

Unifying Local and Global Multimodal Features for Place Recognition in Aliased and Low-Texture Environments

Alberto García-Hernández^{1,2}, Riccardo Giubilato¹, Klaus H. Strobl¹, Javier Civera², and Rudolph Triebel¹

Abstract—Perceptual aliasing and weak textures pose significant challenges to the task of place recognition, hindering the performance of Simultaneous Localization and Mapping (SLAM) systems. This paper presents a novel model, called UMF (standing for Unifying Local and Global Multimodal Features) that 1) leverages multi-modality by cross-attention blocks between vision and LiDAR features, and 2) includes a re-ranking stage that re-orders based on local feature matching the top-k candidates retrieved using a global representation. Our experiments, particularly on sequences captured on a planetary-analogous environment, show that UMF outperforms significantly previous baselines in those challenging aliased environments. Since our work aims to enhance the reliability of SLAM in *all* situations, we also explore its performance on the widely used RobotCar dataset, for broader applicability. Code and models are available at <https://github.com/DLR-RM/UMF>.

I. INTRODUCTION


Simultaneous Localization and Mapping (SLAM) has emerged as a central technology in a multitude of industries including autonomous driving [13], [20], automated construction [33], and agriculture [37], [38]. Its development and adoption have been accelerated by advancements in sensor technologies, including multi-camera setups, RGB-D sensors, and more recently, 3D Light Detection and Ranging (LiDAR) sensors, facilitating the construction of large-scale, dense, high-resolution, and consistent 3D maps.

Many SLAM benchmarks, such as KITTI [1], Oxford RobotCar [6], KAIST [15], and 4Season [28], are focused on autonomous driving in urban environments. While these datasets incorporate substantial challenges, the highly structured scenarios and vehicle-centric perspectives may oversimplify odometry and place recognition tasks. Other datasets like TUM RGB-D [5] and TUM VI [16] offer sequences captured using hand-held stereo or RGB-D cameras primarily indoors. However, these sequences are typically short and re-visit previous locations from similar viewpoints and with limited variations in visual appearance. To address these limitations, synthetic datasets such as ICL-NUIM [21] and TartanAir [7] simulate more general and challenging motions and environments. However, in the quest for more diversity, *unstructured* natural environments still pose the most challenging conditions for visual or LiDAR-based SLAM.

This work was supported by the Helmholtz Association, project ARCHES (contract no. ZT-0033), the Spanish Government (projects PID2021-127685NB-I00 and TED2021-131150B-I00) and the Aragón Government (DGA-T45-23R).

¹Institute of Robotics and Mechatronics, German Aerospace Center (DLR) <firstname>.<surname>@dlr.de

²I3A, Universidad de Zaragoza, Spain jcivera@unizar.es



(a) The LRU rover traversing a Moon-analogue environment of Mt. Etna, Sicily, recording the DLR Planetary Stereo Solid-State LiDAR Inertial (S3LI) dataset [41]. (b) Aligned visual and 3D LiDAR data. Note the challenging texture and geometry for place recognition.

Fig. 1: (a) The LRU rover traversing the Moon-analogue environment of Mt. Etna, Sicily, recording the DLR Planetary Stereo Solid-State LiDAR Inertial (S3LI) dataset [41]. (b) Aligned visual and 3D LiDAR data. Note the challenging texture and geometry for place recognition.

In this paper we propose a novel multimodal place recognition method that we denote as UMF, standing for Unifying local and global Multimodal Features.¹ Our model leverages deep local and global features from visual and LiDAR data, fusing both modalities via cross-attention mechanisms. As our main novelty, we incorporate local feature-based re-ranking to a multimodal setup, showing in our experiments that it leads to a substantial and consistent improvement, in particular in challenging unstructured natural scenes.

II. RELATED WORK

Place recognition is a well-established field with a wide literature, relevant applications, and still significant challenges [3], [9], [10], [30], [34], [40], [50]. In relation to our research, a first categorization can be done based on the input data modality. **Visual place recognition** has largely evolved from traditional hand-crafted features [4], [29] to learned descriptors [11], [14], recently moving towards attention mechanisms [19], [47]. **LiDAR place recognition** emphasizes the geometric properties of environments, generally utilizing point cloud data. PointNet [12], [14] serves as the foundational architecture for subsequent works such as PointNetVLAD [17] and LPD-Net [16]. Attention mechanisms have also been used to improve feature specificity [39], [44]. **Visual-LiDAR place recognition** aims at overcoming the limitations inherent in using either modality alone, combining their strengths for more reliable recognition. Techniques like feature concatenation and attention are used in [32], [43], respectively, which are the main baselines in this area.

¹Our naming builds on the title of Ref. [23].

Attention mechanisms and transformers have shown great utility in several ways. They can serve as patch descriptor filters to spotlight important scene elements [19], [42] or as weight maps that modulate the CNN feature maps to generate global features [46]. This enables models to focus on salient features and ignore irrelevant information.

Patch-NetVLAD [31] and the work of Zheng *et al.* [48] showcase the benefits of **unifying local and global features** for visual place recognition. The latter targets the problem of visual aliasing in unstructured scenes [35], combining SuperPoint and SuperGlue in the context of 6D camera relocalization. In this work, we propose for the first time unifying local and global features for multimodal data.

Self-supervised pretraining refers to an initial training stage across multiple domains that can then be fine-tuned for specific downstream tasks. Here, **contrastive learning methods** like SimCLR [25] and BYOL [26] take on an important role. Likewise, **generative self-supervised methods** allow to learn the underlying dependencies in data. In particular, the masked autoencoders do so by learning to fill artificially corrupted parts of the input data, which can only be successfully performed by producing rich, useful data representations. Unlike contrastive methods, generative methods do not require carefully designed data augmentations or pair constructions, which are especially challenging on point clouds. For visual data, ConvNeXt [52] and Spark [51] employ 3D sparse convolutions to apply CNNs to point clouds. For LiDAR data, however, the application of self-supervised learning presents unique challenges. The inherent sparsity and lack of order in LiDAR point clouds, along with the necessity to capture complex geometric and spatial features, make conventional generative methods less suitable. One of the few successful works for generative, self-supervised pretraining on large point clouds is Occupancy-MAE [45]. It employs voxelization and masking during the training phase, prompting the model to become “voxel-aware,” thereby efficiently leveraging the geometrical and spatial redundancies within the point cloud data. Occupancy-MAE has proven effective in downstream tasks, including 3D object detection and semantic segmentation, even with a high masking ratio of up to 70%. In this work, we employ generative self-supervised pretraining methods in the spirit of the masked autoencoder due to their suitability for multiple data modalities and their lack of dependence on intricate 3D data augmentations. In particular, we build from the Occupancy-MAE due to its proven capability to manage large-scale point clouds with a high masking ratio.

III. UNIFYING LOCAL AND GLOBAL MULTIMODAL FEATURES (UMF)

An overview of our model is showcased in Fig. 2. At its core, the UMF design consists of 1) two distinct branches, one for visual data and the other for LiDAR, that encode the data and extract local features, 2) self- and cross-attention blocks to extract a global representation leveraging both modalities, and 3) re-ranking methods that take the top- k candidates by their global representation and re-order them

based on local feature matching on each modality. Next, we provide further details on the main aspects of the method.

Image and LiDAR Encoding: Firstly, each of these branches transform the respective input data into low-dimensional representations via a Resnet50 backbone and a LiDAR convolutional encoder [45]. Both encoders follow a Feature Pyramid Network (FPN) architecture to extract multi-scale features, thereby capturing both the local fine-grained details and also relevant patterns with wider extent.

Self- and Cross-Attention: Following [43], the first work that incorporated attention in multimodal data for place recognition, our UMF model also incorporates attention to enhance its capability to dynamically focus on different parts of the input data. In the **self-attention** layers, the model assigns different weights of importance to the features within a single modality (either F_{Vision} or F_{LiDAR}), thereby capturing patterns within local and global contexts. This allows the model to identify distinctive patterns within each modality and enhances its ability to recognize places based on a single modality. **Cross-attention** layers, on the other hand, take features from both modalities (F_{Vision} and F_{LiDAR}) as inputs. By interleaving self- and cross-attention layers within our UMF model, it becomes capable of capturing relevant patterns between the two modalities, thereby learning richer scene representations.

Local and Global Features: Building upon [23], our UMF model incorporates both local and global features, integrating fine-grained details along with the global spatial embedding, and extends the approach to the multi-modal case. We utilize transformers with positional encoding for coarse-level fusion, employing both self- and cross-attention. We implement the same approach introduced in DETR [24], which ensures that each element in the feature maps F_{Vision} and F_{LiDAR} has a unique positional encoding such that the transformed features exhibit dependencies on relative positions. The relative position dependency enhances the model spatial awareness, as well as inter- and intra-modality relationships between features in the fusion branch. Furthermore, such a design consideration promotes rotation and viewpoint invariance, enhancing robustness against perspective changes.

As the main novelty of our work, we incorporate re-ranking strategies to multimodal place recognition models. Specifically, we evaluate two strategies for re-ranking based on matching local features, the first one using the so-called Super-features [47] and the second one implementing RANSAC geometric verification.

Super-Features: Super-features [47] were proposed as mid-level scene representations, showing excellent results in place recognition tasks. The Local Super-features Integration Transformer (LIT) is trained via contrastive learning, passing local features through a transformer layer, as summarized in Fig. 3. Specifically, pairs of Super-features are constrained by the contrastive margin loss \mathcal{L}_s , which minimizes the pairwise distance between matching pairs while simultaneously reducing the spatial redundancy of Super-features s within

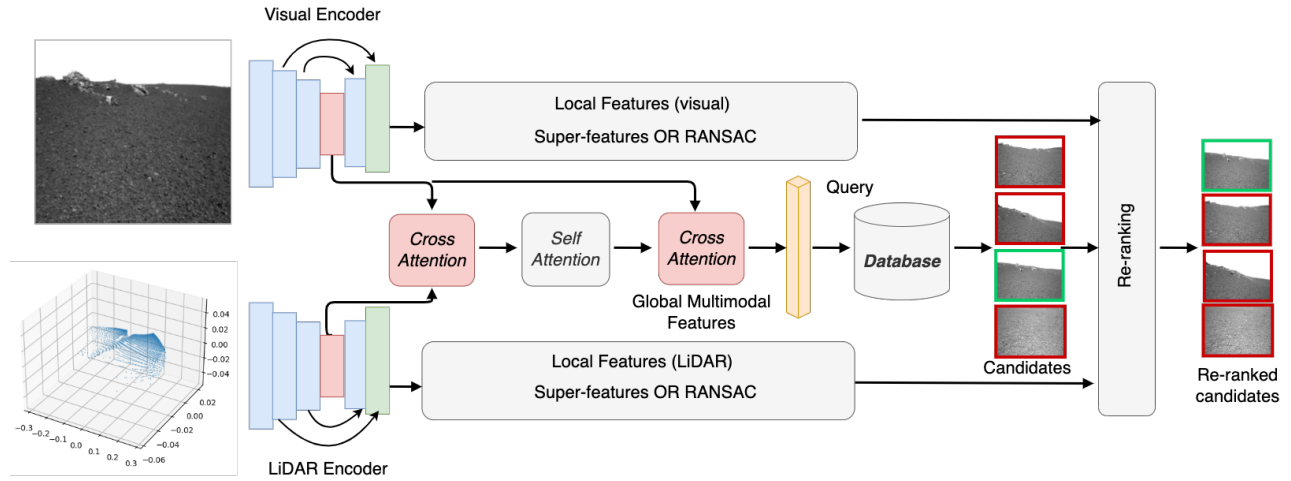


Fig. 2: UMF overview. Each branch encodes each of the inputs independently. The encodings of the individual modalities are used by self- and cross-attention modules into a single global multimodal representation. For each individual data modality, separate branches extract also local features. During inference, we query a database of places with the global multimodal descriptor using a K-Dimensional Tree, the top- k candidates are retrieved via NN-search, and finally they are re-ranked using local features from both modalities. This last stage is the main contribution of our paper.




Fig. 3: Local Super-features extracted with the LIT module for both modalities. Attention maps show the areas where each Super-feature is focused on.

an image:

$$\mathcal{L}_s = \sum_{\mathcal{P}} \left[\|s - s^+\|_2^2 + \sum_n [\mu' - \|s - n\|_2^2]^+ \right], \quad (1)$$

where μ is a margin hyper-parameter. The index n corresponds to the Super-features extracted from all negative images in the training tuple, which are compared against a specific Super-feature s . More explicitly, n serves as a collection of negative samples with the same Super-feature ID as that of s (denoted as $i(s)$).

This process results in a $N \times F$ ordered set of N Super-features of F dimensions. The construction of Super-features involves an iterative attention module, generating a set where each element focuses on a localized and discriminative image pattern. In this paper, we extend the approach of [47] to 3D scenarios, extracting Super-features from images as well as from voxelized point clouds.

To create Super-features that are complementary, they are encouraged to attend to different ($i \neq j$) local features, i.e., to different image locations. To do that, the cosine similarity

between the attention maps of all Super-features of every image is minimized. Let matrix $\alpha = [\tilde{\alpha}_1, \dots, \tilde{\alpha}_N]$ denote the N attention maps after the last iteration of LIT. The attention decorrelation loss is then given by:

$$\mathcal{L}_{\text{attn}}(\mathbf{x}) = \frac{1}{N(N-1)} \sum_{i \neq j} \frac{\tilde{\alpha}_i^\top \cdot \tilde{\alpha}_j}{\|\tilde{\alpha}_i\|_2 \|\tilde{\alpha}_j\|_2}. \quad (2)$$

In our UMF model, we implement a re-ranking mechanism to leverage correspondences at Super-feature level. As mentioned above, for any Super-feature $s \in \mathcal{S}$ we have a function $i(s)$ that returns the Super-feature ID, i.e., $i(s_i) = i, \forall s_i \in \mathcal{S}$. Also, let $n(s, \delta) = \arg \min_{s_i \in \delta} \|s - s_i\|_2$ be the nearest neighbor of s from the set δ .

Given a positive pair of images x, x^+ , and two Super-features $\{s \in \delta, s' \in \delta'\}$ from their respective Super-feature sets $\{\delta, \delta'\}$, we impose the following criteria to consider the Super-feature pair $\{s, s'\}$ eligible:

- 1) Reciprocal nearest neighbors: $s = n(s', \delta)$ and $s' = n(s, \delta')$.
- 2) Pass Lowe's first-to-second neighbor ratio test [2]: $|s - s'|_2 / |s' - n(s', \delta \setminus s)|_2 \geq \tau$.
- 3) Have the same Super-feature ID: $i(s) = i(s')$.

We set $\tau = 0.8$ after empirical considerations.

RANSAC: The Random Sample Consensus (RANSAC) variant of our UMF model emphasizes salient feature selection and geometric verification. Initially, a transformer layer is applied to generate attention maps. These maps then filter and pinpoint salient local features, retaining those that carry significant information for place recognition, as shown in Fig. 4. The local features are further processed using stacked layers of transformer blocks and we use the attention maps to filter as salient those features for which $\mathcal{L}_{\text{attn}} > \delta$, δ being a hyperparameter that we tuned experimentally for both modalities. Similarly to the Super-features, we project




Fig. 4: RANSAC local branch visualization. The resulting attention maps are used to select the salient features.

the local features to a 1-D embedding and use a contrastive loss on each local modality during training.

The RANSAC algorithm is then employed to estimate the geometric transformation between the current observation and candidate locations. This provides a robust approach to matching local features, accommodating the presence of noise and outliers effectively. The spatial consistency score is given by the number of inliers returned when fitting a homography between the two images or computing the rigid transformation between voxel grids, using corresponding keypoints computed by nearest neighbor matching.

A. Training Pipeline

UMF leverages unlabeled data from similar domains, such as the Mars-analogue in Morocco [35], for pretraining. This self-supervised learning approach makes the encoder robust to environmental variations, minimizing the dependency on labeled data and accelerating convergence during fine-tuning for the downstream task. In detail, it leverages the masked autoencoders for both visual and LiDAR modalities, inspired by Spark [51] and Occupancy-MAE [45].

Visual self-supervised pretraining follows a contrastive learning approach on a vast unlabeled dataset. It is designed to discern visually similar yet distinct locations, thereby improving the model’s ability to tackle visual aliasing. A patch-wise masking strategy segments images into non-overlapping square patches, each subject to independent masking according to a predetermined mask ratio. We follow the approach of [51], where the authors propose to assemble all unmasked patches into a sparse image, which is then encoded using sparse convolutions. The encoder f is based on ResNet and the decoder g on U-Net, including three blocks with upsampling layers (see Fig. 5). Subsequently, a “densification” process, based on mask embeddings, produces dense feature and projection layers.

LiDAR self-supervised pretraining adopts the self-supervised pre-training scheme based on masked autoencoders on voxelized clouds presented in Ref. [49]. This encourages the network to reason about high-level semantics to recover the masked occupancy distribution of the 3D scene from a limited number of visible voxels. A binary occupancy

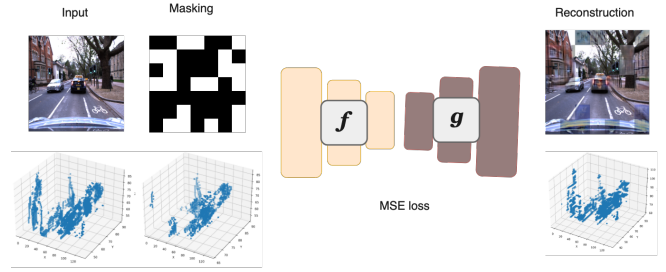


Fig. 5: Illustration of our pre-training on RobotCar. First, masked inputs are encoded by f , followed by the densification process in the decoder g . After pre-training, only the encoder f is used for downstream tasks.

classification loss is calculated using cross-entropy:

$$\mathcal{L}_{\text{LiDAR}} = -\frac{1}{n_{\text{batch}}} \sum_{i=1}^{n_{\text{batch}}} \sum_{j=1}^{n_{\text{voxel}}} \mathbf{T}_j^i \log \mathbf{P}_j^i, \quad (3)$$

where \mathbf{P}_j^i is the predicted occupancy probability of voxel j for the i -th training sample and \mathbf{T}_j^i the ground truth probability indicating whether the voxel contains points.

To address the problem of the range-dependent density of LiDAR points, we employ a range-aware random masking strategy [45]. This separates the occupied voxels into three groups based on their distance from the sensor. We apply a distinct random masking strategy to each group, with decreasing ratio with increasing distance.

Downstream task fine-tuning for place recognition uses a triplet margin loss and batch-hard negative mining, following the MinkLoc approach [32]. The triplets are constructed based on spatial proximity, with a 12 m radius for similarity and over 60 m for dissimilarity. Batch-hard negative mining targets active triplets for effective model refinement. Properly balancing global and modality-specific losses is here of major importance.

IV. EXPERIMENTAL RESULTS

We test the UMF model on two very different datasets: the DLR S3LI Dataset for planetary rover exploration [41] and the Oxford RobotCar dataset for autonomous driving [22].

A. Data Setup

The Oxford RobotCar dataset include diverse driving scenarios across varied weather and lighting conditions. Nonground point clouds are down-sampled to 4096 points, while corresponding RGB images are down-sampled from 1280×960 to 224×224. To enhance data diversity and limit overfitting, we randomly sample from 15 closest RGB images during training, while only one RGB image with the closest timestamp is used during evaluation. Similarity is defined based on their spatial proximity: elements within 10 m are deemed similar, those separated by at least 50 m are dissimilar; and those falling between 10 and 50 m are neutral. The dataset is split into disjoint training (21.7k elements) and test (3k elements) areas based on UTM coordinates,

following the evaluation protocol and the train/test split (baseline) introduced in [32].

The DLR S3LI dataset includes sequences captured with a hand-held sensor setup comprising a solid-state LiDAR and a stereo camera. The planetary environment is affected by extreme visual aliasing and lacks of salient visual or structural features. The dataset was split into training and validation set, using *s3li_loops* and *s3li_traverse_1* for testing, as they contain overlapping areas.

Additional datasets for pre-training and fine-tuning were used to increase the model’s robustness and performance. For unstructured planetary environments, we incorporate the MADMAX dataset [35], the Erfoud dataset [27], and the Long Range Navigation Tests (LRNTs) [18]. To overcome the scarcity of datasets that contain both visual and LiDAR modalities, we also generated synthetic sequences with the OAISYS photorealistic simulator [36].

Dataset	# Samples	Modalities used
MADMAX [35]	11,000	Stereo
Erfoud [27]	7,000	Stereo, LiDAR
LRNTs [18]	5,600	Stereo
OAISYS [36]	10,000	Stereo, LiDAR, Instance, Semantic

TABLE I: Datasets used for pre-training and fine-tuning.

B. Implementation Details

All models were implemented in PyTorch [53] and trained on a compute cluster equipped with 8 NVIDIA RTX 3090 GPUs. For fine-tuning, the learning rate was set to $1e-5$ and reduced by a factor of $1e-1$ upon plateauing. The model was trained for 200 epochs using Adam [8]. The input image size was set to 224×224 . When re-ranking global feature retrieval results with local feature-based matching, the top 25 ranked images from the first stage are considered.

Super-features: The score for Super-features was determined using $\mathcal{L}_{\text{attn}}$. The attention maps were generated at image size of 56×56 and voxel size of $50 \times 50 \times 50$. Super-features were represented as a tensor of size $[N, F]$, and the dimensions for the visual and point cloud features are set to 128 and 32, respectively. The final ranking is based on the number of matching features that satisfy the criteria described in Sec. III.

RANSAC: We use a multi head transformer encoder in our model to process the fine features. The model returns the average of all attention maps and selects an optimal threshold to identify keypoints. The output consists of attention maps for the image and voxel, of sizes $[N, 56, 56]$ and $[N, 50, 50, 50]$, respectively, where N is the number of keypoints. The feature maps for the image and voxel are $[56, 56, 128]$ and $[50, 50, 50, 32]$, respectively. The re-ranking of candidates is based on the total number of inliers, either for one or both modalities: $score = \#inliers_{\text{vision}} + \#inliers_{\text{LiDAR}}$.

Similarity Threshold: The threshold θ plays a pivotal role in our model’s ranking process, acting as a cut-off value to distinguish between similar and dissimilar pairs during re-ranking. Given two samples S_1 and S_2 , with a derived

similarity score $\text{sim}(S_1, S_2)$, they are classified as recognizing the same place if $\text{sim}(S_1, S_2) > \theta$, where θ adjusts the model sensitivity when classifying candidate matches as true or false positives, and it is determined empirically through cross-validation to optimize precision and recall rates.

C. Comparison against Baselines

Table II shows the comparative performance of our UMF models vs. relevant baselines on the S3LI dataset. Our best model **outperforms the best baseline by more than 2% in the three metrics chosen**, which can be attributed mainly to the use of local features for re-ranking. While the LiDAR data has a limited field of view, they still provide a valuable input under challenging conditions by reducing uncertainty. Moreover, the accurate depths play a crucial role in establishing correct positive pairs. Also, the geometric verification using RANSAC shows a significant impact in aliased environments, outperforming other approaches. For an illustrative understanding, Fig. 6 depicts the Super-features attention maps generated by our LIT.

Method	Recall@1	Recall@5	Top 1% recall
Visual			
DBoW2	37.44	66.1	68.12
NetVLAD	67.2	75.5	78.3
MinkLoc++ (Vision)	68.8	77.3	79.2
LiDAR			
PointNet++	48.41	67.77	71.8
MinkLoc++ (LiDAR)	42.4	65.8	69.4
Multimodal			
MinkLoc++	71.4	80.1	85.2
AdaFusion	73.1	82.3	87.2
UMF (only global feat.)	73.5	82.9	87.5
UMF (Super-features)	75	85.1	89.1
UMF (RANSAC)	75.3	85.3	89.5

TABLE II: Comparison against baselines on S3LI (Mt. Etna), for single- and multi-modality.

The evaluation on the RobotCar dataset, in a very different application domain and with different sensor specifications, demonstrates the robustness and adaptability of UMF. Table III shows multimodal results compared against the baselines MinkLoc++ and AdaFusion, in which our UMF model with re-ranking outperforms again both baselines.

Method	Recall@1	Top 1% recall
Multimodal		
MinkLoc++	96.7	99.1
AdaFusion	98.1	99.2
UMF (only global feat.)	97.9	99.1
UMF (Super-features)	98.1	99.1
UMF (RANSAC)	98.3	99.3

TABLE III: Comparison against baselines on RobotCar.

D. Ablation Studies

Re-ranking: To analyze the role of the re-ranking module in our approach, we conduct experiments involving variations of it, such as with and without, for one or all modalities. The improvements after re-ranking shown in Table IV substantiate its importance. On the S3LI dataset, we observe that

Modality	Method	S3LI (Mt. Etna)			RobotCar		
		Recall@1	Recall@5	Top 1% recall	Recall@1	Recall@5	Top 1% recall
Vision	UMF (Super-features)	74.5	84.3	89	98	98.4	99.1
	UMF (RANSAC)	75.1	84.9	89.3	98.1	98.5	99.2
LiDAR	UMF (Super-features)	73.7	83.4	87.5	97.8	98.2	99.1
	UMF (RANSAC)	73.9	83.8	87.8	98	98.3	99.1
Vision+LiDAR	UMF (Super-features)	75	85.1	89.1	98.1	98.5	99.1
	UMF (RANSAC)	75.3	85.3	89.5	98.3	98.8	99.3

TABLE IV: Ablation study on S3LI and RobotCar. Note how the multimodal (Vision+LiDAR) setup presents the best metrics, effectively leveraging both modalities. Observe also how our RANSAC variant consistently outperforms Super-features. Finally, note how the improvement offered by our methods is bigger in S3LI than in the almost saturated RobotCar.

the visual modality outperforms LiDAR. The re-ranking step on LiDAR data shows marginal improvements. The fusion of both modalities, however, does aid in overcoming visual challenges such as poor lighting or aliasing in both datasets. Notably, RANSAC emerged as a clear winner over Super-features within the re-ranking methods. This advantage, however, comes at the expense of an increased computational cost.

The normalized similarity threshold θ defined in Section IV-B has also been adjusted to examine the effectiveness of each variant in the precision-recall curves, see Fig. 7. As expected, RANSAC clearly outperforms the Super-features curve. Despite both approaches offering competitive performance compared to the baseline models without re-ranking, precision deteriorates rapidly, confirming again challenges in S3LI due to the lack of salient features.




Fig. 6: Illustration of the Super-features attention maps generated by the Learning Iterative Transformer (LIT), for S3LI (1st rows) and RobotCar (2nd and 3rd row). The first three Super-features are highlighted, showing the model’s propensity to focus on semantic patterns such as rock formations and terrain structures. The last row shows the attention map extracted from the multi head attention for a pair of RobotCar samples.

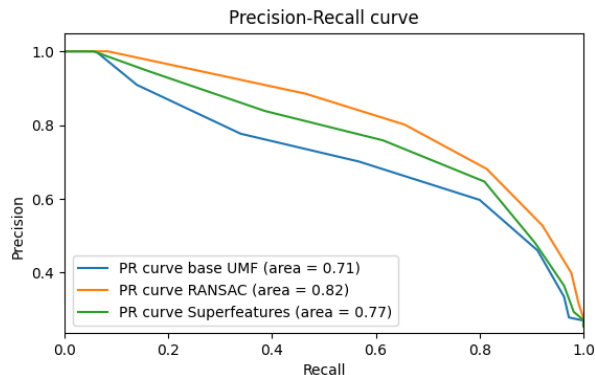


Fig. 7: Precision-recall curves on S3LI (Mt. Etna). We compare the base UMF model with both re-ranking variants.

Pre-training: To assess the impact of pre-training (see Sec. III-A), we perform an ablation study by initializing the weights with either random values or a pre-trained set. Table V shows the important role of pre-training. In our experience, this especially applies for complex models. Also, it reinforces the model’s robustness and generalization by reducing its propensity to overfit on the visual modality.

Method	Recall@1	Recall@5	Top 1% recall
UMF (w/o pre-training)	70.4	81.2	85.9
UMF (w/ pre-training)	73.5	82.9	87.5

TABLE V: Influence of pre-training on S3LI.

V. CONCLUSIONS

In this paper we have presented the “Unifying local and global Multimodal Features” (UMF) model, a novel place recognition method that fuses local and global features of both visual and LiDAR data using transformers and incorporates re-ranking steps based on single-modality local features. We evaluate our UMF against state-of-the-art baselines on two different domains: urban and planetary. UMF shows superior performance in terms of Recall@ N in both domains. In particular, we observed that our model outperforms significantly previous baselines in the extreme conditions of the S3LI dataset (Mt. Etna, Sicily), demonstrating the potential of a sound multimodal fusion for place recognition in challenging (unstructured and aliased) scenes.

REFERENCES

- [1] N. Gelfand *et al.*, “Geometrically stable sampling for the icp algorithm,” in *International Conference on 3-D Digital Imaging and Modeling (3DIM)*, 2003, pp. 260–267.
- [2] D. G. Lowe, “Distinctive image features from scale-invariant keypoints,” *International journal of computer vision*, vol. 60, pp. 91–110, 2004.
- [3] B. Williams, M. Cummins, J. Neira, P. Newman, I. Reid, and J. Tardós, “A comparison of loop closing techniques in monocular slam,” *Robotics and Autonomous Systems*, vol. 57, no. 12, pp. 1188–1197, 2009.
- [4] D. Gálvez-López and J. D. Tardós, “Bags of binary words for fast place recognition in image sequences,” *IEEE Transactions on Robotics*, vol. 28, no. 5, pp. 1188–1197, Oct. 2012, ISSN: 1552-3098. DOI: 10.1109/TRO.2012.2197158.
- [5] J. Sturm *et al.*, “A benchmark for the evaluation of rgb-d slam systems,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 573–580.
- [6] A. Geiger *et al.*, “Vision meets robotics: The kitti dataset,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [7] A. Handa *et al.*, “A benchmark for rgb-d visual odometry, 3d reconstruction and slam,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 1524–1531.
- [8] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [9] E. Garcia-Fidalgo and A. Ortiz, “Vision-based topological mapping and localization methods: A survey,” *Robotics and Autonomous Systems*, vol. 64, pp. 1–20, 2015.
- [10] S. Lowry, N. Sünderhauf, P. Newman, *et al.*, “Visual place recognition: A survey,” *IEEE transactions on robotics*, vol. 32, no. 1, pp. 1–19, 2015.
- [11] R. Arandjelović, P. Gronat, A. Torii, T. Pajdla, and J. Sivic, “NetVLAD: CNN architecture for weakly supervised place recognition,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.
- [12] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” *arXiv preprint arXiv:1612.00593*, 2016.
- [13] G. Bresson *et al.*, “Simultaneous localization and mapping: A survey of current trends in autonomous driving,” *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 3, pp. 194–220, 2017.
- [14] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” *arXiv preprint arXiv:1706.02413*, 2017.
- [15] J. Choi *et al.*, “Kais multi-spectral day/night data set for autonomous and assisted driving,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 3, pp. 934–948, 2018.
- [16] R. Schubert *et al.*, “The tum vi benchmark for evaluating visual-inertial odometry,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 1680–1687.
- [17] M. A. Uy and G. H. Lee, “Pointnetvlad: Deep point cloud based retrieval for large-scale place recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [18] M. Vayugundla, F. Steidle, M. Smisek, M. J. Schuster, K. Bussmann, and A. Wedler, “Datasets of long range navigation experiments in a moon analogue environment on mount etna,” in *ISR 2018; 50th International Symposium on Robotics*, 2018, pp. 1–7.
- [19] Y. Zhu, J. Wang, L. Xie, and L. Zheng, *Attention-based pyramid aggregation network for visual place recognition*, 2018.
- [20] M. Burki *et al.*, “Vizard: Reliable visual localization for autonomous vehicles in urban outdoor environments,” in *IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 1124–1130.
- [21] T. Schops *et al.*, “Bad slam: Bundle adjusted direct rgb-d slam,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [22] 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 *2020 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2020, pp. 6433–6438.
- [23] B. Cao, A. Araujo, and J. Sim, “Unifying deep local and global features for image search,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, Springer, 2020, pp. 726–743.
- [24] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko, “End-to-end object detection with transformers,” in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, Springer, 2020, pp. 213–229.
- [25] T. Chen, S. Kornblith, K. Swersky, M. Norouzi, and G. Hinton, “Big self-supervised models are strong semi-supervised learners,” *arXiv preprint arXiv:2006.10029*, 2020.
- [26] J.-B. Grill, F. Strub, F. Althé, *et al.*, *Bootstrap your own latent: A new approach to self-supervised learning*, 2020. arXiv: 2006.07733 [cs.LG].
- [27] S. Lacroix, A. De Maio, Q. Labourey, *et al.*, “The erfoud dataset: A comprehensive multi-camera and lidar data collection for planetary exploration,” in *15th Symposium on Advanced Space Technologies in Robotics and Automation*, 2020.
- [28] S. Wenzel *et al.*, “4seasons: A cross-season dataset for multi-weather slam in autonomous driving,” in *DAGM German Conference on Pattern Recognition*, Springer, 2020, pp. 404–417.
- [29] C. Campos, R. Elvira, J. J. G. Rodriguez, J. M. M. Montiel, and J. D. Tardós, “ORB-SLAM3: An accurate Open-Source library for visual, Visual-Inertial, and multimap SLAM,” *IEEE Trans. Rob.*, pp. 1–17, 2021.
- [30] S. Garg, T. Fischer, and M. Milford, “Where is your place, visual place recognition?” *arXiv preprint arXiv:2103.06443*, 2021.
- [31] S. Hausler, S. Garg, M. Xu, M. Milford, and T. Fischer, *Patch-NetVLAD: Multi-Scale fusion of Locally-Global descriptors for place recognition*, 2021.
- [32] J. Komorowski, M. Wyszczafiska, and T. Trzcinski, “MinkLoc++: Lidar and monocular image fusion for place recognition,” in *2021 International Joint Conference on Neural Networks (IJCNN)*, Jul. 2021, pp. 1–8.
- [33] S. Mascaro *et al.*, “Towards automating construction tasks: Large-scale object mapping, segmentation, and manipulation,” *Journal of Field Robotics*, vol. 38, no. 5, pp. 684–699, 2021.
- [34] C. Masone and B. Caputo, “A survey on deep visual place recognition,” *IEEE Access*, vol. 9, pp. 19 516–19 547, 2021.
- [35] L. Meyer, M. Smišek, A. Fontan Villacampa, *et al.*, “The madmax data set for visual-inertial rover navigation on mars,” *Journal of Field Robotics*, vol. 38, no. 6, pp. 833–853, 2021.
- [36] M. G. Müller, M. Durner, A. Gawel, W. Stürzl, R. Triebel, and R. Siegwart, “A Photorealistic Terrain Simulation Pipeline for Unstructured Outdoor Environments,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [37] R. Oliveira *et al.*, “Advances in agriculture robotics: A state-of-the-art review and challenges ahead,” *Robotics*, vol. 10, no. 2, p. 52, 2021.
- [38] X. Shu *et al.*, “Slam in the field: An evaluation of monocular mapping and localization on challenging dynamic agricultural environment,” in *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 1761–1771.
- [39] T.-X. Xu, Y.-C. Guo, Y.-K. Lai, and S.-H. Zhang, “TransLoc3D: Point cloud based large-scale place recognition using adaptive receptive fields,” May 2021. arXiv: 2105.11605 [cs.CV].
- [40] X. Zhang, L. Wang, and Y. Su, “Visual place recognition: A survey from deep learning perspective,” *Pattern Recognition*, vol. 113, p. 107 760, 2021.
- [41] R. Giubilato, W. Stürzl, A. Wedler, and R. Triebel, “Challenges of slam in extremely unstructured environments: The dlr planetary stereo, solid-state lidar, inertial dataset,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8721–8728, 2022. DOI: 10.1109/LRA.2022.3188118.
- [42] M.-H. Guo, C.-Z. Lu, Z.-N. Liu, M.-M. Cheng, and S.-M. Hu, *Visual attention network*, 2022.
- [43] H. Lai, P. Yin, and S. Scherer, “AdaFusion: Visual-LiDAR fusion with adaptive weights for place recognition,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12 038–12 045, 2022.
- [44] J. Ma, J. Zhang, J. Xu, R. Ai, W. Gu, and X. Chen, “OverlapTransformer: An efficient and Yaw-Angle-Invariant transformer network for LiDAR-Based place recognition,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6958–6965, Jul. 2022.
- [45] C. Min, X. Xu, D. Zhao, L. Xiao, Y. Nie, and B. Dai, “Occupancy-mae: Self-supervised pre-training large-scale lidar point clouds with masked occupancy autoencoders,” *arXiv e-prints*, 2022.
- [46] R. Wang, Y. Shen, W. Zuo, S. Zhou, and N. Zheng, “TransVPR: Transformer-based place recognition with multi-level attention aggregation,” Jan. 2022. arXiv: 2201.02001 [cs.CV].

- [47] P. Weinzaepfel, T. Lucas, D. Larlus, and Y. Kalantidis, "Learning Super-Features for image retrieval," in *International Conference on Learning Representations*, 2022.
- [48] Y. Zheng, T. Birdal, F. Xia, Y. Yang, Y. Duan, and L. J. Guibas, "6d camera relocalization in visually ambiguous extreme environments," *arXiv preprint arXiv:2207.06333*, 2022.
- [49] G. Hess, J. Jaxing, E. Svensson, D. Hagerman, C. Petersson, and L. Svensson, "Masked autoencoder for self-supervised pre-training on lidar point clouds," in *2023 IEEE/CVF Winter Conference on Applications of Computer Vision Workshops (WACVW)*, IEEE, Jan. 2023. DOI: 10.1109/wacvw58289.2023.00039. [Online]. Available: <https://doi.org/10.1109%2Fwacvw58289.2023.00039>.
- [50] S. Schubert, P. Neubert, S. Garg, M. Milford, and T. Fischer, "Visual place recognition: A tutorial," *arXiv preprint arXiv:2303.03281*, 2023.
- [51] K. Tian, Y. Jiang, Q. Diao, C. Lin, L. Wang, and Z. Yuan, "Designing BERT for convolutional networks: Sparse and hierarchical masked modeling," Jan. 2023. arXiv: 2301.03580 [cs.CV].
- [52] S. Woo, S. Debnath, R. Hu, *et al.*, "ConvNeXt v2: Co-designing and scaling ConvNets with masked autoencoders," Jan. 2023. arXiv: 2301.00808 [cs.CV].
- [53] F. A. research lab, "Pytorch," <https://pytorch.org/>, (date accessed 11-3-2023).