Occupancy Anticipation for Efficient Exploration and Navigation

Santhosh K. Ramakrishnan[1,2], Ziad Al-Halah[1], and Kristen Grauman[1,2]

[1] The University of Texas at Austin, Austin TX 78712, USA
[2] Facebook AI Research, Austin TX 78701, USA
srama@cs.utexas.edu, ziadlhlh@gmail.com, grauman@cs.utexas.edu

Abstract. State-of-the-art navigation methods leverage a spatial memory to generalize to new environments, but their occupancy maps are limited to capturing the geometric structures directly observed by the agent. We propose *occupancy anticipation*, where the agent uses its egocentric RGB-D observations to infer the occupancy state beyond the visible regions. In doing so, the agent builds its spatial awareness more rapidly, which facilitates efficient exploration and navigation in 3D environments. By exploiting context in both the egocentric views and top-down maps our model successfully anticipates a broader map of the environment, with performance significantly better than strong baselines. Furthermore, when deployed for the sequential decision-making tasks of exploration and navigation, our model outperforms state-of-the-art methods on the Gibson and Matterport3D datasets. Our approach is the winning entry in the 2020 Habitat PointNav Challenge. *Project page:* `http://vision.cs.utexas.edu/projects/occupancy_anticipation/`

1 Introduction

In visual navigation, an agent must move intelligently through a 3D environment in order to reach a goal. Visual navigation has seen substantial progress in the past few years, fueled by large-scale datasets and photo-realistic 3D environments [4,9,67,62], simulators [67,31,3,35], and public benchmarks [12,3,35]. Whereas traditionally navigation was attempted using purely geometric representations (i.e., SLAM), recent work shows the power of *learned* approaches to navigation that integrate both geometry and semantics [72,19,50,38,70,11]. Learned approaches operating directly on pixels and/or depth as input can be robust to noise [11,10] and can generalize well on unseen environments [19,35,70,10] —even outperforming pure SLAM given sufficient experience [35].

One of the key factors for success in navigation has been the movement towards complex map-based architectures [19,41,11,10] that capture both geometry [19,11,10] and semantics [19,41,18,22], thereby facilitating efficient policy learning and planning. These learned maps allow an agent to exploit prior knowledge from training scenes when navigating in novel test environments.

Despite such progress, state-of-the-art approaches to navigation are limited to encoding *what the agent actually sees in front of it*. In particular, they build maps

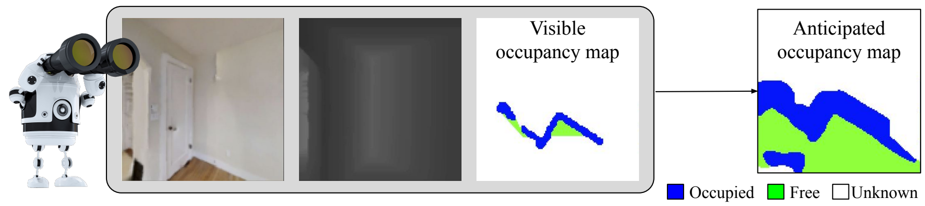

Fig. 1: **Occupancy anticipation:** A robot's perception of the 3D world is limited by its field-of-view and obstacles (the visible map). We propose to anticipate occupancy for unseen regions (anticipated map) by exploiting the context from egocentric views. We then train a deep reinforcement learning agent to move intelligently in a 3D environment, rewarding movements that improve the anticipated map.

of the environment using only the *observed* regions, whether via geometry [11,22] or learning [19,41,18,10]. Thus, while promising, today's models suffer from an important inefficiency: to map a space in the 3D environment as free or occupied, the agent must directly see evidence thereof in its egocentric camera.

Our key idea is to *anticipate occupancy*. Rather than wait to directly observe a more distant or occluded region of the 3D environment to declare its occupancy status, the proposed agent infers occupancy for unseen regions based on the visual context in its egocentric views. For example, in Fig. 1, with only the partial observation of the scene, the agent could infer that it is quite likely that the wall extends to its right, a corridor is present on its left, and the region immediately in front of it is free space. Such intelligent extrapolation beyond the observed space would lead to more efficient exploration and navigation. To achieve this advantage, we introduce a model that anticipates occupancy maps from normal field-of-view RGB(D) observations, while aggregating its predictions over time in tight connection with learning a navigation policy. Furthermore, we incorporate the anticipation objective directly into the agent's exploration policy, encouraging movements in the 3D space that will efficiently yield broader and more accurate inferred occupancy maps.

We validate our approach on Gibson [67] and Matterport3D [9], two 3D environment datasets spanning over 170 real-world spaces with a variety of obstacles and floor plans. Using only RGB(D) inputs to anticipate occupancy, the proposed agent learns to explore intelligently, achieving faster and more accurate maps compared to a state-of-the-art approach for neural SLAM [10], and navigating more efficiently than strong baselines. Furthermore, for navigation under noisy actuation and sensing, our agent improves the state of the art, winning the 2020 Habitat PointNav Challenge [1] by a margin of 6.3 SPL points.

Our main contributions are: (1) a novel occupancy anticipation framework that leverages visual context from egocentric RGB(D) views; (2) a novel exploration approach that incorporates intelligent anticipation for efficient environment mapping, providing better maps in less time; and (3) successful navigation results that improve the state of the art.

2 Related work

Navigation Classical approaches to visual navigation perform passive or active SLAM to reconstruct geometric point-clouds [64,21] or semantic maps [5,49], facilitated by loop closures or learned odometry [7,36,8]. More recent work uses deep learning to learn navigation [72,19,50,38,70,68,53,57] or exploration [42,6,51,26,45] policies in an end-to-end fashion. Explicit *map-based* navigation models [20,41,18,11] usually outperform their implicit counterparts by being more sample-efficient, generalizing well to unseen environments, and even transferring from simulation to real robots [19,10]. However, existing approaches only encode *visible* regions for mapping (i.e., the ground plane projection of the observed or inferred depth). In contrast, our model goes beyond the visible cues and anticipates maps for unseen regions to accelerate navigation.

Layout estimation Recent work predicts 3D Manhattan layouts of indoor scenes given 360 panoramas [73,69,63,66,14]. These methods predict structured outputs such as layout boundaries [73,63], corners [73], and floor/ceiling probability maps [69]. However, they do not extrapolate to unseen regions. FloorNet [33] and Floor-SP [27] use walkthroughs of previously scanned buildings to reconstruct detailed floorplans that may include predictions for the room type, doors, objects, etc. However, they assume that the layouts are polygonal, the scene is fully explored, and that detailed human annotations are available. Our occupancy map representation can be seen as a new way for the agent to infer the layout of its surroundings. Unlike any of the above approaches, our model does not make strict assumptions on the scene structure, nor does it require detailed semantic annotations. Furthermore, the proposed anticipation model is learned jointly with the exploration policy and without human guidance. Finally, unlike prior work, our goal is to accelerate navigation and map creation.

Scene completion Past work in scene completion focuses on pixelwise reconstruction of 360 panoramas with limited glimpses [26,44,45,55], inpainting [43,24,32], and inferring unseen 3D structure and semantics [61,71]. While some methods allow pixelwise extrapolation outside the current field of view (FoV) [45,61,71,25], they do not permit inferences about occluded regions in the scene. Our results show that this limitation is detrimental to successful occupancy estimation (cf. our view extrapolation baseline). SSCNet [60] performs voxelwise geometric and semantic predictions for unseen 3D structures; however, it is computationally expensive, requires voxelwise semantic labels, limits predictions to the agent's FoV, and needs carefully curated viewpoints for training. In contrast, our approach predicts 2D occupancy from egocentric RGB(D) views, and it learns to do so in an active perception setting. Since the agent controls its own camera, the viewpoints tend to be more challenging than those in curated datasets of human-taken photos used in the scene completion literature [60,61,26,44,71].

Occupancy maps In robotics, methods for occupancy focus on building continuous representations of the world [40,47,56], mapping for autonomous driving [23,37,59,34,39], and indoor robot navigation [29,15,58]. Prior extrapolation methods assume wide FoV LIDAR inputs, only exploit geometric cues from top-down views, and demonstrate results in relatively simple 2D floorplans devoid of non-wall obstacles [30,29,15,58]. In contrast, our approach does not require expensive LIDAR sensors. It operates with standard RGB(D) camera inputs, and it exploits both semantic and geometric context from those egocentric views to perform accurate occupancy anticipation. Furthermore, we demonstrate efficient navigation in visually rich 3D environments with challenging obstacles other than walls. Finally, unlike prior work, our anticipation models are learned jointly with a navigation policy that rewards accurate anticipatory mapping.

3 Approach

We propose an occupancy anticipation approach for efficient exploration and navigation. Our model anticipates areas not directly visible to the agent because of occlusion (e.g., behind a table, around a corner) or due to being outside its FoV. The agent's first-person view is provided in the form of RGB-D images (see Fig. 2 left). The goal is to anticipate the occupancy for a fixed region in front of the agent, and integrate those predictions over time as the agent moves about.

Next, we define the task setup and notation, followed by our approach for occupancy anticipation (Sec. 3.1) and a new formulation for exploration that rewards correctly anticipated regions (Sec. 3.2). Then, we explain how our occupancy anticipation model can be integrated into a state-of-the-art approach [10] for autonomous exploration and navigation in 3D environments (Sec. 3.3).

3.1 Occupancy anticipation model

We formulate occupancy anticipation as a pixelwise classification task. The egocentric occupancy is represented as a two-channel top-down map $p \in [0,1]^{2\times V\times V}$ which comprises a local area of $V \times V$ cells in front of the camera. Each cell in the map represents a 25mm $\times$ 25mm region. The two channels contain the probabilities (confidence values) of the cell being occupied and explored, respectively. A cell is considered to be occupied if there is an obstacle, and it is explored if we know whether it is occupied or free. For training, we use the 3D meshes of indoor environments (Sec. 4.1) to obtain the ground-truth local occupancy of a $V \times V$ region in front of the camera, which includes parts that may be occluded or outside the field of view (Fig. 2, bottom right).

Our occupancy anticipation model consists of three main components (Fig. 2):
(1) Feature extraction: Given egocentric RGB-D inputs, we compute:
RGB CNN features: We encode the RGB images using blocks 1 and 2 of a ResNet-18 that is pre-trained on ImageNet, followed by three additional convolution layers that prepare these features to be passed forward with the visible occupancy map. This step extracts a mixture of textural and semantic features.
Depth projection: We estimate a map of occupied, free, and unknown space by setting height thresholds on the point cloud obtained from depth and camera intrinsics [11]. Consistent with past work [11,10], we restrict the projection-based

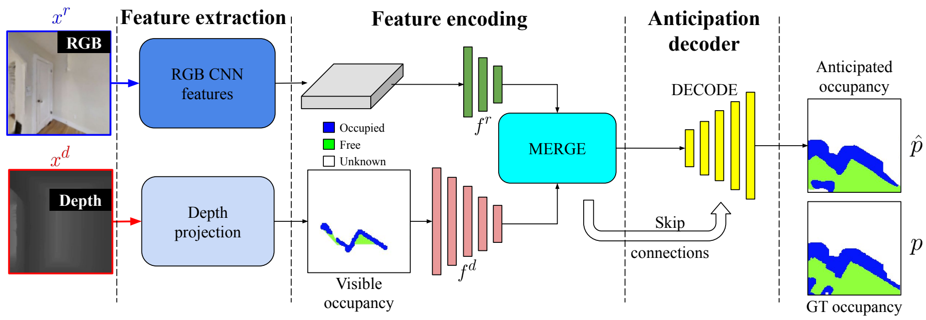

Fig. 2: Our occupancy anticipation model uses RGB(D) inputs to extract features, and processes them using a UNet to anticipate the occupancy. The depth map is projected to the ground plane to obtain the preliminary visible occupancy map. See text.

estimates to points within $\sim$ 3m, the range at which modern depth sensors would provide reliable results. This yields the initial visible occupancy map.

(2) Feature encoding: Given the RGB-D features, we independently encode them using UNet [48] encoders and project them to a common feature space. We encode the depth projection features using a stack of five convolutional blocks which results in features $\boldsymbol{f}^d = f^d_{1:5}$. Since the RGB features are already at a lower resolution, we use only three convolutional blocks to encode them, which results in features $\boldsymbol{f}^r = f^r_{3:5}$. We then combine these features using the MERGE module which contains layer-specific convolution blocks to merge each $[\boldsymbol{f}^r_i, \boldsymbol{f}^d_i]$:

$$\boldsymbol{f} = \text{MERGE}(\boldsymbol{f}^d, \boldsymbol{f}^r). \tag{1}$$

For experiments with only the depth modality, we skip the RGB feature extractor and MERGE layer and directly use the occupancy features obtained from the depth image. For experiments with only the RGB modality, we learn a model to infer the visible occupancy features from RGB (to be defined at the end of Sec. 4.1) and use that instead of the features computed from the depth image.

(3) Anticipation decoding: Given the encoded features $\boldsymbol{f}$, we use a UNet decoder that outputs a $2 \times V \times V$ tensor of probabilities:

$$\hat{p} = \sigma(\text{DECODE}(\boldsymbol{f})), \tag{2}$$

where $\hat{p} \in [0, 1]^{2\times V\times V}$ is the estimated egocentric occupancy and σ is the sigmoid activation function. For training the occupancy anticipation model, we use binary cross entropy loss per pixel and per channel:

$$L = \sum_{i=1}^{V^2} \sum_{j=1}^{2} -\Big[p_{ij} \log \hat{p}_{ij} + (1 - p_{ij}) \log(1 - \hat{p}_{ij})\Big], \tag{3}$$

where p is the ground-truth (GT) occupancy map that is derived from the 3D mesh of training environments (see Sec. S5 in Supp. for details).

So far, we have presented our occupancy anticipation approach supposing a single RGB-D observation as input. However, our model is ultimately used in the context of an embodied agent that moves in the environment and actively collects a sequence of RGB-D views to build a complete map of the environment. Next, we introduce a new reward function that utilizes the agent's anticipation performance to guide its exploration during training.

3.2 Anticipation reward for exploration policy learning

In *visual exploration*, an agent must quickly map a new environment without having a specified target. Prior work on exploration [11,16,10,46] often uses area-coverage—the area seen in the environment during navigation—as a reward function to guide exploration. However, the traditional area-coverage approach is limited to rewarding the agent only for *directly seeing* areas. Arguably, an ideal exploration agent would obtain an accurate and complete map of the environment *without* necessarily directly observing all areas.

Thus, we propose to encourage exploratory behaviors that yield a correctly *anticipated* map. In this case, the occupancy entries in the map need not be obtained via direct agent observations to register a reward; it is sufficient to correctly infer them. In particular, we reward agent actions that yield accurate occupancy predictions for the global environment map, i.e., the number of grid cells where the predicted occupancy matches the layout of the environment.

More concretely, let $\hat{m}_t \in [0,1]^{2\times G\times G}$ be the global environment map obtained by anticipating occupancy for the RGB-D observations $\{x^r_{1:t}, x^d_{1:t}\}$ from time 1 to t, and then geometrically registering the predictions to a single global map based on the agent's pose estimates at each time step (see Fig. 3). Note $G > V$. Let m be the ground-truth layout of the environment. Then, the unnormalized accuracy of a map prediction $\hat{m}$ is measured as follows:

$$\text{Accuracy}(\hat{m}, m) = \sum_{i=1}^{G^2}\sum_{j=1}^{2} \mathbb{1}[\hat{m}_{ij} = m_{ij}], \tag{4}$$

where $\mathbb{1}[\hat{m}_{ij} = m_{ij}]$ is an indicator function that returns one if $\hat{m}_{ij} = m_{ij}$ and zero otherwise. We reward the increase in map accuracy from time $t-1$ to t:

$$R^{anticp}_t = \text{Accuracy}(\hat{m}_t, m) - \text{Accuracy}(\hat{m}_{t-1}, m). \tag{5}$$

This function rewards actions leading to correct global map predictions, irrespective of whether the agent actually *observed* those locations. For example, if the agent correctly anticipates free space behind a table and is rewarded for that, it then learns to avoid spending additional time around tables in the future to observe that space directly. Resources can be instead allocated to visiting more interesting regions that are harder to anticipate. Additionally, this reward

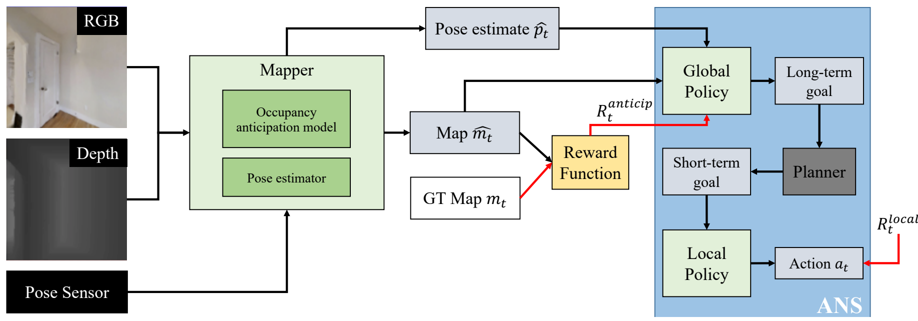

Fig. 3: **Exploration with occupancy anticipation:** We introduce two key upgrades to the original Active Neural SLAM (ANS) model [10] (see text): (1) We replace the projection unit in the mapper with our occupancy anticipation model (see Fig. 2). (2) We replace the area-coverage reward function with the proposed reward (Eqn. 5), which encourages the agent to efficiently explore and build accurate maps through occupancy anticipation. Note that the reward signals (in red) are provided only during training.

provides a better learning signal while training under noisy conditions by accounting for mapping errors arising from noisy pose and map predictions. Thus, our approach encourages more intelligent exploration behavior by injecting our anticipated occupancy idea directly into the agent's sequential decision-making.

3.3 Exploration and navigation with occupancy anticipation

Having defined the core occupancy anticipation components, we now demonstrate how our model can be used to benefit embodied navigation in 3D environments. We consider both exploration (discussed above) and *PointGoal navigation* [52,2], a.k.a PointNav, where the agent must navigate efficiently to a target specified by a displacement vector from the agent's starting position.

For both tasks, we adapt the state-of-the-art Active Neural SLAM (ANS) architecture [10] that previously achieved the best exploration results in the literature and was the winner of the 2019 Habitat PointNav challenge. However, our anticipation model is generic and can be easily integrated with most map-based embodied navigation models [19,11,17].

The ANS model is a hierarchical, modular policy for exploration that consists of a mapper, a planner, a local policy, and a global policy (shown in Fig. 3). Given RGB images, the mapper estimates the egocentric occupancy and agent pose, and then temporally aggregates the maps into a global top-down map using the pose estimates. At regular time intervals Δ, the global policy picks a location on the global map to explore. A shortest-path planner decides what trajectory to take from the current position to the target and picks an intermediate goal (within 1.25m) to navigate to. The local policy then selects actions that lead to the intermediate goal; it gets another intermediate goal upon reaching the current goal. See [10] for details. Critically, and like other prior work, the model of [10] is supervised to generate occupancy estimates based solely on the *visible* occupancy obtained from the egocentric views.

We adapt ANS by modifying the mapper and the reward function. For the mapper, we replace the projection unit from ANS with our anticipation model (see Fig. 3). Additionally, we account for incorrect occupancy estimates in two ways: (1) we filter out high entropy predictions and (2) we maintain a moving average estimate of occupancy at each location in the global map (see Sec. S7 in Supp.). For the reward function, we use the anticipation-based reward presented in Sec. 3.2.

We train the exploration policy with our anticipation model end-to-end, as this allows adapting to the changing distribution of the agent's inputs. Both the local and the global reinforcement learning policies are trained with Proximal Policy Optimization (PPO) [54]. In our model, the reward of the global policy is our anticipation-based reward defined in Eqn. 5. This replaces the traditional area-coverage reward used in ANS and other current models [11,10,46], which rewards the increment in the actual area seen, not the correctly registered area in the map. The reward for the local policy is simply based on the reduction in the distance to the local goal: $R_t^{local} = d_{t-1} - d_t$, where d is the Euclidean distance between the current position and the local goal.

4 Experiments

In the following experiments we demonstrate that 1) our occupancy anticipation module can successfully infer unseen parts of the map (Sec. 4.2) and 2) trained together with an exploration and navigation policy, it accelerates active mapping and navigation in new environments (Sec. 4.3 and Sec. 4.4).

4.1 Experimental setup

We use the Habitat [35] simulator along with Gibson [67] and Matterport3D [9] environments. Each dataset contains around 90 challenging large-scale photo-realistic 3D indoor environments such as houses and office buildings. On average, the Matterport3D environments are larger. Our observation space consists of 128×128 RGB-D observations and odometry sensor readings that denote the change in the agent's pose x, y, θ. Our action space consists of three actions: MOVE-FORWARD by 25cm, TURN-LEFT by $10°$, TURN-RIGHT by $10°$. For navigation, we add a STOP action, which the agent emits when it believes it has reached the goal. We simulate noisy actuation and odometer readings for realistic evaluation (see Sec. S6 in Supp.).

We train our exploration models on Gibson, and then transfer them to Point-Goal navigation on Gibson and exploration on Matterport3D. We use the default train/val/test splits provided for both datasets [35] with disjoint environments across the splits. For evaluation on Gibson, we divide the validation environments into small (area less than 36m^2) and large (area greater than 36m^2) to observe the influence of environment size on results. For policy learning, we use

Method	IoU %			F1 score %		
	free	occ.	mean	free	occ.	mean
all-free	30.1	0	15.1	43.6	0	21.8
all-occupied	0	25.1	12.6	0	39.2	19.6
ANS(rgb)	12.1	14.9	13.5	19.6	24.9	22.5
ANS(depth)	14.5	24.1	19.3	23.1	37.6	30.4
View-extrap.	15.5	26.4	21.0	25.0	40.4	32.7
OccAnt(rgb)	44.4	47.9	46.1	58.2	62.9	60.6
OccAnt(depth)	50.4	**61.9**	56.1	63.8	**75.0**	69.4
OccAnt(rgbd)	**51.5**	61.5	**56.5**	**64.9**	74.8	**69.8**

Table 1: **Occupancy anticipation results** on the Gibson validation set. Our models, OccAnt(·), substantially improve the map quality and extent, showing the advantage of learning to anticipate 3D structures beyond those directly observed.

the Adam optimizer and train on episodes of length 1000 for $1.5-2$ million frames of experience. Please see Sec. S8 in Supp. for more details.

Baselines: We define baselines based on prior work:

- **ANS(rgb)** [10]: This is the state-of-the-art Active Neural SLAM approach for exploration and navigation. We use the original mapper architecture [10], which infers the visible occupancy from RGB.[3]
- **ANS(depth)**: We use depth projection to infer the visible occupancy (similar to [11]) instead of predicting it from RGB.
- **View-extrap.**: We extrapolate an 180° FoV depth map from 90° FoV RGB-D and project it to the top-down view. This is representative of scene completion approaches [61,71]. See Sec. S11 in Supp. for network details.
- **OccAnt(GT)**: This is an upper bound that cheats by using the ground-truth anticipation maps for exploration and navigation.

We implement all baselines on top of the ANS framework. Our goal is to show the impact of our occupancy model, while fixing the backbone navigation architecture and policy learning approach across methods for a fair comparison. We consider three versions of our models based on the input modality:

- **OccAnt(depth)**: anticipate occupancy given the visible occupancy map.
- **OccAnt(rgb)**: anticipate occupancy given only the RGB image. We replace the depth projections in Fig. 2 with the pre-trained ANS(rgb) estimates (kept frozen throughout training).
- **OccAnt(rgbd)**: anticipate occupancy given the full RGB-D inputs.

By default, our methods use the proposed anticipation reward from Sec. 3.2. We denote ablations without this reward as “w/o AR”.

[3] We use our own implementation of ANS since authors’ code was unavailable at the time of our experiments. See Sec. S7 in Supp. for details.

4.2 Occupancy anticipation results

First we evaluate the per-frame prediction accuracy of the mapping models trained during exploration. We evaluate on a separate dataset of images sampled from validation environments in Gibson at uniform viewpoints from discrete locations on a 1m grid, a total of $1,034$ (input, output) samples. This allows standardized evaluation of the mapper, independent of the exploration policy.

To quantify the local occupancy maps' accuracy, we compare the predicted maps to the ground truth. We report the Intersection over Union (IoU) and F1 scores for the "free" and "occupied" classes independently. In addition to the baselines from Sec. 4.1, we add two naive baselines that classify all locations as free (all-free), or occupied (all-occupied).

Table 1 shows the results. Our anticipation models OccAnt are substantially better than all the baselines. Comparing different modalities, OccAnt(depth) is much better than OccAnt(rgb) under all the metrics. This makes sense, as visible occupancy is directly computable from the depth input, but must be inferred for RGB (see Fig. 4). Interestingly, the rgbd models are not better than the depth-only models, likely because (1) geometric cues are more easily learned from depth than RGB, and (2) the RGB encoder contains significantly more parameters and could lead to overfitting. See Table S5 in Supp. for network sizes. Overall, Table 1 demonstrates our occupancy anticipation models successfully broaden the coverage of the map beyond the visible regions.

4.3 Exploration results

Next we deploy our models for visual exploration. The agent is given a limited time budget (T=1000) to intelligently explore and build a 2D top-down occupancy map of a previously unseen environment.

To quantify exploration, we measure both map quality and speed (number of agent actions): (1) **Map accuracy (m^2):** the area in the global map built during exploration (both free and occupied) that matches with the ground-truth layout of the environment. The map is built using predicted occupancy maps which are registered using estimated pose (may be noisy). Note that this is an unnormalized accuracy measure (see Eqn. 4). (2) **IoU:** the intersection over union between that same global map and the ground-truth layout of the environment. (3) **Area seen (m^2):** the amount of free and occupied regions *directly seen* during exploration. The map for this metric is built using ground-truth pose and depth-projections (similar to [11,10]). (4) **Episode steps:** the number of actions taken by the agent. While the first two metrics measure the quality of the created map, the latter two are a function of how (and how long) the agent moved to get that map. Higher accuracy in fewer steps or lower area-seen is better.

All agents are trained on 72 scenes from Gibson under noisy odometry and actuation (see Sec. 4.1), and evaluated on Gibson and Matterport3D under both noisy and noise-free conditions.

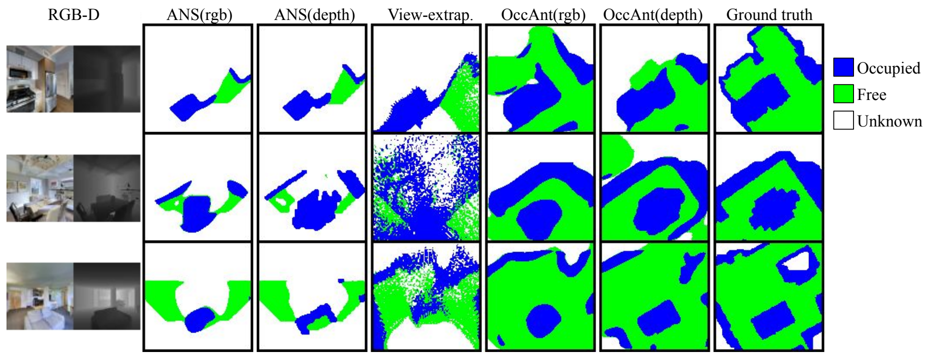

Fig. 4: **Per-frame local occupancy predictions:** First and last columns show the RGB-D input and anticipation ground-truth, respectively. ANS(*) are restricted to only predicting occupancy for visible regions. View-extrap. extrapolates, but is unable to predict occupancy for occluded regions (first row) and struggles to make correct predictions in cluttered scenes (second row). Our model successfully anticipates with either RGB or depth. For example, in the first row, we successfully predict the presence of a corridor and another room on the left. In the second row, we successfully predict the presence of navigable space behind the table. In the third row, we are able to correctly anticipate the free space behind the chair and the corridor to the right.

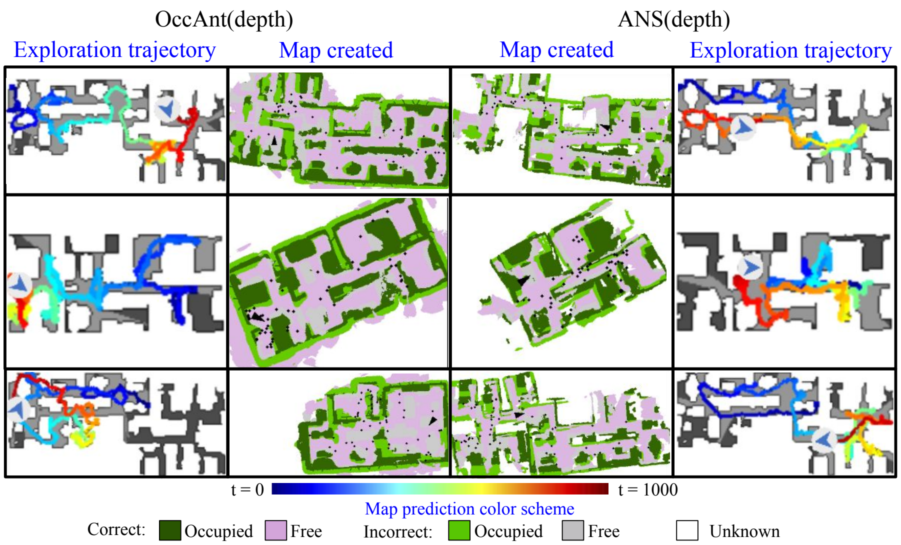

Fig. 5: **Exploration examples:** We compare OccAnt with ANS [10] in Gibson under noisy actuation and odometry. The exploration trajectories and the corresponding maps are shown at the extremes and center, respectively. **Row 1:** Both methods cover similar area, but our method better anticipates the unseen parts with fewer registration errors. **Row 2:** Our method achieves better area coverage and mapping quality whereas the baseline gets stuck in a small room for extended periods of time. **Row 3:** A failure case for our method, where it gets stuck in one part of the house after anticipating that a narrow corridor leading to a different room was occupied.

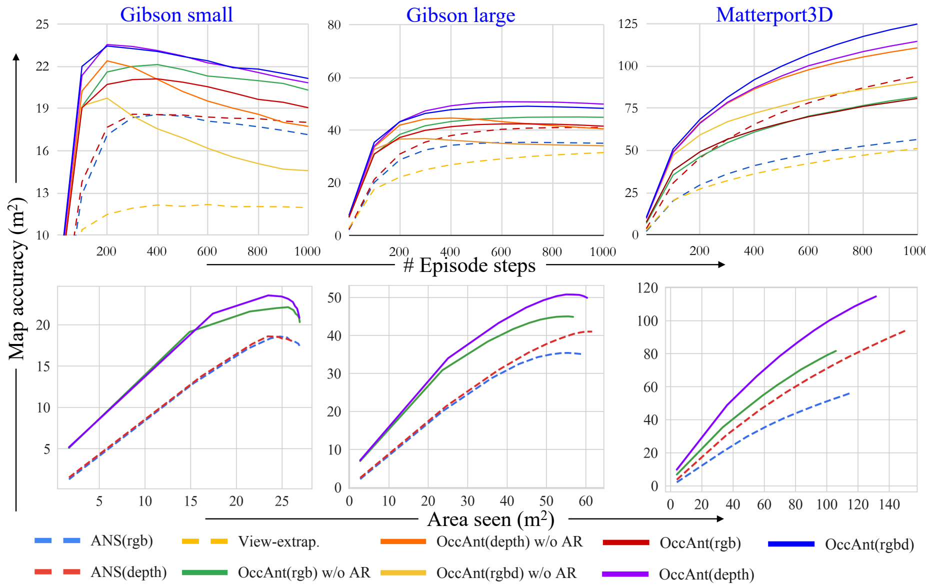

Fig. 6: **Exploration results:** Map accuracy (m^2) as a function of episode duration (top row) and area seen (bottom row) for Gibson (small and large splits) and Matterport3D under noisy conditions (see Sec. S1 in Supp. for noise-free). Higher and steeper curves are better. **Top:** Our OccAnt approach rapidly attains higher map accuracy than the baselines (dotted lines). **Bottom:** OccAnt achieves higher map accuracy for the same area seen (we show the best variants here to avoid clutter). These results show the agent *actively moves better* to explore the environment with occupancy anticipation.

	Noisy test conditions						Noise-free test conditions					
	Gibson small		Gibson large		Matterport3D		Gibson small		Gibson large		Matterport3D	
Method	Map acc.	IoU	Map acc.	IoU	Map acc.	IoU	Map acc.	IoU	Map acc.	IoU	Map acc.	IoU
ANS(rgb) [10]	18.5	55	35.0	47	44.7	18	22.4	76	43.4	64	53.4	23
ANS(depth)	18.5	56	39.4	53	72.5	26	21.4	74	48.0	72	85.9	34
View-extrap.	12.0	26	28.1	27	39.4	14	12.1	27	26.5	27	33.9	13
OccAnt(rgb) w/o AR	21.8	66	44.2	57	65.8	23	22.6	71	45.2	60	64.4	24
OccAnt(depth) w/o AR	20.2	58	44.2	54	92.7	29	**24.9**	**84**	**54.1**	**75**	. **104.7**	**38**.
OccAnt(rgbd) w/o AR	16.9	45	35.6	40	76.3	23	24.8	**84**	52.0	71	98.7	34
OccAnt(rgb)	20.9	62	42.1	54	66.2	22	22.3	70	43.5	58	64.4	22
OccAnt(depth)	**22.7**	**71**	**50.3**	**67**	94.1	**33**	24.8	83	53.1	74	96.5	35
OccAnt(rgbd)	**22.7**	**71**	48.4	62	**99.9**	32	24.5	82	51.0	69	100.3	34
OccAnt(GT)	21.7	67	51.9	63	-	-	26.1	93	65.4	91	-	-

Table 2: **Timed exploration results:** Map quality at T=500 for all models and datasets. See text for details.

Fig. 6 shows the exploration results. Our approach generally outperforms the baselines, improving the map quality more rapidly, whether in terms of time (top row) or area seen (bottom row). When compared on a same-modality basis, we see that OccAnt(rgb) converges much faster than ANS(rgb). Similarly, OccAnt(depth) is able to rapidly improve the map quality and outperforms ANS(depth) on all cases. This apples-to-apples comparison shows that anticipating occupancy leads to much more efficient mapping in unseen environments. Again, using depth generally provides more reliable mapping than pure RGB.

Furthermore, the proposed anticipation reward generally provides significant benefits to map accuracy in the noisy setting (compare our full model to the "w/o AR" models in Fig. 6). While map accuracy generally increases over time for noise-free conditions (see Sec. S1 in Supp.), it sometimes saturates early or even declines slightly over time in the noisy setting as noisy pose estimates accumulate and hurt map registration accuracy. This is most visible in Gibson small (top left plot). However, our anticipatory reward alleviates this decline.

Table 2 summarizes the map accuracy and IoU for all methods at T=500. Our method obtains significant improvements, supporting our claim that occupancy anticipation accelerates exploration and mapping. Additionally, perfect anticipation with the OccAnt(GT) model gives comparably good noisy exploration, and good gains in noise-free exploration (+10-20% IoU). This shows that there is indeed a lot of mileage in anticipating occupancy; our model moves the state-of-the-art towards this ceiling. Fig. 5 shows example exploration trajectories and the final global map predictions on Gibson.

4.4 Navigation results

Next we evaluate the utility of occupancy anticipation for quickly reaching a target. In PointNav [52,2], the agent is given a 2D coordinate (relative to its position) and needs to reach that target as quickly as possible. Following [10], we use noise-free evaluation and directly transfer the mapper, planner, and local policy learned during exploration to this task. In this way, instead of navigating to a point specified by the global policy, the agent has to navigate to a fixed goal location. To evaluate navigation, we use the standard metrics—success rate, success rate normalized by inverse path length (SPL) [2], and time taken. The agent succeeds if it stops within 0.2m of the target under a time budget of $T = 1000$.

Table 3 shows the navigation results on the Gibson validation set. Our approach outperforms the baselines. Thus, not only does occupancy anticipation successfully map the environment, but it also allows the agent to move to a specified goal more quickly by modeling the navigable spaces. This apples-to-apples comparison shows that our idea improves the state of the art for PointNav. As with exploration, using ground truth (GT) anticipation leads to good gains in the navigation performance, and our methods bridge the gap between the prior state of the art and perfect anticipation.

In concurrent work, the DD-PPO approach [65] obtains 0.96 SPL for PointNav, but it requires 2.5 billion frames of experience to do so (and it fails for noisy conditions; see below). To achieve the performance of our method (0.8 SPL in 2M frames), DD-PPO requires more than 50$\times$ the experience. Our sample efficiency can be attributed to explicit mapping along with occupancy anticipation.

Method	SPL %	Success %	Time taken
ANS(rgb) [10]	66.8	87.9	254.109
ANS(depth)	76.8	86.6	226.161
View-extrap.	10.4	33.3	835.556
OccAnt(rgb)	71.2	88.2	223.411
OccAnt(depth)	77.8	91.3	194.751
OccAnt(rgbd)	**80.0**	**93.0**	**171.874**
OccAnt(GT)	89.5	96.0	125.018

Table 3: **PointNav results:** Our approach provides more efficient navigation.

	Test standard			Test challenge		
Rank	Team	SPL %	Success %	Team	SPL %	Success %
1	**OccupancyAnticipation**	19.2	24.8	**OccupancyAnticipation**	20.9	27.5
2	ego-localization [13]	10.4	13.6	ego-localization [13]	14.6	19.2
3	Information Bottleneck	5.0	7.5	DAN [28]	13.2	25.3
4	cogmodel_team	0.8	1.3	Information Bottleneck	6.0	8.8
5	UCULab	0.5	0.8	cogmodel_team	0.7	1.2
6	Habitat Team (DD-PPO) [65]	0.0	0.2	UCULab	0.1	0.2

Table 4: **Habitat Challenge 2020 results:** Our approach is the winning entry.

Finally, we validate our approach on the 2020 Habitat PointNav Challenge [1], which requires the agent to adapt to noisy RGB-D sensors and noisy actuators, and to operate without an odometer. This presents a much more difficult evaluation setup than past work which assumes perfect odometry as well as noise-free sensing and actuation [35,10,65]. See Sec. S13 in Supp. for more details. Table 4 shows the results. Our method won the challenge, outperforming the competing approaches by large margins. While our approach generalizes well to this setting, DD-PPO [65] fails (0 SPL) due to its reliance on perfect odometry.

5 Conclusion

We introduced the idea of occupancy anticipation from egocentric views in 3D environments. By learning to anticipate the navigable areas beyond the agent's actual field of view, we obtain more accurate maps more efficiently in novel environments. We demonstrate our idea both for individual local maps, as well as integrated within sequential models for exploration and navigation, where the agent continually refines its (anticipated) map of the world. Our results clearly demonstrate the advantages on multiple datasets, including improvements to the state-of-the-art embodied AI model for exploration and navigation.

Acknowledgements

UT Austin is supported in part by DARPA Lifelong Learning Machines and the GCP Research Credits Program. We thank Devendra Singh Chaplot for clarifying the implementation details for ANS.

References

1. The Habitat Challenge 2020. https://aihabitat.org/challenge/2020/
2. 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)
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: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
4. Armeni, I., Sax, A., Zamir, A.R., Savarese, S.: Joint 2D-3D-Semantic Data for Indoor Scene Understanding. ArXiv e-prints (Feb 2017)
5. Bao, S.Y., Bagra, M., Chao, Y.W., Savarese, S.: Semantic structure from motion with points, regions, and objects. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 2703–2710. IEEE (2012)
6. Burda, Y., Edwards, H., Pathak, D., Storkey, A., Darrell, T., Efros, A.A.: Large-scale study of curiosity-driven learning. In: arXiv:1808.04355 (2018)
7. Cadena, C., Carlone, L., Carrillo, H., Latif, Y., Scaramuzza, D., Neira, J., Reid, I., Leonard, J.J.: Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. IEEE Transactions on robotics **32**(6), 1309–1332 (2016)
8. Carrillo, H., Reid, I., Castellanos, J.A.: On the comparison of uncertainty criteria for active slam. In: 2012 IEEE International Conference on Robotics and Automation. pp. 2080–2087. IEEE (2012)
9. Chang, A., Dai, A., Funkhouser, T., Nießner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: Learning from rgb-d data in indoor environments. In: Proceedings of the International Conference on 3D Vision (3DV) (2017), matterPort3D dataset license available at: http://kaldir.vc.in.tum.de/matterport/MP_TOS.pdf.
10. Chaplot, D.S., Gupta, S., Gandhi, D., Gupta, A., Salakhutdinov, R.: Learning to explore using active neural mapping. 8th International Conference on Learning Representations, ICLR 2020 (2020)
11. Chen, T., Gupta, S., Gupta, A.: Learning exploration policies for navigation. In: 7th International Conference on Learning Representations, ICLR 2019 (2019)
12. Das, A., Datta, S., Gkioxari, G., Lee, S., Parikh, D., Batra, D.: Embodied question answering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops. pp. 2054–2063 (2018)
13. Datta, S., Maksymets, O., Hoffman, J., Lee, S., Batra, D., Parikh, D.: Integrating egocentric localization for more realistic pointgoal navigation agents. CVPR 2020 Embodied AI Workshop (2020)
14. Dhamo, H., Navab, N., Tombari, F.: Object-driven multi-layer scene decomposition from a single image. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
15. Elhafsi, A., Ivanovic, B., Janson, L., Pavone, M.: Map-predictive motion planning in unknown environments. arXiv preprint arXiv:1910.08184 (2019)
16. Fang, K., Toshev, A., Fei-Fei, L., Savarese, S.: Scene memory transformer for embodied agents in long-horizon tasks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 538–547 (2019)
17. Gan, C., Zhang, Y., Wu, J., Gong, B., Tenenbaum, J.B.: Look, listen, and act: Towards audio-visual embodied navigation. arXiv preprint arXiv:1912.11684 (2019)

18. Gordon, D., Kembhavi, A., Rastegari, M., Redmon, J., Fox, D., Farhadi, A.: Iqa: Visual question answering in interactive environments. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4089–4098 (2018)
19. Gupta, S., Davidson, J., Levine, S., Sukthankar, R., Malik, J.: Cognitive mapping and planning for visual navigation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2616–2625 (2017)
20. Gupta, S., Fouhey, D., Levine, S., Malik, J.: Unifying map and landmark based representations for visual navigation. arXiv preprint arXiv:1712.08125 (2017)
21. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision. Cambridge university press (2003)
22. Henriques, J.F., Vedaldi, A.: Mapnet: An allocentric spatial memory for mapping environments. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8476–8484 (2018)
23. Hoermann, S., Bach, M., Dietmayer, K.: Dynamic occupancy grid prediction for urban autonomous driving: A deep learning approach with fully automatic labeling. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). pp. 2056–2063. IEEE (2018)
24. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. ACM Transactions on Graphics (ToG) **36**(4), 1–14 (2017)
25. Jayaraman, D., Gao, R., Grauman, K.: Shapecodes: self-supervised feature learning by lifting views to viewgrids. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 120–136 (2018)
26. Jayaraman, D., Grauman, K.: Learning to look around: Intelligently exploring unseen environments for unknown tasks. In: Computer Vision and Pattern Recognition, 2018 IEEE Conference on (2018)
27. Jiacheng Chen, Chen Liu, J.W., Furukawa, Y.: Floor-sp: Inverse cad for floorplans by sequential room-wise shortest path. In: The IEEE International Conference on Computer Vision (ICCV) (2019)
28. Karkus, P., Ma, X., Hsu, D., Kaelbling, L.P., Lee, W.S., Lozano-Pérez, T.: Differentiable algorithm networks for composable robot learning. arXiv preprint arXiv:1905.11602 (2019)
29. Katyal, K., Popek, K., Paxton, C., Burlina, P., Hager, G.D.: Uncertainty-aware occupancy map prediction using generative networks for robot navigation. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 5453–5459. IEEE (2019)
30. Katyal, K., Popek, K., Paxton, C., Moore, J., Wolfe, K., Burlina, P., Hager, G.D.: Occupancy map prediction using generative and fully convolutional networks for vehicle navigation. arXiv preprint arXiv:1803.02007 (2018)
31. Kolve, E., Mottaghi, R., Han, W., VanderBilt, E., Weihs, L., Herrasti, A., Gordon, D., Zhu, Y., Gupta, A., Farhadi, A.: AI2-THOR: An Interactive 3D Environment for Visual AI. arXiv (2017)
32. Li, Y., Liu, S., Yang, J., Yang, M.H.: Generative face completion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3911–3919 (2017)
33. Liu, C., Wu, J., Furukawa, Y.: Floornet: A unified framework for floorplan reconstruction from 3d scans. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 201–217 (2018)
34. Lu, C., Dubbelman, G.: Hallucinating beyond observation: Learning to complete with partial observation and unpaired prior knowledge (2019)

35. Manolis Savva*, Abhishek Kadian*, Oleksandr Maksymets*, Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., Parikh, D., Batra, D.: Habitat: A Platform for Embodied AI Research. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
36. Martinez-Cantin, R., De Freitas, N., Brochu, E., Castellanos, J., Doucet, A.: A bayesian exploration-exploitation approach for optimal online sensing and planning with a visually guided mobile robot. Autonomous Robots **27**(2), 93–103 (2009)
37. Mohajerin, N., Rohani, M.: Multi-step prediction of occupancy grid maps with recurrent neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 10600–10608 (2019)
38. Mousavian, A., Toshev, A., Fišer, M., Košecká, J., Wahid, A., Davidson, J.: Visual representations for semantic target driven navigation. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 8846–8852. IEEE (2019)
39. Müller, M., Dosovitskiy, A., Ghanem, B., Koltun, V.: Driving policy transfer via modularity and abstraction. arXiv preprint arXiv:1804.09364 (2018)
40. O'Callaghan, S.T., Ramos, F.T.: Gaussian process occupancy maps. The International Journal of Robotics Research **31**(1), 42–62 (2012)
41. Parisotto, E., Salakhutdinov, R.: Neural map: Structured memory for deep reinforcement learning. arXiv preprint arXiv:1702.08360 (2017)
42. Pathak, D., Agrawal, P., Efros, A.A., Darrell, T.: Curiosity-driven exploration by self-supervised prediction. In: International Conference on Machine Learning (2017)
43. Pathak, D., Krahenbuhl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: Feature learning by inpainting. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2016)
44. Ramakrishnan, S.K., Grauman, K.: Sidekick policy learning for active visual exploration. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 413–430 (2018)
45. Ramakrishnan, S.K., Jayaraman, D., Grauman, K.: Emergence of exploratory look-around behaviors through active observation completion. Science Robotics **4**(30) (2019). https://doi.org/10.1126/scirobotics.aaw6326, `https://robotics.sciencemag.org/content/4/30/eaaw6326`
46. Ramakrishnan, S.K., Jayaraman, D., Grauman, K.: An exploration of embodied visual exploration. arXiv preprint arXiv:2001.02192 (2020)
47. Ramos, F., Ott, L.: Hilbert maps: scalable continuous occupancy mapping with stochastic gradient descent. The International Journal of Robotics Research **35**(14), 1717–1730 (2016)
48. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
49. Salas-Moreno, R.F., Newcombe, R.A., Strasdat, H., Kelly, P.H., Davison, A.J.: Slam++: Simultaneous localisation and mapping at the level of objects. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1352–1359 (2013)
50. Savinov, N., Dosovitskiy, A., Koltun, V.: Semi-parametric topological memory for navigation. arXiv preprint arXiv:1803.00653 (2018)
51. Savinov, N., Raichuk, A., Marinier, R., Vincent, D., Pollefeys, M., Lillicrap, T., Gelly, S.: Episodic curiosity through reachability. arXiv preprint arXiv:1810.02274 (2018)

52. Savva, M., Chang, A.X., Dosovitskiy, A., Funkhouser, T., Koltun, V.: Minos: Multimodal indoor simulator for navigation in complex environments. arXiv preprint arXiv:1712.03931 (2017)
53. Sax, A., Emi, B., Zamir, A.R., Guibas, L., Savarese, S., Malik, J.: Mid-level visual representations improve generalization and sample efficiency for learning visuomotor policies. arXiv preprint arXiv:1812.11971 (2018)
54. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017)
55. Seifi, S., Tuytelaars, T.: Where to look next: Unsupervised active visual exploration on 360 {\deg} input. arXiv preprint arXiv:1909.10304 (2019)
56. Senanayake, R., Ganegedara, T., Ramos, F.: Deep occupancy maps: a continuous mapping technique for dynamic environments (2017)
57. Shen, W.B., Xu, D., Zhu, Y., Guibas, L.J., Fei-Fei, L., Savarese, S.: Situational fusion of visual representation for visual navigation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2881–2890 (2019)
58. Shrestha, R., Tian, F.P., Feng, W., Tan, P., Vaughan, R.: Learned map prediction for enhanced mobile robot exploration. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 1197–1204. IEEE (2019)
59. Sless, L., Cohen, G., Shlomo, B.E., Oron, S.: Self supervised occupancy grid learning from sparse radar for autonomous driving. arXiv preprint arXiv:1904.00415 (2019)
60. Song, S., Yu, F., Zeng, A., Chang, A.X., Savva, M., Funkhouser, T.: Semantic scene completion from a single depth image. Proceedings of 30th IEEE Conference on Computer Vision and Pattern Recognition (2017)
61. Song, S., Zeng, A., Chang, A.X., Savva, M., Savarese, S., Funkhouser, T.: Im2pano3d: Extrapolating 360 structure and semantics beyond the field of view. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3847–3856 (2018)
62. 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)
63. Sun, C., Hsiao, C.W., Sun, M., Chen, H.T.: Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
64. Thrun, S.: Probabilistic robotics. Communications of the ACM **45**(3), 52–57 (2002)
65. Wijmans, E., Kadian, A., Morcos, A., Lee, S., Essa, I., Parikh, D., Savva, M., Batra, D.: Dd-ppo: Learning near-perfect pointgoal navigators from 2.5 billion frames (2020)
66. Wu, W., Fu, X.M., Tang, R., Wang, Y., Qi, Y.H., Liu, L.: Data-driven interior plan generation for residential buildings. ACM Trans. Graph. **38**(6) (Nov 2019). https://doi.org/10.1145/3355089.3356556, `https://doi.org/10.1145/3355089.3356556`
67. Xia, F., Zamir, A.R., He, Z., Sax, A., Malik, J., Savarese, S.: Gibson env: Real-world perception for embodied agents. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9068–9079 (2018), gibson dataset license agreement available at `https://storage.googleapis.com/gibson_material/Agreement%20GDS%2006-04-18.pdf`

68. Yang, J., Ren, Z., Xu, M., Chen, X., Crandall, D., Parikh, D., Batra, D.: Embodied amodal recognition: Learning to move to perceive objects. In: ICCV (2019)
69. Yang, S.T., Wang, F.E., Peng, C.H., Wonka, P., Sun, M., Chu, H.K.: Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3363–3372 (2019)
70. Yang, W., Wang, X., Farhadi, A., Gupta, A., Mottaghi, R.: Visual semantic navigation using scene priors. arXiv preprint arXiv:1810.06543 (2018)
71. Yang, Z., Pan, J.Z., Luo, L., Zhou, X., Grauman, K., Huang, Q.: Extreme relative pose estimation for rgb-d scans via scene completion. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
72. Zhu, Y., Gordon, D., Kolve, E., Fox, D., Fei-Fei, L., Gupta, A., Mottaghi, R., Farhadi, A.: Visual Semantic Planning using Deep Successor Representations. In: Computer Vision, 2017 IEEE International Conference on (2017)
73. Zou, C., Colburn, A., Shan, Q., Hoiem, D.: Layoutnet: Reconstructing the 3d room layout from a single rgb image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2051–2059 (2018)