

Seeing the Un-Scene: Learning Amodal Semantic Maps for Room Navigation

Medhini Narasimhan^{1,2*} Erik Wijmans^{3,1} Xinlei Chen¹
Trevor Darrell² Dhruv Batra^{1,3} Devi Parikh^{1,3}
Amanpreet Singh¹

¹ Facebook AI Research ² University of California, Berkeley ³ Georgia Institute of Technology

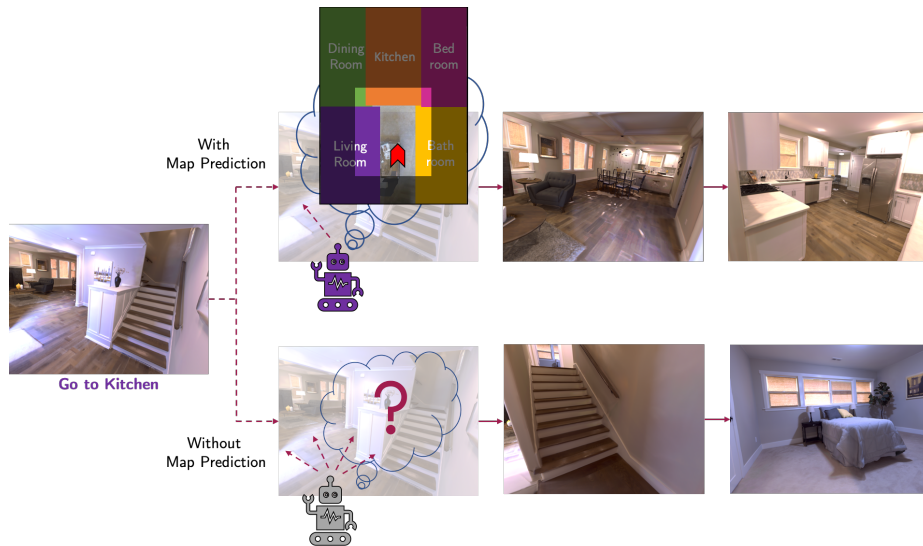


Fig. 1: We train an agent that can navigate to rooms in unseen environments by predicting amodal semantic maps. Here, the agent is placed in the entryway and is asked to navigate to the kitchen. It recognizes a dining room and predicts that the kitchen must be near, even though it has not seen it yet, and heads in that direction. It updates its semantic belief maps as it sees more of the house. Our agent learns to model architectural and stylistic regularities in houses, which helps it find its way in a novel house.

Abstract. We introduce a learning-based approach for room navigation using semantic maps. Our proposed architecture learns to predict top-down belief maps of regions that lie beyond the agent’s field of view while modeling architectural and stylistic regularities in houses. First, we train a model to generate amodal semantic top-down maps indicating beliefs of location, size, and shape of rooms by learning the underlying architectural patterns in houses. Next, we use these maps to predict a

* Work done while an intern at Facebook AI Research. Correspondence to medhini@berkeley.edu

point that lies in the target room and train a policy to navigate to the point. We empirically demonstrate that by predicting semantic maps, the model learns common correlations found in houses and generalizes to novel environments. We also demonstrate that reducing the task of room navigation to point navigation improves the performance further.

Keywords: Embodied AI, Room Navigation

1 Introduction

Humans have an uncanny ability to seamlessly navigate in unseen environments by quickly understanding their surroundings. Consider the example in Fig. 1. You’re visiting a friend’s home for the first time and you want to go to the kitchen. You’re in the entryway and you look around to observe a bedroom in one direction and the dining room in the other direction. A possible, but tedious solution is to head in a random direction and exhaustively search the space until you end up in the kitchen. Another option, and most probably the one you’d pick, is to walk towards the dining room as you are more likely to find the kitchen near the dining room rather than the bedroom. We believe that there are underlying architectural principles which govern the design of houses and our prior knowledge of house layouts helps us to navigate effectively in new spaces. We also improve our predictions of how the house is laid out as we walk around and gather more information. The goal of this work is to elicit a similar behaviour in embodied agents by enabling them to predict regions which lie beyond their field of view. The agent models correlations between the appearance and architectural layout of houses to efficiently navigate in unseen scenes.

Currently, there exist two paradigms for the navigation problem (1) Classical path planning based methods: SLAM-based approaches which first build a geometric map and then use path planning with localization for navigation [18, 25, 36]. (2) Learning-based methods: A policy is learned for a specific task in an environment [17, 46, 47]. In this work, we introduce a learning-based method that unlike previous approaches, predicts an intermediate representation that captures the model’s current belief of the *semantic layout of the house*, beyond the agent’s field of view. [43, 46] represent priors in the form of knowledge graphs or probabilistic relation graphs which capture room location priors. However, we also want to learn other correlations such as estimating the shape of a room by observing parts of it. We choose to model correlations as semantic maps indicating the location, shape, and size of rooms which lie beyond the agent’s field of view. This offers a more flexible representation.

In this work, we develop a novel technique to dynamically generate amodal semantic top-down maps of rooms by learning architectural regularities in houses and to use these predicted maps for room navigation. We define the task of room navigation as navigating to the nearest room of the specified type. For e.g., navigating to the bedroom closest to the starting point. We train an agent using supervision to predict regions on a map that lie beyond its field of view which forces it to develop beliefs about where a room might be present before

navigating to it. The agent constantly updates these beliefs as it steps around the environment. The learned beliefs help the agent navigate in novel environments.

Contributions (1) We introduce a novel learning-based approach for room navigation via amodal prediction of semantic maps. The agent learns architectural and stylistic regularities in houses to predict regions beyond its field of view. (2) Through carefully designed ablations, we show that our model trained to predict semantic maps as intermediate representations achieves better performance on unseen environments compared to a baseline which doesn't explicitly generate semantic top-down maps. (3) To evaluate our approach, we introduce the room navigation task and dataset in the Habitat platform [32].

2 Related Work

Navigation in mobile robotics. Conventional solutions to the navigation problem in robotics are comprised of two main steps: (1) mapping the environment and simultaneously localizing the agent in the generated map (2) path planning towards the target using the generated map. Geometric solutions to the mapping problem include (i) structure from motion and (ii) simultaneous localization and mapping (SLAM) [4, 7, 13, 16, 18, 36]. Various SLAM algorithms have been developed for different sensory inputs available to the agent. Using the generated map, a path can be computed to the target location via several path planning algorithms [25]. These approaches fall under the passive SLAM category where the maps are generated ahead of time. On the other hand, active SLAM research focuses on dynamically controlling the camera for building spatial representations of the environment. Some works formulate active SLAM as Partially Observable Markov Decision Process and use either Bayesian Optimization [27] or Reinforcement Learning [23] to plan trajectories that lead to accurate maps. [8] and [35] use Rao-Blackwellized Particle Filters to choose the set of actions that maximize the information gain and minimize the uncertainty of the predicted maps.

A less studied yet actively growing area of SLAM research is incorporating semantics into SLAM [6, 29, 37, 40]. [40] and [6] use semantic features for improved localization performance and for performing SLAM on dynamic scenes. [3] leverages conceptual spatial knowledge in the form of object co-occurrences for active object search. On the other hand, our work learns semantic maps of room locations for the task of Room Navigation. [11] summarizes the various ways of representing semantic information and using it for indoor navigation. However, all of these techniques rely on sensor data which is highly susceptible to noise [10] and have no mechanism to learn or update beliefs across environments. The limitations and open challenges in SLAM have been outlined in [7]. This motivates learning based methods for navigation, which we describe next.

Learning based methods for navigation. With the motivation of generalizing to novel environments and learning semantic cues, a number of end-to-end learning based approaches have been developed in the recent past [14, 17, 28, 31, 47]. [31] use a topological graph for navigation and [14] propose a memory based policy that uses attention to exploit spatio-temporal dependencies. [28] jointly learn

the goal-driven reinforcement learning problem with auxiliary depth prediction tasks. [42] introduce the RoomNav task in the House3D simulation platform and train a policy using deep deterministic policy gradient [26] to solve the same. [10] focus on building a task-agnostic exploration policy and demonstrate its usefulness for downstream navigation tasks. Most relevant to our work is Cognitive Mapping and Planning (CMP) [17]. It uses a differentiable mapper to learn an egocentric map of the environment and a differentiable planner that uses these maps alongside the goal to output navigational actions. These maps only indicate free space and contain no semantic information. On the other hand, we predict semantic top-down maps indicating the location, shape, and size of rooms in the house. Furthermore, unlike the maps in CMP which indicate regions within the agent’s field of view, our maps model architectural regularities in houses and contain unseen regions which lie beyond the agent’s field of view. [46] uses spatial and visual relationships between objects represented as a knowledge graph [24] for semantic navigation. The agent doesn’t learn new priors or update existing beliefs during training. [43] estimates priors as probabilistic relationship graphs over semantic entities and uses these graphs for planning. However, their graphs don’t capture information regarding size and shape of rooms and patterns in houses. Contrary to previous approaches, our work does not rely on pre-constructed maps or knowledge graphs representing priors. We dynamically learn amodal semantic belief maps which model architectural and stylistic regularities in houses. Further, the agent updates its beliefs as it moves around in an environment.

Vision-Language Navigation (VLN). A different but related task, language guided visual navigation was introduced by [2]. In VLN, an agent follows language instructions to reach a goal in a home. For *e.g.*, “Walk up the rest of the stairs, stop at the top of the stairs near the potted plant”. There are multiple works which attempt to solve this problem [15, 38, 39]. The room navigation task introduced here is different in that the agent doesn’t receive any language based instructions, just the final goal in the form of a room type. In [2], the agent “teleports” from one location to another on sparse pre-computed navigation graph and can never collide with anything. Our Room Navigation task starts from the significantly more realistic setting as in Point Navigation [32] where the agent takes the low-level actions such as move-forward (0.25m), turn-left/right (10 degrees) and learns to avoid collisions. Compared to VLN, the navigation in our task is significantly more challenging. We’d like to highlight that the methods developed for VLN aren’t directly applicable to room navigation as they all rely on intermediate goals in the form of language based instructions.

3 Room Navigation Task

The agent is spawned at a random starting location and orientation in a *novel* environment and is tasked with navigating to a given target room – *e.g.*, “Kitchen.” If there exist multiple rooms of same type, the agent needs to navigate to the room closest to its starting location. We ensure that the agent never starts in a room of the target room type *i.e.* if the agent is in a bedroom the target room




Fig. 2: **Architectural patterns in Matterport 3D.** All homes are scaled and aligned such that the kitchen is in the bottom left corner. The maps confirm the existence of structure in house layouts which can be leveraged by an agent trying to navigate to a room in a new environment.

cannot be a bedroom. Similar to [32], with each step the agent receives an RGB image from a single color vision sensor, depth information from the depth sensor and GPS+Compass that provides the current position and orientation *relative* to the start position. When there’s no GPS information available we only need egomotion, which most robotics platforms provide via IMU sensors or odometry. As in [32, 41], the agent does not have access to any ground truth floor plan map and must navigate using only its sensors. Unlike point navigation [32], room navigation is a semantic task, so GPS+Compass is insufficient to solve the task and only helps in preventing the agent from going around in circles.

4 Room Navigation using Amodal Semantic Maps

We develop a room navigation framework with an explicit mapping strategy to predict amodal semantic maps by learning underlying correlations in houses, and navigating using these maps. Our approach assumes there are architectural and stylistic regularities in houses which allow us to predict regions of a house which lie beyond our field of view. For instance, if we are in the kitchen we can guess that the dining room is adjacent to it and we would be right in most cases.¹ We first verify the existence of such correlations by scaling and aligning all the homes in the Matterport 3D [9] dataset such that the kitchen is in the bottom-left corner. As shown in Fig. 2, we can observe that the concentration of dining rooms is close to the kitchens and the bedrooms are in the opposite corner, away from the kitchens. We believe there exist similar and more subtle correlations (*e.g.* size of the kitchen could be indicative of the number of bedrooms in a house) which our agent can automatically learn while predicting amodal semantic top down maps of regions. We now provide an overview of our approach followed by a detailed explanation of each of its sub-components.

Overview. Our room navigation framework is outlined in Fig. 3. The agent is spawned in a random location and is asked to navigate to a specified target room, $tr \in \mathcal{R}$, where \mathcal{R} is the set of all possible target rooms. With each step t , the agent receives RGB image I_t , Depth D_t and GPS+Compass information. The

¹ We acknowledge that these regularities likely vary across geographies and cultures.

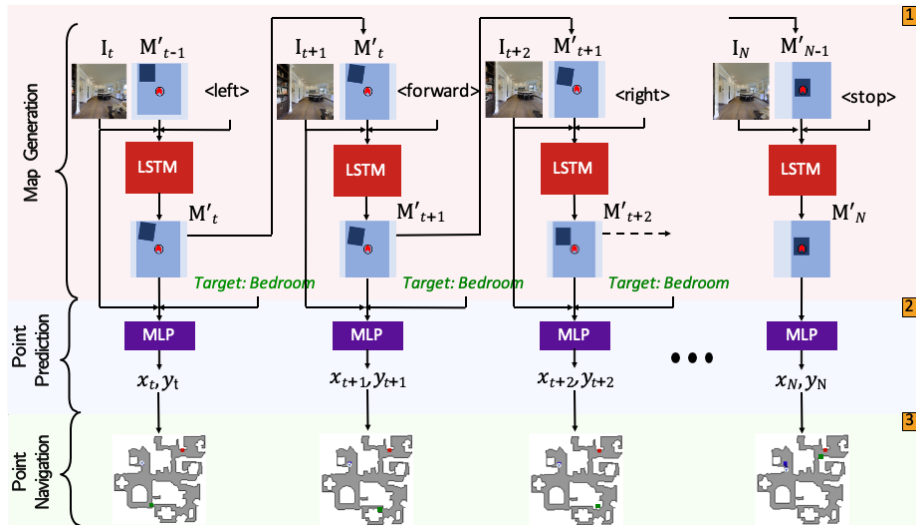


Fig. 3: **Room Navigation using Amodal Semantic Maps:** Our room navigation framework consists of 3 components (1) **Map Generation:** We predict egocentric crops of top down maps indicating rooms in a house. We use a sequence-to-sequence network which takes as input the current image I_t , the previous semantic map M'_{t-1} , and previous a_{t-1} and predicts the current semantic map M'_t . (2) **Point Prediction:** We feed the predicted maps (M'_t) along with I_t and target room ID tr to a network to predict a target point $P_t^{\text{pred}} = (x_t, y_t)$ that lies in the target room. The predicted and ground truth points are shown in green and red respectively on the top down map. (3) **Point Navigation:** We use a pre-trained point navigation policy to navigate to P_t^{pred} . The process is repeated for N steps until the agent calls stop or $N = 500$, whichever happens first.

agent predicts egocentric crops of top down semantic maps indicating the rooms which lie in and beyond its field of view. An example of these maps is shown in Fig. 4. To generate the maps, the agent uses a sequence to sequence (Seq2Seq) network which takes as input the RGB Image from the current time-step I_t , the predicted semantic maps from the previous time step $M'_{t-1,r}^{\text{pred}} \forall r \in \mathcal{R}$, and the previous action a_{t-1} and predicts semantic maps $M'_{t,r}^{\text{pred}}$ for the current time step. The predicted semantic maps $M'_{t,r}^{\text{pred}}$ are fed to a point prediction network which predicts a target point P_t^{pred} lying inside the target room. The agent navigates to the predicted point using a point navigation policy $\pi_{nav}(P_t^{\text{pred}}, D_t)$. The agent updates its beliefs of the predicted semantic maps and the predicted point as it steps around the environment. The episode is deemed successful if the agent stops inside the target room before 500 steps.

Next, we describe the three main components of our model architecture: Map Generation, Point Prediction and Point Navigation, as shown in Fig. 3.




Fig. 4: **Egocentric crops of Semantic Top Down Maps.** We show the regions in a $26m \times 26m$ crop of the ground truth semantic maps w.r.t. agent’s location and orientation. Agents are trained to predict these semantic maps, including regions they have not seen yet, based on regions of the house they have seen so far. Regions colored \blacksquare lie inside the specified room. \blacksquare denotes regions inside the house and not in the room and \square denotes regions outside the house.

4.1 Map Generation

We model the correlations in houses by learning to predict amodal semantic top-down maps. The agent uses the information it has gathered so far to determine where the different rooms in the house are present even before visiting these regions. The maps are crops of top-down maps indicating the type, location, size and shape of rooms and are egocentric to the agent. Fig. 4 shows an example of ground truth maps. The maps have 3 classes: (1) Regions which lie outside the house, (2) Regions which lie in the house but outside the target room, (3) Regions which lie in the target room. We hypothesize that the agent will learn architectural regularities in houses and correlations between the RGB images and layouts in order to predict these maps. We train a map generation network f_{map} , to predict egocentric crops of top down semantic maps for each room $r \in \mathcal{R}$. f_{map} consists of a sequence to sequence network f_{seq} and a decoder network f_{dec} . At each time step t , f_{seq} takes in as input a concatenation of learned representations of the current RGB frame $f_i(I_t)$, the previous action $f_{act}(a_{t-1})$, and the semantic map of the previous time step $f_m(M_{t-1,r})$. f_{seq} outputs a latent representation $h_{t,r}$. $h_{t,r}$ is passed through a parameterized decoder network f_{dec} that resembles the decoder in [30]. f_{dec} upsamples the latent representation using multiple transpose convolutions layers to produce output $M_{t,r}^{pred} \forall r \in \mathcal{R}$. Following [5], during training, we uniformly choose to set $M_{t,r}^{input}$ to be $M_{t-1,r}^{pred}$ (predicted map from previous time step) 50% of the time and $M_{t-1,r}^{GT}$ (ground truth map from the previous time step) the rest of the time. At each step, we use the ground truth semantic map $M_{t,r}^{GT}$ to train f_{seq} and f_{dec} with cross entropy loss, \mathcal{L}_{map} . Eqn. 1-5 describe the exact working of f_{map} . The agent continuously generates maps at each step until it calls `stop` or reaches end of the episode.

$$\mu \sim \text{Uniform}(0, 1) \quad (1)$$

$$M_{t,r}^{\text{input}} = \begin{cases} M_{t-1,r}^{\text{GT}}, & \text{if } \mu > 0.5 \text{ or } t = 0, \\ M_{t-1,r}^{\text{pred}}, & \text{otherwise} \end{cases} \quad (2)$$

$$h_{t,r} = f_{\text{seq}}(f_m(M_t^{\text{input}}), f_{\text{act}}(a_{t-1}), f_i(I_t)) \quad (3)$$

$$M_{t,r}^{\text{pred}} = f_{\text{dec}}(h_{t,r}) \quad (4)$$

$$\mathcal{L}_{\text{map}} = \sum_{r \in \mathcal{R}} \text{CrossEntropy}(M_{t,r}^{\text{pred}}, M_{t,r}^{\text{GT}}) \quad (5)$$

During inference, we feed a random image as the input semantic map for the first time step and use $M_{t,r}^{\text{pred}}$ as $M_{t,r}^{\text{input}}$ for all consecutive steps.

4.2 Point Prediction

The maps predicted by f_{map} are amodal – the agent predicts regions that it has not seen yet. They are however *crops* – the agent does not predict the layout of the entire house. These crops are egocentric to the agent and the target room may not always appear in these maps. For *e.g.*, consider Fig. 4, there exists a bathroom in the house but this region does not appear inside the crop as it does not fall inside the crop w.r.t. the agent’s current location. Inspired by the recent progress in point navigation [32, 41], we reduce the room navigation problem to point navigation. We train a network f_{point} to predict a target point $P_t^{\text{pred}} = (x'_t, y'_t)$ that lies in the target room tr , at each step t of the agent. Similar to Sec. 4.1, we learn representations $f_i(I_t)$ of the RGB image I_t , $f_m(M_{t,r}^{\text{pred}}) \forall r \in \mathcal{R}$ of the predicted semantic maps $M_{t,r}^{\text{pred}}$, and $f_{\text{emb}}(tr)$ which is a one-hot embedding of the target room ID tr . The predicted semantic map representations for the different rooms are combined using Eq. 6,

$$g_M = f_m(M_{t,r}^{\text{pred}}) \odot f_{\text{emb}}(tr) \quad (6)$$

where \odot represents element-wise multiplication. These are then concatenated with the target room ID tr and fed to a multilayer perceptron (MLP) f_{point} which outputs P_t^{pred} as described in Eqn. 7. f_{point} is trained using mean square loss w.r.t. a ground truth target point in the target room, P^{GT} , as shown in Eqn. 8. Sec. 5 describes how this point is chosen.

$$P_t^{\text{pred}} = (x'_t, y'_t) = f_{\text{point}}(g_M, f_i(I_t), f_{\text{emb}}(tr)) \quad (7)$$

$$\mathcal{L}_{\text{point}} = \text{MSELoss}(P_t^{\text{pred}}, P^{\text{GT}}) \quad (8)$$

During inference, the agent predicts a point every $k = 6$ steps. Once the agent completes $K = 60$ steps, the target point is simply fixed and no longer updated. The episode terminates if the agent calls `stop` or reaches $N = 500$ steps.

4.3 Point Navigation

At this stage, we have reduced room navigation to point navigation where the agent needs to navigate to the predicted target point P_t^{pred} . Following the

approach in [41], we train a point navigation policy using Proximal Policy Optimization (PPO) [34] on the dataset of point navigation episodes in [32]. The policy, described in Eqn. 9, is parameterized by a 2-layer LSTM with a 512-dimensional hidden state. It takes three inputs: the previous action a_{t-1} , the predicted target point P_t^{pred} , and an encoding of the Depth input $f_d(D_t)$. We only feed Depth input to the point navigation policy as this was found to work best [32]. The LSTM’s output is used to produce a softmax distribution over the action space and an estimate of the value function.

$$a_t = \pi_{nav}(P_t, f_d(D_t), a_{t-1}) \quad (9)$$

The agent receives terminal reward $r_T = 2.5$, and shaped reward $r_t(a_t) = -\Delta_{\text{geo.dist}} - 0.01$, where $\Delta_{\text{geo.dist}} = d_t - d_{t-1}$ is the change in geodesic distance to the goal by performing action a_t . We then use this pre-trained policy, π_{nav} , to navigate to the predicted point, P_t^{pred} . We also fine-tune π_{nav} on the points predicted by our model and this improves the performance.

To recap, f_{map} generates the semantic map of the space, f_{point} acts as a high-level policy and predicts a point, and the low level point navigation controller π_{nav} predicts actions to navigate to this point.

4.4 Implementation Details

Map Generation. The image and map representations, $f_i(I_t) \in \mathbb{R}^{256}$ and $f_m(M_{t,r}^{\text{input}}) \in \mathbb{R}^{256}$, are obtained by first embedding the input RGB image I_t and semantic map M_t^{input} using a ResNet50 [19] pre-trained on ImageNet [12] followed by a fully connected layer. The action representation $f_{act}(a_t) \in \{0, 1\}^{32}$ is a one-hot embedding of the action a_t .

All three are concatenated to form a 544 dimensional vector which is reduced to a vector of size 512 using a linear layer. This is fed through f_{seq} , a 2 layer LSTM [20] which outputs a hidden state $h_{t,r} \in \mathbb{R}^{512}$. This is passed through f_{dec} which consists of 5 transpose convolutions interleaved with BatchNorm [21] and followed by ReLU. The map generation network is trained using data collected from shortest path trajectories between source and target points in the house. We experimented with multiple semantic map crop sizes ranging from 20m to 40m and found 26m to work the best.

Point Prediction. The input RGB image I_t and semantic maps $M_{t,r}^{\text{input}}$ are embedded using a pre-trained ResNet50 [19], same as before, but $f_m(M_{t,r}^{\text{input}})$ is now a 32 dimensional vector using a fully-connected layer. The target room embedding $f_{emb}(tr) \in \mathbb{R}^{32}$ is a 32 dimensional one-hot encoding of the target room tr . The generated maps for each room type $f_m(M_{t,r}^{\text{input}})$ are multiplied with target embedding $f_{emb}(tr)$ and concatenated to form $g_M \in \mathbb{R}^{160}$. This is concatenated with the RGB image to form a 416 dimensional vector which is passed through f_{point} to obtain P_t^{pred} . We compute a goal vector relative to the agent’s current location which is used by the point navigation module. Similar to map generation, the point prediction network is also trained using data collected from a shortest path trajectories.

Point Navigation Policy. The Depth encoding $f_d(D_t)$ is based on ResNeXt [45] with the number of output channels at every layer reduced by half. As in [41], we replace every BatchNorm layer [21] with GroupNorm [44] to account for highly correlated inputs seen in on-policy RL. As in [41], we use PPO with Generalized Advantage Estimation (GAE) [33] to train the policy network. We set the discount factor γ to 0.99 and the GAE parameter τ to 0.95. Each worker collects (up to) 128 frames of experience from 4 agents running in parallel (all in different environments) and then performs 2 epochs of PPO with 2 mini-batches per epoch. We use Adam [22] with a learning rate of 2.5×10^{-4} . We use DD-PPO [41] to train 64 workers on 64 GPUs.

5 Room Navigation Dataset

Simulator and Datasets. We conduct our experiments in Habitat [32], a 3D simulation platform for embodied AI research. We introduce the room navigation task in the Habitat API and create a dataset of room navigation episodes using scenes from Matterport 3D [9]. We use Matterport 3D as it is equipped with room category and boundary annotations and hence is best suited for the task of Room Navigation. It consists of 61 scenes for training, 11 for validation, and 18 for testing. We only use the subset of 90 buildings which are houses and exclude others such as spas as those locations do not have common room categories with the majority of the dataset. We extract a subset of these scenes which contain at least one of the following room types: Bathroom, Bedroom, Dining Room, Kitchen, Living Room on the first floor. We only use the first floor of the house because, (1) In Matterport3d, the bounding boxes of rooms on different floors overlap at times, *e.g.* the box for a room on the first floor often overlaps with the room right above it on the second floor, making it hard to sample points which lie on the same floor, (2) The floors are uneven, making it difficult to distinguish between the different levels of the house.

Our dataset is comprised of 2.6 million episodes in 32 train houses, 200 episodes in 4 validation houses, and 500 episodes in 10 test houses.

Episode Specification. An episode starts with the agent being initialized at a starting position and orientation that are sampled at random from all navigable positions of the environment [32]. The target room is chosen from R if it is present in the house and is navigable from the starting position. We ensure the start position is not in the target room and has a geodesic distance of at least 4m and at most 45m from the target point in the room. During the episode, the agent is allowed to take up to 500 actions. After each action, the agent receives a set of observations from the active sensors. Statistics of the room navigation episodes can be found in the supplementary.

Evaluation Metric. Similar to [1], we design two evaluation metrics for room navigation - RoomNav Success weighted by (normalized inverse) Path Length (*RoomNav SPL*) and *Success*. An episode is considered a success if the agent stops 0.2m inside the bounds of the specified target room. We use 0.2m as the room boundaries in Matterport 3D sometimes lie outside the room and this

factor ensures the agent has indeed stepped inside the room. As shown in Fig. 5, we compute the geodesic distance from the source point to all the navigable points that lie 0.2m within the bounds of the room and choose a ground-truth target point P^{GT} that is closest to the agent’s start position, *i.e.* has the shortest geodesic distance. RoomNav SPL is similar to the SPL defined in [1]. Let S indicate ‘success’, l be the length of the shortest geodesic distance between start point and P^{GT} defined above, and p be the length of the agent’s path, then RoomNav SPL = $S \frac{l}{\max(l,p)}$. To achieve an SPL of 1, the agent must enter the nearest target room, and call stop when it has stepped 0.2m into the room.

Agent. As in [32], the agent is modeled as a cylinder with diameter 0.2m and height 1.5m. The actions and the sensors are the same as in [32].

6 Results

We design baselines to evaluate the effectiveness of each component of our proposed room navigation framework and to validate our approach. We also report oracle results using ground-truth annotations to establish an upper-bound on the scores that can be achieved by our model.

Table 1 shows the RoomNav SPL and Success scores on the room navigation validation and test sets (for selected baselines). Our room navigation framework described in Sec. 4 achieves an SPL of 0.31 on validation and 0.29 on the test set. Fine-tuning the point navigation policy on points predicted by the point prediction network improves the SPL to 0.35 on validation and 0.33 on test, making this our best performing model.

Vanilla Room Navigation Policy. Here we compare to an approach that does not use semantic maps to model correlations and does not use point navigation. We ablate both the map prediction and point generation components by training room navigation policy from scratch using PPO, similar to the point navigation policy in Sec. 4.3. Instead of a target co-ordinates relative to current state, it takes in the target room ID as input. The agent receives terminal reward $r_T = 2.5 \text{ RoomNav-SPL}$, and shaped reward $r_t(a_t) = -\Delta_{\text{geo.dist}} - 0.01$, where $\Delta_{\text{geo.dist}} = d_t - d_{t-1}$ is the change in geodesic distance to the target room by performing action a_t . The SPL using this baseline is 0.10 on validation and 0.10 on test, significantly worse compared to our approach (SPL 0.35). *This reinforces the effectiveness of our model, specifically the need to generate maps*




Fig. 5: **SPL for Room Navigation.** We compute the geodesic distance from source point to every point inside the room, within 0.2m of the bounds of the room. The point with the smallest geodesic distance is chosen as the *ground truth target point* and the SPL is computed w.r.t. this point.

Table 1: **RoomNav-SPL for our approach, baselines, and oracle methods on test and validation sets of the Room Nav Dataset.** Our proposed model (Map Generation + Point Prediction + PointNav + Fine-tune) achieves 0.35 RoomNav-SPL and outperforms all other baselines.

RoomNav Model	Validation		Test	
	SPL	Success	SPL	Success
Baseline: Random	0.00	0.00	0.00	0.00
Baseline: Vanilla Room Navigation Policy	0.10	0.15	0.10	0.11
Baseline: Map Generation + Room Navigation Policy	0.16	0.17	-	-
Baseline: Point prediction + PointNav	0.17	0.20	-	-
Baseline: Point prediction + PointNav + Fine-tune	0.21	0.23	-	-
Our: Map Generation + Point Prediction + PointNav	0.31	0.35	-	-
Our: Map Generation + Point Prediction + PointNav + Fine-tune	0.35	0.38	0.33	0.36
Oracle: GT Maps + Room Navigation Policy	0.54	0.56	-	-
Oracle: GT Maps + Point Prediction + PointNav	0.61	0.64	-	-
Oracle: GT Maps + Point Prediction + PointNav + Fine-tune	0.67	0.68	-	-
Oracle: GT Point Selection + PointNav	0.83	0.84	0.79	0.82

and use point navigation. Note that this approach mimics an approach that tries to solve room navigation via vanilla (“brute force”) reinforcement learning. **Vanilla Room Navigation Policy with Map Generation.** We ablate the point prediction model in Sec. 3 and train a room navigation policy to navigate to the target room using RGB images and semantic maps. We use the map generator to predict semantic maps for each room type. We then train a policy to navigate to the target room. The policy is similar to the room navigation policy described above and takes four inputs: the previous action, the target room represented as an ID, the predicted semantic maps for all rooms embedded as in Eq. 6 and the Depth encoding. It is trained the same way as the room navigation policy.

This baseline achieves an SPL of 0.16, which is worse by a large margin of 0.2 when compared to our best performing model (room navigation using Map Generation + Point Prediction). *The improved performance of our best method emphasizes the significance of the point prediction and point navigation modules in our best performing model.*

Point Prediction and Point Navigation Policy. We perform room navigation using only the high-level point prediction network and the low-level point navigation controller. We ablate the map generation module and train a modified version of the point prediction network without maps as input. Similar to f_{point} in Eq. 7, it generates $P_t = (x_t, y_t)$ but by using only the image representation $f_i(I_t)$ and target room embedding $f_{emb}(tr)$. The agent then navigates to P_t using the pre-trained point navigation policy as in Sec. 4.3.

This method achieves an SPL of 0.17 when the policy is trained from scratch and an SPL of 0.21 when the policy is fine tuned with points predicted by the point prediction network. Our best model surpasses this by a large margin of ~ 0.15 , which *shows the advantage of using supervision to learn amodal semantic maps that capture correlations.* It also indicates the effectiveness

Table 2: Ablation study of our mapping model and point prediction model.

Mapper	mIoU	Pixel Accuracy %			
		Class-1	Class-2	Class-3	Avg
CNN	25.66	31.84	36.97	5.63	24.81
LSTM(no maps)	32.92	41.44	48.92	10.40	33.59
LSTM(ours)	41.45	56.13	60.92	13.13	43.39

Error %	
Pred Maps + MLP	39.13
GT Maps + MLP	22.91

(a) **Map Generation Models Performance.** Map Generation is a three-way classification problem. Class-1 consists of regions lying outside the house, Class-2 is regions lying inside the house but not in the room and Class-3 is regions lying in the room.

(b) **Point Prediction error using Predicted Maps and Ground Truth Maps.** With ground truth maps, the prediction error is lower, suggesting scope for improvements.

of our map generation network. Since the environments in validation are different from train, we can also conclude that predicting semantic maps allows for better generalization to unseen environments as the RoomNav SPL is a direct indicator of how “quickly” an agent can reach a target room.

Using Ground Truth (GT) Maps. To get a better sense of how well our models can do if we had perfect map generation, we train a few of our baselines with ground truth maps instead of generated maps and report results in Table 1. With GT maps, Vanilla Room Navigation Policy with Map Generation achieves an SPL of 0.54. Adding GT maps to our best model, with and without fine-tuned point navigation policy we achieve SPL of 0.61 and 0.67 respectively. This suggests that there is still a large room for improvements in the Map Prediction module to perform room navigation more effectively. Table 2b reports the prediction error of point prediction model when using generated maps and ground truth maps.

Random. We evaluate using a random agent that takes action randomly among `move_forward`, `turn_left`, and `turn_right` with uniform distribution. It calls `stop` after 60 steps, which is the average length of an episode in our dataset. This achieves a RoomNav SPL of 0 on both test and validation, which implies that our task is difficult and cannot be solved by random walks.

Using GT Point Selection We use the pre-trained point navigation policy defined in Sec. 4.3 to navigate to the ground truth target points P^{GT} in the target room defined in Sec. 5. This achieves an SPL of 0.82 and 0.79 on validation and test respectively. It provides an “upper-bound” on the performance that can be achieved by the room navigation policy, as this indicates the maximum RoomNav SPL that can be achieved by the framework in Sec. 4 if the error on point prediction were 0. These numbers are comparable to the SPL values for point navigation on the Matterport-3D dataset in [32], thus indicating our episodes are at least as difficult as the point navigation episodes in [32].

Map Generation Ablations. We also experimented with different semantic map generation models. The results in Table 2a show that the LSTM map generation model described in Sec. 4 performs best with a mean Intersection-over-Union (mIoU) of 41.45 and an Average Pixel Accuracy of 43.39%. The CNN

only approach predicts a semantic map from each RGB image without maintaining a memory of the previous maps. This performs poorly and has a mIoU of 25.66 and an Average Pixel Accuracy of 24.81%. We train another LSTM model which doesn't use the semantic maps as input at each time step. This has a mIoU of 32.93 and Average Pixel Accuracy of 33.59.

Trajectory Videos. Qualitative results of our model can be found [here](#). The first image in the first row shows the RGB input. The second and third maps in the first row show the location of the agent in allocentric and egocentric views respectively. The last figure on the first row shows two dots, red indicating the ground truth target point in the target room and green showing the predicted point. When only one dot is visible it indicates the predicted and ground-truth points overlap. There are 5 ground-truth semantic maps for each of the 5 room types we consider. The labels at the bottom indicate the room type being predicted. The second row shows the ground truth semantic maps indicating the location of the rooms in the house. The third row shows the maps predicted by our agent. The target room is mentioned at the very bottom, in this case, "Dining Room". As seen in the video, the model dynamically updates the semantic belief maps and predicts the target point with high precision. The agent is able to detect the room its currently present in and also develop a belief of where other rooms lie. The RoomNav-SPL for this episode is 1.0 as the agent successfully reaches the target room following the shortest path. Additional videos can be found [here](#).

7 Conclusion

In this work, we proposed a novel learning-based approach for Room Navigation which models architectural and stylistic regularities in houses. Our approach consists of predicting the top down belief maps containing room semantics beyond the field of view of the agent, finding a point in the specified target room, and navigating to that point using a point navigation policy. Our model's improved performance (SPL) compared to the baselines confirms that learning to generate amodal semantic belief maps of room layouts improves room navigation performance in unseen environments. Our results using ground truth maps indicate that there is a large scope for improvement in room navigation performance by improving the intermediate map prediction step. We will make our code and dataset publicly available.

Acknowledgements. We thank Abhishek Kadian, Oleksandr Maksymets, and Manolis Savva for their help with Habitat, and Arun Mallya and Alexander Sax for feedback on the manuscript. The Georgia Tech effort was supported in part by NSF, AFRL, DARPA, ONR YIPs, ARO PECASE, Amazon. Prof. Darrell's group was supported in part by DoD, NSF, BAIR, and BDD. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the U.S. Government, or any sponsor.

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) 10, 11
2. 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
3. Aydemir, A., Göbelbecker, M., Pronobis, A., Sjö, K., Jensfelt, P.: Plan-based object search and exploration using semantic spatial knowledge in the real world. In: ECMR (2011) 3
4. Bailey, T., Durrant-Whyte, H.: Simultaneous localization and mapping (SLAM): Part ii. IEEE Robotics & Automation Magazine (2006) 3
5. Bengio, S., Vinyals, O., Jaitly, N., Shazeer, N.: Scheduled sampling for sequence prediction with recurrent neural networks. In: Advances in Neural Information Processing Systems (NeurIPS) 7
6. Bowman, S.L., Atanasov, N., Daniilidis, K., Pappas, G.J.: Probabilistic data association for semantic slam. In: International Conference on Robotics and Automation (ICRA) (2017) 3
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 (2016) 3
8. Carlone, L., Du, J., Ng, M.K., Bona, B., Indri, M.: Active slam and exploration with particle filters using kullback-leibler divergence. Journal of Intelligent & Robotic Systems (2014) 3
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. arXiv preprint arXiv:1709.06158 (2017), matterport3D dataset available at "https://niessner.github.io/Matterport/" 5, 10
10. Chen, T., Gupta, S., Gupta, A.: Learning exploration policies for navigation. arXiv preprint arXiv:1903.01959 (2019) 3, 4
11. Crespo, J., Castillo, J.C., Mozos, O.M., Barber, R.: Semantic information for robot navigation: A survey. Applied Sciences (2020) 3
12. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009) 9
13. Durrant-Whyte, H., Bailey, T.: Simultaneous localization and mapping: part i. IEEE Robotics & Automation Magazine (2006) 3
14. Fang, K., Toshev, A., Fei-Fei, L., Savarese, S.: Scene memory transformer for embodied agents in long-horizon tasks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 3
15. Fried, D., Hu, R., Cirik, V., Rohrbach, A., Andreas, J., Morency, L.P., Berg-Kirkpatrick, T., Saenko, K., Klein, D., Darrell, T.: Speaker-follower models for vision-and-language navigation. In: Advances in Neural Information Processing Systems (NeurIPS) (2018) 4
16. Fuentes-Pacheco, J., Ruiz-Ascencio, J., Rendón-Mancha, J.M.: Visual simultaneous localization and mapping: a survey. Artificial Intelligence Review (2015) 3
17. Gupta, S., Davidson, J., Levine, S., Sukthankar, R., Malik, J.: Cognitive mapping and planning for visual navigation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 2, 3, 4

18. Hartley, R., Zisserman, A.: Multiple view geometry in computer vision (2003) [2](#), [3](#)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016) [9](#)
20. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural computation* (1997) [9](#)
21. Ioffe, S., Szegedy, C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167 (2015) [9](#), [10](#)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014) [10](#)
23. Kollar, T., Roy, N.: Trajectory optimization using reinforcement learning for map exploration. *The International Journal of Robotics Research* (2008) [3](#)
24. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision (IJCV)* (2017) [4](#)
25. LaValle, S.M.: *Planning algorithms*. Cambridge University Press (2006) [2](#), [3](#)
26. Lillicrap, T.P., Hunt, J.J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D.: Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971 (2015) [4](#)
27. 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* (2009) [3](#)
28. Mirowski, P., Pascanu, R., Viola, F., Soyer, H., Ballard, A.J., Banino, A., Denil, M., Goroshin, R., Sifre, L., Kavukcuoglu, K., et al.: Learning to navigate in complex environments. arXiv preprint arXiv:1611.03673 (2016) [3](#)
29. Pronobis, A., Jensfelt, P.: Large-scale semantic mapping and reasoning with heterogeneous modalities. In: *International Conference on Robotics and Automation (ICRA)* (2012) [3](#)
30. 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 (MICCAI)* (2015) [7](#)
31. Savinov, N., Dosovitskiy, A., Koltun, V.: Semi-parametric topological memory for navigation. arXiv preprint arXiv:1803.00653 (2018) [3](#)
32. 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. arXiv preprint arXiv:1904.01201 (2019) [3](#), [4](#), [5](#), [8](#), [9](#), [10](#), [11](#), [13](#)
33. Schulman, J., Moritz, P., Levine, S., Jordan, M., Abbeel, P.: High-dimensional continuous control using generalized advantage estimation. arXiv preprint arXiv:1506.02438 (2015) [10](#)
34. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017) [9](#)
35. Stachniss, C., Grisetti, G., Burgard, W.: Information gain-based exploration using rao-blackwellized particle filters. In: *Robotics: Science and Systems* (2005) [3](#)
36. Thrun, S., Burgard, W., Fox, D.: *Probabilistic robotics*. MIT press (2005) [2](#), [3](#)
37. Walter, M.R., Hemachandra, S., Homberg, B., Tellex, S., Teller, S.: Learning semantic maps from natural language descriptions. In: *Robotics: Science and Systems* (2013) [3](#)
38. Wang, X., Huang, Q., Celikyilmaz, A., Gao, J., Shen, D., Wang, Y.F., Wang, W.Y., Zhang, L.: Reinforced cross-modal matching and self-supervised imitation

- learning for vision-language navigation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 4
39. Wang, X., Xiong, W., Wang, H., Yang Wang, W.: Look before you leap: Bridging model-free and model-based reinforcement learning for planned-ahead vision-and-language navigation. In: European Conference on Computer Vision (ECCV) (2018) 4
 40. Wang, Z., Zhang, Q., Li, J., Zhang, S., Liu, J.: A computationally efficient semantic slam solution for dynamic scenes. *Remote Sensing* (2019) 3
 41. 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. In: International Conference on Learning Representations (ICLR) (2020) 5, 8, 9, 10
 42. Wu, Y., Wu, Y., Gkioxari, G., Tian, Y.: Building generalizable agents with a realistic and rich 3d environment. arXiv preprint arXiv:1801.02209 (2018) 4
 43. Wu, Y., Wu, Y., Tamar, A., Russell, S., Gkioxari, G., Tian, Y.: Bayesian relational memory for semantic visual navigation. arXiv preprint arXiv:1909.04306 (2019) 2, 4
 44. Wu, Y., He, K.: Group normalization. In: European Conference on Computer Vision (ECCV) (2018) 10
 45. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 10
 46. Yang, W., Wang, X., Farhadi, A., Gupta, A., Mottaghi, R.: Visual semantic navigation using scene priors. arXiv preprint arXiv:1810.06543 (2018) 2, 4
 47. 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: International Conference on Robotics and Automation (ICRA) (2017) 2, 3