we adopt Recall@1 as the evaluation metric.

## C. Intra-session Place Recognition

For efficient and generalizable evaluation, the algorithm was assessed at every five frame interval, with a revisit criterion defined as a 100 m range, which corresponds to approximately three percent of the radar maximum detection range. All experiments were performed on a system equipped with an Intel® Core™ Ultra 7 165H CPU and 32 GB of RAM.

As shown in Table I and Fig. 6, our algorithm consistently outperforms across all sequences. Notably, in the Twins Island sequence, Referee[11] exhibits comparatively poor performance. This can be attributed to the characteristics of X-band radar in open ocean settings. While Referee is capable of generating robust descriptors in structured and complex environments by leveraging free space regions in the image, its discriminative capability degrades when free space becomes dominant. In this sequence, Referee produces a high number of false alarms in areas lacking distinctive detections.

RadVLAD and FFT-RadVLAD are effective algorithms for W-band radar; however, they exhibited performance degradation when applied to X-band radar. The primary suspected reason is the limited information in X-band radar images, where the simpler pixel composition makes it difficult for the VLAD descriptor to distinguish between different places.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Sequence		Method	Time(ms)	AUC	F1	ROC
MOANA	Single Island	RingKey [6]	0.08	0.89	0.82	0.92
		Referee [11]	0.07	0.73	0.85	0.85
		RadVLAD [10]	9.23	0.59	0.75	0.49
		FFT-RadVLAD [10]	9.46	0.12	0.18	0.38
		Ours	<b>0.04</b>	<b>0.95</b>	<b>0.90</b>	<b>0.96</b>
	Twins Island	RingKey	0.09	0.63	0.56	0.83
		Referee	0.08	0.15	0.16	0.65
		RadVLAD	11.27	0.10	0.27	0.25
		FFT-RadVLAD	11.45	0.06	0.16	0.20
		Ours	<b>0.03</b>	<b>0.84</b>	<b>0.74</b>	<b>0.92</b>
	Complex Island	RingKey	0.31	0.78	0.72	0.80
		Referee	0.21	0.67	0.69	0.72
RadVLAD		31.85	0.07	0.19	0.34	
FFT-RadVLAD		33.13	0.02	0.06	0.17	
Ours		<b>0.09</b>	<b>0.88</b>	<b>0.79</b>	<b>0.88</b>	
Own Coastal	RingKey	0.23	0.74	0.67	0.91	
	Referee	0.21	0.51	0.46	0.74	
	RadVLAD	15.44	0.08	0.18	0.32	
	FFT-RadVLAD	17.37	0.05	0.13	0.28	
	Ours	<b>0.12</b>	<b>0.77</b>	<b>0.68</b>	<b>0.89</b>	

TABLE I: Intra-session PR (Bold: Best, Underline: Runner-up)

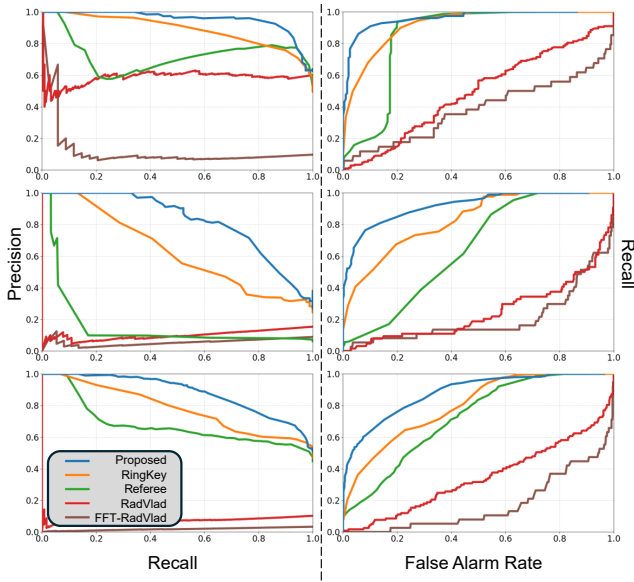


Fig. 6: Precision-recall and ROC curves for intra-session PR on the MOANA dataset (from top to bottom: Single, Twins, Complex) demonstrate that the proposed algorithm achieves the most robust performance among state-of-the-art radar PR methods.

#### D. Inter-session place recognition

For maritime vessels, global localization is a critical component, as GPS signals are frequently subject to interference and vessels often initiate operations from unknown locations. To ensure reliable initialization and enable loop closure for SLAM systems, we evaluate our algorithm with inter-session PR performance. Similar to the results observed in intra-session PR, the RadVLAD and FFT-RadVLAD descriptors were ineffective when applied to X-band radar data. In contrast, the other state-of-the-art algorithms demonstrated reliable performance. Notably, the proposed methods exhibited robustness across all sequences without performance degradation. Detailed results are depicted in Table II and Fig. 8.

1) *Pohang Canal Dataset*: The Pohang Canal Dataset consists of a structured, static environment with minimal dynamic elements, making it challenging to assess the failure case of the algorithms. As reflected in the results, all algorithms demonstrate strong performance across the sequences except for the VLAD descriptors. Furthermore, due to the identical

Target	Query	Method	AUC	F1	ROC	Recall@1
MOANA	Single Island	RingKey	0.87	0.82	0.69	0.70
		Referee	0.91	<b>0.88</b>	<b>0.81</b>	0.70
		RadVLAD	0.08	0.22	0.27	0.13
		FFT-RadVLAD	0.06	0.14	0.37	0.15
		Ours	<b>0.92</b>	<b>0.85</b>	<b>0.76</b>	<b>0.74</b>
	Twins Island	RingKey	0.86	0.87	0.61	0.77
		Referee	0.57	0.53	0.31	0.36
		RadVLAD	0.15	0.32	0.37	0.15
		FFT-RadVLAD	0.05	0.16	0.13	0.11
		Ours	<b>0.95</b>	<b>0.90</b>	<b>0.77</b>	<b>0.81</b>
	Complex Island	RingKey	0.73	0.66	0.72	0.44
		Referee	0.69	<b>0.70</b>	0.75	0.42
		RadVLAD	0.03	0.10	0.20	-
		FFT-RadVLAD	0.02	0.06	0.14	0.06
		Ours	<b>0.80</b>	<b>0.69</b>	<b>0.78</b>	<b>0.45</b>
Own Complex Island	RingKey	0.69	0.77	0.44	0.62	
	Referee	0.73	0.70	<b>0.66</b>	0.48	
	RadVLAD	0.09	0.21	0.38	0.07	
	FFT-RadVLAD	0.06	0.17	0.17	0.09	
	Ours	<b>0.75</b>	<b>0.79</b>	<b>0.53</b>	<b>0.65</b>	
MOANA Own Complex Island	RingKey	0.87	0.84	0.70	<b>0.70</b>	
	Referee	0.77	0.73	0.65	0.53	
	RadVLAD	0.11	0.24	0.38	0.03	
	FFT-RadVLAD	0.05	0.17	0.01	0.03	
	Ours	<b>0.91</b>	<b>0.87</b>	<b>0.82</b>	<b>0.69</b>	
Pohang Canal 04	RingKey	0.94	0.98	0.04	0.96	
	Referee	0.97	<b>0.99</b>	0.08	<b>0.98</b>	
	RadVLAD	0.03	0.10	<b>0.56</b>	0.02	
	FFT-RadVLAD	0.02	0.05	0.37	0.02	
	Ours	<b>0.98</b>	<b>0.98</b>	0.14	0.96	
Pohang Canal 05	RingKey	0.89	<b>0.92</b>	<b>0.61</b>	<b>0.85</b>	
	Referee	0.89	0.90	0.58	0.82	
	RadVLAD	0.02	0.08	0.33	0.02	
	FFT-RadVLAD	0.02	0.08	0.58	0.02	
	Ours	<b>0.91</b>	<b>0.91</b>	0.47	0.84	

TABLE II: Inter-session PR (Bold: Best, Underline: Runner-up)




Fig. 7: Computation time for descriptor retrieval is reduced in the proposed method by constraining the search space to candidates with similar cluster quantities.

route, the evaluation of the ROC curve is ineffective owing to the insufficient number of true negative instances.

2) *MOANA Dataset*: The results for inter-session PR on the MOANA dataset are presented in the top side of Fig. 8. Compared to Referee, the proposed method demonstrates greater robustness on the Twins Island route, where ambiguous data is prevalent. Referee exhibits stable performance only on the simpler and more distinct dataset, Single Island. In comparison to RingKey, the proposed method achieves improved performance across all sequences. For all algorithms, several common false negative cases were observed. This issue arises from the revisit distance criterion, where the 100m threshold is too large, causing post-corner poses that cannot be detected by the sensor to be included in the ground truth pool. To enable a more accurate evaluation of the algorithms, a manually curated or refined ground truth set is required, ensuring that both sensor observations share a common detectable area and fall within the defined distance threshold.

3) *Own Dataset*: A simple loop covering the MOANA Complex Island route was acquired, enabling the evaluation of

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

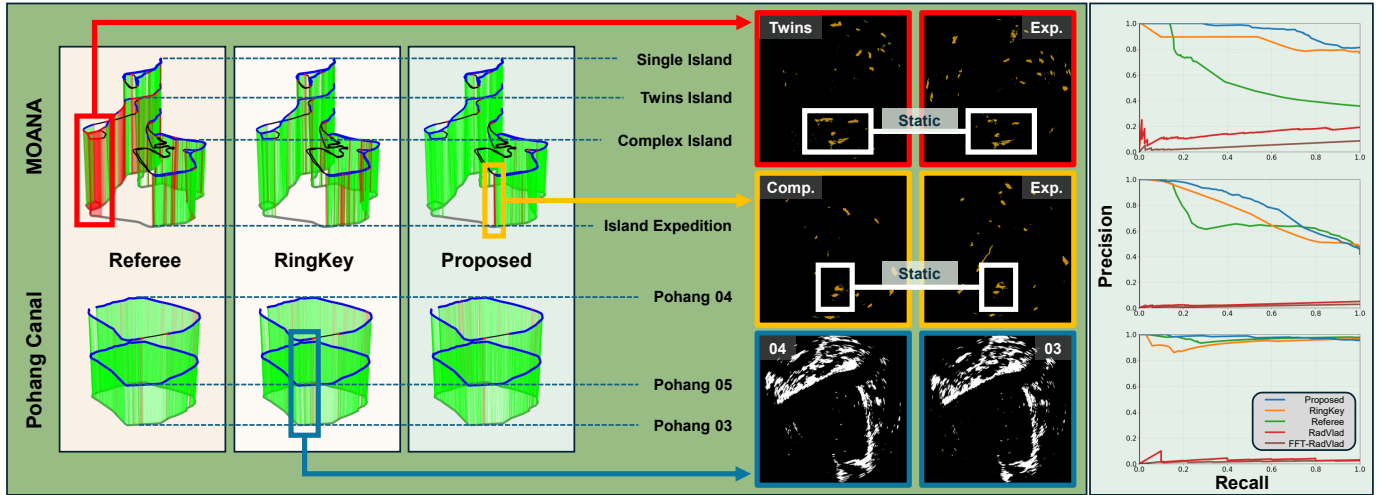


Fig. 8: For the MOANA dataset, the performance differences among the algorithms are evident, with the proposed method demonstrating the highest robustness across all sequences. In contrast, the Pohang sequence yields reliable results for all algorithms. This discrepancy can be attributed to the presence of numerous dynamic objects in the MOANA dataset, whereas the Pohang Canal dataset features predominantly static scenes. These results highlight the robustness of the proposed algorithm in environments with dynamic elements. As illustrated in the right images, only the regions marked by white boxes correspond to static objects in the radar data. Despite this challenge, the proposed method successfully retrieves places in most cases (Red). However, in the most challenging scenario (Yellow), where static information is minimal, all algorithms failed to recognize the location as the same place.

inter-session PR performance on both the MOANA Complex Island and Island Expedition sequences. Despite a time interval of approximately one year between the MOANA dataset and our data, and noticeable changes in the composition of anchored large vessels, the proposed algorithm robustly identified corresponding locations across both sequences.

### E. Computation Time

The computation time of the algorithm is primarily determined by the size of the optimal descriptor. RadVLAD and FFT-RadVLAD require large descriptor sizes, resulting in increased computational costs. In contrast, the proposed method operates effectively with limited information, thereby reducing computation time. Furthermore, as depicted in Fig. 7, applying thresholding based on cluster count minimizes unnecessary candidate searches. These characteristics make the proposed method particularly suitable for long-term navigation, which is common in maritime operations due to the typically extended duration of vessel routes.

## V. ABLATION STUDY

### A. Revisit Distance Threshold

There are no strict lanes for vessels, which results in potentially large distances between revisited locations. As shown in Fig. 9, a 10 m revisit threshold does not yield sufficient ground truth data. This indicates that loop detection in maritime navigation should account for both sensor limitations and environmental constraints. Considering both sensor and environmental conditions, an ablation study was conducted using distance thresholds of 10 m, 50 m, and 200 m. As anticipated, the 10 m threshold caused insufficient ground truth resulting in low detection accuracy across all algorithms. Increasing the threshold to 50 m produced more reliable results; however, it remained inadequate for fully covering ambiguous regions that




Fig. 9: Revisitness map corresponding to different revisit distance thresholds. Cyan-colored points indicate the locations where the vessel revisited. Due to the absence of defined lanes in maritime, maintaining identical routes is challenging. A 100 m distance threshold is considered reasonable for evaluating PR performance.

fall below the minimum detection range of the X-band radar. At a 200 m threshold, the results appeared more reasonable compared to smaller thresholds, but the inclusion of false ground truth led to a negative impact on PR accuracy. From the results of this ablation study, the lateral uncertainty of the X-band radar can be estimated based on the relationship between the distance threshold and PR performance.

### B. Descriptor Size

We employed the azimuth-wise integration of the ellipse polar image as the descriptor, where the vector size corresponds to the resolution of the range-wise information and directly influences the algorithm's computational cost. By varying the descriptor size from 10 to 100, we evaluated the performance of the algorithm. The results indicate that smaller descriptor lengths degrade performance, while beyond a certain size, further increases do not significantly affect the algorithm's effectiveness. To minimize computational complexity, catching an optimal descriptor size is recommended.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper

Criteria	Sequence	Revisit Distance (Referee / RingKey / Proposed) [m]			Descriptor Size			Cluster Threshold		
		10	50	200	10	20	100	1	10	100
AUC	Single Island	0.18 / 0.17 / 0.15	0.63 / 0.76 / 0.82	0.98 / 0.88 / 0.92	0.88	0.94	0.95	0.86	0.96	0.96
	Twins Island	- / 0.06 / 0.06	0.03 / 0.43 / 0.50	0.26 / 0.68 / 0.86	0.79	0.82	0.84	0.71	0.83	0.82
	Complex Island	0.06 / 0.06 / 0.03	0.54 / 0.62 / 0.63	0.69 / 0.86 / 0.92	0.72	0.86	0.88	0.80	0.89	0.88
F1-Score	Single Island	0.32 / 0.36 / 0.40	0.78 / 0.73 / 0.85	0.93 / 0.81 / 0.82	0.81	0.89	0.90	0.76	0.89	0.90
	Twins Island	- / 0.28 / 0.30	0.06 / 0.45 / 0.59	0.28 / 0.61 / 0.77	0.76	0.74	0.72	0.62	0.75	0.74
	Complex Island	0.12 / 0.14 / 0.13	0.59 / 0.56 / 0.59	0.70 / 0.79 / 0.82	0.67	0.77	0.78	0.71	0.79	0.78
ROC	Single Island	0.77 / 0.65 / 0.71	0.83 / 0.88 / 0.93	0.97 / 0.86 / 0.86	0.92	0.95	0.96	0.91	0.97	0.96
	Twins Island	- / 0.62 / 0.77	0.53 / 0.77 / 0.87	0.71 / 0.83 / 0.91	0.91	0.93	0.93	0.87	0.94	0.92
	Complex Island	0.67 / 0.59 / 0.57	0.75 / 0.74 / 0.78	0.70 / 0.85 / 0.90	0.77	0.87	0.88	0.82	0.88	0.88

TABLE III: Results for Ablation Study

### C. Cluster Thresholding

Cluster thresholding is applied to mitigate the influence of dynamic objects appearing in the same location. A threshold value of 1 allows only a single change relative to the base image, while a value of 100 includes nearly all images as candidates. Similar to the descriptor size results, thresholding has minimal impact at larger values but constrains performance at lower thresholds. In this context, determining the optimal threshold is necessary; however, the optimal value varies depending on the characteristics of the environment. An adaptive strategy, where the threshold is increased in regions with frequent dynamic object occurrences and decreased in more static environments, would provide an effective approach.

## VI. CONCLUSION

We proposed a place recognition method tailored for X-band radar, aiming to facilitate the initial stage of autonomous vessel navigation by perceiving the maritime environment. In contrast to conventional radar-based PR approaches, the proposed method addresses the unique characteristics of X-band radar by employing an algorithm designed to suppress the influences of dynamic elements and simplify radar detections with elliptical clusters. In future work, we plan to investigate the challenges induced by vessel perturbations and examine the impact of sensor settings on X-band radar. This will be integrated into our investigation of radar-only SLAM and multi-sensor fusion approaches for robust maritime navigation.

## REFERENCES

- [1] J. Han, Y. Cho, and J. Kim, "Coastal SLAM with marine radar for USV operation in GPS-restricted situations," *IEEE J. of Oceanic Eng.*, 2019.
- [2] H. Jang, W. Yang *et al.*, "Moana: Multi-radar dataset for maritime odometry and autonomous navigation application," *Intl. J. of Robot. Research*, 2025.
- [3] R. Sawada and K. Hirata, "Mapping and localization for autonomous ship using lidar slam on the sea," *Journal of Marine Science and Technology*, vol. 28, no. 2, pp. 410–421, 2023.
- [4] K. Harlow, H. Jang, T. D. Barfoot, A. Kim, and C. Heckman, "A new wave in robotics: Survey on recent mmwave radar applications in robotics," *IEEE Trans. Robot.*, 2024.
- [5] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "Mulran: Multimodal range dataset for urban place recognition," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2020, pp. 6246–6253.
- [6] G. Kim and A. Kim, "Scan context: Egocentric spatial descriptor for place recognition within 3d point cloud map," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2018.
- [7] D. Adolphsson, M. Karlsson, V. Kubelka, M. Magnusson, and H. Andreasson, "TBV radar SLAM—trust but verify loop candidates," *IEEE Robot. and Automat. Lett.*, 2023.
- [8] Z. Hong, Y. Petillot, A. Wallace, and S. Wang, "Radarslam: A robust simultaneous localization and mapping system for all weather conditions," *Intl. J. of Robot. Research*, vol. 41, no. 5, pp. 519–542, 2022.
- [9] H. Jang, M. Jung, and A. Kim, "Raplac: Place recognition for imaging radar using radon transform and mutable threshold," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2023.
- [10] M. Gadd and P. Newman, "Open-radvlad: Fast and robust radar place recognition," in *IEEE Radar Conference*, 2024.
- [11] H. Kim, B. Choi, E. Choi, and Y. Cho, "Referee: Radar-based lightweight and robust localization using feature and free space," *IEEE Robot. and Automat. Lett.*, 2024.
- [12] D. Barnes and I. Posner, "Under the radar: Learning to predict robust keypoints for odometry estimation and metric localisation in radar," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2020, pp. 9484–9490.
- [13] Ş. Săftescu, M. Gadd, D. De Martini, D. Barnes, and P. Newman, "Kidnapped radar: Topological radar localisation using rotationally-invariant metric learning," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2020, pp. 4358–4364.
- [14] J. Yuan, P. Newman, and M. Gadd, "Off the radar: Uncertainty-aware radar place recognition with introspective querying and map maintenance," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2023, pp. 1350–1357.
- [15] H. Yin, R. Chen, Y. Wang, and R. Xiong, "Rall: end-to-end radar localization on lidar map using differentiable measurement model," *IEEE Trans. Intell. Transport. Sys.*, vol. 23, no. 7, pp. 6737–6750, 2021.
- [16] A. Nayak, D. Cattaneo, and A. Valada, "Ralf: Flow-based global and metric radar localization in lidar maps," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2024.
- [17] C. Bibby and I. Reid, "A hybrid slam representation for dynamic marine environments," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2010, pp. 257–264.
- [18] C. H. Schiller, B. Arsenali, D. Maas, and S. Maranó, "Improving marine radar odometry by modeling radar resolution and exploiting additional temporal information," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2022.
- [19] H. Jang, M. Jung, M.-H. Jeon, and A. Kim, "Lodestar: Maritime radar descriptor for semi-direct radar odometry," *IEEE Robot. and Automat. Lett.*, 2024.
- [20] B. Hu, X. Liu, Q. Jing, H. Lyu, and Y. Yin, "Estimation of berthing state of maritime autonomous surface ships based on 3d lidar," *Ocean Engineering*, vol. 251, p. 111131, 2022.
- [21] H. Wang, Y. Yin, Q. Jing, Z. Cao, Z. Shao, and D. Guo, "Berthing assistance system for autonomous surface vehicles based on 3d lidar," *Ocean Engineering*, vol. 291, p. 116444, 2024.
- [22] H. Lu, Y. Zhang, C. Zhang, Y. Niu, Z. Wang, and H. Zhang, "A multi-sensor fusion approach for maritime autonomous surface ships berthing navigation perception," *Ocean Engineering*, vol. 316, p. 119965, 2025.
- [23] H. Samet and M. Tamminen, "Efficient component labeling of images of arbitrary dimension represented by linear bintrees," *IEEE Trans. Pattern Analysis and Machine Intell.*, vol. 10, no. 4, pp. 579–586, 1988.
- [24] D. Chung, J. Kim, C. Lee, and J. Kim, "Pohang canal dataset: A multimodal maritime dataset for autonomous navigation in restricted waters," *Intl. J. of Robot. Research*, vol. 42, no. 12, pp. 1104–1114, 2023.