

# VLPG-Nav: Object Navigation Using Visual Language Pose Graph and Object Localization Probability Maps

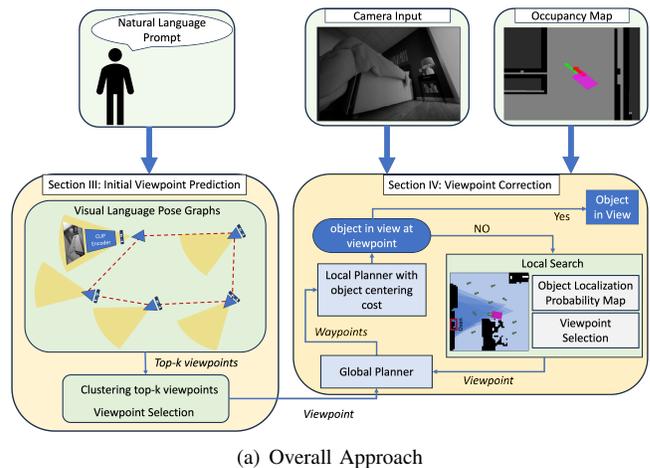
Senthil Hariharan Arul<sup>†\*,‡</sup>, Dhruva Kumar<sup>†</sup>, Vivek Sugirtharaj<sup>†</sup>, Richard Kim<sup>†</sup>, Xuewei (Tony) Qi<sup>†</sup>,  
Rajasimman Madhivanan<sup>†</sup>, Arnie Sen<sup>†</sup>, Dinesh Manocha<sup>‡</sup>

**Abstract**—We present VLPG-Nav, a visual language navigation method for guiding robots to specified objects within household scenes. Unlike existing methods primarily focused on navigating the robot toward objects, our approach considers the additional challenge of centering the object within the robot’s camera view. Our method builds a visual language pose graph (VLPG) that functions as a spatial map of VL embeddings. Given an open-vocabulary object query, we plan a viewpoint for object navigation using the VLPG. Despite navigating to the viewpoint, real-world challenges such as object occlusion, displacement, and the robot’s localization errors can prevent visibility. We build an object localization probability map that leverages the robot’s current observations and prior VLPG. When the object is not visible, the probability map is updated, and an alternate viewpoint is computed. In addition, we propose an object-centering formulation that locally adjusts the robot’s pose to center the object in the camera view. We evaluate the effectiveness of our approach through simulations and real-world experiments, evaluating its ability to successfully view and center the object within the camera’s field of view. VLPG-Nav demonstrates improved performance in locating the object, navigating around occlusions, and centering the object within the robot’s camera view, outperforming selected baselines in the evaluation settings.

## I. INTRODUCTION

Recent advances in motion planning, perception, and deep learning have enabled robots to operate in diverse environments such as homes, factories, hospitals, and shopping malls. With their rising prevalence in day-to-day environments, there is a demand for robots to interact seamlessly with humans and undertake a broader spectrum of tasks<sup>1</sup>. Accordingly, a crucial ability for robots is to comprehend natural language commands and leverage sensor observations to accomplish tasks.

In this regard, the problem of object navigation is especially relevant, where, given a textual prompt about an object, the robot locates and navigates to view the object. However, given the expanse of objects in a typical workspace and the ambiguity in natural language, where different words may be used for the same object, maintaining a database of all the objects (with the aliases) and their locations can be challenging. As a result, open-vocabulary detection methods



**Fig. 1:** (Fig. 1-a) The overall framework of VLPG-Nav. Given an object-related prompt, an initial viewpoint is computed using the visual language pose graph (VLPG) and clustering. Subsequently, we center the object in the camera image through our proposed object-centering cost function. If the object is not in view at the chosen viewpoint, we replan using an object localization probability map constructed using the local occupancy map and VLPG. (Fig. 1-b) A real-world example where the robot is tasked with viewing a *plant*. The figure shows the camera view from the initial viewpoint, and the robot successfully views the plant. (Fig. 1-c) In this case, we obstruct the initial viewpoint with an obstacle. The figure shows the obstructed camera view from the initial viewpoint. (Fig. 1-d) Due to occlusion, the local search identifies an alternative viewpoint, enabling the robot to achieve an unobstructed view of the plant.

are useful for assisting the robot in tackling such ambiguity and efficiently guiding it to the object landmark of interest.

In this paper, we consider the problem of navigating a robot to an object of interest within a household environment. Distinct from the field of object-goal navigation, we consider previously explored environments, and our focus extends beyond merely navigating the robot to the object but also to frame the object within the robot’s camera view. Our focus primarily lies in household robot applications, where memory and computational efficiency are crucial. These robots often operate continuously in the same environments, allowing them to leverage prior knowledge to navigate and view the objects efficiently.

\* Work done during the internship at Amazon.

This work is supported in part by ARO Grants W911NF2310046, W911NF2310352 and US Army Cooperative Agreement W911NF2120076.

<sup>†</sup> Authors are with Amazon Lab126, Sunnyvale, CA 94098, USA. Email: {dhruvkm, vssugirt, richk, qixuewei, rajasimm, senarnie}@amazon.edu

<sup>‡</sup> Authors are with the University of Maryland, College Park, MD 20740, USA. Email: {sarull, dmanocha}@umd.edu

<sup>1</sup><https://www.rd.com/article/future-robots/>

However, relying solely on prior knowledge can be unreliable, especially considering the possibility of object occlusion or displacement. Moreover, household environments are prone to changes, such as the introduction of new appliances, the removal of furniture, or the transformation of one room into another. Additionally, localization errors and sensor noise in affordable home robots can impede the robot’s ability to accurately reach predetermined viewpoints. Given these challenges, integrating current observations with prior knowledge becomes imperative to ensure task success. The method must be adaptable, incorporating recent information into prior knowledge to ensure that the robot not only navigates the current environment effectively but also leverages this updated knowledge for future queries.

### A. Related Works

The area of robot navigation has received significant attention over the past few decades. Classical navigation approaches [1], [2], [3], [4], [5] focus primarily on producing a safe and efficient path to a goal position. In contrast, applications such as visual question answering (VQA) require the robot to capture a specific area or object within its camera’s view for task success, rather than reaching a target goal position. The area of object goal navigation focuses on this problem [6], [7], but using predefined object classes during training can fail to generalize to the spectrum of objects in a typical workspace.

Recently, zero-shot learning methods have shown tremendous potential for computer vision applications such as image recognition [8], [9], object detection [10], [11] and can be utilized to guide object goal navigation tasks. Zhao et al. [12] present a zero-shot object goal visual navigation method that proposes an RL-trained semantic similarity network to prevent over-fitting to seen classes. EmbCLIP [13] leverages a contrastive language image pretraining encoder for visual navigation tasks.

Visual-language models (VLMs) [14], [15], [16], [17], [18], [19] are networks pre-trained on internet-scale data that ground language prompts to visual observations and have shown strong generalization and reasoning ability. These models have been used for object navigation to match the robot’s observation with the natural language command [20], [21], [22], [23]. In LM-Nav [20], VLMs are used to match relevant landmarks in graphs based on the prompt to navigate the robot. CoW [21] uses CLIP and GradCAM [24] to detect the object of interest in the environment, but GradCAM-measured regions of interest can be noisy [22].

VLMaps [22] performs spatial grounding by generating a dense grid map of the environment with the VLM embedding as the third dimension and have the advantage of localizing the object in 3D space accurately. However, memory requirements may pose challenges for low-compute robots, and incorporating recent changes in the environment may not be direct. In [25], they introduce common sense into a frontier-based exploration, which increases efficacy and generalizability. Du et al. [26] consider object goal navigation from a navigation standpoint and propose a history-inspired

navigation learning framework, and agents are observed to update states more effectively.

In many real-world applications, robots repeatedly operate in the same environment, and leveraging pre-existing knowledge can be beneficial for choosing a suitable viewpoint. However, these environments exhibit some level of dynamism, wherein the objects can get displaced, and the surroundings can change and potentially obstruct the view of the object. Consequently, while prior knowledge provides a good initial guess, it isn’t solely reliable. To complete the object navigation task, in addition to the prior knowledge, the robot may need to locally correct and search for the object until it’s in the sensor view.

### B. Main Contributions

We present VLPG-Nav, an object navigation method designed for low-compute household robots. When provided with a textual prompt about an object, our approach navigates the robot to an instance of the object in the environment using prior knowledge of the environment and the robot’s current observation. Given the limited computing and memory power in affordable consumer robots, we desire a memory-efficient method that allows for easy integration of new environmental information. Furthermore, the robot’s final pose should provide a clear and discernible viewpoint for external observers, indicating the specific object the robot is focusing on. Our approach has three main components:

- 1) **Initial Viewpoint Prediction:** We construct a visual language pose graph (VLPG) during environment exploration, which is a graph of the robot’s pose and the corresponding VL embedding. Given a textual prompt, we use VLPG to identify a set of relevant nodes computed based on cosine similarity, and the viewpoints are clustered to compute a suitable viewpoint.
- 2) **Object Centering:** We propose a cost function to locally correct the robot’s pose and center the object in the camera image. Specifically, the cost function orients the robot and moves it toward the object of interest once the initial viewpoint is reached.
- 3) **Local Search:** When an object is undetected at the chosen viewpoint, we propose a local search method to replan an alternative viewpoint to bring the object into view. We propose a probability map based on the VLPG clusters and the local occupancy map, estimating the likelihood of regions where the object can be localized.

We evaluate the effectiveness of our approach using a non-holonomic ground robot navigating toward diverse object goals in both a simulated and real-world home environment. The environment is a furnished 2-bedroom house, complete with common electronic appliances, furniture, plants, and more for a realistic test setting. We highlight the significance of the three components of our method by evaluating them in terms of pixel, and angular error in centering the object in the image space and camera field-of-view. Further, we propose a metric success weighted angular error (SAE) to compare the baselines in terms of task success.

Symbols	Definitions
$x, y$	2-D position in the global frame
$\theta$	Orientation in the global frame
$\mathbf{x}_r$	Robot’s Pose $(x, y, \theta)$
$\mathbf{x}_{vp}$	A viewpoint $(x, y, \theta)$
$f_{ov}(\mathbf{x})$	Field-of-view from pose $\mathbf{x}$
$\mathcal{C}_{vp}$	A set of viewpoints in a cluster
$\mathcal{S}_{vp}$	A set of randomly sampled viewpoints
$p_{occ}$	Local occupancy map for the Robot

TABLE I: Symbols and their definitions

## II. PRELIMINARIES

In this section, we provide an overview of various concepts used in our approach. Table I summarizes the symbols and notations used in the paper.

### A. Object Navigation and Framing

We tackle the challenge of object navigation within household environments, where the goal is for the robot to autonomously navigate and position a specified object within its camera view. Our problem differs from the related field of object-goal navigation methods in that we operate under the assumption of having prior knowledge about the environment. This is reasonable as household robots often operate repeatedly in the same environment. Furthermore, our objective extends beyond mere navigation; we aim to ensure effective capture of the object within the robot’s camera view. The optimal behavior is when the object of interest is brought into the robot’s camera view and centered in the image space. Given that a household environment may contain multiple instances of the same object, we consider viewing any instance of the object as a successful outcome.

### B. Assumption

Natural language instructions from humans are often coarse-grained rather than a single object category. A typical query from a human is likely to be “Did I leave the stove on?” than “View stove.” However, either prompt requires the robot to navigate to the same object.

In this work, we do not consider the task of breaking down the ambiguous, coarse-grained instructions (Did I leave the stove on?) into simpler keyphrases (stove). Given the recent success of Large Language Models (LLMs) (such as GPT 3 [27]), we assume such a model exists and provides the object-specific keyphrase as input to our approach. Moreover, at any timestep  $t$ , the robot can access its 2D position, orientation, camera input, and occupancy grid map of its local environment.

### C. Object Detection

Grad-CAM [24] a class activation mapping (CAM) approach identifies relevant regions in the image affecting the classification score. Similar to [21], we utilize CLIP [14] and Grad-CAM for zero-shot object detection. We use GradCAM to create a saliency map over the input image pertaining to the input text prompt. We threshold the saliency map and construct a bounding box to localize the region of interest in the image. The bounding box information guides the local planner to center the object in the image space and to update our object localization probability map.

### D. Navigation Framework

The robot navigation framework utilizes a global planner for high-level waypoint generation and a local planner for trajectory planning. Given a goal location in the environment, a global planning algorithm (Theta\* [28] in our case) generates a sequence of waypoints to the goal. The local planner, a stochastic MPC method [5] plans the trajectories for the robot to follow the waypoints while avoiding obstacles. The local planner evaluates a set of robot trajectories and selects the one that minimizes the trajectory cost function, given by:

$$J = \sum_{i=1}^N p_{s_i} \cdot J_{progress} + J_{action} + (1 - p_{s_i}) \cdot J_{collision}.$$

For evaluations, a trajectory is split into  $N$  segments. The progress cost measures the distance progress made over the trajectory. The action cost encourages smooth trajectories by limiting large control inputs, and the collision term limits the robot’s velocity in unsafe scenarios. Here,  $p_{s_i}$  is a measure of the probability that the trajectory until the  $i^{th}$  segment remains collision-free. Our proposed object-centering cost is included in the progress term.

### E. Exploration

When deployed in an unknown environment, robots lack knowledge about obstacles and the path to the goal. Exploration enables the robot to gather this information by moving around and constructing a map of the environment. The constructed map could be referred to facilitate subsequent navigation. A popular algorithm is frontier-based exploration [29], where a frontier is a portion of the region separating the explored and unexplored regions which is subsequently explored to map the environment.

## III. INITIAL VIEWPOINT PREDICTION

Robots often operate continuously in the same environment, allowing them to accumulate valuable knowledge that can guide object navigation. To leverage this, we construct a visual language pose graph (VLPG), which represents previously gathered environmental knowledge from exploration. During task execution, the VLPG helps compute a robot pose that is highly likely to provide a view of the object of interest. This section outlines our approach for constructing the VLPG and generating a viewpoint prediction based on an input textual prompt.

### A. Visual Language Pose Graphs (VLPG)

State-of-the-art SLAM systems use pose graphs to localize the robot and construct a spatial representation of the environment [30]. The graph consists of SE(3) robot poses and images captured from the pose as nodes, while the edges are relative pose measurements. Given the comprehensive coverage of the environment during the initial mapping phase, the robot’s path has transited and viewed most of the environment. This encourages using the mapping process to create prior knowledge representative of the environment.

We build on the pose graph by augmenting each node with the VL embeddings of the robot’s camera view from

the pose. We call this visual language pose graph (VLPG). As the robot explores its environment, we augment the VLPG by adding a new node approximately every 100ms.

Using VLPG for the task of viewing objects has multiple advantages:

- **Memory Efficiency:** For an environment of  $\sim 10 \times 10$  sq.m, VLPG is observed to have a graph of size  $\sim 2000 \times 512$ . While, a VL grid map such as in [22] would produce a VL map of size  $200 \times 200 \times 512$  (for a resolution of 0.05 m/cell), which is approximately 20x larger in terms of memory requirements than VLPG. Though VLPG may not localize the objects in the 3D space as accurately as in [22], we can compute a coarse localization by superimposing the camera field of view from poses that view the object (as seen in Fig. 5).
- **Easy Viewpoint Selection:** As VLPG nodes contain the robot’s pose information, it is a collection of viewpoints, which facilitates the selection of object viewpoints. For example, given an object at a higher elevation, a suitable viewpoint is likely farther away from the object to observe it in the field of view. With the VLPG, we automatically obtain this information as it is a collection of camera viewpoints.
- **Recent Environment Knowledge:** During long-term execution, the environment can change; new appliances may emerge, furniture can move around, etc. Hence, prior knowledge should be updated with recent information about the environment. We can directly append new nodes to the VLPG to incorporate recent information. When adding a new node, we compare the CLIP embedding of the robot’s current view with the CLIP embeddings in the VLPG. The new node ( $q^*$ ) is added to the VLPG if it satisfies the following similarity condition for a user-defined threshold ( $\epsilon$ ):

$$\frac{CLIP(q^*) \cdot CLIP(q)}{\|CLIP(q^*)\|_2 \|CLIP(q)\|_2} < \epsilon, \quad \forall q \in VLPG. \quad (1)$$

### B. Generating Initial Guesses

When deployed, the constructed VLPG is leveraged to choose a suitable viewpoint. Given an object of interest, we create two text prompts:

$$\begin{aligned} & \text{[A photo of \{object\},} \\ & \text{A photo of something else].} \end{aligned} \quad (2)$$

We generate a fixed-length embedding for each prompt by passing them through CLIP’s text encoder and compute the cosine similarity scores between the text embedding and every camera view embedding in VLPG. The VLPG nodes resulting in a higher cosine similarity score for prompt 1 have a higher probability of viewing the object.

Though the VLPG node with the highest similarity score could be a suitable viewpoint, it does not provide information on how well it views the object. That is, the object may be only partially in view or could even be a false positive. To address this issue, we use the  $top-k$  nodes with the highest similarity score. If multiple nodes view a region, we can consider the region as having high confidence in localizing

the object. In this regard, we cluster the  $top-k$  nodes to remove noise (false positives) and identify clusters of poses viewing the object with a high probability.

For every pose in the  $top-k$  VLPG nodes, we transform the viewpoint from 3-D  $(x, y, \theta)$  to 4-D  $(x, y, \cos(\theta), \sin(\theta))$  to avoid the discontinuity in orientation values. We utilize the DB-Scan algorithm to cluster the viewpoints and remove noise. The *best guess* viewpoint is chosen as the viewpoint in the largest cluster  $C_{vp}$  whose field of view has the maximum coverage of the region viewed by the other viewpoints in the cluster. The *best guess* is chosen as the viewpoint that maximizes the following cost function.

$$\mathbf{x}_{init} = \max_{\mathbf{x} \in C_{vp}} fov(\mathbf{x}) \cap \left( \bigcup_{\mathbf{y} \in C_{vp}, \mathbf{y} \neq \mathbf{x}} fov(\mathbf{y}) \right) \quad (3)$$

Here,  $\mathbf{x}_{init}$  is the chosen viewpoint. The cost function chooses a viewpoint whose field of view overlaps the most with the union of the area viewed by other viewpoints. The robot navigates towards this chosen pose to view the object.

## IV. VIEWPOINT CORRECTION

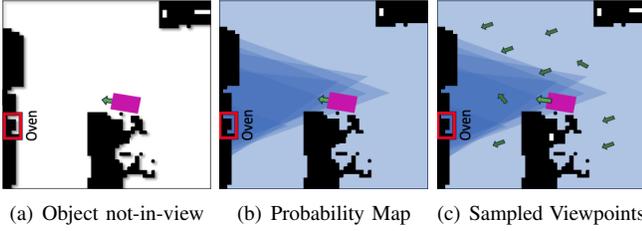
In an ideal scenario, we desire the robot to reach a final viewpoint where the object of interest is clearly visible. It is equally crucial for the robot’s pose to be *legible*, i.e., for an external observer, the robot’s end pose must clearly indicate the chosen object of focus. Thus, as the robot navigates to the chosen viewpoint, we want the object to be in view, centered in the camera view, and occupying a significant portion of the image. However, the initial viewpoint may fail to satisfy this requirement due to localization uncertainty, occlusions, or object displacements. This results in two unsuccessful cases: (1) the object is in view but not centered, and (2) the object is not in view. This section outlines our approach to tackling these unsuccessful cases using current information about the robot’s local environment and prior knowledge from the VLPG.

### A. Object in View

As the robot nears the chosen viewpoint, the local planner uses the object location in the image to correct the robot’s pose to center the object in the camera view. We use CLIP and GradCAM to localize the object in the image frame (Section 3) by constructing a bounding box for the object of interest (when detected). To perform object centering, our approach incorporates an orient and zoom cost into the local planner that aligns the object at the center and zooms in by moving closer to the object. Since the local planner optimizes over a finite time horizon, we require predictions of bounding box coordinates over the horizon. Given the bounding box center at the initial timestep, we use the image Jacobian equations [31] to predict the position of the bounding box centers based on the robot’s velocities over the planning horizon.

The orient cost ( $J_{orient}$ ) minimizes the error between the bounding box center and the image center along the horizontal image axis.

$$J_{orient} = c_x^{t+1^2} - c_x^{t^2} \quad (4)$$



**Fig. 2: Local Search:** We illustrate a scenario in which the robot is tasked with viewing an oven. (Fig. 3-a) The initial viewpoint guess results in the robot’s view of the oven being occluded. (Fig. 3-b) Based on the viewpoint cluster obtained from VLPG and clustering, we compute a probability map for localizing the oven. The obstacles are represented in black; regions with higher intensity blue are more likely to localize the object. (Fig. 3-c) Using the probability map, we sample a set of viewpoints from which a suitable alternative viewpoint is chosen based on our optimization objective.

Given a trajectory, the orient function is defined as the change in squared error value over the trajectory segments, with  $c_x$  representing the error.

To achieve a zoomed image, the robot can move towards the object. We define the following zoom cost function.

$$J_{zoom} = v \cdot \max(\cos \theta_c, 0) \cdot h \cdot \rho \quad (5)$$

$$\rho = \left( \frac{\max(d_{obj} - d_{thresh}, 0)}{d_{max}} \right) \quad (6)$$

Here,  $h$  is the timestep, and  $\theta_c$  is the angle between the robot heading and the line of sight between the robot and the object center (approximated with bounding box center).

When the robot orients towards the object, the  $\cos \theta_c$  component approaches one and the zoom terms become active. Until the distance to the object is lower than the threshold  $d_{thresh}$ , the zoom cost provides a gradient to move toward the object. Ideally, we only need the zoom component if the robot is far from the object.

### B. Object Not-in-view

In certain cases, the initial guess viewpoint may fail to bring the object into the camera’s view, and the robot may need to search the surrounding environment to locate it. In this work, we assume the region around the initial viewpoint to be more likely to localize the object of interest, and we utilize the viewpoint clusters from VLPG to construct an object localization probability map ( $p_l$ ). Consequently, we use the probability map to compute alternative viewpoints to view the higher probability regions.

First, we initialize a probability map  $p_{nl}$ , which is the probability of not localizing the object. The probability map is identical in size to the robot’s local occupancy map ( $p_{occ}$ ), a  $6m \times 6m$  grid map, with its grid cell initialized as

$$p_{nl}(i, j) = \begin{cases} 1, & \text{for } p_{occ}(i, j) = 1 \\ 0.9, & \text{for } else \end{cases}$$

Here,  $i, j$  indicates the grid location, and  $p_{occ}(i, j) = 1$  implies the grid cell is occupied. The global coordinates for a grid location  $(i, j)$  are given by  $(x(i), y(j))$ . Next, for each viewpoint in the cluster from VLPG, we use their camera

fields-of-view to update their probability map  $p_{nl}$ . For each cell  $(i, j)$  in the grid map, we update the probability as,

$$p_{nl}(i, j) = p_{nl}(i, j) \cdot (1 - I(i, j, \mathbf{x}_{vp}) \cdot p_l^*) \quad \forall \mathbf{x}_{vp} \in C_{vp} \text{ and } p_{nl}(i, j) \neq 1, \quad (7)$$

where

$$I(i, j, \mathbf{x}_{vp}) = \begin{cases} 1, & \text{for } (i, j) \in fov(\mathbf{x}_{vp}) \text{ and } p_{nl}(i, j) \neq 1 \\ 0, & else \end{cases}$$

and  $p_l^*$  is a user-defined probability threshold. After accounting for every viewpoint in the cluster, the probability of localizing the object at any grid cell can be computed as

$$p_l(i, j) = 1 - p_{nl}(i, j).$$

Consequently, the object localization probability map ( $p_l$ ) provides a likelihood over the  $6m \times 6m$  grid map for localizing the object of interest based on VLPG. From the probability map, we compute an alternative viewpoint to view the high-likelihood regions. To select an alternative viewpoint, we start by sampling a set of poses  $\mathcal{S}_{vp}$  on the local map. A suitable alternate viewpoint is selected from the samples based on  $p_l$  by minimizing against a cost function. We define the cost function as,

$$J(\mathbf{x}_{vp}) = w_d \cdot J_{distance} + w_q \cdot J_{quality} + w_{obs} \cdot J_{obstacle} \quad (8)$$

Here,  $w_d, w_q, w_{obs}$  are user-defined weights. The distance cost is a measure of how far the new viewpoint is from the current position. The quality cost is a measure of how probable it is to view the object from the re-planned viewpoint. The collision cost encourages having clearance around the robot from the re-planned viewpoint.

$$J_{distance} = (\mathbf{x}_{vp} - \mathbf{x}_r) I (\mathbf{x}_{vp} - \mathbf{x}_r)^T \quad (9)$$

$$J_{quality} = \sum_{(x, y) \in fov(\mathbf{x}_{vp})} 1 - p_l(x, y) \quad (10)$$

$$J_{obstacle} = e^{-d_o^2} \quad (11)$$

Here,  $I$  is the identity matrix and  $d_o$  is the distance to the closest obstacle from the replanned viewpoint.

$$\mathbf{x}_{vp} = \{ \mathbf{x}_{vp} \mid J(\mathbf{x}_{vp}) = \min_{\mathbf{x}'_{vp} \in \mathcal{S}_{vp}} J(\mathbf{x}'_{vp}) \} \quad (12)$$

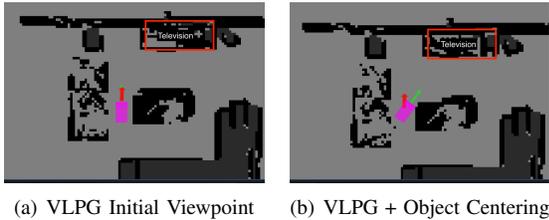
The robot navigates to the chosen viewpoint. If the object is observed, the local planner corrects the robot’s pose to keep the object centered (as discussed in Section 4A), and the task is considered a success. If the object is not in view at the replanned viewpoint, the probability values for the region currently in the camera view are reduced as

$$p_l(x, y) = 0.2 \cdot p_l(x, y) \quad \forall (x, y) \in fov(\mathbf{x}_{vp}).$$

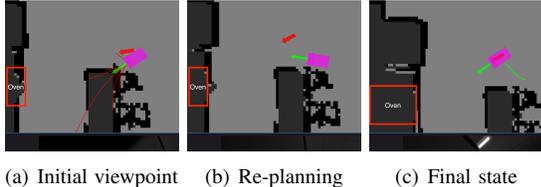
The robot re-computes a new viewpoint using this updated probability map. The replanning process continues until the object is detected or maximum re-computations are reached.

## V. EVALUATIONS

In this section, we outline our experimental setup and highlight the performance of our approach, VLPG-Nav.



**Fig. 3:** The red arrow indicates the initial viewpoint, while the green arrow represents the robot’s final pose. The object centering directs the robot to the object of interest (denoted by red box).



**Fig. 4:** An example demonstrating “Local Search”, where the robot is assigned the task of viewing an *oven*. (Fig. 4(a)) The initial viewpoint guess results in the robot view being occluded. (Fig. 4(b)) The local search identifies a suitable alternative viewpoint (red), which gets a better view of the object. (Fig. 4(c)) The robot then moves to the replanned viewpoint and successfully views the object.

#### A. Experimental Setup

We evaluate our proposed method in both simulation and real-world scenarios, considering a non-holonomic robot and various objects of interest. The simulations are conducted using Gazebo and Unreal Engine, with Unreal providing photo-realistic images for the robot, while a corresponding home environment in Gazebo is used for simulating the robot. Our object localization uses the ResNet50 version of CLIP and utilizes GradCAM to compute bounding boxes of regions of interest. The local planning optimization is implemented using the NLOPT library. For real-world experiments, our mobile ground robot used two Qualcomm QCS605 octa-core ARM processors, each with 2GB RAM and with two cores running at 2.5GHz and six cores running at 1.7GHz. The local search replanning, CLIP-based evaluations for computing the initial viewpoint guess, and object centering are performed on a laptop equipped with an Intel Core-i7 6th generation processor, utilizing sensor inputs from the robot. The computed object bounding boxes and viewpoints are then communicated back to the robot which execute object centering and navigation.

#### B. Evaluation Metrics

Prior object navigation works [22] have typically defined success as the robot reaching within a specified distance threshold of its target object. However, our approach considers success when the robot successfully navigates to the object of interest and centers it within its camera view. This accounts for occlusion and ensures that success is only acknowledged when the object is visible to the robot’s camera. In this regard, we define a new evaluation metric: the Success weighted measure of Angular Error (SAE). This metric quantifies the angular deviation between the robot’s

Object Prompts	Frontier Exploration (Baseline 1)	VLPG (Baseline 2)	VLPG + Obj. Centering (Baseline 3)	VLPG-Nav (Proposed)
Bedside Lamp	-	144	9	9
Cradle	-	33	7	7
Television	-	108	10	10
Beanbag	-	96	8	8
Oven	-	-	-	22
Buddha Statue	15	-	-	-
Bar Stool	-	12	3	3
Coffee Table	-	31	4	4

**TABLE II:** We tabulate the distance of the bounding box centers from the center of the image (in pixels). Since the bounding box information is noisy, we average the bounding box centers and round it to the nearest integer. Frontier exploration found the Buddha statue as it was located near the robot’s start position. Though the proposed method was unable to view the Buddha statue in simulation (as the initial guess was incorrect), it was able to view it in our real-world experiments.

Object Prompts	Frontier Exploration (Baseline 1)	VLPG (Baseline 2)	VLPG + Obj. Centering (Baseline 3)	VLPG-Nav (Proposed)
Bedside Lamp	-	43.86	5.12	5.12
Cradle	-	21.92	15.58	15.58
Television	-	50.66	1.77	1.77
Beanbag	-	24.14	9.34	9.34
Oven	-	-	-	22.2
Dishwasher	-	-	-	-
Buddha Statue	23.31	-	-	-
Bar Stool	-	0.66	4.47	4.47
Coffee Table	-	34.15	28.58	25.58
SAE ( $\uparrow$ )	0.12	0.66	0.73	<b>0.85</b>

**TABLE III:** We tabulate the angle made by the robot’s orientation with the line of sight between the robot and the object’s ground truth positions. This angular value denotes the error in centering the object in the camera field of view. The last row lists the success-weighted measure of angular error (SAE). We highlight the proposed method (VLPG-Nav) results in the best SAE. Frontier exploration was able to view the Buddha statue as it was partly in view from the robot’s starting position. Though the proposed method failed to view the Buddha statue in simulation (as the initial guess was incorrect), it was viewed successfully in our real-world experiments.

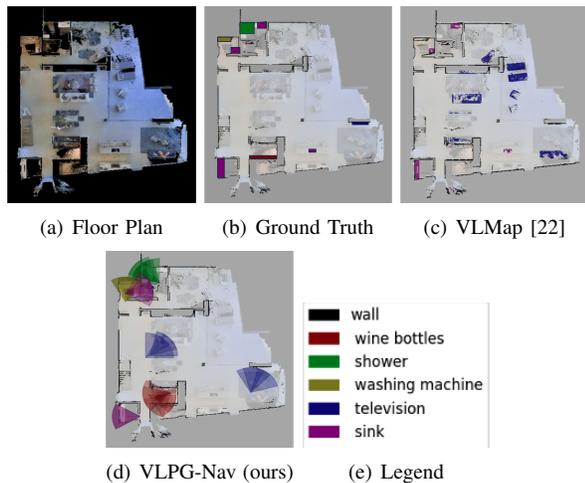
orientation and the line of sight between the robot and the object’s center.

$$SAE = \frac{1}{N} \sum S_i \cdot \exp^{-\left(\frac{\Delta\theta}{\Delta\theta_{max}}\right)^2} \quad (13)$$

Here,  $S_i$  is an indicator function that has a value of 1 when the object is in the camera field of view, else is 0.  $\Delta\theta_{max}$  is set to 90 deg and is a measure of the horizontal field of view of the camera. An SAE of 1 indicates that the robot viewed and centered the object in its camera view in every instance. While an SAE of 0 indicates the robot never succeeds in viewing the object.

#### C. Object Navigation

We consider a household environment where the robot is tasked with observing objects of interest. In our simulation environment, we consider eight randomly selected objects and evaluate our proposed VLPG-Nav method against three baseline approaches. The first baseline adopts a purely exploratory method, utilizing frontier exploration [29] and CLIP [14] + GradCAM [24] to identify and approach the object. The second baseline uses the VLPG to select an initial viewpoint and guide the robot towards it. The third baseline uses VLPG for initial viewpoint selection and combines with our object-centering method to refine the final pose for optimal object viewing. The proposed approach (VLPG-Nav) uses VLPG and object centering and performs a local



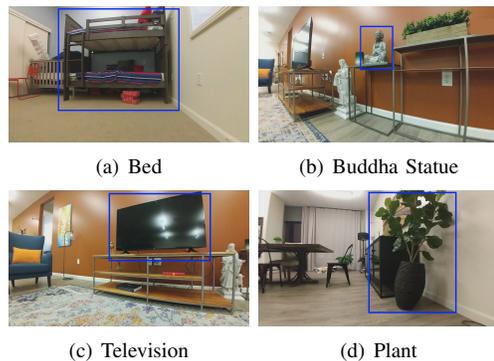
**Fig. 5:** An illustrative example of the viewpoints generation using prior environment knowledge. (Fig. 5(a)) The environment’s floor plan. (Fig. 5(b)) Ground truth locations of the object of interest on the 2D floor plan. (Fig. 5(c)) The 2D object localization computed using VLMap [22]. (Fig. 5(d)) The figure represents the object viewpoint and camera FoV generated by our approach. In this case, the walls (depicted in black) are from the occupancy map and are not computed using the VLPG. (Fig. 5(e)) Depicts the objects of interest with specific colors used on the map. The proposed method improves object localization and consequently object goal navigation since object viewpoints are linked to the SLAM pose graph, which gets optimized over time. Computationally, VLPG is memory efficient, making it effective for deployment on low-compute robots.

search to find the object when the object is not in view. The evaluation serves as an ablation study to emphasize the importance of the three components of our proposed method for reliable object navigation. We regard the execution as a failure when an incorrect object is detected or the object is not detected within a maximum simulation time of 5 minutes.

In Table II, we compare the methods in terms of their ability to view the object head-on. Since our method uses the bounding box information to center the object, we tabulate the error in terms of the pixel distance between the bounding box centroid and the image center along the horizontal image axis. We observe that the object-centering component results in a smaller error than relying purely on the initial viewpoint guess.

Since the bounding boxes are generated using GradCAM, they need not cover the entire object of interest. In Table III, we use the ground truth positions of the robot and the object from the end state to measure the angular deviation between the robot’s line of sight and the object center. On average, we observe that the object-centering cost function helps to center the object in the image. Moreover, we tabulate the SAE values in the table. We observe the proposed method results in the best SAE values compared to the baselines.

In the case of an *oven*, the kitchen counter obstructs the robot’s view of the oven at the end state. This results in failure to view the object in all baseline methods except for the proposed VLPG-Nav. The local search helps identify an alternative viewpoint to bring the oven into view. In the



**Fig. 6: Real-World Experiment:** Snapshots from the robot’s camera captured at the conclusion of the object navigation task. The robot successfully reached the considered objects (Bed, Buddha Statue, Television, Plant). To highlight the objects of interest, bounding boxes have been added to the figures for reference.

case of the Buddha statue, the initial viewpoint guess was in a completely different area of the environment, and local search could not help bring the object into view. As can be observed in Fig. 6-b, the robot navigated successfully to the Buddha statue in our real-world experiment.

In the case of random exploration, many false positives resulted in incorrectly identifying the object. In the case of the Buddha statue, random exploration was successful as the statue was near the robot’s starting position.

#### D. Object Localization

In Fig. 5, we illustrate the viewpoints clusters computed by VLPG (Fig. 5-d) for a set of objects in a sample environment from Matterport 3D dataset [32]. We also show the ground truth and VLMap [22] generated maps as comparisons. In VLMaps, the objects are localized in the 2D map. In contrast, our approach does not directly localize objects in 2D space; instead, we identify a viewpoint in 2D that approximates the object’s location. These viewpoints are then used to construct the probability map for local search.

#### E. Real-world Experiments

We evaluated our approach in real-world scenarios by considering a set of object prompts. Fig. 6 illustrates the robot’s camera view after navigating and framing the different objects. We observe that for the prompts *bed*, *Buddha statue* the objects are centered in the camera view. While for *Television*, *Plant* the object has a small angular deviation between the robot’s orientation and the object’s line of sight. In Fig. 1(b-d), we showcase the behavior when the chosen initial object viewpoint is occluded. We occlude the object *plant*, and our method successfully replans an alternative viewpoint using the object localization probability map.

#### F. Discussion

We highlight that our method is not specific to CLIP and can be replaced with other models such as GLIP [17], FLAVA [16], etc. In our implementation, we utilize CLIP for its comparatively faster online computation as we require the bounding box computation in our object-centering component to run at approximately 10Hz for our local planner.

## VI. CONCLUSIONS

We presented VLP-Nav, a visual language navigation framework for object navigation. We build a visual language pose graph (VLP) as a sparse mapping of the object viewpoints in the environment and is used to predict a relevant viewpoint given an open-vocabulary object prompt. Our proposed pose correction method successfully replans to avoid object occlusions that hinder detections and orients the robot to center the object in the camera view. Our method has some limitations. First, the success of our local search phase depends on the quality of the initial *top-k* viewpoints. Second, CLIP + GradCAM can produce noisy object detections, leading to incorrect attention maps in the image. Additionally, the cost function for selecting an alternative viewpoint during the local search does not account for the robot's non-holonomic dynamics. In future work, we aim to develop improved cost functions for selecting alternative viewpoints within the probability map, with a particular focus on incorporating kinematic constraints during viewpoint re-planning. Additionally, we intend to explore alternative embeddings beyond CLIP for constructing the VLP.

## REFERENCES

- [1] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [2] S. Thrun, M. Bennewitz, W. Burgard, A. Cremers, F. Dellaert, D. Fox, D. Hahnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz, "Minerva: a second-generation museum tour-guide robot," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, vol. 3, 1999, pp. 1999–2005 vol.3.
- [3] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research: The 14th International Symposium ISRR*. Springer, 2011, pp. 3–19.
- [4] S. H. Arul and D. Manocha, "Dcad: Decentralized collision avoidance with dynamics constraints for agile quadrotor swarms," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1191–1198, 2020.
- [5] J. J. Park, C. Johnson, and B. Kuipers, "Robot navigation with model predictive equilibrium point control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 4945–4952.
- [6] A. Pal, Y. Qiu, and H. Christensen, "Learning hierarchical relationships for object-goal navigation," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, J. Kober, F. Ramos, and C. Tomlin, Eds., vol. 155. PMLR, 16–18 Nov 2021, pp. 517–528.
- [7] H. Du, X. Yu, and L. Zheng, "{VTN}et: Visual transformer network for object goal navigation," in *International Conference on Learning Representations*, 2021.
- [8] V. K. Verma, G. Arora, A. Mishra, and P. Rai, "Generalized zero-shot learning via synthesized examples," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 4281–4289.
- [9] M. F. Naem, Y. Xian, F. Tombari, and Z. Akata, "Learning graph embeddings for compositional zero-shot learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 953–962.
- [10] S. Rahman, S. Khan, and F. Porikli, "Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts," in *Asian Conference on Computer Vision*. Springer, 2018, pp. 547–563.
- [11] C. Yan, X. Chang, M. Luo, H. Liu, X. Zhang, and Q. Zheng, "Semantics-guided contrastive network for zero-shot object detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [12] Q. Zhao, L. Zhang, B. He, H. Qiao, and Z. Liu, "Zero-shot object goal visual navigation," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 2025–2031.
- [13] A. Khandelwal, L. Weihs, R. Mottaghi, and A. Kembhavi, "Simple but effective: Clip embeddings for embodied ai," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 14 829–14 838.
- [14] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [15] B. Li, K. Q. Weinberger, S. Belongie, V. Koltun, and R. Ranftl, "Language-driven semantic segmentation," in *International Conference on Learning Representations*, 2022.
- [16] A. Singh, R. Hu, V. Goswami, G. Couairon, W. Galuba, M. Rohrbach, and D. Kiela, "Flava: A foundational language and vision alignment model," *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15 617–15 629, 2021.
- [17] L. H. Li, P. Zhang, H. Zhang, J. Yang, C. Li, Y. Zhong, L. Wang, L. Yuan, L. Zhang, J.-N. Hwang, K.-W. Chang, and J. Gao, "Grounded language-image pre-training," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2022, pp. 10 965–10 975.
- [18] J. Li, D. Li, C. Xiong, and S. Hoi, "Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation," in *International Conference on Machine Learning*. PMLR, 2022, pp. 12 888–12 900.
- [19] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [20] D. Shah, B. Osinski, B. Ichter, and S. Levine, "LM-nav: Robotic navigation with large pre-trained models of language, vision, and action," in *6th Annual Conference on Robot Learning*, 2022.
- [21] S. Y. Gadre, M. Wortsman, G. Ilharco, L. Schmidt, and S. Song, "Clip on wheels: Zero-shot object navigation as object localization and exploration," *ArXiv*, vol. abs/2203.10421, 2022.
- [22] C. Huang, O. Mees, A. Zeng, and W. Burgard, "Visual language maps for robot navigation," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, London, UK, 2023.
- [23] V. S. Dorbala, G. Sigurdsson, R. Piramuthu, J. Thomason, and G. S. Sukhatme, "Clip-nav: Using clip for zero-shot vision-and-language navigation," in *CoRL 2022 Workshop on Language and Robot Learning*, 2022.
- [24] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626.
- [25] K. Zhou, K. Zheng, C. Pryor, Y. Shen, H. Jin, L. Getoor, and X. E. Wang, "Esc: Exploration with soft commonsense constraints for zero-shot object navigation," *arXiv preprint arXiv:2301.13166*, 2023.
- [26] H. Du, L. Li, Z. Huang, and X. Yu, "Object-goal visual navigation via effective exploration of relations among historical navigation states," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 2563–2573.
- [27] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. J. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, "Language models are few-shot learners," *ArXiv*, vol. abs/2005.14165, 2020.
- [28] K. Daniel, A. Nash, S. Koenig, and A. Felner, "Theta\*: Any-angle path planning on grids," *J. Artif. Int. Res.*, vol. 39, no. 1, p. 533–579, sep 2010.
- [29] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. Towards New Computational Principles for Robotics and Automation*, 1997, pp. 146–151.
- [30] F. Dellaert and M. Kaess, *Factor Graphs for Robot Perception*. Now Publishers Inc., August 2017.
- [31] S. Hutchinson, G. Hager, and P. Corke, "A tutorial on visual servo control," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 5, pp. 651–670, 1996.
- [32] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *International Conference on 3D Vision (3DV)*, 2017.