

# LAMP: Implicit Language Map for Robot Navigation

Sibaek Lee , Hyeonwoo Yu , *Member, IEEE*, Giseop Kim , *Member, IEEE*, and Sunwook Choi , *Member, IEEE*

**Abstract**—Recent advances in vision-language models have made zero-shot navigation feasible, enabling robots to interpret and follow natural language instructions without requiring labeling. However, existing methods that explicitly store language vectors in grid or node-based maps struggle to scale to large environments due to excessive memory requirements and limited resolution for fine-grained planning. We introduce LAMP (Language Map), a novel neural language field-based navigation framework that learns a continuous, language-driven map and directly leverages it for fine-grained path generation. Unlike prior approaches, our method encodes language features as an implicit neural field rather than storing them explicitly at every location. By combining this implicit representation with a sparse graph, LAMP supports efficient coarse path planning and then performs gradient-based optimization in the learned field to refine poses near the goal. Our two-stage pipeline of coarse graph search followed by language-driven, gradient-guided optimization is the first application of an implicit language map for precise path generation. This refinement is particularly effective at selecting goal regions not directly observed by leveraging semantic similarities in the learned feature space. To further enhance robustness, we adopt a Bayesian framework that models embedding uncertainty via the von Mises–Fisher distribution, thereby improving generalization to unobserved regions. To scale to large environments, LAMP employs a graph sampling strategy that prioritizes spatial coverage and embedding confidence, retaining only the most informative nodes and substantially reducing computational overhead. Our experimental results, both in NVIDIA Isaac Sim and on a real multi-floor building, demonstrate that LAMP outperforms existing explicit methods in both memory efficiency and fine-grained goal-reaching accuracy, opening new possibilities for scalable, language-driven robot navigation.

**Index Terms**—Vision-based navigation, mapping, path planning, open-vocabulary scene understanding.

## I. INTRODUCTION

**I**N ROBOTICS, map representation has significant value as it enables robots to perceive spatial information and effectively

Received 29 May 2025; accepted 21 September 2025. Date of publication 9 October 2025; date of current version 24 October 2025. This article was recommended for publication by Associate Editor Z. Wu and Editor P. Vasseur upon evaluation of the reviewers' comments. This work was supported by IITP grant funded by the Korea government (MSIT) under Grant RS-2023-00227993 and in part by Sibaek Lee's internship at NAVER LABS. (*Corresponding author: Sunwook Choi.*)

Sibaek Lee is with the Department of Intelligent Robotics, Sungkyunkwan University, Suwon 03063, South Korea, and also with NAVER LABS, Seongnam 13561, South Korea (e-mail: lmjls@skku.edu).

Hyeonwoo Yu is with the Department of Intelligent Robotics, Sungkyunkwan University, Suwon 03063, South Korea (e-mail: hwyu@skku.edu).

Giseop Kim is with the Department of Robotics and Mechatronics Engineering, Daegu Gyeongbuk Institute of Science and Technology, Daegu 42988, South Korea (e-mail: gsk@dgist.ac.kr).

Sunwook Choi is with NAVER LABS, Seongnam 13561, South Korea (e-mail: sunwook.choi@naverlabs.com).

The project page, with additional resources, is available at: <https://lab-of-ai-and-robotics.github.io/LAMP>.

Digital Object Identifier 10.1109/LRA.2025.3619820

perform a wide range of tasks. Traditionally, 3D maps have been constructed by assigning coordinate values, RGB information, and semantic labels to point clouds or voxels [1], [2], [3], [4]. However, with the emergence of visual language models [5], [6], it has become possible to represent space through language embedding vectors, thereby capturing long-tailed objects and abstract concepts without explicit labeling. This shift transforms the environment from a purely geometric construct into a linguistically enriched space, opening the door to more intuitive and meaningful robot-environment interactions. Leveraging the strengths of these language embeddings, they are being applied in various fields such as 3D semantic segmentation [7], [8], [9], path planning [10], [11], scene understanding [12], [13], manipulation [14], extended reality [15], and autonomous driving [16].

However, current language map representations are limited to small environments and encounter significant challenges for large-scale deployment. Since robots must operate indoors and outdoors to perform tasks autonomously, they require effective space interaction in extensive areas. Existing language maps for path planning methods typically store language vectors explicitly in 2D grids [10], [17], [18] or topological graph nodes [11], [19], [20]. As environments scale up, this explicit storage becomes excessively memory and computationally intensive. Some approaches adopt sparse representations to alleviate these constraints, enabling coarse-level path planning but limiting fine-grained navigation. This limitation arises from the inherent difficulty of densely and explicitly storing information on large scales. To address this gap, we propose an implicit language map representation that continuously models language vectors from RGB-only input, facilitating memory-efficient path planning that supports not only coarse navigation but also fine-grained route guidance.

In Fig. 1, (a) the grid-based method leverages depth images and a Visual Language Model [5], [21] to integrate language information into a 3D representation, which is subsequently projected onto a 2D grid map using a top-down approach [10], [17], [18]. This method captures objects well when grid cells are smaller than the objects. However, it struggles to represent smaller objects because of the grid's coarse resolution. Moreover, a dense grid in large environments consumes too much memory, making it unsuitable for robotics. In contrast, (b) the node-based method places nodes at key spatial positions and explicitly stores CLIP vectors and pose information [11], [19], [20]. It uses less memory than grid-based approach and can detect smaller objects at the nodes. However, in large scenes, it is impractical to place enough nodes to capture the fine-grained details required for precise path planning. For instance, Fig. 1 shows that with sparse nodes, the node-based method fails

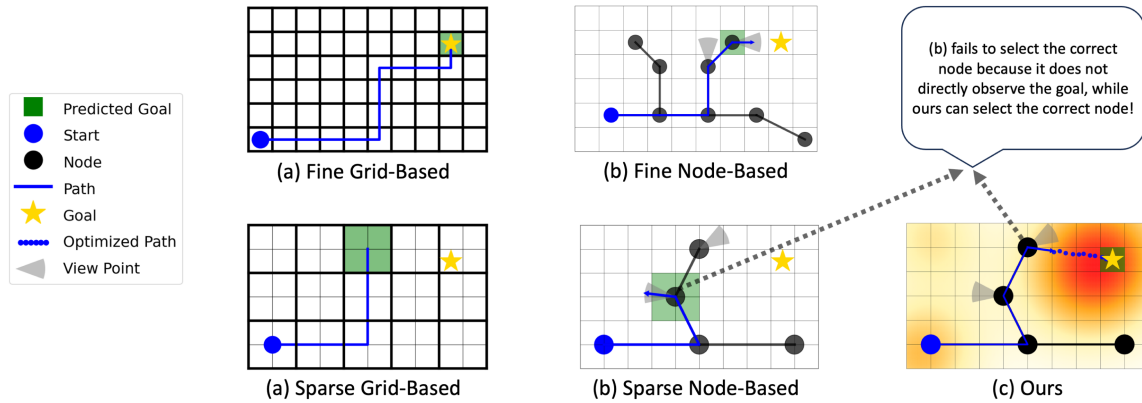


Fig. 1. Comparison of three language map representation methods. (a) The grid-based approach struggles to accurately represent objects at coarse resolutions and requires excessive memory when increasing grid resolution to capture finer details. (b) The node-based approach fails to capture important object details when node spacing is too coarse and cannot guarantee precise path planning. (c) In contrast, our implicit method maintains memory efficiency even at large scales while providing fine-level path guidance.

to locate the goal. Conversely, with dense nodes, the method reaches near the goal but struggles to approach the goal with fine precision. Finally, (c) our implicit method addresses these issues by modeling language information as a continuous function using a neural network. This network infers language vectors in regions not explicitly observed during training, enabling flexible and detailed representations that continuously capture smaller objects. Consequently, the system performs precise path planning without storing vectors at every node, significantly reducing memory requirements. Furthermore, even if a node doesn't directly observe the goal, our method selects nodes near the goal, ensuring guidance to unobserved targets. By maintaining only a sparse set of nodes with pose information and generating language features on demand, our approach achieves memory efficiency and detailed mapping in large environments.

Building on the strengths of our implicit language map, we propose methods to construct and utilize this representation more effectively. Although the implicit map provides a continuous function for language vectors in unobserved areas, mapping camera poses to language vectors in a highly nonlinear manner poses a challenge. Moreover, CLIP vectors often capture multiple objects simultaneously, injecting noise into the language embeddings themselves. To address these issues, we apply a Bayesian approach to our neural network to account for aleatoric uncertainty [22], [23]. By incorporating this uncertainty into our loss function, we enhance the implicit map's generalization to untrained poses. We also introduce a graph sampling strategy for large-scale graphs that utilizes both the uncertainty information derived from our Bayesian model and the network gradients. As the number of nodes increases in large-scale environments, exhaustive exploration becomes impractical, and node selection critically affects task performance [24], [25]. Consequently, graph sampling is essential for efficient large-scale planning and ensures our implicit map remains compact yet effective. By modeling embedding uncertainty via a von Mises–Fisher Bayesian formulation and incorporating semantic sensitivity into node sampling, our system achieves far greater robustness than naive gradient-based optimization alone. We summarize our main contributions of LAMP (Language Map) as follows:

- We introduce LAMP, the first implicit language map leveraging a language-driven continuous field for fine-grained path generation using only RGB images.
- We incorporate a Bayesian approach to improve generalization performance to untrained poses.
- We develop a graph sampling technique that optimizes node selection using language features and uncertainty predictions.

## II. RELATED WORK

**Traditional Mapping Approaches.** Robots have long relied on geometric maps such as occupancy grids [26] and landmark-based [27], [28], [29] representations for navigation. While effective for path planning, these low-level depictions carry no inherent semantic information and the sparse features in SLAM maps are difficult to interpret in human terms without additional context [27], [28], [29], [30]. To make maps more human-meaningful, researchers introduced label-based semantic mapping by annotating regions with object and place names [31], [32], but this approach is confined to pre-defined classes as it depends on explicit annotations or trained detectors. Consequently, conventional maps cannot handle references to novel, long-tail concepts and thus fail to generalize to many unseen objects or descriptors.

**Language-Based Mapping Approaches.** To overcome these limitations, recent work fuses natural language understanding with mapping. Vision–Language Model (VLM) [5], [6] has given rise to approaches where robots use natural language instructions to reach semantically specified targets. For instance, NeRF [7], [8] and Gaussian Splatting [9], [33] integrate language embeddings into scene representations for view synthesis and semantic segmentation, but these approaches are not primarily designed for navigation and cannot scale to large environments due to high memory and rendering costs. Other navigation-oriented methods attempt to address this by explicitly storing language vectors in 2D grids [10], [17], [18] or in topological graph nodes [11], [19], [20], but they still suffer from excessive memory demands and coarse planning resolution in large scale

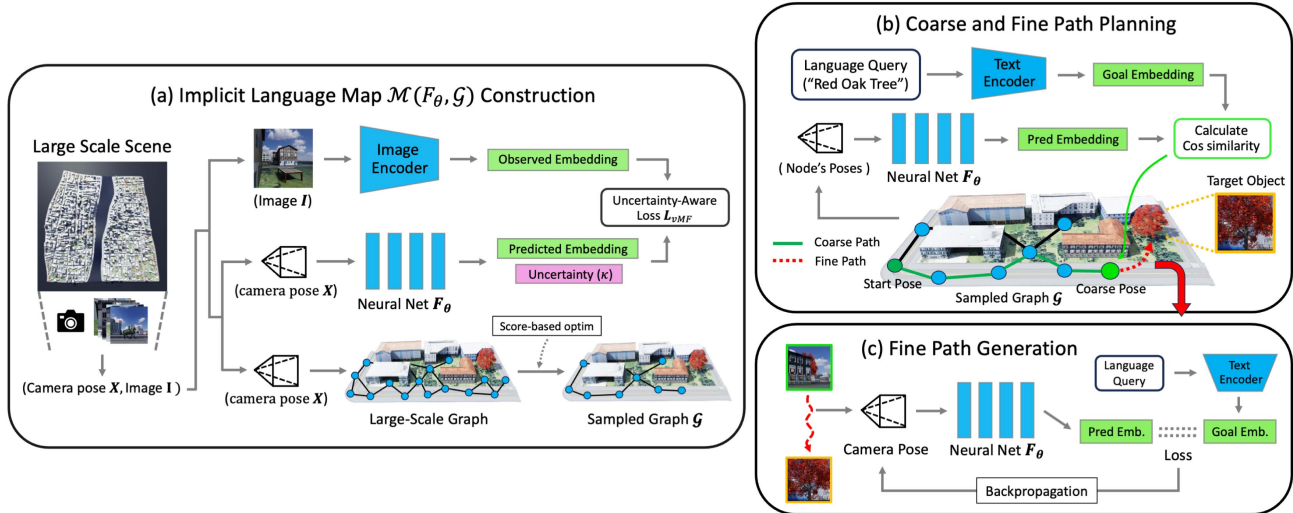


Fig. 2. System Overview. (a) Implicit Language Map Construction: The robot traverses the environment and collects pairs of camera poses  $\mathbf{x}$  and corresponding images  $\mathbf{I}$ . Neural network  $F_{\Theta}$  maps each pose  $\mathbf{x}$  to a language embedding  $\mathbf{z} = F_{\Theta}(\mathbf{x})$ . Since processing the full large-scale topological graph is computationally expensive, we sample the graph  $\mathcal{G}$  using our proposed score-based optimization for coarse planning. (b) Coarse Path Planning: Given a user's natural language query such as "red oak tree", we encode a goal embedding and apply A\* on the sampled graph  $\mathcal{G}$  to obtain a coarse path to the node whose embedding best matches the goal embedding. (c) Fine Path Generation: We then generate the pose using  $F_{\Theta}$  to maximize cosine similarity, moving from the coarse pose to a fine pose that offers a clear view of the target object.

environments. To address this limitation, we introduce the first approach that implicitly represents space for memory efficiency and leverages this implicit map in an optimization framework to generate fine-grained navigation paths.

### III. METHOD

We introduce a map representation that continuously encodes language features within a large-scale space, ensuring memory efficiency and enabling fine-grained path planning. The section is organized into the following parts: Section III-A states the navigation problem. Section III-B describes how we construct the implicit language map, incorporate uncertainty, and prune the topological graph. Section III-C explains how the resulting map is used to generate coarse paths and refine them into precise goal poses.

#### A. Problem Definition

We formalize the navigation task as follows. Let  $\mathcal{M} = \{F_{\Theta}, \mathcal{G}\}$  denote an implicit language map of the environment, where  $F_{\Theta}$  is a neural network that continuously represents language features and  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  is a topological graph providing a global structure for path planning. Navigation begins from an initial pose  $\mathbf{x}_{init}$  within this static environment (i.e.,  $\mathcal{M}$  does not change) and receives a natural language query  $Q$  describing a desired destination. The objective is to generate a twostage path: first, a coarse path  $\gamma_c = (x_0, \dots, x_c)$  obtained by searching over the graph's node set  $\mathcal{V}$  to reach an intermediate goal  $x_c$ ; second, a fine path  $\gamma_f = (x_c, \dots, x_G)$  produced by optimizing  $F_{\Theta}$  to precisely reach the final goal  $x_G$ . In summary, given  $(\mathcal{M}, \mathbf{x}_{init}, Q)$ , the system generates a coarse path  $\gamma_c$  via  $\mathcal{G}$  and then refines it into a fine path  $\gamma_f$  via  $F_{\Theta}$  to obtain an accurate path to the final goal  $\mathbf{x}_G$ .

#### B. Implicit Language Map Pipeline

1) *Implicit Language Map Construction*: To construct  $\mathcal{M}$ , we assume that the robot has traversed the environment extensively, collecting  $(\mathbf{x}, \mathbf{I})$  pairs, where  $\mathbf{x}$  represents the camera pose and  $\mathbf{I}$  is the corresponding image. Inspired by NeRF frameworks [34], we implicitly encode the language embeddings of the scene as viewed from each position, rather than encoding the color and density at the positions themselves. Concretely, let  $\mathbf{t} \in \mathbb{R}^3$  be the 3D coordinate and  $\mathbf{q} \in \mathbb{R}^4$  be a quaternion representing orientation. We concatenate these into  $\mathbf{x} = [\mathbf{t}, \mathbf{q}] \in \mathbb{R}^7$ . Our neural network  $F_{\Theta}$  then maps  $\mathbf{x}$  to a  $d$ -dimensional CLIP embedding:

$$F_{\Theta}(\mathbf{x}) = \mathbf{z} \in \mathbb{R}^d,$$

where  $\mathbf{z}$  captures the language features observed in the input image  $\mathbf{I}$ . By training on  $(\mathbf{x}, \mathbf{I})$  pairs,  $F_{\Theta}$  learns a continuous language representation of the environment using only RGB images.

However, relying solely on  $F_{\Theta}$  for path planning can lead to local minima issues. To address this, we use the topological graph  $\mathcal{G}$ , whose nodes  $v \in \mathcal{V}$  store only pose information  $\mathbf{x}_v$ , from which language embeddings can be generated on demand via  $F_{\Theta}(\mathbf{x}_v)$ . This approach contrasts with previous methods that explicitly store precomputed embeddings at every node, leading to large memory usage. By dynamically generating embeddings through  $F_{\Theta}$ , our method significantly reduces storage while preserving language features.

2) *Bayesian Approach for Robust Embedding Predictions*: Although large-scale data collection allows the robot to learn from a wide set of  $(\mathbf{x}, \mathbf{I})$  pairs, predicting embedding vectors for unseen poses remains challenging. Moreover, CLIP-based embeddings represent the collective semantics of multiple objects within an image, which can introduce noise when focusing on a

specific object. Consequently, accounting for uncertainty is necessary for robust predictions [22], [35]. Since each embedding vector lies on the unit hypersphere in  $\mathbb{R}^d$ , the von Mises-Fisher (vMF) distribution naturally captures its directional uncertainty while preserving the unit norm. We enforce this unit-length condition by  $\ell_2$ -normalising every CLIP feature and network output, so cosine similarity is a true metric and the embeddings reside on the domain of the vMF distribution. To address this issue, we adopt a Bayesian approach based on the vMF distribution [36], [37], wherein each predicted embedding  $\mathbf{z}$  lies on the unit hypersphere in  $\mathbb{R}^d$  with density:

$$p(\mathbf{z} \mid \boldsymbol{\mu}, \kappa) = C_d(\kappa) \exp(\kappa \boldsymbol{\mu}^\top \mathbf{z}), \quad (1)$$

where  $\boldsymbol{\mu} \in \mathbb{R}^d$  is also a unit vector,  $\kappa \geq 0$  is the concentration parameter, and  $C_d(\kappa)$  is the normalization constant.<sup>1</sup> Intuitively,  $\kappa$  governs how tightly  $\mathbf{z}$  concentrates around the mean direction  $\boldsymbol{\mu}$ . To prevent  $\kappa$  from diverging, we place a Gamma prior on it:

$$p(\kappa) = \frac{\beta^\alpha}{\Gamma(\alpha)} \kappa^{\alpha-1} e^{-\beta\kappa}, \quad (2)$$

where  $\alpha$  and  $\beta$  are hyperparameters, and  $\Gamma(\alpha)$  is the gamma function. This prior enforces reasonable bounds on the concentration parameter, mitigating overconfidence in the embedding's direction. To implement this, we define our network  $F_\Theta$  to output the two parameters of the vMF distribution for a given pose  $\mathbf{x}$ : the mean direction  $\boldsymbol{\mu}$  and the concentration  $\kappa$ . Let us denote the network outputs as  $(\boldsymbol{\mu}_\Theta(\mathbf{x}), \kappa_\Theta(\mathbf{x})) = F_\Theta(\mathbf{x})$ . By combining the vMF likelihood in (1) with the Gamma prior in (2), the posterior over the network parameters  $\theta$  is proportional to:

$$p(\theta \mid \mathbf{x}, \mathbf{z}_{\text{obs}}) \propto p(\mathbf{z}_{\text{obs}} \mid F_\Theta(\mathbf{x})) p(\kappa_\Theta(\mathbf{x})), \quad (3)$$

and we train the network by minimizing the negative log-posterior, which serves as our loss function for a given training pair  $(\mathbf{x}, \mathbf{z}_{\text{obs}})$ :

$$\mathcal{L}_{\text{vMF}}(\theta; \mathbf{x}, \mathbf{z}_{\text{obs}}) = -\ln p(\mathbf{z}_{\text{obs}} \mid F_\Theta(\mathbf{x})) - \ln p(\kappa_\Theta(\mathbf{x})). \quad (4)$$

This loss function trains the network to produce parameters  $(\boldsymbol{\mu}_\Theta(\mathbf{x}), \kappa_\Theta(\mathbf{x}))$  that balance closeness to the observed embedding  $\mathbf{z}_{\text{obs}}$  with regularization from the Gamma prior, thereby reducing the influence of noisy data and enhancing robustness in unobserved scenarios. By introducing a Bayesian framework on the outputs of  $F_\Theta$ , our method becomes robust to noisy embeddings and provides reliable predictions in unobserved regions. This enhances path-planning performance and also feeds into the subsequent graph-optimization step.

3) *Node Optimization in Large-Scale Graph*: While our topological graph  $\mathcal{G}$  encodes the global structure for path planning, a large number of nodes can lead to computational inefficiencies during graph search [24], [25]. To address this, we propose a graph sampling method that retains only the most informative nodes, scored by three criteria. Formally, for each node  $v \in \mathcal{V}$ , we define:

$$\text{score}(v) = w_1 s_{\text{vc}}(v) + w_2 s_u(v) + w_3 s_{\text{ss}}(v), \quad (5)$$

<sup>1</sup>  $C_d(\kappa) = \frac{\kappa^{d/2-1}}{(2\pi)^{d/2} I_{d/2-1}(\kappa)}$ , where  $I_\nu$  is the modified Bessel function of the first kind.

where  $s_{\text{vc}}(v)$  is a *View Coverage Score*,  $s_u(v)$  is an *Uncertainty Score*, and  $s_{\text{ss}}(v)$  is a *Semantic Sensitivity Score*; the weights  $w_1$ ,  $w_2$ , and  $w_3$  are hyperparameters that balance the contribution of each score to the overall score.

- *View Coverage Score*  $s_{\text{vc}}(v)$ : For each node  $v$ , we compute the dot product between its quaternion and those of nodes in its twohop neighborhood to measure orientation difference. A lower similarity indicates a novel viewpoint, which enhances coverage in the sampled graph.
- *Uncertainty Score*  $s_u(v)$ : Leverages our Bayesian approach to quantify aleatoric uncertainty. The score is defined as the concentration parameter,  $\kappa$ , of the predicted von Mises-Fisher distribution, where a higher  $\kappa$  indicates greater confidence in the embedding.
- *Semantic Sensitivity Score*  $s_{\text{ss}}(v)$ : Measures the gradient of  $F_\Theta$  with respect to the node's pose  $\mathbf{x}_v$  by computing its norm. A larger gradient suggests that small changes in pose result in substantial variation in the predicted embedding, implying a high semantic contrast. These nodes are valuable for capturing critical transition regions.

Nodes are ranked by their combined score (5), and only the top-scoring nodes are retained in the sampled graph. This selective pruning reduces the computational cost of graph-based planning while preserving coverage of key viewpoints, embedding reliability, and semantic diversity. As a result, our method scales efficiently to complex environments without compromising navigation and exploration capabilities.

### C. Coarse to Fine Path Planning

1) *Coarse Path Planning*: Given a natural language query, we first convert it into a goal embedding  $\mathbf{z}_{\text{goal}} = \text{CLIP}(Q)$ . For each node  $v \in \mathcal{V}$  with a stored pose  $\mathbf{x}_v$ , we obtain an embedding  $F_\Theta(\mathbf{x}_v)$  and compute its cosine similarity to  $\mathbf{z}_{\text{goal}}$ :

$$\text{sim}(\mathbf{z}_{\text{goal}}, F_\Theta(\mathbf{x}_v)), \quad \forall v \in \mathcal{V}.$$

The node with the highest similarity serves as our coarse semantic goal, and its pose is set to  $\mathbf{x}_c$  for the fine-level optimization. We then employ an A\* algorithm to plan a path from the robot's current node to this coarse goal node, achieving coarse-level navigation to the language goal.

2) *Fine Path Generation*: Although  $\mathcal{G}$  provides a coarse path from the most similar node in large-scale environments, additional refinement is necessary to achieve fine-grained goals beyond the graph's node resolution. Therefore, we initiate optimization from the coarse pose  $\mathbf{x}_c$ . To reduce the risk of converging to a poor local optimum, we sample multiple candidate poses near  $\mathbf{x}_c$  and select the pose with the highest similarity to  $\mathbf{z}_{\text{goal}}$ , denoted as  $\mathbf{x}_{\text{best}}$ , which guides the subsequent optimization. We then define a small correction term  $\delta\mathbf{x}$  on top of  $\mathbf{x}_c$  and formulate the following objective:

$$\max_{\delta\mathbf{x}} [\text{sim}(F_\Theta(\mathbf{x}_c + \delta\mathbf{x}), \mathbf{z}_{\text{goal}}) - \lambda_{\text{dist}} \text{dist}(\mathbf{x}_{\text{best}}, \mathbf{x}_c + \delta\mathbf{x})]. \quad (6)$$




Fig. 3. Examples of objects used in our simulation navigation experiments. The top row displays large objects (volume  $\geq 1\text{m}^3$ ) such as statues and a red oak tree, while the bottom row shows smaller objects (volume  $< 1\text{m}^3$ ) such as a Rubik’s cube or a fire alarm, which are harder to detect in a large-scale environment.

Here  $\text{dist}(\mathbf{x}_1, \mathbf{x}_2)$  denotes the Euclidean distance in  $\mathbb{R}^3$ . We solve (6) via a gradient-based method, updating  $\delta\mathbf{x}$  at each step  $t$  by

$$\delta\mathbf{x}_{t+1} = \delta\mathbf{x}_t - \alpha \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}}, \quad (7)$$

where  $\hat{m}_t$  and  $\hat{v}_t$  are Adam’s bias-corrected first and second moment estimates,  $\alpha$  is the learning rate, and  $\epsilon$  is a small constant. This process yields a refined pose  $\mathbf{x}_c + \delta\mathbf{x}^*$  that simultaneously maximizes semantic similarity and remains close to the best sampled pose  $\mathbf{x}_{\text{best}}$ , thus mitigating local minima. By combining coarse-level graph navigation with this guided optimization in  $F_{\Theta}$ , we achieve fine-grained language guidance without sacrificing large-scale planning and memory efficiency.

#### IV. EXPERIMENTS

The purpose of our experiments is to demonstrate that LAMP, our method which implicitly incorporates language information within large-scale scenes, achieves memory efficiency and enables fine-grained zero-shot navigation. In the following subsections, Section IV-A describes the dataset configuration and implementation details, Section IV-B presents the experimental results obtained in the Nvidia Isaac simulation environment along with a discussion, and Section IV-C validates our approach on a real robot by demonstrating navigation across multiple floors of a building using only a very sparse set of nodes.

##### A. Experimental Setup

1) *Dataset Configuration:* To validate our method’s memory efficiency, we evaluate our system on a large-scale environment rather than limiting it to room-scale datasets such as MP3D [38]. Therefore, for our simulation experiments, we use the NVIDIA Isaac Sim simulator along with the City Tower Demo 3D Models Pack, which covers about  $1.6\text{ km} \times 1.8\text{ km}$ . We integrate assets from the NVIDIA library to set up a zero-shot navigation task. As shown in Fig. 3, these assets are divided into two groups according to size: objects with a volume of  $1\text{ m}^3$  or greater are labeled as easy cases, while those with a smaller volume are hard cases. The easy cases include large items like statues and trees that are visible from many viewpoints, making them easier to navigate to. In contrast, hard cases include smaller items, such as a Rubik’s cube or a fire alarm, whose compact size makes

them difficult to detect in a large-scale environment. We assume that the robot explores the environment for an extended period to collect data. Since the data are obtained directly through the robot’s own exploration, the gathered pose data ensure the robot’s traversability. We collect five million pose–image pairs to provide diverse viewpoints for a refined map. In our real-world experiments, we deployed a robot equipped with six cameras to navigate each floor of the building. On each floor, the robot collected between 10,000 and 20,000 pose–image pairs, ensuring comprehensive coverage of the environment. The building consists of 28 floors, each spanning approximately  $100\text{ m} \times 70\text{ m}$ .

2) *Implementation Details:* In implementing our approach, we follow the MLP architecture used in NeRF [34] by incorporating positional encoding and skip connections to capture necessary variations. Our model outputs a 512-dimensional language embedding vector along with a concentration parameter to indicate uncertainty. During training, we model uncertainty by placing a Gamma prior on the concentration parameter with  $\alpha = 2$  and  $\beta = 0.5$ . As described in Section III-B3, key nodes within the graph are selected using a heuristic weighting scheme with  $w_1 = 1$ ,  $w_2 = 1$ , and  $w_3 = 0.5$ . For the fine path generation objective (6), the distance penalty hyperparameter  $\lambda_{\text{dist}}$  is set to 5. We assess navigation performance using three metrics. First, *success rate* is computed considering only the top 1% of the predictions; a trial is deemed successful if the robot ends up within 20 m of the center of an object. Next, the *Success weighted by Path Length (SPL)* metric evaluates navigation efficiency by penalizing unnecessarily long paths. Finally, for successful trials, we calculate the *goal-to-distance* measure by averaging the remaining distance from the target.

##### B. Simulation Experiments and Analysis

1) *Comparison of Language Map Representation Methods for Object Goal Tasks:* As shown in Table I, when using the same amount of memory for map representation, our method outperforms existing approaches in terms of success rate, success weighted by path length, and goal-to-distance metrics. As summarized in Table II, the grid-based method requires approximately 1,000 times more memory than our method to achieve comparable performance by setting the grid size to 40 cm. Even with this increased memory usage, the grid-based approach captures large objects but fails to detect smaller ones. In contrast, the node-based method needs about 70 times more memory than our method to reach a similar success rate, yet its performance in the goal-to-distance metric still falls short of that achieved by our implicit method. This performance stems from our two-stage pipeline. The initial coarse stage achieves a success rate of 42–67%, after which the fine-grained optimization stage consistently reduces the goal distance by about 50%. For easy cases, this optimization always moves closer to the target, and for hard cases, it succeeds in 90% of cases. Furthermore, while the on-demand embedding generation and fine-path optimization introduce a computational cost, our method completes a full navigation query in under one second, with inference times measured on an NVIDIA RTX 4090 GPU, as shown in Tables I and II.

TABLE I  
QUANTITATIVE COMPARISON OF LANGUAGE MAP REPRESENTATIONS (UNIFORM MEMORY SETTING)

Method	Memory(GB)	Time (s)	SR (Easy)	SPL (Easy)	GDist (Easy)	SR (Hard)	SPL (Hard)	GDist (Hard)
Grid-based (Sparse)	0.055	0.0001	0.08	0.07	15.50	0.0	0.0	–
Node-based (Sparse)	0.057	0.0066	0.41	0.36	12.69	0.21	0.17	4.14
Ours (Implicit)	0.057	0.8041	<b>0.67</b>	<b>0.62</b>	<b>6.36</b>	<b>0.42</b>	<b>0.38</b>	<b>1.74</b>

TABLE II  
QUANTITATIVE COMPARISON OF LANGUAGE MAP REPRESENTATIONS (MEMORY-INTENSIVE SETTING)

Method	Memory(GB)	Time (s)	SR (Easy)	SPL (Easy)	GDist (Easy)	SR (Hard)	SPL (Hard)	GDist (Hard)
Grid-based (Dense)	56.34	0.1036	0.83	0.80	2.44	0.36	0.33	1.27
Node-based (Dense)	3.962	0.3395	0.67	0.61	7.51	0.47	0.41	2.38
Ours (Implicit)	<b>0.057</b>	0.8041	0.67	0.62	6.36	0.42	0.38	1.74

This experiment uses the NVIDIA’s City Tower Demo 3D Models Pack scene, which covers  $1.6\text{ km} \times 1.8\text{ km}$ . In the Grid-based (Sparse) method, the grid cell is set to 12m, while the Grid-based (Dense) method utilizes a grid size of 0.4m. The Node-based (Sparse) map comprises 30,000 nodes, and the Node-based (Dense) map contains 2,000,000 nodes. Success Rate (SR) and Success weighted by Path Length (SPL) are reported as proportions. GDist denotes the remaining Goal-to-Distance in metres.

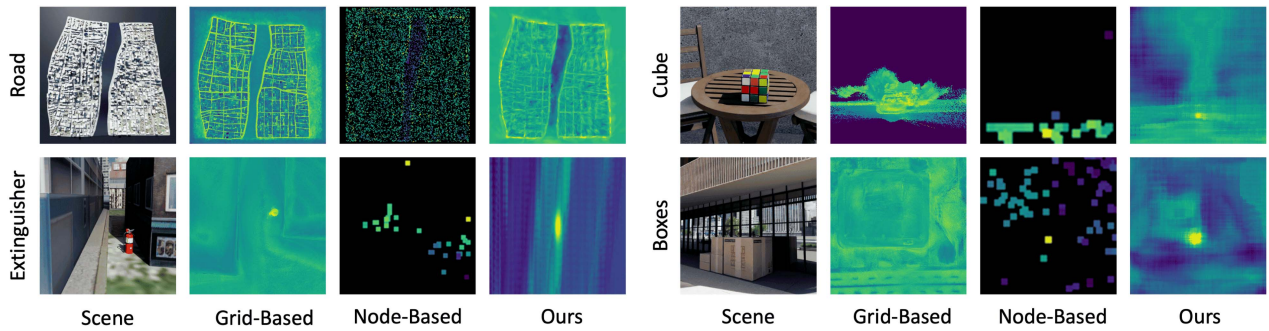


Fig. 4. Visualization of each language map representation in the near-goal region of NVIDIA’s City Tower Demo 3D Models Pack scene using the viridis colormap. The Node-based method utilizes a dense setting, while the Grid-based method employs a dense setting for Road and Cube scenes and an extremely dense setting (5 cm grid size) for Extinguisher and Boxes scenes.

In Fig. 4, we present the visualization results of four different methods across distinct scenes: Road, Cube, Extinguisher, and Boxes. For the Road and Cube scenes, the grid-based method employs a dense setting with a grid size of 40 cm. However, for a more precise analysis, we adopt a 5 cm grid size for the Extinguisher and Boxes scenes, in line with the room-scale settings described in [17]. Due to memory constraints in city-scale environments, the 5 cm grid size is applied only within specific spatial regions during these experiments.

In the Road scene, all methods effectively represent the spatial layout of the road, as its continuous structure is straightforward to capture. In the Cube scene, however, the grid-based method fails as its 40 cm grid resolution is too coarse to capture the smaller cube, causing its representation to be mixed with surrounding objects. In the Extinguisher scene, the node-based method fails because it does not directly observe the goal, whereas our method correctly identifies the target by leveraging the implicit network to infer its presence from the surrounding context. Finally, in the Boxes scene, the grid-based method is hindered by z-axis projection artifacts, while the node-based method detects the boxes but fails to plan a precise path. Our method addresses these issues, accurately detecting the objects and enabling fine-grained goal reaching through optimization.

TABLE III  
COMPARISON OF NODE SELECTION METHODS

Method	$S_{vc}$	$S_u$	$S_{ss}$	SR	SPL	GDist
RN [25]	–	–	–	0.48	0.44	4.34
RDN [25]	–	–	–	0.45	0.40	5.31
RW [39]	–	–	–	0.48	0.43	4.13
FF [40]	–	–	–	0.45	0.41	4.79
Ours	✓	×	×	0.48	0.43	4.28
Ours	✓	✓	×	<b>0.51</b>	<b>0.47</b>	4.36
Ours	✓	×	✓	0.48	0.44	4.66
Ours	✓	✓	✓	<b>0.51</b>	<b>0.47</b>	<b>4.05</b>

$S_{vc}$ ,  $S_u$ , and  $S_{ss}$  refer to the View Coverage Score, Uncertainty Score, and Semantic Sensitivity Score, respectively, as detailed in Section III-B.

Overall, these results demonstrate that our method consistently maintains high performance across diverse scenarios by efficiently managing memory and effectively leveraging implicit information, thereby outperforming both grid-based and node-based approaches in challenging environments.

2) *Performance Comparison of Node Selection Methods:* Selecting appropriate nodes is essential in large-scale graphs to reduce computational complexity while preserving structural properties [24], [25]. In Table III, we compare our method



Fig. 5. Visualization of real-world experiments. ●: Start pose, representing the initial position from which navigation is initiated. ●: Coarse goal pose, selected based on its language embedding similarity to the target object. ●: Optimized goal pose, which is obtained through optimization of the implicit language map, ensuring that the target object is accurately brought into view.

against traditional sampling strategies: Random Node (RN) and Random Degree Node (RDN) [25], which select nodes uniformly at random and based on degree, respectively; Random Walk (RW) [39], which randomly traverses neighboring nodes; and Forest Fire (FF) [40], which simulates a spreading process to capture connectivity patterns. In addition, this experiment serves as an ablation study to validate the impact of each component in our proposed scoring mechanism. Unlike conventional methods that rely solely on edge information, our approach leverages implicit information from each node in large graphs, specifically using a neural network to derive an Uncertainty Score ( $S_u$ ) and language embeddings to guide the sampling process. The comparison between rows with and without the  $S_u$  term reveals that explicitly modeling uncertainty via our Bayesian framework provides a notable boost in performance, particularly in reducing the final goal to distance (GDist). By integrating the View Coverage Score ( $S_{vc}$ ), Uncertainty Score ( $S_u$ ), and Semantic Sensitivity Score ( $S_{ss}$ ), our method improves SR and SPL, demonstrating the effectiveness of incorporating learned features at the node level to enhance graph sampling beyond classical strategies.

### C. Real-World Experiments and Analysis

We conduct real-world experiments using the M2 robot [41] for zero-shot object goal navigation. Our real-world experiments evaluate our method’s precision using a sparse graph and its ability to navigate towards goals not directly observed by any node. To do this, we set approximately 20–30 nodes sparsely on each floor. During navigation, the robot first moves toward a target node based on the given natural language command and then refines its position through optimization. Specifically, when navigating to “drinks”, no node in the map directly observes the drinks; however, as shown in Fig. 5, a node near the drinks is selected, and through subsequent optimization, the robot adjusts its position to view the drinks. The implicit map learns from a large number of pose-image pairs and continuously represents coordinates in the language feature space, allowing nodes that have not directly observed the goal to possess language vectors similar to those of the goal. Furthermore, during the process of moving to the visible location of the goal, the method utilizes gradients in the continuous language space,

TABLE IV  
REAL-WORLD NAVIGATION PERFORMANCE

Method	Success (%)	GDist (m)	Time (s)
Explicit Baseline	90.0	5.19	1.82
Ours (LAMP)	90.0	1.89	3.84

Real-world trial results. The explicit baseline uses pre-computed CLIP features for each node. The time reflects the total end-to-end duration for planning each query.

enabling path planning to the actual visible location of the goal. These results demonstrate that our method can effectively represent space, allowing the robot to precisely locate target objects without relying on memory-intensive explicit language maps. To quantitatively substantiate these findings, we report the key performance metrics from our real-world trial in Table IV, where LAMP significantly outperforms the explicit baseline in goal-reaching accuracy.

However, this approach has some limitations. Its effectiveness depends on the VLM’s accuracy. For instance, if other objects share similar visual features with the target, the model may mistakenly select wrong nodes. Moreover, when the target’s appearance is weak or ambiguous, it might struggle to identify the correct nodes. Therefore, enhancing the VLM’s performance and selecting objects with distinct and clear features are crucial for improving navigation accuracy.

## V. CONCLUSION

We introduce LAMP, an implicit language map for large-scale, zero-shot navigation using only RGB images. Beyond a simple implicit representation, LAMP is the first to leverage a language-driven continuous field for fine-grained path generation. Our approach combines a continuous neural network with a sparse topological graph, achieving efficient memory usage by dynamically generating language embeddings for path planning rather than storing them explicitly. Additionally, we enhance pose refinement by sampling poses near the coarse estimate and adding a distance penalty to avoid local minima. We also model embedding uncertainty in a Bayesian framework and use learned node-level features for efficient graph sampling. Simulation experiments demonstrate that LAMP surpasses traditional

grid- and node-based approaches in diverse scenarios, offering improved or comparable performance with substantially reduced memory requirements. Real-world experiments show that continuously encoding language features enables LAMP to detect and navigate toward objects not explicitly observed during mapping.

## REFERENCES

- [1] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3D mapping framework based on octrees," *Auton. Robots*, vol. 34, pp. 189–206, 2013.
- [2] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 834–849.
- [3] J. Zhang et al., "LOAM: LiDAR odometry and mapping in real-time," *Robot.: Sci. Syst.*, vol. 2, no. 9, pp. 1–9, 2014.
- [4] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 652–660.
- [5] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [6] M. Caron et al., "Emerging properties in self-supervised vision transformers," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2021, pp. 9650–9660.
- [7] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik, "LERF: Language embedded radiance fields," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2023, pp. 19729–19 739.
- [8] F. Engelmann, F. Manhardt, M. Niemeyer, K. Tateno, M. Pollefeys, and F. Tombari, "OpenNeRF: Open set 3D neural scene segmentation with pixel-wise features and rendered novel views," 2024, *arXiv:2404.03650*.
- [9] M. Qin, W. Li, J. Zhou, H. Wang, and H. Pfister, "LangSplat: 3D language Gaussian splatting," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 20051–20060.
- [10] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, "Cows on pasture: Baselines and benchmarks for language-driven zero-shot object navigation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23171–23181.
- [11] D. Shah et al., "LM-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *Proc. Conf. Robot Learn.*, 2023, pp. 492–504.
- [12] D. Azuma, T. Miyayoshi, S. Kurita, and M. Kawanabe, "ScanQA: 3D question answering for spatial scene understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 19129–19139.
- [13] P. Cascante-Bonilla, H. Wu, L. Wang, R. S. Feris, and V. Ordonez, "SimVQA: Exploring simulated environments for visual question answering," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 5056–5066.
- [14] K. Huang, Y. Han, J. Wu, F. Qiu, and Q. Tang, "Language-driven robot manipulation with perspective disambiguation and placement optimization," *IEEE Robot. Automat. Lett.*, vol. 7, no. 2, pp. 4188–4195, Apr. 2022.
- [15] R. Kurai et al., "Magicitem: Dynamic behavior design of virtual objects with large language models in a commercial metaverse platform," *IEEE Access*, vol. 13, pp. 19132–19143, 2025.
- [16] K. M. Jatavallabhula et al., "Conceptfusion: Open-set multimodal 3D mapping," 2023, *arXiv:2302.07241*.
- [17] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2023, pp. 10608–10615.
- [18] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Audio visual language maps for robot navigation," in *Proc. Int. Symp. Exp. Robot.*, 2023, pp. 105–117.
- [19] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard, "Hierarchical open-vocabulary 3D scene graphs for language-grounded robot navigation," in *Proc. 1st Workshop Vis.-Lang. Models Navigat. Manipulation*, 2024.
- [20] S. H. Arul et al., "VLP-NAV: Object navigation using visual language pose graph and object localization probability maps," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2024, pp. 7625–7632.
- [21] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, "Language-driven semantic segmentation," 2022, *arXiv:2201.03546*.
- [22] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 5580–5590.
- [23] S. Lee, K. Kang, S. Ha, and H. Yu, "Bayesian NeRF: Quantifying uncertainty with volume density for neural implicit fields," *IEEE Robot. Automat. Lett.*, vol. 10, no. 3, pp. 2144–2151, Mar. 2025.
- [24] S. Zhang, R. Xu, and Y. Quan, "Large graph layout optimization based on vision and computational efficiency: A survey," *Vis. Intell.*, vol. 1, no. 1, 2023, Art. no. 14.
- [25] J. Leskovec and C. Faloutsos, "Sampling from large graphs," in *Proc. 12th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2006, pp. 631–636.
- [26] A. Elfes, "Using occupancy grids for mobile robot perception and navigation," *Computer*, vol. 22, no. 6, pp. 46–57, 1989.
- [27] Davison, "Real-time simultaneous localisation and mapping with a single camera," in *Proc. 9th IEEE Int. Conf. Comput. Vis.*, 2003, pp. 1403–1410.
- [28] M. Montemerlo et al., "FastSLAM: A factored solution to the simultaneous localization and mapping problem," *AAAI/IAAI*, vol. 593598, pp. 593–598, 2002.
- [29] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [30] C. Cadena et al., "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1309–1332, Dec. 2016.
- [31] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3D semantic mapping with convolutional neural networks," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2017, pp. 4628–4635.
- [32] R. Dubé, A. Cramariuc, D. Dugas, J. Nieto, R. Siegwart, and C. Cadena, "Segmap: 3D segment mapping using data-driven descriptors," 2018, *arXiv:1804.09557*.
- [33] Q. Peng et al., "3D vision-language Gaussian splatting," 2024, *arXiv:2410.07577*.
- [34] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Commun. ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [35] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [36] M. A. Hasnat, J. Bohné, J. Milgram, S. Gentic, and L. Chen, "Von mises-fisher mixture model-based deep learning: Application to face verification," 2017, *arXiv:1706.04264*.
- [37] K. Nakadai, S. Masaki, R. Kojima, O. Sugiyama, K. Itoyama, and K. Nishida, "Sound source localization based on Von-Mises-Bernoulli deep neural network," in *Proc. IEEE/SICE Int. Symp. System Integration*, 2020, pp. 658–663.
- [38] A. Chang et al., "Matterport3D: Learning from RGB-D data in indoor environments," 2017, *arXiv:1709.06158*.
- [39] R.-H. Li, J. X. Yu, L. Qin, R. Mao, and T. Jin, "On random walk based graph sampling," in *Proc. IEEE 31st Int. Conf. Data Eng.*, 2015, pp. 927–938.
- [40] S. H. Lee, P.-J. Kim, and H. Jeong, "Statistical properties of sampled networks," *Phys. Rev. E—Stat. Nonlinear Soft Matter Phys.*, vol. 73, no. 1, 2006, Art. no. 016102.
- [41] Naver Labs, "M2: 3D high-precision mapping robot for digital twin data creation," Naver Labs Website. Accessed: Mar. 2, 2022. [Online]. Available: <https://www.naverlabs.com/en/storyDetail/234>