

# OMCL: Open-vocabulary Monte Carlo Localization

Evgenii Kruzhkov<sup>1</sup>, Raphael Memmesheimer<sup>1</sup>, and Sven Behnke<sup>1</sup>

**Abstract**—Robust robot localization is an important prerequisite for navigation, but it becomes challenging when the map and robot measurements are obtained from different sensors. Prior methods are often tailored to specific environments, relying on closed-set semantics or fine-tuned features. In this work, we extend Monte Carlo Localization with vision-language features, allowing OMCL to robustly compute the likelihood of visual observations given a camera pose and a 3D map created from posed RGB-D images or aligned point clouds. These open-vocabulary features enable us to associate observations and map elements from different modalities, and to natively initialize global localization through natural language descriptions of nearby objects. We evaluate our approach using Matterport3D and Replica for indoor scenes and demonstrate generalization on SemanticKITTI for outdoor scenes. The code is accessible at <https://github.com/AIS-Bonn/omcl>.

**Index Terms**—Localization, Semantic Scene Understanding, Mapping.

## I. INTRODUCTION

LOCATION is a fundamental problem in robotics, allowing robots to estimate their position and orientation within an environment. Traditional approaches utilize single or fused sensor modalities [1], [2], [3]. Cross-modal localization [4] further increases flexibility for heterogeneous systems and facilitates map reuse across platforms, while incorporating semantic awareness enhances robot autonomy [5], [6]. The semantics can be incorporated through object detection [7], semantic segmentation [8], or vision-language features [9], [10]. This letter introduces the Open-vocabulary Monte Carlo Localization (OMCL) framework, which extends Monte Carlo Localization (MCL) with vision-language features [10]. These abstract features enable affordable camera-only localization in 3D maps created from different sensors, like RGB-D cameras or LiDAR. Fig. 1 demonstrates the language map and localization process of the proposed framework.

Our approach stores features learned by contrastive language-image pretraining (CLIP [10]) in a 3D map. The positions of features are represented within a spatial language map, while RGB input is processed to extract the features that describe the current observation. Ray tracing is employed to estimate the correlation between the observations and the map allowing for pose estimation in 3D environments. We propose a technique to convert pre-existing maps to the compatible representation for our localization framework, compare it to

Manuscript received: August, 16, 2025; Revised November, 14, 2025; Accepted December, 14, 2025.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers' comments.

<sup>1</sup> All authors are with the Autonomous Intelligent Systems group, Computer Science Institute VI – Intelligent Systems and Robotics – and the Center for Robotics and the Lamarr Institute for Machine Learning and Artificial Intelligence, University of Bonn, Germany; [ekruzhkov@ais.uni-bonn.de](mailto:ekruzhkov@ais.uni-bonn.de)

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

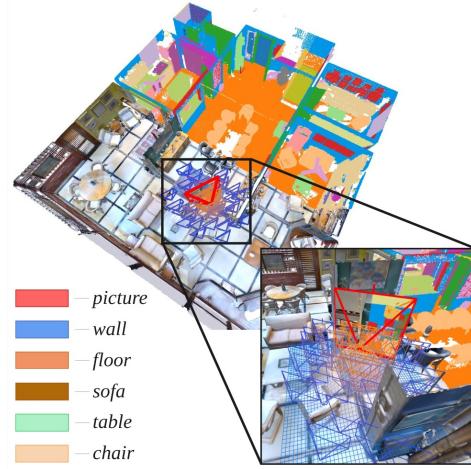


Fig. 1. OMCL particles are sampled on a Language Map storing open-vocabulary features. Each particle represents a candidate camera pose. Particles are weighted according to how well the VLM-processed RGB input matches the ray-traced map value at the relevant location. The red particle denotes the estimated pose (weighted mean). The Language Map is colored by similarity to the prompted labels.

state-of-the-art methods, and conduct a comprehensive ablation study. Our contributions include the following:

- **Language-grounded localization.** We present OMCL, a novel localization framework that grounds pose estimation in language features and accelerates global localization via open-vocabulary prompts.
- **Cross-modal sensor usage.** OMCL includes a mapping module that constructs a unified, sparse language map, enabling sensors of different modalities to be used for mapping and localization (e.g., RGB-D or point clouds for mapping and RGB for localization).
- **Generalization.** OMCL is compatible with independently constructed point clouds, enabling reuse of existing maps, and it generalizes across indoor and outdoor environments.

OMCL offers an effective, language-feature guided solution for localization. It achieves state-of-the-art performance compared to existing baselines and opens up new possibilities for natural language-guided and cross-modal localization. A robot can be localized by non-expert users without the necessity of comprehending the underlying map representation through the use of natural language descriptions in OMCL.

## II. RELATED WORKS

**Semantic Localization.** We focus on visual continuous localization within 3D maps constructed from cross-modal sensors [11], [12]. Unlike Loc-NeRF [13], which relies on rendered visual consistency for pose tracking, we achieve sensor cross-modality through semantic consistency. Prior methods,

SeDAR [14] and SIMP [7], use MCL for long-term indoor semantic localization on floor plans; in contrast, we propose a more general map representation and a semantic-consistency measurement model. We evaluate against SeDAR [14] and SIMP [7] by adapting our map representation to their floor-plan views.

SemLoc [15] and Zhang *et al.* [16] explore semantic consistency for localization in outdoor environments. Unlike these methods, we rely purely on semantic consistency in a multi-modal feature space, without additional geometry-based optimization terms or custom semantic segmentation models. CMRNext [4] is a state-of-the-art camera-to-LiDAR map-matching baseline for outdoor localization, which, in contrast to our approach, relies on a fine-tuned neural network model.

Unlike our approach, place recognition (PR) methods [17] estimate the pose from a single image, whereas we handle continuous image streams and pose tracking, potentially requiring multiple observations before convergence. However, MCL allows using the same map for both global localization and subsequent pose tracking [18]. While PR can be used for initialization, it requires additional PR-specific infrastructure; in contrast, we propose a map-native, prompt-augmented initialization.

**Semantic Mapping.** Many recent works focus on open-set semantic mapping of indoor environments followed by 3D semantic segmentation. SAM3D [19] creates 3D scene masks from SAM [20]. OpenScene [5] proposes a method to directly predict point-wise CLIP [10] features for input point clouds. ConceptGraphs [21] constructs a graph-based representation with incorporated semantic feature vectors. ConceptFusion [22], implemented on top of  $\nabla$ SLAM [23], produces unordered multi-modal maps. LiLMaps [24] investigates sequential visual–language mapping for implicit representations, reducing the memory footprint of such dense maps. Unlike these mapping-focused works, we provide options for constructing the map from different sensing modalities and couple it with a ray-tracing-based, semantic-aware probabilistic localization module.

The recent OVO [25] mapping approach uses Gaussian-SLAM [26] and ORB-SLAM2 [27] but does not exploit the semantic information added to the map to improve pose estimation. Concurrent to our work, RayFronts [28] similarly constructs an ordered map by averaging visual–language features from multiple directions. However, they focus on multi-modal open-set querying and beyond-range semantic classification. Similar to the aforementioned works, we construct our map representation from externally estimated poses, while tight integration with simultaneous localization and mapping (SLAM) [29], [30] is a possible future improvement.

**Vision Language Models.** We employ contrastive Vision-Language Models (VLMs) [10], [5], [31], [28], [32], for the localization task. VLMs and large language models (LLMs) are widely used for related navigation [33], [6] and scene understanding [34] tasks. We use pixel-wise visual–language features (such as CLIP [10]) produced by the image encoder as open semantic representations in the maps. By default, our framework employs LSeg [31] and OpenScene [5] encoders, but we demonstrate flexibility with respect to the choice of

visual–language backbone by switching to X-Decoder [32] and NARADIO [28].

We don't compete with the aforementioned approaches; instead, our method creates a unified visual–language map from cross-modal sensors and evaluates its downstream impact on localization.

### III. METHOD

The proposed pipeline is illustrated in Fig. 2. The map representation constructed and used by our framework is the Octree Language Map. We focus on storing visual–language features [10], [31], [5] in the map to investigate their potential for advancing perception systems. Section III-A covers Octree Language Map implementation details and, after that, describes its construction from different inputs based on pixel-wise visual encoders for RGB-D input and OpenScene [5] for point clouds.

In Section III-B we address the visual-only localization in 3D maps. Although multimodal sensor configurations (e.g., RGB-D cameras, LiDAR) can be employed for mapping, the subsequent localization requires only visual data and map-scaled odometry (Motion Model), under the assumption that the features in the map are consistent with those extracted from the input RGB image.

We focus on mapped visual–language features that are semantically grounded using open-set semantic prompts.

#### A. Mapping

Throughout the mapping process, we generate a volumetric language map (Octree Language Map), a sparse octree-based structure, in which each map voxel is linked to the corresponding  $F$ -dimensional visual–language feature [10], [31], [5]. The Octree Language Map offers compact storage at fine-grained resolution and provides efficient ray-tracing functionality [35]. It retains only features that are mutually different by at least a cosine distance threshold  $\tau$ . Each voxel in the map stores the index of its corresponding feature in the database  $\text{Features}_{\text{DB}}$ :

$$\begin{aligned} \text{Features}_{\text{DB}} &= \{f_i \in \mathbb{R}^n, \mid d(f_i, f_j) > \tau; i, j \in \mathbb{N}\}, \\ d(A, B) &= 1 - \frac{A \cdot B}{\|A\| \|B\|}, \text{ where } A, B \in \mathbb{R}^n. \end{aligned} \quad (1)$$

We provide two mapping approaches: one for posed RGB-D images and another for fused point clouds (Fig. 2). The produced Octree Language Map is unified.

*Input Option 1.* Unless stated otherwise, we use the LSeg [31] visual–language model to extract pixel-wise features and project them onto the Octree Language Map using depth measurements and known camera poses. The mapping is sequential, and features projected into the same voxel are cumulatively averaged with the existing ones. Once all data have been processed, the mapped features form  $\text{Features}_{\text{DB}}$  according to Eq. (1). The first features that satisfy Eq. (1) are added to  $\text{Features}_{\text{DB}}$ , while subsequent ones are replaced by their cosine-closest counterparts from  $\text{Features}_{\text{DB}}$ .

*Input Option 2.* We process the aggregated point cloud map with the OpenScene 3D distillation model [5] to predict visual–language features for each point in a feed-forward

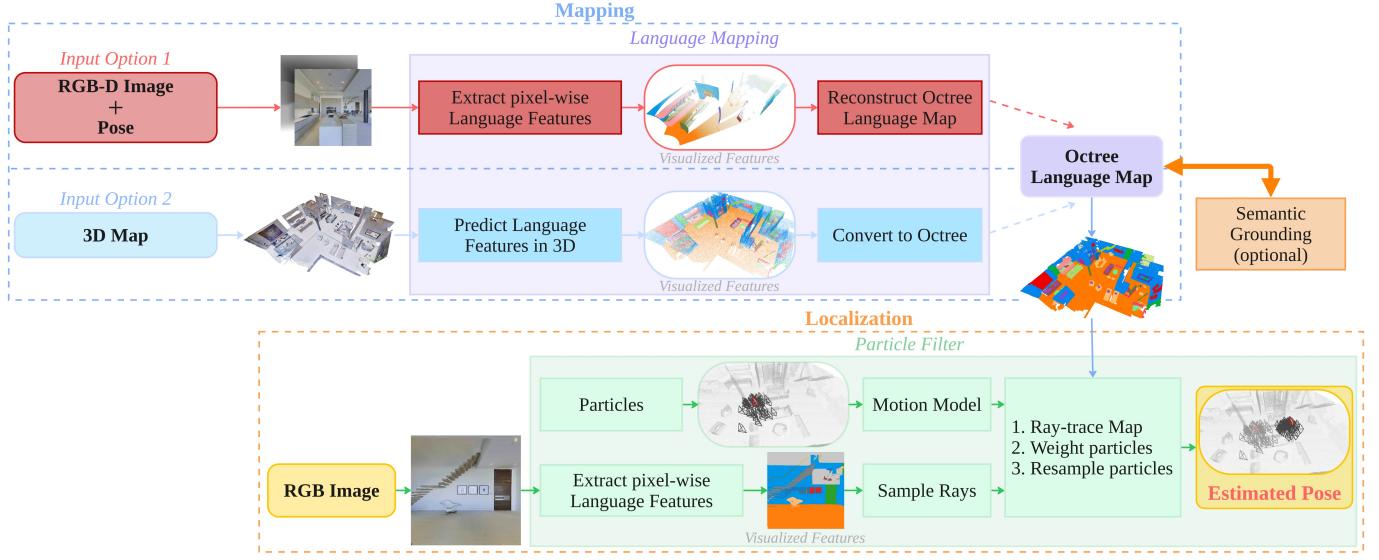


Fig. 2. **Mapping:** We propose two options to create Octree Language Maps. *Input Option 1*: OMCL derives language features from RGB images and reconstructs a 3D map from them using the corresponding volumetric data (depth images, LiDAR measurements, etc.). *Input Option 2*: Language features are directly predicted on precomputed 3D point clouds for each point and subsequently converted into the octree representation. **Localization:** A particle filter uses an RGB image as the only input, weighting particles by the discrepancy between language features extracted from the input image and those ray-traced from the Octree Language Map. Our stratified ray sampling strategy compensates for the imbalance between different object instance sizes in the image. All features are colored for visualization purpose only.

manner. The Octree Language Map is then formed by averaging features within voxels and constructing  $\text{Features}_{\text{DB}}$  as input option 1. The described approach is suitable for converting existing 3D maps into our Octree Language Map representation.

*Semantic Grounding.* Each element of the constructed  $\text{Features}_{\text{DB}}$  can be viewed as an automatically created semantic class. We propose to ground  $\text{Features}_{\text{DB}}$  using user-defined open-set semantic prompts. In practice, we employ the text encoder corresponding to the mapping model (CLIP [10] for LSeg [31] and OpenScene [5]) and redefine  $\text{Features}_{\text{DB}}$  using the features of the provided open-set semantic classes. The feature indices stored in the voxels of the Octree Language Map are remapped based on the highest cosine similarity between their original associated features and the redefined  $\text{Features}_{\text{DB}}$ . Old features that are too far in cosine distance from the provided open-set semantic classes are discarded along with their corresponding map voxels. Grounding reduces the memory footprint by decreasing the number of stored features. It increases the discriminability of the remaining features and allows the user to specify which semantic classes are represented in the map for downstream localization.

### B. Localization

We use RGB images and odometry (Motion Model) to localize within the language maps employing MCL. We initially assign uniformly distributed weights to each particle and sample particles with random poses around the probable starting position. When a new RGB image is received, the motion model is applied to all particles to predict their new poses. A coarse odometry from any built-in sensor can be employed. The image is processed using the language-driven semantic segmentation model [31] to extract per-pixel features.

To form the measurement model of our MCL and assign new weights for the particles, we evaluate the consistency between the extracted features and the map (Fig. 2 bottom).

*Observation Likelihood.* If the camera is exactly at a particle's pose, the features extracted from the image should match those stored in the Octree Language Map where viewing rays hit a surface. We employ camera intrinsics and each particle's predicted pose to form the rays in world coordinates. Ray-tracing [35] on the map finds the first voxel hit by each ray, from which we retrieve the stored language feature. Let  $\gamma$  and  $\varphi$  denote the voxel features and pixel features corresponding to the same rays, respectively. The weight  $w_i^t$  of particle  $i$  at time  $t$  is then estimated as:

$$w_i^t = \frac{w_i^{t-1} \max(\mathcal{L}_i, 0)}{\sum_i (w_i^{t-1} \max(\mathcal{L}_i, 0))}, \quad \mathcal{L}_i = \frac{1}{N} \sum_{j=1}^N \frac{\varphi_j \cdot \gamma_j}{\|\varphi_j\| \|\gamma_j\|}, \quad (2)$$

where  $N$  represents the total number of rays per particle.

$\mathcal{L}_i$  is, in effect, the averaged cosine similarity between image and map features. The more similar the observations and the map are, the higher the value of  $\mathcal{L}_i$  and, consequently, the higher the corresponding particle weight  $w_i^t$ . The final camera pose is then estimated as the average of particles that are close to the most likely particle, using their corresponding weights  $w_i^t$ .

*Stratified Ray Sampling.* We reduce computational costs by sampling and using only  $N$  pixels and their corresponding features from the image. However, large surfaces (e.g., walls, floor) occupy most pixels, while smaller but distinctive objects cover far fewer pixels, causing many objects to be underrepresented or missed when sampled uniformly. Since accurate pose estimation relies on preserving information from all objects, we introduce sampling masks for the image, where each mask corresponds to a cluster of pixels with similar visual-

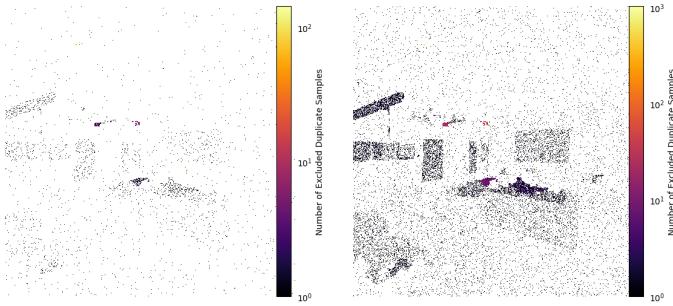


Fig. 3. Sampled pixels for images of resolution  $540 \times 540$ . Both images employ the same sampling masks. The left image corresponds to  $2^8$  samples per cluster and the right one to  $2^{11}$ . Small clusters have a higher sampling density, but the total number of samples is less for them because the duplicates are discarded.

language features. From each such feature-consistent cluster, we uniformly sample an equal number of pixels, discarding duplicates (Fig. 3). The same samples are employed for all particles. The proposed stratified per-cluster sampling improves pose estimation accuracy, as confirmed by our ablation study in Section IV-D.

To form the clusters, we use the map features  $\text{Features}_{\text{DB}}$  as centroids because, according to Eq. (2), localization assumes that map and image features are correlated and  $\text{Features}_{\text{DB}}$  already stores only distinct features (see Eq. (1)). The clusters are then formed from the extracted image features assigned to the centroids with which they produce the highest cosine similarity. Unless stated otherwise, we use full  $\text{Features}_{\text{DB}}$  for clustering, though random subsets of centroids may be used to limit the number of clusters and reduce computation. Both  $\text{Features}_{\text{DB}}$  and the clusters are created automatically, requiring no manual labeling for mapping or localization.

### C. Prompt-augmented Initialization

Standard approaches for global localization either require geometric coordinates as input for initialization or sample particles across the entire environment. However, sampling across the whole environment is inefficient, while geometric coordinates are difficult for non-expert users to interpret or provide, and localization may diverge if the initialization is incorrect. Moreover, global localization that is decoupled from the subsequent pose tracking step may require storing additional information within the map itself.

We propose using textual prompts  $\theta$  to describe probable initial positions. As demonstrated in Fig. 4, the prompt is an open set of natural language words that describe the surrounding environment. There is no restrictions on word content, as each word is encoded into a feature using a text encoder (CLIP [10] for LSeg [31] and OpenScene [5]).

First, the proposed initialization accepts natural language words as input, taking a step toward full natural language-based localization. Second, since the input is text, it is naturally compatible with the output of large language models, enabling initialization with AI agents. Finally, instead of relying on a single initial location, we initialize particles by sampling them uniformly over all locations that match the given prompt.



Fig. 4. An example of initial locations for the global localization based on the user prompt. The red spots correspond to the prompt **(toilet, mirror, towel, sink)** and the green one to **(table, chair, picture, door, tv monitor)**. OMCL particles can be initialized nearby the prompt matching spots instead of the random locations.

Through direct integration with MCL, the initialization subsequently converges to a precise pose, as described in Section III-B, based on multiple consecutive observations, and then MCL continues to track the pose. The same Octree Language Map is employed for both global localization and following pose tracking.

We segment the scene into two classes: floor and surroundings. We obtain the floor voxels by comparing all stored features with the feature corresponding to the word “floor”, extracted by the text encoder. For each floor voxel, we then find surrounding features  $\bar{\gamma}$  within a radius  $R$ . Next, for each word in the prompt  $\theta$ , we count the number of these surrounding features whose cosine similarity with that word exceeds the threshold  $\rho$ :

$$\mathcal{V}_m = \sum_n (\text{CosineSimilarity}(\bar{\gamma}_{n \times F}, \bar{\theta}_{m \times F}) > \rho), \quad (3)$$

where  $n$  is the number of neighbor surroundings,  $m$  is the number of words in the prompt, and  $F$  is the feature dimensionality.

Finally, we estimate the floor voxel–prompt alignment ratio  $s$  as a measure of how well the floor voxel matches the prompt, by requiring that at least  $k$  of its surrounding voxels correspond to each word:

$$s = \frac{1}{m} \sum_m (\mathcal{V}_m \geq k), \quad 0 \leq k \leq m, \quad (4)$$

where  $s$  is equal to 1 for floor voxels that match all words. The particles are uniformly sampled above “floor” voxels that match all words.

## IV. EXPERIMENTS

We provide complete evaluation of our localization in medium-size indoor environments on Matterport3D dataset [37], comparing against similar approaches and using publicly available sequences provided by VLMaps [6]. We further analyze how localization performance depends on measurement–map consistency.

Using the Replica dataset [38] for smaller, room-sized environments, we obtain localization accuracy aligned with that observed on medium-sized environments and additionally

TABLE I  
ABSOLUTE TRAJECTORY ERROR (APE) FOR LOCALIZATION ON MATTERPORT3D DATASET.

	RMSE [m]	STD [m]	Mean [m]	Median [m]	min RMSE [m]	max RMSE [m]	SSE [m <sup>2</sup> ]	Completed Scens
2D	OMCL <sub>GT</sub>	0.11	0.07	0.09	0.08	0.18	17.38	10/10
	<b>OMCL<sub>LSeg</sub></b>	<b>0.15</b>	<b>0.09</b>	<b>0.13</b>	<b>0.12</b>	<b>0.11</b>	<b>0.28</b>	<b>33.07</b>
	OMCL <sub>auto<sup>LSeg</sup></sub>	0.24	0.15	0.18	0.15	0.14	0.52	95.0
	OMCL <sub>Open\LSeg</sub>	0.36	0.24	0.27	0.22	0.15	0.62	164.07
	SeDAR <sub>GT</sub> <sup>*</sup> [14]	0.79	0.46	0.64	0.54	0.24	2.33	1891.2
	SIMP <sub>GT</sub> <sup>*</sup> [7]	1.39	1.03	0.79	0.94	0.45	2.81	2898.15
3D	OMCL <sub>GT</sub>	0.15	0.09	0.12	0.1	0.1	30.45	10/10
	<b>OMCL<sub>LSeg</sub></b>	<b>0.2</b>	<b>0.1</b>	<b>0.18</b>	<b>0.17</b>	<b>0.17</b>	<b>0.41</b>	<b>59.53</b>
	OMCL <sub>Open\LSeg</sub>	0.42	0.25	0.31	0.25	0.22	0.75	223.32
	RTAB-Map [36]	5.23	2.97	4.30	3.6	0.31	8.14	25781.93

TABLE II  
OMCL PARAMETERS.

Parameter	Value	Parameter	Value
Discrepancy threshold $\tau$	0.02	Language feature size $F$	512
Number of particles	1024	Surrounding radius $R$	2 [m]
Number of rays	2048	Similarity threshold $\rho$	0.9
Map resolution	0.02 [m]	Matching criterion $k$	500

assess the semantic coverage of our Octree Language Map relative to open-vocabulary baselines. For the large-scale evaluation, we employ the SemanticKITTI urban driving dataset [39], demonstrating generalization capabilities of the proposed approach.

Unless otherwise noted, we use the parameters summarized in Table II. We use the lower index notation OMCL<sub>Mapping\Localization</sub> to indicate the specific model used for mapping and localization, respectively. For instance, OMCL<sub>GT</sub> means that ground truth semantic images are used for both mapping and localization; while OMCL<sub>Open\LSeg</sub> means OpenScene [5] is used for mapping and LSeg [31] for localization. OMCL always operates in the continuous feature-embedding space, processing GT labels with a corresponding text encoder.

By default, we ground Features<sub>DB</sub> as described in Section III-A with a prompt consisting of the datasets' semantic class labels. For reference purposes, we provide evaluation for the ungrounded OMCL<sub>auto</sub> version. All results are obtained using an Nvidia RTX 3090 GPU equipped with 24GB of memory.

#### A. Long-Term Indoor Localization

**Localization Evaluation: Table I.** We compare our approach with SeDAR [14] and SIMP [7] which, similar to us, employ MCL and measure semantic consistency for long-term indoor localization (Section II). For fair comparison, all approaches are evaluated on the same Octree Language Map which was created from ground truth data (GT) and projected to a floor plan map (2D) by averaging the data along the vertical dimension. We denote baselines adapted to the common map as SeDAR\* and SIMP\*. The same noisy odometry is applied to all methods: ground-truth poses are perturbed with zero-mean Gaussian noise with standard deviations of 0.10 m (translation) and 6° (rotation). For reference, the average per-step ground truth motion is 0.13 m in translation and 8° in rotation. Particles are initialized around the starting pose with

TABLE III  
MEASUREMENTS-MAP CONSISTENCY EVALUATION.

	Accuracy [%]	Precision [%]	Recall [%]	IoU [%]
Matterport3D [37]				
OMCL <sub>GT</sub>	86.75	78.96	78.95	67.46
<b>OMCL<sub>LSeg</sub></b>	<b>82.74</b>	<b>59.74</b>	<b>54.30</b>	<b>43.35</b>
OMCL <sub>auto<sup>LSeg</sup></sub>	80.99	59.13	51.24	41.28
OMCL <sub>Open\LSeg</sub>	62.89	38.25	42.91	26.95
Replica [38]				
OMCL-NAR	96.53	23.87	46.8	17.1

standard deviations of 0.3 m and 17°. RTAB-Map [36] is employed as the evaluation baseline of our full pose estimation on the Octree Language Map in 3D. We use the RGB-D variant of RTAB-Map to avoid the degeneracies of monocular SLAM in low-parallax, in-place rotation segments widely present in the dataset.

For completeness, we report the full APE statistics [40] in Table I, averaged over all sequences, and provide the corresponding measurement-map consistency evaluation in Table III. We evaluate both OMCL configurations shown in (Fig. 2). Under “Input Option 1”, we use OMCL<sub>GT</sub> and OMCL<sub>LSeg</sub>. Under “Input Option 2”, OMCL<sub>Open\LSeg</sub> employs the OpenScene [5] 3D distillation model for the mapping and, because OpenScene does not accept image input, we use LSeg for localization. This is feasible because both models operate in the same CLIP feature space.

#### Measurements-Map Consistency Evaluation: Table III.

We report accuracy, precision, recall and IoU between the map and the observed features during localization, averaged over all poses and sequences for different datasets. To compute these metrics, we treat each feature stored in the Octree Language Map (i.e., each element of Features<sub>DB</sub>) as an independent semantic class. For each per-pixel image feature, we identify its corresponding map feature via ray casting and count a correct correspondence if that map feature attains the highest cosine similarity among all features in Features<sub>DB</sub>.

The evaluation quantifies a lower bound on consistency between the map and measurements during localization and assesses both the map's ability to preserve information and the image encoder's consistency across viewing directions. It is a lower-bound because correspondences are discretized as correct/incorrect, whereas localization operates in a continuous

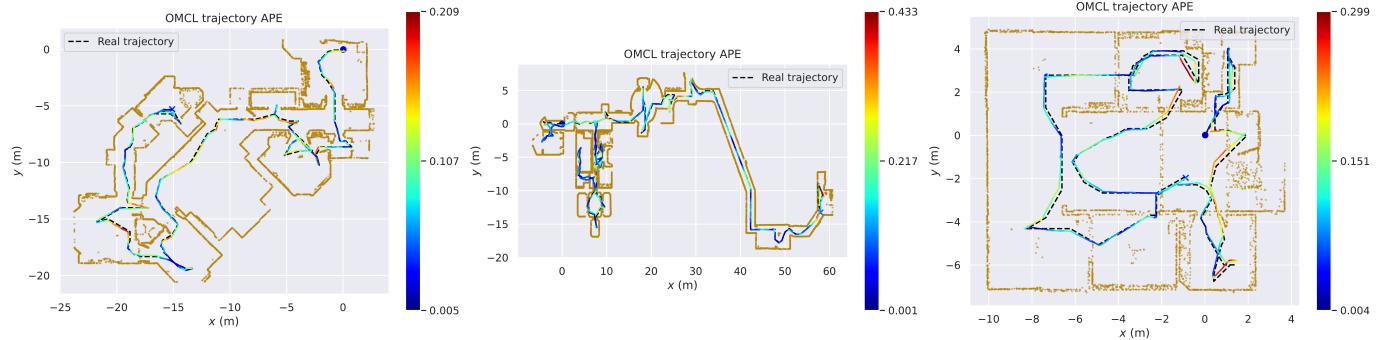


Fig. 5. Example trajectories performed by OMCL on the Matterport3D dataset, with APE indicated by color for each segment. Map projections are shown in brown. The plots demonstrate the performance in scenarios that involve loopy paths, corridors, and long monotonic trajectories.

TABLE IV  
REPLICA 3D SEMANTIC AND LOCALIZATION EVALUATION.

	All		Tail		Camera pose / ATE RMSE [cm]
	mIoU	mAcc	mIoU	mAcc	
OpenScene [5]	15.9	24.6	1.5	6.7	N/A
ConceptGraphs [21]	16.7	33.7	4.4	26.8	N/A
RayFronts [28]	27.7	54.5	17.6	41.3	N/A
OVO [25]	27.1	39.1	12.1	19.6	<b>0.6 – 1.9</b>
<b>OMCL<sub>NAR</sub> (ours)</b>	<b>32.1</b>	<b>56.2</b>	<b>21.4</b>	<b>43.9</b>	<u>35</u>

feature-similarity space. Note that localization depends solely on this consistency and not on the absolute correctness of the map’s semantics.

**Results and Comparison: Tables I and III.** The metrics reported in Table III align with the localization results in Table I. The metrics of OMCL<sub>LSeg</sub> are close to those of OMCL<sub>GT</sub> and both outperform SeDAR<sub>GT</sub><sup>\*</sup> and SIMP<sub>GT</sub><sup>\*</sup>. OMCL<sub>Open\LSeg</sub> lacks sufficient mapping consistency to complete all test sequences; nevertheless, it still outperforms the baselines and successfully executes most sequences. Moreover, OMCL<sub>Open\LSeg</sub> can create the Octree Language Map directly from existing point clouds. The ungrounded OMCL<sub>LSeg</sub><sup>auto</sup> achieves metrics comparable to OMCL<sub>LSeg</sub>.

Larger gaps between the Mean and RMSE in Table I indicate more significant deviations from the real pose. The geometry-based RTAB-Map [36] loses tracking in narrow, feature-poor areas. Although it subsequently relocates, the accumulated errors prevent accurate recovery of the full trajectory. Fig. 5 illustrates the performance of OMCL in a variety of scenarios, including loopy paths, corridors, and monotonic trajectories.

### B. Datasets Generalization

**Replica.** The Replica dataset [38] contains photo-realistic 3D indoor scene reconstructions at room and building scale. Following the OVO [25] evaluation procedure, we estimate the semantic map classification quality and report localization accuracy on the resulting estimated map in Table IV. We employ the NARADIO [28] image and text encoder for mapping and localization. Ground-truth odometry is perturbed with zero-mean Gaussian noise at 20% of the per-step motion. The classification accuracies are reported for all classes (All) and for the most challenging, low-frequency one-third subset

TABLE V  
ATE ON SEQUENCE 00 OF THE KITTI DATASET.

	Translation [m]		Rotation [°]		Fine-tuning data
	Mean	Std	Mean	Std	
Stereo DSO [41]	7.23	5.0	1.99	0.86	<b>no training</b>
SemLoc [15]	1.49	–	–	–	other urban-driving
Zhang <i>et al.</i> [16]	0.58	–	–	–	KITTI
Pi-Long [42]	9.88	5.39	6.24	3.61	<b>no fine-tuning</b>
<b>OMCL<sub>X-Dec</sub>(our)</b>	<b>0.52</b>	<b>0.26</b>	<b>1.77</b>	<b>0.26</b>	<b>no fine-tuning</b>
CMRNext [4]	0.11	0.06	0.25	0.15	KITTI + other urban

(Tail [25]). The camera poses in OVO [25] are estimated using Gaussian-SLAM [26] and ORB-SLAM2 [27], with ATE RMSEs of 0.6 cm and 1.9 cm, respectively.

OMCL<sub>NAR</sub> achieves the best semantic mapping quality due to its high-resolution mapping, while the small improvement over similar RayFronts is likely due to implementation details and design choices. The largest improvement is observed in the less frequent classes (Tail). Although the measured semantic quality is high, the subsequent localization on the map is stable but not very accurate (about 35 cm). This correlates with the low measurement-map consistency metric estimated for the used dataset-encoder (Replica-NARADIO) pair in Table III. Replica is smaller but also semantically denser than the Matterport3D dataset, which makes it more challenging to perceive visual features on the map equally well from all viewing directions. Precision and IoU in Table III correlate with localization accuracy across the Matterport3D and Replica datasets.

**SemanticKITTI.** SemanticKITTI [39] contains semantically annotated 3D LiDAR from urban driving. As in SemLoc, we rely on the reconstructed semantic point cloud and evaluate localization on existing maps. OMCL<sub>X-Dec</sub> employs X-Decoder [32] as the visual localization model and language encoder, with odometry provided by Stereo DSO. We use 512 particles, 4096 rays per image, and an Octree Language map resolution of 0.05 m. Following CMRNext [4], we report the averaged localization metrics on Sequence 00 in Table V.

OMCL<sub>X-Dec</sub> outperforms other semantic-consistency-based localization methods in accuracy, while requiring no fine-tuning, leading to a significant improvement in the estimated odometry. CMRNext is the state-of-the-art camera-to-LiDAR map matching approach, but it requires fine-tuning to the

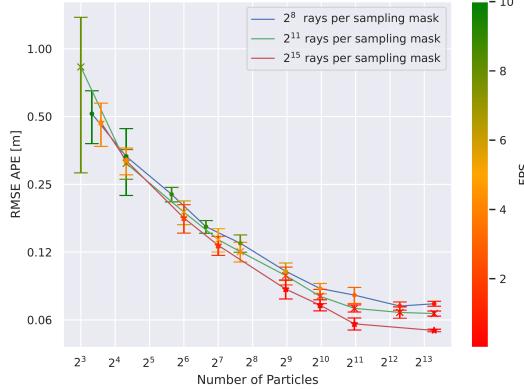


Fig. 6. Impact of the number of particles on localization accuracy and FPS for varying numbers of sampled rays.

data and was trained on the remaining KITTI sequences. We also compare against the feed-forward Pi-Long approach [42] with loop closure, which, similar to ours, performs monocular localization and does not require data-specific fine-tuning.

### C. Prompt-augmented Initialization

In order to assess the proposed initialization, we randomly distribute particles across the map and report the average number of OMCL steps required to achieve the designated localization accuracy in Table VI. The single OMCL step corresponds to the weighing of the particles and then resampling. The prompts are randomly generated from the lists of probable words for each starting location, with each prompt containing three to five words. We evaluate on the sequences of Matterport3D dataset, and localization can be considered sufficient once it reaches 0.2 m, since this matches the accuracy of the subsequent pose tracking (Table I).

Table VI shows that the maximum number of steps needed for localization is 44, with fast convergence once 1 m accuracy is achieved. On average, localization requires 29 steps to determine the camera pose, and a minimum of 24 steps is needed to reach 0.1 m accuracy.

Enhancing global localization with the user prompt according to Section III-C allows us to accelerate localization convergence. OMCL<sub>LSeg</sub><sup>0.1 m / prompt</sup> used 14 steps on average to achieve an accuracy of 0.1 m with the fastest case taking 3 steps, accelerating localization convergence by factors of 2.7 and 8, respectively.

### D. Ablation Studies

We demonstrate the impacts of different parameters and ray sampling strategies in Figs. 6 and 7. The plots are constructed by running one trajectory ten times with OMCL<sub>LSeg</sub>. According to Fig. 7, our stratified ray sampling strategy (Section III-B) outperforms the one with an equal distribution of the rays among the observable categories in both the APE metric and the FPS for the high number of rays. Moreover, our implementation consumes less GPU memory, allowing it to handle a larger number of rays. Compared with uniform sampling over the full image, our stratified approach can work even with a small number of sampled rays because it tries to

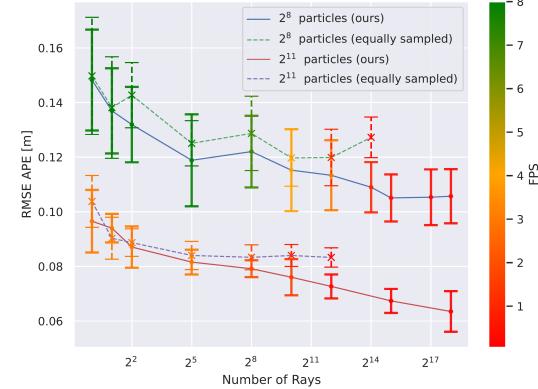


Fig. 7. Comparison between our stratified ray sampling strategy and equal distribution of rays among observable categories, for different particle counts.

TABLE VI  
GLOBAL LOCALIZATION STEPS BEFORE CONVERGENCE.

	mean [st.]	std [st.]	min [st.]	max [st.]
OMCL <sub>LSeg</sub> <sup>2 m</sup>	19	8	11	39
OMCL <sub>LSeg</sub> <sup>1 m</sup>	22	8	16	41
OMCL <sub>LSeg</sub> <sup>0.5 m</sup>	25	8	16	42
OMCL <sub>LSeg</sub> <sup>0.2 m</sup>	29	8	17	44
OMCL <sub>LSeg</sub> <sup>0.1 m</sup>	38	7	24	45
<b>OMCL<sub>LSeg</sub><sup>0.1 m / prompt</sup></b>	<b>14</b>	<b>5</b>	<b>3</b>	<b>23</b>

The steps indicate the number of OMCL iterations needed to reach the specified localization accuracy.

cover all different observable object instances. It tends to cover the full image without duplicated samples when increasing the number of rays. The APE variance does not decrease gradually with an increase in rays. Keeping the number of particles constant, using more rays improves the APE by about 10% compared to 2<sup>8</sup> rays.

Fig. 6 demonstrates that the increase in the number of particles reduces the variance of APE. It is noticeable that even with a small number of particles, OMCL can still perform pose estimation. In the beginning, the particles have a higher impact on APE compared to the rays (Fig. 7), but the rays have a smaller initial APE error. Increasing the number of particles beyond 2<sup>11</sup> leads to a 10% improvement in APE, but the system may become less stable to the track loss. Together, Figs. 6 and 7 help to select parameters with corresponding APE and processed frames per second (FPS) to satisfy the application needs. We do not include the timing of the used visual model in the presented FPS values because it varies for different models and can be parallelized.

## V. CONCLUSION

In this paper, we introduced OMCL, a framework for vision-based Monte Carlo localization in open-vocabulary 3D semantic maps. The maps are created from posed RGB-D images or 3D point clouds and store CLIP features in an octree, enabling cross-modal sensor setups for mapping and localization. A textual prompt can be used to initialize localization. OMCL computes the likelihood of localization hypotheses based on the consistency of open-vocabulary features extracted from the current input image and the corresponding map features

that are retrieved by ray casting. OMCL is flexible with respect to the choice of the feature extractor and benefits from stratified ray sampling. Our method generalizes across room-scale indoor, large-scale indoor, and outdoor environments. We provide a measurements-map consistency analysis and ablation studies of the proposed framework.

OMCL lacks online map updates and requires map-scaled odometry. Future work on tighter integration with a SLAM backend could potentially address both issues.

## ACKNOWLEDGMENT

This work has been partially funded by the Federal Ministry of Research, Technology and Space of Germany (BMFTR) under grant no. 01IS22094A WestAI and within the Robotics Institute Germany, grant no. 16ME0999.

## REFERENCES

- [1] N. Krombach, D. Droschel, S. Houben, and S. Behnke, “Feature-based visual odometry prior for real-time semi-dense stereo SLAM,” *Robotics and Autonomous Systems*, vol. 109, pp. 38–58, 2018.
- [2] D. Droschel and S. Behnke, “Efficient continuous-time SLAM for 3D Lidar-based online mapping,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 5000–5007.
- [3] H. Cao and S. Behnke, “SLCF-Net: Sequential LiDAR-camera fusion for semantic scene completion using a 3D recurrent U-Net,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2024, pp. 2767–2773.
- [4] D. Cattaneo and A. Valada, “CMRNext: Camera to LiDAR matching in the wild for localization and extrinsic calibration,” *IEEE Transactions on Robotics (T-RO)*, vol. 41, pp. 1995–2013, 2025.
- [5] S. Peng, K. Genova *et al.*, “OpenScene: 3D scene understanding with open vocabularies,” in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023, pp. 815–824.
- [6] C. Huang, O. Mees, A. Zeng, and W. Burgard, “Visual language maps for robot navigation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 10 608–10 615.
- [7] N. Zimmerman, M. Sodano, E. Marks, J. Behley, and C. Stachniss, “Constructing Metric-semantic maps using floor plan priors for long-term indoor localization,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2023, pp. 1366–1372.
- [8] C. Yu, Z. Liu, X. Liu, F. Xie, Y. Yang *et al.*, “DS-SLAM: A semantic visual SLAM towards dynamic environments,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2018, pp. 1168–1174.
- [9] N. Muhammad, C. Paxton, L. Pinto, S. Chintala, and A. Szelam, “CLIP-Fields: Weakly supervised semantic fields for robotic memory,” in *Robotics: Science and Systems (RSS)*, 2023.
- [10] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh *et al.*, “Learning transferable visual models from natural language supervision,” in *38th Int. Conf. on Machine Learning (ICML)*, vol. 139, 2021, pp. 8748–8763.
- [11] E. Yudin, P. Karpyshev *et al.*, “CloudVision: DNN-based visual localization of autonomous robots using prebuilt LiDAR point cloud,” in *IEEE Vehicular Technology Conference (VTC)*, 2023, pp. 1–6.
- [12] L. Zhang, Y. Tao, J. Lin, F. Zhang, and M. Fallon, “Visual localization in 3D maps: Comparing point cloud, mesh, and NeRF representations,” *arXiv:2408.11966*, 2024.
- [13] D. Maggio, M. Abate, J. Shi, C. Mario, and L. Carlone, “Loc-NeRF: Monte Carlo localization using neural radiance fields,” in *IEEE Int. Conference on Robotics and Automation (ICRA)*, 2023, pp. 4018–4025.
- [14] O. Mendez, S. Hadfield, N. Pugeault, and R. Bowden, “SeDAR-Semantic detection and ranging: Humans can localise without LiDAR, can robots?” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 6053–6060.
- [15] S. Liang, Y. Zhang, R. Tian, D. Zhu, L. Yang, and Z. Cao, “SemLoc: Accurate and robust visual localization with semantic and structural constraints from prior maps,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 4135–4141.
- [16] C. Zhang, H. Zhao, C. Wang *et al.*, “Cross-modal monocular localization in prior LiDAR maps utilizing semantic consistency,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2023, pp. 4004–4010.
- [17] B. Zhao, L. Yang, M. Mao, H. Bao, and Z. Cui, “PNeRFLoc: Visual localization with point-based neural radiance fields,” in *AAAI Conference on Artificial Intelligence (AAAI)*, vol. 38, no. 7, 2024, pp. 7450–7459.
- [18] N. Atanasov, M. Zhu, K. Daniilidis, and G. J. Pappas, “Localization from semantic observations via the matrix permanent,” *The International Journal of Robotics Research (IJRR)*, vol. 35, no. 1-3, pp. 73–99, 2016.
- [19] Y. Yang, X. Wu, T. He, H. Zhao, and X. Liu, “SAM3D: Segment anything in 3D scenes,” *arXiv:2306.03908*, 2023.
- [20] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao *et al.*, “Segment anything,” in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 4015–4026.
- [21] Q. Gu, A. Kuwajerwala, S. Morin *et al.*, “ConceptGraphs: Open-vocabulary 3D scene graphs for perception and planning,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2024, pp. 5021–5028.
- [22] K. M. Jatavallabhula, A. Kuwajerwala, Q. Gu, M. Omama, G. Iyer, S. Saryazdi, T. Chen *et al.*, “ConceptFusion: Open-set multimodal 3D mapping,” in *Robotics: Science and Systems XIX (RSS)*, 2023.
- [23] K. M. Jatavallabhula, G. Iyer, and L. Paull, “gradSLAM: Dense SLAM meets automatic differentiation,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 2130–2137.
- [24] E. Kruzhkov and S. Behnke, “LiLMaps: Learnable implicit language maps,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025, pp. 7711–7720.
- [25] T. B. Martins, M. R. Oswald, and J. Civera, “Open-vocabulary online semantic mapping for SLAM,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 10, no. 11, pp. 11 745–11 752, 2025.
- [26] V. Yugay, Y. Li, T. Gevers, and M. Oswald, “Gaussian-SLAM: Photo-realistic dense SLAM with Gaussian splatting,” *arXiv:2312.10070*, 2023.
- [27] R. Mur-Artal and J. D. Tardós, “ORB-SLAM2: an open-source SLAM system for monocular, stereo, and rgbd cameras,” *IEEE Transactions on Robotics (T-RO)*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [28] O. Alama, A. Bhattacharya, H. He, S. Kim, Y. Qiu, W. Wang, C. Ho, N. Keetha *et al.*, “RayFronts: Open-set semantic ray frontiers for online scene understanding and exploration,” *arXiv:2504.06994*, 2025.
- [29] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, “Fusion++: Volumetric object-level SLAM,” in *Int. Conf. on 3D Vision (3DV)*. IEEE, 2018, pp. 32–41.
- [30] J. Stückler, N. Biresev, and S. Behnke, “Semantic Mapping Using Object-Class Segmentation of RGB-D Images,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 3005–3010.
- [31] B. Li, K. Q. Weinberger, S. J. Belongie, V. Koltun, and R. Ranftl, “Language-driven semantic segmentation,” in *10th International Conference on Learning Representations (ICLR)*, 2022.
- [32] X. Zou, Z.-Y. Dou, J. Yang, Z. Gan, L. Li, C. Li, X. Dai, H. Behl, J. Wang, L. Yuan *et al.*, “Generalized decoding for pixel, image, and language,” in *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 15 116–15 127.
- [33] B. Chen, F. Xia, B. Ichter, K. Rao *et al.*, “Open-vocabulary queryable scene representations for real world planning,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2023, pp. 11 509–11 522.
- [34] H. Ha and S. Song, “Semantic Abstraction: Open-world 3D scene understanding from 2D vision-language models,” in *Conference on Robot Learning (CoRL)*, 2022.
- [35] K. M. Jatavallabhula, E. J. Smith, J. Lafleche, C. F. Tsang, A. Rozantsev *et al.*, “Kaolin: A PyTorch library for accelerating 3D deep learning research,” *CoRR*, vol. abs/1911.05063, 2019.
- [36] Labb  , Mathieu and Michaud, Fran  ois, “RTAB-Map as an open-source Lidar and visual SLAM library for large-scale and long-term online operation,” *Journal of Field Robotics (JFR)*, vol. 36, no. 2, pp. 416–446, 2019.
- [37] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner *et al.*, “Matterport3D: Learning from RGB-D data in indoor environments,” *Int. Conf. on 3D Vision (3DV)*, 2017.
- [38] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma *et al.*, “The Replica dataset: A digital replica of indoor spaces,” *arXiv:1906.05797*, 2019.
- [39] J. Behley, M. Garbade, A. Milioti, J. Quenzel, S. Behnke, J. Gall, and C. Stachniss, “Towards 3D LiDAR-based semantic scene understanding of 3D point cloud sequences: The SemanticKITTI dataset,” *International Journal of Robotics Research (IJRR)*, vol. 40, no. 8-9, 2021.
- [40] M. Grupp, “evo: Python package for the evaluation of odometry and SLAM.” <https://github.com/MichaelGrupp/evo>, 2017.
- [41] R. Wang, M. Schworer, and D. Cremers, “Stereo DSO: Large-scale direct sparse visual odometry with stereo cameras,” in *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 3903–3911.
- [42] K. Deng, Z. Ti, J. Xu, J. Yang, and J. Xie, “VGGT-Long: Chunk it, loop it, align it – pushing VGGT’s limits on kilometer-scale long RGB sequences,” 2025. [Online]. Available: <https://arxiv.org/abs/2507.16443>