

WildCross: A Cross-Modal Large Scale Benchmark for Place Recognition and Metric Depth Estimation in Natural Environments

Joshua Knights^{1,2}, Joseph Reid¹, Kaushik Roy¹, David Hall¹, Mark Cox¹, Peyman Moghadam^{1,2}

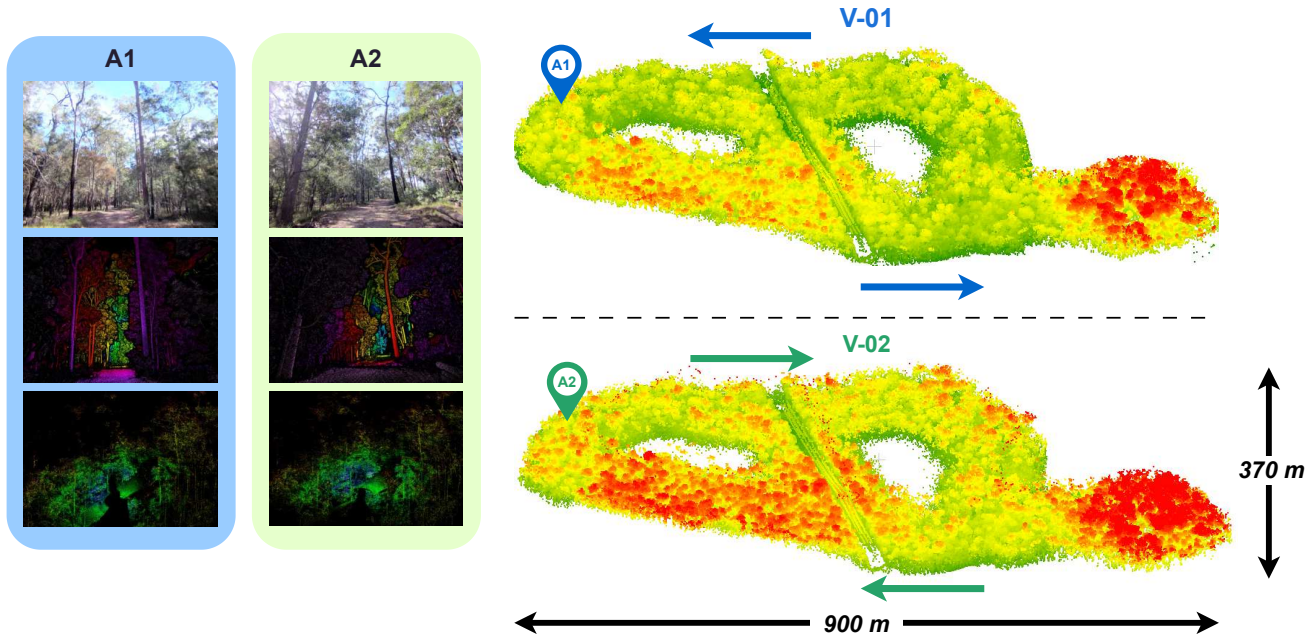


Fig. 1: The global maps of two sequences from *WildCross*. The left panels show the RGB images (top), annotated depth images (middle), and lidar submaps (bottom) at locations A1 and A2. These correspond to revisits of the same location from opposite directions across different sessions. *WildCross* presents a challenging new benchmark for cross-modal place recognition and metric depth estimation, with eight traversals covering diverse viewpoints in two large-scale forests.

Abstract—Recent years have seen a significant increase in demand for robotic solutions in unstructured natural environments, alongside growing interest in bridging 2D and 3D scene understanding. However, existing robotics datasets are predominantly captured in structured urban environments, making them inadequate for addressing the challenges posed by complex, unstructured natural settings. To address this gap, we propose *WildCross*, a cross-modal benchmark for place recognition and metric depth estimation in large-scale natural environments. *WildCross* comprises over 476K sequential RGB frames with semi-dense depth and surface normal annotations, each aligned with accurate 6DoF poses and synchronized dense lidar submaps. We conduct comprehensive experiments on visual, lidar, and cross-modal place recognition, as well as metric depth estimation, demonstrating the value of *WildCross* as a challenging benchmark for multi-modal robotic perception tasks. We provide access to the code repository and dataset at <https://csiro-robotics.github.io/WildCross>.

I. INTRODUCTION

Autonomous robots are increasingly deployed in unstructured and natural environments for applications such as agriculture, environmental monitoring, and search and res-

cue [6]–[8]. However, progress in robotic navigation and perception tasks remains heavily dependent on public datasets, given the high cost and logistical challenges of large-scale field trials. Benchmarks such as KITTI [9] and Oxford RobotCar [10] have been instrumental in advancing the field, but they are predominantly captured in structured urban or indoor settings [9]–[11]. In contrast, natural environments are characterized by irregular terrain, dense vegetation, narrow trails, and complex occlusions, rendering existing datasets insufficient for evaluating robotic autonomy in environments where it is most urgently required. Concurrently, the robotics and computer vision communities are placing increasing emphasis on bridging 2D and 3D scene understanding, exemplified by recent advances in learning-based 3D reconstruction [12]–[14] and cross-modal place recognition [15]–[18]. To support these developments, datasets must provide accurate ground truth across both 2D and 3D modalities under the added complexity of natural scenes. To this end, we present *WildCross*, a large-scale multi-modal benchmark designed to advance cross-modal place recognition and metric depth estimation in natural environments.

We derive a novel benchmark dataset from Wild-Places [3] by extending it in two key directions. First, we regenerate

¹ CSIRO Robotics, Data61, CSIRO, Australia. E-mail: firstname.lastname@csiro.au

² Queensland University of Technology (QUT), Australia.

Name	Supported Tasks				Diversity		
	VPR	LPR	CMPR	Depth Est.	ViewPoint	Temporal	Scene
Nordland [1]	✓	✗	✗	✗	*	***	***
RELLIS-3D [2]	✓	✓	✓	✗	*	*	*
Wild-Places [3]	✗	✓	✗	✗	***	***	***
Oxford Forest [4]	✗	✓	✗	✗	***	**	***
BotanicGarden [5]	✓	✓	✓	✗	**	*	*
WildCross (Ours)	✓	✓	✓	✓	***	***	***

TABLE I: Comparison of existing datasets in natural environments with support for relocalization. Datasets are compared based on the range of supported tasks and their diversity across viewpoint, temporal, and scene variations between revisits.

the original sequences with accurate camera poses, enabling large-scale training and evaluation using RGB video data alongside synchronized lidar submaps. Second, we develop an annotation pipeline that produces semi-dense depth images by combining accumulated point cloud maps with robust point visibility estimation. This enables, for the first time, reliable benchmarking of metric depth estimation in challenging natural environments alongside visual, lidar, and cross-modal place recognition. The resulting dataset comprises over 476K sequential high-resolution RGB frames with corresponding semi-dense depth and surface normal annotations, alongside accurate 6-DoF poses and extrinsics for synchronization to the lidar submaps in the Wild-Places dataset. Spanning eight traversals over 14 months and incorporating diverse viewpoints, *WildCross* establishes a challenging and versatile benchmark for visual, lidar, and cross-modal place recognition, as well as metric depth estimation in unstructured environments. In summary, the contributions of this paper are as follows:

- 1) We present *WildCross*, a large-scale multi-modal benchmark for natural environments. *WildCross* contains over 476K sequential RGB frames with corresponding semi-dense depth annotations and surface normal annotations, each with accurate 6DoF ground truth poses.
- 2) We demonstrate a new method for generating semi-dense depth and surface normal annotations for each image in our dataset through a combination of leveraging a dense accumulated point cloud map, accurate camera poses, and a robust point visibility estimation pipeline to ensure a high degree of accuracy in our depth annotations, providing a valuable benchmark for metric depth estimation in natural environments
- 3) We conduct extensive experiments on *WildCross* by evaluating several state-of-the-art methods for visual and cross-modal place recognition, as well as metric depth estimation. Our results show that leading methods struggle in these challenging natural environments, underscoring the difficulty of the *WildCross* benchmark and opening new opportunities for future research.

II. RELATED WORK

A. Place Recognition in Natural Environments

Place recognition is essential for the safe and reliable deployment of autonomous robots in complex environments, enabling both loop closure detection in simultaneous localization and mapping (SLAM), and re-localization in previously visited environments. Progress in this field is largely driven by the availability of high-quality datasets for training and

evaluation, particularly in VPR, where many recent state-of-the-art approaches [19]–[21] leverage large-scale pre-training datasets such as GSV-Cities [22], Google Landmarks v2 [23], and SF-XL [24]. More recently, there has also been growing interest in cross-modal place recognition (CMPR) [15]–[18], [25], which requires datasets with synchronized visual and lidar data for training and evaluation.

Recent state-of-the-art place recognition approaches continue to demonstrate strong generalization across urban benchmarks for VPR [26]–[28], LPR [9], [29], [30], and CMPR [9], [10]. This progress has been driven in part by the growing adoption of large-scale pre-trained foundation models [19] and task-specific pre-training strategies [21]. However, even methods that benefit from such large-scale pre-training often degrade sharply under the severe domain shift between urban and natural environments. These limitations motivate the development of new large-scale datasets that capture the complexity and variability of unstructured natural environments.

Table I provides an overview of existing datasets with support for place recognition evaluation in natural environments. Nordland [1] offers a challenging VPR benchmark through data collected from a train-mounted camera along a fixed route across multiple seasonal changes; however, it lacks lidar data and provides no viewpoint diversity (*i.e.*, reverse revisits) due to the fixed trajectory of the train. Wild-Places [3] and Oxford Forest [4] provide large-scale benchmarks for LPR in forest environments with both intra-sequence and multi-session revisits, but they lack support for VPR or CMPR due to the absence of camera data. REllIS-3D [2] includes camera and lidar recordings of several traversals of unpaved trails around a university campus, but the limited scale and diversity of its sequences restrict its utility as a place recognition benchmark. While BotanicGarden [5] provides synchronized image and lidar data from multiple traversals, its restriction to a single, small-scale botanical garden results in limited scene, viewpoint, and temporal diversity. In contrast, *WildCross* introduces sequential RGB frames synchronized to the dense lidar submaps from the Wild-Places [3] dataset for eight traversals across multiple natural environments, with long-term revisits and diverse viewpoints, establishing a challenging new benchmark for VPR and CMPR.

B. Metric Depth Estimation

Metric depth estimation has recently made significant progress, driven by the emergence of 3D foundation models [14], [31], which achieve strong generalization across diverse visual domains. While these advances have been trans-




Fig. 2: *WildCross* overview (a) RGB Image, (b) Depth Image, (c) Depth Overlay, (d) Surface Normal, (e) Lidar Submap.

formative for computer vision, their impact is particularly critical in robotics, where reliable depth estimation underpins core capabilities such as SLAM [32], [33], cross-modal place recognition [18], and 3D structure-from-motion [12]–[14]. Despite this progress, a key challenge in training metric depth estimation models is acquiring accurate ground-truth annotations in real-world environments. Synthetic datasets such as VirtualKITTI [34] can generate dense pixel-wise annotations across diverse conditions, yet models trained on them often struggle with the sim-to-real domain gap [31]. Outdoor benchmarks such as KITTI Depth [35] alleviate this issue by producing semi-dense depth maps through lidar accumulation, but these datasets only represent urban scenes.

In contrast, *WildCross* provides semi-dense sequential depth annotations in unstructured natural environments. By leveraging accumulated global point cloud maps and robust visibility estimation to remove occluded points (see Section III-D), *WildCross* enables training and evaluation of metric depth estimation under the complex conditions of natural scenes, where irregular terrain and dense vegetation continue to challenge state-of-the-art models.

III. WILDCROSS BENCHMARK

WildCross leverages the raw data from the Wild-Places [3] LPR dataset and extends it into a cross-modal benchmark for place recognition and metric depth estimation through two main advances, complementary to WildScenes [36], which focuses on 2D and 3D semantic segmentation in the same natural environments. Firstly, we reprocess the original traversals to produce sequential RGB frames at 15Hz with accurate 6DoF ground truth poses synchronized with dense 3D lidar submaps in the same environment. Secondly, we introduce an annotation pipeline that generates semi-dense metric depth maps with surface normals for every RGB frame through leveraging the accumulated global point cloud and 6DoF camera poses for each image, using point visibility estimation to remove occluded points from each annotated frame. These contributions make *WildCross* a powerful benchmark for evaluating performance on the tasks of visual and cross-modal place recognition, in addition to metric depth estimation in complex unstructured natural

environments. Section III-A describes the sequences, while Sections III-B, III-C, and III-D cover the generation of RGB and lidar data, the pose estimation, and the creation of semi-dense depth and surface normal annotations, respectively.

A. Sequence Information

For consistency with Wild-Places [3] and WildScenes [36], we adopt the notation V-XX and K-XX to denote sequence XX on the Venman and Karawatha trajectories, respectively. Table II summarizes the per-sequence and overall statistics of the dataset, including the number of submaps, image frames, intra-sequence revisits, and total traversal distance for each sequence. The traversals follow a consistent pattern: in both environments, Sequence 02 corresponds to the reverse trajectory of Sequence 01, Sequence 03 follows an alternate extended route, and Sequence 04 repeats the route of Sequence 01. As we show in Section V-A, state-of-the-art VPR approaches continue to face significant challenges in successfully relocalizing under reverse revisits. The repeated traversals in *WildCross*, including reverse and alternate trajectories with substantial overlap, therefore provide a basis for evaluating both intra- and inter-sequence re-localization under challenging revisit conditions.

B. Color Images and Submap Generation


To obtain our color images, we extract camera frames from the raw video output of the forward-facing camera of the sensor payload at 15Hz. These are then rectified using the distortion parameters obtained after sensor calibration. To generate our lidar submaps, we follow the approach of [3] and generate a global accumulated point cloud map with corresponding sensor trajectory calculated using a lidar-inertial SLAM [37]. We then use the SLAM trajectory to generate 3D submaps from the dense global map, sampling all points within a 30m radius and a 1s time window around the position of the sensor payload every 0.5s along the global trajectory.

C. Ground Truth Poses

For each image frame \mathcal{I} in *WildCross* we provide ground-truth 6DoF poses in the world frame, represented as $T(t) = \{q(t), x(t)\}$, where $q(t) \in \mathbb{R}^4$ denotes rotation as a unit quaternion and $x(t) \in \mathbb{R}^3$ denotes the translation of the sensor origin at time t . To generate the poses, we interpolate

Sequence	Distance	Camera / Depth Images			Submaps			
		All	Revisits	%	All	Revisits	%	
Venman	01	2.64km	35.3K	10.6K	30.03	4.7K	1.2K	25.69
	02	2.64km	34.1K	8.3K	24.25	4.6K	980	21.51
	03	4.59km	63.8K	8.7K	13.70	8.5K	906	10.70
	04	2.81km	43.1K	8.0K	18.57	5.7K	937	16.33
Karawatha	01	5.14km	66.1K	5.6K	8.43	8.8K	553	6.27
	02	5.66km	75.6K	6.6K	8.70	10.0K	522	5.18
	03	6.27km	114.2K	38.4K	33.64	15.2K	2.5K	16.73
	04	3.17km	43.8K	9.3K	21.35	5.8K	1.0K	18.59
Total	33km	476K	95.5K	-	63.3K	8.7K	-	

TABLE II: Sequence statistics for *WildCross*. “Revisits” denote intra-sequence loop closures.



(a) RGB Image (b) Noisy Depth (c) Ours

Fig. 3: Impact of visibility estimation. (a) RGB Image, (b) Naïve projection of global 3D points produces noisy depth maps with occluded points. (c) Our visibility pipeline removes these, yielding higher-quality depth.

between poses in the SLAM trajectory by applying spherical linear interpolation for rotation and linear interpolation for translation:

$$T(t) = \left\{ q(t) = q_2 (q_2^{-1} \cdot q_1)^u, x(t) = (1 - u)x_1 + ux_2 \right\}, \quad (1)$$

$$u = \frac{t - t_1}{t_2 - t_1}, \quad (2)$$

where (q_1, x_1, t_1) and (q_2, x_2, t_2) are the SLAM poses immediately before and after timestamp t . Poses for the camera images are then transformed into camera coordinate frames using the extrinsic transform between the SLAM and camera frames, and the poses for multiple sessions in the same environment are aligned by using ICP to align the global point cloud maps produced by the SLAM for each sequence.

D. Metric Depth and Surface Normal Generation

We provide sequential, semi-dense metric depth and surface normal annotations for each camera frame. Each pixel in the depth annotation, denoted as $D(u, v)$ at image coordinate (u, v) , represents the distance to the nearest object intersected by the ray originating from the camera center and passing through pixel (u, v) in 3D space. Likewise, the surface normal $SN(u, v)$ at image coordinate (u, v) represents the surface normal vector at the same point with respect to the camera coordinate frame. Depth images are generated by projecting 3D points from the global point cloud onto the image plane using the camera projection function and the extrinsic transform relating the camera and

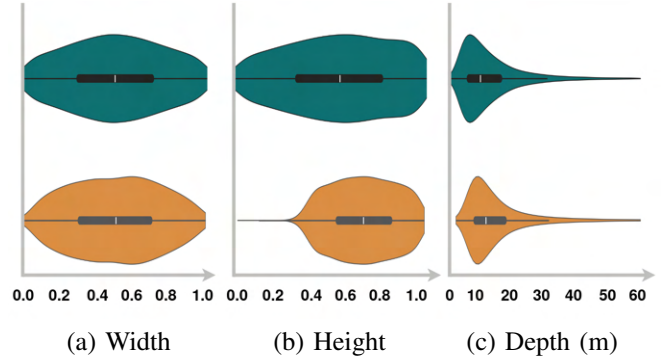


Fig. 4: Depth distribution for *WildCross* (●) vs. KITTI Annotated Depth (●) [35]. Violin plots are computed from 1% subsamples of both datasets. Width and height are normalized with respect to image sizes.

lidar coordinate frames. Importantly, not all 3D points in the point cloud are visible in the captured image (see Figure 3). Determining visibility is challenging because each 3D point has infinitesimal volume, making it unlikely for multiple points to lie along the same ray passing through the camera center. To identify the visible points in an image, we compute the surface normal of each point in the point cloud and select the points with a normal oriented toward the camera. Points that project outside the image bounds or lie behind the camera are also discarded. Finally, we eliminate points belonging to surfaces occluded by other surfaces. The surface normal of each point $p_i \in \mathbb{R}^3$ is estimated by computing the eigendecomposition of its local neighborhood, defined as all points within 0.5m of p_i . The eigenvector corresponding to the smallest eigenvalue is selected as the surface normal, after ensuring it is oriented towards the observation location of the point p_i . These surface normals, oriented in the camera frame, are also used to generate the surface normal image after filtering of points for occlusions, as is done for depth.

To address the problem of occlusion, we utilize the generalized hidden point removal (GHPR) operator proposed in [38], [39]. The operator consists of two steps: i) apply a spherical reflection to the filtered point cloud, which has the camera position as its origin, and ii) calculate the convex hull of the spherical reflected point cloud. The spherical reflection of a 3D point $p \in \mathbb{R}^3$ is defined by the function $F(p; \gamma)$:

$$F(p; \gamma) = \begin{cases} p \|p\|_2^{-1} f(\|p\|_2; \gamma) & \|p\|_2 \neq 0 \\ 0 & \|p\|_2 = 0, \end{cases} \quad (3)$$

where the kernel function $f(d; \gamma)$ is a monotonically decreasing function of the distance $d > 0$. In this work, we use the exponential inversion kernel for its scale-invariant properties:

$$f(d; \gamma) = d^\gamma \quad \gamma < 0. \quad (4)$$

The spherical reflection function F has the property that a point which is close to the camera is transformed to a location that is far away from the camera, and vice versa. Therefore, we can determine which points are observed in the camera image by selecting those that lie on the convex hull of the spherically reflected point cloud. We perform

Split	V/K-01	V/K-02	V/K-03	V/K-04	Lidar		Camera	
					Train	Test	Train	Test
01	Test	Train	Train	Train	49.8K	13.5K	374.6K	101.4K
02	Train	Test	Train	Train	48.7K	14.6K	366.3K	109.7K
03	Train	Train	Test	Train	39.7K	23.6K	298.0K	177.9K
04	Train	Train	Train	Test	51.8K	11.5K	389.1K	86.9K

TABLE III: Train/test cross-fold splits for *WildCross*.

a statistical analysis comparing the depth data provided by *WildCross* to that of the annotated depth maps in the KITTI dataset [35], which were produced through the projection of accumulated lidar point clouds but in a structured urban environment. We present the distribution of non-zero pixels from depth images as violin plots shown in Fig 4. As the datasets have different image sizes, the distributions of pixels across the image coordinates are provided in terms of relative coordinates within the image, while depth values are given in meters. The most significant difference in the distribution between the datasets can be seen in the distribution of points along the height axis of the images, where the top portion of the image is unlabeled in the KITTI dataset due to the limited vertical Field-of-View (FoV) of its lidar, whereas our depth images provides a much denser depth along the height axis of the images which is crucial in natural environment.

IV. EXPERIMENTS

A. Training and Testing Splits

We observed that the training splits introduced for LPR in Wild-Places [3] are not well suited to training large-scale VPR and CMPR networks. To address this, we propose a cross-fold training and evaluation setup for benchmarking VPR and CMPR performance on *WildCross*. In this setup, training and evaluation follow a four-fold cross-split design. In each split, the sequences with the same index (e.g., Split-1, V-01, and K-01) from both environments are held out together for evaluation, while the remaining sequences are used for training. Table III reports the number of training and testing samples in each split. During evaluation, the held-out sequences are used for intra-sequence place recognition and also serve as queries for inter-sequence recognition, with the training sequences acting as the database.

B. Visual Place Recognition (VPR)

We evaluate four state-of-the-art methods for VPR: NetVLAD [40], MixVPR [41], DINOv2-SALAD (SALAD) [19], and Bag-of-Queries (BoQ) [20]. Unlike LPR, positive training pairs for VPR cannot be formed solely using a distance threshold. The limited field-of-view of the camera and the presence of reverse revisits within and across sequences can result in false positives, where two images selected as a pair share little or no visual overlap. To mitigate this, and inspired by [24], we define positive training pairs in *WildCross* as images whose camera poses are within 5m distance and 15° bearing of each other, and negative pairs are those separated by more than 50m. At evaluation, a retrieved image is considered a correct match if its pose lies within 25m of the query. We report results under two evaluation settings: zero-shot and fine-tuned. In the **zero-shot** setting, each method is

evaluated on *WildCross* using its best released pretrained model, without any additional training on *WildCross*. This setup reflects the scenario where models pretrained on large-scale, predominantly urban, datasets are applied directly to natural environments, thereby measuring their out-of-domain generalization ability.

In the **fine-tuned** setting, the same methods are further fine-tuned on the *WildCross* training splits before evaluation. This measures their in-domain performance once exposed to data from natural environments. Reporting both settings highlights the gap between cross-domain generalization (urban-to-natural) and in-domain adaptation, providing a comprehensive benchmark of VPR in natural environments.

C. Cross-Modal Place Recognition (CMR)

Cross-modal place recognition (CMPR) aims to localize across different sensing modalities, such as retrieving lidar submaps given visual queries. This task is particularly challenging in natural environments, where structural complexity and viewpoint variation exacerbate the difficulty of aligning cross-modal features. To explore this task on *WildCross*, we evaluate LIP-Loc [17], a baseline learning-based CMPR method. Within our cross-fold training and testing regime, we evaluate the inter-sequence CMPR performance using the images from the unseen sequences as queries and lidar submaps from all sequences in an environment as databases. When trained with the original batched contrastive loss from [17] we observed a collapse in the feature space, which we attribute to the lack of any mechanisms in the LIP-Loc loss to deal with false negative examples in the batch (i.e., negative samples with high real-world similarity). To address this limitation, we replace the contrastive formulation with a modified cross-entropy loss as follows:

$$\mathcal{L}_{\text{LIP-Loc}} = \mathcal{L}_{\text{CE}} \left(\frac{\exp(v_i^\top v_i^+)}{\exp(v_i^\top v_i^+) + \sum_{j \notin nn_i} \exp(v_i^\top v_j)} \right), \quad (5)$$

where v_i, v_i^+, v_j represent global feature vectors for a training sample and its corresponding positive and negative samples in the batch respectively, nn_i denotes the set of non-negatives in the batch for v_i (i.e., samples located within 50m in the real world), and $\mathcal{L}_{\text{CE}}(x) = -x \log(x)$ ($x = -x \log(x)$) is the cross-entropy loss. We use a positive threshold of 25m during evaluation. While the original LIP-Loc framework employed ResNet50 on ImageNet-pretrained visual transformers as its backbone, we extend this setup by incorporating more powerful recent pretraining approaches. Specifically, we evaluate DINOv2 [42] and DINOv3 [43], using the ViT-S encoder and extracting the class token, followed by a fully connected layer, as the global feature representation. Note, we report the best configuration results found experimentally for each method, with DINOv2 having a frozen backbone and DINOv3 having an unfrozen backbone when fine-tuning.

D. Lidar Place Recognition (LPR)

To facilitate comparison with purely 3D place recognition approaches, we evaluate three learning-based LPR methods:

Method	V-01		V-02		V-03		V-04		K-01		K-02		K-03		K-04		Average		
	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	
Zero-shot	NetVLAD [40]	47.24	51.91	52.44	64.70	7.040	14.65	38.98	47.91	47.09	59.10	55.70	73.86	13.76	23.34	51.98	55.67	39.28	48.89
	MixVPR [41]	51.86	57.39	44.60	50.63	7.33	13.80	44.60	50.84	74.89	81.88	69.53	85.18	24.88	32.34	57.12	59.58	46.85	53.96
	SALAD [19]	50.64	57.01	44.60	49.06	13.05	19.58	43.72	48.53	78.48	83.50	54.71	68.16	26.67	32.78	54.93	59.31	45.85	52.24
	BoQ [20]	54.65	62.48	48.22	55.88	11.62	17.69	46.66	52.47	80.99	84.22	45.06	55.62	26.61	31.36	55.41	58.19	46.15	52.24
Fine-tuned	NetVLAD [40]	68.95	72.16	69.89	74.05	18.09	22.72	59.90	66.02	83.86	87.53	86.32	91.57	26.82	32.35	56.32	59.96	58.77	63.30
	MixVPR [41]	66.16	68.52	68.86	72.78	17.80	25.01	52.90	57.34	84.66	87.89	87.99	89.59	30.40	38.22	60.87	65.04	58.71	63.05
	SALAD [19]	69.04	72.68	72.24	80.33	23.30	29.48	58.96	64.65	86.37	89.33	88.98	91.26	34.05	38.83	61.13	62.53	61.76	66.14
	BoQ [20]	70.22	74.61	72.84	77.67	28.68	38.87	62.40	68.71	87.62	89.51	90.81	92.93	32.26	38.05	60.49	62.90	63.17	67.91

TABLE IV: Intra-sequence VPR results on *WildCross* for zero-shot and fine-tuned networks.

Method (Backbone)	Venman		Karawatha		Average		
	R1	R5	R1	R5	R1	R5	
Zero-Shot	NetVLAD [40]	25.86	43.94	16.15	29.73	21.00	36.84
	MixVPR [41]	54.10	61.41	35.73	44.31	44.92	52.86
	SALAD [19]	57.49	64.49	41.27	50.14	49.38	57.32
	BoQ [20]	61.62	67.98	45.89	54.98	53.76	61.48
Fine-tuned	NetVLAD [40]	64.31	67.49	46.94	52.43	55.63	59.96
	MixVPR [41]	65.30	68.58	50.24	55.80	57.77	62.19
	SALAD [19]	68.54	71.86	54.29	59.86	61.41	65.86
	BoQ [20]	68.66	72.01	55.07	60.37	61.87	66.19

TABLE V: Inter-Sequence VPR Results on *WildCross* for zero-shot and fine-tuned networks.

MinkLoc3Dv2 [44], LoGG3D-Net [45], and HOTFormer-Loc [46]. For training, we define positive and negative pairs using distance thresholds of 3m and 20m, respectively. At evaluation, a retrieved submap is considered a correct match if its pose is within 3m of the query.

E. Metric Depth Estimation

WildCross also supports training and evaluation for metric depth estimation. We evaluate this task using DepthAnythingV2 [31] as a representative state-of-the-art baseline. For this experiment, sequences V-01 and K-01 are held out for testing, K-02 is used for validation, and the remaining sequences are used for training. We report three common metrics: threshold accuracy (δ_1), which measures the percentage of the predicted pixels whose depth differs from ground truth by no more than 25%, Absolute Relative Error (AbsRel), which quantifies the average relative difference between predicted and true depths, and Root Mean Square Error (RMSE), which measures the overall deviation of predictions from the ground truth. As with VPR, we report results under both zero-shot and fine-tuned settings. In the zero-shot setting, we directly evaluate the released model trained on KITTI [9] and VirtualKitti [34], thereby assessing Out-Of-Domain (OOD) generalization from urban to natural environments. In the fine-tuned setting, we adapt the model to *WildCross* depth images, measuring its in-domain performance.

V. RESULTS

A. Visual Place Recognition

Tables IV and V summarize intra- and inter-sequence VPR performance on *WildCross*. Fine-tuning consistently improves performance across all methods, showcasing the value of training on in-domain natural environment data. Nevertheless, even the strongest overall method, BoQ [20], achieves only 64.45% R1 for intra-sequence and 61.44% for inter-sequence evaluation. In comparison, on established urban benchmarks such as Pittsburgh [26] and MSLS [28],

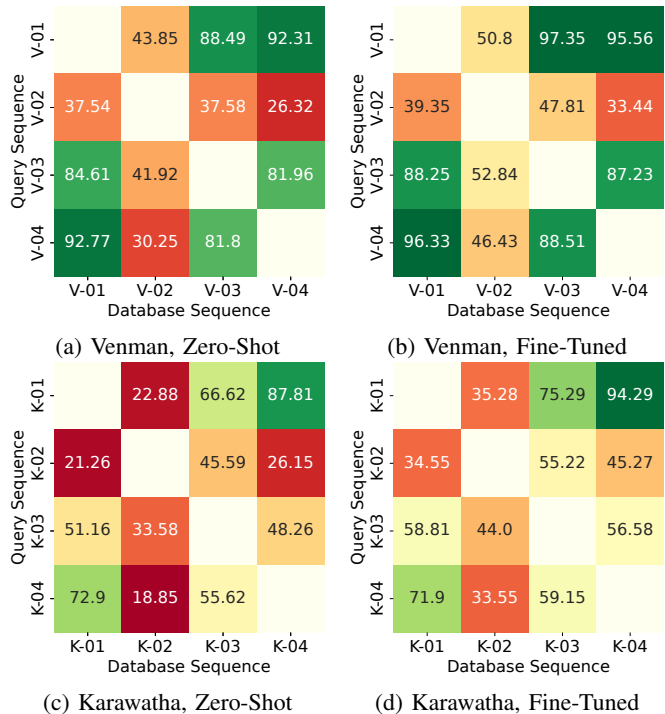


Fig. 5: Cross-sequence VPR R1 on *WildCross*. Reverse revisit sequences V-02 and K-02 significantly degrade performance, underscoring the challenge of viewpoint diversity.

the same method exceeds 90% R1.

One of the primary challenges in our benchmark derives from the prevalence of reverse revisits in both intra- and inter-sequence evaluation. As shown in Figure 5, inter-sequence Recall@1 drops substantially when K-02 or V-02 appear as either query or database sequences; these correspond to the reverse trajectories described in Section III-A. Similarly, Table IV shows that the sequences with the highest number of intra-sequence reverse revisits (V-03 and K-03) also yield the lowest performance, with the best method achieving only 28.68% and 36.57% R1, respectively. These results highlight that unstructured natural environments, particularly reverse revisits, remain a persistent challenge for state-of-the-art VPR methods. This highlights the value of *WildCross* as a benchmark for advancing visual place recognition under these difficult and underexplored conditions.

B. Lidar Place Recognition

Tables VI and VII present intra- and inter-sequence LPR results on the new cross-fold splits introduced in *WildCross*. For intra-sequence evaluation, state-of-the-art methods achieve strong performance, with R1 scores above 90% on

Method	V-01		V-02		V-03		V-04		K-01		K-02		K-03		K-04		Average	
	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5	R1	R5
MinkLoc3Dv2 [44]	96.53	100.0	96.02	100.0	59.27	93.27	89.11	99.04	71.97	98.73	98.28	100.00	57.80	82.68	86.84	99.35	81.98	96.64
LoGG3D-Net [45]	94.54	99.75	94.08	99.69	83.22	97.57	86.34	95.09	98.55	99.82	95.02	98.85	57.28	93.66	63.30	92.40	84.04	97.10
HOTFormerLoc [46]	97.35	99.67	96.84	99.90	50.88	84.55	92.85	99.47	97.29	99.64	97.32	100.00	56.50	68.58	95.46	99.54	85.56	93.92

TABLE VI: Intra-sequence LPR results on *WildCross*.

Method	Venman		Karawatha		Average	
	R1	R5	R1	R5	R1	R5
MinLoc3Dv2 [44]	90.67	98.88	78.58	92.33	84.62	95.61
LoGG3D-Net [45]	84.39	94.19	72.09	86.02	78.24	90.11
HOTFormerLoc [46]	91.60	99.09	78.11	91.62	84.85	95.36

TABLE VII: Inter-Sequence LPR Results on *WildCross*.

Method (Backbone)	Venman		Karawatha		Average	
	R1	R5	R1	R5	R1	R5
LIP-Loc (ResNet50)	40.16	54.45	34.25	48.91	37.20	51.68
LIP-Loc* (DINOv2-s)	52.55	62.71	45.26	57.40	48.90	60.06
LIP-Loc* (DINOv3-s)	56.54	63.19	48.16	57.06	52.35	60.12

TABLE VIII: Cross-Modal Place Recognition Results on *WildCross*. * ViT-S for the pretrained model backbone.

average. This indicates that LPR methods can effectively handle revisits within the same sequence in natural environments. In contrast, inter-sequence evaluation remains more challenging, with all LPR methods achieving average R1 scores below 86%. This shows that long-term and multi-session generalization in unstructured environments remains challenging.

C. Cross-Modal Place Recognition

Table VIII reports CMPR results on *WildCross*. Performance across all configurations remains limited, with the best result obtained by LIP-Loc using DINOv3 pretraining, which achieves an average R1 score of 51.42%. This indicates that cross-modal retrieval in unstructured natural environments is particularly challenging.

We also observe that the choice of backbone has a substantial effect on performance. Substituting the ImageNet-pretrained ResNet50 with transformer-based backbones yields consistent improvements, with DINOv3 providing an approximately 15% increase in average R1 compared to the ResNet50 baseline. While this trend reflects the benefits of stronger visual pretraining, the overall performance remains considerably below that observed in VPR and LPR tasks. Progress will likely require approaches that explicitly address the domain gap between 2D image features and 3D structural representations, rather than relying on backbone improvements alone. By providing large-scale cross-modal data together with depth ground truth, *WildCross* offers the basis for training and evaluating more complex models that can better exploit cross-modal information for place recognition in natural environments.

D. Metric Depth Estimation

Table IX reports zero-shot and fine-tuned metric depth prediction results on *WildCross*. Fine-tuning with our depth annotations consistently improves the performance of DepthAnythingV2 [31] across all backbones, with larger ViT backbones yielding stronger results. In the zero-shot setting, however, larger models perform worse, with RMSE

		Method (Backbone)	$\delta_1 \uparrow$	AbsRel \downarrow	RMSE \downarrow
Zero-Shot		DepthAnythingV2 (ViT-S)	0.284	0.558	7.651
		DepthAnythingV2 (ViT-B)	0.222	0.769	7.915
		DepthAnythingV2 (ViT-L)	0.074	1.478	13.734
Fine-tuned		DepthAnythingV2 (ViT-S)	0.746	0.172	3.412
		DepthAnythingV2 (ViT-B)	0.766	0.167	3.289
		DepthAnythingV2 (ViT-L)	0.789	0.157	3.150

TABLE IX: Zero-Shot and Fine-tuned DepthAnythingV2 [31] results on *WildCross*.

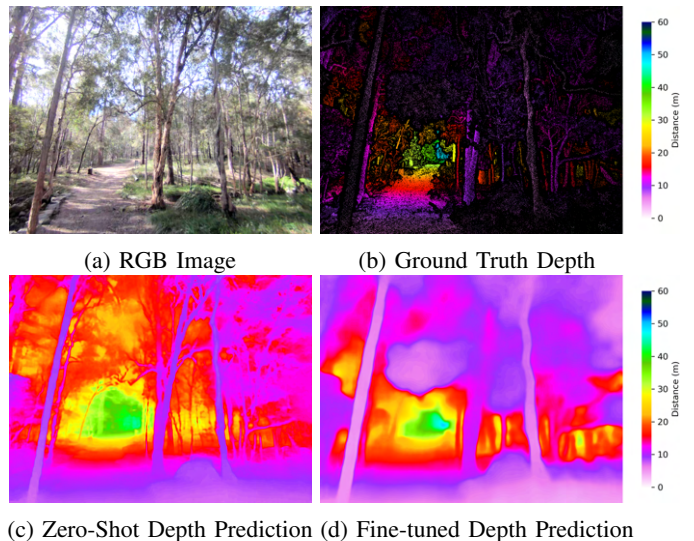


Fig. 6: Zero-shot vs. fine-tuned depth predictions. Fine-tuning improves scale but reduces detail.

increasing by more than 5m from ViT-S to ViT-L. Qualitative results in Figure 6 show that fine-tuning improves the overall scale of the predictions but reduces fine-grained detail compared to the pretrained model. Adapting depth prediction methods to the high-frequency structure of natural environments, characterized by foliage, irregular terrain, and fine textures, remains a key challenge compared to urban scenes dominated by walls and planar surfaces.

Beyond spatial accuracy, temporal consistency is critical for robotics, as frame-to-frame flickering undermines reliable perception and navigation. To the best of our knowledge, *WildCross* is the first large-scale benchmark with sequential depth ground truth in natural environments, providing a foundation for systematic evaluation and advancement of temporally consistent metric depth estimation for robotic systems.

VI. CONCLUSION

In this paper, we introduced *WildCross*, a large-scale benchmark for cross-modal place recognition and metric depth estimation in natural environments. The dataset comprises over 476K sequential RGB frames with semi-dense depth and surface normal annotations, each aligned with accurate 6DoF poses and synchronized dense lidar

submaps. Alongside describing an annotation pipeline for generating sequential depth ground truth, we evaluated a range of state-of-the-art methods for visual, lidar, and cross-modal place recognition, as well as metric depth estimation, and illustrated that even leading approaches struggle under these conditions. By capturing the complexity of natural environments, dense vegetation, irregular terrain, and diverse viewpoints, *WildCross* highlights open research gaps where current methods struggle. In particular, reverse revisits in place recognition and temporally consistent depth estimation remain unresolved problems that are central to robust robotic autonomy. We hope *WildCross* inspires future research into bridging 2D and 3D perception and developing methods capable of reliable operation in complex natural environments.

ACKNOWLEDGMENT

The authors would like to acknowledge support from the CRC-P Round 16 in partnership with Emesent.

REFERENCES

- [1] N. Sünderhauf, P. Neubert, and P. Protzel, “Are we there yet? Challenging SeqSLAM on a 3000 km journey across all four seasons,” in *IEEE Int. Conf. Robot. Autom.*, 2013.
- [2] P. Jiang, P. Osteen *et al.*, “RELLIS-3D Dataset: Data, Benchmarks and Analysis,” in *IEEE Int. Conf. Robot. Autom.*, 2021, pp. 1110–1116.
- [3] J. Knights, K. Vidanapathirana *et al.*, “Wild-Places: A Large-Scale Dataset for Lidar Place Recognition in Unstructured Natural Environments,” in *IEEE Int. Conf. Robot. Autom.*, 2023, pp. 11 322–11 328.
- [4] H. Oh, N. Chebrolu *et al.*, “Evaluation and Deployment of LiDAR-based Place Recognition in Dense Forests,” in *IEEE/RSJ Int. Conf. Intell. Robot. Syst.*, 2024, pp. 12 824–12 831.
- [5] Y. Liu, Y. Fu *et al.*, “BotanicGarden: A High-Quality Dataset for Robot Navigation in Unstructured Natural Environments,” *IEEE Robot. Autom. Lett.*, vol. 9, no. 3, pp. 2798–2805, 2024.
- [6] L. F. Oliveira, A. P. Moreira, and M. F. Silva, “Advances in Agriculture Robotics: A State-of-the-Art Review and Challenges Ahead,” *Robotics*, vol. 10, no. 2, p. 52, 2021.
- [7] M. V. Malladi, N. Chebrolu *et al.*, “DigiForests: A Longitudinal LiDAR Dataset for Forestry Robotics,” in *IEEE Int. Conf. Robot. Autom.*, 2025, pp. 1459–1466.
- [8] Y. Blei, M. Krawez *et al.*, “CloudTrack: Scalable UAV tracking with cloud semantics,” in *IEEE Int. Conf. Robot. Autom.*, 2025, pp. 15 893–15 899.
- [9] A. Geiger, P. Lenz *et al.*, “Vision meets Robotics: The KITTI Dataset,” *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [10] W. Maddern, G. Pascoe *et al.*, “1 Year, 1000km: The Oxford RobotCar Dataset,” *Int. J. Robot. Res.*, vol. 36, no. 1, pp. 3–15, 2017.
- [11] N. Silberman, D. Hoiem *et al.*, “Indoor Segmentation and Support Inference from RGBD Images,” in *Eur. Conf. Comput. Vis.*, 2012, pp. 746–760.
- [12] S. Wang, V. Leroy *et al.*, “DUS3R: Geometric 3D Vision Made Easy,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 20 697–20 709.
- [13] V. Leroy, Y. Cabon, and J. Revaud, “Grounding Image Matching in 3D with MAST3R,” in *Eur. Conf. Comput. Vis.*, 2024, pp. 71–91.
- [14] J. Wang, M. Chen *et al.*, “VGGT: Visual Geometry Grounded Transformer,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2025, pp. 5294–5306.
- [15] X. Cai, Y. Wang *et al.*, “VOloc: Visual Place Recognition by Querying Compressed Lidar Map,” in *IEEE Int. Conf. Robot. Autom.*, 2024, pp. 10 192–10 199.
- [16] Z. Zhao, H. Yu *et al.*, “Attention-Enhanced Cross-Modal Localization Between Spherical Images and Point Clouds,” *IEEE Sensors Journal*, vol. 23, no. 19, pp. 23 836–23 845, 2023.
- [17] S. Shubodh, M. Omama *et al.*, “Lip-loc: Lidar image pretraining for cross-modal localization,” in *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, 2024, pp. 948–957.
- [18] H. Xu, H. Liu *et al.*, “C2L-PR: Cross-Modal Camera-to-LiDAR Place Recognition via Modality Alignment and Orientation Voting,” *IEEE Transactions on Intelligent Vehicles*, 2024.
- [19] S. Izquierdo and J. Civera, “Optimal transport aggregation for visual place recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.
- [20] A. Ali-Bey, B. Chaib-draa, and P. Giguere, “BoQ: A Place is Worth a Bag of Learnable Queries,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 17 794–17 803.
- [21] S. Hausler and P. Moghadam, “Pair-vpr: Place-aware pre-training and contrastive pair classification for visual place recognition with vision transformers,” *IEEE Robot. Autom. Lett.*, vol. 10, no. 4, pp. 4013–4020, 2025.
- [22] A. Ali-bey, B. Chaib-draa, and P. Giguere, “GSV-Cities: Toward appropriate supervised visual place recognition,” *Neurocomputing*, vol. 513, pp. 194–203, 2022.
- [23] T. Weyand, A. Araujo *et al.*, “Google Landmarks Dataset v2: A Large-Scale Benchmark for Instance-Level Recognition and Retrieval,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 2575–2584.
- [24] G. Berton, C. Masone, and B. Caputo, “Rethinking Visual Geolocalization for Large-Scale Applications,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022, pp. 4878–4888.
- [25] J. Knights, S. B. Laina *et al.*, “SOLVR: Submap oriented lidar-visual re-localisation,” in *IEEE Int. Conf. Robot. Autom.*, 2025, pp. 6387–6393.
- [26] A. Torii, J. Sivic *et al.*, “Visual Place Recognition with Repetitive Structures,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2013, pp. 883–890.
- [27] A. Torii, R. Arandjelovic *et al.*, “24/7 place recognition by view synthesis,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2015, pp. 1808–1817.
- [28] F. Warburg, S. Hauberg *et al.*, “Mapillary Street-Level Sequences: A Dataset for Lifelong Place Recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020, pp. 2626–2635.
- [29] G. Kim, Y. S. Park *et al.*, “Mulran: Multimodal range dataset for urban place recognition,” in *IEEE Int. Conf. Robot. Autom.*, 2020, pp. 6246–6253.
- [30] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, “University of Michigan North Campus long-term vision and lidar dataset,” *Int. J. Robot. Res.*, vol. 35, no. 9, pp. 1023–1035, 2016.
- [31] L. Yang, B. Kang *et al.*, “Depth Anything V2,” *Adv. Neural Inform. Process. Syst.*, vol. 37, pp. 21 875–21 911, 2024.
- [32] S. Zhang, L. Zheng, and W. Tao, “Survey and Evaluation of RGB-D SLAM,” *IEEE Access*, vol. 9, pp. 21 367–21 387, 2021.
- [33] N. Keetha, J. Karhade *et al.*, “SplaTAM: Splat, Track & Map 3D Gaussians for Dense RGB-D SLAM,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024, pp. 21 357–21 366.
- [34] A. Gaidon, Q. Wang *et al.*, “Virtual Worlds as Proxy for Multi-Object Tracking Analysis,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 4340–4349.
- [35] J. Uhrig, N. Schneider *et al.*, “Sparsity Invariant CNNs,” in *International Conference on 3D Vision (3DV)*, 2017.
- [36] K. Vidanapathirana, J. Knights *et al.*, “WildScenes: A benchmark for 2D and 3D semantic segmentation in large-scale natural environments,” *Int. J. Robot. Res.*, vol. 44, no. 4, pp. 532–549, 2025.
- [37] M. Ramezani, K. Khosoussi *et al.*, “Wildcat: Online continuous-time 3d lidar-inertial slam,” *arXiv preprint arXiv:2205.12595*, 2022.
- [38] S. Katz and A. Tal, “On the Visibility of Point Clouds,” in *IEEE Int. Conf. Comput. Vis.*, 2015, pp. 1350–1358.
- [39] P. Vechevsky, M. Cox *et al.*, “Colourising point clouds using independent cameras,” *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3575–3582, 2018.
- [40] R. Arandjelovic, P. Gronat *et al.*, “NetVLAD: CNN architecture for weakly supervised place recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016, pp. 5297–5307.
- [41] A. Ali-Bey, B. Chaib-Draa, and P. Giguere, “MixVPR: Feature Mixing for Visual Place Recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023, pp. 2998–3007.
- [42] M. Oquab, T. Darcet *et al.*, “DINOv2: Learning Robust Visual Features without Supervision,” *arXiv preprint arXiv:2304.07193*, 2023.
- [43] O. Siméoni, H. V. Vo *et al.*, “DINOv3,” *arXiv preprint arXiv:2508.10104*, 2025.
- [44] J. Komorowski, “Improving Point Cloud Based Place Recognition with Ranking-based Loss and Large Batch Training,” in *Int. Conf. Pattern Recog.*, 2022, pp. 3699–3705.
- [45] K. Vidanapathirana, M. Ramezani *et al.*, “LoGG3D-Net: Locally Guided Global Descriptor Learning for 3D Place Recognition,” in *IEEE Int. Conf. Robot. Autom.*, 2022, pp. 2215–2221.
- [46] E. Griffiths, M. Haghighat *et al.*, “HOTFormerLoc: Hierarchical Oc-tree Transformer for Versatile Lidar Place Recognition Across Ground and Aerial Views,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2025, pp. 6648–6658.