

# SHeRLoc: Synchronized Heterogeneous Radar Place Recognition for Cross-Modal Localization

Hanjun Kim<sup>1</sup>, Minwoo Jung<sup>2</sup>, Woosong Yang<sup>2</sup> and Ayoung Kim<sup>2\*</sup>

**Abstract**—Despite the growing adoption of radar in robotics, the majority of research has been confined to homogeneous sensors, overlooking the integration and cross-modality challenges inherent in heterogeneous radar. This leads to significant difficulties in generalizing across diverse radar types, with modality-aware approaches that could leverage the complementary strengths of heterogeneous radar remaining unexplored. To bridge these gaps, we propose SHeRLoc, the first deep network tailored for heterogeneous radar, which utilizes radar cross-section polar matching to align multimodal radar data. Our hierarchical optimal transport-based feature aggregation generates rotationally robust multi-scale descriptors. By employing FFT-similarity-based data mining and adaptive margin-based triplet loss, SHeRLoc enables FOV-aware metric learning. SHeRLoc achieves an order of magnitude improvement in heterogeneous radar place recognition, increasing recall@1 from below 0.1 to 0.9, and paves the way for cross-modal localization.

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

## I. INTRODUCTION

PLACE recognition (PR) is a cornerstone of robust localization in autonomous driving, enabling vehicles to identify previously visited locations. Traditionally, cameras and LiDAR have dominated PR due to their rich data representations [1, 2]. However, their susceptibility to adverse conditions has shifted attention to radar, which offers unparalleled robustness by penetrating fog, rain, and snow [3, 4]. Early radar-based PR mostly leveraged 360° spinning radars, establishing a foundation for robust localization [5–9]. Building on this foundation, compact system-on-a-chip (SoC) radars have been integrated into PR tasks [10–12]. More recently, 4D radars, capable of capturing range, azimuth, elevation, and radial velocity, have also been adopted in radar PR [13, 14].

Despite the growing adoption of radar in robotics, the majority of research has been confined to homogeneous sensor systems, overlooking the integration and cross-modality challenges inherent in heterogeneous radar technologies. Notably,

Manuscript received: June 16, 2025; Revised: September 12, 2025; Accepted: October 8, 2025.

This letter was recommended for publication by Associate Editor S. Behnke and Editor T. Asfour upon evaluation of the reviewers' comments. This work was supported in part by the Technology Innovation Program (1415187329, 20024355, Development of autonomous driving connectivity technology based on sensor-infrastructure cooperation) funded by MOTIE, Korea, and in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00241758).

<sup>1</sup>H. Kim is with the Dept. of Future Automotive Mobility, SNU, Seoul, S. Korea hanjun815@snu.ac.kr

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

Project Page: <https://sites.google.com/view/radar-sherloc>

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

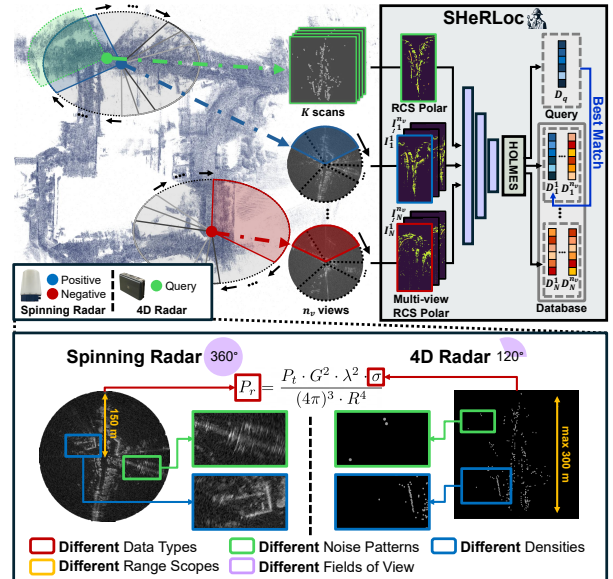


Fig. 1: SHeRLoc generates  $n_v$  views from spinning radar scans to align with the narrow FOV of 4D radar, while bridging the modality gap across heterogeneous radars through RCS polar matching.

different radar systems operate under distinct sensing principles, leading to substantial variations in noise characteristics, point density, and field of view (FOV). As a result, methods designed for homogeneous radar systems fail to generalize across diverse radar configurations. Moreover, radar types are suited for different tasks: spinning radar, with its dense 360° measurements, is ideal for building comprehensive mapping databases, while 4D radar, prevalent in modern vehicles, excels in real-time dynamic sensing and is well-suited for query generation. Despite these complementary properties, existing approaches do not account for the differences between radar modalities, leaving the problem of heterogeneous radar PR largely unresolved. These challenges highlight the need for a new approach: *modality-aware heterogeneous radar PR*.

In this paper, we propose SHeRLoc, Synchronized Heterogeneous Radar Place Recognition for Cross-Modal Localization. Despite the inherent challenges of heterogeneous radars, as illustrated in Fig. 1, SHeRLoc enables radar-agnostic learning through Radar Cross Section (RCS) synchronization, where RCS reflects the object's material and geometry. By employing multi-view polar bird's-eye view (BEV), SHeRLoc effectively addresses FOV differences and achieves rotation invariance. Additionally, we present HOLMES, a feature aggregation network that addresses radar speckle noise and multipath issues through a learnable weight matrix and ghostbin, while ensuring stable optimal transport with an

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

adaptive entropy-regularized Sinkhorn algorithm. Unlike high-dimensional single-scale descriptors [1, 15], which may lack global scene context, HOLMES integrates both local RCS patterns and global structural information into a compact descriptor. Our FFT-similarity-based data mining and adaptive margin-based triplet loss enable spatially and viewpoint-aware metric learning. By overcoming sparsity and noise patterns, SHerLoc addresses the kidnapped robot problem with 4D radar on a spinning radar map, while also being applicable to Frequency modulated continuous wave (FMCW) LiDAR.

Our contributions can be summarized as follows:

- We present SHerLoc, a deep network designed to tackle the complexities of heterogeneous radars, including sparsity, diverse FOV, and varying data dimensions. To the best of our knowledge, this is the first approach specifically designed for heterogeneous radar systems.
- To address the modality gap, we propose a novel RCS matching and multi-view polar projection, enabling radar-agnostic learning. Our preprocessing pipeline effectively removes dynamic objects, clutter, and ground return.
- We propose a hierarchical optimal transport-based feature aggregation method for rotationally robust descriptors, overcoming the loss of global context in traditional high-dimensional descriptors. Our FFT-similarity-based data mining and adaptive margin-based triplet loss ensure spatially and viewpoint-aware metric learning.
- We evaluate SHerLoc against state-of-the-art (SOTA) methods under diverse scenarios (homogeneous, heterogeneous, single-/multi-session, and lidar-to-radar) and demonstrate superior performance in extreme conditions (e.g., crowded, snow, random rotation, and zero-shot). We open-source the code for the radar robotics community.

## II. RELATED WORK

### A. Radar Place Recognition

Spinning radar and SoC radar are the two primary sensor types in radar PR. Early radar PR mainly utilized spinning radar, leveraging their  $360^\circ$  FOV and relatively dense representations. Radar Scan Context [4] extended the LiDAR-based Scan Context [2] for rotation invariance, while Kidnapped Radar [6] introduced a NetVLAD-based approach. As time efficiency became critical, RaPlace [7] and Open-RadVLAD [8] adopted the Fourier Transform, and ReFeree [9] proposed a compact one-dimensional descriptor. These advancements established a foundation for radar PR but faced challenges when the sensor output is sparse, as seen in SoC radar.

The adoption of SoC radars in PR tasks has rendered traditional spinning radar-based methods less applicable due to the narrower FOV and sparser data of SoC radars, necessitating new approaches. Notably, AutoPlace [10] increased the FOV by using five automotive radars and overcame sparsity by utilizing a spatial-temporal encoder. However, it relies on high-dimensional descriptors and requires an additional reranking process. mmPlace [11] expands the FOV by leveraging a rotating platform and concatenating heatmaps, while SPR [12] proposes a lightweight method using a single radar scan. However, both methods are vulnerable in dynamic environments as

they ignore velocity information. Recently, TransLoc4D [13] employs transformer architectures to integrate geometry, intensity, and velocity from 4D radar, but remains susceptible to rotation due to its reliance on a single viewpoint. Hilger et al. [14] considered both co- and counter-directional viewpoints; nevertheless, this approach is still limited to just two viewpoints, and the challenge of PR across heterogeneous radar with distinct data characteristics remains unsolved.

### B. Challenges in Heterogeneous Sensor Place Recognition

Heterogeneous sensor PR requires seamless integration of sensors with distinct data characteristics. RaLF [16] retrieves spinning radar queries from existing LiDAR maps, while Get to the Point [17] leverages overhead imagery for LiDAR PR and metric localization. However, both approaches struggle to handle differing FOVs. Similarly, Cattaneo et al. [18] proposed a LiDAR-Camera PR model without FOV considerations, resulting in  $360^\circ$  point clouds that inevitably include irrelevant scene information outside the RGB camera's FOV. ModaLink [19] addressed this by using camera pose priors to crop point clouds, but this cropping approach is impractical since the required extrinsic calibration is rarely available. In response, LCPR [20] integrated sensor data using panoramic views, while Yao et al. [15] introduced xNetVLAD that accounts for FOV differences. However, these methods require separate networks for each modality. Recently, HeLiOS [21] introduced overlap-based learning with guided-triplet loss for heterogeneous LiDAR PR, though it is limited by LiDAR's high point density and minimal differences between heterogeneous LiDAR scans. In contrast to previous methods, we propose a novel RCS polar matching to achieve sensor-agnostic learning, overcoming radar's sparsity and distinct noise patterns.

## III. METHODS

As illustrated in Fig. 2, SHerLoc transforms 4D radar scans  $\mathcal{S}_{4D}$  and spinning radar scans  $\mathcal{S}_{spin}$  into RCS polar BEV images denoted as  $I_{4D}, I_{spin} \in \mathbb{R}^{H \times W}$ . Here,  $H$  corresponds to a maximum range  $\rho_{max}$ , and  $W$  corresponds to an azimuth span  $\phi$ . The synchronized data is then processed by a feature extraction network  $\mathcal{G}$ , and local features  $F_{4D}, F_{spin}$  are aggregated into global descriptors  $\mathcal{D}_{4D}, \mathcal{D}_{spin}$  using HOLMES  $\mathcal{H}$ , respectively.

### A. Preprocessing

1) *Dynamic Removal with Ego-velocity Estimation:* Each 4D radar scan at time  $k$ ,  $\mathcal{S}_{4D,k} = \{(x_i, y_i, z_i, v_i^d, \sigma_i)\}_{i=1}^N \in \mathbb{R}^{N \times 5}$ , comprises  $N$  points with 3D coordinates  $(x_i, y_i, z_i)$ , Doppler velocity  $v_i^d$ , and RCS  $\sigma_i$ . We employ the 3-Point RANSAC-LSQ [22] to estimate the ego-velocity and apply a refinement function  $f_{removal}$ , which filters out moving objects, clutter, and ground returns. The refined scan is defined as  $\mathcal{S}'_{4D,k} = f_{removal}(\mathcal{S}_{4D,k}) = \{(x_i, y_i, z_i, v_i^d, \sigma_i) \in \mathcal{S}_{4D,k} \mid v_i^d < \tau_v, z_i \geq \tau_z, \sigma_i \geq \tau_\sigma\}$ , where each  $\tau$  denotes a threshold.

2) *FOV Matching with Multi-view Polar Projection:* Prior works [15, 23] address varying FOV issues in panoramic front-view settings using sliding window cropping. However, such approaches are unsuitable for spinning radar lacking elevation

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

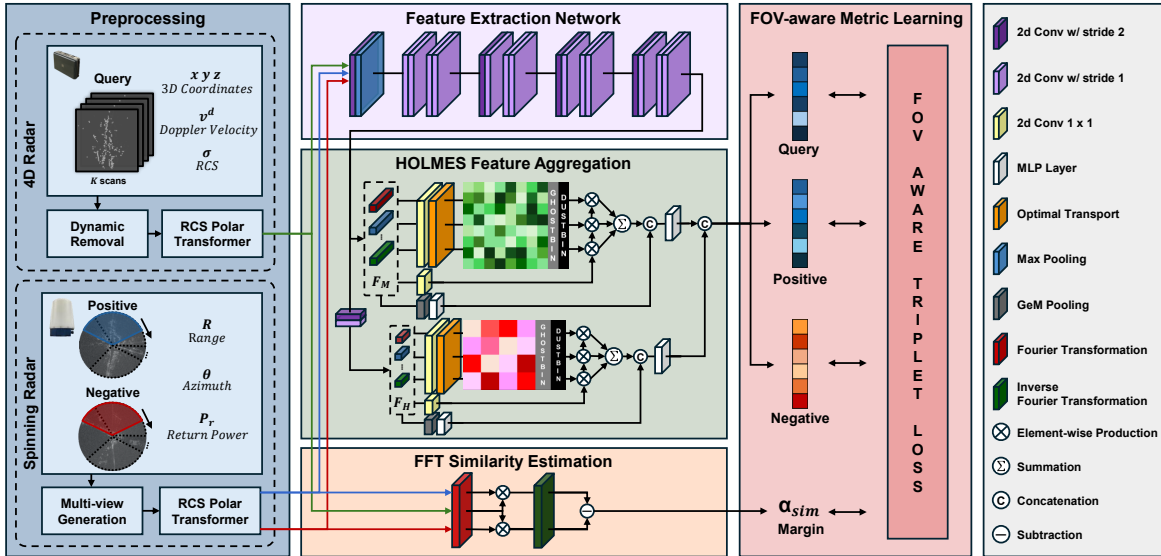


Fig. 2: The overall pipeline of SHerLoc. RCS polar images  $I_{4D}$  and  $I_{spin}$  are generated from heterogeneous radars and processed through a shared feature extraction network  $\mathcal{G}$ . Multi-level features  $F_M$  and  $F_H$  are aggregated into global descriptors  $\mathcal{D}$  using HOLME.

measurements. Additionally, the 4D radar has a sector-shaped FOV in Cartesian BEV, leaving empty regions in the rectangular image and thereby degrading performance. In contrast, polar BEV provides a consistent rectangular representation of overlapping regions between heterogeneous radars.

For the 4D radar,  $\mathcal{S}'_{4D,k}$  is projected into RCS polar BEV image  $I_{4D,k} \in \mathbb{R}^{H \times W}$  using a polar mapping function  $f_{polar} : \mathcal{S}'_{4D,k} \rightarrow I_{4D,k}(r, \theta)$ , where  $r$  and  $\theta$  are defined as:

$$r = \left[ \sqrt{x^2 + y^2} \right] \frac{H}{\rho_{max}}, \quad \theta = \left[ 1 - 2\phi^{-1} \arctan\left(\frac{y}{x}\right) \right] \frac{W}{2}. \quad (1)$$

Each pixel in  $I_{4D,k}$  encodes RCS in a scaled form, consistent with the spinning radar convention of quantizing return power in half-dB steps (e.g., 100 corresponds to 50 dB). To form a temporally aggregated 4D radar image  $I_{4D,t}$  at time  $t$ , we apply max-pooling over  $K$  consecutive frames:

$$I_{4D,t}(r, \theta) = \max_{k \in \{t-K+1, \dots, t\}} I_{4D,k}(r, \theta). \quad (2)$$

For the spinning radar, we rescale  $\mathcal{S}_{spin,t} \in \mathbb{R}^{N_r \times N_a}$ , with  $N_r$  range bins and  $N_a$  azimuth bins, to  $\mathcal{S}'_{spin,t} \in \mathbb{R}^{H \times 3W}$  to account for its three times larger FOV. We then extract  $n_v$  sub-views  $\{M_{spin,t}^j\}_{j=1}^{n_v}$  from  $\mathcal{S}'_{spin,t}$  using a sliding window of size  $H \times W$  along the azimuthal axis. Each sub-view with azimuth offset  $\Delta$  between consecutive views is given by:

$$M_{spin,t}^j = \mathcal{S}'_{spin,t}[:, \Delta(j-1) : \Delta(j-1) + W]. \quad (3)$$

This multi-view generation function  $f_{multiview} : \mathcal{S}'_{spin,t} \rightarrow \{M_{spin,t}^j \in \mathbb{R}^{H \times W}\}_{j=1}^{n_v}$  ensures  $360^\circ$  FOV coverage, supporting rotation-invariance, detailed in Section III-E.

3) *RCS Matching*: We match RCS values  $\sigma$  with return power  $P_r$  using the classical radar equation, formulated as:

$$P_r = \frac{P_t \cdot G^2 \cdot \lambda^2 \cdot \sigma}{(4\pi)^3 \cdot R^4}, \quad (4)$$

where  $P_t$  is the transmitted power,  $G$  is the antenna gain,  $\lambda$  is the wavelength, and  $R$  is the range. An object at height  $h$  above the ground and slant range  $R$  forms an angle  $\alpha$  given by  $\sin \alpha = h/R$ , or equivalently,  $R = h \csc \alpha$ . Inspired by air-surveillance radars that maintain constant received power

regardless of range [24], modern spinning radars adopt a cosecant-squared beam profile, where the antenna gain satisfies  $G \propto R^2$ . Assuming constant  $P_t$  and  $\lambda$ , the RCS in the radar equation (4) is simplified on the dBsm scale as:

$$\sigma_{dBsm} = P_r [\text{dB}] + C, \quad (5)$$

where  $\sigma_{dBsm} = 10 \log_{10}(\sigma)$ , with the reference RCS defined as  $\sigma_{ref} = 1 \text{ m}^2$ , and  $C = 40 \log_{10}(R) + 30 \log_{10}(4\pi) - 20 \log_{10}(G \cdot \lambda) - P_t [\text{dB}]$  is a correction constant.

To synchronize  $M_{spin,t}$  with  $I_{4D,t}$ , where pixels represent  $P_r$  and  $\sigma$  in half-dB steps, respectively, we precompute a correction term  $C_{corr}$  only once per radar pair to align  $I_{4D,i}$  with the most similar subview  $M_{spin,i}^{\max}$ . Then, each  $M_{spin,t}^j$  is processed by an RCS matching module  $f_{rcs} : M_{spin,t}^j \rightarrow I_{spin,t}^j$  to produce the final RCS polar BEV representation:

$$I_{spin,t}^j = f_{rcs}(M_{spin,t}^j) = M_{spin,t}^j + C_{corr}, \quad (6)$$

where  $C_{corr}$  is obtained by optimizing a per-frame correction term  $k_i$  to minimize the Huber loss. Let the set of valid pixel coordinates be:  $\Omega_i = \{(r, \theta) \mid M_{spin,i}^{\max}(r, \theta) \neq 0 \text{ and } I_{4D,i}(r, \theta) \neq 0\}$ , and define the Huber loss as:

$$L_\delta(r) = \begin{cases} \frac{1}{2} r^2, & \text{if } |r| \leq \delta, \\ \delta (|r| - \frac{1}{2} \delta), & \text{if } |r| > \delta. \end{cases} \quad (7)$$

Then, The loss for each image pair  $i$  can be formulated as:

$$L_i(k_i) = \frac{1}{|\Omega_i|} \sum_{(r, \theta) \in \Omega_i} L_\delta(I_{4D,i}(r, \theta) - (M_{spin,i}^{\max}(r, \theta) + k_i)). \quad (8)$$

To ensure smooth variation of  $k_i$ , we introduce a regularization term with factor  $\lambda$  in the joint optimization problem:

$$\{k_i^*\}_{i=1}^N = \underset{\{k_i\}_{i=1}^N}{\operatorname{argmin}} \left( \sum_{i=1}^N L_i(k_i) + \lambda \sum_{j=2}^N (k_j - k_{j-1})^2 \right). \quad (9)$$

Then, we compute  $C_{corr} = \frac{1}{N} \sum_{i=1}^N k_i^*$ . As shown in Fig. 3, 4D radar scans  $\mathcal{S}_{4D}$  and spinning radar scans  $\mathcal{S}_{spin}$  are synchronized into an RCS polar BEV representation.

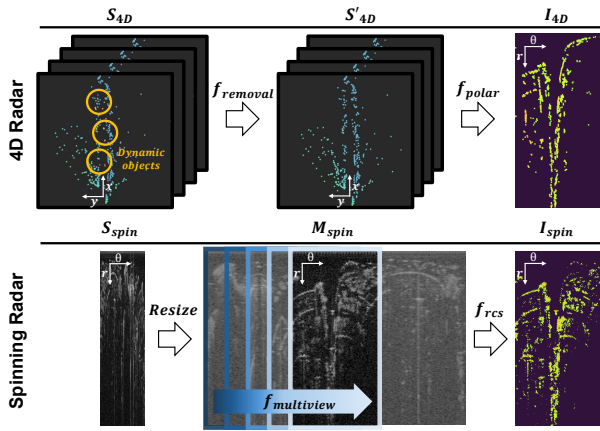


Fig. 3: The pipeline of RCS polar synchronization. The turbo colormap is applied to  $I_{4D}$  and  $I_{spin}$  for visualization clarity.

### B. Feature Extraction Network

The preprocessed  $I_{4D}, I_{spin} \in \mathbb{R}^{H \times W}$  are fed simultaneously into a ResNet-based architecture [25],  $\mathcal{G} : I_{4D}, I_{spin} \rightarrow F_{4D}, F_{spin}$ , which generate the feature maps  $F_{4D}, F_{spin} \in \mathbb{R}^{C \times H/32 \times W/32}$ . Traditional heterogeneous PR methods require separate encoders for each sensor [16] or a backbone network for independent learning [15, 18]. However, due to the shared RCS polar BEV format, our approach allows the use of a common network across radars, unlike previous methods.

### C. Multi-scale Descriptors with Optimal Transport

To achieve a compact integration of complementary feature hierarchies, we propose HOLMES, Hierarchical Optimal transport for Locally aggregated Multi-scale descriptors with adaptive Entropy-regularized Sinkhorn-algorithm. Previously, SALAD [26] fuses local features with a DINOv2 CLS token, while HeLiOS [21] combines them with GeM-pooled global features. However, both result in high-dimensional descriptors, whereas HeLiOS-S, a smaller variant of HeLiOS, sacrifices expressiveness and global context. We address these issues with a hierarchical framework that combines local RCS patterns and global context. Mid-level features form a local feature matrix  $\mathbf{F}_M \in \mathbb{R}^{n_M \times d_{f_M}}$ , and high-level features, processed through an additional CNN layer, form  $\mathbf{F}_H \in \mathbb{R}^{n_H \times d_{f_H}}$ , where  $n$  is the number of spatial locations and  $d_f$  is the feature dimension. A convolution layer predicts a score matrix  $\mathbf{S} \in \mathbb{R}^{n_M \times m_M}$  for  $\mathbf{F}_M$ , where  $m_M$  denotes the number of clusters. Inspired by GhostVLAD [27], we extend SALAD's dustbin with a *ghostbin* to handle noise points, resulting in a score matrix  $\tilde{\mathbf{S}} \in \mathbb{R}^{n_M \times (m_M + 2)}$ . The Sinkhorn algorithm [28] optimizes feature-to-cluster assignments, yielding a refined matrix  $\mathbf{R} \in \mathbb{R}^{n_M \times m_M}$  by normalizing  $\exp(\tilde{\mathbf{S}})$  and dropping dustbin and ghostbin. We introduce an adaptive entropy regularization as:

$$\text{reg} = 1 + 2 \cdot \tanh \left( \frac{\frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2}{2(\mu + \epsilon)} \right), \quad (10)$$

where  $x_i$  is RCS feature,  $\mu$  is the mean and  $N$  is the number of elements. We compute the variance relative to  $\mu$  and apply  $\tanh(\cdot)$  to bound the regularization smoothly. High variance induces smooth matching for stability, while low variance leads to precise matching. Then, the aggregated feature matrix  $\mathbf{V} \in \mathbb{R}^{m_M \times l_M}$  with the cluster dimension  $l_M$  is computed as:

$$V_{j,k} = \sum_{i=1}^n R_{i,k} \cdot \bar{\mathbf{F}}_{i,j}, \quad (11)$$

where  $\bar{\mathbf{F}}$  is the feature derived by applying a dimensionality reduction to  $\mathbf{F}$ . Meanwhile, GeM pooling and MLP layers produce a global representation  $\mathbf{G} \in \mathbb{R}^{s_M}$ . Concatenating  $\mathbf{V}$  and  $\mathbf{G}$  yields  $\mathbf{g} \in \mathbb{R}^{m_M \cdot l_M + s_M}$ , which is then transformed by a learnable weight matrix  $\mathbf{H} \in \mathbb{R}^{(m_M \cdot l_M + s_M) \times d_M}$  into the mid-level descriptor  $\mathcal{D}_M \in \mathbb{R}^{d_M}$ . Applying the same process to  $\mathbf{F}_H$  produces a high-level descriptor  $\mathcal{D}_H \in \mathbb{R}^{d_H}$ . By concatenating the mid-level descriptor, the HOLMES  $\mathcal{H}$  generates the final descriptor  $\mathcal{D} = \mathcal{D}_M \oplus \mathcal{D}_H \in \mathbb{R}^{d_M + d_H}$ .

### D. FOV-aware Metric Learning

1) *FOV-aware Data Mining*: We introduce a FOV-aware data mining approach that selects positive and negative samples based on view similarity. For efficient similarity computation, we leverage the Convolution Theorem, performing computations in the frequency domain [29]. Concretely, the cross-correlation between two scans  $I_1$  and  $I_2$  is defined as:

$$\text{Sim}(I_1, I_2) = \frac{\max \left[ \mathcal{F}^{-1} \left( \mathcal{F}(I_1) \cdot \mathcal{F}(I_2) \right) \right]}{\|I_1\|_2 \cdot \|I_2\|_2}, \quad (12)$$

where  $\mathcal{F}$  and  $\mathcal{F}^{-1}$  denote the Fourier and inverse Fourier transforms, respectively. The complexity is reduced from  $\mathcal{O}(N^2)$  to  $\mathcal{O}(N \log N)$ , making data mining time-efficient.

We select the multi-view set  $\{I_{spin,t}^j\}_{j=1}^{n_v}$  with the closest timestamp to the corresponding  $I_{4D,t}$ . Among these multi-views, the scan  $I_{spin,t}^j$  that has the highest similarity is labeled as a positive. In contrast, scans beyond the 25 m radius are considered negative, even if their FOV partially overlap, ensuring both spatially and viewpoint-wise meaningful.

2) *FOV-aware Triplet Loss*: To train the network with the mined samples, we adopt a FoV-aware triplet loss formulation. The objective is to bring positive samples closer in the embedding space while pushing negatives apart, with a margin that dynamically adapts to the degree of similarity:

$$\mathcal{L}_{\text{triplet}} = \max \left( d(x_i^q, x_j^p) - \min_{n=1}^N d(x_i^q, x_j^n) + \alpha_{\text{sim}}, 0 \right), \quad (13)$$

where  $x_i^q$  is the query pair,  $x_j^p$  is the positive pair,  $x_j^n$  are negative pair, and  $d(\cdot)$  denotes the  $L_2$  distance in the embedding space. The adaptive margin  $\alpha_{\text{sim}}$  is defined as:

$$\alpha_{\text{sim}} = \gamma (\text{Sim}(\mathbf{q}, \mathbf{p}) - \text{Sim}(\mathbf{q}, \mathbf{n})), \quad (14)$$

where  $\gamma$  is a scale factor. By incorporating similarity into the margin, the loss becomes sensitive to FOV overlap. If a negative has a similar view to the query, the margin becomes smaller, making the model focus on harder negatives.

### E. Rotation Invariance and Intentional Translation Variance

As shown in Fig. 4, spatially distant negatives with overlapping FOV produce similar features when using Cartesian BEV, as seen in [10, 16]. Also, nearby positives may exhibit feature mismatches due to small rotations. To address this, we leverage  $f_{\text{polar}}$ ,  $f_{\text{multiview}}$ ,  $\mathcal{G}$ , and  $\mathcal{H}$  to achieve rotation invariance and intentional translation variance, thereby maintaining discriminative power in long-range radar PR. See RING++ [30] for more details about invariance and equivariance.

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

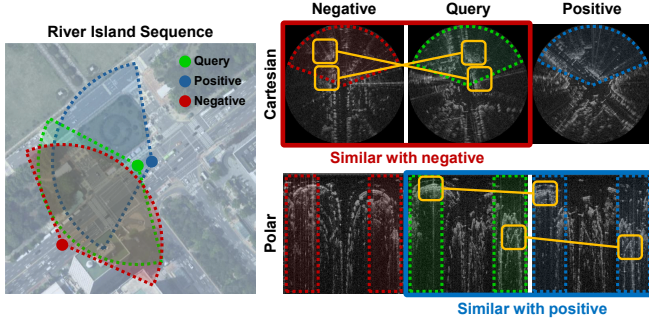


Fig. 4: For a far but overlapping negative and a nearby but rotated positive, Cartesian BEV yields higher similarity with the negative due to rotation variance. In contrast, polar BEV produces higher similarity with the positive, demonstrating robustness to rotation.

1) *Pipeline Invariance*: A rotated scan  $I'(r, \theta) = I(r, \theta + \delta\theta)$  shifts multi-view azimuths cyclically. Then,  $\mathcal{G}$  outputs  $F'_j(r, \theta)$ , and  $\mathcal{H}$  produces an invariant descriptor. With polar BEV ensuring translation variance, the pipeline is rotation-invariant and intentionally translation variant, satisfying  $\mathcal{H}(\mathcal{G}(f_{\text{multiview}}(T_{\delta\theta}S(r, \theta))) = \mathcal{H}(\mathcal{G}(f_{\text{multiview}}(S(r, \theta))))$ .

2) *Feature Extraction Equivariance*: Since a rotation transformation  $T_{\delta\theta}$  corresponds to a shift in the azimuth dimension and  $f_{\text{multiview}}$  operates uniformly across  $\theta$ , the output of  $f_{\text{multiview}}$  is shifted by the same  $\delta\theta$ :  $f_{\text{multiview}}(T_{\delta\theta}S(r, \theta)) = T_{\delta\theta}f_{\text{multiview}}(S(r, \theta))$ . Similarly, the CNN-based network  $\mathcal{G}$  preserves translation-equivariance up to edge effects [31]. Therefore, both  $f_{\text{multiview}}$  and  $\mathcal{G}$  maintain the translation-equivariance property. Theoretically, we have to generate infinite multi-views, and aliasing may occur due to downsampling. However, robust performance is achieved with a limited number of multi-views, detailed in Section IV-D.

3) *HOLMES Invariance*:  $\mathcal{H}$  processes feature maps  $F$  to produce  $\mathcal{D} = \mathcal{D}_M \oplus \mathcal{D}_H$ . GeM Pooling produces a global feature  $\mathbf{G}$ , invariant to translations, as the summation is order-independent. A convolutional layer produces a score matrix  $\mathbf{S}$  for Optimal Transport, and the Sinkhorn algorithm computes assignments  $\mathbf{R}$ . The aggregated feature  $\mathbf{V}_{j,k} = \sum_{i=1}^n R_{i,k} \cdot F_{i,j}$  is invariant, as summation reorders with translated inputs  $F'_j(r, \theta)$ . Ghostbin and dustbin discard non-informative points, preserving invariance by excluding spatial dependencies. For multi-scale processing, high-level features follow the same process. Concatenation  $(\mathcal{D}_M \oplus \mathcal{D}_H)$  remains invariant to cyclic shifts. For translated feature maps  $F'_j(r, \theta)$ ,  $\mathcal{H}$  yields  $\mathbf{D}'_M \oplus \mathbf{D}'_H = \mathcal{H}(F'_j(r, \theta)) = \mathcal{H}(F_j(r, \theta)) = \mathcal{D}_M \oplus \mathcal{D}_H$ . Thus, our proposed HOLMES network is translation-invariant.

## IV. EXPERIMENTAL RESULTS

### A. Implementation Details

We trained SHERLoc on a GeForce RTX 3090 using the AdamW optimizer with CosineAnnealingLR scheduler, starting with a learning rate of 1e-3 and decaying to a minimum of 1e-5 over 20 epochs. For the polar BEV projection, the resolution of a polar image with a maximum range  $\rho_{\text{max}} = 150$  m and an azimuth span  $\phi = 120^\circ$  is  $384 \times 192$ . The number of scans  $K$  is set to 5 for (2); the multi-view offset  $\Delta$  is set to  $10^\circ$  (16 pixels); and  $\lambda$  is set to 0.1 for (9). For HOLMES,  $(m_M, l_M, s_M, d_M)$  are set to (64, 256, 256, 256), and

$(m_H, l_H, s_H, d_H)$  are set to (16, 64, 64, 64), respectively. We utilize one positive and five negative samples for (13), and  $\gamma$  is set to 1.0 for (14). For SHERLoc-S, a lightweight variant, only the mid-level descriptor  $\mathcal{D}_M \in \mathbb{R}^{256}$  is retained.

### B. Heterogeneous Radar Dataset and Evaluation Metrics

To evaluate SHERLoc, a dataset must include both 4D radar and spinning radar. Existing public datasets, such as Oxford Radar RobotCar [3] and MulRan [4], are limited to spinning radar, while NTU4DRadLM [32] and SNAIL Radar [33] focus solely on 4D radar. In contrast, HeRCULES [34] is the only publicly available heterogeneous range sensor dataset, integrating 4D radar, spinning radar, and FMCW LiDAR, making it uniquely suited for heterogeneous radar PR evaluation.

We evaluate performance using Recall@K (R@K) and Average Recall@K (AR@K). A retrieval is deemed correct if the predicted location lies within a 5 m radius of the ground-truth. In addition, we employ the Precision-Recall (PR) curve to assess overall retrieval performance. The PR curve is obtained by sweeping a threshold  $\tau$  over the descriptor distance  $d(x_i^q, x_j^m)$  between query and map embeddings. For each  $\tau \in [0, 2]$ , uniformly sampled at 1000 intervals, a retrieved pair is considered a match if  $d(x_i^q, x_j^m) < \tau$ .

### C. Comparison with State-of-the-Art Methods

As the first heterogeneous radar PR model, SHERLoc was compared against homogeneous PR SOTA models, including spinning radar PR methods Radar Scan Context [4], RaPlace [7], RadVLAD, and FFT-RadVLAD [8], as well as SoC radar PR models AutoPlace [10], TransLoc4D [13], and LiDAR PR method MinkLoc3Dv2 [35]. For fair comparison, we used 5-scan aggregation for 4D radar and applied RCS matching for spinning radar across all models. Since methods like RaPlace [7] and AutoPlace [10] operate on BEV images, they were included in the heterogeneous radar PR evaluation. Also, all network-based models were trained on the Mountain 01-03, Bridge 01, Stream 02, and Parking Lot 03-04, while validation was conducted on Parking Lot 01-02 of HeRCULES.

1) *Heterogeneous Radar Single-session PR*: For single-session evaluation, spinning radar data were used as the database and 4D radar data as queries across the Sports Complex, Library, and River Island sequences. As shown in Table I, SHERLoc consistently outperforms all baselines, while the lightweight SHERLoc-S also surpasses all SOTA methods. Furthermore, as illustrated in Fig. 5, SHERLoc achieves robust performance even in challenging scenarios such as night, snow, and dynamic object-rich environments.

2) *Heterogeneous Radar Multi-session PR*: For multi-session evaluation, spinning radar data from sequence 01 were used as the database, and 4D radar data from sequences 02 and 03 were used as queries. Quantitative results are presented in Table II, with qualitative analysis shown in Fig. 6 and recall@N curves depicted in Fig. 7. We have demonstrated the necessity of heterogeneous radar PR models beyond homogeneous radar systems by showing that existing SOTA methods fail to address the inherent challenges of heterogeneous radar.

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

TABLE I: Performance Comparison for Single-session Place Recognition with Heterogeneous Radar in Various Challenging Conditions (☀: Clear, 🌆: Dusk, 🌃: Night, ☁: Cloud, ❄: Snow, 🚗: dynamic object-rich, **Bold**: Best, Underline: Second Best)

Methods	Sports Complex						Library						River Island					
	01 ☀		02 🌃		03 ❄		01 ☀		02 🌃		03 ❄		01 🌆🚗		02 🌆🚗		03 🌆🚗	
	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%
Radar SC [4]	0.047	0.135	0.042	0.061	0.044	0.064	0.040	0.060	0.041	0.063	0.042	0.066	0.010	0.020	0.004	0.012	0.005	0.030
RaPlace [7]	0.042	0.148	0.056	0.156	0.025	0.152	0.035	0.069	0.040	0.090	0.016	0.123	0.002	0.152	0.000	0.024	0.004	0.162
RadVLAD [8]	0.019	0.334	0.051	0.222	0.022	0.426	0.011	0.237	0.014	0.295	0.012	0.245	0.001	0.217	0.066	0.543	0.006	0.367
FFT-RadVLAD [8]	0.040	0.344	0.014	0.186	0.022	0.377	0.023	0.287	0.014	0.273	0.011	0.229	0.002	0.330	0.026	0.533	0.001	0.488
Autoplace [10]	0.015	0.198	0.046	0.098	0.017	0.186	0.010	0.129	0.031	0.131	0.018	0.122	0.003	0.051	0.001	0.009	0.001	0.021
SHeRLoc-S	<u>0.857</u>	<u>0.924</u>	<b>0.975</b>	<u>0.985</u>	<u>0.945</u>	<b>0.981</b>	<u>0.866</u>	<u>0.917</u>	<u>0.925</u>	<u>0.950</u>	<u>0.850</u>	<u>0.891</u>	<b>0.899</b>	<b>0.965</b>	<b>0.868</b>	<b>0.963</b>	<u>0.850</u>	<u>0.945</u>
SHeRLoc	<b>0.900</b>	<b>0.936</b>	<u>0.962</u>	<b>0.987</b>	<b>0.958</b>	<u>0.980</u>	<b>0.881</b>	<b>0.936</b>	<b>0.938</b>	<b>0.964</b>	<b>0.868</b>	<b>0.912</b>	<u>0.880</u>	<u>0.957</u>	<u>0.860</u>	<u>0.952</u>	<b>0.858</b>	<b>0.959</b>

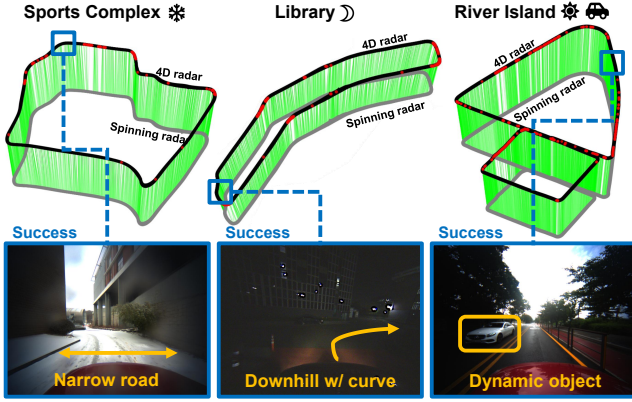


Fig. 5: Trajectory from Sports Complex, Library, and River Island sequences, with green indicating true matching pairs, high-lighting SHeRLoc’s robustness in challenging scenarios.




Fig. 6: Qualitative results of SOTA methods on a challenging query. Green indicates successful retrieval, and red indicates failure.

3) *Homogeneous Radar Multi-session PR*: SHeRLoc was compared with Autoplace [10], MinkLoc3Dv2 [35], and TransLoc4D [13] for 4D radar PR, and with RaPlace [7], RadVLAD, and FFT-RadVLAD [8] for spinning radar PR. The results are reported in Table II and Fig. 8, confirming SHeRLoc’s superior performance despite not being designed for homogeneous radar. For 4D radar PR, it demonstrates top performance across nearly all sequences. For spinning radar PR, SHeRLoc achieves an AR@1 of 0.993, surpassing all SOTA methods. Furthermore, PR curves indicate that SHeRLoc achieves high precision and recall for both radar types under diverse conditions, demonstrating robust performance for both homogeneous and heterogeneous radar.

#### D. Robustness against Random Rotation

In real-world scenarios, random yaw angle rotations between the database and query are common, making rotation robustness crucial. Therefore, we conducted experiments

TABLE II: Performance Comparison for Multi-session Place Recognition with Heterogeneous Radar and Homogeneous Radar

Methods	Sports Complex				Library				AR@1	
	01 - 02	01 - 03	01 - 02	01 - 03	01 - 02	01 - 03	01 - 02	01 - 03		
	R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%		
Heterogeneous	RaPlace [7]	0.024	0.056	0.016	0.057	0.021	0.044	0.011	0.023	0.018
	RadVLAD [8]	0.011	0.118	0.005	0.214	0.012	0.138	0.011	0.073	0.010
	FFT-RadVLAD [8]	0.006	0.121	0.007	0.182	0.015	0.217	0.016	0.136	0.011
	Autoplace [10]	0.007	0.071	0.012	0.070	0.022	0.129	0.007	0.119	0.012
	SHeRLoc-S	<u>0.796</u>	<u>0.890</u>	<u>0.580</u>	<u>0.696</u>	<b>0.822</b>	<b>0.892</b>	<b>0.618</b>	<b>0.757</b>	<u>0.704</u>
SHeRLoc	<b>0.812</b>	<b>0.893</b>	<b>0.650</b>	<b>0.759</b>	<u>0.817</u>	<u>0.887</u>	<u>0.610</u>	<u>0.743</u>	<b>0.722</b>	
4D	Autoplace [10]	0.799	0.967	0.725	0.945	0.812	0.986	0.619	0.901	0.738
	MinkLoc3Dv2 [35]	0.837	<b>0.982</b>	0.743	<b>0.977</b>	0.725	0.981	0.619	0.963	0.735
	TransLoc4D [13]	0.833	<u>0.970</u>	<u>0.804</u>	<u>0.976</u>	0.801	<u>0.991</u>	0.676	0.940	0.779
	SHeRLoc-S	<u>0.866</u>	0.945	0.803	0.925	0.914	0.986	<u>0.908</u>	<b>0.988</b>	0.873
	SHeRLoc	<b>0.904</b>	0.961	<b>0.843</b>	0.949	<b>0.950</b>	<b>0.995</b>	<b>0.923</b>	<b>0.988</b>	<b>0.905</b>
Spinning	RaPlace [7]	0.943	0.962	0.927	0.989	0.998	1.000	0.906	0.955	0.944
	RadVLAD [8]	<b>0.996</b>	<b>1.000</b>	<u>0.981</u>	<b>1.000</b>	<u>0.999</u>	<b>1.000</b>	<u>0.975</u>	<b>1.000</b>	0.988
	FFT-RadVLAD [8]	<b>0.996</b>	<b>1.000</b>	<b>0.985</b>	<b>1.000</b>	<u>0.999</u>	<b>1.000</b>	<u>0.979</u>	<b>1.000</b>	0.990
	SHeRLoc-S	0.963	<u>0.993</u>	0.962	0.981	<b>1.000</b>	<b>1.000</b>	0.936	<u>0.997</u>	0.965
	SHeRLoc	<u>0.993</u>	<b>1.000</b>	<u>0.981</u>	<u>0.999</u>	<u>0.999</u>	<b>1.000</b>	<b>0.998</b>	<b>1.000</b>	<b>0.993</b>

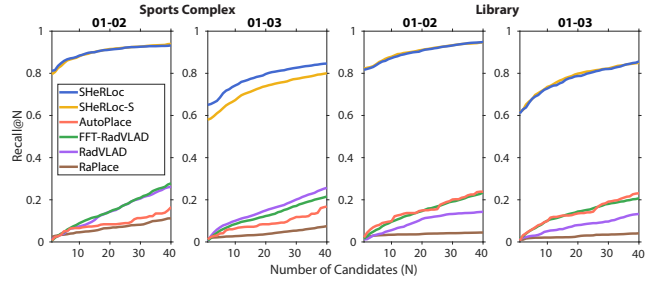


Fig. 7: R@N curves for multi-session PR with heterogeneous radar. SHeRLoc outperforms all methods when retrieving 40 candidates.

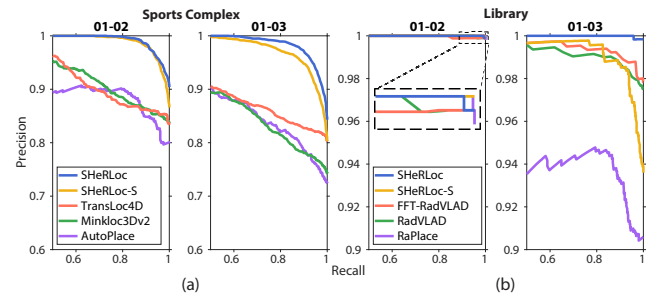


Fig. 8: Precision-Recall curve for homogeneous radar multi-session PR: (a) 4D radar-based PR results on the Sports Complex, (b) Spinning radar-based PR results on the Library.

with random yaw angle rotations on database scans within a 360° range. Through polar domain multi-view generation and the rotation-invariant design of the HOLMES feature aggregation network, SHeRLoc demonstrates superior rotation robustness compared to other methods. As shown in Table III, comparing AR@1 with SOTA methods on Sports Complex, Library, and River Island, both SHeRLoc and SHeRLoc-S consistently achieve the best performance. This robust performance, despite angle variations, demonstrates that a limited number of multi-views is sufficient for achieving robust rotation invariance.

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

TABLE III: Performance Comparison under Random Rotation

Methods	Sports Complex		Library		River Island	
	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%
Radar SC [4]	0.045	0.087	0.041	0.063	0.007	0.021
RaPlace [7]	0.030	0.204	0.130	0.229	0.002	0.191
RadVLAD [8]	0.011	0.318	0.012	0.249	0.016	0.402
FFT-RadVLAD [8]	0.019	0.310	0.012	0.234	0.019	0.429
Autoplace [10]	0.026	0.161	0.020	0.127	0.001	0.122
SHeRLoc-S	0.905	0.959	0.828	0.893	0.710	0.880
SHeRLoc	<b>0.918</b>	<b>0.964</b>	<b>0.858</b>	<b>0.920</b>	<b>0.733</b>	<b>0.904</b>

TABLE IV: Zero-shot Performance on Unseen Datasets (R@1)

Methods	MulRan		Oxford Radar		AR@1
	DCC 01	KAIST 03	#1 to #3	#2 to #3	
Radar SC [4]	0.683	0.896	0.655	0.348	0.646
RaPlace [7]	0.641	0.932	0.666	0.437	0.669
RadVLAD [8]	0.507	0.833	0.813	0.693	0.712
FFT-RadVLAD [8]	0.706	0.948	0.666	0.427	0.687
SHeRLoc-S	<b>0.732</b>	<b>0.961</b>	<b>0.855</b>	<b>0.748</b>	<b>0.824</b>
SHeRLoc	<u>0.708</u>	<b>0.963</b>	<b>0.900</b>	<b>0.890</b>	<b>0.865</b>

TABLE V: Performance Comparison for Place Recognition with Spinning Radar as Database and FMCW LiDAR as Query

Methods	Sports Complex		Library		River Island	
	AR@1	AR@1%	AR@1	AR@1%	AR@1	AR@1%
Radar SC [4]	0.141	0.188	0.046	0.060	0.033	0.050
Autoplace [10]	0.017	0.105	0.014	0.059	0.004	0.041
Radar-to-LiDAR [36]	<u>0.656</u>	0.832	0.468	0.769	0.229	0.530
SHeRLoc-S	<b>0.948</b>	<b>0.971</b>	<b>0.895</b>	<b>0.945</b>	<b>0.932</b>	<b>0.983</b>
SHeRLoc	<b>0.948</b>	<b>0.977</b>	<b>0.899</b>	<b>0.931</b>	<b>0.952</b>	<b>0.987</b>

## E. Zero-shot Generalization Performance Evaluation

We evaluated SHeRLoc’s zero-shot generalization on unseen datasets, including MulRan [4] and the Oxford Radar RobotCar [3]. While HeRCULES [34] employs the Navtech RAS6 spinning radar, MulRan uses the earlier Navtech CIR204-H, and Oxford Radar RobotCar uses the Navtech CTS350-X. These hardware differences induce substantial domain shifts in terms of noise characteristics, resolution, and detection range. For MulRan, we conducted single-session PR on the DCC 01 and KAIST 03 sequences. For Oxford Radar RobotCar, we performed multi-session PR using 2019-01-18 (#3) as the database and 2019-01-10 (#1) and 2019-01-16 (#2) as queries. As shown in Table IV, SHeRLoc achieves strong recall even under challenging zero-shot settings with domain shifts.

## F. Heterogeneous Range Sensor Place Recognition

We evaluated heterogeneous range sensor PR performance using FMCW LiDAR queries and a spinning radar database, comparing it with Radar Scan Context [4], Autoplace [10], and Radar-to-LiDAR [36]. We utilized radial velocity in a manner analogous to 4D radar, and replaced RCS with reflectivity  $\rho$ , which satisfies the following LiDAR equation:

$$I \propto \frac{P_t D_r^2 \eta^2 \rho \cos \alpha_i}{4R^2}, \quad (15)$$

where  $I$  is the LiDAR intensity,  $P_t$  is the transmitted power,  $D_r$  is the receiver aperture diameter,  $\eta$  is the system transmission factor,  $\alpha_i$  is the angle of incidence. Our preprocessing aligns the distributions of FMCW LiDAR and spinning radar, as shown in Fig. 9, and Table V confirms that SHeRLoc remains robust across heterogeneous range sensors.

TABLE VI: Ablation Study on SHeRLoc Components to Evaluate Individual Contributions

B	S	P	R	A	Sports Complex		Library					
					01 - 02		01 - 03		01 - 02		01 - 03	
					R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%
✓	×	×	×	×	0.009	0.015	0.009	0.013	0.019	0.032	0.016	0.035
✓	✓	×	×	×	0.013	0.072	0.009	0.032	0.011	0.022	0.008	0.031
✓	✓	✓	×	×	0.659	0.792	0.516	0.633	0.657	0.749	0.497	0.665
✓	✓	✓	✓	×	0.737	0.840	0.529	0.625	0.751	0.836	0.514	0.679
×	✓	✓	✓	✓	0.558	0.748	0.295	0.426	0.511	0.631	0.199	0.350
✓	✓	✓	✓	✓	<b>0.812</b>	<b>0.893</b>	<b>0.650</b>	<b>0.759</b>	<b>0.817</b>	<b>0.887</b>	<b>0.610</b>	<b>0.743</b>

TABLE VII: Ablation Study on Feature Aggregation Method

Methods	Size	Sports Complex		Library					
		01 - 02		01 - 03		01 - 02		01 - 03	
		R@1	R@1%	R@1	R@1%	R@1	R@1%	R@1	R@1%
NetVLAD [1]	32768	0.378	0.598	0.151	0.254	0.302	0.474	0.102	0.218
MAC [37]	512	0.700	0.846	0.404	0.534	0.525	0.657	0.258	0.514
SPoC [38]	512	0.282	0.517	0.112	0.210	0.155	0.258	0.036	0.150
GeM [39]	512	0.486	0.717	0.234	0.384	0.324	0.475	0.091	0.197
SALAD [26]	8448	0.371	0.611	0.170	0.295	0.330	0.497	0.094	0.235
HOLMES-S*	<b>256</b>	0.769	0.882	0.616	0.730	<b>0.840</b>	<b>0.909</b>	0.587	0.756
HOLMES-S	<b>256</b>	0.796	0.890	0.580	0.696	0.822	0.892	<b>0.618</b>	<b>0.757</b>
HOLMES	<b>320</b>	<b>0.812</b>	<b>0.893</b>	<b>0.650</b>	<b>0.759</b>	0.817	0.887	<u>0.610</u>	0.743

HOLMES-S\*: HOLMES-S without ghostbin and adaptive entropy regularization.




Fig. 9: Data distribution of each sensor and the synchronized data distribution after SHeRLoc’s RCS polar domain matching.

## G. Ablation Study

1) *BSPRA Components*: We compared SHeRLoc variants that incorporate the following modules: **B** (Backbone sharing), **S** (Scan aggregation), **P** (Polar projection), **R** (RCS matching), **A** (Adaptive margin loss). Table VI demonstrates how each component contributes to the performance. In the Cartesian domain, performance is significantly poor, and applying **S** alone fails to yield significant improvement. However, combining **S** with **P**, as in our approach, provides an average improvement of 57.2% in recall@1, with **R** further enhancing performance by 5.05%. As illustrated in Fig. 9, although each sensor’s raw data has different distributions, using **S**, **P**, and **R** leads to very similar data distributions. Finally, **A** contributes an additional 8.95% increase, resulting in an overall improvement of 71.2%. Moreover, **B** encourages learning a shared feature space, which enhances cross-modal generalization and avoids redundant parameterization.

2) *Feature Aggregation*: We evaluated HOLMES against NetVLAD [1], MAC [37], SPoC [38], GeM [39], and SALAD [26]. As shown in Table VII, widely used methods like GeM pooling are vulnerable to noisy and sparse radar data. Additionally, NetVLAD and SALAD require large descriptor sizes, leading to longer processing times and higher storage demands. In contrast, HOLMES-S, with a compact 256-dimensional descriptor, achieves accurate PR performance through its radar-specific design, including ghostbin and adaptive entropy regularization. Moreover, HOLMES, with a descriptor size of 320 dimensions, further improves performance by employing multi-scale descriptors. As a result, HOLMES outperforms other methods, while requiring less storage and generating each descriptor in just 13ms.

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

V. CONCLUSION

In this work, we introduce SHerLoc, the first deep network for heterogeneous radar PR. SHerLoc employs a novel preprocessing method to transform radar data into synchronized RCS polar BEV representations and hierarchical optimal transport-based multi-scale descriptors. Our FOV-aware data mining and adaptive margin-based triplet loss facilitate spatially and viewpoint-wise meaningful metric learning. Evaluations on diverse sequences demonstrate that SHerLoc achieves robust performance in challenging environments and is applicable to both heterogeneous and homogeneous radar systems, as well as heterogeneous range sensor systems. As the first study on heterogeneous radar PR, SHerLoc opens new opportunities for cross-modal place recognition, enabling efficient localization of vehicles with automotive 4D radars on pre-built spinning radar maps.

REFERENCES

- [1] R. Arandjelovic, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, "Netvlad: Cnn architecture for weakly supervised place recognition," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recog.*, 2016, pp. 5297–5307.
- [2] 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.
- [3] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The oxford radar robotcar dataset: A radar extension to the oxford robotcar dataset," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2020, pp. 6433–6438.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] H. Jang, M. Jung, and A. Kim, "Rplace: Place recognition for imaging radar using radon transform and mutable threshold," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2023.
- [8] M. Gadd and P. Newman, "Open-radvlad: Fast and robust radar place recognition," in *2024 IEEE Radar Conference*, 2024, pp. 1–6.
- [9] 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.
- [10] K. Cai, B. Wang, and C. X. Lu, "Autoplace: Robust place recognition with single-chip automotive radar," in *Proc. IEEE Intl. Conf. on Robot. and Automat.*, 2022, pp. 2222–2228.
- [11] C. Meng, Y. Duan, C. He, D. Wang, X. Fan, and Y. Zhang, "mmplace: Robust place recognition with intermediate frequency signal of low-cost single-chip millimeter wave radar," *IEEE Robot. and Automat. Lett.*, 2024.
- [12] D. C. Herraes, L. Chang, M. Zeller, L. Wiesmann, J. Behley, M. Heindingsfeld, and C. Stachniss, "Spr: Single-scan radar place recognition," *IEEE Robot. and Automat. Lett.*, 2024.
- [13] G. Peng, H. Li, Y. Zhao, J. Zhang, Z. Wu, P. Zheng, and D. Wang, "Transloc4d: Transformer-based 4d radar place recognition," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recog.*, 2024, pp. 17595–17605.
- [14] M. Hilger, V. Kubelka, D. Adolfsson, R. Becker, H. Andreasson, and A. J. Lilienthal, "Introspective loop closure for slam with 4d imaging radar," *arXiv preprint arXiv:2503.02383*, 2025.
- [15] G. Yao, X. Li, L. Fu, and Y. Pan, "Monocular visual place recognition in lidar maps via cross-modal state space model and multi-view matching," *arXiv preprint arXiv:2410.06285*, 2024.
- [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, pp. 5097–5103.
- [17] T. Y. Tang, D. De Martini, and P. Newman, "Get to the point: Learning lidar place recognition and metric localisation using overhead imagery," *Proceedings of Robotics: Science and Systems*, 2021, 2021.
- [18] D. Cattaneo, M. Vaghi, S. Fontana, A. L. Ballardini, and D. G. Sorrenti, "Global visual localization in lidar-maps through shared 2d-3d embedding space," in *2020 IEEE International Conference on Robotics and Automation*. IEEE, 2020, pp. 4365–4371.
- [19] W. Xie *et al.*, "Modalink: Unifying modalities for efficient image-to-pointcloud place recognition," in *Proc. IEEE/RSJ Intl. Conf. on Intell. Robots and Sys.*, 2024, pp. 3326–3333.
- [20] Z. Zhou, J. Xu, G. Xiong, and J. Ma, "Lcpr: A multi-scale attention-based lidar-camera fusion network for place recognition," *IEEE Robotics and Automation Letters*, vol. 9, no. 2, pp. 1342–1349, 2023.
- [21] M. Jung, S. Jung, H. Gil, and A. Kim, "Helios: Heterogeneous lidar place recognition via overlap-based learning and local spherical transformer," *arXiv preprint arXiv:2501.18943*, 2025.
- [22] C. Doer and G. F. Trommer, "An ekf based approach to radar inertial odometry," in *2020 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*. IEEE, 2020, pp. 152–159.
- [23] Z. Shi, H. Shi, K. Yang, Z. Yin, Y. Lin, and K. Wang, "Panovpr: towards unified perspective-to-equirectangular visual place recognition via sliding windows across the panoramic view," in *Proc. IEEE Intell. Transport. Sys. Conf.*, 2023, pp. 1333–1340.
- [24] Z.-C. Hao and M. He, "Developing millimeter-wave planar antenna with a cosecant squared pattern," *IEEE Trans. Antennas and Prop.*, vol. 65, no. 10, pp. 5565–5570, 2017.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recog.*, 2016, pp. 770–778.
- [26] S. Izquierdo and J. Civera, "Optimal transport aggregation for visual place recognition," in *Proc. IEEE Conf. on Comput. Vision and Pattern Recog.*, 2024, pp. 17658–17668.
- [27] Y. Zhong, R. Arandjelović, and A. Zisserman, "Ghostvlad for set-based face recognition," in *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part II 14*. Springer, 2019, pp. 35–50.
- [28] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," *Advances in Neural Information Processing Sys. Conf.*, vol. 26, 2013.
- [29] P. Duhamel and M. Vetterli, "Fast fourier transforms: a tutorial review and a state of the art," *Signal processing*, vol. 19, no. 4, pp. 259–299, 1990.
- [30] X. Xu, S. Lu, J. Wu, H. Lu, Q. Zhu, Y. Liao, R. Xiong, and Y. Wang, "Ring++: Roto-translation invariant gram for global localization on a sparse scan map," *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4616–4635, 2023.
- [31] T. Cohen and M. Welling, "Group equivariant convolutional networks," in *International conference on machine learning*. PMLR, 2016, pp. 2990–2999.
- [32] J. Zhang *et al.*, "Ntu4dradm: 4d radar-centric multi-modal dataset for localization and mapping," in *Proc. IEEE Intell. Transport. Sys. Conf.*, 2023, pp. 4291–4296.
- [33] J. Huai, B. Wang, Y. Zhuang, Y. Chen, Q. Li, and Y. Han, "Snail radar: A large-scale diverse benchmark for evaluating 4d-radar-based slam," *The International Journal of Robotics Research*, p. 02783649251329048, 2024.
- [34] H. Kim *et al.*, "Hercules: Heterogeneous radar dataset in complex urban environment for multi-session radar slam," in *Proc. of the IEEE International Conference on Robotics and Automation*, 2025.
- [35] J. Komorowski, "Improving point cloud based place recognition with ranking-based loss and large batch training," in *2022 26th international conference on pattern recognition*. IEEE, 2022, pp. 3699–3705.
- [36] H. Yin, X. Xu, Y. Wang, and R. Xiong, "Radar-to-lidar: Heterogeneous place recognition via joint learning," *Frontiers in Robotics and AI*, vol. 8, p. 661199, 2021.
- [37] G. Toliás, R. Sicre, and H. Jégou, "Particular object retrieval with integral max-pooling of cnn activations," *arXiv preprint arXiv:1511.05879*, 2015.
- [38] A. Babenko and V. Lempitsky, "Aggregating deep convolutional features for image retrieval," *arXiv preprint arXiv:1510.07493*, 2015.
- [39] F. Radenović, G. Toliás, and O. Chum, "Fine-tuning cnn image retrieval with no human annotation," *IEEE Trans. Pattern Analysis and Machine Intell.*, vol. 41, no. 7, pp. 1655–1668, 2018.