

Embodied Agents for Efficient Exploration and Smart Scene Description

Roberto Bigazzi, Marcella Cornia, Silvia Cascianelli, Lorenzo Baraldi, Rita Cucchiara

Abstract—The development of embodied agents that can communicate with humans in natural language has gained increasing interest over the last years, as it facilitates the diffusion of robotic platforms in human-populated environments. As a step towards this objective, in this work, we tackle a setting for visual navigation in which an autonomous agent needs to explore and map an unseen indoor environment while portraying interesting scenes with natural language descriptions. To this end, we propose and evaluate an approach that combines recent advances in visual robotic exploration and image captioning on images generated through agent-environment interaction. Our approach can generate smart scene descriptions that maximize semantic knowledge of the environment and avoid repetitions. Further, such descriptions offer user-understandable insights into the robot’s representation of the environment by highlighting the prominent objects and the correlation between them as encountered during the exploration. To quantitatively assess the performance of the proposed approach, we also devise a specific score that takes into account both exploration and description skills. The experiments carried out on both photorealistic simulated environments and real-world ones demonstrate that our approach can effectively describe the robot’s point of view during exploration, improving the human-friendly interpretability of its observations.

I. INTRODUCTION

Over the past few years, advances on visual navigation and machine learning have shown that training navigation policies with reinforcement learning in simulated photorealistic environments [1]–[4] for millions or even billions of frames [5] allows agents to generalize their navigation skills to unseen environments without the need of acquiring any previous knowledge. A plethora of exploration rewards and strategies has been presented in the representation learning community and, generally, those strategies can be adapted to work on board of smart autonomous agents and improve their perception of the surrounding, and consequently, the desired behavior. A popular approach for such agents follows a modular, hierarchical navigation policy design such as the one proposed by Chaplot *et al.* [6], which optimizes a reward function that encourages exploration [7]–[10]. Embodied agents trained on sufficiently photorealistic simulators are then deployable onto real-world robotic platforms [11]–[13], thus contributing to the widespread of service robots.

* This work has been supported by the “Fit for Medical Robotics” (Fit4MedRob) project, funded by the Italian Ministry of University and Research, by “FAR Dipartimentale 2021-2022” of the Department of Engineering “Enzo Ferrari” of the University of Modena and Reggio Emilia, and by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 955778 for project “Personalized Robotics as Service Oriented Applications” (PERSEO).

R. Bigazzi, M. Cornia, S. Cascianelli, L. Baraldi, R. Cucchiara are with the University of Modena and Reggio Emilia, Italy {firstname.lastname}@unimore.it

At the same time, research connecting robotics, vision, and natural language processing has attracted interest and has led to the emergence of robotic vision and language tasks, such as vision-and-language navigation [14]–[17] and embodied question answering [18]–[21]. However, using language to facilitate the human understanding of the behavior and perception of a robotic agent is still an under-explored path [22]–[25], in which further research steps are needed to reach a seamless interaction between humans and robots. In this respect, image captioning approaches, whose goal is to generate a natural language description of a given image, have seldom been employed in navigation and exploration settings. In such approaches, images can be represented by using convolutional neural networks to extract global features [26], grids of features [27], or features for image regions containing visual entities [28]. In most cases, attention mechanisms are applied to enhance the visual input representation. More recent approaches employ fully-attentive Transformer-like architectures [29] as visual encoders, which can also be applied directly to image patches [30], [31]. The image representation is then used to condition a recurrent neural network [26], [32], [33] or a Transformer-based [34], [35] language model that generates the caption.

The capability of image captioning approaches to produce natural descriptions of everyday images is, however, insufficient when working in an embodied setting in which the agent continuously moves inside of an environment and needs to portray interesting scenes. In such a scenario, indeed, not all the images observed by the agent are semantically relevant as they might contain uninteresting objects or scenes (*e.g.* when facing a wall) or there might be low variety between temporally consecutive observations. Hence, a naive captioning of such images would result in uninteresting descriptions and avoidable repetitions. To overcome this, the captions should be generated only when the scene observed is worth being described to a human.

In this work, we take a step forward with respect to the above-mentioned limitations and propose a complete pipeline for efficient exploration and mapping which can, at the same time, provide user-understandable representations of the perceived environment in the form of natural language descriptions. We jointly integrate and propose visual exploration strategies, a state-of-the-art approach for image captioning and smart description policies in an embodied setting, with the final aim of improving human understanding of robotic perception. Also, we devise a novel metric, called episode description score (ED-S), that evaluates the exploration and the ability of covering objects in the environment avoiding repetitions. We extensively test the performance of

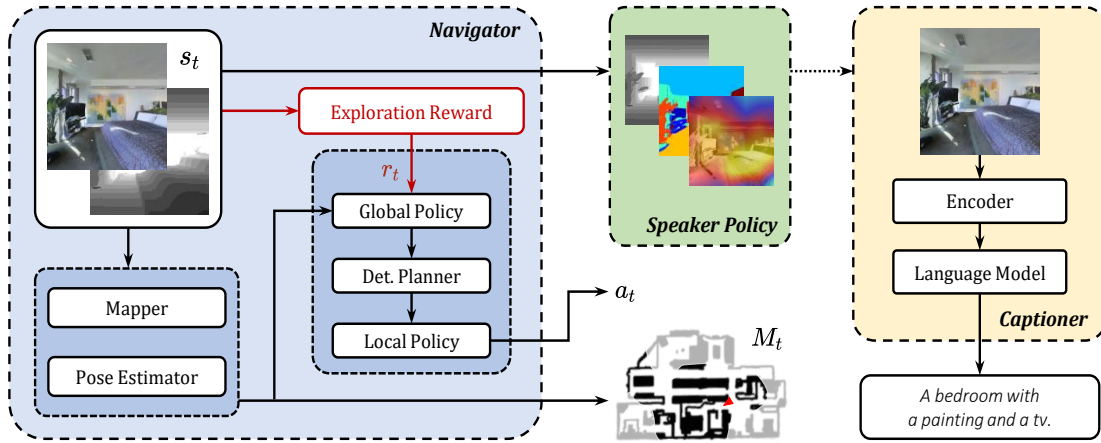


Fig. 1: Overview of the approach for smart scene description, comprising a navigator, a speaker policy, and a captioner.

the proposed approach in comparison with different baselines on both Gibson [1] and Matterport3D [2] datasets. Finally, while our approach is trained and evaluated in simulation, the proposed architecture is designed for the final deployment on a real robotic platform, as we show in the video accompanying the submission.

II. PROPOSED APPROACH

Our proposed architecture is composed of three main components: a *navigator*, in charge of the exploration, a *captioner*, that describes interesting scenes, and the *speaker policy* that decides when the captioner should be activated. An overview of our complete architecture is shown in Fig. 1.

A. Navigator

The exploration capabilities of the agent are strictly dependant on the performance of the navigation module, therefore relying on a proper navigation approach is of fundamental importance. Following recent literature on embodied visual navigation [6], [7], [36], we devise a hierarchical policy coupled with a learned neural occupancy mapper and a pose estimator. The hierarchical policy sets long and short-term navigation goals, while the neural mapper builds an occupancy grid map representation of the environment and the pose estimator locates the agent on such map.

Mapper. In order to track explored and unexplored regions of the environment over time, the mapper is a fundamental component. Moreover, a neural-based mapper allows to infer regions occupancy beyond the observable area in front of the agent, facilitating the planning phase [7]. The output of the mapper is a $M \times M \times 2$ global map of the environment m_t that keeps track of the non-traversable space in its first channel and the area explored by the agent in the second one. At each time step, the mapper processes the RGB-D observation $s_t = (s_t^{rgb}, s_t^d)$ coming from the agent and predicts a $L \times L \times 2$ egocentric local map l_t representing the state in front of the agent. RGB images are encoded using a ResNet-18 [37] followed by a UNet encoder [38], while depth observations are encoded using only a UNet encoder. RGB and depth features are concatenated and fed to a CNN to merge the two modalities. Finally, a UNet decoder processes the merged

features to predict the local map l_t . At every timestep t , the local map l_t is transformed using the estimated pose of the agent ω_t and registered to the global map m_t with a moving average. The global map is initially empty and is built incrementally with the exploration of the environment.

Pose Estimator. Relying on a global map requires a robust pose estimator in order to build geometrically coherent and precise maps. Indeed, an inaccurate pose estimate would rapidly diverge from the ground-truth pose, and loop closure is inapplicable if previous knowledge of the environment is not available. Furthermore, directly using the pose sensor of the robot is not sufficient since sensor noise, slipping wheels, and collisions with obstacles would not be accounted for. The adopted approach uses the difference between consecutive pose sensor readings $\Delta\omega'_t = \omega'_t - \omega'_{t-1}$ as a first estimate of the motion of the agent, where $\omega'_t = (x'_t, y'_t, \theta'_t)$, with (x'_t, y'_t) being the coordinates on the map, and θ'_t the orientation of the agent. In order to correct eventual inaccuracies, we use local maps l_t, l_{t-1} extracted from the respective observations as feedback. The local map l_{t-1} is rototranslated with respect to the current position of the agent using $\Delta\omega'_t$. Transformed l_{t-1} and l_t are concatenated and fed to a CNN to output a corrected displacement $\Delta\omega_t$. At every timestep, $\Delta\omega_t$ is used to compute the pose of the agent with respect to the pose at the previous step:

$$\omega_t = \omega_{t-1} + \Delta\omega_t \quad \text{where} \quad \omega_t = (x_t, y_t, \theta_t). \quad (1)$$

Without loss of generality, we consider the agent starting from $\omega_0 = (0, 0, 0)$, *i.e.* the center of the global map m_t .

Navigation Policy. The navigation policy adopts a hierarchical structure as used in [6], [7], [36]. Specifically, the navigation policy comprehends three modules: a high-level global policy, a deterministic planner, and an atomic local policy. The hierarchical policy is adopted to decouple high-level and low-level concepts like moving across rooms and avoiding obstacles. It samples a goal coordinate on the map, while the deterministic planner uses the global goal to compute a local goal in close proximity of the agent. The local policy then predicts actions to reach the local goal.

The global policy takes as input an enriched version of




Fig. 2: Qualitative exploration trajectories of different navigation agents on the same episode.

the current global map m_t . A map encoding of the current position of the agent and a map with already visited locations are concatenated to m_t to form a $M \times M \times 4$ map m_t^+ . The enriched map m_t^+ is both max-pooled and cropped with respect to the position of the agent to a dimension $G \times G \times 4$. Then, the resulting two versions of m_t^+ are stacked together to form the final 8-channel input of the global policy which is processed to sample a point on a $G \times G$ grid. The output of the global policy is converted to a coordinate on the global map m_t , that is the global goal g_t . The global policy is trained with reinforcement learning using the global reward r_t^{global} .

The deterministic planner adopts the A* algorithm to compute a feasible trajectory from the current position of the agent to the global goal using the current state of the map m_t . A point on the trajectory within 1.25m from the agent is extracted to form the local goal p_t .

The local policy takes as input the current RGB observation s_t^{rgb} as well as the relative displacement of the local goal p_t from agent's position ω_t , and predicts the atomic action needed to reach the local goal. The output of the local policy corresponds to one of the following atomic actions: *move forward 0.25m*, *turn left 10°*, and *turn right 10°*. This policy is trained with a reward r_t^{local} that encourages the decrease in the geodesic distance between agent and local goal:

$$r_t^{local} = d(m_{t-1}, p_{t-1}, \omega_{t-1}) - d(m_t, p_t, \omega_t), \quad (2)$$

where $d(\cdot)$ returns the geodesic distance using the position of the agent, the local goal, and the global map to account for possible obstacles on the way.

B. Exploration Rewards

We compare various global exploration rewards such as curiosity [39], coverage [40], anticipation [7], and impact [9]. All the considered methods obtain the reward by exploiting visual input sensors only. Exemplar exploration trajectories resulting from the different rewards are reported in Fig. 2.

Curiosity. The curiosity reward adopts two additional neural networks that learn the environment dynamics. A forward model is trained to predict the encoding of the future RGB observation given the encoding of the current observation and action, and an inverse model is trained to infer the action a_t performed between consecutive observations $(s_t^{rgb}, s_{t+1}^{rgb})$. These models are trained minimizing the following losses: $L_{fwd} = 0.5 \cdot \|\hat{\phi}(s_{t+1}^{rgb}) - \phi(s_{t+1}^{rgb})\|_2^2$ and $L_{inv} = y_t \log \hat{a}_t$, where $\hat{\phi}(s^{rgb})$ and $\phi(s^{rgb})$ denote predicted and ground-truth encodings of the observation s^{rgb} , y is the one-hot encoding of the ground-truth action a , and \hat{a} is the predicted action probability distribution. The global reward for the curiosity-driven model is given by the error of the forward dynamics

model prediction during the navigation:

$$r_t^{global} = \frac{\eta}{2} \left\| \hat{\phi}(s_{t+1}^{rgb}) - \phi(s_{t+1}^{rgb}) \right\|_2^2, \quad (3)$$

where η is a normalizing term set to 0.01.

Coverage. The coverage-based reward maximizes the information gathered at each timestep, being it the number of objects or landmarks reached or area seen. In this work, we consider the area seen definition, as proposed in [6]:

$$r_t^{global} = AS_t - AS_{t-1}, \quad (4)$$

where AS indicates the number of pixels explored in the ground-truth map.

Anticipation. The occupancy anticipation reward [7] aims to maximize accuracy in the prediction of the map including occluded unseen areas, *i.e.*,

$$r_t^{global} = \text{Acc}(m_t, m^*) - \text{Acc}(m_{t-1}, m^*), \quad (5)$$

$$\text{Acc}(m, m^*) = \sum_{i=1}^{M^2} \sum_{j=1}^2 \mathbb{1}[m_{ij} = m_{ij}^*], \quad (6)$$

where m is the predicted global map, m^* is the ground-truth global map, and $\mathbb{1}[\cdot]$ is the indicator function.

Impact. The impact reward encourages actions that modify agent's internal representation of the environment, with impact at timestep t that is measured as the l_2 -norm of the encodings of the two consecutive states $\phi(s_t)$ and $\phi(s_{t+1})$. However, using the formulation of impact as it is, could lead to trajectory cycles with high impact but low exploration. To overcome this issue, Raileanu *et al.* [41] uses the state visitation count $N(s_t)$ to scale the reward. Unfortunately in our setting, the concept of the visitation count is not directly applicable, due to the continuous space of the photo-realistic environment. Hence, we adopt and evaluate the impact-based methods proposed in [9]. Such methods formalize a pseudo-visitation count $\hat{N}(s_t)$ in continuous environments with two different approaches: grid and density model estimation. The final global reward for the impact-driven model becomes:

$$r_t^{global} = \left\| \phi(s_{t+1}^{rgb}) - \phi(s_t^{rgb}) \right\|_2 / \sqrt{\hat{N}(s_t)}, \quad (7)$$

where $\phi(s^{rgb})$ and $\hat{N}(s_t)$ are the encoding and the estimated pseudo-visitation count at timestep t .

C. Captioner

The goal of the captioning module is that of modeling an autoregressive distribution probability $p(\mathbf{w}_t | \mathbf{w}_{\tau < t}, \mathbf{V})$, where \mathbf{V} is an image captured from the agent and $\{\mathbf{w}_t\}_t$ is the sequence of words comprising the generated caption. This is usually achieved by training a language model conditioned




Fig. 3: A sample of agent observation and corresponding images used by the speaker policy to trigger the captioner.

on visual features to mimic ground-truth descriptions. For multimodal fusion, we employ an encoder-decoder Transformer [29] architecture. Each layer of the encoder employs multi-head self-attention (MSA) and feed-forward layers, while each layer of the decoder employs multi-head self- and cross-attention (MSCA) and feed-forward layers. For enabling text generation, sequence-to-sequence attention masks are employed in each self-attention layer of the decoder.

To obtain the set of visual features \mathbf{V} for an image, our model employs a visual encoder that is pre-trained to match vision and language (*i.e.* CLIP [42]). Compared to using features extracted from object detectors [28], [43], our strategy is beneficial in terms of both computational efficiency and feature quality. The visual descriptors $\mathbf{V} = [v_1, v_2, \dots, v_N]$ are encoded via bi-directional attention in the encoder, while the token embeddings of the caption $\mathbf{W} = [w_1, w_2, \dots, w_L]$ are inputs of the decoder, where N and L indicate the number of visual embeddings and caption tokens, respectively. The overall network operates according to the following schema:

$$\begin{aligned} \text{encoder} \quad & \tilde{v}_i = \text{MSA}(v_i, \mathbf{V}), \\ \text{decoder} \quad & \mathbf{O}_{w_i} = \text{MSCA}(w_i, \tilde{\mathbf{V}}, [w_1, w_2, \dots, w_i]), \end{aligned} \quad (8)$$

where \mathbf{O} is the network output, $\text{MSA}(x, \mathbf{Y})$ is a self-attention with x mapped to query and \mathbf{Y} mapped to key-values, $\text{MSCA}(x, \mathbf{Y}, \mathbf{Z})$ indicates a self-attention with x as query and \mathbf{Z} as key-values, followed by cross-attention with x as query and \mathbf{Y} as key-values, and $\tilde{\mathbf{V}} = [\tilde{v}_1, \tilde{v}_2, \dots, \tilde{v}_N]$. We omit feed-forward layers and the dependency between consecutive layers for ease of notation.

D. Speaker Policy

While exploring the environment, the agent sees various RGB observations. Even if the agent was navigating efficiently, the majority of the observations would be overlapped with each other, and the same objects would be observed at multiple consecutive timesteps. Since the agent should describe only relevant scenes during exploration and avoid uninformative captions or unnecessary repetitions, a component that controls caption generation becomes necessary. We thus introduce a speaker policy which is responsible for triggering the captioner depending on the current view. We compare three approaches that exploit different modalities: a *depth-based* policy, an *object-based* policy, and a *visual activation-based* policy. An example of the considered modalities for the same observation is reported in Fig. 3.

Depth-based Policy. High mean depth values indicate a larger area observed by the agent, and potentially, a richer

scene to be described. Instead, when the field of view of the agent is occluded by an obstacle, the mean depth value of the observation s_t^d is typically low. Therefore, the depth-driven policy uses the current depth observation s_t^d and computes its mean value. The captioner is activated if the mean depth value is above a predetermined threshold D .

Object-based Policy. Considering that the description of the scene will concentrate on relevant objects, the object-driven policy uses the number of relevant objects in the RGB observation s_t^{rgb} to decide if the captioner should generate the description. Specifically, the captioner is triggered only if at least a number O of objects are being observed in the scene since using observations with multiple objects allows a larger variety of generated captions.

Visual Activation-based Policy. Another possible strategy to implement the speaker policy entails exploiting the activation maps of the same visual encoder used by the captioner (which is a CLIP-like [42] encoder, as detailed in Sec. II-C). Such a speaker policy is more closely related to the captioning module and provides a means to interpret the image regions that are more relevant to the agent. In particular, in this work, we consider a CNN-based visual encoder and thus take the feature maps from the last convolutional block, projected into a d -dimensional vector. This vector is then averaged, and the speaker policy is activated if its average is above a certain threshold A , thus indicating the presence of sufficient semantic content in the image.

III. EXPERIMENTS

A. Implementation and Training Details

Navigator. All the exploration models are trained for 5M frames on Gibson Dataset [1] environments using Habitat simulator [3]. The evaluation is performed using the test split of Matterport3D (MP3D) dataset [2] and the validation split of Gibson Tiny dataset, because they contain object annotations that are used to evaluate the generated captions.

The RGB-D input to the components of the navigator is resized to 128×128 pixels, and the global map size is $M = 2001$ for MP3D and $M = 961$ and Gibson environments. The size of the local map predicted by the mapper is $L = 101$, and each pixel in the maps describes a 5×5 cm² of the environment. Regarding the global policy, the grid size used for the prediction of the global goal is $G = 240$, and the global goal is sampled every $N_G = 25$ timesteps. Both the global and local policies are trained using PPO algorithm [44] with a learning rate of 2.5×10^{-4} , while the mapper and the pose estimator use a learning rate of 10^{-3} . Episode length is set to 500 and 1000 for the training and evaluation phases, respectively.

Speaker Policy. Since the MP3D dataset has richer object annotations and larger environments than the Gibson dataset, we compare two different sets of threshold values depending on the evaluation dataset for the depth- and object-based policies for triggering the captioner. In particular, the threshold values for MP3D dataset are $D = (1.0, 2.0, 3.0)$ and $O = (1, 3, 5)$ for depth- and object-based policies. Depth and

object thresholds are $\mathbf{D} = (1.0, 1.5, 2.0)$ and $\mathbf{O} = (1, 2, 3)$ for Gibson dataset. On the other hand, for the activation-based criterion, we use the same set of threshold values for both the evaluation datasets, *i.e.* $\mathbf{A} = (4.5, 5.0, 5.5)$.

Captioner. As training and evaluation dataset, we employ COCO [45] following the splits defined in [26]. To improve the generalization abilities of the model, we also train a variant on a combination of 35.7M images taken from both human-collected datasets (*i.e.* COCO [45]) and web-collected sources (*i.e.* SBU [46], CC3M [47], CC12M [48], WIT [49], and a subset of YFCC100M [50]).

We consider three configurations of the captioner, varying the number of decoding layers l , model dimensionality d , and the number of attention heads h : Tiny ($l = 3, d = 384, h = 6$), Small ($l = 6, d = 512, h = 8$), and Base ($l = 12, d = 768, h = 12$). For all models, we employ CLIP-ViT-L/14 [42] as visual feature extractor and three layers in the visual encoder. To assess the effectiveness of CLIP-based features, we also consider a variant of the Tiny model that employs region-based visual features, extracted from Faster R-CNN [28], [51]. We train all captioning variants with cross-entropy loss using LAMB [52] as optimizer. We employ the learning rate scheduling strategy proposed in [29], with a warmup of 6,000 iterations and a batch size equal to 1,080. We additionally fine-tune the models with the SCST strategy [32], by using the Adam optimizer [53], a fixed learning rate equal to 5×10^{-6} , and a batch size of 80.

B. Evaluation Protocol

As the task of smart scene description requires both exploration and description capabilities, for evaluation we use exploration and captioning metrics, as well as a novel score specifically devised for the task.

Navigator. As for the performance of the navigation module, we express them in terms of metrics that are commonly used for evaluating embodied exploration agents. In particular, we consider the intersection over union between the ground-truth map of the environment and the map reconstructed by the agent (IoU), the extent of correctly mapped area (*i.e.* the map accuracy Acc), and the extent of environment area visited by the agent (*i.e.* the area seen AS), both expressed in m^2 .

Captioner. For evaluating the performance of the captioning module on the COCO dataset, we consider the standard image captioning metrics BLEU-4 [54], METEOR [55], ROUGE [56], CIDEr [57], and SPICE [58].

Episode Description Score. Different from standard captioning settings, where ground-truth captions are available for the images, in our setting such information is not available. However, the considered 3D environments datasets come with annotations of the objects in the scene, which can be exploited for performance evaluation. In particular, based on the objects in the scene, we define the **soft-coverage score** (Cov) and the **diversity score** (Div), to evaluate the ability of the agent to mention all the relevant objects in the scene and to produce interesting, non-repetitive descriptions, respectively. The first is computed by considering the intersection

Model	Gibson Val			MP3D Test		
	IoU	Acc	AS	IoU	Acc	AS
Curiosity [8]	0.528	66.19	102.59	0.368	130.34	186.67
Coverage [6]	0.608	73.69	102.66	0.417	146.16	195.03
Anticipation [7]	0.706	81.13	102.22	0.494	157.02	177.14
Impact (Grid) [9]	0.738	82.91	104.16	0.519	164.26	185.13
Impact (DME) [9]	0.694	79.47	105.03	0.496	167.58	205.02

TABLE I: Navigation results on Gibson Val and MP3D Test.

	Train Ims	BLEU-4	METEOR	ROUGE	CIDEr	SPICE
Region-based ^{tiny}	112k	37.7	28.3	57.6	124.8	21.9
CLIP-based ^{tiny}	112k	40.6	30.0	59.9	139.4	23.9
CLIP-based ^{small}	112k	40.9	30.4	60.1	141.5	24.5
CLIP-based ^{base}	112k	41.4	30.2	60.2	142.0	24.0
CLIP-based ^{base}	35.7M	42.9	31.4	61.5	149.6	25.0

TABLE II: Captioning results on the COCO test set.

score between the set of nouns in the produced caption and the set of categories of the relevant objects in the scene. By “relevant object” we mean those whose area covers at least 10% of the total image area, and thus, can be more useful to identify a scene. The latter score is defined as the intersection over union between the sets of nouns mentioned in two consecutively generated captions.

Additionally, we measure the agent’s overall **loquacity** (Loq) as the number of times it is activated by the speaker policy, normalized by the episode length. In other words, the loquacity can be seen as the inverse of the average number of navigation steps between two consecutive captions. Moreover, we resort to the recently-proposed **CLIP score** [59] (CLIP-S), in its unpaired definition, to evaluate the alignment between the agent’s view and the generated caption.

To evaluate the overall system on each episode, we define an ad-hoc score to measure the concept coverage of the generated descriptions, which is an important aspect of the task. The proposed **episode description score** (ED-S) reflects the ability of the robot to produce sufficient descriptions in strategic moments, so that the maximum amount of information collected in the environment is covered. The rationale is that it should capture the ability of the agent to mention all the relevant landmarks (objects and rooms) when needed, without unnecessary repetitions, making the description more useful and interesting. ED-S is defined as:

$$\text{ED-S} = \overline{\text{CLIP-S}} \cdot \text{IoU}(\mathbf{N}, \mathbf{O}) \cdot \%AS, \quad (9)$$

where $\overline{\text{CLIP-S}}$ is mean of the CLIP scores of all the captions produced during the episode. Moreover, \mathbf{N} is the list of nouns in all the captions produced during the episode and \mathbf{O} is the list of objects in the environment. The intersection-over-union operator IoU is implemented via the Jonker-Volgenant linear assignment algorithm [60]. Finally, $\%AS$ is the percentage of the total environment area visited by the agent. At the dataset level, the ED-S is given by the average of the scores obtained in the dataset episodes.

C. Experimental Results

Navigation Results. First, we compare the different exploration approaches alone on the MP3D and Gibson datasets.

	Loq	COCO only				COCO + Web-collected			
		Cov	Div	CLIP-S	ED-S	Cov	Div	CLIP-S	ED-S
Always	100.00	0.864	0.352	0.670	0.119	0.862	0.348	0.692	0.120
Depth									
D \geq 1.0	83.26	0.868	0.335	0.670	0.140	0.865	0.338	0.690	0.140
D \geq 1.5	55.24	0.871	0.323	0.664	0.203	0.868	0.330	0.683	0.204
D \geq 2.0	27.38	0.793	0.293	0.629	0.250	0.780	0.304	0.650	0.257
Object									
O \geq 1	41.73	0.793	0.314	0.663	0.222	0.784	0.332	0.682	0.225
O \geq 2	21.55	0.703	0.289	0.645	0.219	0.697	0.307	0.664	0.220
O \geq 3	7.58	0.416	0.232	0.549	0.107	0.410	0.260	0.561	0.105
Activation									
A \geq 4.5	87.79	0.866	0.340	0.672	0.134	0.864	0.343	0.691	0.134
A \geq 5.0	51.13	0.828	0.349	0.674	0.223	0.827	0.348	0.691	0.220
A \geq 5.5	2.20	0.133	0.153	0.455	0.038	0.140	0.153	0.464	0.040

TABLE III: Episode description results on Gibson Tiny Val.

	Loq	COCO only				COCO + Web-collected			
		Cov	Div	CLIP-S	ED-S	Cov	Div	CLIP-S	ED-S
Always	100.00	0.768	0.363	0.648	0.172	0.771	0.348	0.687	0.179
Depth									
D \geq 1.0	89.05	0.765	0.352	0.648	0.180	0.767	0.341	0.687	0.187
D \geq 2.0	45.06	0.751	0.317	0.637	0.155	0.750	0.311	0.668	0.160
D \geq 3.0	15.98	0.317	0.161	0.338	0.030	0.317	0.151	0.360	0.031
Object									
O \geq 1	75.82	0.754	0.340	0.635	0.190	0.756	0.333	0.670	0.196
O \geq 3	46.57	0.700	0.310	0.605	0.168	0.701	0.310	0.634	0.172
O \geq 5	19.90	0.616	0.255	0.533	0.106	0.614	0.254	0.553	0.107
Activation									
A \geq 4.5	82.05	0.765	0.348	0.641	0.106	0.767	0.337	0.676	0.107
A \geq 5.0	46.28	0.754	0.350	0.636	0.153	0.757	0.341	0.667	0.158
A \geq 5.5	1.28	0.325	0.118	0.347	0.015	0.328	0.116	0.362	0.016

TABLE IV: Episode description results on MP3D Test.

The results of this analysis are reported in Table I. The best agent in terms of the area seen (AS) is the impact-based method using density model estimation. In particular, this approach is able to efficiently explore both Gibson and MP3D datasets, giving its best in large environments. In fact, the small 0.87m² margin over the second best approach on Gibson becomes 9.99m² in the larger MP3D environments. Moreover, this method is still competitive in terms of IoU, being also the best in terms of Acc on the MP3D test split. In light of these results, we use the impact-based navigator with DME as the navigator of the overall approach.

Captioning Results. Then, we evaluate the performance of the captioner alone on the COCO dataset. The results of this analysis are reported in Table II. It can be observed that the CLIP-based variants are the best-performing ones, with a noticeable advantage over the region-based captioner. This confirms the representative power of CLIP features. The Base variant has also been trained on additional image-caption pairs from web-collected sources, which further increases its performance. It is worth mentioning that these results are in line with those of state-of-the-art captioners (*e.g.* [43], [61]). In light of these results, we use the CLIP-based Base variant as the captioner of the overall approach.

Episode Description Results. Finally, we compare variants of the overall approach using different speaking policies with different threshold values, and use as reference a dummy policy according to which the captioning module is always activated. The results are reported in Tables III and IV. It




Fig. 4: Sample observations and corresponding captions generated by our model.

can be noticed that the captioner trained on web-collected sources performs better than the variant trained on COCO only in terms of all metrics, suggesting its superior generalization capabilities and thus, suitability to be employed in an embodied setting. However, to evaluate on the overall task, the proposed ED-S score is more informative than the other metrics, which can nonetheless be used in combination with the ED-S to gain additional insights on the agents' behaviour. In fact, the values of all metrics but the ED-S are comparable in both datasets, while the ED-S is on average higher on Gibson: this is due to the fact that Gibson has on average smaller and less cluttered spaces, which can be more easily fully explored (higher values of the Cov on Gibson confirm this intuition). This trend is further confirmed by the fact that on the Gibson dataset the speaking policy must ensure the Loc being in a specific range (roughly between 20 and 80) to obtain the best ED-S scores, while on the wider spaces of MP3D, speaking policies ensuring a higher Loc lead to better performance. Qualitative examples of the output of our approach on selected observations are reported in Fig. 4.

Real-World Deployment. Exploration agents trained on the photorealistic environments of the Habitat simulator and general-purpose captioners allow the deployment of our approach to the real world, using a LoCobot platform [62]. For the deployment, the captioner is left untouched, whilst we modify the camera parameters of the navigator, such as camera height, the field of view, and depth sensor range to match the real-world setting. Furthermore, the deployed agent is trained by adding noise models fitted to mimic the LoCobot camera noise over the observations retrieved from the simulator. As the last step, we apply the correction presented by [12] to correct noisy real-world depth observations. In the video accompanying the submission, we show the agent exploring and describing a real-world apartment.

IV. CONCLUSION

In this work, we have presented an embodied exploration agent whose internal representation of the environment can be interpreted by non-expert users. This is achieved by equipping the agent with the ability to produce a natural language description of the observed scene when this is deemed interesting according to a speaking policy. The experimental results show that the proposed approach is a viable solution to gain insights into the perception and navigation capabilities of embodied agents. Moreover, the generalization capabilities of the modules adopted allow real-world deployment without major redesigns.

REFERENCES

- [1] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese, "Gibson Env: Real-world perception for embodied agents," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [2] 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," in *Proceedings of the International Conference on 3D Vision*, 2017.
- [3] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, et al., "Habitat: A Platform for Embodied AI Research," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [4] R. Bigazzi, F. Landi, S. Cascianelli, M. Cornia, L. Baraldi, and R. Cucchiara, "Embodied navigation at the art gallery," in *Proceedings of the International Conference on Image Analysis and Processing*, 2022.
- [5] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, "DD-PPO: Learning Near-Perfect PointGoal Navigators from 2.5 Billion Frames," in *Proceedings of the International Conference on Learning Representations*, 2019.
- [6] D. S. Chaplot, D. Gandhi, S. Gupta, A. Gupta, and R. Salakhutdinov, "Learning to Explore Using Active Neural SLAM," in *Proceedings of the International Conference on Learning Representations*, 2019.
- [7] S. K. Ramakrishnan, Z. Al-Halah, and K. Grauman, "Occupancy Anticipation for Efficient Exploration and Navigation," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [8] S. K. Ramakrishnan, D. Jayaraman, and K. Grauman, "An Exploration of Embodied Visual Exploration," *International Journal of Computer Vision*, vol. 129, no. 5, pp. 1616–1649, 2021.
- [9] R. Bigazzi, F. Landi, S. Cascianelli, L. Baraldi, M. Cornia, and R. Cucchiara, "Focus on Impact: Indoor Exploration with Intrinsic Motivation," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2985–2992, 2022.
- [10] F. Landi, R. Bigazzi, M. Cornia, S. Cascianelli, L. Baraldi, and R. Cucchiara, "Spot the difference: A novel task for embodied agents in changing environments," in *Proceedings of the International Conference on Pattern Recognition*, 2022.
- [11] A. Kadian, J. Truong, A. Gokaslan, A. Clegg, E. Wijmans, S. Lee, M. Savva, S. Chernova, and D. Batra, "Sim2Real Predictivity: Does evaluation in simulation predict real-world performance?" *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6670–6677, 2020.
- [12] R. Bigazzi, F. Landi, M. Cornia, S. Cascianelli, L. Baraldi, and R. Cucchiara, "Out of the Box: Embodied Navigation in the Real World," in *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, 2021.
- [13] J. Truong, S. Chernova, and D. Batra, "Bi-directional domain adaptation for sim2real transfer of embodied navigation agents," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2634–2641, 2021.
- [14] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [15] F. Landi, L. Baraldi, M. Cornia, M. Corsini, and R. Cucchiara, "Multimodal Attention Networks for Low-Level Vision-and-Language Navigation," *Computer Vision and Image Understanding*, vol. 210, p. 103255, 2021.
- [16] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the Nav-Graph: Vision-and-Language Navigation in Continuous Environments," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [17] P. Anderson, A. Shrivastava, J. Truong, A. Majumdar, D. Parikh, D. Batra, and S. Lee, "Sim-to-Real Transfer for Vision-and-Language Navigation," in *Proceedings of the Conference on Robot Learning*, 2021.
- [18] A. Das, S. Datta, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, "Embodied Question Answering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2018.
- [19] A. Das, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, "Neural Modular Control for Embodied Question Answering," in *Proceedings of the Conference on Robot Learning*, 2018.
- [20] E. Wijmans, S. Datta, O. Maksymets, A. Das, G. Gkioxari, S. Lee, I. Essa, D. Parikh, and D. Batra, "Embodied Question Answering in Photorealistic Environments With Point Cloud Perception," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [21] L. Yu, X. Chen, G. Gkioxari, M. Bansal, T. L. Berg, and D. Batra, "Multi-Target Embodied Question Answering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [22] S. Cascianelli, G. Costante, T. A. Ciarfuglia, P. Valigi, and M. L. Fravolini, "Full-GRU Natural Language Video Description for Service Robotics Applications," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 841–848, 2018.
- [23] M. Cornia, L. Baraldi, and R. Cucchiara, "SMARt: Training Shallow Memory-aware Transformers for Robotic Explainability," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2020.
- [24] R. Bigazzi, F. Landi, M. Cornia, S. Cascianelli, L. Baraldi, and R. Cucchiara, "Explore and Explain: Self-supervised Navigation and Recounting," in *Proceedings of the International Conference on Pattern Recognition*, 2020.
- [25] S. Tan, D. Guo, H. Liu, X. Zhang, and F. Sun, "Embodied scene description," *Autonomous Robots*, vol. 46, pp. 21–43, 2022.
- [26] A. Karpathy and L. Fei-Fei, "Deep visual-semantic alignments for generating image descriptions," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2015.
- [27] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *Proceedings of the International Conference on Machine Learning*, 2015.
- [28] P. Anderson, X. He, C. Buehler, D. Teney, M. Johnson, S. Gould, and L. Zhang, "Bottom-up and top-down attention for image captioning and visual question answering," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017.
- [30] W. Liu, S. Chen, L. Guo, X. Zhu, and J. Liu, "CPTN: Full Transformer Network for Image Captioning," *arXiv preprint arXiv:2101.10804*, 2021.
- [31] M. Cornia, L. Baraldi, G. Fiameni, and R. Cucchiara, "Universal Captioner: Inducing Content-Style Separation in Vision-and-Language Model Training," *arXiv preprint arXiv:2111.12727*, 2022.
- [32] S. J. Rennie, E. Marcheret, Y. Mroueh, J. Ross, and V. Goel, "Self-critical sequence training for image captioning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017.
- [33] L. Huang, W. Wang, J. Chen, and X.-Y. Wei, "Attention on Attention for Image Captioning," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [34] M. Cornia, M. Stefanini, L. Baraldi, and R. Cucchiara, "Meshed-Memory Transformer for Image Captioning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [35] X. Zhang, X. Sun, Y. Luo, J. Ji, Y. Zhou, Y. Wu, F. Huang, and R. Ji, "RSTNet: Captioning with Adaptive Attention on Visual and Non-Visual Words," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [36] S. K. Ramakrishnan, D. S. Chaplot, Z. Al-Halah, J. Malik, and K. Grauman, "PONI: Potential Functions for ObjectGoal Navigation with Interaction-free Learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2016.
- [38] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention*, 2015.
- [39] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *Proceedings of the International Conference on Machine Learning*, 2017.
- [40] D. S. Chaplot, D. P. Gandhi, A. Gupta, and R. R. Salakhutdinov, "Object Goal Navigation using Goal-Oriented Semantic Exploration," in *Advances in Neural Information Processing Systems*, 2020.

- [41] R. Raileanu and T. Rocktäschel, "RIDE: Rewarding impact-driven exploration for procedurally-generated environments," in *Proceedings of the International Conference on Learning Representations*, 2021.
- [42] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning Transferable Visual Models From Natural Language Supervision," in *Proceedings of the International Conference on Machine Learning*, 2021.
- [43] P. Zhang, X. Li, X. Hu, J. Yang, L. Zhang, L. Wang, Y. Choi, and J. Gao, "VinVL: Revisiting visual representations in vision-language models," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [44] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal Policy Optimization Algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [45] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft COCO: Common Objects in Context," in *Proceedings of the European Conference on Computer Vision*, 2014.
- [46] V. Ordonez, G. Kulkarni, and T. Berg, "Im2Text: Describing Images Using 1 Million Captioned Photographs," in *Advances in Neural Information Processing Systems*, 2011.
- [47] P. Sharma, N. Ding, S. Goodman, and R. Soicuc, "Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2018.
- [48] S. Changpinyo, P. Sharma, N. Ding, and R. Soicuc, "Conceptual 12M: Pushing Web-Scale Image-Text Pre-Training To Recognize Long-Tail Visual Concepts," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [49] K. Srinivasan, K. Raman, J. Chen, M. Bendersky, and M. Najork, "WIT: Wikipedia-based Image Text Dataset for Multimodal Multilingual Machine Learning," in *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021.
- [50] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li, "YFCC100M: The new data in multimedia research," *Communications of the ACM*, vol. 59, no. 2, pp. 64–73, 2016.
- [51] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: towards real-time object detection with region proposal networks," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [52] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh, "Large Batch Optimization for Deep Learning: Training BERT in 76 minutes," in *Proceedings of the International Conference on Learning Representations*, 2019.
- [53] D. Kingma and J. Ba, "Adam: a method for stochastic optimization," in *Proceedings of the International Conference on Learning Representations*, 2015.
- [54] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, "BLEU: a method for automatic evaluation of machine translation," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2002.
- [55] S. Banerjee and A. Lavie, "METEOR: An automatic metric for MT evaluation with improved correlation with human judgments," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics Workshops*, 2005.
- [56] C.-Y. Lin, "Rouge: A package for automatic evaluation of summaries," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics Workshops*, 2004.
- [57] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, "CIDEr: Consensus-based Image Description Evaluation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2015.
- [58] P. Anderson, B. Fernando, M. Johnson, and S. Gould, "SPICE: Semantic Propositional Image Caption Evaluation," in *Proceedings of the European Conference on Computer Vision*, 2016.
- [59] J. Hessel, A. Holtzman, M. Forbes, R. L. Bras, and Y. Choi, "CLIP-Score: A Reference-free Evaluation Metric for Image Captioning," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2021.
- [60] R. Jonker and A. Volgenant, "A shortest augmenting path algorithm for dense and sparse linear assignment problems," *Computing*, vol. 38, no. 4, pp. 325–340, 1987.
- [61] X. Li, X. Yin, C. Li, P. Zhang, X. Hu, L. Zhang, L. Wang, H. Hu, L. Dong, F. Wei, *et al.*, "Oscar: Object-semantic aligned pre-training for vision-language tasks," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [62] "LoCoBot: An Open Source Low Cost Robot," <https://locobot-website.netlify.com>.