This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Towards Disturbance-Free Visual Mobile Manipulation

Tianwei Ni^{1,2,*} Kiana Ehsani² Luca Weihs^{2,†} Jordi Salvador^{2,†}

¹Université de Montréal & Mila – Quebec AI Institute ²PRIOR @ Allen Institute for AI

tianwei.ni@mila.quebec {kianae, lucaw, jordis}@allenai.org
 * Work was primarily done during internship at AI2 † Equal advising

Abstract

Deep reinforcement learning has shown promising results on an abundance of robotic tasks in simulation, including visual navigation and manipulation. Prior work generally aims to build embodied agents that solve their assigned tasks as quickly as possible, while largely ignoring the problems caused by collision with objects during interaction. This lack of prioritization is understandable: there is no inherent cost in breaking virtual objects. As a result, "well-trained" agents frequently collide with objects before achieving their primary goals, a behavior that would be catastrophic in the real world. In this paper, we study the problem of training agents to complete the task of visual mobile manipulation in the ManipulaTHOR environment while avoiding unnecessary collision (disturbance) with objects. We formulate disturbance avoidance as a penalty term in the reward function, but find that directly training with such penalized rewards often results in agents being unable to escape poor local optima. Instead, we propose a two-stage training curriculum where an agent is first allowed to freely explore and build basic competencies without penalization, after which a disturbance penalty is introduced to refine the agent's behavior. Results on testing scenes show that our curriculum not only avoids these poor local optima, but also leads to 10% absolute gains in success rate without disturbance, compared to our state-of-the-art baselines. Moreover, our curriculum is significantly more performant than a safe RL algorithm that casts collision avoidance as a constraint. Finally, we propose a novel disturbance-prediction auxiliary task that accelerates learning. ¹

1. Introduction

Advances in deep reinforcement learning (RL) for embodied agents has led to significant progress in visual navigation [61] [108, 47, [100, 80, 96, [10, [104]] and manipulation [43, [106, 24, 98, 35], [21, 72]. In this paper, we focus on

the relatively new embodied-AI problem of **visual mobile manipulation** [54, 83, [101], in particular the ArmPoint-Nav task proposed by Ehsani *et al.* [16] set in the simulated mobile manipulation framework ManipulaTHOR. In Arm-PointNav, the goal is to bring an object to a goal location. An agent must perform navigation and manipulation jointly by navigating to an object of interest, picking the object up using its attached 6-DOF robotic arm, and then carrying the object to a target location. Combining navigation and manipulation, especially when agents are expected to generalize to novel scenes and objects, is a challenging but important step towards building generally capable household robotic agents. Similarly, as for other embodied AI tasks, RL methods attain respectable performance for ArmPoint-Nav, with a baseline achieving a success rate of 62% [16].

These relatively high success rates for ArmPointNav come, however, with a significant caveat: success only requires that an agent manages to bring an object to a goal and entirely ignores whether the agent collides with other objects during interaction. Ignoring collision avoidance when measuring success is not exclusive to ArmPointNav, but ubiquitous in pure navigation [96, 104, 14], pure manipulation [52, 43], and mobile manipulation [87, 92] tasks, even despite there being collision detection capabilities in relevant simulators [88, 83]. Furthermore, the popular SPL metric (success weighed by path length) [4] may even encourage collision, as it rewards agents for taking shortcuts that may result in collisions. With the introduction of mobile manipulation, where agents have more chances to interact with objects, the catastrophic impact of collisions becomes too obvious to ignore: Fig. 1 shows a prototypical example where an agent "successfully" completes the Arm-PointNav task, but disturbs many objects in the scene. In fact, the ArmPointNav baseline of Ehsani et al. has only a 30% success rate without disturbance (collision) with other objects. This means that even a (relatively) high success rate should give little confidence that a policy can be safely deployed in the real world.

Unlike in the robotics community, where the importance of safety and collision avoidance are deeply ingrained [20].

¹Project page is at <u>https://sites.google.com/view/</u> disturb-free



Figure 1. Success with disturbance. A robotic ManipulaTHOR 16 agent attempts to complete an ArmPointNav task where it must pick up and move one object from one side of a kitchen to another. A well-trained agent (in white) successfully moves the target object (a silver pot) from a source location (stove burner) to a target location (dining table). Despite this success, the agent disturbs several objects (including a coffee maker and a toaster) when moving its arm and body. Both the coffee maker and toaster are pushed off the countertop and the coffee maker is further pushed across the kitchen, a catastrophic outcome in the real world. The figures show both RGB egocentric and top-down views, while the agent only has access to egocentric observations.

51. 69], the incentives for ignoring collisions in simulated embodied AI tasks are clear: (1) there is no inherent cost of collision in *simulation* and (2) accurate collision detectors are computationally expensive and so, as speed is one of the great advantages of simulation, environments frequently simplify or ignore the impact of collisions altogether [80] 47]. As algorithms improve and task success rates increase, this practice of ignoring collisions will prevent real world deployment where the costs of breaking objects, damaging robots, and harming humans are unacceptable [19].

In this paper, towards enabling safe deployment of real robots, we propose to train embodied agents while prioritizing what we call disturbance avoidance: namely, we require that agents avoid moving any object that is not directly related to the agents' goal. For flexibility, we only consider object positions at the start and end of an episode. For example, this allows an agent to temporarily move an object that would otherwise prevent it from reaching its goal. While collision avoidance and disturbance avoidance may, at first, seem synonymous, the problem of collision avoidance is strictly more difficult than disturbance avoidance. To see this, note that not all collisions need result in the positions of objects being changed. For instance, an agent might run into a wall (a collision) but, as the wall is not moved by this collision, the scene is undisturbed. Indeed, disturbance avoidance can be thought of as "visible collision avoidance": collisions that do not result in a visual change in the environment are ignored. We study disturbance avoidance, instead of collision avoidance, for three reasons: (1) disturbance can be efficiently measured in simulation as it requires only checking how the pose of objects has changed between time-steps, while collision detection requires measuring contact forces; (2) our agents take as input purely egocentric visual inputs and thus have a limited capacity to learn about phenomena that they cannot directly observe; (3) disturbance avoidance allows for the temporary movement of the objects so long as they are eventually moved back to their original positions, a practical necessity for some tasks.

A standard approach for encouraging safe behavior in RL agents is simply to penalize unsafe behavior [41, 77]. [89] [53]; in this vein, we modify the standard ArmPointNav reward structure by introducing a penalty for object disturbance. Perhaps surprisingly, in practice, training agents *from scratch* using this new reward structure results in extremely unstable learning. Indeed, in many cases, agents trained with this reward structure learn not to disturb any objects by terminating early without reaching the goal: a bad local optimum. We hypothesize that this object disturbance penalty discourages early exploration and thus results in a highly suboptimal policy.

Inspired by this empirical finding, we propose a simple but effective **two-stage training curriculum**: we first train the agent with original reward (without penalty), and then fine-tune the agent with penalized reward. During the first stage, the agent can learn to solve the task *with* disturbance as it has enough freedom to explore, while in the second stage the agent can learn to adjust its behavior to avoid disturbance when solving the task.

In what follows, we first focus on the original ArmPoint-Nav task with its original reward and success criteria. For this task, we show how several critical design decisions result in dramatic improvements over the existing state-ofthe-art [16]: an 11.1% absolute increase in success rate (SR) on testing scenes with novel objects (same evaluation setting below) after 20M training frames. With 45M train-

Methods	Navi- gation?	Manipu- lation?	Avoid Collision?	Framework
108	1	X	X	MDP
27.61.96.104	1	×	×	POMDP
77	1	×	1	Supervised
40, 39	1	×	1	MDP
70, 91, 64, 107	X	1	×	Supervised
<u>[65, 66]</u>	X	1	1	Supervised
52, 43, 106, 24, 84	X	1	×	MDP
58	X	1	1	POMDP
<mark>53,</mark> 98	1	1	1	MDP
<mark>92,</mark> 87	1	1	×	MDP
16	1	1	×	POMDP
Ours	1	1	✓	POMDP

Table 1. Data-driven methods for visual navigation and manipulation. We classify the methods by the studied problem (navigation, manipulation, or both, *i.e.*, mobile manipulation), whether they **learn to** avoid collision (disturbance), and their adopted algorithmic frameworks (supervised learning, MDP [6], and POMDP [5]).

ing frames, our improved model further attains a success rate of 82.7% but, critically, only achieves a 35.5% success rate without disturbance (SRwoD). We then move on to our main focus: disturbance avoidance in ArmPointNav. We find that with the same 45M frames budget, our two-stage training is much more effective than training from scratch, with 80.1% vs 18.0% in SR, and 46.5% vs 10.5% in SRwoD. In other words, two-stage training avoids performance degradation and achieves higher SRwoD than our improved baseline using the original objective (46.5% vs 35.5%). It also outperforms PPO-Lagrangian [76], a popular safe RL algorithm, by over 30% in SRwoD. Finally, we propose a new supervised auxiliary task that requires the agent to predict how its actions will disturb the environment and show that co-training with the auxiliary task can accelerate learning and increase final performance when compared to using no (or self-supervised) auxiliary tasks.

In summary, in this paper we present the following contributions: (1) we propose a disturbance avoidance objective for embodied RL agents, (2) we introduce a state-of-theart model for the original ArmPointNav task with extensive ablative experiments, (3) we provide strong empirical evidence that our two-stage curriculum can lead to agents which avoid disturbance while retaining high success rates, and (4) we offer a new auxiliary disturbance-prediction task which accelerates learning.

2. Related Work

Visual navigation and manipulation. A long history of work exists for both (stationary) tabletop manipulation using fixed-based arms [45] [74, [36], pure embodied navigation [61, [108], and mobile manipulation [102, [46, [97, [32]] [86, [62], [54]]. Much of this prior work, especially traditional methods for manipulation [81, [71, [8], [44, [73], [75], [9], [15], [60].

38, 85, requires ground-truth knowledge of objects and the environment (such as their geometry) and becomes computationally expensive in high-dimensional settings. Datadriven methods relax these assumptions and enable agents to act from visual inputs. We summarized the data-driven related work in Table 1 from several perspectives. Our work focuses on ArmPointNav [16], a visual mobile manipulation task where agents must navigate in the ManipulaTHOR environment to find an object, pick it up using an attached arm, and then bring the object to a new location. The ManipulaTHOR environment consists of a set of visually complex scenes, supports navigation with a mobile robot, and allows for object manipulation with a 6-DOF arm through clutter. ArmPointNav follows the more general POMDP framework that only has egocentric depth observations and 3D goal sensors without other state information. Our method is built upon a recurrent model-free RL baseline provided by ManipulaTHOR, and focuses on disturbance avoidance discussed below.

Collision/disturbance avoidance. In safety-critical domains, collision avoidance is extremely important. Classic methods in motion planning [20, 51, 69, 59, 13] and path planning [44, 33, 11], which provide safety guarantees, require privileged information of obstacles to avoid collision and are hard to scale to partially-observable visually complex settings. Data-driven methods learn to avoid collision from data with less privileged information (see the "collision" column in Table I for a summary). Deep-RL methods, which can learn collision avoidance from interaction within environments, can be divided into model-free and model-based approaches. Model-free methods simply introduce collision penalties in the reward function, in the form of constants when facing collision [89, 57, 98, 53, 18] or being proportional to the distances to the nearest obstacles [42, 79, 17, 39]. Model-based methods [41, 40, 39] explicitly learn a collision prediction model and use it for policy search. Our work studies disturbance, a subset of collision, where an object is moved by some distance, because disturbance is easier to compute in simulation and allows temporary displacement. We first consider a modelfree setting (adding a disturbance penalty) and show that it can perform well when paired with our two-stage training curriculum. We then add a model-based component (a disturbance prediction auxiliary task) to accelerate learning.

Constrained MDPs and safe RL. An alternative formulation of collision avoidance is to frame avoidance as *constraint* (not a fixed penalty in reward) during policy optimization, *i.e.*, frame the problem as a constrained MDP [3]. Algorithms solving constrained MDPs are frequently employed in the domain of safe RL. A popular approach in this area is to employ Lagrange multipliers in various RL algorithms [76, 55, 90, [29] allowing for the adaptive penalization of unwanted behavior. Despite being adaptive

and agnostic to the reward scale, Lagrangian methods have been shown to be sensitive to the initialization of the multipliers [1]. In our experiments, we find our approach to be more performant than a competing, Lagrangian-based, safe RL baseline.

Transfer and curriculum learning in RL. Transfer learning is widely used in deep learning to transfer knowledge from a source domain to a target domain [31], 22, [105]. Transfer learning is also popular in RL [43], such as continual RL setting [78], 37]. Our two-stage training methodology reframes the single task learning (success without disturbance) into "curriculum learning" [7], 67] on a task sequence, where the first task (stage) is to succeed with disturbance. We show that this curriculum formulation is far more effective than direct learning on the final (hard) task, as learning on the early task is much easier and bootstraps learning the final task.

Auxiliary tasks in RL. Auxiliary tasks, co-trained with the main task (maximizing the total rewards) on the shared model weights, have been shown to have the promising capability to improve sample efficiency and asymptotic performance in visual RL. Supervised auxiliary tasks provide extra information to the policy via external signals, such as depth maps [61, 94] and game internal states [48], while self-supervised/unsupervised auxiliary tasks use existing information as signals, such as auto-encoders [49, 28, 103], forward [23] and inverse dynamics [68], reward prediction [34], and contrastive learning [25, 26, 104, 50]. Our work introduces a supervised auxiliary task that predicts the "disturbance distance", namely, how greatly an agent's action will disturb the environment. Since this disturbance distance is one component of our reward, this task can be viewed as distilling knowledge of the reward's composition to the agent.

3. Towards Disturbance-Free ArmPointNav

As preliminaries, we first introduce our model architecture (Sec. 3.1) for ArmPointNav. Then we formulate the concept of disturbance avoidance along with a new, corresponding, reward structure for ArmPointNav (Sec. 3.2). Next, we define a new auxiliary task, disturbance prediction, which could improve sample efficiency in training (Sec. 3.3). Finally, we discuss the training techniques on the new objective, and introduce our two-stage training curriculum (Sec. 3.4) that is most critical to final performance.

3.1. Preliminaries: Model Architecture

For the task of ArmPointNav, given only egocentric depth observations and goal coordinates, an agent must interact with its environment so as to to pick up a target object with its arm and then navigate to place that target object in a goal location. See App. \overline{A} for more details.

We now describe our model architecture for ArmPoint-Nav, based on the original baseline proposed by Ehsani *et al.* [16]. At a timestep $t \ge 0$, the model takes as input the current egocentric depth observation o_t , which is passed through a modified ResNet18 [30] [96] to produce the embedding $z_t = \text{ResNet}(o_t)$. The model then passes this embedding z_t along with an embedding of distance to goal g_t , the encoding of previous action a_{t-1} , and the belief state b_{t-1} from the previous timestep, to a single-layer GRU [12] to produce the current belief state $b_t = \text{GRU}(b_{t-1}, \text{ResNet}(o_t), g_t, a_{t-1})$. Finally, following an actor-critic formulation, a linear layer is applied to the belief state b_t to produce the agent's policy $\pi(b_t)$ (i.e, a distribution over actions) and an estimate of the value of the agent's current state $V(b_t)$.

Fig. 2 shows a summary of this architecture, and the new design choices we made compared to the original baseline. This agent is trained using the PPO algorithm [82, 96] to maximize the discounted cumulative rewards $\sum_{t=0}^{T} \gamma^t r_t$, with $\gamma = 0.99$ and T = 200. The reward function r_t is defined in Eq. 1 in [16].

3.2. Formulation of Disturbance Avoidance

Now we focus on the goal of training agents that are able to avoid disturbance. One of the simplest strategies for discouraging unwanted behavior in RL is to incorporate a penalty into the reward function for this behavior. We consequently add a penalty to the original reward function r_t (Eq. 1 in [16]) to define a new reward function r'_t :

$$r'_t = r_t + \lambda_{\text{disturb}} (d^{\text{objects}}_{t-1} - d^{\text{objects}}_t) , \qquad (1)$$

where $\lambda_{\text{disturb}} > 0$ is a coefficient controlling the magnitude of the disturbance penalty, and d_t^{objects} is the sum of 3D Euclidean distances of all objects (except the target object) from their initial locations at time t.

The **disturbance-free objective** is now defined as the discounted cumulative sum of these new rewards:

$$\sum_{t=0}^{T} \gamma^{t} r_{t}' = \sum_{t=0}^{T} \gamma^{t} r_{t} + \gamma^{t} \lambda_{\text{disturb}} (d_{t-1}^{\text{objects}} - d_{t}^{\text{objects}}) .$$
⁽²⁾

Notice that if $\gamma = 1$ then the above sum telescopes to simply equal $(\sum_{t=0}^{T} r_t) - \lambda_{\text{disturb}} d_T^{\text{objects}}$. Thus, up to discounting, the disturbance-free objective can be interpreted as the original ArmPointNav objective with a soft constraint on the final total disturbance distance. This new objective discourages the agent from ending the episode with objects out of their original positions, but allows the agent to temporarily move objects as long as the agent eventually moves them back into their original locations. This behavior emerges in training (see Sec. [4.2]).

3.3. Disturbance Prediction as an Auxiliary Task

Beyond passing knowledge indirectly via a model-free method (*i.e.*, with disturbance penalty as in Eq. (2), to im-



Figure 2. Our model architecture for ArmPointNav. We make several improvements to the existing SoTA baseline from [16]. These improvements include: replacing the existing shallow CNN with a ResNet18, adding an embedding of the previous action $a_{t-1} \in \mathcal{A}$ to the agent's inputs, and using polar, instead of Cartesian, goal coordinates $g_t \in \mathbb{R}^3$. These changes greatly increase performance (see Sec. [4.1]). Building on the work of Ye *et al.* [104], we also enable support for training with arbitrary self-supervised auxiliary tasks given current belief b_t for sample efficiency. Finally, we add a specific auxiliary task head for our novel next-step disturbance prediction task $\hat{c}_{t+1} \in \mathbb{R}^{|\mathcal{A}|}$.

prove sample efficiency, we also consider a model-based approach that explicitly requires the agent to predict object disturbance as an auxiliary task. Formally, given its current belief b_t about the environment and the action a_t it has decided to enact, the agent must predict the probability \hat{c}_{t+1} of current action disturbing the environment. To produce this probability estimate in practice, we use an MLP denoted as Disturb:

$$Disturb(b_t, a_t) = \hat{c}_{t+1} \in [0, 1].$$
 (3)

We obtain the ground-truth *binary* disturbance signal c_{t+1} from the simulated environment: it is calculated by thresholding the current change in the disturbance distance:

$$c_{t+1} = \mathbb{1}\left(d_{t+1}^{\text{objects}} - d_t^{\text{objects}} \ge \tau\right) \in \{0, 1\}, \quad (4)$$

where $\tau\approx 0.001$ m in our case.

As most (\approx 90%) actions taken by the agent during training do not result in disturbance, there is a significant class imbalance for this task formulation. To mitigate this imbalance, we leverage the Focal loss [56], a modified cross entropy loss designed for class-imbalanced prediction. For every step t in training, we compute the total loss for our agent as the sum of the usual actor, critic, and entropy losses from PPO as well as the auxiliary disturbance prediction loss $\mathcal{L}_{\text{Focal}}(\hat{c}_t, c_t)$.

Our auxiliary task works in concert with our new objective: our auxiliary task can directly teach the agent to recognize the types of actions that result in disturbance, while our new objective encourages the agent to avoid such actions.

3.4. Two-Stage Training Curriculum



Figure 3. Our two-stage training curriculum. r_t and r'_t (Eq. 1) are the original and new reward function, respectively. The main model refers to the model components other than those of auxiliary tasks. Co-training with auxiliary task is optional.

We train an agent using RL on the disturbance-free objective (Eq. 2), with an option of co-training any auxiliary task (*e.g.*, Sec. 3.3). Directly training to optimize for the new objective *from scratch* is straightforward but, in practice, suffers from extreme instability with some "trained" models achieving near-zero success rates (see Sec. 4.2). As the disturbance penalty discourages the agent from interacting with the objects, we hypothesize that it may also prevent the agent from sufficiently exploring potential strategies before settling on a conservative approach: giving up partway

through an episode to avoid potential object disturbance.

Inspired by the empirical findings and work in curriculum learning, we propose a two-staged training curriculum for learning with the disturbance-free objective (see Fig. 3). In the first stage, the agent is trained with the original reward function r_t (potentially co-trained with auxiliary tasks). As the agent is not penalized for causing disturbance, this allows the agent to freely explore and thus, as we show in our experiments, enables the agent to learn a policy with a high success rate. In the second stage, we fine-tune the agent from the previously learned model in the first stage using our new reward function r'_t (Eq. 1). In this stage, the agent learns to refine its behavior to better avoid disturbance, without sacrificing performance. Intuitively, we decompose the disturbance-free objective r' into a task sequence, with the first task being r and the second being r'. The first task is easier to learn from scratch, and its goal is closely aligned with the second task.

4. Experiments

We evaluate our method on the ArmPointNav task within the manipulation framework ManipulaTHOR [16] set within AI2-THOR simulator [47]. Following [16], in all experiments we train our agents on the 19 training scenes, tune model hyperparameters (and report the ablation study results) on the 5 validation scenes, and finally report our best-validation-model results on the 5 testing scenes. Each scene has 720 data points (episodes). Due to space constraints, we report the agent performance when faced with *novel* objects (360 data points) in the main paper and report results on the less-challenging *seen* objects subset in App. [D]. For details on the ManipulaTHOR environment, see App. [A].

We focus on two primary metrics, **Success Rate (SR)** and **Success Rate without Disturbance (SRwoD)**², measuring the original objective and the disturbance-free objective, respectively. For each episode, we consider it *successful without disturbance* if it is successful and the *final disturbance distance* $d_T^{objects}$ is less than a threshold (≈ 0.01 m). We use AllenAct [93] as our training framework. We provide more implementation details, such as hyperparameters, in App. B All training code and model weights will be made open-source. Qualitative videos can be found in the supplementary materials.

In Sec. 4.1, we evaluate several small but critical design decisions that allow us to dramatically improve the SR of the baseline model from Ehsani *et al.* In Sec. 4.2, we describe how our two-stage training approach from Sec. 3.4 allows for large gains in SRwoD without sacrificing SR in the final disturbance-free setting.

Visual Encoder	Normalized Advantage?	Previous Action?	Goal Coordinate	SR (%)	SRwoD (%)
CNN	1	X	Cartesian	55.8	12.3
CNN	×	X	Cartesian	59.9	16.7
CNN	×	1	Cartesian	52.8	12.7
CNN	1	1	Polar	68.9	18.3
CNN	×	X	Polar	64.2	18.3
CNN	×	1	Polar	66.5	15.3
ResNet	1	X	Cartesian	60.2	13.1
ResNet	×	X	Cartesian	62.4	14.3
ResNet	×	1	Cartesian	63.6	15.9
ResNet	1	1	Polar	58.4	13.3
ResNet	×	X	Polar	67.7	14.6
ResNet	×	1	Polar	73.6	18.1

Table 2. Ablating our improved baselines on validation scenes. Using a ResNet18 visual encoder, un-normalized advantages, polar goal coordinates, and adding previous actions, can significantly increase the performance of the baselines on validation scenes with novel objects. Reported metrics include Success Rate (SR) and Success Rate without Disturbance (SRwoD). The final improved baseline (last row) outperforms the original baseline from [16] (first row).

4.1. Improved Baseline for ArmPointNav

As described in Sec. 3.1, we improve and extend the original architecture for ArmPointNav. Below are the details of our decision choices. First, inspired by recent works [96, 95, 104, 61], we replace the simple visual encoder with a modified ResNet18 [30], use un-normalized advantage estimation in PPO [82],³ and add previous actions as an input to the GRU model. Moreover, we replace the original *unsigned* Cartesian coordinates (|x|, |y|, |z|) in ArmPointNav by polar coordinates $(\rho, \alpha, \beta)^4$ as relative 3D goal coordinates.

Table 2 shows how these various design decisions impact the model performance on the *validation* set. All models are trained for 20M simulation steps, following Ehsani *et al.* Combining all the above modifications, we obtain a new baseline model (last row) that greatly outperforms the previous SoTA (row $\#1)^{5}$ by 17% absolute points in SR in the same training setting.

4.2. Results in the Disturbance-Free Setting

Before moving to our main results for the disturbancefree setting, we describe one additional design decision: we have enlarged the original action space $\mathcal{A}_{\text{small}}$ in ArmPoint-Nav into $\mathcal{A}_{\text{large}}$ so as to include camera and arm rotation actions. We enlarge the action space in this way as, in a qualitative analysis of model failures, we found that agents

 $^{^{2}}$ Note that this metric is abbreviated as SRwD in the original paper 16. However we find "SRwoD" to be less ambiguous.

³Please refer to batch-wise normalization in advantage in [95].

⁴where $x = \rho \cos \alpha \sin \beta$, $y = \rho \sin \alpha \sin \beta$, $z = \rho \cos \alpha$ following the standard coordinate conversion.

⁵As we use a newer, more physically accurate, version of AI2-THOR [47] and smaller batch sizes, the results of our re-implemented baseline are slightly lower than those in the original paper (SR: 61.7% v.s. 62.1%; SRwoD: 29.8% v.s. 32.7%).

Stand Damand	Initial	Frames	A ur Tealr	SR (%)		SRwoD (%)		
Stage Reward			Aux Task	Mean	IQM	Mean	IQM	
Ι	r	scratch	20M	None (Original)	61.7	-	29.8	-
Ι	r	scratch	20M	None (New)	73.3	-	31.7	-
Ι	r	scratch	20M	CPC A [25, 104]	74.1	-	31.9	-
Ι	r	scratch	20M	Inv. Dyn. [68 , 104]	76.8	-	35.0	-
Ι	r	scratch	20M	Disturb (Ours)	78.3	-	34.0	-
Ι	r	scratch	45M	None (New)	82.7	82.1	35.5	35.3
Ι	r	scratch	45M	CPC A	81.4	81.6	36.2	36.9
Ι	r	scratch	45M	Inv. Dyn.	67.8	80.9	29.4	34.9
Ι	r	scratch	45M	Disturb	83.5	82.7	37.2	36.4
Ι	r'	scratch	45M	None (New)	18.0	4.8	10.5	3.0
Ι	r'	scratch	45M	CPC A	18.2	3.6	11.1	2.1
Ι	r'	scratch	45M	Inv. Dyn.	30.4	25.9	18.4	15.6
Ι	r'	scratch	45M	Disturb	1.4	1.3	0.9	0.8
PPO-L	agrangian [76 ($\lambda_0 = 1.0$)	45M	None (New)	30.8	36.6	15.2	18.3
PPO-L	agrangian ($\lambda_0 = 15.0)$	45M	None (New)	0.0	0.0	0.0	0.0
II	r'	finetune	20M+25M	None (New)	80.1	79.9	46.5	45.9
II	r'	finetune	20M+25M	CPC A	79.1	78.9	46.7	46.6
II	r'	finetune	20M+25M	Inv. Dyn.	79.6	79.8	46.9	47.1
Π	r'	finetune	20M+25M	Disturb	81.3	81.4	47.1	46.6

Table 3. Main results on testing scenes with novel objects using the large action space A_{large} . Each method is labeled by its stage in our curriculum (Fig. 3.4), the reward it received (*r* for original reward; *r'* for new reward), the weight initialization (from scratch or fine-tuned), number of training frames, and what auxiliary task it used. For none auxiliary task, "original" refers to the original baseline, and "new" refers to our improved variant. "Mean" column shows the averages over 5 random seeds while "IQM" column shows the averages over the 3 seeds with median performance, to reduce the effect of outliers, suggested by [2] (see Fig. 6) for plots with confidence intervals).

often appeared to disturb objects in part due to (1) a lack of degrees of freedom in their arm movement and (2) an inability to change their camera's viewing angle to see objects they might disturb. We also ablate this decision in App. D and find that it indeed has an impact on performance.

Table **3** summarizes our main results (using A_{large}) when evaluating models on testing scenes with novel objects. We consider 4 training scenarios, corresponding to the 4 main blocks in the table: (**Block 1**) training from scratch with original objective r for 20M frames; (**Block 2**) training from scratch with original objective r for 45M frames; (**Block 3**) training from scratch with new objective r' for 45M frames; and (**Block 5**) fine-tuning from Block 1 with r' for 25M frames (45M total).

Each row shows the average result over **5** seeds for reproducibility. In Block 3 and 5, we use a fixed penalty coefficient $\lambda_{disturb} = 15.0$ for r' after tuning on the validation set (see ablation study in App. D). We also run a common safe RL baseline, PPO-Lagrangian [76] (Block 4), with two initial multiplier values λ_0 (see App. C for details).

Auxiliary tasks can improve sample efficiency (Block 1 and 2). For comparison, we consider three different auxiliary tasks: our proposed disturbance prediction task (Sec. 3.3) and two self-supervised tasks, Inverse Dynamics Prediction [68, [104] and the contrastive CPC | A

method [25] 104]. As seen by examining the results in Block 1 (which shows results after 20M training steps), sample efficiency improves when co-training with CPC|A and our task. This gain in sample efficiency does not, however, lead to substantial gains in performance after 45M training steps (Block 2). The agents using our disturbance prediction task perform the best in both regimes, although with a narrower advantage after 45M training steps. This demonstrates that an auxiliary task can indeed enhance *sample efficiency*, but not necessarily asymptotic performance. In our initial experiments we found that combining multiple auxiliary tasks did not meaningfully improve results and so we report using only a single auxiliary task.

Training from scratch learns to stop early with poor success rate (Block 2 and 3). Now we move on to the disturbance-free setting. The simplest way to train an agent that avoids disturbance is to directly train the policy with the new reward r' (Eq. []) from scratch (*i.e.*, Block 3). However, even when co-training auxiliary tasks, the agents simply fail to learn a reasonable policy in most seeds, with much worse average SR (and also SRwoD) when compared to Block 2 (trained with r). In fact, 16 out of 20 trials in Block 3 totally fail, with <10% SR. We investigate these failed trials and find that these agents pick up the target objects in 94.1% of episodes, but only choose to terminate the episode *shortly*

after the pick-up within 7.1 steps (recall that the total horizon T = 200). This means that these agents learn **a bad local optimum**: pick up the target object to get the pick-up reward bonus, and then immediately terminate the episode to avoid any disturbance penalty. Such degradation in SR can be explained as a side effect of disturbance avoidance as suggested in Sec. [3.4]

Two-stage training achieves higher SRwoD without sacrificing SR (Blocks 2, 3, and 5), and much better than PPO-Lagrangian (Block 4), with impressive robustness. Our two-stage training curriculum (Block 5) allows agents to avoid degradation in SR, compared to training from scratch (Block 3), while also achieving a much higher SRwoD by $\sim 10\%$ with similar SR, compared to training on the original reward (Block 2). It is also significantly more performant than the safe RL baseline, PPO-Lagrangian, which is sensitive to the initial coefficient. In the appendix Fig. 6 we show that our approach is also more robust to different seeds than training from scratch or using PPO-Lagrangian. Note that all agents were trained for 45M frames. Our proposed curriculum is highly effective and robust, easy to implement, and can be used with auxiliary task co-training. Because of these advantages, as embodied-AI tasks begin to take disturbance avoidance more seriously in an effort to enable real-world deployment, we expect that our approach's simplicity and robustness will be its great advantages: any researcher in embodied AI can easily leverage our training approach to enable disturbance avoidance in their models.

Emergence of temporary displacement. As noted previously, our reward structure r' allows agents to disturb objects, so long as they eventually move those objects back to their original positions (approximately). Perhaps surprisingly, we find that this behavior emerges in testing scenes: in a quantitative analysis, we find that the agents in Block 5 (Table 3) learn to temporarily move the other objects (and then recover their positions) in $\approx 5\%$ of episodes.

Success rate at various disturbance thresholds. To show the detailed results on the relationship between disturbance distance and success rate, we plot the success without disturbance curves in Fig. 4 for our best performing models. We can see that Block 5 agents indeed outperform Block 2 agents in all the disturbance distance thresholds less than 2.5 meters, reaching around 80% success rate when the disturbance distance threshold is less than 1 meter or at most 2 objects are disturbed.

Qualitative results. Fig. 5 shows an example of how our disturbance-free approach can achieve better performance on the task of ArmPointNav. See our supplementary materials for more qualitative results.

5. Conclusion

This paper highlights the importance of disturbance (collision) avoidance to embodied AI, with a focus on visual



Figure 4. Success rate without disturbance curves. Each method is labeled by its training setting (Block 2 or 5) and the auxiliary task it uses, all trained with 45M frames. The x-axis in the left figure is the disturbance distance $(d_T^{\text{oljects}}, \text{DD} \text{ in short})$, and y-axis is the % of episodes that are both successful and have DD lower than the x value. The intersection of the vertical line at DD = 0.01 and each curve is \approx SRwoD. Similarly, the right figure uses the number of disturbed objects as its x-axis.



Figure 5. Qualitative improvements. The task is to pick up the mug on the table starting from the stove. Our disturb-free agent observes the chairs and plans accordingly to avoid disturbing the scene. It successfully picks up the object. In contrast, our baseline agent disturbs the chairs with its body and the bowl on the table with its arm. The steps it spent colliding with objects prevent it from picking up the mug before the episode finishes.

mobile manipulation. We first formalize the objective of disturbance avoidance for RL agents, and then provide extensive evidence that our two-stage training curriculum is much more effective and robust than training from scratch on that objective, leading to state-of-the-art performance on success rates without disturbance in ArmPointNav task. Moreover, we propose a new auxiliary task of disturbance prediction to improve sample efficiency. Although we evaluate our method only in ArmPointNav task due to the scarcity of benchmarks, we believe that all the components of our method, including the disturbance-free objective, the two-stage training curriculum, and the new auxiliary task, are general, and could be applied to other tasks to accelerate safe deployment of robots in the real world.

Acknowledgement

We thank Aniruddha Kembhavi and Roozbeh Mottaghi for their insightful feedback on the early draft of the paper.

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *Proceedings* of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, 2017. 4,15
- [2] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron C Courville, and Marc Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in Neural Information Processing Systems*, 34, 2021. 7.16
- [3] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999. 3 [4]
- [4] Peter Anderson, Angel X. Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, and Amir Roshan Zamir. On evaluation of embodied navigation agents. *arXiv preprint arXiv:1807.06757*, 2018.
- [5] Karl Johan Åström. Optimal control of markov processes with incomplete state information i. *Journal of Mathematical Analysis and Applications*, 10:174–205, 1965. 3
- [6] Richard Bellman. A markovian decision process. *Journal* of mathematics and mechanics, 6(5):679–684, 1957. 3
- [7] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the* 26th annual international conference on machine learning, pages 41–48, 2009. [4]
- [8] Dmitry Berenson, Rosen Diankov, Koichi Nishiwaki, Satoshi Kagami, and James Kuffner. Grasp planning in complex scenes. In 2007 7th IEEE-RAS International Conference on Humanoid Robots, November 29th - December 1st, Pittsburgh, PA, USA, 2007. 3
- [9] Oliver Brock and Oussama Khatib. Mobile manipulation: Collision-free path modification and motion coordination. In *Proceedings of the 2nd International Conference on Computational Engineering in Systems Applications*. Citeseer, 1998. 3
- [10] Devendra Singh Chaplot, Ruslan Salakhutdinov, Abhinav Gupta, and Saurabh Gupta. Neural topological SLAM for visual navigation. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, 2020. [1]
- [11] Andrea Cherubini and François Chaumette. Visual navigation with obstacle avoidance. In 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2011, San Francisco, CA, USA, September 25-30, 2011, 2011.
- [12] Junyoung Chung, Çaglar Gülçehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555, 2014.
- [13] Shreyansh Daftry, Sam Zeng, J. Andrew Bagnell, and Martial Hebert. Introspective perception: Learning to predict failures in vision systems. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2016, Daejeon, South Korea, October 9-14, 2016, 2016.

- [14] Matt Deitke, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Dustin Schwenk, Eli VanderBilt, Matthew Wallingford, Luca Weihs, Mark Yatskar, and Ali Farhadi. Robothor: An open simulation-to-real embodied AI platform. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, 2020. []
- [15] Mehmet Remzi Dogar and Siddhartha S. Srinivasa. A framework for push-grasping in clutter. In *Robotics: Sci*ence and Systems VII, University of Southern California, Los Angeles, CA, USA, June 27-30, 2011, 2011.
- [16] Kiana Ehsani, Winson Han, Alvaro Herrasti, Eli Vander-Bilt, Luca Weihs, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Manipulathor: A framework for visual object manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, 2021. [12] [3] [4] [5] [6] [14]
- [17] Michael Everett, Yu Fan Chen, and Jonathan P. How. Collision avoidance in pedestrian-rich environments with deep reinforcement learning. *IEEE Access*, 2021. 3
- [18] Shumin Feng, Bijo Sebastian, and Pinhas Ben-Tzvi. A collision avoidance method based on deep reinforcement learning. *Robotics*, 2021. 3
- [19] Fabrizio Flacco, Torsten Kröger, Alessandro De Luca, and Oussama Khatib. Depth space approach to human-robot collision avoidance. In *IEEE International Conference on Robotics and Automation, ICRA 2012, 14-18 May, 2012, St. Paul, Minnesota, USA*, 2012.
- [20] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 4(1):23–33, 1997.
 3
- [21] Chuang Gan, Jeremy Schwartz, Seth Alter, Martin Schrimpf, James Traer, Julian De Freitas, Jonas Kubilius, Abhishek Bhandwaldar, Nick Haber, Megumi Sano, Kuno Kim, Elias Wang, Damian Mrowca, Michael Lingelbach, Aidan Curtis, Kevin T. Feigelis, Daniel M. Bear, Dan Gutfreund, David D. Cox, James J. DiCarlo, Josh H. McDermott, Joshua B. Tenenbaum, and Daniel L. K. Yamins. Threedworld: A platform for interactive multi-modal physical simulation. arXiv preprint arXiv:2007.04954, 2020. [1]
- [22] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In 2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014, 2014.
- [23] Karol Gregor, Danilo Jimenez Rezende, Frederic Besse, Yan Wu, Hamza Merzic, and Aäron van den Oord. Shaping belief states with generative environment models for RL. In Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada, 2019. 4
- [24] Marcus Gualtieri and Robert Platt Jr. Learning 6-dof grasping and pick-place using attention focus. In 2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings, 2018. [1]

- [25] Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Bernardo A. Pires, Toby Pohlen, and Rémi Munos. Neural predictive belief representations. *arXiv preprint arXiv:1811.06407*, 2018. 4, 7, 17
- [26] Zhaohan Daniel Guo, Bernardo Ávila Pires, Bilal Piot, Jean-Bastien Grill, Florent Altché, Rémi Munos, and Mohammad Gheshlaghi Azar. Bootstrap latent-predictive representations for multitask reinforcement learning. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, 2020. 4
- [27] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2616–2625, 2017. 3
- [28] David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, 2018. 4
- [29] Sehoon Ha, Peng Xu, Zhenyu Tan, Sergey Levine, and Jie Tan. Learning to walk in the real world with minimal human effort. *arXiv preprint arXiv:2002.08550*, 2020. 3
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016, 2016.
 [4] 6
- [31] Geoffrey E Hinton and Ruslan R Salakhutdinov. Reducing the dimensionality of data with neural networks. *science*, 313(5786):504–507, 2006. 4
- [32] Armin Hornung, Mike Phillips, Edward Gil Jones, Maren Bennewitz, Maxim Likhachev, and Sachin Chitta. Navigation in three-dimensional cluttered environments for mobile manipulation. In *IEEE International Conference on Robotics and Automation, ICRA 2012, 14-18 May, 2012, St. Paul, Minnesota, USA*, 2012. 3
- [33] Yong Koo Hwang and Narendra Ahuja. A potential field approach to path planning. *IEEE Trans. Robotics Autom.*, 1992. 3
- [34] Max Jaderberg, Volodymyr Mnih, Wojciech Marian Czarnecki, Tom Schaul, Joel Z. Leibo, David Silver, and Koray Kavukcuoglu. Reinforcement learning with unsupervised auxiliary tasks. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017. 4
- [35] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark and learning environment. *IEEE Robotics Autom. Lett.*, 2020. 1
- [36] Haiyang Jin, Qing Chen, Zhixian Chen, Ying Hu, and Jianwei Zhang. Multi-leapmotion sensor based demonstration for robotic refine tabletop object manipulation task. *CAAI Trans. Intell. Technol.*, 2016. 3
- [37] Ryan Julian, Benjamin Swanson, Gaurav Sukhatme, Sergey Levine, Chelsea Finn, and Karol Hausman. Never stop

learning: The effectiveness of fine-tuning in robotic reinforcement learning. In 4th Conference on Robot Learning, CoRL 2020, 16-18 November 2020, Virtual Event / Cambridge, MA, USA, 2020. [4]

- [38] Leslie Pack Kaelbling and Tomás Lozano-Pérez. Unifying perception, estimation and action for mobile manipulation via belief space planning. In *IEEE International Conference on Robotics and Automation, ICRA 2012, 14-18 May,* 2012, St. Paul, Minnesota, USA, 2012. 3
- [39] Gregory Kahn, Pieter Abbeel, and Sergey Levine. BADGR: an autonomous self-supervised learning-based navigation system. *IEEE Robotics Autom. Lett.*, 2021. 3
- [40] Gregory Kahn, Adam Villaflor, Bosen Ding, Pieter Abbeel, and Sergey Levine. Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation. In 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018, 2018.
- [41] Gregory Kahn, Adam Villaflor, Vitchyr Pong, Pieter Abbeel, and Sergey Levine. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017. 2.3
- [42] Gregory Kahn, Tianhao Zhang, Sergey Levine, and Pieter Abbeel. PLATO: policy learning using adaptive trajectory optimization. In 2017 IEEE International Conference on Robotics and Automation, ICRA 2017, Singapore, Singapore, May 29 - June 3, 2017, 2017. 3
- [43] Dmitry Kalashnikov, Alex Irpan, Peter Pastor, Julian Ibarz, Alexander Herzog, Eric Jang, Deirdre Quillen, Ethan Holly, Mrinal Kalakrishnan, Vincent Vanhoucke, and Sergey Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. arXiv preprint arXiv:1806.10293, 2018. [13] [4]
- [44] Oussama Khatib. Real-time obstacle avoidance for manipulators and mobile robots. In *Autonomous robot vehicles*, pages 396–404. Springer, 1986. 3
- [45] Oussama Khatib. Inertial properties in robotic manipulation: An object-level framework. Int. J. Robotics Res., 1995. 3
- [46] Oussama Khatib. Mobile manipulation: The robotic assistant. *Robotics Auton. Syst.*, 1999. 3
- [47] Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-THOR: an interactive 3d environment for visual AI. arXiv preprint arXiv:1712.05474, 2017. [] [2] 6 [14]
- [48] Guillaume Lample and Devendra Singh Chaplot. Playing FPS games with deep reinforcement learning. In Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA, 2017. 4
- [49] Sascha Lange and Martin A. Riedmiller. Deep auto-encoder neural networks in reinforcement learning. In International Joint Conference on Neural Networks, IJCNN 2010, Barcelona, Spain, 18-23 July, 2010, 2010.
- [50] Michael Laskin, Aravind Srinivas, and Pieter Abbeel. CURL: contrastive unsupervised representations for reinforcement learning. In *Proceedings of the 37th Interna-*

tional Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, 2020. [4]

- [51] Steven M LaValle and James J Kuffner Jr. Randomized kinodynamic planning. *The international journal of robotics research*, 20(5):378–400, 2001. [2] 3
- [52] Sergey Levine, Peter Pastor, Alex Krizhevsky, Julian Ibarz, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *Int. J. Robotics Res.*, 2018. [1]
- [53] Chengshu Li, Fei Xia, Roberto Martín-Martín, and Silvio Savarese. HRL4IN: hierarchical reinforcement learning for interactive navigation with mobile manipulators. In 3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 November 1, 2019, Proceedings, 2019.
 [2] [3]
- [54] Zhijun Li, Ting Zhao, Fei Chen, Yingbai Hu, Chun-Yi Su, and Toshio Fukuda. Reinforcement learning of manipulation and grasping using dynamical movement primitives for a humanoidlike mobile manipulator. *IEEE/ASME Transactions on Mechatronics*, 23(1):121–131, 2017. [1] 3
- [55] Qingkai Liang, Fanyu Que, and Eytan Modiano. Accelerated primal-dual policy optimization for safe reinforcement learning. *arXiv preprint arXiv:1802.06480*, 2018. 3
- [56] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision, ICCV* 2017, Venice, Italy, October 22-29, 2017, 2017. [5], [14]
- [57] Pinxin Long, Tingxiang Fan, Xinyi Liao, Wenxi Liu, Hao Zhang, and Jia Pan. Towards optimally decentralized multirobot collision avoidance via deep reinforcement learning. In 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018, 2018. 3
- [58] Jeffrey Mahler and Ken Goldberg. Learning deep policies for robot bin picking by simulating robust grasping sequences. In *1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings,* 2017. 3
- [59] Anirudha Majumdar and Russ Tedrake. Funnel libraries for real-time robust feedback motion planning. *The International Journal of Robotics Research*, 36(8):947–982, 2017.
 3
- [60] Andrew T. Miller, Steffen Knoop, Henrik I. Christensen, and Peter K. Allen. Automatic grasp planning using shape primitives. In Proceedings of the 2003 IEEE International Conference on Robotics and Automation, ICRA 2003, September 14-19, 2003, Taipei, Taiwan, 2003. 3
- [61] Piotr Mirowski, Razvan Pascanu, Fabio Viola, Hubert Soyer, Andy Ballard, Andrea Banino, Misha Denil, Ross Goroshin, Laurent Sifre, Koray Kavukcuoglu, Dharshan Kumaran, and Raia Hadsell. Learning to navigate in complex environments. In 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017.
 [1] [3] [4] [6]
- [62] Mayank Mittal, David Hoeller, Farbod Farshidian, Marco Hutter, and Animesh Garg. Articulated object interaction

in unknown scenes with whole-body mobile manipulation. *arXiv preprint arXiv:2103.10534*, 2021. 3

- [63] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin A. Riedmiller. Playing atari with deep reinforcement learning. arXiv preprint arXiv:1312.5602, 2013. 14
- [64] Douglas Morrison, Juxi Leitner, and Peter Corke. Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018,* 2018.
- [65] Arsalan Mousavian, Clemens Eppner, and Dieter Fox. 6dof graspnet: Variational grasp generation for object manipulation. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, 2019. 3
- [66] Adithyavairavan Murali, Arsalan Mousavian, Clemens Eppner, Chris Paxton, and Dieter Fox. 6-dof grasping for target-driven object manipulation in clutter. In 2020 IEEE International Conference on Robotics and Automation, ICRA 2020, Paris, France, May 31 - August 31, 2020, 2020. 3
- [67] Sanmit Narvekar, Bei Peng, Matteo Leonetti, Jivko Sinapov, Matthew E Taylor, and Peter Stone. Curriculum learning for reinforcement learning domains: A framework and survey. arXiv preprint arXiv:2003.04960, 2020. 4
- [68] Deepak Pathak, Pulkit Agrawal, Alexei A. Efros, and Trevor Darrell. Curiosity-driven exploration by selfsupervised prediction. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017, 2017.* [4] [7] [17]
- [69] Theodore J. Perkins and Andrew G. Barto. Lyapunov design for safe reinforcement learning. J. Mach. Learn. Res., 2002. 2.3
- [70] Lerrel Pinto and Abhinav Gupta. Supersizing selfsupervision: Learning to grasp from 50k tries and 700 robot hours. In 2016 IEEE International Conference on Robotics and Automation, ICRA 2016, Stockholm, Sweden, May 16-21, 2016, 2016. 3
- [71] Domenico Prattichizzo and Jeffrey C Trinkle. Grasping. In Springer handbook of robotics, pages 955–988. Springer, 2016. 3
- [72] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. Virtualhome: Simulating household activities via programs. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, 2018.
- [73] Sean Quinlan and Oussama Khatib. Elastic bands: Connecting path planning and control. In Proceedings of the 1993 IEEE International Conference on Robotics and Automation, Atlanta, Georgia, USA, May 1993, 1993.
- [74] Krisnawan Rahardja and Akio Kosaka. Vision-based binpicking: recognition and localization of multiple complex objects using simple visual cues. In Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS 1996, November 4-8, 1996, Osaka, Japan, 1996. 3

- [75] Nathan D. Ratliff, Matthew Zucker, J. Andrew Bagnell, and Siddhartha S. Srinivasa. CHOMP: gradient optimization techniques for efficient motion planning. In 2009 IEEE International Conference on Robotics and Automation, ICRA 2009, Kobe, Japan, May 12-17, 2009, 2009.
- [76] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking Safe Exploration in Deep Reinforcement Learning. 2019. 3, 7, 14, 17
- [77] Charles Richter and Nicholas Roy. Safe visual navigation via deep learning and novelty detection. In *Robotics: Sci*ence and Systems XIII, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, July 12-16, 2017, 2017. 2.3
- [78] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. arXiv preprint arXiv:1606.04671, 2016. 4
- [79] Bianca Sangiovanni, Angelo Rendiniello, Gian Paolo Incremona, Antonella Ferrara, and Marco Piastra. Deep reinforcement learning for collision avoidance of robotic manipulators. In 16th European Control Conference, ECC 2018, Limassol, Cyprus, June 12-15, 2018, 2018.
- [80] Manolis Savva, Jitendra Malik, Devi Parikh, Dhruv Batra, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, and Vladlen Koltun. Habitat: A platform for embodied AI research. In 2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019, 2019. [1] [2]
- [81] John Schulman, Jonathan Ho, Alex X. Lee, Ibrahim Awwal, Henry Bradlow, and Pieter Abbeel. Finding locally optimal, collision-free trajectories with sequential convex optimization. In *Robotics: Science and Systems IX, Technische Universität Berlin, Berlin, Germany, June 24 - June 28, 2013*, 2013. 3
- [82] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347, 2017. [4] [6]
- [83] Bokui Shen, Fei Xia, Chengshu Li, Roberto Martín-Martín, Linxi Fan, Guanzhi Wang, Shyamal Buch, Claudia D'Arpino, Sanjana Srivastava, Lyne P. Tchapmi, Micael E. Tchapmi, Kent Vainio, Li Fei-Fei, and Silvio Savarese. igibson, a simulation environment for interactive tasks in large realistic scenes. arXiv preprint arXiv:2012.02924, 2020. 1
- [84] Shuran Song, Andy Zeng, Johnny Lee, and Thomas A. Funkhouser. Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations. *IEEE Robotics Autom. Lett.*, 2020. 3
- [85] Mike Stilman and James J Kuffner. Navigation among movable obstacles: Real-time reasoning in complex environments. *International Journal of Humanoid Robotics*, 2(04):479–503, 2005. 3
- [86] Freek Stulp, Andreas Fedrizzi, Lorenz Mösenlechner, and Michael Beetz. Learning and reasoning with action-related places for robust mobile manipulation. J. Artif. Intell. Res., 2012. 3

- [87] Charles Sun, Jedrzej Orbik, Coline Devin, Brian Yang, Abhishek Gupta, Glen Berseth, and Sergey Levine. Relmm: Practical RL for learning mobile manipulation skills using only onboard sensors. arXiv preprint arXiv:2107.13545, 2021. [], 3
- [88] Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John Turner, Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Habitat 2.0: Training home assistants to rearrange their habitat. *arXiv preprint arXiv:2106.14405*, 2021. 1
- [89] Lei Tai, Giuseppe Paolo, and Ming Liu. Virtual-to-real deep reinforcement learning: Continuous control of mobile robots for mapless navigation. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017, 2017. 2] 3
- [90] Chen Tessler, Daniel J. Mankowitz, and Shie Mannor. Reward constrained policy optimization. In 7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019, 2019.
- [91] Ulrich Viereck, Andreas ten Pas, Kate Saenko, and Robert Platt Jr. Learning a visuomotor controller for real world robotic grasping using simulated depth images. In 1st Annual Conference on Robot Learning, CoRL 2017, Mountain View, California, USA, November 13-15, 2017, Proceedings, 2017. 3
- [92] Cong Wang, Qifeng Zhang, Qiyan Tian, Shuo Li, Xiaohui Wang, David Lane, Yvan R. Petillot, and Sen Wang. Learning mobile manipulation through deep reinforcement learning. Sensors, 2020. [1] 3
- [93] Luca Weihs, Jordi Salvador, Klemen Kotar, Unnat Jain, Kuo-Hao Zeng, Roozbeh Mottaghi, and Aniruddha Kembhavi. Allenact: A framework for embodied AI research. arXiv preprint arXiv:2008.12760, 2020. 6, 14
- [94] Erik Wijmans, Samyak Datta, Oleksandr Maksymets, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, and Dhruv Batra. Embodied question answering in photorealistic environments with point cloud perception. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, 2019. 4
- [95] Erik Wijmans, Irfan Essa, and Dhruv Batra. How to train pointgoal navigation agents on a (sample and compute) budget. arXiv preprint arXiv:2012.06117, 2020. 6
- [96] Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, and Dhruv Batra. DD-PPO: learning near-perfect pointgoal navigators from 2.5 billion frames. In 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020, 2020. [13] [4] [6] [14]
- [97] Jason Andrew Wolfe, Bhaskara Marthi, and Stuart J. Russell. Combined task and motion planning for mobile manipulation. In *Proceedings of the 20th International Confer-*

ence on Automated Planning and Scheduling, ICAPS 2010, Toronto, Ontario, Canada, May 12-16, 2010, 2010. 3

- [98] Jimmy Wu, Xingyuan Sun, Andy Zeng, Shuran Song, Johnny Lee, Szymon Rusinkiewicz, and Thomas A. Funkhouser. Spatial action maps for mobile manipulation. In *Robotics: Science and Systems XVI, Virtual Event / Cor*valis, Oregon, USA, July 12-16, 2020, 2020. [1] 3
- [99] Yuxin Wu and Kaiming He. Group normalization. In Proceedings of the European conference on computer vision (ECCV), pages 3–19, 2018. 14
- [100] Fei Xia, Amir Roshan Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, 2018. 1
- [101] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020, 2020. [1]
- [102] Yoshio Yamamoto and Xiaoping Yun. Coordinating locomotion and manipulation of a mobile manipulator. *IEEE Trans. Autom. Control.*, 1994. 3
- [103] Denis Yarats, Amy Zhang, Ilya Kostrikov, Brandon Amos, Joelle Pineau, and Rob Fergus. Improving sample efficiency in model-free reinforcement learning from images. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, 2021.*
- [104] Joel Ye, Dhruv Batra, Erik Wijmans, and Abhishek Das. Auxiliary tasks speed up learning pointgoal navigation. *arXiv preprint arXiv:2007.04561*, 2020. [1] [3] [4] [5] [6] [7]
- [105] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How transferable are features in deep neural networks? In Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada, 2014. 4
- [106] Andy Zeng, Shuran Song, Stefan Welker, Johnny Lee, Alberto Rodriguez, and Thomas A. Funkhouser. Learning synergies between pushing and grasping with selfsupervised deep reinforcement learning. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2018, Madrid, Spain, October 1-5, 2018, 2018. [1] 3
- [107] Andy Zeng, Shuran Song, Kuan-Ting Yu, Elliott Donlon, Francois Robert Hogan, Maria Bauzá, Daolin Ma, Orion Taylor, Melody Liu, Eudald Romo, Nima Fazeli, Ferran Alet, Nikhil Chavan Dafle, Rachel Holladay, Isabella Morona, Prem Qu Nair, Druck Green, Ian J. Taylor, Weber Liu, Thomas A. Funkhouser, and Alberto Rodriguez.

Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching. In 2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018, 2018.

[108] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J. Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In 2017 IEEE International Conference on Robotics and Automation, ICRA 2017, Singapore, Singapore, May 29 - June 3, 2017, 2017. [1] [3]