

Zero-shot Active Visual Search (ZAVIS): Intelligent Object Search for Robotic Assistants

Jeongeun Park¹, Taerim Yoon¹, Jejoon Hong², Youngjae Yu³, Matthew Pan⁴, and Sungjoon Choi¹

Abstract—In this paper, we focus on the problem of efficiently locating a target object described with *free-form text* using a mobile robot equipped with vision sensors (e.g., an RGBD camera). Conventional active visual search predefines a set of objects to search for, rendering these techniques restrictive in practice. To provide added flexibility in active visual searching, we propose a system where a user can enter target commands using *free-form text*; we call this system *Zero-shot Active Visual Search (ZAVIS)*. ZAVIS detects and plans to search for a target object inputted by a user through a semantic grid map represented by static landmarks (e.g., desk or bed). For efficient planning of object search patterns, ZAVIS considers commonsense knowledge-based co-occurrence and predictive uncertainty while deciding which landmarks to visit first. We validate the proposed method with respect to SR (success rate) and SPL (success weighted by path length) in both simulated and real-world environments. The proposed method outperforms previous methods in terms of SPL in simulated scenarios, and we further demonstrate ZAVIS with a Pioneer-3AT robot in real-world studies.

I. INTRODUCTION

In this paper, we focus on the problem of efficiently locating a target object using a mobile robot equipped with vision sensors (e.g., an RGB-D camera). This task is often referred to as Active Visual Search (AVS) [1] and is considered to be an essential component for successful social, service, or rescue robots. A robot's ability to search for a person or an object in a cluttered environment can be beneficial for many applications, including searching for a cellphone in an apartment for the owner, conducting a rescue mission to locate an injured person, removing hazardous materials at the scene of an accident, or even searching for evidence at the scene of a crime.

Many of the existing work [2]–[4] on AVS rely on the existence of a predefined set of object categories. However, this reliance places severe restrictions on AVS's applicability outside of the lab. Rather, real-world applications would be better served by AVS if it could accept object queries using *free-form text* (e.g., "red wallet") and without dependence on predefined object categories. The task of searching for the target from a free-form text is often referred to as zero-shot object goal navigation [5]–[7].

¹Jeongeun Park, Taerim Yoon, and Sungjoon Choi are with the Department of Artificial Intelligence, Korea University, Seoul, Korea {baro0906, taerimyo, sungjoon-choi}@korea.ac.kr.

²Jejoon Hong is with the Department of Mechanical Engineering, Korea University, Seoul, Korea hdavid0510@korea.ac.kr.

³Youngjae Yu is with the Department of Artificial Intelligence, Yonsei University, Seoul, Korea yjy@yonsei.ac.kr.

⁴Matthew Pan is with the Department of Electrical and Computer Engineering, Queens University, Kingston, Canada matthew.pan@queensu.ca.

Similar tasks have been handled on the ALFRED challenge, which is also based on free-form text-based commands. The ALFRED (Action Learning From Realistic Environments and Directives) challenge [8] is a set of robotics benchmarks that aims to solve a given task with a high-level language goal (e.g., "rinse off a mug and place it in the coffee maker") and low-level instructions (e.g., "walk to the coffee maker on the right" and "wash the mug in the sink"), containing both navigation and object manipulation tasks. Then again, low-level instructions are costly as they rely on step-by-step human planning. In this work, we focus on an active search of the target given as free-form text without any detailed instructions.

We propose an object search method that uses free-form text target commands, which we call Zero-shot Active Visual Search (ZAVIS). ZAVIS can detect novel objects and plan search patterns to allow a robot to efficiently search for the target object based on a semantic grid map represented by static landmarks (e.g., desk or bed). ZAVIS utilizes an open-set object detector to obtain the bounding box of unseen objects; this collection forms the set of candidate objects that are compared to the target object provided by a user in language form by using pre-trained vision-language models (i.e., CLIP [9]). ZAVIS also includes a robot planning approach for object search: an efficient landmark-based planner equipped with a commonsense knowledge model to leverage practical prior knowledge [2].

II. RELATED WORK

Active visual search (AVS) consists of a number of components including perception [10], [11] and planning [12], [13]. The perception module (e.g., the detection or segmentation module) enables robots to detect and locate objects which can be either target objects or landmark objects. The planning module plans and obtains a high-level trajectory or waypoints that effectively search for the target object. In this work, we focus on both components in an open-set setting, where the label of the target object may not be included in the training datasets. Thus, in this section, we provide existing work on perception and planning algorithms for active search.

Traditionally, AVS has been tackled with a detection module and a planning module that work in a closed-set setting [2], [4], [12], [13], where the label of the target object is included in the training dataset. Correlational Object Search (COS-POMDP) [4] formulates the task as a partially observable Markov decision process (POMDP) with a correlational observation model. Unfortunately, this limits the target objects searchable by a robot. In order to overcome this problem, there has been an attempt to solve the zero-shot object goal navigation [5]–[7]. Gadre et al. [6] leverages

Grad-CAM [14] of CLIP [9] to locate the novel target objects and uses FBE [15] to explore to search for target objects. Our work stands along with these in that we also assume AVS in an open-set setting.

III. PROBLEM FORMULATION

Zero-shot Active Visual Search describes the goal of searching for a target object that a user refers to in a free-form text. We call our task zero-shot since the target object is unseen during training and does not appear in predefined knowledge graphs. For this task, we assume that the location of the robot is accessible (i.e., GPS or wheel odometry), which is a common assumption adopted in other work [3], [4], [6]. The robot is also able to observe the world through RGB-D images and LiDAR sensor data. We further assume all of the objects in the scene are static and the target objects are small enough to co-exist with the larger static landmark objects, following the same hypothesis from previous work [3], [4].

The task is composed of two parts, planning and object detection, which are represented as modules in our system. Within the planning module, we define the target object as o^t and the set of landmark names as \mathcal{O} . Landmark objects can be classified as either known \mathcal{O}_k or unknown \mathcal{O}_u during training while the target object is unknown. A set of N viewpoints of a landmark is denoted using the notation $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^N$ with $\mathbf{p}_i = \{o_i, x_i^v, y_i^v, \theta_i^v, I_i, \mathbf{b}_i\}$ where o_i is a landmark object class, x_i^v, y_i^v, θ_i^v is a viewpoint position and orientation, I_i is a raw input image and \mathbf{b}_i is a bounding box from object detector corresponding to the object class o_i .

IV. PROPOSED METHOD

In this section, we discuss how our ZAVIS system tackles the active visual search problem. We explain the broad approach first, followed by details.

A. Zero-shot Active Visual Search

ZAVIS takes an RGB-D image, LiDAR data, and the robot's location as inputs. In addition, the target object and the set of landmark object names are also given as inputs in the free-form text (presumably given by a user to the robot via an interface). The system's output is a candidate set of patches with a location projected on the map. The overall procedure of the algorithm is shown in Figure 1. The robot scans the landmarks, calculates the order of visiting landmarks, navigates to the viewpoint of the selected landmark or explores, detects the target, and asks for feedback.

The algorithm starts with scanning for landmarks, the robot detects a landmark by rotating the pan-tilt camera and extends the free space based on the LiDAR sensor. Then the location of the landmark is projected onto an internal map by transforming the depth information of the detected landmark. For efficient search, target detection is also conducted concurrently during the scan phase. Second, the waypoint generator calculates the order of the visiting landmark, in which the waypoint is a pose of the robot to view the landmark detected previously or explore. The

module is based on map information, knowledge prior, and semantic detection uncertainty. Given the sequence of waypoints, the robot visits the landmarks and searches for the target object using the proposed object detection module on various viewpoints obtained by the pan-tilt camera. After detecting a set of target candidate objects, the robot asks the human if the target is in the candidate set. If the robot gains positive feedback, the search is complete; otherwise continues. If the target is not found and the robot has visited all of the detected landmarks, the robot explores by moving to the nearest unexplored space, called the frontier. When the robot reaches the frontier, the robot reverts to the scanning phase and reiterates the procedure mentioned above.

B. Object Detection

In order to search for objects given as free-form text, we need to detect the objects whose label is not in the training dataset. We divide the problem into two stages: open-set object detection and text-image matching.

In the first stage, an open-set object detection approach is used with the aim of detecting both known and unknown labels trained only with a fixed set of labels. In this stage, the novel class is detected as a predefined 'unknown' class. We follow the baseline model used in FasterRCNN [11] with one additional class $K + 1$ called 'unknown'. Yet, we need to address two issues in order to detect novel objects: localizing the objects that are unknown and classifying these novel objects as members of the 'unknown' class. For identifying the unseen object, we utilize the pretrained CLIP [9] model to make a pseudo annotation for the potential unseen object. For classifying the novel object as 'unknown', we account for the uncertainty obtained by a mixture of logit networks.

In order to localize the unknown object, we make pseudo annotations of the unknown class based on the RoI (Region of Interest) prediction of a Region Proposal Network (RPN) [11]. We follow the prompt engineering setting¹ of the CLIP model [9] to measure the general objectness [10] of the RoIs. Next, we measure the CLIP objectness score by computing the distances between texts that correspond to the background (first three prompts from \mathcal{S}), then select the top- k patches with higher CLIP objectness scores than a certain threshold, 0.9. The CLIP objectness score of the cropped patch B is defined as follows:

$$o(B) = 1 - \frac{\sum_{s_i \in \mathcal{S}} f(B) \cdot g(s_i)}{\sum_{s_j \in \mathcal{S}} f(B) \cdot g(s_j)} \quad (1)$$

where $b = 3$ is a number of prompts corresponding to a background, f is an image encoder and g is a text encoder in CLIP [9]. We linearly increase the number of pseudo annotations for every iteration during training, which is denoted as k . In particular, we set $k = \text{int}(3 * (\text{iter}/\text{max iter})) + 1$. The pseudo annotation trains the whole network (i.e., RoI heads and RPN) for learning the unknown objects.

¹The prompt set $\mathcal{S} = \{"a photo of "o | o \in \{"background", "road scene", "house scene", "animal", "fashion accessory", "transport", "traffic sign", "home appliance", "food", "sport equipment", "furniture", "office supplies", "electronics", "kitchenware"\}\}$

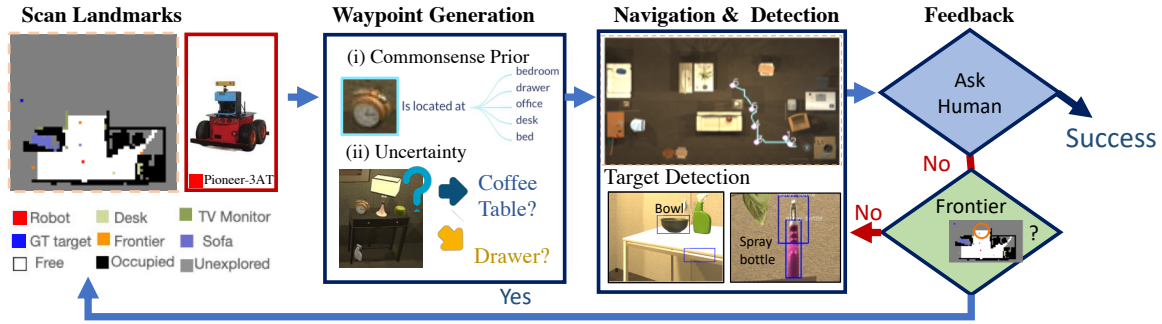


Fig. 1. The overall procedure of ZAVIS. We present ZAVIS framework, object search based on free-form text. The framework consists of scanning, waypoint generation, navigation, and detection. The framework iteratively asks the human to compare the candidate patch set.




Fig. 2. Proposed Open-set Object Detector. Our model classifies the unknown object by an uncertainty-aware classification head, designed with a mixture of logit networks. Furthermore, we utilize the pre-trained CLIP model to generate the pseudo annotation for the potential unknown objects.

We use an uncertainty measure to categorize the novel object to 'unknown' class. Although we have an additional class $K + 1$ for unknown labels, we found it was not sufficient to capture all of the unknown objects. Rather, we utilize the mixture of logit network (MLN) [16] on the RoI classification head to measure the uncertainty. Due to the noisy pseudo annotation mentioned above, MLN is suitable as it can robustly learn the target.

The output of the MLN head is composed of mixture logit μ and mixture weight π . Then, the epistemic uncertainty is obtained by the disagreement of mixtures, and the aleatoric uncertainty is calculated by the weighted summation of the entropy of the logits. Both epistemic uncertainty and aleatoric uncertainty are calculated as follows:

$$\sigma_e^2 = \sum_{j=1}^J \left(\sum_{c=1}^{K+2} \pi_j \left\| \mu_j^{(c)} - \sum_{m=1}^J \pi_m \mu_m^{(c)} \right\|^2 \right) \quad (2)$$

$$\sigma_a^2 = \sum_{j=1}^J \pi_j \cdot \sum_{c=1}^{K+2} [-\mu_j^{(c)} \log \mu_j^{(c)}] \quad (3)$$

where J is the number of mixtures, and $K+2$ is total number of the class labels including background and unknown.

The loss function of the RoI classification head is as follows.

$$\mathcal{L}_{cls}(\mu, \pi, c^*) = \sum_{j=1}^J \pi_j l(\mu_j, c^*) \quad (4)$$

where l is a cross-entropy loss, c is a class label, \mathbf{b} is a corresponding bounding box, and the ground-truth label is

$\{c^*, \mathbf{b}^*\}$. After training, we utilize the total uncertainty, i.e., the sum of epistemic and aleatoric uncertainty, to detect the unseen classes. We model the Gaussian distribution of total uncertainty on the training set; if the CDF of estimated total uncertainty is bigger than 0.9, we change the label of the patch to the 'unknown' class.

The second stage in our object detection system is text-image matching. As CLIP embeds text and image features in the same space, we can use the distance between the embedding of text and image as a metric. We first make text embedding $g(o)$ with target object o , with the prompt mentioned in CLIP "A photo of [object]". With open-set detection result $\{c_i, \mathbf{b}_i\}_{i=1}^n$, we crop the image based on a bounding box corresponding to an unknown class. We then measure the dot product of patches, the similarity between the image and text, to calculate the matching score. The equation for the matching score is as follows:

$$m(o, I, \mathbf{b}_u) = f(C(I, \mathbf{b}_u)) \cdot g(o) \quad (5)$$

where f is the image encoder, g is the text encoder, C is a crop function, and \mathbf{b}_u is an unknown class bounding box.

The detection module is used in both landmark detection and target object detection. In the case where detected landmarks belong to a 'known class', we directly use the corresponding class c_i . If this is not the case, we crop the bounding box of 'unknown' class and identify its class by conducting text-image matching with \mathcal{O}_u . In the case of target detection, as the target object assumes to be unseen during training, we conduct a text-image matching with bounding boxes whose label is 'unknown'. If the matching score is bigger than a certain threshold m_t , we define the patch matches with the text.

C. Waypoint Generator

The robot searches for a target object by sequentially visiting waypoints which are generated based on the current map information and prior knowledge, (i.e., co-occurrence between landmark and target object). In order to obtain the waypoints, the viewpoints \mathcal{P} for each landmark are first generated. The viewpoint is the position and orientation to view the landmark. Each viewpoint is calculated based on the grid map M , based on world map coordinate frames, with a predefined set of orientations. We utilize the heuristic that

viewpoints must be located far enough to view the object and must be in the robot’s reachable space.

The waypoints are generated with a greedy algorithm based on the cost function. The cost function between the currently selected viewpoint \mathbf{p}_c and the next viewpoint candidate \mathbf{p}_n , is defined as co-occurrence and distance, which is as follows.

$$c(\mathbf{p}_c, \mathbf{p}_n) = \sqrt{(x_c - x_n)^2 + (y_c - y_n)^2} + \lambda_1(1 + 1^{-3} - p(o^t|o_n)) + \lambda_2\sigma_c(I_n, \mathbf{b}_n) \quad (6)$$

where the $p(o^t|o_n)$ is co-occurrence measure of the landmark, λ_1 is a weight of a co-occurrence and λ_2 is a weight of a semantic uncertainty $\sigma_c(I_n, \mathbf{b}_n)$. Measurement of the co-occurrence and the semantic uncertainty is discussed in detail in section IV-D.

By the greedy algorithm, the candidate set of \mathbf{p}_n is initialized as a set of viewpoints during scanning (\mathcal{P}), and the next viewpoint is selected, which has the minimum cost. If the co-occurrence of the viewpoint is less than threshold t_c or the uncertainty is higher than t_u , the corresponding viewpoint is skipped. The viewpoint is removed from the candidate set after it is selected.

D. Knowledge Prior & Uncertainty

For efficient search, we also leverage prior knowledge, i.e., co-occurrence between query and landmark objects and the uncertainty of detected objects. Ambiguity in a language is known to be inherent, leading to a semantic uncertainty of image-to-text matching. This ambiguity of text-image matching could lead to false text labels. For this reason, we add semantic uncertainty of unseen landmark objects to the cost function of the searching algorithm shown in Eq 6.

For measuring the co-occurrence ($p(o^t|o_j, o_j \in \mathcal{O}_u)$) between the landmark object and target object, we utilize the pre-trained commonsense knowledge model. Specifically, we use the commonsense transformer (COMET) [17] BART trained on the ATMOIC 2020 dataset [18] to generate the possible locations where the target object would be located, (e.g., [remote control] [located at] [TV]). We made an input of COMET s_o as "[Target name] [AtLocation] [GEN]", distilling up to 20 words generations from the COMET $\{G(s_o)_i\}_{i=1}^{20}$. Then, we estimate the maximum word similarity over the generated words and the name of landmarks, which can be leveraged as a knowledge-based co-occurrence score between the target and the landmark.

$$p(o^t|o) = \max_i \ell(w(o^t), w(G(s_o)_i))_{i=1}^{20} \quad (7)$$

where ℓ denotes cosine similarity, and w is for the word to vector embedding [19].

To measure semantic uncertainty ($\sigma_c p_{next}$), we use the output entropy of a text-image matching model (i.e., CLIP). We define the probability that a certain patch of an image would belong to a specific landmark name o_i as follows:

$$p(o_i, I, \mathbf{b}) = \frac{\exp(m(o_i, I, \mathbf{b}))}{\sum_{j=1}^{L'} \exp(m(o_j, I, \mathbf{b}))} \quad (8)$$

TABLE I
OPENSET DETECTION RESULT

Method	WI	U-Recall	U- Precision	K-mAP
Faster RCNN [11]	-	-	-	56.35
ORE [23]	0.04893	9.390	5.248	55.83
ZAVIS (Ours)	0.03999	19.06	2.279	53.67




Fig. 3. Predictions of proposed open-set object detector and text-image matching result. The first column is the result of open-set detection and the rest results from text-image matching. Note that book, red book, trash can, Redbull, alarm clock, glasses, cup, ipod is unseen during training and given as a language form in the text-image matching phase.

where $o_i \in \mathcal{O}_u$, and m is defined on equation 5. Then we calculate the entropy of the probability, defining it as a semantic uncertainty.

$$\sigma_c = - \sum_{i=1}^{L'} p(o_i, I, \mathbf{b}) \log p(o_i, I, \mathbf{b}) \quad (9)$$

V. EXPERIMENTS

In this section, we discuss the experimental results of the proposed method. We first validate the proposed open-set object detector on the vision dataset [20], [21]. We then quantitatively evaluate the proposed method in the simulation environment, AI2Thor simulation [22]. Finally, we qualitatively demonstrate the proposed method can be applicable to real-world scenarios and analyze the results.

A. OpenSet Object Detection

We validate the proposed open-set object detector with respect to wilderness impact (WI) [24], recall of unknown class (U-Recall), the precision of the unknown class (U-Precision), and mean average precision of the known class

TABLE II
HYPERPARAMETERS & OBJECT SETTINGS

	RoboThor	IThor	RealWorld
λ_1	1	3	20
λ_2	0.05	0.05	10
t_c	0.2	0.2	0.2
t_u	2.5	2.5	1
m_t	29	29	29
target	RemoteControl, Laptop, Book, Apple, CD, Pot, Bowl, AlarmClock, TeddyBear, CellPhone, SprayBottle, Pillow		Book, Cup, Laptop, Cellphone, Tumbler
Known Landmark	Tv monitor, Sofa, Diningtable		Tv monitor, Diningtable
Unknown Landmark	Armchair, Side table, Coffee table', Desk, Bed, Drawer		Table, Side table, Coffee Table, Desk

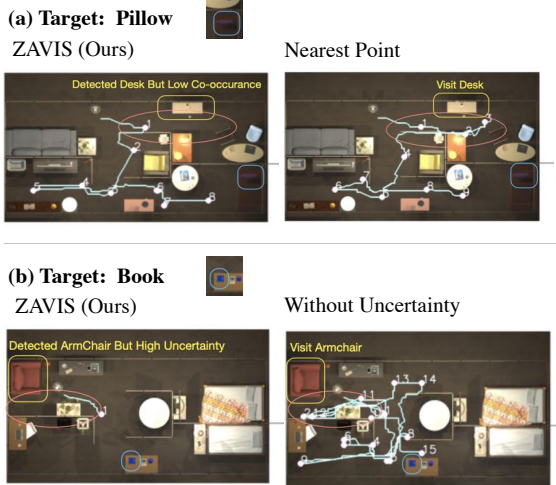


Fig. 4. Examples of trajectory generated in ablation studies on AI2THOR. (Top) compares the path between studies without the co-occurrence module. Bottom compares the path between other methods without considering the uncertainty.

TABLE III
COMPARISON WITH PREVIOUS WORK

scene	method	SPL	SR (%)
RoboThor	COS-POMDP [4] †	0.0693	14.38
	CoW [6]	0.1154	51.91
	ZAVIS (Ours)	0.3462	91.80
IThor-Bedroom	COS-POMDP [4] †	0.2334	49.21
	CoW [6]	0.093	41.09
	ZAVIS (Ours)	0.3017	81.42
IThor-Livingroom	COS-POMDP [4] †	0.1410	29.16
	CoW [6]	0.00088	14.70
	ZAVIS (Ours)	0.1969	75.00

(K-mAP). WI measures how much the model is confused with unknown and known labels, and U-Recall measures how much the detectors capture the unknown objects with respect to unknown objects in the test set. All of the detectors are trained on the Pascal VOC 2012 dataset [20] and the subset MS COCO dataset [21] that has the same class label as the Pascal VOC dataset. The result is shown in Table I, and the qualitative results are in Figure 3.

As the baseline Faster RCNN [11] is not designed for open-set detection, we can not measure any metric related to the open-set setting. Compared with an open-set detector ORE [23], our proposed detector has a lower WI 0.00894 and higher U-recall 9.67. Although the U-Precision of our method is lower, we observe that the model with higher recall better fits the object-searching problem. Because the algorithm is designed to ask for humans after target detection, the success depends on the ability to capture the target object regarding false prediction.

B. Simulational Result

We conducted experiments of ZAVIS in simulated 3D scenes using the IThor and RoboThor [25] environments within the open-source AI2Thor framework, a tool for visual AI research [22]. Using the IThor environment, we validate the proposed method in multiple bedrooms and living room

TABLE IV
ABLATION STUDIES

Ablations		SPL	SR (%)	#
Detectors	FasterRCNN [11]	0.2600	75.95	9.48
	ORE [23]	0.2835	85.79	7.71
Co-occurrence	Nearest Point	0.3293	89.61	7.05
	Web-based [3]	0.3210	89.00	7.36
	Word Similarity [19]	0.3374	90.71	6.56
Uncertainty	w/o Uncertainty	0.3292	87.91	8.66
	ZAVIS (Ours)	0.3462	91.80	6.65

scene types, each having 30 scenes. Additionally, we also use a total of 15 scenes from the RoboThor environment to help validate ZAVIS. In our experiments, however, we mainly use RoboThor as the environment has diverse multi-room scenes, whereas IThor is limited to containing single or double rooms. We use an RRT planner [26] for the navigation, and the hyperparameters and object settings are described in Table II.

As the proposed method asks humans whether the target is in the image, success is defined as the condition where the candidate set contains the ground-truth target object in the simulated environment. Particularly, We denote the target is detected if the IoU (intersection over union) between ground-truth objects bounding box and the candidate bounding box is bigger than 0.3 or IoA (intersection over area) bigger than 0.5. We define a failing condition to be when no patch corresponds to the ground truth image while the robot is moving more than 50m.

We evaluate the proposed method with two commonly used evaluation metrics for active visual search, SPL (Success weighted by Path Length) [27] and SR (success rate). We benchmark two methods, COS-POMDP [4] and CoW [6]. As COS-POMDP [4] does not fit for zero-shot setting on both the detector and correlation model, we changed the detector to Faster RCNN with the CLIP model and web-based correlation model [4]. We evaluate with the same success criterion with human feedback to all of the compared methods.

The experimental results on the simulation environment are shown in Table III. The method with † is adjusted for the zero-shot setting. The proposed method outperforms those of all previous work by a significant margin with an average gap of 0.173 in SPL and 49.33% in SR. Thus, we observe that the proposed detection method can locate unseen objects better than the Grad-CAM-based method based on success rate. In addition, while CoW [6] directly searches for the objects based on FBE [15] without any landmarks, our proposed ZAVIS method searches for both target objects and co-occurrences with landmark objects for efficient search. We have also observed that the proposed method outperforms algorithms designed for non-zero-shot AVS with the adjustment to the zero-shot setting.

We further conduct ablation studies to show the effect of the proposed detector, co-occurrence, and uncertainty. We experimented on various object detectors, FasterRCNN [11] [9] and ORE [23]. In addition, we evaluate the various co-occurrence measure, web-based knowledge prior [3], word similarity [19], and without any co-occurrence measure. The results are summarized in Table IV,

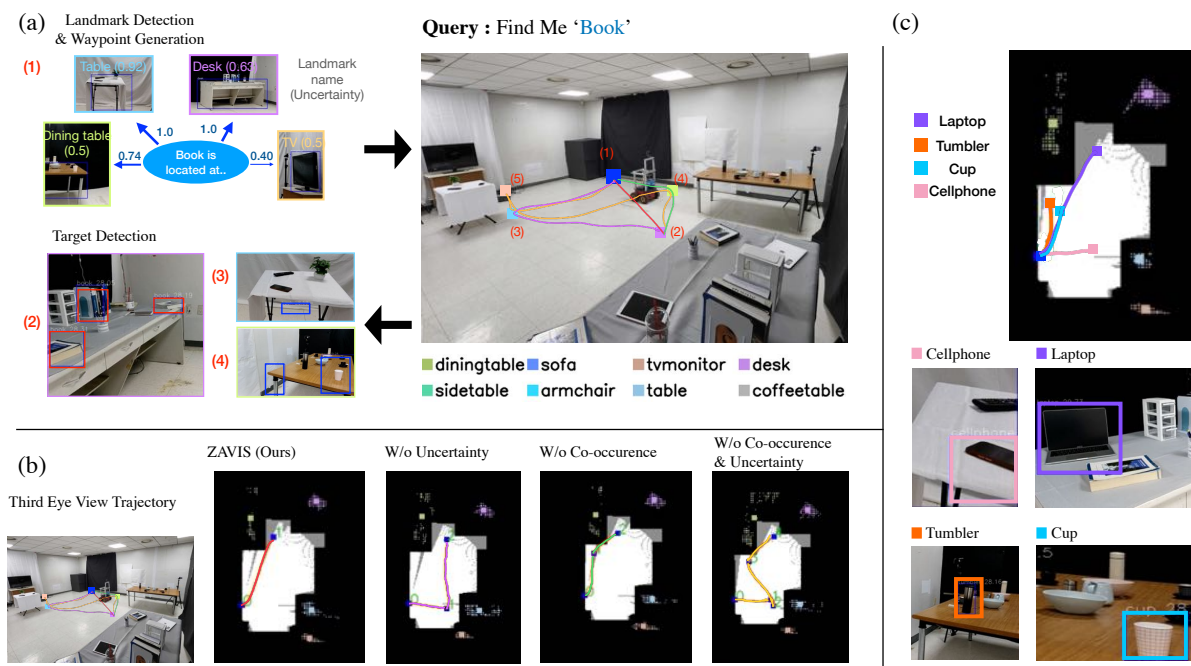


Fig. 5. Real-world Demonstration. (a) Explains the demonstration on the target `book`. (b) shows the trajectory for four scenarios, the proposed method, without semantic uncertainty, without co-occurrence, and without both. (c) shows a various trajectory and detection result with target `book`, `notebook`, `tumbler`, `cup`, `cellphone`

where # denotes the average number of waypoints. Figure 4 shows an example of the trajectory. The robot with the target as `pillow` skips to visit the `desk` as it has low co-occurrence, while the method that does not utilizes co-occurrence visits the `desk` as shown in (a). In addition, in the case of the `pillow` target (b), the robot does not visit the `armchair` due to high uncertainty, inducing an efficient trajectory.

C. Real-world Demonstration

In this section, we discuss the qualitative results of a real-world pilot experiment where ZAVIS was implemented on a robot. We used the Pioneer-3AT robot with a pan-tilt stereo camera (OAKD-Pro) and LiDAR sensor (Sick LMS100). A DWA planner (dynamic window approach) [28] is implemented for the local planning phase to consider the motion dynamics and the localization is based on Adaptive Monte Carlo Localization algorithm with a prebuilt map. The target and landmark objects are shown in Table II.

Figure 5-(a) shows the demonstration in which the robot is searching for the `book`. The robot plans to visit the `desk` first because the `desk` has a high co-occurrence score and low semantic uncertainty. Although the `table` has high co-occurrence and is closer, the robot decides to visit the `table` after the `desk` due to high uncertainty. After the decision, the robot moves to the landmark `desk` and detects the `book`, and asks the human if the `book` exists in the detected patch set.

We then show the effect of the proposed co-occurrence measure and semantic uncertainty measure with the target `book` shown in Figure 5-(b), the estimated uncertainty and co-occurrence are illustrated in Figure 5-(a). Without the

uncertainty measure, the robot visits closer landmarks with high co-occurrence `table` (1.0), then visits the `desk`, because the `table` is closer to the `desk` and did not adopt the semantic uncertainty. Without a co-occurrence measure, the robot plans to visit `dining table` first as it has low uncertainty and is close to its current position. Without both of these measures, the robot only plans based on the distance, leading to the longest trajectory. In addition, we demonstrate the proposed method on various target objects shown in Figure 5-(c).

VI. CONCLUSION

In this paper, we have introduced a system that enables a robot to actively move around the environment to search and locate the target object that is unknown to the robot and proposed Zero-shot Active Visual Search (ZAVIS). Firstly, we have introduced a novel open-set object detector that can detect unseen objects during training and expand it for zero-shot detection with a text-image matching model. Secondly, we utilized the commonsense-based language model to gain prior knowledge regarding co-occurrences between landmark and target objects. Lastly, we have shown that considering the uncertainty from the text-image matching model helps efficient search. As part of this study, we have validated the effectiveness and utility of ZAVIS during experiments performed in both simulated and real-world environments.

ACKNOWLEDGMENT

This work was supported by Samsung Electronics (IO201230-08278-01), and the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2.220451.01).

REFERENCES

- [1] Yiming Ye and John K Tsotsos. Sensor planning for 3d object search. *Computer Vision and Image Understanding*, 73(2):145–168, 1999.
- [2] Ying Zhang, Guohui Tian, Jiaying Lu, Mengyang Zhang, and Senyan Zhang. Efficient dynamic object search in home environment by mobile robot: A priori knowledge-based approach. *IEEE Transactions on Vehicular Technology*, 68(10):9466–9477, 2019.
- [3] Zhen Zeng, Adrian Röfer, and Odest Chadwicke Jenkins. Semantic linking maps for active visual object search. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1984–1990. IEEE, 2020.
- [4] Kaiyu Zheng, Rohan Chitnis, Yoonchang Sung, George Konidaris, and Stefanie Tellex. Towards optimal correlational object search. In *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2022.
- [5] Qianfan Zhao, Lu Zhang, Bin He, Hong Qiao, and Zhiyong Liu. Zero-shot object goal visual navigation. *arXiv preprint arXiv:2206.07423*, 2022.
- [6] Samir Yitzhak Gadre, Mitchell Wortsman, Gabriel Ilharco, Ludwig Schmidt, and Shuran Song. Clip on wheels: Zero-shot object navigation as object localization and exploration. *arXiv preprint arXiv:2203.10421*, 2022.
- [7] Apoorv Khandelwal, Luca Weihs, Roozbeh Mottaghi, and Aniruddha Kembhavi. Simple but effective: Clip embeddings for embodied ai. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14829–14838, 2022.
- [8] Mohit Shridhar, Jesse Thomason, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. Alfred: A benchmark for interpreting grounded instructions for everyday tasks. In *Proc. of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 10740–10749, 2020.
- [9] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proc. of the International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- [10] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [11] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [12] Yuchen Xiao, Sammie Katt, Andreas ten Pas, Shengjian Chen, and Christopher Amato. Online planning for target object search in clutter under partial observability. In *Proc. of the International Conference on Robotics and Automation (ICRA)*, pages 8241–8247. IEEE, 2019.
- [13] Yiming Ye and John K Tsotsos. Sensor planning in 3d object search: its formulation and complexity. In *Proc. of the International Symposium on Artificial Intelligence and Mathematics, Florida, USA*, 1996.
- [14] Ramprasaath R Selvaraju, Abhishek Das, Ramakrishna Vedantam, Michael Cogswell, Devi Parikh, and Dhruv Batra. Grad-cam: Why did you say that? In *Proc. of the Conference on Neural Information Processing Systems (NIPS) Workshop*, 2016.
- [15] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proc. of the second international conference on Autonomous agents*, pages 47–53, 1998.
- [16] Jeongeun Park, Seungyoun Shin, Sangheum Hwang, and Sungjoon Choi. Elucidating robust learning with uncertainty-aware corruption pattern estimation. *Pattern Recognition*, 138:109387, 2023.
- [17] Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. Comet: Commonsense transformers for automatic knowledge graph construction. In *Proc. of the Association for Computational Linguistics (ACL)*, 2019.
- [18] Jena D Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, and Yejin Choi. (comet-) atomic 2020: on symbolic and neural commonsense knowledge graphs. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 35, pages 6384–6392, 2021.
- [19] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- [20] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proc. of the European conference on computer vision (ECCV)*, pages 740–755. Springer, 2014.
- [22] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017.
- [23] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proc. of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5830–5840, 2021.
- [24] Akshay Dhamija, Manuel Gunther, Jonathan Ventura, and Terrance Boulton. The overlooked elephant of object detection: Open set. In *Proc. of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1021–1030, 2020.
- [25] Matt Deitke, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Dustin Schwenk, Eli VanderBilt, Matthew Wallingford, et al. Robothor: An open simulation-to-real embodied ai platform. In *Proc. of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3164–3174, 2020.
- [26] Steven M LaValle and James J Kuffner. Rapidly-exploring random trees: Progress and prospects: Steven m. lavalle, iowa state university, a james j. kuffner, jr., university of tokyo, tokyo, japan. *Algorithmic and Computational Robotics*, pages 303–307, 2001.
- [27] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [28] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 4(1):23–33, 1997.