

This is the peer reviewed version of the following article:

Embodied Navigation at the Art Gallery / Bigazzi, Roberto; Landi, Federico; Cascianelli, Silvia; Cornia, Marcella; Baraldi, Lorenzo; Cucchiara, Rita. - 13231:(2022), pp. 739-750. (21st International Conference on Image Analysis and Processing, ICIAP 2022 Lecce, Italy May 23-27, 2022) [10.1007/978-3-031-06427-2_61].

Springer Science and Business Media Deutschland GmbH

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

13/12/2025 21:45

(Article begins on next page)

Embodied Navigation at the Art Gallery

Roberto Bigazzi^[0000-0002-6457-1860], Federico Landi^[0000-0003-2092-1934],
Silvia Cascianelli^[0000-0001-7885-6050], Marcella Cornia^[0000-0001-9640-9385],
Lorenzo Baraldi^[0000-0001-5125-4957], and Rita Cucchiara^[0000-0002-2239-283X]

University of Modena and Reggio Emilia, Modena, Italy
{name.surname}@unimore.it

Abstract. Embodied agents, trained to explore and navigate indoor photorealistic environments, have achieved impressive results on standard datasets and benchmarks. So far, experiments and evaluations have involved domestic and working scenes like offices, flats, and houses. In this paper, we build and release a new 3D space with unique characteristics: the one of a complete art museum. We name this environment ArtGallery3D (AG3D). Compared with existing 3D scenes, the collected space is ampler, richer in visual features, and provides very sparse occupancy information. This feature is challenging for occupancy-based agents which are usually trained in crowded domestic environments with plenty of occupancy information. Additionally, we annotate the coordinates of the main points of interest inside the museum, such as paintings, statues, and other items. Thanks to this manual process, we deliver a new benchmark for PointGoal navigation inside this new space. Trajectories in this dataset are far more complex and lengthy than existing ground-truth paths for navigation in Gibson and Matterport3D. We carry on extensive experimental evaluation using our new space for evaluation and prove that existing methods hardly adapt to this scenario. As such, we believe that the availability of this 3D model will foster future research and help improve existing solutions.

Keywords: Embodied AI · Visual Navigation · Sim2Real.

1 Introduction

In recent years, Embodied AI has benefited from the introduction of rich datasets of 3D spaces and new tasks, ranging from exploration to PointGoal or Image-Goal navigation [9,26]. Such availability of 3D data allows to train and deploy modular embodied agents, thanks to powerful simulation platforms [23]. Despite the high number of available spaces, though, the topology and nature of the different scenes have low variance. Indeed, many environments represent apartments, offices, or houses. In this paper, we take a different path and collect and introduce the 3D space of an art gallery.

Current agents for embodied exploration feature a modular approach [5,10,20]. While the agents are trained for embodied exploration using deep reinforcement learning, this hierarchical paradigm allows for great adaptability on downstream

tasks. Hence, models trained to explore the Gibson dataset can solve Point-Goal navigation with satisfactory accuracy under the appropriate hypotheses. Furthermore, accurate and realistic simulating platforms such as Habitat [23] facilitate the deployment in the real world of the trained agents [7,15]. While agent architectures and simulating platforms are possible sources of improvement, there is a third important direction of research that regards the availability of 3D scenes to train and test the different agents. Indeed, the nature of the different environments influences the variety of tasks that the agent can learn and perform.

In this work, we contribute to this third direction by collecting and presenting a previously unseen type of 3D space, *i.e.*, a museum. This new environment for embodied exploration and navigation, named ArtGallery3D (AG3D), presents unique features when compared to flats and offices. First, the dimension of the rooms drastically increases, and the same goes for the size of the building itself. In our 3D model, some rooms are as big as 20×15 meters, while the floor hosting the art gallery spans a total of 2,000 square meters. However, dimensions are not the only difference with current available 3D spaces. As a second factor, the presented gallery is incredibly rich in visual features, offering multiple paintings, sculptures, and rare objects of historical and artistic interest. Every item represents a unique point of interest, and this is in contrast to traditional scenes where all elements have approximately the same visual relevance. Finally, the museum has sparse occupancy information. Many agents count on depth information to plan short-term displacements. However, when placed in the middle of an open empty hall, depth information is less informative. In our challenging 3D scene, the agent must learn to combine RGB and depth information and not be overconfident on immediately available knowledge on the occupancy map. All these challenges make our newly-proposed 3D space a valuable asset for current and future research.

Together with the 3D model of the museum, we present a dataset for embodied exploration and navigation. For the navigation task, we annotate the position of most of the points of interest in the museum. Examples include numerous paintings, sculptures, and other relevant objects. Finally, we present an experimental analysis including the performance of existing architectures on this novel benchmark and a discussion of potential future research directions made possible by the presence of the collected 3D space.

2 Related Work

Both autonomous robotics [5,14] and embodied AI [8,11,18,6,12,21] have recently witnessed a boost of interest, which has been enabled by the release of photorealistic 3D simulated environments. In such environments, algorithms for intelligent exploration and navigation can be developed safely and more quickly than in the real-world, before being easily deployed on real robotic platforms [15,7,2]. Among the datasets of spaces, the most commonly used are MP3D [9], Gibson [26], HM3D [22], and Replica [25]. These datasets mainly contain house-like

and office-like environments, with some environments taken from shops, garages, churches, and restaurants. Rooms in such environments are generally cluttered, and thus, rich of landmarks and texture information that the agent can exploit while navigating. In contrast, the presented AG3D environment has been collected in a museum, with larger, uncluttered spaces.

Algorithms developed in the simulated environments are typically trained with deep reinforcement learning, both for exploration and navigation tasks. The exploration task, which is often tackled to allow other downstream navigation tasks [27,28], consists in letting an agent equipped with visual sensors (*i.e.*, RGB-D cameras) freely navigate the environment to gather as much information as possible, usually in the form of an occupancy map. To this end, intrinsic rewards have been proposed, which can be based on novelty, curiosity, reconstruction enabling, and coverage [21,11,20]. For the navigation tasks, the agent is deployed in an unknown environment (*i.e.*, no map provided) and given some assignments in visual or textual form. These tasks include PointGoal navigation [1], where the robot is expected to reach a coordinates-specified goal, ImageGoal navigation [29], where the robot must reach an observation point in the environment that matches an image-specified goal, and ObjectGoal navigation [4], where the robot is asked to get to any instance of a label-specified object in the environment. Other related tasks involve embodied question answering [13] and vision-and-language navigation [3,16,17], where the robot must follow a natural language instruction to reach the goal. The environment presented in this paper is used for the exploration task and for the PointNav task. In this latter case, we define a variant in which the goal is expressed in terms of both coordinates and orientation.

3 ArtGallery3D (AG3D) Dataset

Existing datasets for indoor navigation comprise 3D acquisitions of different types of buildings, ranging from private houses, that cover the majority of the scenes, to offices and shops. Nevertheless, the focus of these datasets is on private spaces and there is low variance in terms of dimension and contained objects. In fact, to the best of our knowledge, among the publicly available datasets, no acquired indoor environment is composed of large rooms with a low occupied/free space ratio as in a museum. To overcome this deficiency in current literature we release a new indoor dataset for exploration and navigation captured inside a museum environment, called AG3D¹.

Acquisition. To build the 3D model of the art gallery, we employ a Matterport camera² and related software. This technology is the same employed to collect Matterport3D and HM3D datasets of spaces [9,22] and is particularly suitable to capture indoor photorealistic environments. We place the camera in the physical environment and capture a 360° RGB-D image of the surrounding. Then, we

¹ The dataset has been collected at the Galleria Estense museum of Modena and can be found at <https://github.com/aimagelab/ag3d>.

² <https://matterport.com/it/cameras/pro2-3d-camera>



Fig. 1. On the left: a view of the 3D model of the acquired environment. On the right: images captured during the acquisition of the scene.

repeat the same process after moving the camera approximately 1.5 meters away. Using consecutive panoramic acquisitions, the software is able to compute the 3D geometry of the space using depth information and the correspondences between the same keypoints in different acquisitions. To capture the entire museum, we make 232 different scans. Thanks to the high number of acquisitions, we are able to reproduce fine geometric and visual details of the original space (see Fig. 1). The resulting 3D model consists of more than 1430 m² of navigable space.

Dataset details. The proposed dataset allows two different tasks: exploration and navigation. Episodes for the exploration task include starting position and orientation of the agent which are sampled uniformly over the entire navigable space. The navigation dataset, instead, extends traditional PointGoal navigation where episodes are defined with a starting pose and a goal coordinate, including an additional final orientation vector. Conceptually, we can consider this setting as the link between PointGoal navigation and ImageGoal navigation since the goal is to rotate the agent towards a precise objective/scene, specifying the goal using coordinates instead of an image. We name this new setting PointGoal++ navigation (PointNav++). To create the navigation dataset we annotate 147 points of interest mostly consisting of paintings and statues. The annotated goal position is around 1 meter in front of the artwork and the goal orientation vector is directed to its center. For each point of interest, we define three episodes with different difficulties based on the geodesic distance between start and goal positions: easy ($< 15\text{m}$), medium ($> 15\text{m}$), and difficult ($> 30\text{m}$). In particular, thanks to the dimension of the acquired environment, each difficult episode has a geodesic distance larger than the longest path of MatterPort3D and Gibson datasets. A comparison of the geodesic distance distribution of the episodes of various available PointGoal navigation datasets is presented in Fig. 2. The introduction of AG3D enables the evaluation of agents on long navigation episodes which were previously not possible and highlights the inaccuracy of components of the architecture that accumulate error over time. The exploration task dataset contains 500k, 100, 1000 episodes respectively for training, validation, and test, while the PointNav++ dataset includes 411 annotated navigation episodes.

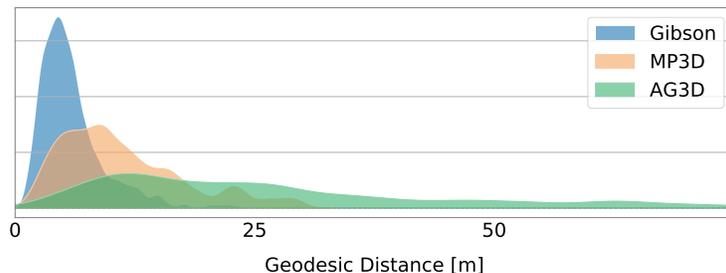


Fig. 2. Comparison of the distribution of the geodesic distances from starting position to goal position of the episode for different datasets.

4 Architecture

We provide an experimental analysis comparing recently proposed approaches on the devised environment, both for exploration and PointNav++ tasks. The evaluated methods are consistent with recent literature on embodied AI [10,20,5] and adopt an architecture shown in Fig. 3, which is composed of a neural mapper, a pose estimator, and a hierarchical navigation policy. The mapper generates a representation of the environment while the agent moves, the pose estimator is in charge of locating the agent in the environment, and the policy is responsible for the movement capabilities of the agent. The core difference between the evaluated approaches resides in the navigation policy, as described in the following. For further details, we refer the reader to the original papers.

4.1 Mapper

The mapper module incrementally builds an occupancy grid map of the environment in parallel with the navigation task. At each timestep, the RGB-D observations (s_t^{rgb}, s_t^d) coming from the visual sensors are processed to extract a $L \times L \times 2$ agent-centric map m_t where the channels indicate, respectively, the occupancy and exploration state of the currently observed region, and each pixel of the map describes the state of an area of 5×5 cm. The RGB observation is encoded using a ResNet-18 followed by a UNet, while the depth observation is encoded using another UNet. The features extracted from the two modalities are combined using CNNs at different levels of the output of the two UNet encoders and a final UNet decoder is used to process the combined features to retrieve the resulting local map m_t . Following the method proposed in [20], our mapper is not limited to predicting the occupancy map of the visible space but tries to infer also occluded and not visible regions of the local map. The global level map of the environment M_t has a dimensionality of $G \times G \times 2$, where $G > L$, and is built using local maps m_t step-by-step. At each timestep the pose of the agent x_t is used to apply a rototranslation to the local map, then, the transformed local map is finally registered to the global map M_t with a moving average.

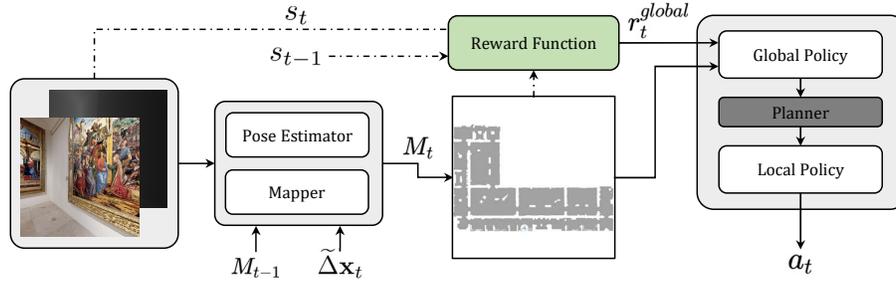


Fig. 3. Overall architecture of the models employed for exploration and navigation on AG3D.

4.2 Pose Estimator

In order to create a coherent representation of the environment during navigation, a precise and robust pose estimation needs to be achieved. To address problems like noise in the sensors and collisions with obstacles, we adopt a pose estimator and avoid the direct use of sensor readings. The pose estimator computes the pose of the agent $\mathbf{x}_t = (x_t, y_t, \theta_t)$ where (x_t, y_t) and θ_t are its position and orientation in the internal representation of the environment. The pose estimate \mathbf{x}_t is computed in an incremental way, adding the displacement $\Delta\mathbf{x}_t$ caused by the action a_t to the current pose estimate \mathbf{x}_t . In order to retrieve a first noisy estimate of the displacement $\Delta\mathbf{x}_t$, we use the difference between consecutive readings of the pose sensor $(\tilde{\mathbf{x}}_{t-1}, \tilde{\mathbf{x}}_t)$. To account for errors in the sensor displacement, consecutive local maps (m_{t-1}, m_t) coming from the mapper are used as feedback; m_{t-1} is reprojected using $\tilde{\Delta\mathbf{x}}_t$ to the same point of view of m_t and the concatenation of the transformed m_{t-1} and m_t is processed using a CNN to retrieve the final robust pose displacement $\Delta\mathbf{x}_t$. At each timestep $\Delta\mathbf{x}_t$ is used to compute the pose of the robot \mathbf{x}_t :

$$\mathbf{x}_t = \mathbf{x}_{t-1} + \Delta\mathbf{x}_t, \quad (1)$$

where we assume $\mathbf{x}_0 = (0, 0, 0)$ without loss of generality and \mathbf{x}_0 corresponds to the center of the map M_t with the agent facing north.

4.3 Navigation Policy

The navigation policy is the module that determines the movement of the agent in the environment. Its hierarchical design is required in order to allow the agent to uncouple high-level navigation concepts, such as moving across different rooms, and low-level concepts, like obstacle avoidance. The navigation policy is defined by a three-component module consisting of a global farsighted policy, a deterministic planner, and a local policy for atomic action inference.

Global Policy. The global policy is the high-level component of the navigation policy and is responsible for extracting a long-term goal on the global map g_t . The global policy takes as input an enriched $G \times G \times 4$ current global map M_t^+

retrieved stacking the two-channel global map M_t , the one-hot representation of the current position on the map, and the map of the already visited states. M_t^+ is in parallel cropped with respect to the position of the agent and max-pooled to a lower dimensionality $H \times H \times 4$. These two versions of M_t^+ are stacked together to obtain the final $H \times H \times 8$ input of the global policy. A CNN is used to sample a point of a $H \times H$ grid that is converted to a goal position on the global map g_t . The global policy is trained with reinforcement learning using PPO [24] to maximize different rewards used in literature.

In the experiments, we employ and compare different reward methods, namely Coverage, Anticipation, and Curiosity. The Coverage reward [10,21] maximizes the information gathered at each time-step, expressed in terms of the area seen. The Anticipation reward [20] is defined by comparing the predicted local occupancy map with the ground-truth considering also occluded areas. The Curiosity reward [19] encourages the agent towards areas that maximize the prediction error of a model trained to predict future states, thus improving the learning of the dynamics of the environment.

Planner. Given the global goal on the map, the planner has the task of computing a short-term goal on the map that the agent should reach. We employ an A* algorithm on the global map M_t to plan a path from the current position of the agent to the global goal and a local goal l_t is computed on the obtained trajectory within a distance D from the agent.

Local Policy. The local policy is the module that allows the movement of the agent in the environment and its objective is to reach the local goal l_t determined by the planner. The input of the local policy, formed by the relative displacement from the position of the agent to the local goal l_t and the current RGB observation s_t^{rgb} , is processed to compute an atomic action a_t . The available actions are: *move ahead 0.25m*, *turn left 10°*, *turn right 10°*, with the addition of a *stop* action when performing the navigation task. During training with reinforcement learning, the reward of the local policy r_t^{local} encourages the agent to reduce the distance from the local goal:

$$r_t^{local}(\mathbf{x}_t, \mathbf{x}_{t+1}) = d(\mathbf{x}_t) - d(\mathbf{x}_{t+1}), \quad (2)$$

where $d(\mathbf{x}_t)$ is the euclidean distance between the agent and the local goal l_t at timestep t . Following the hierarchical design, the global goal is sampled every N_G timesteps, while the local goal is reset if a new global goal is sampled, if the previous local goal is found to be in an occupied area, or if the previous local goal has been reached.

5 Experiments

We perform experiments on the proposed dataset comparing various models trained with different global rewards on another dataset, with models trained from scratch or finetuned on AG3D on exploration and PointNav++ to evaluate the performance gap between these approaches and highlight the difference between the characteristics of AG3D compared to other datasets. A sample episode of PointGoal++ navigation of ArtGallery3D is shown in Fig. 4.

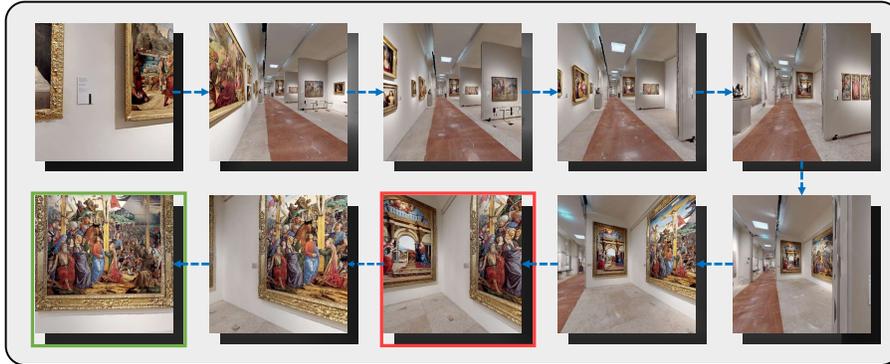


Fig. 4. An episode of PointNav++ in AG3D where consecutive frames have a distance of 10 timesteps approximately. The red frame indicates the stop action in the traditional PointNav task. The green frame corresponds to the stop action in PointNav++.

5.1 Experimental Setting

Evaluation protocol.

The baselines are trained with Coverage, Anticipation, and Curiosity rewards on the Gibson dataset for $\approx 5\text{M}$ frames corresponding to 12 GPU-days on NVIDIA V100. The best performing approach among the baseline is also both trained from scratch and finetuned, but since high-quality textures and memory occupancy of AG3D do not allow training with the same number of environments in parallel as Gibson, we trained the model from scratch on AG3D with the same GPU time for $\approx 2.8\text{M}$ frames, while the finetuned model is trained for $\approx 1\text{M}$ additional frames.

For the exploration task we evaluate the following metrics: **IoU** (Intersection-over-Union) between the map built during at end of the episode and the ground-truth map. **Acc** measures the correctly reconstructed map in m^2 . **AS** indicates the area seen by the agent during exploration (in m^2). **FIoU**, **OIoU**, **FAS**, and **OAS** measure, respectively, IoU and area seen for free and occupied portions of the environment. **TE** and **AE** are the translation and angular error between estimated and ground-truth pose measured respectively in meter and degrees. PointGoal++ navigation is evaluated considering these metrics: **D2G** (Distance to Goal) and **OE** (Orientation Error) are the mean geodesic distance to the goal and the mean orientation error at the end of the episode. The orientation error is computed considering the vector between the center of the artwork and the position of the agent as ground-truth. **SR** (Success Rate) is the percentage of episodes terminated successfully. In PointNav++ the agent needs to be within 0.2 meters to the goal and with an orientation error lower than 10 degrees. **PGSR** (PointNav Success Rate) and **ASR** (Angular Success Rate) consider only one component of SR; respectively D2G and OE. **SPL** and **SoftSPL** are success rates weighted on the length of the trajectory of the agent.

Implementation details. The experiments are performed extracting 128×128 RGB-D observations from the acquired 3D model using the Habitat simulator.

Table 1. Exploration results over the 100 episodes of the AG3D validation split in noise-free and noisy conditions.

Model	Training	IoU \uparrow	FIoU \uparrow	OIoU \uparrow	Acc \uparrow	AS \uparrow	FAS \uparrow	OAS \uparrow	TE \downarrow	AE \downarrow
Noise-Free										
Anticipation [20]	Gibson	0.163	0.170	0.157	294.4	290.6	258.3	32.3	0.0	0.0
Curiosity [19]	Gibson	0.175	0.184	0.166	317.9	317.5	281.7	35.8	0.0	0.0
Coverage [10,21]	Gibson	0.214	0.237	0.191	403.1	384.3	341.1	43.2	0.0	0.0
Coverage [10,21]	AG3D	0.219	0.239	0.200	400.6	354.6	316.6	38.0	0.0	0.0
Coverage [10,21]	Gibson+AG3D	0.296	0.313	0.278	531.8	470.2	418.1	52.1	0.0	0.0
Noisy										
Anticipation [20]	Gibson	0.144	0.157	0.131	269.6	281.7	249.9	31.9	0.48	2.95
Curiosity [19]	Gibson	0.119	0.151	0.086	251.1	307.3	272.8	34.5	2.62	15.99
Coverage [10,21]	Gibson	0.148	0.203	0.093	327.0	380.7	337.9	42.8	2.98	12.54
Coverage [10,21]	AG3D	0.144	0.200	0.088	320.0	356.4	317.3	39.2	2.65	13.35
Coverage [10,21]	Gibson+AG3D	0.191	0.266	0.116	427.3	461.9	413.2	48.7	2.68	10.16

The maximum length of the exploration episodes during training is set to $T = 500$. Regarding the mapping process, we set $L = 101$ and $G = 2881$ for the local and global map dimensionalities. The action space grid $H \times H$ of the global policy is 240×240 . The maximum distance of the local goal l_t from the position of the agent is $D = 0.5\text{m}$ for exploration and $D = 0.25\text{m}$ for PointNav++.

5.2 Experimental Results

Exploration results. As a first experiment, in Table 1 we compare the considered models on the exploration task on the AG3D validation set. Each exploration episode has a length of $T = 1000$ timesteps during which the agent has to disclose the initially unknown environment. Among the baselines trained only on the Gibson dataset, the Coverage-based model achieves the best results in terms of IoU and Area Seen in both noise-free and noisy settings. The model trained with Coverage from scratch obtains competitive results even using fewer training frames (2.8M vs. 5M), showing the importance of adapting the models to AG3D. This conclusion is supported by the fact that the model trained on Gibson and finetuned for 1M frames on AG3D achieves the best results on noise-free and noisy settings, with a significant margin on the second-best model. In both settings, the performance gap in terms of Area Seen (85.9m^2 and 81.2m^2) and IoU (0.082 and 0.043) between the best models trained only on Gibson dataset and using AG3D denotes the need of adapting the weight of the models to the different visual characteristics and occupation of AG3D.

PointNav++ results. Moving on to the navigation task, models trained on exploration substitute the global goal with a fixed goal specified by the navigation episode. Experimental results on PointNav++, shown in Table 2, present a similar trend as on the exploration task. In fact, the Coverage model has the best results in terms of SPL and Success Rate related metrics among the models trained only on Gibson. Moving to the Coverage models trained on AG3D, in the noise-free setting, the model trained from scratch achieves the best results even in comparison to the finetuned counterpart which is trained with more than double the total observations (2.8M vs 6M). This behavior can be explained

Table 2. PointNav++ results on the AG3D navigation episodes under noise-free and noisy settings.

Model	Training	SPL \uparrow	SoftSPL \uparrow	SR \uparrow	PNSR \uparrow	ASR \uparrow	Steps \downarrow	D2G \downarrow	OE \downarrow
Noise-Free									
Anticipation [20]	Gibson	0.697	0.780	0.803	0.873	0.808	364.3	4.131	12.2
Curiosity [19]	Gibson	0.625	0.706	0.732	0.803	0.732	416.4	7.934	17.0
Coverage [10,21]	Gibson	0.760	0.838	0.876	0.954	0.883	314.6	0.700	5.2
Coverage [10,21]	AG3D	0.805	0.875	0.898	0.973	0.908	270.3	0.268	4.8
Coverage [10,21]	Gibson+AG3D	0.793	0.873	0.883	0.964	0.891	273.1	0.323	5.0
Noisy									
Anticipation [20]	Gibson	0.211	0.788	0.224	0.255	0.387	338.6	3.152	32.2
Curiosity [19]	Gibson	0.225	0.655	0.243	0.275	0.341	446.3	9.746	38.7
Coverage [10,21]	Gibson	0.228	0.783	0.243	0.260	0.392	348.6	3.165	34.1
Coverage [10,21]	AG3D	0.235	0.832	0.248	0.273	0.445	306.2	2.420	28.8
Coverage [10,21]	Gibson+AG3D	0.373	0.853	0.399	0.443	0.543	283.8	1.430	19.8

by the performance of its mapper that is trained for more frames using visual observation from AG3D (2.8M vs 1M) and extracts a more detailed map sacrificing robustness and generalization. Accordingly, in the noisy setting, the higher robustness of the Coverage-based model finetuned on AG3D regains the first place with a noteworthy margin on the other models, while the Coverage model trained from scratch goes down to the second position in terms of SPL and SR. As in the case of the exploration task, the performance gap between models trained on Gibson and using AG3D (0.045 and 0.145 for SPL in noise-free and noisy settings) stresses the importance of adapting the parameters to the features extracted from AG3D. Moreover, it is worth noting that the gap of the best model from noise-free to noisy navigation (0.432 for SPL) is a consequence of the length of the navigation episodes of AG3D, and the difficulty of performing precise lengthy trajectories in the presence of noise. This is an interesting aspect that the AG3D dataset offers for exploration in future works.

6 Conclusion

In this work, we introduced the AG3D photorealistic 3D dataset for embodied exploration and PointGoal navigation tasks. The dataset has been collected in an art gallery, which features larger and more uncluttered spaces compared to most of the environments available in commonly used benchmark datasets. For the PointNav task, we propose a variant that is more suitable to the type of environment in the AG3D dataset. The variant entails not only reaching the specified coordinates, as in standard PointNav but also assuming a specified orientation. We also present an experimental comparison of state-of-the-art approaches on the devised dataset, which can serve as baselines for future research on embodied AI tasks performed in museum-like environments.

Acknowledgment

This work has been supported by “Fondazione di Modena” and the “European Training Network on PErsonalized Robotics as SErvice Oriented applications” (PERSEO) MSCA-ITN-2020 project.

References

1. Anderson, P., Chang, A., Chaplot, D.S., Dosovitskiy, A., Gupta, S., Koltun, V., Kosecka, J., Malik, J., Mottaghi, R., Savva, M., et al.: On evaluation of embodied navigation agents. arXiv preprint arXiv:1807.06757 (2018)
2. Anderson, P., Shrivastava, A., Truong, J., Majumdar, A., Parikh, D., Batra, D., Lee, S.: Sim-to-real transfer for vision-and-language navigation. In: CoRL (2021)
3. Anderson, P., Wu, Q., Teney, D., Bruce, J., Johnson, M., Sünderhauf, N., Reid, I., Gould, S., van den Hengel, A.: Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments. In: CVPR (2018)
4. Batra, D., Gokaslan, A., Kembhavi, A., Maksymets, O., Mottaghi, R., Savva, M., Toshev, A., Wijmans, E.: Objectnav revisited: On evaluation of embodied agents navigating to objects. arXiv preprint arXiv:2006.13171 (2020)
5. Bigazzi, R., Landi, F., Cascianelli, S., Baraldi, L., Cornia, M., Cucchiara, R.: Focus on Impact: Indoor Exploration with Intrinsic Motivation. RA-L (2022)
6. Bigazzi, R., Landi, F., Cornia, M., Cascianelli, S., Baraldi, L., Cucchiara, R.: Explore and Explain: Self-supervised Navigation and Recounting. In: ICPR (2020)
7. Bigazzi, R., Landi, F., Cornia, M., Cascianelli, S., Baraldi, L., Cucchiara, R.: Out of the Box: Embodied Navigation in the Real World. In: CAIP (2021)
8. Cascianelli, S., Costante, G., Ciarfuglia, T.A., Valigi, P., Fravolini, M.L.: Full-gru natural language video description for service robotics applications. RA-L **3**(2), 841–848 (2018)
9. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D Data in Indoor Environments. In: 3DV (2017)
10. Chaplot, D.S., Gandhi, D., Gupta, S., Gupta, A., Salakhutdinov, R.: Learning To Explore Using Active Neural SLAM. In: ICLR (2019)
11. Chen, T., Gupta, S., Gupta, A.: Learning Exploration Policies for Navigation. In: ICLR (2019)
12. Cornia, M., Baraldi, L., Cucchiara, R.: SMArT: Training Shallow Memory-aware Transformers for Robotic Explainability. In: ICRA (2020)
13. Das, A., Datta, S., Gkioxari, G., Lee, S., Parikh, D., Batra, D.: Embodied question answering. In: CVPR (2018)
14. Irshad, M.Z., Ma, C.Y., Kira, Z.: Hierarchical Cross-Modal Agent for Robotics Vision-and-Language Navigation. In: ICRA (2021)
15. Kadian, A., Truong, J., Gokaslan, A., Clegg, A., Wijmans, E., Lee, S., Savva, M., Chernova, S., Batra, D.: Sim2Real Predictivity: Does evaluation in simulation predict real-world performance? RA-L **5**(4), 6670–6677 (2020)
16. Krantz, J., Wijmans, E., Majumdar, A., Batra, D., Lee, S.: Beyond the nav-graph: Vision-and-language navigation in continuous environments. In: ECCV (2020)
17. Landi, F., Baraldi, L., Cornia, M., Corsini, M., Cucchiara, R.: Multimodal Attention Networks for Low-Level Vision-and-Language Navigation. CVIU **210**, 103255 (2021)
18. Niroui, F., Zhang, K., Kashino, Z., Nejat, G.: Deep Reinforcement Learning Robot for Search and Rescue Applications: Exploration in Unknown Cluttered Environments. RA-L **4**(2), 610–617 (2019)
19. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: ICML (2017)
20. Ramakrishnan, S.K., Al-Halah, Z., Grauman, K.: Occupancy Anticipation for Efficient Exploration and Navigation. In: ECCV (2020)

21. Ramakrishnan, S.K., Jayaraman, D., Grauman, K.: An Exploration of Embodied Visual Exploration. *IJCV* **129**(5), 1616–1649 (2021)
22. Ramakrishnan, S.K., Gokaslan, A., Wijmans, E., Maksymets, O., Clegg, A., Turner, J.M., Undersander, E., Galuba, W., Westbury, A., Chang, A.X., Savva, M., Zhao, Y., Batra, D.: Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In: *NeurIPS* (2021), <https://openreview.net/forum?id=-v40uqNs5P>
23. Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., et al.: Habitat: A Platform for Embodied AI Research. In: *ICCV* (2019)
24. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal Policy Optimization Algorithms. *arXiv preprint arXiv:1707.06347* (2017)
25. Straub, J., Whelan, T., Ma, L., Chen, Y., Wijmans, E., Green, S., Engel, J.J., Mur-Artal, R., Ren, C., Verma, S., Clarkson, A., Yan, M., Budge, B., Yan, Y., Pan, X., Yon, J., Zou, Y., Leon, K., Carter, N., Briales, J., Gillingham, T., Mueggler, E., Pesqueira, L., Savva, M., Batra, D., Strasdat, H.M., Nardi, R.D., Goesele, M., Lovegrove, S., Newcombe, R.: The Replica Dataset: A Digital Replica of Indoor Spaces. *arXiv preprint arXiv:1906.05797* (2019)
26. Xia, F., Zamir, A.R., He, Z., Sax, A., Malik, J., Savarese, S.: Gibson env: Real-world perception for embodied agents. In: *CVPR* (2018)
27. Ye, J., Batra, D., Das, A., Wijmans, E.: Auxiliary tasks and exploration enable ObjectNav. In: *ICCV* (2021)
28. Ye, J., Batra, D., Wijmans, E., Das, A.: Auxiliary Tasks Speed Up Learning Point Goal Navigation. In: *CoRL* (2021)
29. Zhu, Y., Mottaghi, R., Kolve, E., Lim, J.J., Gupta, A., Fei-Fei, L., Farhadi, A.: Target-driven visual navigation in indoor scenes using deep reinforcement learning. In: *ICRA* (2017)