

LiLMaps: Learnable Implicit Language Maps

Evgenii Kruzhkov

Autonomous Intelligent Systems, Computer Science Institute VI
University of Bonn, Germany

ekruzhkov@ais.uni-bonn.de

Sven Behnke

Autonomous Intelligent Systems, Computer Science Institute VI – Intelligent Systems and Robotics
Center for Robotics and the Lamarr Institute for Machine Learning and Artificial Intelligence

University of Bonn, Germany

behnke@cs.uni-bonn.de

Abstract

One of the current trends in robotics is to employ large language models (LLMs) to provide non-predefined command execution and natural human-robot interaction. It is useful to have an environment map together with its language representation, which can be further utilized by LLMs. Such a comprehensive scene representation enables numerous ways of interaction with the map for autonomously operating robots. In this work, we present an approach that enhances incremental implicit mapping through the integration of vision-language features. Specifically, we (i) propose a decoder optimization technique for implicit language maps which can be used when new objects appear on the scene, and (ii) address the problem of inconsistent vision-language predictions between different viewing positions. Our experiments demonstrate the effectiveness of LiLMaps and solid improvements in performance.

1. Introduction

Classic robotic maps are commonly used for estimating distances to obstacles and costs of motions in navigation and localization tasks. However, more comprehensive tasks, as well as natural human-robot interaction, may require a deeper understanding of the environment, and thus imply more advanced map representations. For example, visual-language navigation is the task where a robot must interpret a natural language command from a non-expert user and proceed towards the goal according to the command. The environment might be unknown in advance, but the robot still must navigate in the shortest possible time.




Figure 1. Reconstructed implicit language map built with LiLMaps. Semantic colors are assigned based on the similarity of reconstructed language features and CLIP [33] encodings of semantic categories from the Matterport3D dataset [5].

From this example, it is clear that a map that allows one to easily find correlation between the given language command and a partially or fully mapped environment can be more useful than a pure obstacle costmap.

In this work, we make a step towards creation of efficient yet compact natural language environment representations and introduce Learnable implicit Language Maps (LiLMaps). We chose an implicit representation because of its ability to compactly represent the data and for the possi-

bility of further detailed reconstruction.

Recent studies in implicit mapping demonstrate outstanding results in geometry reconstruction. While in these studies, geometry decoders can be easily pre-trained or even trained in the first few iterations, in our work we demonstrate that LiLMaps performs better compared to the pre-trained language decoders because some language features can be poorly represented in them. Moreover, pre-training the decoder to represent every possible language feature could make its structure and training process significantly more complicated, and it can reduce its flexibility for different applications.

Another challenging problem that frequently appears in incremental language learning is the inconsistency of measurements taken from different viewing positions. Precise range sensors, such as LiDARs and RGB-D cameras, are commonly used in implicit mapping, but usually they do not provide contradictory measurements. However, in visual-language navigation tasks, information about the environment is often derived from RGB images. Vision-language features extracted from RGB images may have many sources of inconsistency: a painting can be recognized as a wall from a greater distance; a bed object can be misclassified as a sofa at different angles of view; objects on image borders and occluded objects might not be visible enough to provide correct features; inaccurate detection on the object edge can spoil features of the objects behind them; etc.

LiLMaps focuses on incremental implicit language mapping, i.e., when new observations become available incrementally, one-by-one. This is a typical condition for SLAM, and LiLMaps can be integrated into existing implicit SLAM approaches with minimal changes. We achieve this with the following key techniques that are presented in this work:

- **Adaptive Language Decoder Optimization** dynamically updates the decoder to new discovered language features in the environment providing flexible and sufficient coverage of language representations.
- **Measurements Update Strategy** adjusts incoming measurements reusing accumulated and implicitly stored knowledge about the environment to reduce measurements inconsistency.

Our experiments show that LiLMaps enables incremental vision-language environment exploration with just a small overhead.

2. Related Works

Language and Vision-Language Models. Dynamic execution of natural language commands has been an active research topic for a long time [2, 13]. Previous works often propose custom environment representations that are

difficult to reuse in real applications. Recently, large language models (LLMs) have received increased attention in this field. The advantage of LLMs is their ability to be applied to a wide range of tasks. LLMs can enhance robot abilities to understand and execute natural language commands [14, 29]. In addition, LLMs have been shown to be successful in guiding object grasping [9, 23], navigation [1, 6, 15, 34, 37] and scene understanding [7, 11, 31]. Vision-Language Models (VLMs) [10, 20, 22, 35], which extract language information about the environment from the provided images, are frequently used in conjunction with LLMs. For instance, CLIP [33] is one of the most widely used models capable of mapping images into a natural language space.

Often, VLMs transfer only single or batched images into a natural language vector space. However, some tasks may benefit from mapping the entire environment into the language space. For example, VLMaps [15] suggests enhancing 2D maps with language features and then demonstrates that navigation and detection tasks can be solved directly on these enhanced maps. Although VLMaps can be used alongside simultaneous localization and mapping (SLAM), its performance depends on localization and mapping quality, and the method is limited to building only 2D language maps. On the other hand, OpenScene [31] can generate maps as 3D point clouds with corresponding language features. However, the method performs best in batch-like operations, where all RGB images of the environment, their poses, and 3D point clouds are available in advance, meaning no real-time data is processed. Another approach, ConceptFusion [16], demonstrates that language features can be fused into 3D maps using traditional SLAM approaches. SAM3D [41] projects SAM segmentation masks [19] into 3D and creates 3D scene masks.

Implicit Representations. Implicit representations [3, 24, 26, 30, 36] have gained popularity for their compactness and ability to achieve high-resolution reconstructions. They demonstrate great capabilities in environment mapping. iMap [38] uses an RGB-D sensor to perform a real-time SLAM task. Nice-SLAM [46] extended the possible sizes of the mapped environment. SHINE-Mapping [43] demonstrated implicit mapping of outdoor environments. Recently, works based on Gaussian splatting [17] demonstrated exceptional results [24, 28, 44, 45]. In addition to geometry, implicit maps demonstrate a successful reconstruction of semantic information [42], physical properties [12], and visual features [25].

Integration of language features into implicit representation is an actively researched topic. LERF [18] studies the fusion of language features into Radiance Fields, but is limited to small scenes. LangSplat [32] uses Gaussian Splatting [17] to achieve higher precision and training speed. However, LangSplat demonstrates only the reconstruction

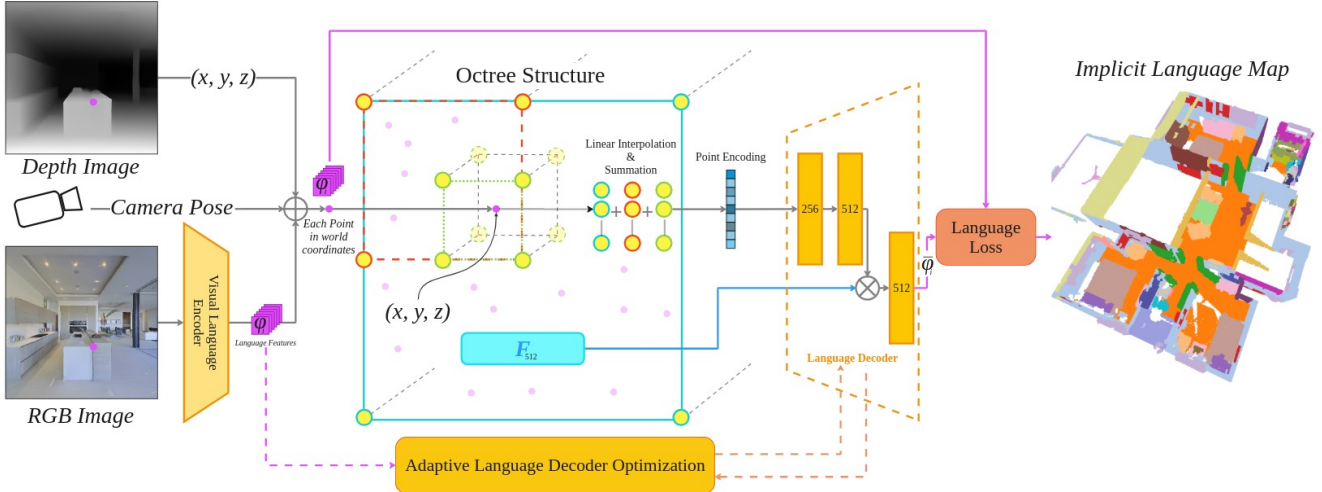


Figure 2. Implicit language mapping. Vision-language features φ are extracted from the RGB image. The corresponding points of the depth image are projected to the world coordinate system. Each point can be encoded using its coordinates and octree: the coordinates are used to find the corresponding octree voxels (blue, red, green); learnable features stored in the voxels’ corners are interpolated and summed, producing the point encoding. F vectors are stored only in the voxels of the coarse octree level (blue). The language decoder reconstructs the language feature $\hat{\varphi}$ in the spatial coordinates of the point based on its encoding and the vector F . Language loss optimizes the learnable features and F vectors. After optimization, the language map can be reconstructed in arbitrary spatial coordinates. The language detector is optimized independently of the implicit mapping (Sec. 3.2).

of small table-sized scenes and does not consider incremental mapping.

Implicit representations are highly dependent on the type of encoding they use. The encodings employed in the original NeRF work [26] were able to generalize the predictions [12, 25], but were constrained by the limited size of the environments. Subsequent works successfully increased training and reconstruction speeds [27, 36]. Gaussian Splatting [17] is currently one of the most popular encodings due to its speed, simplicity, and high quality reconstruction. However, some works may benefit from structured encodings such as grid-based [40], octree-based feature volumes [39], or combined ones [21].

Compared to the works discussed above, our approach can build large-scale 3D implicit language maps and can be seamlessly integrated with implicit SLAM methods. LiLMaps use a sparse octree-based representation [39] to store learnable features, but our method is not tied to any particular representation and can be adapted to others with minimal effort.

3. Method

We address the task of building an implicit vision-language representation along with environment mapping. Sec. 3.1 describes our model architecture used in LiLMaps. During incremental mapping, the future observed objects and their encoded representations are unknown in advance, which makes it challenging to train a language decoder

to represent all language features. To address this issue, in Sec. 3.2 we propose the adaptive language decoder optimization strategy that can effectively adjust the decoder to new language features while retaining previously observed ones. In this work, we employ a visual language encoder but pixel-wise language features are often inconsistent between frames. We address this issue in Sec. 3.3.

3.1. LiLMaps Architecture

Fig. 2 shows the architecture of the proposed approach. Input data for our pipeline are point clouds associated with CLIP language features (language point clouds), as well as camera poses estimated by any external SLAM method. We produce language point clouds by extracting language features from an RGB image. Extracted language features are projected to the point clouds in the world coordinate system using the corresponding depth image and camera pose. The extraction of language features can be done using per-pixel visual language encoders such as LSeg, OpenSeg, Segment-Anything-CLIP [22], etc. For example, VLMaps [15] utilizes LSeg for this purpose. In this study, the visual language encoder is treated as an external module, and improving its performance is not our focus.

Our goal is to enable implicit vision-language mapping under the conditions of environment exploration when future measurements are not available. We use the octree structure as positional encoding to build the implicit representation. Unless otherwise specified, we consistently use three different levels of the octree to store the features.

Each level of the octree is made up of voxels, and each voxel from a higher level can encompass multiple voxels from lower, more detailed levels. It should be noted that we use a sparse octree representation, meaning voxels are only present where observations have been made. When the point clouds are projected to the world coordinates, we find the corresponding voxels in the octree for each point. Each voxel holds learnable features at its corners. These features are shared among voxels that have common corners.

The language features have a high dimensionality and a straightforward solution to encode them in the octree is to increase the size of the learnable features stored in the corners. However, storing high-dimensional learnable features consumes a significant amount of memory. Instead, we suggest storing one high-dimensional learnable feature vector F per voxel of the first (coarse) octree level, while keeping the corner features low-dimensional.

To train the implicit representation, we apply the cosine similarity loss between the features decoded from the octree $\bar{\varphi}$ and the vision-language features φ of the input point cloud:

$$\mathcal{L}_{vl} = -\frac{1}{N} \sum_{i=1}^N \text{CosineSimilarity}(\varphi_i, \bar{\varphi}_i), \quad (1)$$

where N is the number of points with vision-language features in the point cloud.

Note that we encode every new available measurement into the learnable features using (1), but the weights of the language decoder are not updated with this loss. The optimization of the decoder is described in Sec. 3.2.

The decoder reconstructs the language feature $\bar{\varphi}$ in spatial coordinates using feature vector F and corners features of the corresponding voxels. Before feeding to the decoder, the corner features are linearly interpolated into the reconstruction point and summed across all octree levels. The decoder consists of three fully connected layers. The first two layers expand the dimensions of the corner features and produce the element-wise scaling vector for F , which is then multiplied by it and passed to the last fully connected layer, which outputs the predicted language feature $\bar{\varphi}$.

3.2. Adaptive Language Decoder Optimization

Equation (1) uses our MLP-based language decoder to predict vision-language features $\bar{\varphi}$ based on the encodings stored in the octree. However, the new data can contain features that have not been observed before. In this case, the decoder weights must be updated to be able to reconstruct new features without forgetting the old ones, but re-training of the whole previously mapped environment is computationally expensive.

We propose adaptive language decoder optimization in Algorithm 1. When a new point cloud with vision-language features arrives, we perform decoder optimization before the optimization of octree features described in Sec. 3.1.

Algorithm 1 Adaptive Language Decoder Optimization

```

1: global names
2:   LDec                                ▷ Language decoder
3:   kFeatures ← ∅                        ▷ Known features
4:   Encodings ← ∅                        ▷ Learnable Parameter
5:   FVectors ← ∅                         ▷ Learnable Parameter
6:   τ                                    ▷ Cosine similarity threshold
7:   LOSS                                  ▷ Cosine similarity loss
8: end global names
9: procedure OPTIMIZE(inFeatures)
10:  uniqueFeatures ← UNIQUE(inFeatures, τ)
11:  newFeatures ← UNKNOWN(uniqueFeatures,
                          kFeatures, τ)
12:  if len(newFeatures) = 0 then
13:    return                                ▷ No optimization if no new features
14:  end if
15:  encodings ← RANDOM(len(newFeatures), m)
16:  if len(FVectors) = 0 then
17:    fvectors ← RANDOM(1, L)
18:  else
19:    fvectors ← MEAN(FVectors, dim = 0,
                    keepdim = true)
20:  end if
21:  optimizer ← ADAM([encodings, fvectors, LDec])
22:  allEncodings ← Encodings ∪ encodings
23:  allFVectors ← FVectors ∪ fvectors
24:  allFeatures ← kFeatures ∪ newFeatures
25:  for  $i \leftarrow 0$  to  $N_{\text{opt}}$  do
26:     $\bar{\varphi} \leftarrow$  LDec(allEncodings, allFVectors)
27:    fvectors1 ← SHUFFLE(allFVectors)
28:    loss ← LOSS(allFeatures,  $\bar{\varphi}$ )
29:    loss $_F$  ← LOSS(fvectors1, allFVectors)
30:    optimizer.optimize(loss + loss $_F$ )
31:  end for
32:  kFeatures ← allFeatures                ▷ Update features
33:  Encodings ← allEncodings              ▷ Update encodings
34:  FVectors ← allFVectors                ▷ Update FVectors
35: end procedure

```

The proposed optimization operates with a language decoder (Line 2) and learnable parameters. The learnable parameters are the inputs that are directly forwarded to the decoder. In our work (Fig. 2), the learnable parameters are F vectors (Line 5) and the interpolated and summed point encoding (Line 4). Note that Algorithm 1 is not limited to our network architecture and can be adapted to other implicit representations by simply replacing the corresponding learnable parameters.

Firstly, we extract only unique language features (Line 10) from all available ones in the input point cloud and then filter out already known features (Line 11). In both cases, cosine similarity and a predefined threshold τ

Parameter	Symbol	Value
Similarity threshold	τ	0.02
Used octree levels		8,9,10
Fine level resolution		0.05 [m]
Learnable features size	m	16
F vectors size	L	512
Iterations per decoder optimization	N_{opt}	100
Iterations per mapping loss Eq. (1)		100

Table 1. LiLMaps parameters and their values in the experiments.

are used to estimate similarity.

For the new features (unobserved features without duplicates), we initialize the learnable parameters: the encodings and F vectors (Lines 15, 17 and 19). Note that the initialized encodings (Line 15) correspond to the linearly interpolated and summed features of the corners of the octree (Point Encoding in Fig. 2). We initialize only a single F feature vector for all new language features (Line 17) because the F feature vector stored in the coarse octree level may be used for points with different language features (Sec. 3.1). We also regularize new F vectors by enforcing them to be similar to existing ones (Lines 27 and 29). However, this regularization is optional and can be omitted or changed to any other regularization required by the corresponding implicit representation.

The proposed adaptive optimization approach optimizes the decoder for unobserved language features and finds the learnable parameters for them. Only the decoder and new learnable parameters are optimized (Line 21). The already known features and their encodings are not optimized, but used to prevent forgetting (Lines 22 to 24). After optimization, we update the list of known vision-language features and their learnable parameters (Lines 32 to 34).

The proposed optimization efficiently stores only a small number of known features for replay (Line 3), as demonstrated in experiments (Fig. 6). It enables fast decoder optimization using vectorization. Moreover, the language decoder and the learnable parameters are optimized only if new language features are observed.

3.3. Measurement Update Strategy

During incremental mapping, new observations added to the map should not corrupt previous measurements. However, vision-language features predicted by the visual encoder may not be consistent between frames. VLMaps [15] averaged the language features of the objects received from different views. We note that the averaging can be done in a recursive form:

$$\mathcal{A}_n = \frac{1}{N} \sum_{i=1}^N \varphi_i = \frac{n-1}{n} \mathcal{A}_{n-1} + \frac{\varphi_n}{n}, \text{ with } \mathcal{A}_0 = 0. \quad (2)$$



Figure 3. *Left*: Environments reconstructed without measurement update; *Middle*: Ground Truth; *Right*: Environments reconstructed with measurement update.

In this work, we propose to use as target for training in Eq. (1) a weighted average φ_n^* between observations φ_n and the features $\bar{\varphi}_{n-1}$ already stored in the map:

$$\varphi_n^* = \frac{n-1}{n} \bar{\varphi}_{n-1} + \frac{\varphi_n}{n}, \text{ with } \bar{\varphi}_0 = 0. \quad (3)$$

This averaging is especially useful for noisy measurements such as vision-language features because they may significantly vary with distance to the objects or point of view (Sec. 4.1). Mapping with Eq. (3) forces the map to store all observations similarly to [15]. However, we observed better results when not all previous data are stored and decided to use exponential smoothing instead of averaging:

$$\varphi_n^* = \alpha \bar{\varphi}_{n-1} + (1-\alpha) \varphi_n \text{ with } \bar{\varphi}_0 = 0, \quad (4)$$

where α is set dynamically to higher values if new measurements φ_n are more different from the previously optimized map features $\bar{\varphi}_{n-1}$ and lower otherwise:

$$\alpha = \frac{\text{CosineSimilarity}(\varphi_i, \bar{\varphi}_i)}{0.5 + \text{CosineSimilarity}(\varphi_i, \bar{\varphi}_i)}. \quad (5)$$

4. Experiments

In the experiments, we validate that our method can be used for incremental implicit mapping of language features. We use depth, semantic, and RGB images provided by [15] through the Habitat simulator using Matterport3D [5]. Matterport3D provides ground truth meshes with each face assigned to a class label. To get ground truth point clouds with the corresponding language features, we sample the meshes and encode their labels using CLIP [33]. During all experiments, we project depth images into 3D using the current camera pose to obtain input point clouds. The parameters we use during the experiments are summarized in Tab. 1.

4.1. Mapping Quality

We evaluate accuracy, recall, precision, and intersection over union for our implicit language map. Accuracy

Sequence \ Approach	5LpN3gDmAk7_1				YmJkqBEsHnH_1				gTV8FGcVJC9_1				jh4fc5c5qoQ_1				JmbYfDe2QKZ_2			
	A	mR	mP	mIoU	A	mR	mP	mIoU	A	mR	mP	mIoU	A	mR	mP	mIoU	A	mR	mP	mIoU
LiLMaps* _{GT}	97	97	96	93	98	93	92	89	98	98	95	94	98	96	88	85	95	96	92	89
LiLMaps _{GT}	97	97	93	91	96	94	96	90	97	98	93	92	98	98	92	89	95	96	90	87
LiLMaps* _{SEM}	84	77	70	57	84	77	82	65	85	85	84	73	83	78	73	61	78	77	77	63
LiLMaps _{SEM}	88	86	75	66	90	82	87	72	90	90	86	79	90	84	82	71	85	88	82	74
OpenScene [31]	68	45	67	36	63	50	77	41	61	49	60	36	77	52	59	39	56	51	67	41
LiLMaps* _{Lseg}	64	29	46	21	52	44	58	31	65	48	59	32	59	32	44	21	53	41	50	33
LiLMaps _{Lseg}	68	37	57	26	57	56	60	34	70	50	63	33	73	39	58	27	56	42	56	34
VLMaP _{2D} ^{mean}	28	-	-	19	28	-	-	19	28	-	-	19	28	-	-	19	28	-	-	19

Table 2. Language mapping quality evaluation: accuracy (A), recall (mR), precision (mP) and mean IoU (mIoU) in [%].




Figure 4. *Left*: Language map produced by OpenScene 3D [31]; *Middle*: Ground Truth; *Right*: Language map created by LiLMaps.

is defined as the number of points with correctly reconstructed language features divided by the total number of points. The values are compared with the OpenScene 3D model [31] trained on Matterport3D [5]. The results of random sequences are presented in Tab. 2. To validate the measurement update strategy, we also present results with deactivated measurement update, marked by an asterisk (LiLMaps*).

For LiLMaps_{GT} and LiLMaps*_{GT}, the language features of a measurement are obtained directly from the closest points of the ground truth point cloud. This allows us to estimate upper-bound performance with close-to-ideal input data. However, simulation measurements and ground truth points sampled from uneven meshes do not always match perfectly. As a result, some input points may have wrong language features or do not have language features at all.

LiLMaps_{SEM} and LiLMaps*_{SEM} denote experiments in which language features are extracted from semantic images. Simulated semantic images have incorrect labels due to mesh discontinuities and on object edges. This allows us to test LiLMaps when the input data are closer to the real ones, e.g. when the data are imprecise and inconsistent between frames.

Our approach demonstrates the best performance when used with the GT data. The introduction of the Measurements Update does not change the performance signifi-



Figure 5. Language map incrementally created with our adaptive optimization. *Bottom Left*: A region mapped in the beginning. *Bottom Right*: The same region after the mapping is completed. All initially mapped objects remain unchanged.

cantly, as the input features are precise and consistent between frames in this case.

As expected, LiLMaps_{SEM} and LiLMaps*_{SEM} yield worse results due to the inconsistency of the input data, but both still outperform OpenScene [31]. LiLMaps_{SEM} outperforms LiLMaps*_{SEM} due to the proposed measurement update technique which addresses potential data inconsistencies. Fig. 3 shows maps learned with activated and deactivated measurement update procedure. Enabling measure-

Decoders	F1-Score				
	↑100% – 90%	↓90% – 80%	↓↓80% – 70%	↓↓↓70% – 50%	↓↓↓↓50% – 0%
LiLMaps _{Adaptive} ^{simple}	20	2	1 (picture)	0	0
LiLMaps _{Pretrained} ^{simple}	20	2	0	1 (towel)	0
OpenScene _{HEAD} ^{MT}	16	4	2	0	1 (objects)
OpenScene _{HEAD} ^{SC}	16	5	1	0	1 (objects)
OpenScene _{HEAD} ^{NS}	9	8	1	3	2 (picture, shelving)

Table 3. Number of classes falling into different F1-score ranges for compared decoders. A method is considered to be better if a larger number of classes are listed in the first column with F1 > 90%.

ments update results in a cleaner final map, which is crucial for object detection and navigation.

Fig. 4 compares a LiLMaps reconstruction with the prediction of the OpenScene 3D model. Despite OpenScene being trained on the same dataset, it struggles with certain labels and completely misses labels such as "TV monitor", "appliances", "stool", whereas our approach achieves high accuracy for these labels and does not completely miss objects.

Tab. 2 compares our approach combined with the LSeg model (LiLMaps_{LSeg} and LiLMaps*_{LSeg}) and VLMaps (VLMaps_{2D}^{mean}) with metrics reported in [15]. LSeg frequently misses objects (e.g., segments a painting as a wall) or provides wrong language features (e.g., detects a bed as a sofa), which significantly influence the final results. Improving the quality of per-pixel language segmentation is beyond the scope of this research, however. In all cases, our 3D reconstructed language maps yield better results than the mean results reported in VLMaps [15] for their 2D maps.

4.2. Adaptive Language Decoder Optimization

We demonstrate the impact of our Adaptive Optimization Strategy on sequence 5LpN3gDmAk7_1 with GT labels. We compare the performance of the decoder trained with our Adaptive Optimization with other decoders built from pre-trained models. We chose OpenScene [31] 3D model’s head as the pre-trained decoder because it can predict language features for arbitrary input point clouds. For fair comparison, we changed our language decoder architecture (LiLMaps^{simple}) to match the architecture of OpenScene’s head. This head does not allow us to use the learnable vectors F , however, and therefore the results in Tab. 3 are presented when only the learnable corner features with $m = 96$ are optimized.

Tab. 3 summarizes number of classes with distinct F1-score qualities using different optimization types: decoder trained with the proposed adaptive optimization (LiLMaps_{adaptive}^{simple}); decoder pre-trained with the proposed




Figure 6. Number of language features stored for the adaptive optimization with varying feature similarity thresholds τ . *Blue*: Language features are extracted from ground truth (GT) data; *Orange*: Language features are extracted by LSeg; *Red line*: Total number of different GT classes presented in the scene.

optimization (LiLMaps_{pretrained}^{simple}); and pre-trained and fixed headings of OpenScene [31], trained on different datasets (Matterport3D [5], nuScenes [4], ScanNet [8]).

To obtain LiLMaps_{pretrained}^{simple} we extract all available labels from the scene, convert them to language features using CLIP and use our adaptive language decoder optimization with all features at once. The additional possibility of pre-training the decoder in advance without any real measurements may be useful for some applications. Our adaptive language decoder optimization allows to adjust pre-trained decoders online if necessary, but for this experiment, we do not update pre-trained models (LiLMaps_{pretrained}^{simple}, OpenScene_{HEAD}^{MT}, OpenScene_{HEAD}^{SC}, OpenScene_{HEAD}^{NS}) during the mapping.

The final results of LiLMaps_{adaptive}^{simple} are similar to those of LiLMaps_{pretrained}^{simple} because every time a new object is observed, the corresponding language features are included in the Adaptive Optimization and the decoder of

LiLMaps^{simple}_{adaptive} is updated to represent new features without forgetting the old ones. Adaptive and pre-trained LiLMaps models may have minor differences in the results due to their initial states and optimization processes being different. In Fig. 5 we demonstrate that the proposed adaptive optimization can incrementally extend the decoder to represent new features without catastrophic forgetting of language features observed in the beginning.

The results of Tab. 3 show that the proposed adaptive optimization LiLMaps^{simple}_{Adaptive} performs better than pre-trained and fixed models. Our adaptive optimization fits the model to a specific scene while pre-trained decoders (in this case OpenScene’s heads) are trained for general language prediction. If a model trained for general prediction is used, then some language features of the environment may be poorly represented in it (e.g. picture and shelving in Tab. 3), while well-represented features may be irrelevant for the specific scene. This can be seen in the results of OpenScene^{NS}_{HEAD}. OpenScene^{MT}_{HEAD} is trained on Matterport3D [5] and has better results. OpenScene^{SC}_{HEAD} is trained on ScanNet [8] which is similar to Matterport3D that explains the similar results. However, OpenScene^{NS}_{HEAD} pre-trained on NuScenes [4] has significant degradation in the results because the NuScenes environment is more different from Matterport3D. Moreover, our adaptive language decoder optimization allows one to build a custom decoder architecture while employing pre-trained models could restrict available architecture options.

We analyze different values of the threshold τ used to extract unique and unknown features from all input features. Fig. 6 shows the final number of features that were considered distinguished and were involved in the optimization at the end of mapping. Lower τ values lead to a larger number of features, but they are still memory efficient. For comparison, the stored features are collected from hundreds of high-resolution images, but their final number is less than 0.5% of the number of pixels in a single image with resolution 640×480. During all tests, adaptive language decoder optimization operated at a rate of 4 frames per second (fps). To achieve real-time performance, adaptive optimization can be executed in parallel to mapping.

5. Conclusion

In this work, we presented an implicit language mapping approach called LiLMaps. We address the problem of unseen language features that appear during the mapping process and the problem of inconsistencies between frames. Currently, LiLMaps is the only approach capable of large-scale incremental implicit language mapping. It can be used alone and enables a variety of interactions with the environment, for instance, 3D language-based object detection (Fig. 7). Additionally, it can be integrated into existing implicit mapping approaches, introducing only slight over-

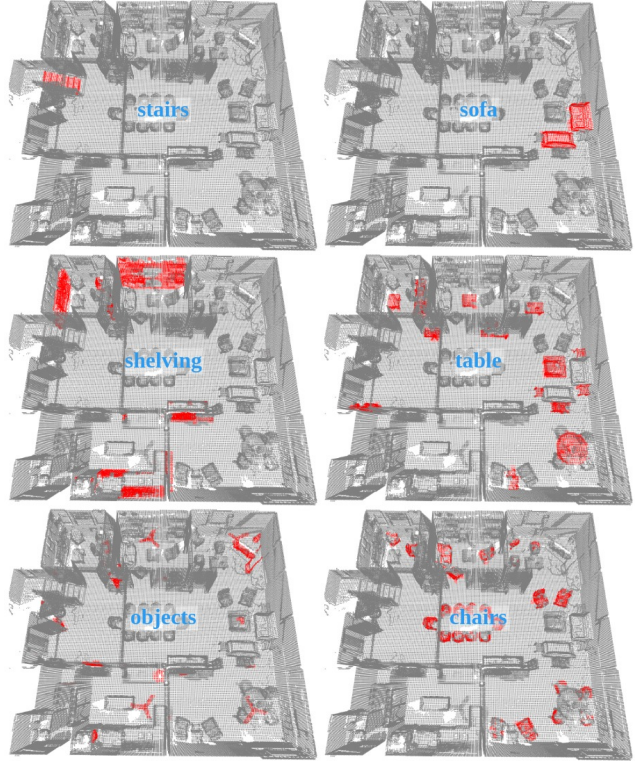


Figure 7. 3D language-based object detection performed on our language map. LiLMaps creates an implicit language map which is reconstructed and queried for different objects. The highest correspondences between the reconstructed language map and corresponding request (blue) are highlighted in red.

head.

We evaluated LiLMaps on the public dataset commonly used in related works. Based on the results, we outperform similar works in terms of language mapping quality. However, LiLMaps significantly depends on the quality of visual language features produced by the encoder, which is considered to be an external module in our study. The importance of this dependency is reduced by the proposed Measurement Update strategy that handles inconsistency between frames. We demonstrated that LiLMaps can adapt to the environment outperforming pre-trained decoders. The proposed Adaptive Optimization demonstrates the ability to prepare decoders given arbitrary language features without the need for actual observations.

Acknowledgment

This research was funded by the German Federal Ministry of Education and Research (BMBF) in the project WestAI – AI Service Center West, grant no. 01IS22094A.

References

- [1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as I can, not as I say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. 2
- [2] Peter Anderson, Qi Wu, Damien Teney, Jake Bruce, Mark Johnson, Niko Sünderhauf, Ian Reid, Stephen Gould, and Anton Van Den Hengel. Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3674–3683, 2018. 2
- [3] Michael Bloesch, Jan Czarnowski, Ronald Clark, Stefan Leutenegger, and Andrew J Davison. CodeSLAM – learning a compact, optimisable representation for dense visual SLAM. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2560–2568, 2018. 2
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multi-modal dataset for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11621–11631, 2020. 7, 8
- [5] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3D: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, pages 667–676, 2017. 1, 5, 6, 7, 8
- [6] Boyuan Chen, Fei Xia, Brian Ichter, Kanishka Rao, Keerthana Gopalakrishnan, Michael S Ryoo, Austin Stone, and Daniel Kappler. Open-vocabulary queryable scene representations for real world planning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 11509–11522, 2023. 2
- [7] William Chen, Siyi Hu, Rajat Talak, and Luca Carlone. Leveraging large language models for robot 3D scene understanding. *arXiv preprint arXiv:2209.05629*, 2022. 2
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2432–2443, 2017. 7, 8
- [9] Murtaza Dalal, Tarun Chiruvolu, Devendra Chaplot, and Ruslan Salakhutdinov. Plan-Seq-Learn: Language model guided RL for solving long horizon robotics tasks. In *12th International Conference on Learning Representations (ICLR)*, 2024. 2
- [10] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European Conference on Computer Vision (ECCV)*, pages 540–557. Springer, 2022. 2
- [11] Huy Ha and Shuran Song. Semantic Abstraction: Open-world 3D scene understanding from 2D vision-language models. In *Conference on Robot Learning (CoRL)*, volume 205 of *Proceedings of Machine Learning Research*, pages 643–653. PMLR, 2022. 2
- [12] Iain Houghton, Edgar Sucar, André Mouton, Edward Johns, and Andrew J. Davison. Real-time mapping of physical scene properties with an autonomous robot experimenter. In *Conference on Robot Learning (CoRL)*, volume 205 of *Proceedings of Machine Learning Research*, pages 118–127. PMLR, 2022. 2, 3
- [13] Yicong Hong, Zun Wang, Qi Wu, and Stephen Gould. Bridging the gap between learning in discrete and continuous environments for vision-and-language navigation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15439–15449, 2022. 2
- [14] Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim, Yaqi Xie, Tianyi Zhang, Zhibo Zhao, et al. Toward general-purpose robots via foundation models: A survey and meta-analysis. *arXiv preprint arXiv:2312.08782*, 2023. 2
- [15] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2023. 2, 3, 5, 7
- [16] Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd. Omama, Ganesh Iyer, Soroush Saryazdi, Tao Chen, Alaa Maalouf, Shuang Li, Nikhil Varma Keetha, Ayush Tewari, Joshua B. Tenenbaum, Celso Miguel de Melo, K. Madhava Krishna, Liam Paull, Florian Shkurti, and Antonio Torralba. ConceptFusion: Open-set multimodal 3D mapping. In *Robotics: Science and Systems XIX (RSS)*, 2023. 2
- [17] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)*, 42(4):139–1, 2023. 2, 3
- [18] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. LERF: Language embedded radiance fields. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 19729–19739, 2023. 2
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026, 2023. 2
- [20] Boyi Li, Kilian Q. Weinberger, Serge J. Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. In *10th International Conference on Learning Representations (ICLR)*, 2022. 2
- [21] Jiaze Li, Zhengyu Wen, Luo Zhang, Jiangbei Hu, Fei Hou, Zhebin Zhang, and Ying He. GS-Octree: Octree-based 3D Gaussian splatting for robust object-level 3D reconstruction under strong lighting. *Computer Graphics Forum (CGF)*, 43(7):i–xxii, 2024. 3
- [22] Ming-Feng Li. Per-pixel features: Mating segment-anything with CLIP, 2023. 2, 3
- [23] Corey Lynch, Ayzaan Wahid, Jonathan Tompson, Tianli Ding, James Betker, Robert Baruch, Travis Armstrong, and Pete Florence. Interactive language: Talking to robots in real time. *IEEE Robotics and Automation Letters (RA-L)*, 2023. 2
- [24] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting SLAM. In *IEEE/CVF Confer-*

- ence on Computer Vision and Pattern Recognition (CVPR), pages 18039–18048, 2024. 2
- [25] Kirill Mazur, Edgar Sucar, and Andrew J Davison. Feature-realistic neural fusion for real-time, open set scene understanding. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 8201–8207, 2023. 2, 3
- [26] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. NeRF: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 2, 3
- [27] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022. 3
- [28] Jens Naumann, Binbin Xu, Stefan Leutenegger, and Xingxing Zuo. NeRF-VO: Real-time sparse visual odometry with neural radiance fields. *IEEE Robotics and Automation Letters (RA-L)*, 9(8):7278–7285, 2024. 2
- [29] Yoshiki Obinata, Naoaki Kanazawa, Kento Kawaharazuka, Iori Yanokura, Soonhyo Kim, Kei Okada, and Masayuki Inaba. Foundation model based open vocabulary task planning and executive system for general purpose service robots. *arXiv preprint arXiv:2308.03357*, 2023. 2
- [30] Joseph Ortiz, Alexander Clegg, Jing Dong, Edgar Sucar, David Novotný, Michael Zollhöfer, and Mustafa Mukadam. iSDF: Real-time neural signed distance fields for robot perception. In *Robotics: Science and Systems XVIII (RSS)*, 2022. 2
- [31] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. OpenScene: 3D scene understanding with open vocabularies. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–824, 2023. 2, 6, 7
- [32] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. LangSplat: 3D language Gaussian splatting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20051–20060, 2024. 2
- [33] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, pages 8748–8763. PMLR, 2021. 1, 2, 5
- [34] Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. In *NeurIPS Foundation Models for Decision Making Workshop (FMDM)*, 2022. 2
- [35] Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. Perceptual grouping in contrastive vision-language models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5548–5561, 2023. 2
- [36] Christian Reiser, Songyou Peng, Yiyi Liao, and Andreas Geiger. KiloNeRF: Speeding up neural radiance fields with thousands of tiny MLPs. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 14335–14345, 2021. 2, 3
- [37] Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. LLM-Planner: Few-shot grounded planning for embodied agents with large language models. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2998–3009, 2023. 2
- [38] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. iMAP: Implicit mapping and positioning in real-time. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6229–6238, 2021. 2
- [39] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11358–11367, 2021. 3
- [40] Jingwen Wang, Tymoteusz Bleja, and Lourdes Agapito. GO-Surf: Neural feature grid optimization for fast, high-fidelity RGB-D surface reconstruction. In *International Conference on 3D Vision (3DV)*, pages 433–442. IEEE, 2022. 3
- [41] Yunhan Yang, Xiaoyang Wu, Tong He, Hengshuang Zhao, and Xihui Liu. SAM3D: Segment anything in 3D scenes. *arXiv preprint arXiv:2306.03908*, 2023. 2
- [42] Shuaifeng Zhi, Edgar Sucar, Andre Mouton, Iain Houghton, Tristan Laidlow, and Andrew J Davison. iLabel: Interactive neural scene labelling. *arXiv preprint arXiv:2111.14637*, 2021. 2
- [43] Xingguang Zhong, Yue Pan, Jens Behley, and Cyrill Stachniss. SHINE-Mapping: Large-scale 3D mapping using sparse hierarchical implicit neural representations. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 8371–8377, 2023. 2
- [44] Liyuan Zhu, Yue Li, Erik Sandström, Konrad Schindler, and Iro Armeni. LoopSplat: Loop closure by registering 3D Gaussian splats. *arXiv preprint arXiv:2408.10154*, 2024. 2
- [45] Siting Zhu, Renjie Qin, Guangming Wang, Jiuming Liu, and Hesheng Wang. SemGauss-SLAM: Dense semantic Gaussian splatting slam. *arXiv preprint arXiv:2403.07494*, 2024. 2
- [46] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. NICE-SLAM: Neural implicit scalable encoding for SLAM. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12786–12796, 2022. 2