

PIRLNav: Pretraining with Imitation and RL Finetuning for OBJECTNAV

Ram Ramrakhya¹ Dhruv Batra^{1,2} Erik Wijmans¹ Abhishek Das²

¹Georgia Institute of Technology ²FAIR, Meta AI

¹{ram.ramrakhya, dbatra, etw}@gatech.edu ²abhshkdz@meta.com

Abstract

We study ObjectGoal Navigation— where a virtual robot situated in a new environment is asked to navigate to an object. Prior work [1] has shown that imitation learning (IL) on a dataset of human demonstrations achieves promising results. However, this has limitations – 1) IL policies generalize poorly to new states, since the training mimics actions not their consequences, and 2) collecting demonstrations is expensive. On the other hand, reinforcement learning (RL) is trivially scalable, but requires careful reward engineering to achieve desirable behavior. We present a two-stage learning scheme for IL pretraining on human demonstrations followed by RL-finetuning. This leads to a PIRLNav policy that advances the state-of-the-art on OBJECTNAV from 60.0% success rate to 65.0% (+5.0% absolute).

Using this IL→RL training recipe, we present a rigorous empirical analysis of design choices. First, we investigate whether human demonstrations can be replaced with ‘free’ (automatically generated) sources of demonstrations, e.g. shortest paths (SP) or task-agnostic frontier exploration (FE) trajectories. We find that IL→RL on human demonstrations outperforms IL→RL on SP and FE trajectories, even when controlled for the same IL-pretraining success on TRAIN, and even on a subset of VAL episodes where IL-pretraining success favors the SP or FE policies. Next, we study how RL-finetuning performance scales with the size of the IL pre-training dataset. We find that as we increase the size of the IL-pretraining dataset and get to high IL accuracies, the improvements from RL-finetuning are smaller, and that 90% of the performance of our best IL→RL policy can be achieved with less than half the number of IL demonstrations. Finally, we analyze failure modes of our OBJECTNAV policies, and present guidelines for further improving them.

1. Introduction

Since the seminal work of Winograd [2], designing embodied agents that have a rich understanding of the environment they are situated in, can interact with humans (and other agents) via language, and the environment via actions has

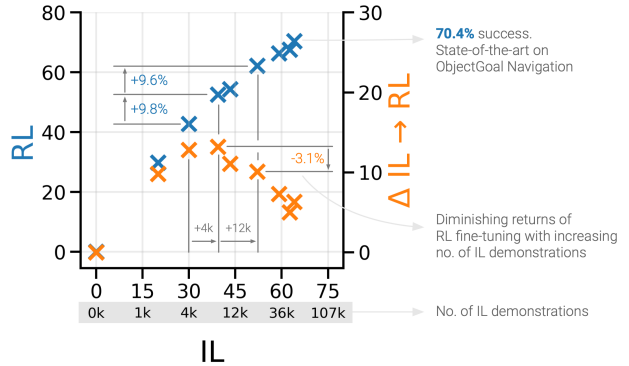


Figure 1. OBJECTNAV success rates of agents trained using imitation learning (IL) vs. IL-pretraining followed by reinforcement learning (RL) (in blue). RL from scratch (i.e. IL=0) fails to get off-the-ground. With more IL demonstrations, IL success increases, and it transfers to even higher RL-finetuning success. But the difference between RL-finetuning vs. IL-pretraining success (in orange) plateaus and starts to decrease beyond a certain point, indicating diminishing returns with each additional IL demonstration.

been a long-term goal in AI [3–12]. We focus on ObjectGoal Navigation [13, 14], wherein an agent situated in a new environment is asked to navigate to any instance of an object category (‘find a plant’, ‘find a bed’, etc.); see Fig. 2. OBJECTNAV is simple to explain but difficult for today’s techniques to accomplish. First, the agent needs to be able to ground the tokens in the language instruction to physical objects in the environment (e.g. what does a ‘plant’ look like?). Second, the agent needs to have rich semantic priors to guide its navigation to avoid wasteful exploration (e.g. the microwave is likely to be found in the kitchen, not the washroom). Finally, it has to keep track of where it has been in its internal memory to avoid redundant search.

Humans are adept at OBJECTNAV. Prior work [1] collected a large-scale dataset of 80k human demonstrations for OBJECTNAV, where human subjects on Mechanical Turk tele-operated virtual robots and searched for objects in novel houses. This first provided a human baseline on OBJECTNAV of 88.9% success rate on the Matterport3D (MP3D) dataset [15]¹ compared to 35.4% success rate of the best

¹On VAL split, for 21 object categories, and a maximum of 500 steps.



Figure 2. OBJECTNAV trajectories for policies trained with IL+RL on 1) Human Demonstrations, 2) Shortest Paths, and 3) Frontier Exploration Demonstrations.

performing method [1]. This dataset was then used to train agents via imitation learning (specifically, behavior cloning) of human demonstrations.

While this approach achieved state-of-art results (35.4% success rate on MP3D VAL dataset), it has two clear limitations. First, imitation learning (behavior cloning) is known to suffer from poor generalization to out-of-distribution states not seen during training, since the training emphasizes imitating actions not accomplishing their goals. Second and more importantly, it is expensive and thus not scalable. Specifically, Ramrakhya *et al.* [1] collected 80k demonstrations on 56 scenes in Matterport3D Dataset, which took ~ 2894 hours of human teleoperation and \$50k dollars. A few months after [1] was released, a new higher-quality dataset called HM3D-Semantics v0.1 [16] became available with 120 annotated 3D scenes, and a few months after that HM3D-Semantics v0.2 added 96 additional scenes. Scaling Ramrakhya *et al.*’s approach to continuously incorporate new scenes involves replicating that entire effort again and again. On the other hand, training with reinforcement learning (RL) is trivially scalable once annotated 3D scans are available. However, as demonstrated in Maksymets *et al.* [17], RL requires careful reward engineering, the reward function typically used for OBJECTNAV actually *penalizes* exploration (even though the task requires it), and the existing RL policies overfit to the small number of available environments.

Our primary technical contribution is PIRLNav, an approach for pretraining with IL and finetuning with RL for OBJECTNAV. IL pretrained policies provide a reasonable starting point for ‘bootstrapping’ RL and make the optimization easier than learning from scratch. In fact, we show that IL pretraining even unlocks RL with sparse rewards. Sparse rewards are simple (do not involve any reward engineering) and do not suffer from the unintended consequences described above. However, learning from scratch with sparse

rewards is typically out of reach since most random action trajectories result in no positive rewards.

While combining IL and RL has been studied in prior work [18–22], the main technical challenge in the context of modern neural networks is that imitation pretraining results in weights for the policy (or actor), but not a value function (or critic). Thus, naively initializing a new RL policy with these IL-pretrained policy weights often leads to catastrophic failures due to destructive policy updates early on during RL training, especially for actor-critic RL methods [23]. To overcome this challenge, we present a two-stage learning scheme involving a critic-only learning phase first that gradually transitions over to training both the actor and critic. We also identify a set of practical recommendations for this recipe to be applied to OBJECTNAV. This leads to a PIRLNav policy that advances the state-the-art on OBJECTNAV from 60.0% success rate (in [24]) to 65.0% (+5.0%, 8.3% relative improvement).

Next, using this IL+RL training recipe, we conduct an empirical analysis of design choices. Specifically, an ingredient we investigate is whether human demonstrations can be replaced with ‘free’ (automatically generated) sources of demonstrations for OBJECTNAV, *e.g.* (1) shortest paths (SP) between the agent’s start location and the closest object instance, or (2) task-agnostic frontier exploration [25] (FE) of the environment followed by shortest path to goal-object upon observing it. We ask and answer the following questions:

1. ‘Do human demonstrations capture any unique OBJECTNAV-specific behaviors that shortest paths and frontier exploration trajectories do not?’ Yes. We find that IL / IL+RL on human demonstrations outperforms IL / IL+RL on shortest paths and frontier exploration trajectories respectively. When we control the number of demonstrations from each source such that IL success on TRAIN is the same, RL-finetuning when initialized from IL

on human demonstrations still outperforms the other two.

2. ‘How does performance after RL scale with IL dataset size?’ We observe diminishing returns from RL-finetuning as we scale IL dataset size. This suggests, by effectively leveraging the trade-off curve between size of pretraining dataset size vs. performance after RL-Finetuning, we can achieve closer to state-of-the-art results without investing into a large dataset of IL demonstrations.

3. ‘Does IL on frontier exploration demonstrations present similar scaling behavior as IL on human demonstrations?’

No. We find that as we scale frontier exploration demonstrations past 70k trajectories, the performance plateaus.

Finally, we present an analysis of the failure modes of our OBJECTNAV policies and present a set of guidelines for further improving them. Our policy’s primary failure modes are: a) Dataset issues: comprising of missing goal annotations, and navigation meshes blocking the path, b) Navigation errors: primarily failure to navigate between floors, c) Recognition failures: where the agent does not identify the goal object during an episode, or confuses the specified goal with a semantically-similar object.

2. Related Work

ObjectGoal Navigation. Prior works on OBJECTNAV have used end-to-end reinforcement learning (RL) [17, 26, 27], modular learning [24, 28, 29], and imitation learning [1, 30]. Works that use end-to-end RL have proposed improved visual representations [26, 31], auxiliary tasks [27], and data augmentation techniques [17] to improve generalization to unseen environments. Improved visual representations include object relation graphs [31] and semantic segmentations [26]. Ye *et al.* [27] use auxiliary tasks like predicting environment dynamics, action distributions, and map coverage in addition to OBJECTNAV and achieve promising results. Maksymets *et al.* [17] improve generalization of RL agents by training with artificially inserted objects and proposing a reward to incentivize exploration.

Modular learning methods for OBJECTNAV have also emerged as a strong competitor [24, 28, 32]. These methods rely on separate modules for semantic mapping that build explicit structured map representations, a high-level semantic exploration module that is learned through RL to solve the ‘where to look?’ subproblem, and a low-level navigation policy that solves ‘how to navigate to (x, y) ?’.

The current state-of-the-art methods on OBJECTNAV [1, 30] make use of imitation learning (IL) on a large dataset of 80k human demonstrations. with a simple CNN+RNN policy architecture. In this work, we improve on them by developing an effective approach to finetune these imitation-pretrained policies with RL.

Imitation Learning and RL Finetuning. Prior works have considered a special case of learning from demonstration data. These approaches initialize policies trained using be-

havior cloning, and then fine-tune using on-policy reinforcement learning [18, 20–22, 33, 34]. On classical tasks like cart-pole swing-up [18], balance, hitting a baseball [33], and underactuated swing-up [34], demonstrations have been used to speed up learning by initializing policies pretrained on demonstrations for RL. Similar to these methods, we also use a on-policy RL algorithm for finetuning the policy trained with behavior cloning. Rajeswaran *et al.* [20] (DAPG) pretrain a policy using behavior cloning and use an augmented RL finetuning objective to stay close to the demonstrations which helps reduce sample complexity. Unfortunately DAPG is not feasible in our setting as it requires solving a systems research problem to efficiently incorporate replaying demonstrations and collecting experience online at our scale. [20] show results of the approach on a dexterous hand manipulation task with a small number of demonstrations that can be loaded in system memory and therefore did not need to solve this system challenge. This is not possible in our setting, just the 256×256 RGB observations for the 77k demos we collect would occupy over 2 TB memory, which is out of reach for all but the most exotic of today’s systems. There are many methods for incorporating demonstrations/imitation learning with off-policy RL [35–39]. Unfortunately these methods were not designed to work with recurrent policies and adapting off-policy methods to work with recurrent policies is challenging [40]. See the Appendix A for more details. The RL finetuning approach that demonstrates results with an actor-critic and high-dimensional visual observations, and is thus most closely related to our setup is proposed in VPT [21]. Their approach uses Phasic Policy Gradients (PPG) [41] with a KL-divergence loss between the current policy and the frozen pretrained policy, and decays the KL loss weight ρ over time to enable exploration during RL finetuning. Our approach uses Proximal Policy Gradients (PPO) [42] instead of PPG, and therefore does not require a KL constraint, which is compute-expensive, and performs better on OBJECTNAV.

3. OBJECTNAV and Imitation Learning

3.1. OBJECTNAV

In OBJECTNAV an agent is tasked with searching for an instance of the specified object category (*e.g.*, ‘bed’) in an unseen environment. The agent must perform this task using only egocentric perceptions. Specifically, a RGB camera, Depth sensor², and a GPS+Compass sensor that provides location and orientation relative to the start position of the episode. The 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. An episode is considered successful if the agent stops within 1m Euclidean distance of the goal object

²We don’t use this sensor as we don’t find it helpful.

within 500 steps and is able to view the object by taking turn actions [14].

We use scenes from the HM3D-Semantics v0.1 dataset [16]. The dataset consists of 120 scenes and 6 unique goal object categories. We evaluate our agent using the train/val/test splits from the 2022 Habitat Challenge³.

3.2. OBJECTNAV Demonstrations

Ramrakhya *et al.* [1] collected OBJECTNAV demonstrations for the Matterport3D dataset [15]. We begin our study by replicating this effort and collect demonstrations for the HM3D-Semantics v0.1 dataset [16]. We use Ramrakhya *et al.*'s Habitat-WebGL infrastructure to collect 77k demonstrations, ~ 2378 human hours.

3.3. Imitation Learning from Demonstrations

We use behavior cloning to pretrain our OBJECTNAV policy on the human demonstrations we collect. Let $\pi_{\theta}^{BC}(a_t | o_t)$ 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, 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 optimal parameters are

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

We use inflection weighting [43] to adjust the loss function to upweight timesteps where actions change (*i.e.* $a_{t-1} \neq a_t$). Our **ObjectNav policy** architecture is a simple CNN+RNN model from [30]. To encode RGB input ($i_t = \text{CNN}(I_t)$), we use a ResNet50 [44]. Following [30], the CNN is first pre-trained on the Omnidata starter dataset [45] using the self-supervised pretraining method DINO [46] and then fine-tuned during OBJECTNAV training. The GPS+Compass inputs, $P_t = (\Delta x, \Delta y, \Delta z)$, and $R_t = (\Delta \theta)$, are passed through fully-connected layers $p_t = \text{FC}(P_t)$, $r_t = \text{FC}(R_t)$ to embed them to 32-d vectors. Finally, we convert the object goal category to one-hot and pass it through a fully-connected layer $g_t = \text{FC}(G_t)$, resulting in a 32-d vector. All of these input features are concatenated to form an observation embedding, and fed into a 2-layer, 2048-d GRU at every timestep to predict a distribution over actions a_t - formally, given current observations $o_t = [i_t, p_t, r_t, g_t]$, $(h_t, a_t) = \text{GRU}(o_t, h_{t-1})$. To reduce overfitting, we apply color-jitter and random shifts [47] to the RGB inputs.

4. RL Finetuning

Our motivation for RL-finetuning is two-fold. First, finetuning may allow for higher performance as behavior cloning

is known to suffer from a train/test mismatch – when training, the policy sees the result of taking ground-truth actions, while at test-time, it must contend with the consequences of its own actions. Second, collecting more human demonstrations on new scenes or simply to improve performance is time-consuming and expensive. On the other hand, RL-finetuning is trivially scalable (once annotated 3D scans are available) and has the potential to reduce the amount of human demonstrations needed.

4.1. Setup

The RL objective is to find a policy $\pi_{\theta}(a|s)$ that maximizes expected sum of discounted future rewards. Let τ be a sequence of object, action, reward tuples (o_t, a_t, r_t) where $a_t \sim \pi_{\theta}(\cdot | o_t)$ is the action sampled from the agent's policy, and r_t is the reward. For a discount factor γ , the optimal policy is

$$\pi^* = \arg \max_{\pi} \mathbb{E}_{\tau \sim \pi} [R_T], \text{ where } R_T = \sum_{t=1}^T \gamma^{t-1} r_t. \quad (2)$$

To solve this maximization problem, actor-critic RL methods learn a state-value function $V(s)$ (also called a critic) in addition to the policy (also called an actor). The critic $V(s_t)$ represents the expected value of returns R_t when starting from state s_t and acting under the policy π , where returns are defined as $R_t = \sum_{i=t}^T \gamma^{i-t} r_i$. We use DD-PPO [48], a distributed implementation of PPO [42], an on-policy RL algorithm. Given a θ -parameterized policy π_{θ} and a set of rollouts, PPO updates the policy as follows. Let $\hat{A}_t = R_t - V(s_t)$, be the advantage estimate and $p_t(\theta) = \frac{\pi_{\theta}(a_t | o_t)}{\pi_{\theta_{\text{old}}}(a_t | o_t)}$ be the ratio of the probability of action a_t under current policy and under the policy used to collect rollouts. The parameters are updated by maximizing:

$$J^{PPO}(\theta) = \mathbb{E}_t \left[\min(p_t(\theta) \hat{A}_t, \text{clip}(p_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t) \right] \quad (3)$$

We use a sparse success reward. Sparse success is simple (does not require hyperparameter optimization) and has fewer unintended consequences (*e.g.* Maksymets *et al.* [17] showed that typical dense rewards used in OBJECTNAV actually *penalize* exploration, even though exploration is necessary for OBJECTNAV in new environments). Sparse rewards are desirable but typically difficult to use with RL (when initializing training from scratch) because they result in nearly all trajectories achieving 0 reward, making it difficult to learn. However, since we pretrain with IL, we do not observe any such pathologies.

4.2. Finetuning Methodology

We use the behavior cloned policy π_{θ}^{BC} weights to initialize the actor parameters. However, notice that during behavior

³<https://aihabitat.org/challenge/2022/>

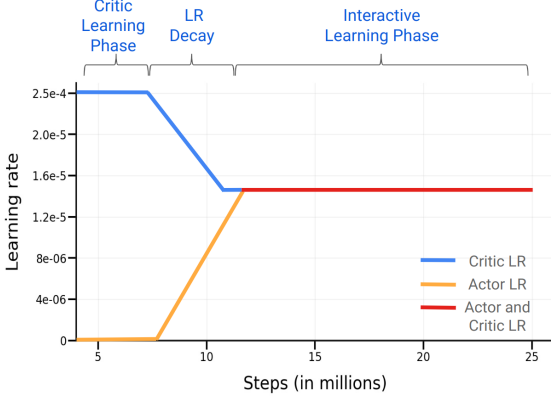


Figure 3. Learning rate schedule for RL Finetuning.

cloning we do not learn a critic nor is it easy to do so – a critic learned on human demonstrations (during behavior cloning) would be overly optimistic since all it sees are successes. Thus, we must learn the critic from scratch during RL. Naively finetuning the actor with a randomly-initialized critic leads to a rapid drop in performance⁴ (see Fig. 8) since the critic provides poor value estimates which influence the actor’s gradient updates (see Eq.(3)). We address this issue by using a two-phase training regime:

Phase 1: Critic Learning. In the first phase, we rollout trajectories using the frozen policy, pre-trained using IL, and use them to learn a critic. To ensure consistency of rollouts collected for critic learning with RL training, we sample actions (as opposed to using argmax actions) from the pre-trained IL policy: $a_t \sim \pi_\theta(s_t)$. We train the critic until its loss plateaus. In our experiments, we found 8M steps to be sufficient. In addition, we also initialize the weights of the critic’s final linear layer close to zero to stabilize training.

Phase 2: Interactive Learning. In the second phase, we unfreeze the actor RNN⁵ and finetune both actor and critic weights. We find that naively switching from phase 1 to phase 2 leads to small improvements in policy performance at convergence. We gradually decay the critic learning rate from 2.5×10^{-4} to 1.5×10^{-5} while warming-up the policy learning rate from 0 to 1.5×10^{-5} between 8M to 12M steps, and then keeping both at 1.5×10^{-5} through the course of training. See Fig. 3. We find that using this learning rate schedule helps improve policy performance. For parameters that are shared between the actor and critic (*i.e.* the RNN), we use the lower of the two learning rates (*i.e.* always the actor’s in our schedule). To summarize our finetuning methodology:

- First, we initialize the weights of the policy network with the IL-pretrained policy and initialize critic weights close to zero. We freeze the actor and shared weights. The only

⁴After the initial drop, the performance increases but the improvements on success are small.

⁵The CNN and non-visual observation embedding layers remain frozen. We find this to be more stable.

learnable parameters are in the critic.

- Next, we learn the critic weights on rollouts collected from the pretrained, frozen policy.
- After training the critic, we warmup the policy learning rate and decay the critic learning rate.
- Once both critic and policy learning rate reach a fixed learning rate, we train the policy to convergence.

4.3. Results

Comparing with the RL-finetuning approach in VPT [21].

We start by comparing our proposed RL-finetuning approach with the approach used in VPT [21]. Specifically, [21] proposed initializing the critic weights to zero, replacing entropy term with a KL-divergence loss between the frozen IL policy and the RL policy, and decay the KL divergence loss coefficient, ρ , by a fixed factor after every iteration. Notice that this prevents the actor from drifting too far too quickly from the IL policy, but does not solve uninitialized critic problem. To ensure fair comparison, we implement this method within our DD-PPO framework to ensure that any performance difference is due to the fine-tuning algorithm and not tangential implementation differences. Complete training details are in the Appendix B.3. We keep hyperparameters constant for our approach for all experiments. Table 1 reports results on HM3D VAL for the two approaches using 20k human demonstrations. We find that PIRLNav achieves +2.1% Success compared to VPT and comparable SPL.

Method	Success (\uparrow)	SPL (\uparrow)
1) IL	52.0%	20.6%
2) IL \rightarrow RL-FT w/ VPT [21]	60.1%	29.1%
3) PIRLNav (Ours)	62.2%	28.7%

Table 1. Comparison with VPT on HM3D VAL [16, 32]

Method	Success (\uparrow)	SPL (\uparrow)
1) IL	52.0%	20.6%
2) IL \rightarrow RL-FT	54.8%	29.1%
3) IL \rightarrow RL-FT (+ Critic Learning)	55.5%	26.7%
4) IL \rightarrow RL-FT (+ Critic Learning, Critic Decay)	59.6%	26.5%
5) IL \rightarrow RL-FT (+ Critic Learning, Actor Warmup)	58.7%	25.8%
6) PIRLNav	62.2%	28.7%

Table 2. RL-finetuning ablations on HM3D VAL [16, 32]

Ablations. Next, we conduct ablation experiments to quantify the importance of each phase in our RL-finetuning approach. Table 2 reports results on the HM3D VAL split for a policy IL-pretrained on 20k human demonstrations and RL-finetuned for 300M steps, complete training details are in Appendix B.4. First, without a gradual learning transition (row 2), *i.e.* without a critic learning and LR decay phase, the policy improves by 2.8% on success and 8.5% on SPL. Next, with only a critic learning phase (row 3), the policy improves by 3.5% on success and 6.1% on SPL. Using an LR decay schedule only for the critic after the critic learn-

ing phase improves success by 7.6% and SPL by 5.9%, and using an LR warmup schedule for the actor (but no critic LR decay) after the critic learning phase improves success by 6.7% and SPL by 5.3%. Finally, combining everything (critic-only learning, critic LR decay, actor LR warmup), our policy improves by 10.2% on success and 8.1% on SPL.

Method	TEST-STD		TEST-CHALLENGE	
	Success (↑)	SPL (↑)	Success (↑)	SPL (↑)
1) Stretch [24]	60.0%	34.0%	56.0%	29.0%
2) ProcTHOR-Large [49]	54.0%	32.0%	-	-
3) Habitat-Web [1]	55.0%	22.0%	-	-
4) DD-PPO [50]	26.0%	12.0%	-	-
5) Populus A.	66.0%	32.0%	60.0%	30.0%
6) ByteBOT	68.0%	37.0%	64.0%	35.0%
7) PIRLNav ⁶	65.0%	33.0%	65.0%	33.0%

Table 3. Results on HM3D TEST-STANDARD and TEST-CHALLENGE [16, 50]. Unpublished works submitted only to the OBJECTNAV leaderboard have been grayed out.

ObjectNav Challenge 2022 Results. Using our overall two-stage training approach of IL-pretraining followed by RL-finetuning, we achieve state-of-the-art results on OBJECTNAV – 65.0% success and 33.0% SPL on both the TEST-STANDARD and TEST-CHALLENGE splits. Table 3 compares our results with the top-4 entries to the Habitat OBJECTNAV Challenge 2022 [50]. Our approach outperforms Stretch [24] on success rate on both TEST-STANDARD and TEST-CHALLENGE and is comparable on SPL (1% worse on TEST-STANDARD, 4% better on TEST-CHALLENGE). ProcTHOR [49], which uses 10k procedurally-generated environments for training, achieves 54% success and 32% SPL on TEST-STANDARD split, which is 11% worse at success and 1% worse at SPL than ours. For sake of completeness, we also report results of two unpublished entries uploaded to the leaderboard – Populus A. and ByteBOT. Unfortunately, there is no associated report yet with these entries, so we are unable to comment on the details of these approaches, or even whether the comparison is meaningful.

5. Role of demonstrations in IL→RL transfer

Our decision to use human demonstrations for IL-pretraining before RL-finetuning was motivated by results in prior work [1]. Next, we examine if other cheaper sources of demonstrations lead to equally good IL→RL generalization. Specifically, we consider 3 sources of demonstrations:

Shortest paths (SP). These demonstrations are generated by greedily sampling actions to fit the geodesic shortest path to the nearest navigable goal object, computed using the ground-truth map of the environment. These demonstrations do not capture any exploration, they only capture success at the OBJECTNAV task via the most efficient path.

Task-Agnostic Frontier Exploration (FE) [24]. These

Training demonstrations	Success (↑)	SPL (↑)
Shortest paths (240k)	6.4%	5.0%
Frontier exploration (70k)	44.9%	21.5%
Human demonstrations (77k)	64.1%	27.1%

Table 4. Performance on HM3D VAL with imitation learning on SP, FE, and HD demonstrations. The size of each demonstration dataset is picked such that total steps of experience is $\sim 12M$.

are generated by using a 2-stage approach: 1) Exploration: where a task-agnostic strategy is used to maximize exploration coverage and build a top-down semantic map of the environment, and 2) Goal navigation: once the goal object is detected by the semantic predictor, the developed map is used to reach it by following the shortest path. These demonstrations capture OBJECTNAV-agnostic exploration.

Human Demonstrations (HD) [1]. These are collected by asking humans on Mechanical Turk to control an agent and navigate to the goal object. Humans are provided access to the first-person RGB view of the agent and tasked to reach within 1m of the goal object category. These demonstrations capture human-like OBJECTNAV-specific exploration.

5.1. Results with Imitation Learning

Using the IL approach described in Sec. 3.3, we train on SP, FE, and HD demonstrations. Since these demonstrations vary in trajectory length (*e.g.* SP are significantly shorter than FE), we collect $\sim 12M$ steps of experience with each method. That amounts to 240k SP, 70k FE, and 77k HD demonstrations respectively. Our IL results are summarized in Table 4. IL on 240k SP demonstrations leads to 6.4% success and 5.0% SPL. We believe this poor performance is due to an imitation gap [51], *i.e.* the shortest path demonstrations are generated with access to privileged information (ground-truth map of the environment) which is not available to the policy during training. Without a map, following the shortest path in a new environment to find a goal object is not possible. IL on 70k FE demonstrations achieves 44.9% success and 21.5% SPL, which is significantly better than IL on shortest paths (+38.5% success, +16.5% SPL). Finally, IL on 77k HD obtains the best results – 64.1% success and 27.1% SPL. These trends suggest that task-specific exploration (captured in human demonstrations) leads to much better generalization than task-agnostic exploration (FE) or shortest paths (SP).

5.2. Results with RL Finetuning

Using the IL-pretrained policies on SP, FE, and HD demonstrations as initialization, we RL-finetune each using our approach described in Sec. 4. These results are summarized in Fig. 4. Perhaps intuitively, the trends after RL-finetuning follow the same ordering as IL-pretraining, *i.e.* RL-finetuning from IL on HD > FE > SP. But there are two factors that could be leading to this ordering after RL-finetuning – 1) inconsistency in performance at initialization (*i.e.* IL on HD

⁶The approach is called “BadSeed” on the HM3D leaderboard: eval.ai/web/challenges/challenge-page/1615/leaderboard/3899




Figure 4. OBJECTNAV performance on HM3D VAL with IL-pretraining on shortest path (SP), frontier exploration (FE), and human demonstrations (HD), followed by RL-finetuning from each.




Figure 5. IL and RL performance for shortest paths (SP), frontier exploration (FE), and human demonstrations (HD) with equal IL training success on HM3D TRAIN (left) and VAL (right).

is already better than IL on FE), and 2) amenability of each of these initializations to RL-finetuning (*i.e.* is RL-finetuning from HD init better than FE init?).

We are interested in answering (2), and so we control for (1) by selecting IL-pretrained policy weights across SP, FE, and HD that have equal performance on a subset of TRAIN = $\sim 48.0\%$ success. This essentially amounts to selecting IL-pretraining checkpoints for FE and HD from earlier in training as $\sim 48.0\%$ success is the maximum for SP.

Fig. 5 shows the results after IL and RL-finetuning on a subset of the HM3D TRAIN and on HM3D VAL. First, note that at IL-pretraining TRAIN success rates are equal ($= \sim 48.0\%$), while on VAL FE is slightly better than HD followed by SP. We find that after RL-finetuning, the policy trained on HD still leads to higher VAL success (66.1%) compared to FE (51.3%) and SP (43.6%). Notice that RL-finetuning from SP leads to high TRAIN success, but low VAL success, indicating significant overfitting. FE has smaller TRAIN-VAL gap after RL-finetuning but both are worse than HD, indicating underfitting. These results show that learning to imitate human demonstrations equips the agent with navigation strategies that enable better RL-finetuning generalization compared to imitating other kinds of demonstrations, even when controlled for the same IL-pretraining accuracy.

Results on SP-favoring and FE-favoring episodes. To further emphasize that imitating human demonstrations is key to good generalization at OBJECTNAV, we created two subsplits from the HM3D VAL split that are adversarial to HD performance – SP-favoring and FE-favoring. The

Training demonstrations	IL Success (\uparrow)	RL-FT Success (\uparrow)
1) SP	5.2%	34.8%
2) HD	0.0%	57.2%
3) FE	26.3%	43.0%
4) HD	0.0%	57.2%

Table 5. Results on SP-favoring and FE-favoring splits.

SP-favoring VAL split consists of episodes where IL on SP achieved a higher performance compared to IL on HD, *i.e.* we select episodes where IL on SP succeeded but IL on HD did not or both IL on SP and IL on HD failed. Similarly, we also create an FE-favoring VAL split using the same sampling strategy biased towards IL on FE. Next, we report the performance of RL-finetuned from IL on SP, FE, and HD on these two evaluation splits in Table 5. On both SP-favoring and FE-favoring, IL on HD is at 0% success (by design), but after RL-finetuning, is able to significantly outperform RL-finetuning from the respective IL on SP and FE policies.

5.3. Scaling laws of IL and RL

In this section, we investigate how IL-pretraining and RL-finetuning performance scales with no. of IL demonstrations. **Human demonstrations.** We create HD subsplits ranging in size from $2k$ to $77k$ episodes, and IL-pretrain policies with the same set of hyperparameters on each split. Then, for each, we RL-finetune from the best-performing checkpoint. The resulting IL and RL success on HM3D VAL *vs.* no. of HD episodes is plotted in Fig. 1. Similar to [1], we see promising scaling behavior with more IL demonstrations.

Interestingly, as we increase the size of the IL pretraining dataset and get to high IL accuracies, the improvements from RL-finetuning decrease. *E.g.* at $20k$ IL demonstrations, the IL \rightarrow RL improvement is 10.1% success, while at $77k$ IL demonstrations, the improvement is 6.3%. Furthermore, with $35k$ IL-pretraining demonstrations, the RL-finetuned success is only 4% worse than RL-finetuning from $77k$ IL demonstrations (66.4% *vs.* 70.4%). Both of these suggest that by effectively leveraging the trade-off curve between the size of the IL-pretraining dataset *vs.* performance gains after RL-finetuning, it may be possible to achieve close to state-of-the-art results without large investments in demonstrations.

How well does FE Scale? In Section 5.1, we showed that IL on human demonstrations outperforms IL on both shortest paths and frontier exploration demonstrations, when controlled for the same amount of training experience. In contrast to human demonstrations however, collecting shortest paths and frontier exploration demonstrations is cheaper, which makes scaling these demonstration datasets easier. Since IL performance on shortest paths is significantly worse even with 3x more demonstrations compared to FE and HD ($240k$ SP *vs.* $70k$ FE and $77k$ HD demos, Sec. 5.1), we focus on scaling FE demonstrations. Fig. 6 plots performance on HM3D VAL against FE dataset size. We created splits




Figure 6. Success on ObjectNav HM3D VAL split vs. no. of frontier exploration demonstrations for training.




Figure 7. Failure modes of our best IL+RL OBJECTNAV policy

ranging in size from 10k to 150k. Increasing the dataset size doesn’t consistently improve performance and saturates after 70k demonstrations, suggesting that simply generating more FE demonstrations is unlikely to help. We hypothesize that the saturation in performance is because these demonstrations do not capture task-specific exploration.

6. Failure Modes

To better understand the failure modes of our IL+RL OBJECTNAV policies, we manually annotate 592 failed HM3D VAL episodes from our best OBJECTNAV agent. See Fig. 7. We include videos of various failures modes in the supplement. The most common failure modes are:

Missing Annotations (27%): Episodes where the agent navigates to the correct goal object category but the episode is counted as a failure due to missing annotations in the data.

Inter-Floor Navigation (21%): The object is on a different floor and the agent fails to climb up/down the stairs.

Recognition Failure (20%): The agent sees the object in its field of view but fails to navigate to it.

Last Mile Navigation [52] (12%). Repeated collisions against objects or mesh geometry close to the goal object preventing the agent from reaching close to it.

Navmesh Failure (9%). Hard-to-navigate meshes blocking the path of the agent. *E.g.* in one instance, the agent fails to

climb stairs because of a narrow nav mesh on the stairs.

Looping (4%). Repeatedly visiting the same location and not exploring the rest of the environment.

Semantic Confusion (5%). Confusing the goal object with a semantically-similar object. *E.g.* ‘armchair’ for ‘sofa’.

Exploration Failure (2%). Catch-all for failures in a complex navigation environment, early termination, semantic failures (*e.g.* looking for a chair in a bathroom), *etc.*

As can be seen in Fig. 7, most of the failures (~36%) are due to issues in the OBJECTNAV dataset – 27% due to missing object annotations + 9% due to holes / issues in the navmesh. 21% failures are due to the agent being unable to climb up/down stairs. We believe this happens because climbing up / down stairs to explore another floor is a difficult behavior to learn and there are few episodes that require this. Oversampling inter-floor navigation episodes during training can help address this. Another failure mode is failing to recognize the goal object – 20% failures where the object is in the agent’s field of view but it does not navigate to it, and 5% where the agent navigates to another semantically-similar object. Advances in the visual backbone and object recognition can help address these. To this end, prior works [1, 24] have used explicit semantic segmentation modules to recognize objects at each step of navigation. Incorporating this within the IL+RL training pipeline could help. 11% failures are due to last mile navigation, suggesting that equipping the agent with better goal-distance estimators could help. Finally, only ~6% failures are due to looping and lack of exploration, which is promising!

7. Conclusion

To conclude, we propose PIRLNav, an approach to combine imitation (IL) and reinforcement learning (RL) for OBJECTNAV, wherein we pretrain a policy with IL on 77k human demonstrations and then finetune it with RL, leading to state-of-the-art results on OBJECTNAV (65% success, 5% improvement over previous best). Next, using this IL+RL training recipe, we present a thorough empirical study of the impact of different demonstration datasets used for IL-pretraining on downstream RL-finetuning performance. We show that IL / IL+RL on human demonstrations outperforms IL / IL+RL on shortest paths and frontier exploration trajectories, even when we control for same IL success on TRAIN. We also show that as we scale the pretraining dataset size for IL and get to higher IL success rates, the improvements from RL-finetuning start to diminish. Finally, we characterize the failure modes of our best OBJECTNAV agent, and find that the largest sources of error are 1) dataset annotation noise (missing goal annotations, incorrect navigation meshes), and inability of the agent to 2) navigate across floors, and 3) recognize the correct goal object.

References

- [1] R. Ramrakhyia, E. Undersander, D. Batra, and A. Das, “Habitat-web: Learning embodied object-search strategies from human demonstrations at scale,” in *CVPR*, 2022. 1, 2, 3, 4, 6, 7, 8, 13
- [2] T. Winograd, “Understanding natural language,” *Cognitive Psychology*, 1972. 1
- [3] L. Smith and M. Gasser, “The development of embodied cognition: six lessons from babies,” *Artificial life*, vol. 11, no. 1-2, 2005. 1
- [4] K. M. Hermann, F. Hill, S. Green, F. Wang, R. Faulkner, H. Soyer, D. Szepesvari, W. Czarnecki, M. Jaderberg, D. Teplyashin, *et al.*, “Grounded language learning in a simulated 3D world,” *arXiv preprint arXiv:1706.06551*, 2017. 1
- [5] F. Hill, K. M. Hermann, P. Blunsom, and S. Clark, “Understanding grounded language learning agents,” *arXiv preprint arXiv:1710.09867*, 2017. 1
- [6] D. S. Chaplot, K. M. Sathyendra, R. K. Pasumarthi, D. Rajagopal, and R. Salakhutdinov, “Gated-attention architectures for task-oriented language grounding,” in *AAAI*, 2018. 1
- [7] 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
- [8] U. Jain, L. Weihs, E. Kolve, M. Rastegari, S. Lazebnik, A. Farhadi, A. G. Schwing, and A. Kembhavi, “Two body problem: Collaborative visual task completion,” in *CVPR*, 2019. 1
- [9] A. Das, *Building agents that can see, talk, and act*. PhD thesis, Georgia Institute of Technology, 2020. 1
- [10] J. Abramson, A. Ahuja, I. Barr, A. Brussee, F. Carnevale, M. Cassin, R. Chhaparia, S. Clark, B. Damoc, A. Dudzik, *et al.*, “Imitating interactive intelligence,” *arXiv preprint arXiv:2012.05672*, 2020. 1
- [11] L. Weihs, A. Kembhavi, K. Ehsani, S. M. Pratt, W. Han, A. Herrasti, E. Kolve, D. Schwenk, R. Mottaghi, and A. Farhadi, “Learning generalizable visual representations via interactive gameplay,” in *ICLR*, 2021. 1
- [12] C. Lynch, A. Wahid, J. Tompson, T. Ding, J. Betker, R. Baruch, T. Armstrong, and P. Florence, “Interactive language: Talking to robots in real time,” *arXiv preprint arXiv:2210.06407*, 2022. 1
- [13] 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
- [14] 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. 1, 4
- [15] 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. 1, 4
- [16] K. Yadav, R. Ramrakhyia, S. K. Ramakrishnan, T. Gervet, J. Turner, A. Gokaslan, N. Maestre, A. X. Chang, D. Batra, M. Savva, *et al.*, “Habitat-matterport 3d semantics dataset,” *arXiv preprint arXiv:2210.05633*, 2022. 2, 4, 5, 6, 14
- [17] 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. 2, 3, 4
- [18] S. Schaal, “Learning from demonstration,” in *NIPS*, 1996. 2, 3
- [19] A. Das, G. Gkioxari, S. Lee, D. Parikh, and D. Batra, “Neural Modular Control for Embodied Question Answering,” in *CoRL*, 2018. 2
- [20] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, “Learning complex dexterous manipulation with deep reinforcement learning and demonstrations,” in *RSS*, 2018. 2, 3, 12
- [21] B. Baker, I. Akkaya, P. Zhokhov, J. Huizinga, J. Tang, A. Ecoffet, B. Houghton, R. Sampedro, and J. Clune, “Video pretraining (vpt): Learning to act by watching unlabeled online videos,” *arXiv preprint arXiv:2206.11795*, 2022. 2, 3, 5, 13, 14
- [22] A. Gupta, V. Kumar, C. Lynch, S. Levine, and K. Hausman, “Relay policy learning: Solving long horizon tasks via imitation and reinforcement learning,” in *CoRL*, 2019. 2, 3
- [23] I. Uchendu, T. Xiao, Y. Lu, B. Zhu, M. Yan, J. Simon, M. Bennice, C. Fu, C. Ma, J. Jiao, S. Levine, and K. Hausman, “Jump-start reinforcement learning,” *arXiv preprint arXiv:2204.02372*, 2022. 2

- [24] D. S. Chaplot, D. Gandhi, A. Gupta, and R. Salakhutdinov, "Object goal navigation using goal-oriented semantic exploration," in *NeurIPS*, 2020. 2, 3, 6, 8
- [25] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, 1997. 2
- [26] A. Mousavian, A. Toshev, M. Fiser, J. Kosecka, A. Wahid, and J. Davidson, "Visual representations for semantic target driven navigation," in *ICRA*, 2019. 3
- [27] J. Ye, D. Batra, A. Das, and E. Wijmans, "Auxiliary Tasks and Exploration Enable ObjectNav," in *ICCV*, 2021. 3
- [28] Y. Liang, B. Chen, and S. Song, "Sscnav: Confidence-aware semantic scene completion for visual semantic navigation," 2020. 3
- [29] S. K. Ramakrishnan, D. S. Chaplot, Z. Al-Halah, J. Malik, and K. Grauman, "PONI: Potential Functions for ObjectGoal Navigation with Interaction-free Learning," in *CVPR*, 2022. 3
- [30] K. Yadav, R. Ramrakhya, A. Majumdar, V.-P. Berges, S. Kuhar, D. Batra, A. Baevski, and O. Maksymets, "Offline visual representation learning for embodied navigation," 2022. 3, 4
- [31] W. Yang, X. Wang, A. Farhadi, A. Gupta, and R. Motlaghi, "Visual semantic navigation using scene priors," in *ICLR*, 2019. 3
- [32] S. K. Ramakrishnan, D. Jayaraman, and K. Grauman, "An exploration of embodied visual exploration," *arXiv preprint arXiv:2001.02192*, 2020. 3, 5, 14
- [33] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural Networks*, vol. 21, no. 4, pp. 682–697, 2008. 3
- [34] J. Kober and J. Peters, "Policy search for motor primitives in robotics," in *NeurIPS*, 2008. 3
- [35] A. Nair, M. Dalal, A. Gupta, and S. Levine, "Accelerating online reinforcement learning with offline datasets," *arXiv preprint arXiv:2006.09359*, 2020. 3, 12
- [36] Y. Lu, K. Hausman, Y. Chebotar, M. Yan, E. Jang, A. Herzog, T. Xiao, A. Irpan, M. Khansari, D. Kalashnikov, and S. Levine, "Aw-opt: Learning robotic skills with imitation and reinforcement at scale," in *CoRL*, 2021. 3, 12
- [37] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, *et al.*, "Scalable deep reinforcement learning for vision-based robotic manipulation," in *CoRL*, 2018. 3, 12
- [38] X. B. Peng, A. Kumar, G. Zhang, and S. Levine, "Advantage-weighted regression: Simple and scalable off-policy reinforcement learning," *arXiv preprint arXiv:1910.00177*, 2019. 3, 12
- [39] Q. Wang, J. Xiong, L. Han, p. sun, H. Liu, and T. Zhang, "Exponentially weighted imitation learning for batched historical data," in *NeurIPS*, 2018. 3, 12
- [40] S. Kapturowski, G. Ostrovski, J. Quan, R. Munos, and W. Dabney, "Recurrent experience replay in distributed reinforcement learning," in *ICLR*, 2019. 3, 13
- [41] K. Cobbe, J. Hilton, O. Klimov, and J. Schulman, "Phasic policy gradient," *arXiv preprint arXiv:2009.04416*, 2020. 3
- [42] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017. 3, 4, 13
- [43] 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. 4
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *CVPR*, 2016. 4
- [45] A. Eftekhar, A. Sax, J. Malik, and A. Zamir, "Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans," in *ICCV*, 2021. 4
- [46] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, "Emerging properties in self-supervised vision transformers," in *ICCV*, 2021. 4
- [47] D. Yarats, R. Fergus, A. Lazaric, and L. Pinto, "Mastering visual continuous control: Improved data-augmented reinforcement learning," *arXiv preprint arXiv:2107.09645*, 2021. 4
- [48] E. Wijmans, A. Kadian, A. Morcos, S. Lee, I. Essa, D. Parikh, M. Savva, and D. Batra, "DD-PPPO: Learning near-perfect pointgoal navigators from 2.5 billion frames," in *ICLR*, 2020. 4
- [49] M. Deitke, E. VanderBilt, A. Herrasti, L. Weihs, J. Salvador, K. Ehsani, W. Han, E. Kolve, A. Farhadi,

A. Kembhavi, and R. Mottaghi, “Proctor: Large-scale embodied ai using procedural generation,” in *NeurIPS*, 2022. 6

- [50] H. Team, “Habitat challenge, 2022.” <https://aihabitat.org/challenge/2022>, 2020. 6
- [51] L. Weihs, U. Jain, I.-J. Liu, J. Salvador, S. Lazebnik, A. Kembhavi, and A. Schwing, “Bridging the imitation gap by adaptive insubordination,” in *NeurIPS*, 2021. the first two authors contributed equally. 6
- [52] J. Wasserman, K. Yadav, G. Chowdhary, A. Gupta, and U. Jain, “Last-mile embodied visual navigation,” in *CoRL*, 2022. 8
- [53] V. Kumar and E. Todorov, “Mujoco haptix: A virtual reality system for hand manipulation,” in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, 2015. 12
- [54] M. Hausknecht and P. Stone, “Deep Recurrent Q-Learning for Partially Observable MDPs,” in *AAAI*, 2015. 12, 13
- [55] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, “Playing atari with deep reinforcement learning,” *arXiv preprint arXiv:1312.5602*, 2013. 12
- [56] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, “High-dimensional continuous control using generalized advantage estimation,” in *ICLR*, 2016. 13

A. Prior work in RL Finetuning

A.1. DAPG [20]

Preliminaries. Rajeswaran *et al.* [20] proposed DAPG, a method which incorporates demonstrations in RL, and thus quite relevant to our methodology. DAPG first pretrains a policy using behavior cloning then finetunes the policy using an augmented RL objective (shown in Eq. (4)). DAPG proposes to use different parts of demonstrations dataset during different stages of learning for tasks involving sequence of behaviors. To do so, they add an additional term to the policy gradient objective:

$$g_{aug} = \sum_{(s,a) \in \tau \sim \pi_\theta} \nabla_\theta \log_{\pi_\theta}(a|s) A^\pi(s,a) + \sum_{(s,a) \in \tau \sim \mathcal{T}} \nabla_\theta \log_{\pi_\theta}(a|s) w(s,a) \quad (4)$$

Here $\tau \sim \pi_\theta$ is a trajectory obtained by executing the current policy, $\tau \sim \mathcal{T}$ denotes a trajectory obtained by replaying a demonstration, and $w(s,a)$ is a weighting function to alternate between imitation and reinforcement learning. DAPG uses a heuristic weighting scheme to set $w(s,a)$ to decay the auxiliary objective:

$$w(s,a) = \lambda_0 \lambda_1^k \max_{(s',a') \in \tau \sim \pi_\theta} A^{\pi_\theta}(s',a') \forall (s,a) \quad (5)$$

where λ_0 and λ_1 are hyperparameters and k is the update iteration counter. The decaying weighting term λ_1^k is used to avoid biasing the gradient towards the demonstrations data towards the end of training.

Implementation Details. [20] showed results of using DAPG on dexterous hand manipulation tasks for object relocation, in-hand manipulation, tool use, *etc.* To train the policy with behavior cloning, they use 25 demonstrations for each task gathered using the Mujoco HAPTIX system [53]. The small size of the demonstrations dataset and the observation input allows DAPG to load the demonstrations dataset in system memory which makes it feasible to compute the augmented RL objective shown above.

Challenges in adopting [20]’s setup. Compared to [20], our setup uses high-dimensional visual input (256×256 RGB observations) and 77k OBJECTNAV demonstrations for training. Following DAPG’s training implementation, storing the visual inputs for 77k demonstrations in system memory would require 2TB, which is significantly higher than what is possible on today’s systems. An alternative is to leverage on-the-fly demonstration replay during RL training. However, efficiently incorporating demonstration replay with experience collection online requires solving a systems research problem. Naively switching between online experience collection using the current policy and replay demonstrations

would require 2x the current experience collection time, overall hurting the training throughput.

A.2. Feasibility of Off-Policy RL finetuning

There are several methods for incorporating demonstrations with off-policy RL [35–39]. Algorithm 1 shows the general framework of off-policy RL (finetuning) methods.

Algorithm 1 General framework of off-policy RL algorithm

Require: π_θ : Policy, B : replay buffer, N : Rounds, I : Policy Update Iterations
for $k = 1$ to N **do**
 Trajectory $\tau \leftarrow \text{Rollout } \pi_\theta(\cdot|s)$ to collect trajectory $\{(s_1, a_1, r_1, h_1), \dots, (s_T, a_T, r_T, h_T)\}$
 $B \leftarrow \{B\} \cup \{\tau\}$
 $\pi_\theta \leftarrow \text{TrainPolicy}(\pi_\theta, B)$ for I iterations
end for

Unfortunately, most of these methods use feedforward state encoders, which is ill-posed for partially observable settings. In partially observable settings, the agent requires a state representation that combines information about the state-action trajectory so far with information about the current observation, which is typically achieved using a recurrent network.

To train a recurrent policy in an off-policy setting, the full state-action trajectories need to be stored in a replay buffer to use for training, including the hidden state h_t of the RNN. The policy update requires a sequence input for multiple time steps $[(s_t, a_t, r_t, h_t), \dots, (s_{t+l}, a_{t+l}, r_{t+l}, h_{t+l})] \sim \tau$ where l is sampled sequence length. Additionally, it is not obvious how the hidden state should be initialized for RNN updates when using a sampled sequence in the off-policy setting. Prior work DRQN [54] compared two training strategies to train a recurrent network from replayed experience:

1. **Bootstrapped Random Updates.** The episodes are sampled randomly from the replay buffer and the policy updates begin at random steps in an episode and proceed only for the unrolled timesteps. The RNN initial state is initialized to zero at the start of the update. Using randomly sampled experience better adheres to DQN’s [55] random sampling strategy, but, as a result, the RNN’s hidden state must be initialized to zero at the start of each policy update. Using zero start state allows for independent decorrelated sampling of short sequences which is important for robust optimization of neural networks. Although this can help RNN to learn to recover predictions from an initial state that mismatches with the hidden state from the collected experience but it might limit the ability of the network to rely on it’s recurrent state and exploit long term temporal correlations.
2. **Bootstrapped Sequential Updates.** The full episode replays are sampled randomly from the replay buffer and

the policy updates begin at the start of the episode. The RNN hidden state is carried forward throughout the episode. Eventhough this approach avoids the problem of finding the correct initial state it still has computational issues due to varying sequence length for each episode, and algorithmic issues due to high variance of network updates due to highly correlated nature of the states in the trajectory.

Even though using bootstrapped random updates with zero start states performed well in Atari which is mostly fully observable, R2D2 [40] found using this strategy prevents a RNN from learning long-term dependencies in more memory critical environments like DMLab. [40] proposed two strategies to train recurrent policies with randomly samples sequences:

1. **Stored State.** In this strategy, the hidden state is stored at each step in the replay and use it to initialize the network at the time of policy updates. Using stored state partially remedies the issues with initial recurrent state mismatch in zero start state strategy but it suffers from ‘representational drift’ leading to ‘recurrent state staleness’, as the stored state generated by a sufficiently old network could differ significantly from a state from the current policy.
2. **Burn-in.** In this strategy the initial part of the replay sequence is used to unroll the network and produce a start state (‘burn-in period’) and update the network on the remaining part of the sequence.

While R2D2 [40] found a combination of these strategies to be effective at mitigating the representational drift and recurrent state staleness, this increases computation and requires careful tuning of the replay sequence length m and burn-in period l .

Both [40, 54] demonstrate the issues associated with using a recurrent policy in an off-policy setting and present approaches that mitigate issues to some extent. Applying these techniques for Embodied AI tasks and off-policy RL finetuning is an open research problem and requires empirical evaluation of these strategies.

B. Training Details

B.1. Imitation Learning

We use a distributed implementation of behavior cloning by [1] for our imitation learning pretraining. Each worker collects 64 frames of experience from 8 environments parallelly by replaying actions from the demonstrations dataset. We then perform a policy update using supervised learning on 2 mini batches. For all of our IL experiments, we train the policy for 500M steps on 64 GPUs using Adam optimizer with a learning rate 1.0×10^{-3} which is linearly decayed after each policy update. Tab. 6 details the default hyperparameters used in all of our training runs.

Parameter	Value
Number of GPUs	64
Number of environments per GPU	8
Rollout length	64
Number of mini-batches per epoch	2
Optimizer	Adam
Learning rate	1.0×10^{-3}
Weight decay	0.0
Epsilon	1.0×10^{-5}
DDPIL sync fraction	0.6

Table 6. Hyperparameters used for Imitation Learning.

Parameter	Value
Number of GPUs	16
Number of environments per GPU	8
Rollout length	64
PPO epochs	2
Number of mini-batches per epoch	2
Optimizer	Adam
Weight decay	0.0
Epsilon	1.0×10^{-5}
PPO clip	0.2
Generalized advantage estimation	True
γ	0.99
τ	0.95
Value loss coefficient	0.5
Max gradient norm	0.2
DDPPO sync fraction	0.6

Table 7. Hyperparameters used for RL finetuning.

B.2. Reinforcement Learning

To train our policy using RL we use PPO with Generalized Advantage Estimation (GAE) [56]. We use a discount factor γ of 0.99 and set GAE parameter τ to 0.95. We do not use normalized advantages. To parallelize training, we use DD-PPO with 16 workers on 16 GPUs. Each worker collects 64 frames of experience from 8 environments parallelly and then performs 2 epochs of PPO update with 2 mini batches in each epoch. For all of our experiments, we RL finetune the policy for 300M steps. Tab. 7 details the default hyperparameters used in all of our training runs.

B.3. RL Finetuning using VPT

To compare with RL finetuning approach proposed in VPT [21] we implement the method in DD-PPO framework. Specifically, we initialize the critic weights to zero, replace the entropy term in PPO [42] with a KL-divergence loss between the frozen IL policy and RL policy, and decay the KL divergence loss coefficient, ρ , by a fixed factor after every

iteration. This loss term is defined as:

$$L_{kl_penalty} = \rho \text{KL}(\pi_{\theta}^{BC}, \pi_{\theta}) \quad (6)$$

where π_{θ}^{BC} is the frozen IL policy, π_{θ} is the current policy, and ρ is the loss weighting term. Following, VPT [21] we set ρ to 0.2 at the start of training and decay it by 0.995 after each policy update. We use learning rate of 1.5×10^{-5} without a learning rate decay for our VPT [21] finetuning experiments.

B.4. RL Finetuning Ablations




Figure 8. A policy pretrained on the OBJECTNAV task is used as initialization for actor weights and critic weights are initialized randomly for RL finetuning using DD-PPO. The policy performance immediately starts dropping early on during training and then recovers leading to slightly higher performance with further training.

Method	Success (\uparrow)	SPL (\uparrow)
1) IL	52.0%	20.6%
2) IL \rightarrow RL-FT	54.8%	29.1%
3) IL \rightarrow RL-FT (+ Critic Learning)	55.5%	26.7%
4) IL \rightarrow RL-FT (+ Critic Learning, Critic Decay)	59.6%	26.5%
5) IL \rightarrow RL-FT (+ Critic Learning, Actor Warmup)	58.7%	25.8%
6) PIRLNav	62.2%	28.7%

Table 8. RL-finetuning ablations on HM3D VAL [16, 32]

For ablations presented in Sec. 4.3 of the main paper (also shown in Tab. 8) we use a policy pretrained on 20k human demonstrations using IL and finetuned for 300M steps using hyperparameters from Tab. 7. We try 3 learning rates (1.5×10^{-4} , 2.5×10^{-4} , and 1.5×10^{-5}) for both IL \rightarrow RL (row 2) and IL \rightarrow RL (+ Critic Learning) (row 3) and we report the results with the one that works the best. For PIRLNav we use a starting learning rate of 2.5×10^{-4} and decay it to 1.5×10^{-5} , consistent with learning rate schedule of our best performing agent. For ablations we do not tune learning rate parameters of PIRLNav, we hypothesize tuning the parameters would help improve performance.

We find IL \rightarrow RL (row 2) works best with a smaller learning

rate but the training performance drops significantly early on, due to the critic providing poor value estimates, and recovers later as the critic improves. See Fig. 8. In contrast when using proposed two phase learning setup with the learning rate schedule we do not observe a significant drop in training performance.

C. Code

Source code and model weights for the project will be publicly released upon acceptance of the paper.