

Habitat-Web: Learning Embodied Object-Search Strategies from Human Demonstrations at Scale

Ram Ramrakhya¹ Eric Undersander² Dhruv Batra^{1,2} Abhishek Das²

¹Georgia Institute of Technology ²Meta AI Research

¹{ram.ramrakhya,dbatra}@gatech.edu ²{undersander,abshkdz}@fb.com

Abstract

We present a large-scale study of imitating human demonstrations on tasks that require a virtual robot to search for objects in new environments – (1) ObjectGoal Navigation (e.g. ‘find & go to a chair’) and (2) PICK&PLACE (e.g. ‘find mug, pick mug, find counter, place mug on counter’). First, we develop a virtual teleoperation data-collection infrastructure – connecting Habitat simulator running in a web browser to Amazon Mechanical Turk, allowing remote users to teleoperate virtual robots, safely and at scale. We collect 80k demonstrations for OBJECTNAV and 12k demonstrations for PICK&PLACE, which is an order of magnitude larger than existing human demonstration datasets in simulation or on real robots. Our virtual teleoperation data contains 29.3M actions, and is equivalent to 22.6k hours of real-world teleoperation time, and illustrates rich, diverse strategies for solving the tasks. Second, we use this data to answer the question – how does large-scale imitation learning (IL) (which has not been hitherto possible) compare to reinforcement learning (RL) (which is the status quo)? On OBJECTNAV, we find that IL (with no bells or whistles) using 70k human demonstrations outperforms RL using 240k agent-gathered trajectories. This effectively establishes an ‘exchange rate’ – a single human demonstration appears to be worth ~ 4 agent-gathered ones. More importantly, we find the IL-trained agent learns efficient object-search behavior from humans – it peeks into rooms, checks corners for small objects, turns in place to get a panoramic view – none of these are exhibited as prominently by the RL agent, and to induce these behaviors via contemporary RL techniques would require tedious reward engineering. Finally, accuracy vs. training data size plots show promising scaling behavior, suggesting that simply collecting more demonstrations is likely to advance the state of art further. On PICK&PLACE, the comparison is starker – IL agents achieve $\sim 18\%$ success on episodes with new object-receptacle locations when trained with 9.5k human demonstrations, while RL agents fail to get beyond 0%. Overall, our work provides compelling evidence for investing in large-scale imitation learning. Project page: ram81.github.io/projects/habitat-web.

1. Introduction

General-purpose robots that can perform a diverse set of embodied tasks in a diverse set of environments *have* to be good at visual exploration. Consider the canonical example of asking a household robot, ‘Where are my keys?’. To answer this (assuming the robot does not remember the answer from memory), the robot would have to search the house, often guided by intelligent priors – e.g. peeking into the washroom or kitchen might be sufficient to be reasonably sure the keys are not there, while exhaustively searching the living room might be much more important since keys are more likely to be there. While doing so, the robot has to internally keep track of where all it has been to avoid redundant search, and it might also have to interact with objects, e.g. check drawers and cabinets in the living room (but not those in the washroom or kitchen!).

This example illustrates fairly sophisticated exploration, involving a careful interplay of various implicit objectives (semantic priors, exhaustive search, efficient navigation, interaction, etc.). Many recent tasks of interest in the embodied AI community – e.g. ObjectGoal Navigation [1, 2], rearrangement [3, 4], language-guided navigation [5, 6] and interaction [7], question answering [8–12] – involve some flavor of this visual exploration. With careful reward engineering, reinforcement learning (RL) approaches to these tasks have achieved commendable success [13–17]. However, engineering the ‘right’ reward function so that the learned policy exhibits desired behavior is unintuitive and frustrating (even for domain experts), expensive (requiring multiple rounds of retraining under different rewards), and not scalable to new tasks or behaviors. For complex tasks (e.g. object rearrangement or tasks specified in open-ended natural language), RL from scratch may not even get off the ground.

In this work, we advance the alternative research agenda of imitation learning [18] – i.e. collecting a large dataset of human demonstrations (that implicitly capture intelligent behavior we wish to impart to our agents) and learning policies directly from these human demonstrations.

First, we develop a safe scalable virtual teleoperation data-collection infrastructure – connecting the Habitat simulator

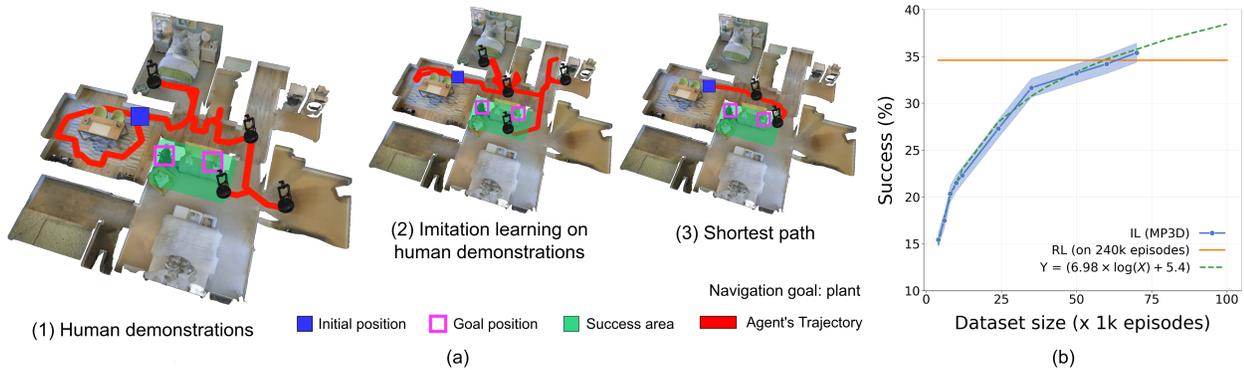


Figure 1. a) Example OBJECTNAV 1) human demonstration, 2) agent trained on human demonstrations, and 3) shortest path. Notice how humans demonstrate sophisticated exploration behavior to succeed at this task in unseen environments, which is hard to engineer into the right reward for an RL agent and is unlikely to be captured in shortest path demonstrations. An agent trained on human demonstrations learns this exploration and object-search behavior. b) Success on the OBJECTNAV MP3D-VAL split vs. no. of human demonstrations for training.

running in a browser to Amazon Mechanical Turk (AMT). We develop this in way that enables collecting human demonstrations for a variety of tasks being studied within the Habitat [19, 20] ecosystem (e.g. PointNav [2], OBJECTNAV [1, 2], ImageNav [21], VLN-CE [6], MultiON [22], etc.).

We use this infrastructure to collect human demonstration datasets for 2 tasks requiring visual search – 1) ObjectGoal Navigation (e.g. ‘find & go to a chair’) and 2) PICK&PLACE (e.g. ‘find mug, pick mug, find counter, place on counter’). In total we collect 92k human demonstrations, 80k demonstrations for OBJECTNAV and 12k demonstrations for PICK&PLACE. In contrast, the largest existing datasets have 3-10k human demonstrations in simulation [23–25] or on real robots [26, 27], an order of magnitude smaller. This virtual teleoperation data contains 29.3M actions, which is equivalent to 22, 600 hours of real-world teleoperation time assuming a LoCoBot motion model from [28] (details in appendix (Sec. A.3)). The first thing this data provides is a ‘human baseline’ with sufficiently tight error-bars to be taken seriously. On the OBJECTNAV validation split, humans achieve $93.7 \pm 0.1\%$ success and $42.5 \pm 0.5\%$ Success Weighted by Path Length (SPL) [2] (vs. 34.6% success and 7.9% SPL for the 2021 Habitat ObjectNav Challenge winner [15]). The success rate (93.7%) suggests that this task is largely doable for humans (but not 100%). The SPL (42.5%) suggests that even humans need to explore significantly.

Beyond scale, the data is also rich and diverse in the strategies that humans use to solve the tasks. Fig. 1 shows an example trajectory of an AMT user controlling a LoCoBot looking for a ‘plant’ in a new house – notice the peeking into rooms, looping around the dining table – all of which is (understandably) absent from the shortest path to the goal.

We use this data to answer the question – how does large-scale imitation learning (IL) (which has not been hitherto possible) compare to large-scale reinforcement learning (RL) (which is the status quo)? On OBJECTNAV, we find that IL

(with no bells or whistles) using only 70k human demonstrations outperforms RL using 240k agent-gathered trajectories. This effectively establishes an ‘exchange rate’ – a single human demonstration appears to be worth ~ 4 agent-gathered ones. More importantly, we find the IL-trained agent learns *efficient object-search behavior* – as shown in Fig. 1 and Sec. 7. The IL agent learns to mimic human behavior of peeking into rooms, checking corners for small objects, turning in place to get a panoramic view – none of these are exhibited as prominently by the RL agent. Finally, the accuracy vs. training-data-size plot (Fig. 1b) shows promising scaling behavior, suggesting that simply collecting more demonstrations is likely to advance the state of art further. On PICK&PLACE, the comparison is even starker – IL-agents achieve $\sim 18\%$ success on episodes with new object-receptacle locations when trained with 9.5k human demonstrations, while RL agents fail to get beyond 0%.

On both tasks, we find that demonstrations from humans are essential; imitating shortest paths from an oracle produces neither accuracy nor the strategic search behavior. In hindsight, this is perfectly understandable – shortest paths (e.g. Fig. 1(a3)) do not contain any exploration but the task requires the agent to explore. Essentially, a shortest path is inimitable, but imitation learning is invaluable. Overall, our work provides compelling evidence for investing in large-scale imitation learning of human demonstrations.

2. Related work

Embodied Demonstrations from Humans. Prior expert demonstration datasets for embodied tasks combining vision and action (and optionally language) can be broadly categorized into either consisting of shortest-path trajectories from a planner with privileged information [5, 7, 8, 29], or consisting of human-provided trajectories [23–25]. While some works in the former collect natural language data from hu-

mans [5, 7], we contend that collecting navigation data from humans is equally crucial. Datasets with human-provided navigation trajectories are typically small. TEACH [23], CVDN [24] and WAY [25] have $<10k$ episodes, while the EmbodiedQA [8] dataset has ~ 700 human-provided episodes – all prohibitively small for training proficient agents. A key contribution of our work is a scalable web-based infrastructure for collecting human navigation and interaction demonstrations, that is easily extensible to *any* task situated in the Habitat [19] simulator, including language-based tasks. We have collected $\sim 13x$ more demonstrations (in total $92k$) compared to prior works.

Exploration. Learning how to explore an environment to gather sufficient information for use in downstream tasks has a rich history [30]. Curiosity-based approaches typically use reinforcement learning to maximize intrinsic rewards that capture the surprise or state prediction error of the agent [31–33]. State visitation count rewards are also popular for learning exploration [34, 35]. We refer the reader to Ramakrishnan *et al.* [36] for a review of exploration objectives for embodied agents. For improving exploration in OBJECTNAV specifically, SemExp [17] made use of a modular policy for semantic mapping and path planning, Ye *et al.* [15] used time-decaying state visitation count reward, and Maksymets *et al.* [16] used area coverage reward.

Most relatedly, Chen *et al.* [37] used ~ 700 human navigation trajectories from the EmbodiedQA dataset [8] (ignoring the questions) to learn task-independent exploration using imitation learning. We likewise train agents via imitation learning on human demonstrations, but rather than encouraging task-agnostic exploration, we consider human demonstrations to be a rich *task-specific* mix of exploration and efficient navigation, that simple architectures without explicit mapping and planning modules can be trained on.

3. Habitat-WebGL Infrastructure

To be able to train agents via imitation learning on human demonstrations, we first need a reliable pipeline to collect human demonstrations at scale. To this end, we develop a web-based setup to connect the Habitat simulator [19, 20] to AMT users, building on the work of Newman *et al.* [38].

Interface. Fig. 2 shows a screenshot of the interface an AMT user interacts with to complete a data collection task. This web application renders assets from Habitat-Sim running on the user’s browser via WebGL. All data collection in this work was done in Matterport3D [39] scans, but any Habitat-compatible asset may be used in future. Users can see the agent’s first-person RGB view, and can move around and grab / release objects using keyboard controls. On the task page, users are provided an instruction and details about keyboard controls to complete the task. For OBJECTNAV, we provide an instruction of the form ‘*Find and go to the <goal_object_category>*’. For tasks requiring interaction

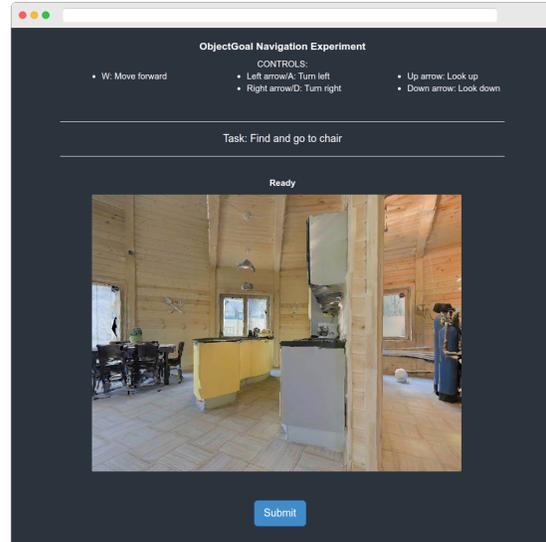


Figure 2. Screenshot of our Amazon Mechanical Turk interface for collecting OBJECTNAV demonstrations. Users are provided the agent’s first-person view of the environment and an instruction such as “Find and go to chair”. They can make the agent look around and move in the environment via keyboard controls, and can submit the task upon successful navigation by clicking ‘Submit’.

with objects (*e.g.* PICK&PLACE), we highlight the object under the user’s gaze by drawing a 3D bounding box around it (pointed to by a crosshair as in video games). In our initial pilots, we found this to improve user experience when grabbing objects instead of users having to guess when objects are available to be picked up. When an object is successfully grabbed, it disappears from the first-person view and immediately appears in the ‘inventory’ area on the task interface. When a grabbed object is released, it is dropped at the center of the user’s screen where the crosshair would be pointing to. If the crosshair points to a distance, the object is dropped on the floor from a height at a distance of $1m$ from the agent’s location. Upon completion, users submit the task by clicking ‘Submit’. At this point, the sequence of keyboard actions, agent, and object states are recorded in our backend server.

Habitat simulator and PsiTurk. Our Habitat-WebGL application is developed in Javascript, and allows us to access all C++ simulator APIs through Javascript bindings. This lets us use the full set of simulation features available in Habitat. To simulate physics, we use the physics APIs from Habitat 2.0 [20], including rigid body dynamics support (C++ APIs exposed as Javascript bindings). Our interface executes actions entered by users every 50ms (rendering 20 frames per second) and then steps physics for 50ms in the simulator. All of our tasks on AMT are served using PsiTurk and an NGINX reverse proxy, and all data stored in a MySQL database. We use PsiTurk to manage the tasks as it provides us with useful helper functions to log task-related metadata, as well as launch and approve tasks.

See Section A.6 for details on how we validate human-submitted AMT tasks and ensure data quality.

4. Tasks and Datasets

Using our web infrastructure, we collect demonstration datasets for two embodied tasks – OBJECTNAV [1, 2] and PICK&PLACE, an instantiation of object rearrangement [3].

4.1. ObjectGoal Navigation

In the ObjectGoal Navigation (OBJECTNAV) task, an agent is tasked with navigating to an instance of a specified object category (*e.g.* ‘chair’) in an unseen environment. The agent does not have access to a map of the environment and must navigate using an RGBD camera and a GPS+Compass sensor which provides location and orientation information relative to the start of the episode. The agent also receives the goal object category ID as input. The full action space is discrete and consists of MOVE_FORWARD (0.25m), TURN_LEFT (30°), TURN_RIGHT (30°), LOOK_UP (30°), LOOK_DOWN (30°), and STOP actions. For the episode to be considered successful, the agent must stop within 1m Euclidean distance of the goal object within a maximum of 500 steps and be able to turn to view the object from that end position [40].

Human Demonstrations (OBJECTNAV-HD). We collect 70k demonstrations on the 56 training scenes from Matterport3D [39] following the standard splits defined in [2, 39]. For each scene, we collect ~ 59 demonstration episodes for each unique goal object category with a randomly set start location of the human demonstrator for each episode. This amounts to an average of ~ 1250 demonstrations per scene. Additionally, we collect 10k demonstrations on 25 training scenes from Gibson [41]. For each Gibson scene, we collect ~ 66 demonstration episodes for each unique goal object category. This amounts to ~ 396 demonstrations per scene. Similar to when training artificial agents, humans can view first-person RGB on the task interface, but unlike artificial agents, humans do not get access to Depth and GPS+Compass. We assume humans are sufficiently proficient at inferring depth and odometry from vision, to the extent required to accomplish the goal. In total, we collect 80k OBJECTNAV demonstrations amounting to $\sim 19.5M$ steps of experience, each episode averaging 243 steps.

Shortest Path Demonstrations. To compare against prior embodied datasets of shortest paths [5, 7, 8, 29] and to demonstrate the unique advantage of human demonstrations, we also generate a dataset of shortest paths. The analysis in this section was performed on a subset of 35k demonstrations of OBJECTNAV-HD (collected in first phase). These demonstrations are generated by greedily fitting actions to follow the geodesic shortest path to the nearest navigable goal object viewpoint. Since shortest paths are (by design) shorter than human demonstrations (average 67 vs. 243 steps per demonstration), we compensate by generating a larger

number of shortest paths to roughly match the steps with 35k human demonstrations (7.6M steps from 114k shortest paths vs. 8.4M steps from 35k human demonstrations).

Analysis. Table 8a reports statistics of our human and shortest path demonstration datasets. Recall that an episode is considered a failure if the target object is not found within 500 navigation steps. Under this definition, humans fail on 11.1% training set episodes; they fail on 0% episodes if we relax the step-limit. Surprisingly, SPL for humans is 39.9% for training split episodes, significantly lower than 94.9% for shortest paths underscoring the difficulty in searching for objects in in unseen environments.

We additionally report two metrics to demonstrate that the OBJECTNAV task requires significant exploration. Occupancy Coverage (OC) measures percentage of total area covered by the agent when navigating. To compute OC, we first divide the map into voxel grids of $2.5m \times 2.5m \times 2.5m$ and increment a counter for each visited voxel. Sight Coverage (SC) measures the percentage of total navigable area visible to the agent in its field of view (FOV) during an episode. To compute SC, we project a mask on the top-down map of the environment using the agent’s FOV, that is iteratively updated at every step to update the area seen by the agent. OC and SC metrics for human demonstrations show that humans traverse 3-4x and observe 2x the area of the environment when performing this task compared to shortest paths.

Fig. 8b,c show episode length and action histograms for human and shortest path demonstrations. Human demonstrations are longer (average ~ 243 vs. ~ 67 steps per demonstration) and have a slightly more uniform action distribution.

4.2. Object Rearrangement – PICK&PLACE

In the pick-and-place task (PICK&PLACE), an agent must follow an instruction of the form ‘Place the <object> on the <receptacle>’, without being told the location of the <object> or <receptacle> in a new environment. The agent must explore and navigate to the object, pick it up, explore and navigate to the receptacle, and place the previously picked-up object on it. Similar to OBJECTNAV, agents are not equipped with a map of the environment, and only have access to an RGBD camera and a GPS+compass sensor. At a high level, PICK&PLACE can be thought of as a natural extension of OBJECTNAV, performing it twice in the same episode – once to find the specified object and again to find the specified receptacle – delimited by grab and release actions. For object interaction, we use the ‘magic pointer’ abstraction defined in [3]. If the agent is not holding any object, the grab/release action will pick the object pointed to by its crosshair (at the center of its viewpoint) if within 1.5m of the object. If the agent is already holding an object, the grab/release action will drop the object at the crosshair location. If there is no drop-off point within 1.5m in the direction of the crosshair, the object will be dropped on the floor 1m

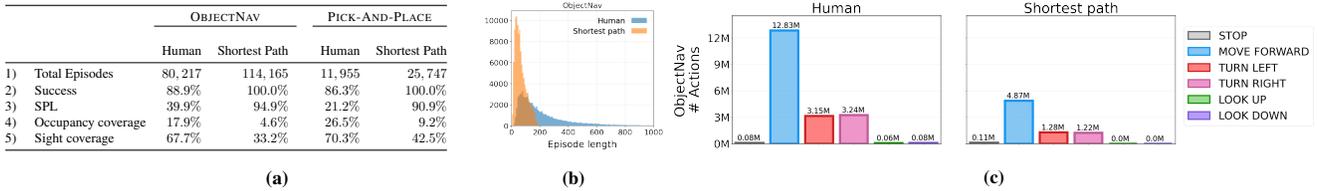


Figure 3. (a) Dataset statistics for human demonstrations vs. shortest paths for OBJECTNAV and PICK&PLACE. Coverage metrics are computed on subset of 1000 episodes. (b) Comparison of episode lengths and action histograms for human demonstrations vs. shortest paths. Human demonstrations are longer and have a more uniform action distribution than shortest paths.

in front of the agent. The full action space is discrete and consists of MOVE_FORWARD ($0.15m$), MOVE_BACKWARD ($0.15m$), TURN_LEFT (5°), TURN_RIGHT (5°), LOOK_UP (5°), LOOK_DOWN (5°), GRAB_RELEASE, NO_OP (step physics $50ms$), and STOP. For the episode to be considered successful, the agent must place the object on top of the receptacle – *i.e.* the object center should be at a height greater than the receptacle center, and within $0.7m$ of the receptacle object center – within 1500 steps. We picked this $0.7m$ threshold distance between the object and receptacle based on pilots on AMT. $0.7m$ was sufficiently strict for avoiding false positives in the collected demonstrations where users are able to submit the task without necessarily placing the object on top of the receptacle.

Human Demonstrations (PICK&PLACE-HD). We collect human demonstrations for PICK&PLACE on 9 scenes from Matterport3D [39]. In each episode, objects and receptacles are instantiated by randomly sampling from 457 possible object-receptacle pairs. We initialize the object and receptacle at randomly sampled locations in the environment, and collect 3 demonstrations for each object-receptacle pair. The agent, object, and receptacle locations are randomized across all episodes (including the 3 we collect for each object-receptacle pair). In total, we have 457×3 unique object-receptacle-agent position initializations per scene, amounting to $457 \times 3 \times 9 = \sim 12k$ demonstrations, which is $\sim 11.5M$ steps in experience, each episode averaging 932 steps.

Shortest Path Demonstrations. Similar to OBJECTNAV, we generate shortest path demonstrations for PICK&PLACE. These demonstrations are generated by first using the geodesic shortest-path follower to the object, then using a heuristic action planner to face and pick up the object, then following the geodesic shortest-path to the receptacle, and again using a heuristic action planner to drop the object on the receptacle. We generated 25.7k shortest path demonstrations for PICK&PLACE, each averaging 342 steps, amounting to a total of ~ 8.8 million steps of experience.

Analysis. Table 8a reports statistics for human and shortest path demonstrations. Similar to OBJECTNAV, humans have significantly lower SPL, and 2x higher occupancy and sight coverage compared to shortest paths, suggesting the need for exploration. Comparing episode lengths and action histograms (see appendix (Sec. A.1.1) for figure), human demonstrations are longer and make use of all 9 actions. Interestingly, humans often use the MOVE_BACKWARD action to backtrack, which the shortest path agents do not use

(by design), instead of turning 180° and moving forward. This behavior does not appear in OBJECTNAV shortest path demonstrations because there is just one target object, and so the geodesic shortest path would never involve backtracking or making 180° turns.

5. Imitation Learning from Demonstrations

We use behavior cloning to learn a policy from demonstrations. Let $\pi_\theta(a_t | o_t)$ denote a policy parametrized by θ that maps observations o_t to a distribution over actions a_t . Let τ denote a trajectory consisting of state, observation, human action tuples: $\tau = (s_0, o_0, a_0, \dots, s_T, o_T, a_T)$ and $\mathcal{T} = \{\tau^{(i)}\}_{i=1}^N$ denote a dataset of human demonstrations. The learning problem can be summarized as:

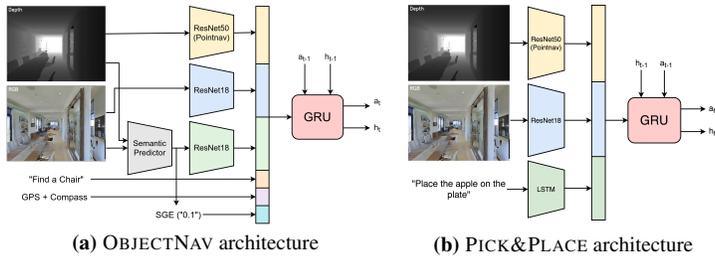
$$\theta^* = \arg \min_{\theta} \sum_{i=1}^N \sum_{(o_t, a_t) \in \tau^{(i)}} -\log(\pi_\theta(a_t | o_t)) \quad (1)$$

Inflection weighting introduced in Wijmans *et al.* [9], adjusts the loss function to upweight timesteps where actions change (*i.e.* $a_{t-1} \neq a_t$). Specifically, the inflection weighting loss coefficient is computed as total no. of actions in the dataset divided by the total no. of inflection points, and this coefficient is multiplied with the loss at each inflection timestep where $a_{t-1} \neq a_t$. This approach was found to be useful for tasks like navigation with long sequences of the same actions, *e.g.* several ‘forward’ actions when navigating corridors [9]. We use inflection weighting in all our experiments and found it to help over vanilla behavior cloning.

Our **base policy** is a simple CNN+RNN architecture. We first embed all sensory inputs using feed-forward modules. For RGB, we use a randomly initialized ResNet18 [42]. For depth, we use a ResNet50 that was pretrained on PointGoal navigation using DD-PPO [13]. Then these RGB and depth features (and optionally other task-specific features) are concatenated and fed into a GRU [43] to predict a distribution over actions a_{t+1} . Task-specific architectural choices over this base policy are described in the next sections.

5.1. OBJECTNAV

Fig. 4a shows our OBJECTNAV architecture. Similar to Anand *et al.* [44], we feed in RGBD inputs of size 640×480 passed through a 2×2 -AVGPOOL layer to reduce the resolution (performing low-pass filtering + downsampling). The agent also has a GPS+Compass sensor, which provides location and orientation relative to start of the episode.



Method	Success (†)	SPL (†)
1) RL (ExploreTillSeen) [16]	20.0%	6.5%
2) RL (ExploreTillSeen + THDA) [16]	28.4%	11.0%
3) RL (Red Rabbit) [15]	34.6%	7.9%
4) IL w/ Shortest Paths	4.4%	2.2%
5) IL w/ 35k Human Demos (similar #steps as row 4)	31.6%	8.5%
6) IL w/ 50k Human Demos	32.4%	9.1%
7) IL w/ 50k Human Demos (includes 10k THDA [16])	33.2%	9.5%
8) IL w/ 70k Human Demos	35.4%	10.2%
9) IL w/ 70k Human Demos (includes 10k Gibson [41])	33.9%	9.7%
10) IL w/ 80k Human Demos	33.8%	9.9%
11) Humans	93.7%	42.5%

(c) OBJECTNAV results on MP3D-VAL

Figure 4. Our policy architectures for a) OBJECTNAV and b) PICK&PLACE. Both are simple CNN+RNN networks that embed and concatenate all sensory inputs, which are then fed into a GRU to predict actions. c) OBJECTNAV results on the MP3D VAL split [2, 39].

GPS+Compass inputs are passed through fully-connected layers to embed them to 32-d vectors. In addition to RGBD and GPS+Compass, following Ye *et al.* [15], we use two additional semantic features – semantic segmentation (SemSeg) of the input RGB and a ‘Semantic Goal Exists’ (SGE) scalar which is the fraction of the visual input occupied by the goal category. These semantic features are computed using a pretrained and frozen RedNet [45] that was pretrained on SUN RGB-D [46] and finetuned on 100k randomly sampled front-facing views rendered in the Habitat simulator. Finally, we also feed in the object goal category embedded into a 32-d vector. All of these input features are concatenated to form an observation embedding, and fed into a 2-layer, 512-d GRU at every timestep. We train this policy for $\sim 400M$ steps ($= \sim 21$ epochs on $\sim 70k$ demonstration episodes). We evaluate checkpoints at every $\sim 15M$ steps for the last 50M steps of training, and report metrics for checkpoints with the highest success on the validation split.

5.2. PICK&PLACE

Fig. 4b shows our PICK&PLACE architecture. We feed in RGBD inputs of size 256×256 . In addition to RGBD observations, the policy gets as input language instructions of the form ‘Place the <object> on the <receptacle>’ encoded using a single-layer LSTM [47]. RGBD and instruction features are concatenated to form an observation embedding, which is fed into a 2-layer, 512-d GRU at every timestep. We train this policy for $\sim 90M$ steps ($= \sim 10$ epochs on $\sim 9.5k$ demonstration episodes). We evaluate checkpoints at every $\sim 10M$ steps during training, and report metrics for checkpoints with the highest success on the validation split.

6. Experiments & Results

6.1. OBJECTNAV

Table 4c reports results on the MP3D VAL split for several baselines. First, we compare our approach with two state-of-the-art RL approaches from prior work. Maksymets *et al.* [16] (row 1) train their policy using a reward structure that breaks OBJECTNAV into two subtasks – exploration and direct navigation to goal object once it is spotted. This

agent gets a positive reward for maximizing area coverage until it sees the goal object. It then receives a navigation reward to minimize distance-to-object. This policy achieves 20.0% success and 6.5% SPL (row 1), which is 15.4% worse on success and 3.7% worse on SPL compared to behavior cloning on 70k human demonstrations (row 8). [16] then combine this reward structure with Treasure Hunt Data Augmentation (THDA) – inserting arbitrary 3D target objects in the scene to augment the set of training episodes. With THDA, this achieves 28.4% success and 11.0% SPL (row 2), 7.0% worse and 0.8% better respectively than our approach (row 8). Ye *et al.* [15] (row 3) train their policy with a combination of exploration and distance-based navigation rewards, and their representations with several auxiliary tasks (*e.g.* inverse dynamics and predicting map coverage). This achieves 34.6% success and 7.9% SPL (row 3), which is 0.8% worse on success and 2.3% worse on SPL than our approach (row 8). IL on a dataset of shortest paths achieves 4.4% success and 2.2% SPL (row 4), significantly worse than training on 35k human demonstrations (31.6% success, 8.5% SPL). Recall that comparison of shortest path demonstrations was done with a subset of 35k OBJECTNAV-HD demonstrations that were collected in the first phase of the project. Next, we also collected 10k human demonstrations on OBJECTNAV episodes generated in THDA fashion – *i.e.* asking humans to find randomly inserted objects. Notice that this involves pure exhaustive search, since there are no semantic priors that humans can leverage in this setting. An IL agent trained on 10k THDA demonstrations combined with the original 40k demonstrations achieves 33.2% success and 9.5% SPL (row 7) which is 0.8% better on success and 0.4% better on SPL than 50k non-THDA demonstrations (row 6), *i.e.* adding these THDA demonstrations with exhaustive search behavior helps. We also collected 10k demonstrations on Gibson [41] OBJECTNAV episodes to compare effect of different scene datasets. An agent trained on 10k Gibson demonstrations combined with 60k MP3D demonstrations achieves 33.9% success and 9.7% SPL (row 9), which is 1.5% worse on success and 0.5% worse on SPL compared to when we use MP3D-only demonstrations (row 8).

Finally, we also benchmark human performance on the

Method	Success (\uparrow)	SPL (\uparrow)
1) IL w/ Vision	0.0%	0.0%
2) IL w/ Semantic Input	22.7%	6.1%
3) IL w/ RGBD + Semantic Input	31.6%	8.5%

Table 1. ObjectNav ablation results on the MP3D VAL split [2, 39].

MP3D VAL split – 93.7% success, 42.5% SPL (row 11).

ObjectNav Sensor Ablations. Table 1 reports results on the MP3D VAL split for various ablations of our approach trained on 35k human demonstrations. First, without any visual input (row 1), *i.e.* no RGBD and semantic inputs, the agent fails to learn anything (0% success, 0% SPL). Second, without SemSeg and SGE features (and keeping only RGB and Depth features) to the policy, performance drops by 8.9% success and 2.4% SPL (row 2 vs. 3).

Habitat ObjectNav Challenge Results. Table 2 compares our results with prior approaches from the 2020 and 2021 Habitat Challenge leaderboards. Our approach (IL w/ 70k demonstrations) achieves 27.8% success and 9.9% SPL (row 8), outperforming prior RL-trained counterparts – 3.3% better success, 3.5% better SPL than Red Rabbit (6-Act Base) [15] (row 5), and 6.7% better success, 1.1% better SPL than ExploreTillSeen + THDA [16] (row 7).

Team / Method	Success (\uparrow)	SPL (\uparrow)
1) DD-PPO baseline [13, 15]	6.2%	2.1%
2) Active Exploration (Pre-explore)	8.9%	4.1%
3) SRCB-robot-sudoer	14.4%	7.5%
4) SemExp [48]	17.9%	7.1%
5) Red Rabbit (6-Act Base) [15]	24.5%	6.4%
6) Red Rabbit (6-Act Tether) [15]	21.1%	8.1%
7) ExploreTillSeen + THDA [16]	21.1%	8.8%
8) IL w/ 70k Human Demos	27.8%	9.9%

Table 2. Results on Habitat ObjectNav Challenge TEST-STD [49].

Performance vs. Dataset size. To investigate scaling behavior, we plot VAL success against the size of the human demonstrations dataset in Fig. 1b. We created splits of the human demonstrations’ dataset of increasing sizes, from 4k to 70k, and trained models with the same set of hyperparameters on each split. All hyperparameters were picked early in the course of the data collection (on the 4k and 12k subsplits) and fixed for later experiments. So VAL performance in the small-data regime may be an optimistic estimate and in the large data regime a pessimistic estimate. True scaling behavior may be even stronger. Increasing dataset size consistently improves performance and has not yet saturated, suggesting that simply collecting more demonstrations is likely to lead to further gains.

Sample Efficiency. Fig. 5 plots VAL success against no. of training steps of experience (in millions) in Fig. 5a and against *unique* steps of experience in Fig. 5b. Recall that IL involves ~ 21 epochs on a static dataset of $\sim 70k$ demos, while RL (from [15]) gathers unique agent-driven trajectories

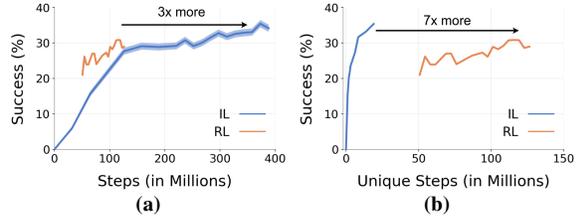


Figure 5. Comparing RL and IL on (a) VAL success vs. no. of training steps, and (b) VAL success vs. no. of unique training steps. This distinguishes between an IL agent that learns from a static dataset vs. an RL agent that gathers unique trajectories on-the-fly.

Method	VAL		TEST	
	Success % (\uparrow)	SPL % (\uparrow)	Success % (\uparrow)	SPL % (\uparrow)
1) IL w/ Shortest Paths	1.9	1.8	1.7%	1.6
2) IL w/ Human Demos	17.6 ± 0.8	9.7 ± 0.3	17.5	9.8
3) Humans	87.2	21.8%	89.1	21.9

4) IL w/ Shortest Paths	1.3	1.2	1.1	1.0
5) IL w/ Human Demos	15.9 ± 0.2	8.4 ± 0.4	15.1	8.3
6) Humans	85.0	21.0	86.1	20.5

7) IL w/ Shortest Paths	–	–	0.2	0.3
8) IL w/ Human Demos	–	–	8.3	4.1
9) Humans	–	–	94.9	20.5

Table 3. Pick-and-place results on splits constructed with unseen initializations in seen environments (1-3), with unseen instructions (4-6), and with unseen environments (7-9).

on-the-fly. Fig. 5a shows that IL behaves like supervised learning (as expected) with improvements coming from long training schedules; unfortunately, this means that wall-clock training times are not lower than RL. Fig. 5b shows that IL requires 7x fewer unique steps of experience to outperform success and is thus much more sample-efficient.

Zero-shot results on Gibson [41] are in Section A.2.

6.2. PICK&PLACE

Results. We report results in Table 3 across three evaluation splits. 1) New Initializations: new locations of objects and receptacles. This tests generalization to unseen locations in seen environments. 2) New Instructions: compositionally novel object-receptacle combinations of objects and receptacles individually seen during training. 3) New Environments: generalization to 2 scenes held out from training. Similar to OBJECTNAV and as described in Section 4, we also report results with shortest paths. Again, these paths are significantly shorter (average 342 vs. 932 steps per demonstration) and hence, we generate a larger dataset of 25.7k episodes roughly matching the cumulative steps of experience with human demonstrations (8.8M shortest path steps vs. 11.5M human steps). Training on 9.5k human demonstrations achieves 17.5% success, 9.8% SPL on new object-receptacle initializations (row 2). Across splits, training on shortest paths hurts success by 8-16%. Going to new object-receptacle pairs, success drops by 2.4% (row 5 vs. 2), and then going to new environments further hurts success by 6.8% (row 8 vs. 5). We also trained an RL policy with the exploration and distance-based rewards from [16], but it failed to get beyond

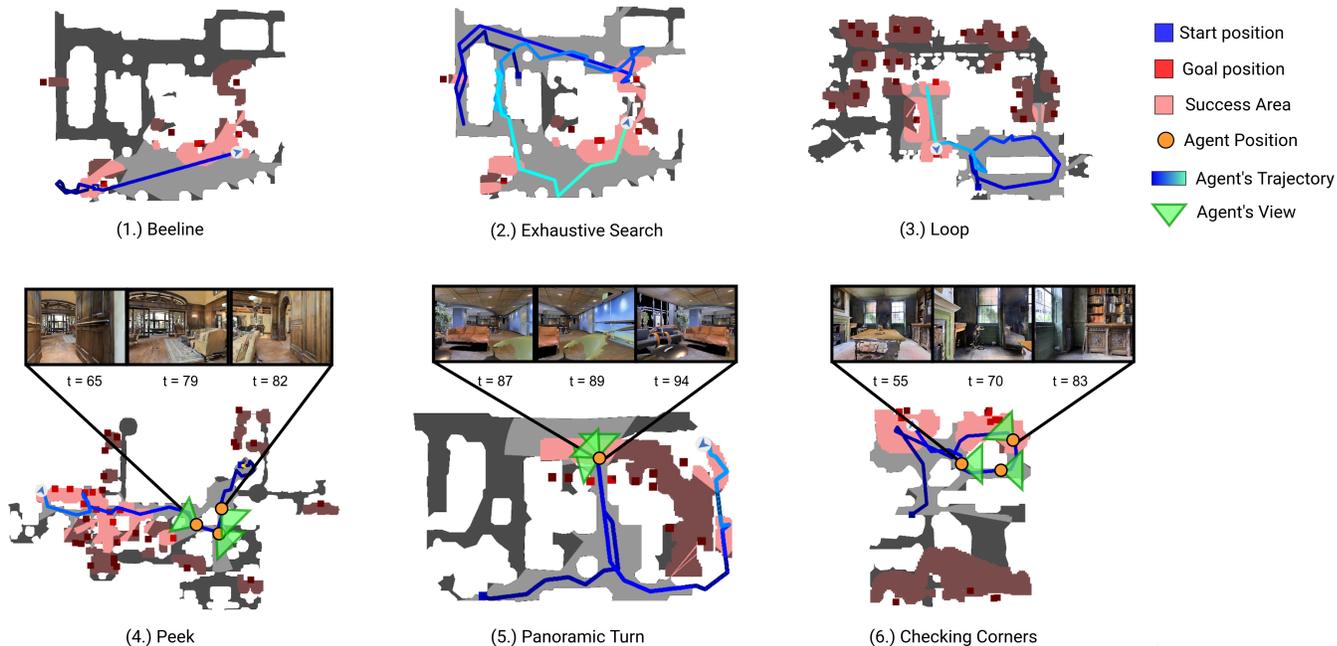


Figure 6. Visualizations of different learnt agent behaviors. Best viewed in the video at [ram81.github.io/projects/habitat-web](https://github.com/ram81/projects/habitat-web).

0% success on new object-receptacle initializations. See the appendix (Sec. A.1.2) for training details.

Performance vs. Dataset size. Similar to OBJECTNAV, we trained policies on 2.5k to 9.5k subsets of our PICK&PLACE data, and found that performance continues to improve with more data. Figure in appendix (Sec. A.1.3).

7. Characterizing Learned Behaviors

To characterize the behaviors learnt by our best IL agents, we first sample 300 validation OBJECTNAV episodes for each method and manually categorize the behavior observed. A subset of observed behaviors are visualized in Fig. 13. Our agents demonstrate sophisticated object-search behaviors *e.g.* peeking into rooms to maximize sight coverage (SC), instead of occupancy coverage (OC), checking corners of rooms for small objects, beelining to goal object once seen, exhaustive search (ES), turning in place to get a panoramic view (PT), and looping back to recheck some areas. Amusingly, unlike shortest path / RL agents, these IL agents also stand idle and ‘look around’ *i.e.* turn in place, like humans. Table 4 quantifies these behaviors. See appendix (Sec. A.4) for details on how these were computed. Agents trained with IL on human demonstrations have higher coverage (both occupancy and sight), peeking behavior, panoramic turns, beelines, and exhaustive search than RL. RL-trained agents achieve higher average Goal Room Time Spent (GRTS) – *i.e.* time spent in the room containing the target object – but also have significantly higher variance in GRTS across scenes compared to IL agents. See appendix (Sec. A.4) for a per-scene breakdown of GRTS as well as histograms of

Method	OC (%)	SC (%)	GRTS (%)	Peeks (%)	PT (%)	Beeline (%)	ES (%)
1) IL w/ shortest paths	4.2±1.1	31.2±3.2	20.5±4.3	3.0±1.9	0.0±0.0	0.0±5.2	10.1±3.5
2) IL w/ human demos	21.4±1.8	72.1±3.5	22.4±4.1	19.6±4.4	4.3±2.3	10.3±5.1	55.3±5.6
3) RL [15]	14.6±1.6	66.6±5.1	27.7±8.5	9.7±5.5	0.0±0.0	0.1±2.2	49.0±7.0
4) Humans	15.4±1.6	70.3±3.4	-	13.8±3.9	5.1±2.4	23.6±4.8	52.1±5.6

Table 4. Quantifying semantic exploration behaviors for IL agents trained on shortest paths (row 1) and human demonstrations (row 2), the Red Rabbit RL agent [15] (row 3), and humans (row 4).

time spent in each room (instead of just target room) when searching for a target object. We also discuss limitations of our approach in the appendix Sec. A.7.

8. Conclusion

We developed the infrastructure to collect human demonstrations at scale and using this, trained imitation learning (IL) agents on 92k+ human demonstrations for OBJECTNAV and PICK&PLACE. On OBJECTNAV, we found that IL using 70k human demonstrations outperforms RL using 240k agent-gathered trajectories, and on PICK&PLACE, IL agents get to ~18% success while RL fails to get beyond 0%. Qualitatively, we found that IL agents pick up on sophisticated object-search behavior implicitly captured in human demonstrations, much more prominently than RL agents. Overall, we believe our work makes a compelling case for investing in large-scale imitation learning of human demonstrations.

Acknowledgements. We thank Devi Parikh for help with idea conceptualization. The Georgia Tech effort was supported in part by NSF, ONR YIP, and ARO PECASE. 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] M. Savva, A. X. Chang, A. Dosovitskiy, T. Funkhouser, and V. Koltun, "Minos: Multimodal indoor simulator for navigation in complex environments," *arXiv preprint arXiv:1712.03931*, 2017. 1, 2, 4
- [2] P. Anderson, A. X. Chang, D. S. Chaplot, A. Dosovitskiy, S. Gupta, V. Koltun, J. Kosecka, J. Malik, R. Mottaghi, M. Savva, and A. R. Zamir, "On evaluation of embodied navigation agents," *arXiv preprint arXiv:1807.06757*, 2018. 1, 2, 4, 6, 7
- [3] D. Batra, A. X. Chang, S. Chernova, A. J. Davison, J. Deng, V. Koltun, S. Levine, J. Malik, I. Mordatch, R. Mottaghi, M. Savva, and H. Su, "Rearrangement: A Challenge for Embodied AI," *arXiv preprint arXiv:2011.01975*, 2020. 1, 4
- [4] L. Weihs, M. Deitke, A. Kembhavi, and R. Mottaghi, "Visual room rearrangement," in *CVPR*, 2021. 1
- [5] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, I. Reid, S. Gould, and A. van den Hengel, "Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments," in *CVPR*, 2018. 1, 2, 3, 4
- [6] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee, "Beyond the nav-graph: Vision-and-language navigation in continuous environments," in *ECCV*, 2020. 1, 2
- [7] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox, "ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks," in *CVPR*, 2020. 1, 2, 3, 4
- [8] A. Das, S. Datta, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, "Embodied Question Answering," in *CVPR*, 2018. 1, 2, 3, 4
- [9] E. Wijmans, S. Datta, O. Maksymets, A. Das, G. Gkioxari, S. Lee, I. Essa, D. Parikh, and D. Batra, "Embodied Question Answering in Photorealistic Environments with Point Cloud Perception," in *CVPR*, 2019. 1, 5
- [10] A. Das, F. Carnevale, H. Merzic, L. Rimell, R. Schneider, J. Abramson, A. Hung, A. Ahuja, S. Clark, G. Wayne, and F. Hill, "Probing emergent semantics in predictive agents via question answering," in *ICML*, 2020. 1
- [11] A. Das, *Building agents that can see, talk, and act*. PhD thesis, Georgia Institute of Technology, 2020. 1
- [12] L. Yu, X. Chen, G. Gkioxari, M. Bansal, T. L. Berg, and D. Batra, "Multi-target embodied question answering," in *CVPR*, 2019. 1
- [13] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, "DD-PPO: Learning near-perfect pointgoal navigators from 2.5 billion frames," in *ICLR*, 2020. 1, 5, 7, 12
- [14] J. Ye, D. Batra, E. Wijmans, and A. Das, "Auxiliary tasks speed up learning pointgoal navigation," in *CoRL*, 2020. 1
- [15] J. Ye, D. Batra, A. Das, and E. Wijmans, "Auxiliary Tasks and Exploration Enable ObjectNav," in *ICCV*, 2021. 1, 2, 3, 6, 7, 8, 12
- [16] O. Maksymets, V. Cartillier, A. Gokaslan, E. Wijmans, W. Galuba, S. Lee, and D. Batra, "THDA: Treasure Hunt Data Augmentation for Semantic Navigation," in *ICCV*, 2021. 1, 3, 6, 7
- [17] D. S. Chaplot, D. Gandhi, A. Gupta, and R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," in *NeurIPS*, 2020. 1, 3
- [18] S. Schaal, "Learning from demonstration," in *NIPS*, 1996. 1
- [19] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik, et al., "Habitat: A platform for embodied AI research," in *ICCV*, 2019. 2, 3, 13
- [20] A. Szot, A. Clegg, E. Undersander, E. Wijmans, Y. Zhao, J. Turner, N. Maestre, M. Mukadam, D. S. Chaplot, O. Maksymets, A. Gokaslan, V. Vondrus, S. Dharur, F. Meier, W. Galuba, A. Chang, Z. Kira, V. Koltun, J. Malik, M. Savva, and D. Batra, "Habitat 2.0: Training home assistants to rearrange their habitat," in *NeurIPS*, 2021. 2, 3
- [21] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, and A. Farhadi, "Target-driven visual navigation in indoor scenes using deep reinforcement learning," in *ICRA*, 2017. 2
- [22] S. Wani, S. Patel, U. Jain, A. X. Chang, and M. Savva, "MultiON: Benchmarking Semantic Map Memory using Multi-Object Navigation," in *NeurIPS*, 2020. 2
- [23] A. Padmakumar, J. Thomason, A. Shrivastava, P. Lange, A. Narayan-Chen, S. Gella, R. Pira-muthu, G. Tur, and D. Hakkani-Tur, "TEACh: Task-driven Embodied Agents that Chat," *arXiv preprint arXiv:2110.00534*, 2021. 2, 3

- [24] J. Thomason, M. Murray, M. Cakmak, and L. Zettlemoyer, "Vision-and-dialog navigation," in *CoRL*, 2020. 2, 3
- [25] M. Hahn, J. Krantz, D. Batra, D. Parikh, J. M. Rehg, S. Lee, and P. Anderson, "Where are you? localization from embodied dialog," *arXiv preprint arXiv:2011.08277*, 2020. 2, 3
- [26] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín, "What matters in learning from off-line human demonstrations for robot manipulation," in *CoRL*, 2021. 2
- [27] F. Ebert, Y. Yang, K. Schmeckpeper, B. Bucher, G. Georgakis, K. Daniilidis, C. Finn, and S. Levine, "Bridge data: Boosting generalization of robotic skills with cross-domain datasets," *arXiv preprint arXiv:2109.13396*, 2021. 2
- [28] J. Krantz, A. Gokaslan, D. Batra, S. Lee, and O. Maksymets, "Waypoint models for instruction-guided navigation in continuous environments," in *ICCV*, 2021. 2, 13
- [29] D. Gordon, A. Kembhavi, M. Rastegari, J. Redmon, D. Fox, and A. Farhadi, "IQA: Visual Question Answering in Interactive Environments," in *CVPR*, 2018. 2, 4
- [30] J. Schmidhuber, "A possibility for implementing curiosity and boredom in model-building neural controllers," in *Proc. of the international conference on simulation of adaptive behavior: From animals to animats*, 1991. 3
- [31] B. C. Stadie, S. Levine, and P. Abbeel, "Incentivizing exploration in reinforcement learning with deep predictive models," *arXiv preprint arXiv:1507.00814*, 2015. 3
- [32] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *ICML*, 2017. 3
- [33] Y. Burda, H. Edwards, D. Pathak, A. Storkey, T. Darrell, and A. A. Efros, "Large-scale study of curiosity-driven learning," *arXiv preprint arXiv:1808.04355*, 2018. 3
- [34] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, "Unifying count-based exploration and intrinsic motivation," in *NeurIPS*, 2016. 3
- [35] H. Tang, R. Houthoofd, D. Foote, A. Stooke, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel, "# exploration: A study of count-based exploration for deep reinforcement learning," in *NeurIPS*, 2017. 3
- [36] S. K. Ramakrishnan, D. Jayaraman, and K. Grauman, "An exploration of embodied visual exploration," *arXiv preprint arXiv:2001.02192*, 2020. 3
- [37] T. Chen, S. Gupta, and A. Gupta, "Learning exploration policies for navigation," in *ICLR*, 2019. 3
- [38] B. Newman, K. Carlberg, and R. Desai, "Optimal assistance for object-rearrangement tasks in augmented reality," *arXiv preprint arXiv:2010.07358*, 2020. 3
- [39] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3D: Learning from RGB-D Data in Indoor Environments," in *3DV*, 2017. MatterPort3D dataset license: http://kaldir.vc.in.tum.de/matterport/MP_TOS.pdf. 3, 4, 5, 6, 7, 13
- [40] D. Batra, A. Gokaslan, A. Kembhavi, O. Maksymets, R. Mottaghi, M. Savva, A. Toshev, and E. Wijmans, "Objectnav revisited: On evaluation of embodied agents navigating to objects," *arXiv preprint arXiv:2006.13171*, 2020. 4
- [41] F. Xia, A. R. Zamir, Z. He, A. Sax, J. Malik, and S. Savarese, "Gibson env: Real-world perception for embodied agents," in *CVPR*, 2018. Gibson dataset license agreement available at storage.googleapis.com/gibson_material/Agreement_GDS_06-04-18.pdf. 4, 6, 7, 13
- [42] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *CVPR*, 2016. 5, 12
- [43] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," in *EMNLP*, 2014. 5
- [44] A. Anand, E. Belilovsky, K. Kastner, H. Larochelle, and A. Courville, "Blindfold baselines for embodied qa," *arXiv preprint arXiv:1811.05013*, 2018. 5
- [45] J. Jiang, L. Zheng, F. Luo, and Z. Zhang, "RedNet: Residual Encoder-Decoder Network for indoor RGB-D Semantic Segmentation," *arXiv preprint arXiv:1806.01054*, 2018. 6
- [46] S. Song, S. P. Lichtenberg, and J. Xiao, "Sun RGB-D: A RGB-D scene understanding benchmark suite," in *CVPR*, 2015. 6

- [47] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, 1997. 6, 12
- [48] D. S. Chaplot, D. Gandhi, A. Gupta, and R. Salakhutdinov, “Object goal navigation using goal-oriented semantic exploration,” in *NeurIPS*, 2020. 7, 13
- [49] H. Team, “Habitat challenge, 2020.” <https://aihabitat.org/challenge/2020>, 2020. 7
- [50] A. Mousavian, A. Toshev, M. Fiser, J. Kosecka, A. Wahid, and J. Davidson, “Visual representations for semantic target driven navigation,” in *ICRA*, 2019. 13
- [51] D. S. Chaplot, S. Gupta, D. Gandhi, A. Gupta, and R. Salakhutdinov, “Learning to explore using active neural mapping,” in *ICLR*, 2020. 13