

# QuASH: Using Natural-Language Heuristics to Query Visual-Language Robotic Maps

Matti Pekkanen, Francesco Verdoja, and Ville Kyrki

**Abstract**—Embeddings from Visual-Language Models are increasingly utilized to represent semantics in robotic maps, offering an open-vocabulary scene understanding that surpasses traditional, limited labels. Embeddings enable on-demand querying by comparing embedded user text prompts to map embeddings via a similarity metric. The key challenge in performing the task indicated in a query is that the robot must determine the parts of the environment relevant to the query.

This paper proposes a solution to this challenge. We leverage natural-language synonyms and antonyms associated with the query within the embedding space, applying heuristics to estimate the language space relevant to the query, and use that to train a classifier to partition the environment into matches and non-matches. We evaluate our method through extensive experiments, querying both maps and standard image benchmarks. The results demonstrate increased queryability of maps and images. Our querying technique is agnostic to the representation and encoder used, and requires limited training.

## I. INTRODUCTION

Robots must understand the semantic and geometric properties of their environment to perform complex tasks. While the ability of maps to capture semantics has increased with developments in computer vision, most semantic segmentation methods are trained to segment measurements into specified categories.

Visual-Language Models (VLMs) are networks that jointly train a visual and language encoder to learn a mapping from text and image inputs into a common visual-language latent space. Modern, efficient transformer architectures [1], such as CLIP [2], are trained with large data sets, which promises to enable these models to effectively have an *open vocabulary*, meaning visual semantics can be matched to natural language sentences instead of manually chosen categories, allowing richer representations of semantics and helping to overcome the difficulty of adapting to new environments.

Lately, many types of maps [3]–[7] using the embeddings of VLMs, *i.e.*, *latent-semantic maps* [8], have been proposed. The idea is that the visual perceptions of the robot are embedded with the visual encoder of a VLM, and these latent representations are stored as semantics in the map. These embeddings are then queried on demand by supplying the user query as text, embedding it, and comparing the query embedding to the embeddings in the map using a similarity

This work was supported by the Research Council of Finland (decision 354909). The authors acknowledge the use of the MIDAS infrastructure of the Aalto School of Electrical Engineering.

M. Pekkanen, F. Verdoja, and V. Kyrki are with the School of Electrical Engineering, Aalto University, Espoo, Finland. {firstname.lastname}@aalto.fi

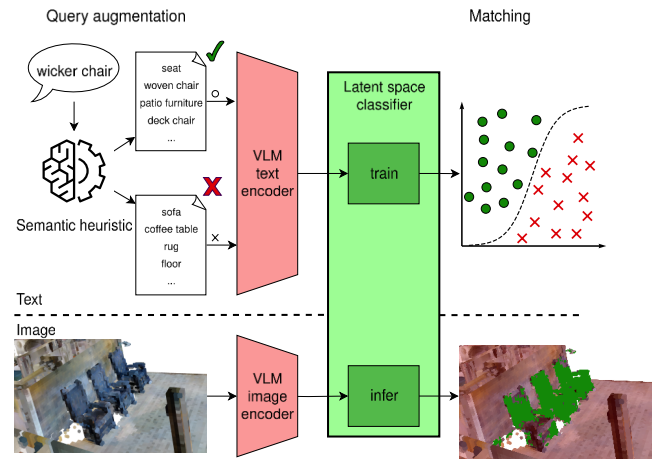


Fig. 1: Our method, QuASH, works by augmenting a single query with a semantic heuristic to acquire a set of semantic synonyms and antonyms to train a classifier in the latent space of a VLM. The classifier enables non-linear classification in the latent space, which in turn enables accurate estimation of the region in an image or a map that corresponds to the query.

metric, typically cosine similarity. Despite these advances, the key challenge remains: the selection of relevant matches.

Many methods have been proposed for matching a query to a latent-semantic map, *e.g.*, by contrasting the query embedding with embeddings of a negative query [4], or by using a threshold on a similarity metric [6]. While methods exist that can utilize direct similarities [9], most methods ultimately require a decision threshold: some parts of the map are considered matches, while others are not. However, the matching performance of the state-of-the-art is still relatively low, making this an open problem.

In this work, we propose a solution to this problem. We begin by formalizing the problem in terms of identifying and matching sets between the visual, language, and latent spaces. This allows us to formalize the process of querying as the process of classifying the latent-space embeddings. Furthermore, we show how the existing state-of-the-art solutions fit into our formalization.

We propose *Query Augmentation using Semantic Heuristics (QuASH)*, a method to query the embeddings, shown in Figure 1, by augmenting the query with a set of semantically close words, as well as producing a set of counterexamples, then training an off-the-shelf classifier on these sets. Our method is not dependent on explicit thresholds, does not

require expensive neural network training, and works with any encoder. We show that the proposed method outperforms the comparison with a single complementary query in both standard image benchmarks and in maps. Finally, we conduct extensive ablation studies of querying latent-semantic maps.

The main contributions of this paper are:

- i) A formalization of the problem of matching language queries to visual objects, which generalizes most common approaches in the state-of-the-art.
- ii) QuASH: a novel encoder-agnostic method<sup>1</sup> for querying VLM-embeddings. The method is evaluated in the context of querying maps and images.
- iii) Extensive ablation experiments on querying VLM-embeddings.

## II. RELATED WORK

In this work, we are concerned with the *queryability* of embeddings [8], which differs from two closely related tasks: open-vocabulary object detection and semantic segmentation.

Most state-of-the-art methods implicitly estimate the region in the embedding space corresponding to the query and use one of two approaches to find matches. First, thresholding the cosine similarities [6], [10] produces effectively a conic volume centered on the query  $q$ , shown in Figure 2a. This does not account for the variation in the relevance of the dimensions to  $q$ .

Furthermore, by using a single query  $q$ , the method cannot accurately estimate the extent of the region. The embeddings in the embedding space are located on a very narrow cone [11] on the surface of a hypersphere, which is so narrow that the embeddings are effectively on an  $(N - 1)$ -dimensional hyperplane. This means that the embeddings acquired by thresholding are in an  $(N - 1)$ -dimensional isotropic hypersphere.

Second, by comparing the embeddings to the query and a negative query, such as “other” [3], [4], [9], [12]. This produces an implicit estimate of the region by linearly splitting the embedding space into two regions by a hyperplane, shown in Figure 2b.

Unlike this work, the state-of-the-art methods always produce isotropic and connected estimates of the region. Given that the directions have semantic meaning, with most queries, we expect the directions not to be uniformly important with respect to the query. To estimate the shape of the region, in this work, we estimate the relative importance of the dimensions of the embedding space with respect to  $q$  from the spread of the sampled embeddings.

Other proposed ways to solve the thresholding problem are to train an additional neural network to classify the queries [13]–[16], use an Large Language Model (LLM) [7], [17]–[19], or fine-tune the VLM to the query [20].

Alternatively, many methods do not estimate the region at all, but instead use the whole set of similarities [5], [9], [21]–[25]. In this work, we aim to query the maps directly, without having to train a neural network.

<sup>1</sup><https://aalto-intelligent-robotics.github.io/quash/>




Fig. 2: The state-of-the-art solutions implicitly classify the embedding space linearly by thresholding cosine similarities (a) [6], or by using a single complementary embedding (b) [4]. By using a Support Vector Machine (SVM) with a non-linear kernel, our proposed method (c) can classify the embeddings non-linearly. Embeddings matching the query are visualized in red, and those not matching are visualized in blue. A diamond represents the query and its complement, and dashed lines represent the thresholds.

## III. PROBLEM STATEMENT

We consider the problem of determining a target object referred to by the query  $q$  defined in a language space  $\mathcal{L}$ . The search space is the internal world representation of the robot, which we denote as the visual space  $\mathcal{V}$ , *e.g.*, a map of the environment or an image from the camera. Elements  $v \in \mathcal{V}$  are either the pixels or voxels representing a portion of the space. We assume  $q$  to correspond to a subset  $\mathcal{V}_Q \subset \mathcal{V}$  representing the extent of the target object. The problem is, therefore, to find the mapping  $f_q : q \mapsto \mathcal{V}_Q$ .

The underlying assumption behind the existence of  $f_q$  shares some commonalities to that behind VLMs, which are designed to embed elements of  $\mathcal{L}$  and  $\mathcal{V}$  into the shared embedding space  $\mathcal{E}$  using a pair of embedding functions  $\varphi : \mathcal{L} \rightarrow \mathcal{E}$  and  $\varrho : \mathcal{V} \rightarrow \mathcal{E}$  such that for a textual description  $l \in \mathcal{L}$  and a corresponding visual point  $v \in \mathcal{V}$ ,  $\varphi(l) = \varrho(v)$ .

Therefore, given the VLM  $M = (\varphi, \varrho)$ , our problem corresponds to the problem of finding the set  $\mathcal{R}_Q \subset \mathcal{E} = \{\forall v \in \mathcal{V}_Q : \varrho(v)\}$  that contains all embeddings matching the object described by the query  $q$ .

## IV. METHODS

To solve this problem, we present QuASH. First, we formalize the process of querying latent-semantic maps in Section IV-A. Second, we present state-of-the-art solutions within our formalization in Section IV-B, and finally, propose a querying method by augmenting the query using semantic heuristics in Section IV-C.

### A. Formalization of the querying

Let  $\mathcal{Q} \subset \mathcal{L}$  be the preimage of  $\mathcal{R}_Q$  under  $\varphi$  and  $\bar{\mathcal{Q}} \subset \mathcal{L}$  be the preimage of  $\mathcal{E} \setminus \mathcal{R}_Q$  under  $\varphi$ . We assume  $q \in \mathcal{Q}$ . The set  $\mathcal{Q}$  contains the set of all semantically close words to  $q$ . To preserve intuition, we hereafter refer to them as *semantic synonyms*, notwithstanding that they are not literal synonyms.  $\bar{\mathcal{Q}}$ , on the other hand, contains all negative queries of the query  $q$ , which we will refer to as *semantic antonyms* for the aforementioned reason.




Fig. 3: The intuition behind our method is that, as the textual descriptions and visual features are embedded into the same space  $\mathcal{E}$ , training a classifier in the text domain will result in a good classifier in the visual domain. Instead of using just a single example, *i.e.*, the query  $q$ , represented as a black dot, we additionally sample semantically close words, represented as circles, to form a set  $\mathcal{Q}$ , as well as a non-matching queries  $\bar{q} \in \bar{\mathcal{Q}}$ , represented as crosses, from the language space  $\mathcal{L}$ . Using these, we aim to capture the extent of the region  $\mathcal{R}_Q$  in the embedding space  $\mathcal{E}$ . Knowing  $\mathcal{R}_Q$  enables finding the visual region  $\mathcal{V}_Q$  corresponding to the query  $q$ .

The common approach in literature is to estimate  $\mathcal{R}_Q$  using some classification function  $\psi_Q : \mathcal{E} \rightarrow \{0, 1\}$ ,  $\psi_Q = \mathbb{1}_{\mathcal{R}_Q}$ . To do this, most approaches rely on the implicit assumption that  $\psi_Q$  can be learned by embedding samples from  $\mathcal{Q}$  and  $\bar{\mathcal{Q}}$  using  $\varphi$  [3], [4], [6], [9], [10], [12].

More specifically, given a map or image  $\mathcal{V}_m \subset \mathcal{V}$ , with corresponding embeddings  $\mathcal{E}_m \subset \mathcal{E}$ , the embedding of the query  $e_q = \varphi(q)$  is used as representative sample from  $\mathcal{R}_Q$ , and an the embedding  $e_{\bar{q}} = \varphi(\bar{q})$  of an arbitrary antonym  $\bar{q} \in \bar{\mathcal{Q}}$  (*e.g.*, the word "other") is used as representative sample of  $\mathcal{E} \setminus \mathcal{R}_Q$ .

### B. Querying latent-semantic maps: the state of the art

The classifier  $\psi_Q$  is then often either  $d_{\cos}(\mathcal{E}_m, e_q) < \theta$  [6], [10], where cosine distance  $d_{\cos}$  is thresholded with  $\theta$ , yielding a hyperconic section of  $\mathcal{E}$ , shown in Figure 2a, or

$$\mathbb{1}_{e_q} \left( \operatorname{argmin}_{x \in \{e_q, e_{\bar{q}}\}} d_{\cos}(\mathcal{E}_m, x) \right), \quad (1)$$

where the embeddings  $\mathcal{E}_m$  with  $e_q$  and  $e_{\bar{q}}$  are compared, splitting  $\mathcal{E}$  with a hyperplane [3], [4], [9], [12], shown in Figure 2b.

Provided that the directions in the embedding space have semantic meaning, we expect that the directions of the embedding space have varying degrees of relevance with respect to  $q$ , which means that the subset  $\mathcal{R}_Q$  is neither necessarily isotropic nor a connected space, reducing the effectiveness of distance-based classifiers. We argue that

estimating the extent of  $\mathcal{R}_Q$  from a pair of points  $e_q$  and  $e_{\bar{q}}$  is not sufficient.

### C. Querying latent-semantic maps: QuASH

In practice, it is not feasible to acquire  $\mathcal{Q}$  or  $\bar{\mathcal{Q}}$ . We assume that  $\mathcal{Q}$  contains semantically similar phrases to the query. This means that they are "close" in natural language, and therefore, the intuition behind our method, shown in Figure 3, is to use a natural-language semantic heuristic to estimate the augmented query set  $\mathcal{Q}$ .

Sampling from  $\mathcal{Q}$  we acquire the semantic synonyms  $\mathcal{S}_q$  and sampling  $\bar{\mathcal{Q}}$  the semantic antonyms  $\mathcal{A}_q$ . The embedded samples of synonyms and antonyms form the training data set, with which a classifier  $\psi_Q$  is learned in the latent space  $\mathcal{E}$ . The classifier enables the estimation of the volume of  $\mathcal{R}_Q \in \mathcal{E}$ . If a non-linear classifier is used, a non-linear estimate of  $\mathcal{R}_Q$  can be acquired, as shown in Figure 2c.

The method has two key design choices. First, finding a suitable heuristic function  $h(q) = (\mathcal{S}_q, \mathcal{A}_q)$ , as the performance of the classifier depends on the quality of the samples. Second, finding a suitable classifier  $\psi_Q$  that can learn the region  $\mathcal{R}_Q$  effectively from the samples. We discuss the choices for  $h$  in Section V-D.3, and for  $\psi_Q$  in Section V-C.

Our method works with the following steps:

- 1) Given the textual query  $q$ , we create a set of  $N$  strings  $\mathcal{S}_q \sim \mathcal{Q}$ , semantic synonyms to  $q$ , and  $M$  strings  $\mathcal{A}_q \sim \bar{\mathcal{Q}}$  that are semantic antonyms to  $q$ . Additionally, we add to  $\mathcal{A}_q$  a set of generic negative queries found in the literature: "other", "background", "nothing".
- 2) Given an embedding function  $\varphi$ , we embed all strings  $\mathcal{E}_{\mathcal{S}_q} = \varphi(\mathcal{S}_q \cup q)$ , and  $\mathcal{E}_{\mathcal{A}_q} = \varphi(\mathcal{A}_q \cup \mathcal{C})$ .
- 3) We train a classifier  $\psi_Q$  with positive examples  $\mathcal{E}_{\mathcal{S}_q}$  and negative examples  $\mathcal{E}_{\mathcal{A}_q}$ .
- 4) Given a map  $\mathcal{M}$  where each item  $m_i$  contain an embedding  $e_i$ , we acquire a binary mask over the map  $\mathcal{V}_Q = \psi_Q(\mathcal{M})$

The decision boundary of  $\psi_Q$  forms the estimate of  $\mathcal{R}_Q$ .

## V. EXPERIMENTS

The main question that we aim to answer with the experiments is: Does using the natural-language heuristics improve the visual classification task?

To answer this question, we evaluated our proposed querying method in a series of experiments. We employ the evaluation method described in [8] to assess *queryability*.

### A. Baseline

In all experiments, we use as the *baseline* a method proposed in [4], where the embedding of each voxel is compared with the embeddings of the query and the complement, the string "other".

The baseline uses *prompt engineering* (PE), where both the query and the complement are inserted into a list of 63 strings, such as "There is  $q$  in the scene" or "a photo of  $q$ ". All the resulting strings are embedded, and the means of the embeddings are used as the query and the complement. The original method additionally post-processes the query results

and evaluates in 2D, but these steps decrease performance; therefore, our baseline method does not use them.

The embeddings that are closer to the query than to the complement with respect to the cosine similarity metric are chosen as the results. We evaluate the effect of prompt engineering in an ablation study, presented at Section V-D.1.

### B. Data set and metrics

In all map querying experiments, 10 maps, the same as in [4], were created from the Matterport3D data set [26], containing ground truth semantic segmentation of 40 labels  $L$ , using a 3D version of VLMs [4], [27], where the method works similarly, but is extended to 3D. We use the original proposed parametrization. We use two encoders, LSeg [28] and OpenSeg [29], with pre-trained weights provided by the original authors.

We removed seven negative or aggregate classes: “void”, “unlabeled”, “misc”, “objects”, “seating”, “furniture”, and “appliances”, from  $L$  to acquire query set  $L_q$  of 34 labels.

In all experiments, following [8], we compute precision, recall, F1-score, and Intersection over Union (IoU) as metrics and report the average over the data set. As the areas of interest are relatively small in our experiments, true negatives dominate accuracy, making it an unsuitable metric.

### C. Querying method

To test our proposed method of querying VLM embeddings, we evaluated querying maps, described in Section V-C.2, as well as performed a set of standard image benchmarks, described in Section V-C.1.

In our method, we prompted<sup>2</sup> GPT-4o to produce a list of 20 words that are: synonyms, closely related words, subclasses, superclasses, related objects, *etc.*, of the query to acquire  $S_q$  and similarly 20 antonyms to acquire  $A_q$ , which we used as the heuristic for  $Q$ .

As the classifier  $\psi$ , we evaluated all of the applicable off-the-shelf classifiers from the scikit-learn library, and decided to use SVM with Radial Basis Function (RBF) kernel with Euclidean and cosine similarity metrics, as it had the best performance.

1) *Image benchmarks*: First, we tested our method against the baseline in two widely used image benchmarks: COCO 80 and Pascal Context 459. Additionally, we use the set of images used in the map generation as a data set, denoted “Matterport”. In these experiments, only LSeg was used as the encoder.

2) *Querying maps*: Second, we tested our method in querying maps created as described in Section V-B.

We compared our method with the baseline. Additionally, in this experiment, we compare to another method which uses the best information available, which we will refer to as “ground truth”. We use this method to separate the evaluation of the generation of the  $Q$  and its use. The method uses the whole set of 1659  $L_{sup}$  classes present in the Matterport3D data set as the heuristic for  $Q$ .  $L_{sup}$  contains more specific

TABLE I: Results of image benchmark experiment with LSeg

Dataset	Method	F1	IoU	Precision	Recall
COCO	baseline	0.680	0.515	0.536	<b>0.929</b>
	Our SVM	<b>0.786</b>	<b>0.647</b>	<b>0.696</b>	0.901
	Our C-SVM	0.771	0.627	0.672	0.905
PC 459	baseline	0.654	0.486	0.540	<b>0.831</b>
	Our SVM	0.667	0.501	<b>0.693</b>	0.643
	Our C-SVM	<b>0.728</b>	<b>0.572</b>	0.680	0.783
Matterport	baseline	0.544	0.374	0.448	0.691
	Our SVM	<b>0.621</b>	<b>0.451</b>	<b>0.626</b>	0.616
	Our C-SVM	0.579	0.407	0.483	<b>0.724</b>

subclasses of the superclasses that compose the set  $L$ . For example,  $L$  contains the superclass “chair”, and  $L_{sup}$  contains the subclasses “armchair”, “dining chair”. So, for each query term in  $L$ , we selected up to 20 random specific subclasses of that query as synonyms and, as antonyms, 100 random other subclasses that do not correspond to the query.

### D. Ablation studies

1) *Prompt engineering*: We tested the effect of prompt engineering, as described in Section V-A, on the baseline method and our method.

2) *Antonyms*: We evaluated the baseline method with a set of antonyms of the query found in the literature: “other” [4], [12], [28], “object” [3], “things” [3], “stuff” [3], “texture” [3], “background” [30], and a zero vector [30], as well as “nothing” and “others”.

3) *Language distribution*: We evaluated four different methods of creating the query sets  $Q$  and  $\bar{Q}$ : LLM, lexicon, adjectives, and ground truth.

The “LLM” is the default implementation of our method, as described in Section V-C. The “ground truth” set is described in Section V-C.2.

In the “lexicon” method, we select 20 random words from the English lexicon for synonyms and 20 random words for antonyms. In the “adjectives” method, we produce a static list of 44 adjectives across seven types, *e.g.*, different colors, materials, and textures. For each query  $q$ , we add to  $S_q$  all singular adjectives (*e.g.*, color +  $q$ ), 10 randomly combined adjectives from two types (*e.g.*, color + material +  $q$ ), and five randomly combined adjectives from three types, totaling 59 random variations of the query. As antonyms, a similar set is created by using “other” instead of the query.

## VI. RESULTS

### A. Image benchmarks

From the results of the image benchmark experiments, presented in Table I, it can be seen that our method yields increases of 15% in F1-score on COCO, 9% in Pascal Context 459, and 15% in the Matterport data set. This supports the hypothesis that the knowledge of the distribution in language space is beneficial for visual classification.

Notably, recall is higher in the baseline. This is likely because the method of comparing the query with a single

<sup>2</sup><https://aalto-intelligent-robotics.github.io/quash/>



Fig. 4: Examples of predictions from the COCO data set. The original image is shown on the top row, our method using an SVM classifier is on the second row, and the baseline is on the third row. Our method is generally slightly more accurate in contouring objects, as shown in the first column, and can distinguish related objects better, shown in the second column, and is less prone to overprediction, as shown in the third and fourth columns.

antonym splits the embedding space in two by a hyperplane. This leads to half of the embedding space being predicted as positive, which can encompass the correct region in its entirety; however, it conversely yields a high number of false positives, as reflected in the low precision. Our method can alleviate this problem by providing a more accurate estimate of the correct volume in the embedding space.

Examples of predictions from the COCO data set are shown in Figure 4. While most of the predictions are very similar, our method performs more accurately in contouring the objects, as shown in the first column. Another common occurrence is that the baseline misclassifies visually similar and semantically close objects, such as a refrigerator and a microwave oven. In contrast, our method can distinguish between these concepts, as shown in the second column. At times, the single complement is so far from the image embeddings that the baseline dramatically extends the prediction to unrelated areas, shown in the third and fourth columns.

### B. Querying maps

From the results of the map querying experiment, presented in Table II, it can be seen that, with LSeg, our method using ground truth synonyms cannot significantly improve over the baseline, and our method with LLM-generated synonyms performs worse than the baseline. This highlights

the sensitivity of the choice of the synonyms and antonyms. Still, ultimately, this result is congruent with the hypothesis that with a good approximation of  $Q$ , the queryability results do improve.

Example queries from the data set are shown in Figure 5, where predicted binary masks from our method with the ground truth synonyms and antonyms, along with the baseline, are shown. From the Figure, we can see that the differences in LSeg are very minor. However, with OpenSeg, the baseline exhibits a high number of false positives, reflected in the low precision, exemplified in Figure 5f and 5g, whereas our method finds the correct voxels from the map in corresponding Figures 5b and 5c. The precision-recall curve of our method could be further optimized, as our method can sometimes suffer from reduced recall, as shown in the example of a failure case in Figure 5d.

When using OpenSeg, our method outperforms the baseline with 227% increase in F1-score. Even with LLM-generated synonyms and antonyms, our method yields 78% improvement on the F1-score. This significant increase suggests the potential for using encoders other than LSeg in robotic applications, where LSeg has performed the best with existing querying methods so far.

While OpenSeg appears to be more sensitive to thresholding, leading to lower overall results, the vocabulary seems to

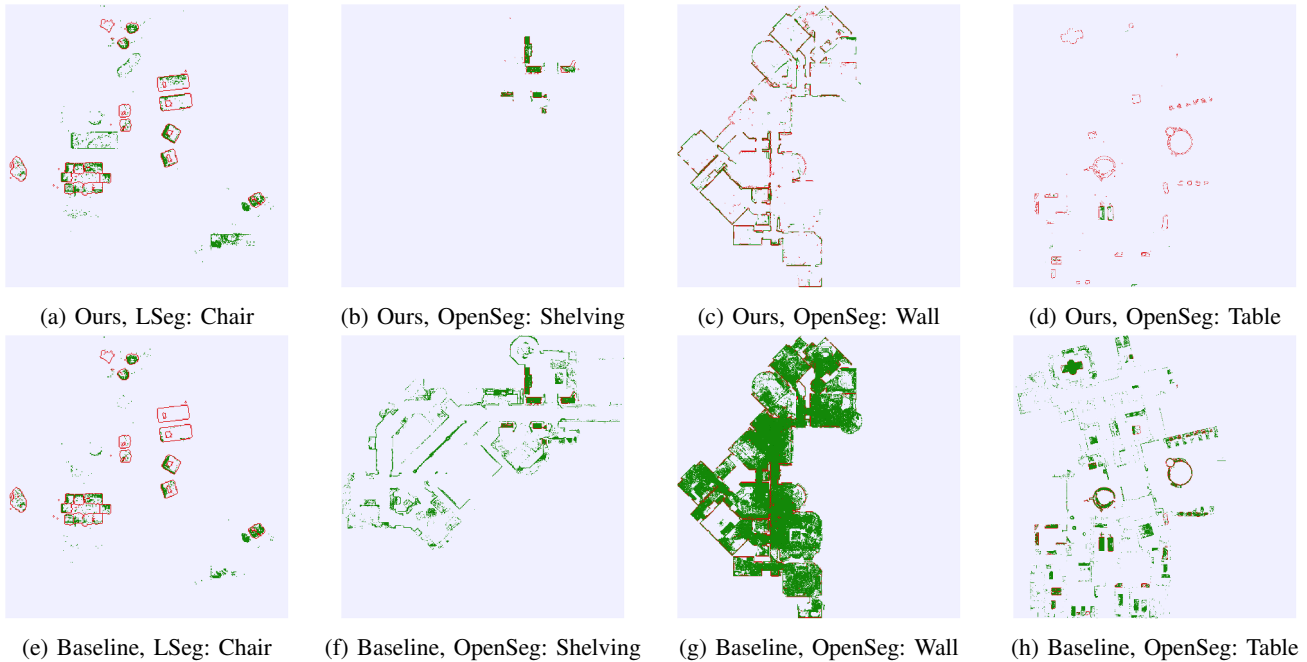


Fig. 5: Examples of map query experiment results. Our method with “ground truth” synonyms is shown in the top row, and the baseline is shown in the bottom row, with the same query displayed in each column. With LSeg (a), (e), the results are very similar. On the other hand, with OpenSeg, we can clearly see the difference, particularly in the significantly improved precision of our method. However, the drawback is that our method suffers from reduced recall, as evident in the failure cases in (d) and (h). The predictions are shown in green, and the outline of the ground truth in red.

be broader [8]. Although OpenSeg results are overall worse than LSeg, the results indicate that our method is effective at classifying embeddings from different encoders.

Interestingly, the results differ significantly between the Matterport data set and the maps. This implies that a modality gap exists between the maps and the images. The maps aggregate the measurements originating from the image data set, so the aggregation is likely the leading cause of this modality gap.

In VLMaps, the environment is represented as a voxel map, and therefore multiple embeddings are aggregated in each voxel. The mean of the embeddings is used as the aggregate. This can cause problems, especially in cases where the embeddings are multi-modal within the cell. Typically, each mode results from a set of measurements in which an object with a specific semantics is identified. However, the aggregate embedding produced by the mean of the distribution may not accurately reflect the original semantics, but may be a new concept altogether.

We evaluated different aggregation methods for the voxels, including the geometric median, selecting a representative sample, and training a transformer encoder to aggregate the embeddings. However, the geometric median yielded only 4% increase in F1-score, and therefore, cannot address this problem.

### C. Ablation studies

1) *Prompt engineering*: The ablation results from prompt engineering are presented in Table III. The prompt engineer-

TABLE II: Comparison of map querying methods in maps

Encoder	Method	F1	IoU	Precision	Recall
LSeg	baseline	0.561	0.395	0.536	0.589
	our ground truth	<b>0.574</b>	<b>0.408</b>	<b>0.567</b>	0.583
	our LLM	0.396	0.252	0.289	<b>0.656</b>
OpenSeg	baseline	0.136	0.073	0.073	<b>0.926</b>
	our ground truth	<b>0.445</b>	<b>0.296</b>	<b>0.622</b>	0.369
	our LLM	0.242	0.139	0.154	0.586

ing yields benefits only on the baseline with LSeg, as it was proposed for that purpose.

Interestingly, prompt engineering does not yield benefits with OpenSeg, whereas in the original work, prompt engineering was evaluated, but only in conjunction with ensembling, making a direct comparison difficult. This suggests that not all heuristics used for improving queries are directly generalizable across encoders. With our method, the prompt engineering did not yield improved performance.

This is understandable, as our method already uses more embeddings to estimate the distribution  $\mathcal{Q}$ . From this experiment, we can see that the majority of the gap between the baseline and our method is due to the prompt engineering. Prompt engineering aims to generalize the specific query to capture better the region  $\mathcal{Q}$ .

Even if ultimately only one embedding is used, the mean of the 63 prompts is better located in the embedding space

TABLE III: The effect of prompt engineering (PE) in map querying

Method	Encoder	PE	F1	Precision	Recall
baseline	LSeg	yes	<b>0.561</b>	<b>0.536</b>	0.589
		no	0.407	0.301	<b>0.651</b>
baseline	OpenSeg	yes	0.136	0.073	0.926
		no	0.136	0.074	0.931
C-SVM	LSeg	yes	0.396	0.289	0.656
		no	0.396	0.289	0.656
C-SVM	OpenSeg	yes	0.242	0.154	0.586
		no	0.242	0.154	0.586

TABLE IV: The comparison of antonyms in map querying with baseline

Encoder	Antonym	F1	Precision	Recall
LSeg	other	<b>0.561</b>	<b>0.536</b>	0.589
	others	0.546	0.500	0.604
	background	0.463	0.363	0.649
	texture	0.460	0.361	0.657
	stuff	0.451	0.346	0.651
	nothing	0.393	0.279	0.680
	object	0.358	0.240	0.709
	things	0.346	0.231	0.698
	zero vector*	0.048	0.025	<b>1.000</b>
	OpenSeg	other	0.136	0.073
others		0.127	0.068	0.936
background		0.131	0.071	0.935
texture		<b>0.193</b>	<b>0.110</b>	0.870
stuff		0.114	0.061	0.947
nothing		0.185	0.104	0.890
object		0.118	0.063	0.934
things		0.123	0.066	0.939
zero vector*		0.065	0.033	<b>0.991</b>

\* Prompt engineering not used

to capture  $\mathcal{R}_Q$ , and therefore this result is congruent to our hypothesis.

2) *Antonyms*: The results of the ablation experiment with different antonyms are presented in Table IV. With LSeg, “other” is the best single word. This is expected, as it is used in LSeg as the background label, and subsequently proposed to be used in VLMaps.

Interestingly, the distribution of the results differs between LSeg and OpenSeg; in LSeg, there is a greater spread between any two words, where the interval is 0.215, whereas in OpenSeg, the interval is only 0.079, without taking into account the zero vectors, which perform significantly worse than any other antonym with both encoders.

3) *Synonyms*: The results of the ablation experiment in terms of synonym generation are presented in Table V. A suitable set of synonyms is essential to match the targets accurately, as evident from the low precision of the lexicon, which serves as the random baseline. However, to get the best results, a suitable antonym set is needed in conjunction with the synonyms. This is highlighted, *e.g.*, with the adjective antonyms reducing the accuracy of the ground truth synonyms.

TABLE V: The comparison of different query set generation methods in map querying with our method with LSeg

Synonyms	Antonyms	F1	Precision	Recall
ground truth	ground truth	<b>0.576</b>	0.565	<b>0.587</b>
adjectives	ground truth	<b>0.575</b>	<b>0.583</b>	0.568
adjectives	LLM	0.555	0.534	0.581
ground truth	LLM	0.528	0.469	0.608
adjectives	lexicon	0.518	0.438	0.639
LLM	ground truth	0.462	0.476	0.469
LLM	LLM	0.396	0.289	0.656
ground truth	lexicon	0.378	0.261	0.697
adjectives	adjectives	0.362	0.246	0.700
ground truth	adjectives	0.279	0.174	0.740
LLM	lexicon	0.236	0.141	0.729
LLM	adjectives	0.142	0.078	0.859
lexicon	adjectives	0.049	0.025	0.995
lexicon	lexicon	0.048	0.025	0.994
lexicon	ground truth	0.040	0.036	0.047
lexicon	LLM	0.024	0.013	0.116

The ground truth provides the best overall results, as none of the synonyms are other classes present in the evaluation dataset, and antonyms correspond to the other classes present in the environment, which effectively enhances the precision of the detections. These results are the best overall query results in maps and support the hypothesis that knowledge of language distribution improves visual classification results.

While the adjective method performs well, it cannot be used in querying specific object instances. In this experiment, we evaluate the queryability of generic classes (*e.g.*, “chair”, “table”, “painting”), which are more effectively captured by adjectives that expand the region in the embedding space. However, when a particular object is queried (*e.g.*, “the red metal chair”), the adjectives force the generalization of the query, therefore making it unusable in an open-vocabulary context. However, adjectives perform poorly as antonyms, as they are in the form adjective + “other”, which is semantically meaningless.

The LLM is the most versatile method, as it can produce synonyms for any word. The downside is the unreliability of the predictions as well as the presence of other classes in the synonyms, which reduces the prediction accuracy.

## VII. CONCLUSION

In this work, we argue that using a single query embedding and a single antonym is not sufficient to perform visual classification on VLM embeddings in maps or images. Leveraging the semantic knowledge to heuristically sample from the language space improves visual classification results due to the shared latent visual-language embedding space of VLMs. We show this with the proposed method, where the cosine similarity metric is replaced by a classifier, yielding improved results in image data sets and maps when using OpenSeg.

Notably, our method yields a significant improvement when utilizing OpenSeg as the encoder. This opens the possibility of using other encoders besides the widely adopted LSeg. While the embedding space of OpenSeg is more

sensitive to thresholding, the latent space seems to have a more expansive vocabulary, making it more suitable for open-vocabulary querying.

In our ablation experiments, we explore a range of commonly used heuristics for querying embeddings. We verify that prompt engineering phrases have to be designed for a particular encoder. We examine the various negative examples used in the literature and observe different behaviors between the tested encoders. This leads to the conclusion that these heuristics do not readily generalize beyond their designed environments.

These results are applicable in a wide range of robotic applications, as our methods are not dependent on a particular encoder or even a modality of the visual space, whether maps or images. Many of the recent methods proposed for robotic object retrieval today threshold the cosine similarity metric or contrast the query with a single negative example. Future work could be devoted to extending these methods by replacing thresholding with a classifier, possibly yielding significantly improved results, while retaining their other distinctive benefits.

However, several open questions remain, particularly in the open-vocabulary setting, where it remains open how to find the best possible query set. In this work, we evaluated the queryability of generic classes, but how can the queryability be evaluated in an open-vocabulary setting? Furthermore, as the majority of the novel computer vision methods have originally been proposed to work in the image domain, how to bridge the modality gap between images and maps? These questions need answers to bring robotic mapping into the era of open vocabulary semantics.

## REFERENCES

- [1] A. Vaswani *et al.*, “Attention Is All You Need,” in *Advances in Neural Inf. Process. Syst. 30 (NeurIPS)*, Long Beach, CA, USA, Dec. 2017, pp. 6000–6010.
- [2] A. Radford *et al.*, “Learning Transferable Visual Models From Natural Language Supervision,” in *Proc. Int. Conf. on Mach. Learn. (ICML)*, Virtual, Jul. 2021, pp. 8748–8763.
- [3] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik, “LERF: Language Embedded Radiance Fields,” in *Proc. IEEE/CVF Int. Conf. on Comput. Vis. (ICCV)*, Paris, France, Oct. 2023, pp. 19 672–19 682.
- [4] C. Huang, O. Mees, A. Zeng, and W. Burgard, “Visual Language Maps for Robot Navigation,” in *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, London, United Kingdom, 5 2023, pp. 10 608–10 615.
- [5] J. Yu *et al.*, “Language-Embedded Gaussian Splats (LEGS): Incrementally Building Room-Scale Representations with a Mobile Robot,” in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS)*, Abu Dhabi, United Arab Emirates, Oct. 2024, pp. 13 326–13 332.
- [7] Q. Gu *et al.*, “ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning,” in *Proc. IEEE Int. Conf. on Robot. and Automat. (ICRA)*, Yokohama, Japan, 5 2024, pp. 5021–5028.
- [8] M. Pekkanen, F. Verdoja, T. Mihaylova, and V. Kyrki, “Do Visual-Language Grid Maps Capture Latent Semantics?” in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS)*, Hangzhou, China, Oct. 2025.
- [9] R.-Z. Qiu *et al.*, “Learning Generalizable Feature Fields for Mobile Manipulation,” 2024, arXiv:2403.07563.
- [6] K. M. Jatavallabhula *et al.*, “ConceptFusion: Open-set Multimodal 3D Mapping,” in *Proc. Robotics: Sci. and Syst. (RSS)*, Daegu, Republic of Korea, Jul. 2023.
- [10] N. Patel, P. Krishnamurthy, and F. Khorrami, “RAZER: Robust Accelerated Zero-Shot 3D Open-Vocabulary Panoptic Reconstruction with Spatio-Temporal Aggregation,” 2025, arXiv:2505.15373.
- [11] W. Liang, Y. Zhang, Y. Kwon, S. Yeung, and J. Zou, “Mind the Gap: Understanding the Modality Gap in Multi-modal Contrastive Representation Learning,” 2022, arXiv:2203.02053.
- [12] M. Wei, T. Wang, Y. Chen, H. Wang, J. Pang, and X. Liu, “OVExp: Open Vocabulary Exploration for Object-Oriented Navigation,” 2024, arXiv:2407.09016.
- [13] Z. Liu *et al.*, “A Neural Representation Framework with LLM-Driven Spatial Reasoning for Open-Vocabulary 3D Visual Grounding,” 2025, arXiv:2507.06719.
- [14] Q. Li, J. Sun, L. An, Z. Su, H. Zhang, and Y. Liu, “SemanticSplat: Feed-Forward 3D Scene Understanding with Language-Aware Gaussian Fields,” 2025, arXiv:2506.09565.
- [15] Z. Wang *et al.*, “CRIS: CLIP-Driven Referring Image Segmentation,” in *Proc. IEEE/CVF Conf. on Comput. Vis. and Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 11 676–11 685.
- [16] T. Luddecke and A. Ecker, “Image Segmentation Using Text and Image Prompts,” in *Proc. IEEE/CVF Conf. on Comput. Vis. and Pattern Recognit. (CVPR)*, New Orleans, LA, USA, Jun. 2022, pp. 7076–7086.
- [17] W. Zhang, T. Li, and Y. Liu, “MAG-Nav: Language-Driven Object Navigation Leveraging Memory-Reserved Active Grounding,” 2025, arXiv:2508.05021.
- [18] S. Linok and G. Naumov, “Open-Vocabulary Indoor Object Grounding with 3D Hierarchical Scene Graph,” 2025, arXiv:2507.12123.
- [19] Z. Zhong, Y. He, P. Li, F. Yu, and F. Ma, “A Language-Driven Navigation Strategy Integrating Semantic Maps and Large Language Models,” in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS)*, Abu Dhabi, United Arab Emirates, Oct. 2024, pp. 9753–9760.
- [20] N. H. Chapman, F. Dayoub, W. Browne, and C. Lehnert, “QueryAdapter: Rapid Adaptation of Vision-Language Models in Response to Natural Language Queries,” 2025, arXiv:2502.18735.
- [21] Y. Yuan and A. Nuechter, “Uni-Fusion: Universal Continuous Mapping,” *IEEE Trans. on Robot. (T-RO)*, vol. 40, pp. 1373–1392, Jan. 2024.
- [22] Z. Yan *et al.*, “Dynamic Open-Vocabulary 3D Scene Graphs for Long-term Language-Guided Mobile Manipulation,” *IEEE Robot. and Automat. Lett.*, vol. 10, no. 5, pp. 4252–4259, 5 2025.
- [23] S. Sharma *et al.*, “Semantic Mechanical Search with Large Vision and Language Models,” in *Proc. Conf. on Robot Learn. (CoRL)*, Atlanta, GA, USA, Nov. 2023, pp. 971–1005.
- [24] S. Peng, K. Genova, C. M. Jiang, A. Tagliasacchi, M. Pollefeys, and T. Funkhouser, “OpenScene: 3D Scene Understanding with Open Vocabularies,” in *Proc. IEEE/CVF Conf. on Comput. Vis. and Pattern Recognit. (CVPR)*, Vancouver, BC, Canada, Jun. 2023, pp. 815–824.
- [25] J. Kim, Y. Park, H. J. Yoon, and B.-T. Zhang, “OV-MAP : Open-Vocabulary Zero-Shot 3D Instance Segmentation Map for Robots,” in *Proc. IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS)*, Abu Dhabi, United Arab Emirates, Oct. 2024, pp. 13 780–13 786.
- [26] A. Chang *et al.*, “Matterport3D: Learning From RGB-D Data in Indoor Environments,” in *Proc. Int. Conf. on 3D Vis. (3DV)*, Qingdao, China, Oct. 2017, pp. 667–676.
- [27] C. Huang, O. Mees, A. Zeng, and W. Burgard. VLMaps. [Online]. Available: <https://github.com/vlmaps/vlmaps>
- [28] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, “Language-driven Semantic Segmentation,” in *Proc. Int. Conf. on Learn. Representations (ICLR)*, Virtual, Apr. 2022.
- [29] G. Ghiasi, X. Gu, Y. Cui, and T.-Y. Lin, “Scaling Open-Vocabulary Image Segmentation with Image-Level Labels,” in *Proc. Eur. Conf. on Comput. Vis. (ECCV)*, Tel Aviv, Israel, Jul. 2022, pp. 540–557.
- [30] D. Miller, N. Sünderhauf, A. Kenna, and K. Mason, “Open-Set Recognition in the Age of Vision-Language Models,” in *Proc. Eur. Conf. on Comput. Vis. (ECCV)*, Milano, Italy, Sep. 2024, pp. 1–18.