

ROBUSTNAV: Towards Benchmarking Robustness in Embodied Navigation

Prithvijit Chattopadhyay^{1,2*}Judy Hoffman¹Roозbeh Mottaghi^{2,3}Aniruddha Kembhavi^{2,3}¹Georgia Tech²PRIOR @ Allen Institute of AI³University of Washington

{prithvijit3, judy}@gatech.edu

{roozbehm, anik}@allenai.org

prior.allenai.org/projects/robustnav

Abstract

As an attempt towards assessing the robustness of embodied navigation agents, we propose **ROBUSTNAV**, a framework to quantify the performance of embodied navigation agents when exposed to a wide variety of visual – affecting RGB inputs – and dynamics – affecting transition dynamics – corruptions. Most recent efforts in visual navigation have typically focused on generalizing to novel target environments with similar appearance and dynamics characteristics. With **ROBUSTNAV**, we find that some standard embodied navigation agents significantly underperform (or fail) in the presence of visual or dynamics corruptions. We systematically analyze the kind of idiosyncrasies that emerge in the behavior of such agents when operating under corruptions. Finally, for visual corruptions in **ROBUSTNAV**, we show that while standard techniques to improve robustness such as data-augmentation and self-supervised adaptation offer some zero-shot resistance and improvements in navigation performance, there is still a long way to go in terms of recovering lost performance relative to clean “non-corrupt” settings, warranting more research in this direction. Our code is available at <https://github.com/allenai/robustnav>.

1. Introduction

A longstanding goal of the artificial intelligence community has been to develop algorithms for embodied agents that are capable of reasoning about rich perceptual information and thereby accomplishing tasks by navigating in and interacting with their environments. In addition to being able to exhibit these capabilities, it is equally important that such embodied agents are able to do so in a robust and generalizable manner.

A major challenge in Embodied AI is to ensure that agents can generalize to environments with different *appearance statistics* and *motion dynamics* than the environment used for training those agents. For instance, an agent

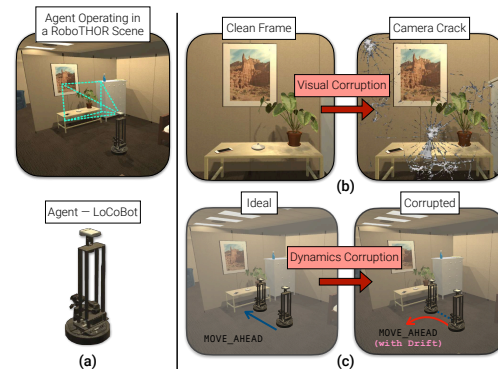


Figure 1. **ROBUSTNAV**. (a) A navigation agent pretrained in clean environments is asked to navigate to targets in unseen environments in the presence of (b) *visual* and (c) *dynamics* based corruptions. Visual corruptions (ex. camera crack) affect the agent’s egocentric RGB observations while Dynamics corruptions (ex. drift in translation) affect transition dynamics in the unseen environment. that is trained to navigate in “sunny” weather should continue to operate in rain despite the drastic changes in the appearance, and an agent that is trained to move on carpet should decidedly navigate when on a hardwood floor despite the discrepancy in friction. While a potential solution may be to calibrate the agent for a specific target environment, it is not a scalable one since there can be enormous varieties of unseen environments and situations. A more robust, efficient and scalable solution is to equip agents with the ability to autonomously adapt to new situations by interaction without having to train for every possible target scenario. Despite the remarkable progress in Embodied AI, especially in embodied navigation [59, 46, 48, 54, 7], most efforts focus on generalizing trained agents to unseen environments, but critically assume similar appearance and dynamics attributes across train and test environments.

As a first step towards assessing general purpose robustness of embodied agents, we propose **ROBUSTNAV**, a framework to quantify the performance of embodied navigation agents when exposed to a wide variety of common visual (*vis*) and dynamics (*dyn*) corruptions – artifacts that affect the egocentric RGB observations and transition dy-

*Part of the work done when PC was a research intern at AI2.

namics, respectively. We envision ROBUSTNAV as a testbed for adapting agent behavior across different perception and actuation properties. While assessing robustness to changes (stochastic or otherwise) in environments has been investigated in the robotics community [32, 13, 14, 21], the simulated nature of ROBUSTNAV enables practitioners to explore robustness against a rich and very diverse set of changes, while inheriting the advantages of working in simulation – speed, safety, low cost and reproducibility.

ROBUSTNAV consists of two widely studied embodied navigation tasks, Point-Goal Navigation (POINTNAV) [2] and Object-Goal Navigation (OBJECTNAV) [4] – the tasks of navigating to a goal-coordinate in a global reference frame or an instance of a specified object, respectively. Following the standard protocol, agents learn using a set of training scenes and are evaluated within a set of held out test scenes, but differently, ROBUSTNAV test scenes are subject to a variety of realistic *visual* and *dynamics* corruptions. These corruptions can emulate real world scenarios such as a malfunctioning camera or drift (see Fig. 1).

As zero shot adaptation to test time corruptions may be out of reach for our current algorithms, we provide agents with a fixed “calibration budget” (number of interactions) within the target world for unsupervised adaptation. This mimics a real-world analog where a shipped robot is allowed to adapt to changes in the environment by executing a reasonable number of unsupervised interactions. Post calibration, agents are evaluated on the two tasks in the corrupted test environments using standard navigation metrics.

Our extensive analysis reveals that both POINTNAV and OBJECTNAV agents experience significant degradation in performance across the range of corruptions, particularly when multiple corruptions are applied together. We show that this degradation reduces in the presence of a clean depth sensor suggesting the advantages of incorporating multiple sensing modalities, to improve robustness. We find that data augmentation and self-supervised adaptation strategies offer some zero-shot resistance and improvement over degraded performance, but are unable to fully recover this gap in performance. Interestingly, we also note that visual corruptions affect embodied tasks differently from static tasks like object recognition – suggesting that visual robustness should be explored within an embodied task. Finally, we analyze several interesting behaviors our agents exhibit in the presence of corruptions – such as increase in the number of collisions and inability to terminate episodes successfully.

In summary, our contributions include: (1) We present ROBUSTNAV – a framework for benchmarking and assessing the robustness of embodied navigation agents to visual and dynamics corruptions. (2) Our findings show that present day navigation agents trained in simulation underperform severely when evaluated in corrupt target environments. (3) We systematically analyze the kinds of mis-

takes embodied navigation agents make when operating under such corruptions. (4) We find that although standard data-augmentation techniques and self-supervised adaptation strategies offer some improvement, much remains to be done in terms of fully recovering lost performance.

ROBUSTNAV provides a fast framework to develop and test robust embodied policies, before they can be deployed onto real robots. While ROBUSTNAV currently supports navigation heavy tasks, the supported corruptions can be easily extended to more tasks, as they get popular within the Embodied AI community.

2. Related Work

Visual Navigation. Tasks involving navigation based on egocentric visual inputs have witnessed exciting progress in recent years [48, 10, 24, 8, 19, 9]. Some of the widely studied tasks in this space include POINTNAV [2], OBJECTNAV [4] and goal-driven navigation where the target is specified by a goal-image [59]. Approaches to solve POINTNAV and OBJECTNAV can broadly be classified into two categories – (1) learning neural policies end-to-end using RL [53, 57, 46, 48, 54] or (2) decomposing navigation into a mapping (building a semantic map) and path planning stage [6, 7, 25, 43]. Recent research has also focused on assessing the ability of policies trained in simulation to transfer to real-world robots operating in physical spaces [33, 12].

Robustness Benchmarks. Assessing robustness of deep neural models has received quite a bit of attention in recent years [30, 45, 31, 3]. Most relevant and closest to our work is [30], where authors show that computer vision models are susceptible to several synthetic visual corruptions, as measured in the proposed ImageNet-C benchmark. In [34, 39], authors study the effect of similar visual corruptions for semantic segmentation and object detection on standard static benchmarks. ROBUSTNAV integrates several visual corruptions from [30] and adds ones such as low-lighting and crack in the camera-lens, but within an embodied scenario. Our findings (see Sec. 5) show that visual corruptions affect embodied tasks differently from static tasks like object recognition. In [50], authors repurpose the ImageNet validation split to be used as a benchmark for assessing robustness to natural distribution shifts (unlike the ones introduced in [30]) and [17] identifies statistical biases in the same. Recently, [29] proposes three extensive benchmarks assessing robustness to image-style, geographical location and camera operation. Adversarial robustness of embodied agents has also been explored in [37], where the authors utilize spatio-temporal context to craft adversarial perturbations to alter 3D properties (shape and texture) of objects in Embodied QA [10] and Embodied Visual Recognition [56].

Real-world RL Suite. Efforts similar to ROBUSTNAV have been made in [16], where authors formalize 9 different challenges holding back RL from real-world use – including actuator delays, high-dimensional state and action spaces,

latency, and others. In contrast, ROBUSTNAV focuses on challenges in the visually rich domains and complexities associated with visual observation. Recently, Habitat [48] also introduced actuation (from [40]) and visual noise models for navigation tasks. In contrast, ROBUSTNAV is designed to benchmark robustness of models against a variety of visual and dynamics corruptions (7 vis and 4 dyn corruptions for both POINTNAV and OBJECTNAV).

Adapting Visio-Motor Policies. Significant progress has been made in the problem of adapting policies trained with RL from a source to a target environment. Unlike ROBUSTNAV, major assumptions involved in such transfer settings are either access to task-supervision in the target environment [23] or access to paired data from the source and target environments [22, 51]. Domain Randomization (DR) [1, 44, 36, 41] is another common approach to train policies robust to various environmental factors. Notably, [36] perturbs features early in the visual encoders of the policy network so as to mimic DR and [41] selects optimal DR parameters during training based on sparse data obtained from the the real world. In absence of task supervision, another common approach is to optimize self-supervised objectives in the target [54, 47] and has been used to adapt policies to visual disparities (see Sec. 5) in new environments [26]. To adapt to changes in transition dynamics, a common approach is to train on a broad family of dynamics models and perform system-identification (ex. with domain classifiers [18]) in the target environment [55, 58]. [33, 12] studies the extent to which embodied navigation agents transfer from simulated environments to real-world physical spaces. Among these, we investigate two of the most popular approaches – self-supervised adaptation [26] and aggressive data augmentation and measure if they can help build resistance to vis corruptions.

3. ROBUSTNAV

We present ROBUSTNAV, a benchmark to assess the robustness of embodied agents to common visual (vis) and dynamics (dyn) corruptions. ROBUSTNAV is built on top of ROBOTHOR [11]. In this work, we study the effects corruptions have on two kinds of embodied navigation agents – namely, POINTNAV (navigate to a specified goal coordinate) and OBJECTNAV (navigate to an instance of an object category). While we restrict our experiments to navigation, in practice, our vis and dyn corruptions can also be extended to other embodied tasks that share the same modalities, for instance tasks involving interacting with objects.

In ROBUSTNAV, agents are trained within the training scenes and evaluated on “corrupt” unseen target scenes. Corruptions in target scenes are drawn from a set of predefined vis and dyn corruptions. As is the case with any form of modeling of corruptions (or noise) in simulation [32, 11], there will always be an approximation error when the vis and dyn corruptions are compared to their real world coun-




Figure 2. **Visual Corruptions.** Visual corruptions ROBUSTNAV supports in the unseen target environments. Top-left shows a clean RGB frame and rest show corrupted versions of the same. Our aim is to ensure that the ROBUSTNAV benchmark acts as a stepping stone towards the larger goal of obtaining robust agents, ready to be deployed in real world.

To adapt to a corrupt target scene, we provide agents with a “calibration budget” – an upper bound on the number of interactions an agent is allowed to have with the target environment without any external task supervision. This is done to mimic a real-world analog where a shipped robot is allowed to adapt to changes in the environment by executing a reasonable number of unsupervised interactions. Our definition of the calibration-budget is based on the number of steps it takes an agent to reasonably recover degraded performance in the most severely corrupted environments when finetuned under complete supervision (see Table. 3) – set to $\sim 166k$ steps for all our experiments. We attempt to understand if self-supervised adaptation approaches [26] improve performance when allowed to adapt under this calibration budget (see Sec. 5, resisting corruptions). We now describe in detail the vis and dyn corruptions present in ROBUSTNAV.

Visual Corruptions. Visual corruptions are artifacts that degrade the navigation agent’s egocentric RGB observation (see Fig. 2). We provide seven visual corruptions within ROBUSTNAV, four of which are drawn from the set of corruptions and perturbations proposed in [30] – Spatter, Motion Blur, Defocus Blur and Speckle Noise; realistic corruptions that one might expect to see on a real robot. Spatter emulates occlusion in images due to particles of dirt, water droplets, etc. residing on the camera lens. Motion Blur emulates blur in images due to jittery movement of the robot. Defocus Blur occurs when the RGB image is out of focus. Speckle Noise emulates granular interference that inherently exists in and degrades the quality of images obtained by the camera (modeled as additive noise with the noise being proportional to the original pixel intensity). Each of these corruptions can manifest at five levels of severity indicating increase in the extent of visual degradation ($1 \rightarrow 5$).

In addition to these, we also add low-lighting (low-lighting conditions in the target environment, has associated severity levels $1 \rightarrow 5$), lower-FOV (agents operating with a lower camera field of view compared to the one used during training, $79^\circ \rightarrow 39.5^\circ$) and camera-crack (a randomized

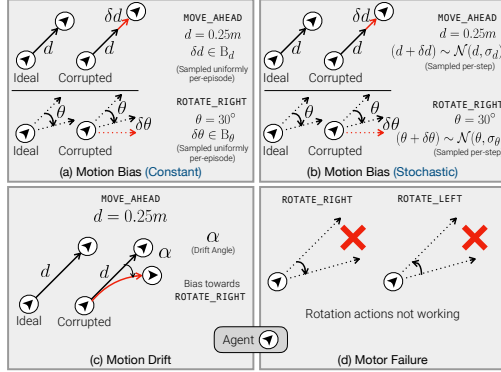


Figure 3. **Dynamics Corruptions.** We show the kinds of dynamics corruptions supported in ROBUSTNAV. Motion Bias (C & S) are modeled to mimic friction. Motion Drift models a setting where translation actions have a slight bias towards rotating right (or left). In Motor Failure, the one of the rotation actions fail.

crack in the camera lens). For camera-crack, we use fixed random seeds for the 15 validation scenes which dictate the location and kind of crack on the camera lens.

Dynamics Corruptions. Dynamics corruptions affect the transition dynamics of the agents in the target environment (see Fig. 3). We consider three classes of dynamics corruptions – Motion Bias, Motion Drift and Motor Failure. Our dyn corruptions are motivated from and in line with the well-known systematic and/or stochastic drifts (due to error accumulation) and biases in robot motion [35, 5, 20, 42].

A common dynamics corruption observed in the real world is friction. Unfortunately ROBOTHOR does not yet natively support multiple friction zones within a scene, as may be commonly observed in a real physical environment (for instance the kitchen floor in a house may have smooth tiles while the bedroom may have rough hardwood floors). In lieu of this, we present the Motion Bias corruption. In the absence of this corruption, the `move_ahead` action moves an agent forward by 0.25m, and rotation `rotate_left` and `rotate_right` actions rotate an agent by 30° left and right respectively. Motion Bias can induce either (a) a *constant* bias drawn uniformly per-episode from $B_d = \{\pm 0.05, \pm 0.1, \pm 0.15\}$ m or $B_\theta = \{\pm 5^\circ, \pm 10^\circ, \pm 15^\circ\}$ or (b) *stochastic* translation and rotation amounts drawn per-step from $\mathcal{N}(0.25\text{m}, 0.1\text{m})$ and $\mathcal{N}(30.0^\circ, 10^\circ)$ respectively.¹

Motion Drift models a setting where an agent’s translation movements in the environment include a slight bias towards turning left or right. Specifically, the `move_ahead` action, instead of moving an agent forward 0.25m in the direction of its heading (intended behavior), drifts towards the left or right directions stochastically (for an episode) by $\alpha = 10^\circ$ and takes it to a location which deviates in a

¹(a) Motion Bias (C) is intended to model scene-level friction, different floor material in the target environment; (b) Motion Bias (S) is intended to model high and low friction zones in a scene. Including more sophisticated models of friction is in the feature roadmap for ROBUSTNAV.

direction perpendicular to the original heading by a max of $\sim 0.043\text{m}$. Motor-failure is the setting where either the `rotate_left` or the `rotate_right` actions malfunction throughout an evaluation episode.

With the exception of Motion-Bias (S) – the *stochastic* version – the agent also operates under standard actuation noise models as calibrated from a LoCoBot in [12]. Recently, PyRobot [40] has also introduced LoCoBot calibrated noise models that demonstrate strafing and drifting. While we primarily rely on the noise models calibrated in [11], for completeness, we also include results with the PyRobot noise models.

Tasks. ROBUSTNAV consists of two major embodied navigation tasks – namely, POINTNAV and OBJECTNAV. In POINTNAV, an agent is initialized at a random spawn location and orientation in an environment and is asked to navigate to target coordinates specified relative to the agent’s position. The agent must navigate based **only** on sensory inputs from an RGB (or RGB-D) and a GPS + Compass sensor. An episode is declared successful if the agent stops within 0.2m of the goal location (by intentionally invoking an `end` action). In OBJECTNAV, an agent is instead asked to navigate to an instance of a specified object category (e.g., *Television*, 1 out of total 12 object categories) given **only** ego-centric sensory inputs – RGB or RGB-D. An episode is declared successful if the agent stops within 1.0m of the target object (by invoking an `end` action) and has the target object in it’s egocentric view. Due to the lack of perfect localization (no GPS + Compass sensor) and the implicit need to ground the specified object within its view, OBJECTNAV may be considered a harder task compared to POINTNAV – also evident in lower OBJECTNAV performance (Table. 2).

Metrics. We report performance in terms of the following well established navigation metrics reported in past works – Success Rate (SR) and Success Weighted by Path Length (SPL) [2]. SR indicates the fraction of successful episodes. SPL provides a score for the agent’s path based on how close it’s length is to the shortest path from the spawn location to the target. If $\mathbb{I}_{\text{success}}$ denotes whether an episode is successful (binary indicator), l is the shortest path length, and p is the agent’s path length then $\text{SPL} = \mathbb{I}_{\text{success}} \frac{l}{\max(l, p)}$

Scenes. ROBUSTNAV is built on top of the ROBOTHOR scenes [12]. ROBOTHOR consists of 60 training and 15 validation environments based on indoor apartment scenes drawn from different layouts. To assess robustness in the presence of corruptions, we evaluate on 1100 (and 1095) episodes of varying difficulties (easy, medium and hard)² for POINTNAV (and OBJECTNAV) across the 15 val scenes.

Benchmarking. Present day embodied navigation agents are typically trained without any corruptions. However, we

²Based on shortest path lengths – (1) POINTNAV: 0.00 – 2.28 for easy, 2.29 – 4.39 for medium, 4.40 – 9.61 for hard; (2) OBJECTNAV: 0.00 – 1.50 for easy, 1.51 – 3.78 for medium , 3.79 – 9.00 for hard.

Corruptions	Top-1 Acc. \uparrow	Top-5 Acc. \uparrow
1 Clean	69.76	89.08
2 Camera Crack [†]	57.71 \pm 5.82	80.27 \pm 4.54
3 Lower FOV*	45.44	69.53
4 Low Lighting	35.76	58.54
5 Spatter	19.73	39.34
6 Motion Blur	10.11	22.66
7 Defocus Blur	9.39	22.25
8 Speckle Noise	7.79	18.84

Table 1. **ImageNet Performance Degradation.** Degradation in classification performance on the ImageNet validation split under visual corruptions for ResNet-18 [28] trained on ImageNet (used as the agent’s visual encoder). Corruptions in 2-8 are present ROBUSTNAV. *Since mimicking lower FOV requires access to camera intrinsics, unavailable for static datasets, we mimic the same by aggressive center-cropping. [†]For camera-crack, we report performance over all possible crack settings present in ROBUSTNAV.

anticipate that researchers may incorporate corruptions as augmentations at training time to improve the robustness of their algorithms in order to make progress on our ROBUSTNAV framework. For the purposes of fair benchmarking, we recommend that future comparisons using ROBUSTNAV do not draw from the set of corruptions reserved for the target scenes – ensuring the corruptions encountered in the target scenes are indeed “unseen”.

4. Experimental Setup

Agent. Our POINTNAV agents have 4 actions available to them – namely, **move_ahead** (0.25m), **rotate_left** (30°), **rotate_right** (30°) and **end**. The action **end** indicates that the agent believes that it has reached the goal, thereby terminating the episode. During evaluation, we allow an agent to execute a maximum of 300 steps – if an agent does not call **end** within 300 steps, we forcefully terminate the episode. For OBJECTNAV, in addition to the aforementioned actions, the agent also has the ability to **look_up** or **look_down** – indicating change in the agent’s view 30° above or below the forward camera horizon. The agent receives 224×224 sized ego-centric observations (RGB or RGB-D). All agents are trained under LoCoBot calibrated actuation noise models from [12] – $\mathcal{N}(0.25\text{m}, 0.005\text{m})$ for translation and $\mathcal{N}(30^\circ, 0.5^\circ)$ for rotation. Our agent architectures (akin to [53]) are composed of a CNN head to process input observations followed by a recurrent (GRU) policy network (more details in Sec. 1.3 of supplementary).

Training. We train our agents using DD-PPO [53] – a decentralized, distributed and synchronous version of the Proximal Policy Optimization (PPO) [49] algorithm. If $R = 10.0$ denotes the terminal reward obtained at the end of a successful episode (with $\mathbb{I}_{\text{success}}$ being an indicator variable denoting whether an episode was successful), Δ_t^{Geo} denotes the change in geodesic distance to target at timestep t from $t - 1$ and $\lambda = -0.01$ denotes a slack penalty to encourage efficiency, then the reward received by the agent at time-step t can be expressed as,

$$r_t = \underbrace{R \cdot \mathbb{I}_{\text{success}}}_{\text{success reward}} - \underbrace{\Delta_t^{\text{Geo}}}_{\text{reward shaping}} + \underbrace{\lambda}_{\text{slack reward}}$$

We train our agents using the AllenAct [52] framework.

5. Results and Findings

In this section, we show that the performance of POINTNAV and OBJECTNAV agents degrades in the presence of corruptions (see Table. 2). We first highlight how vis corruptions affect static vision and embodied navigation tasks differently (see Table 1). Following this, we analyze behaviors that emerge in these agents when operating in the presence of vis, dyn, and vis+dyn corruptions. Finally, we investigate whether standard data-augmentation and self-supervised adaptation [26] techniques help recover the degraded performance (see Table 3).

5.1. Degradation in Performance

We now present our findings regarding degradation in performance relative to agents being evaluated in clean (no corruption) target environments (row 1 in Table. 2).

Visual corruptions affect static and embodied tasks differently. In Table 1, we report object recognition performance for models trained on the ImageNet [15] train split and evaluated on the corrupt validation splits. In Table 2, we report performance degradation of POINTNAV and OBJECTNAV agents under corruptions (row 1, clean & rows 2-8 corrupt). It is important to note that the nature of tasks (one-shot prediction vs sequential decision making) are different enough that the difficulty of corruptions for classification may not indicate the difficulty of corruptions for navigation. We verify this hypothesis by comparing results in Tables 1 and 2 – for instance, corruptions which are severe for classification (Defocus Blur and Speckle Noise) are not as severe for POINTNAV-RGB agents in terms of relative drop from clean performance. Additionally, for MaskRCNN [27] trained on AI2-THOR images, we note that detection (segmentation)³ mAP^{0.5:0.95} drops from 62.93 (66.29) to 7.96 (8.64) and 6.56 (6.68) for Spatter (S5) and Low-Lighting (S5), respectively – unlike rows 2 & 8 in Table 2, where Spatter appears to be much severe compared to Low-Lighting. This difference in relative degradation suggests that that techniques for visual adaptation or robustness in static settings may not transfer out-of-the-box to embodied tasks, warranting more research in this direction.

Not all corruptions are equally bad. While we note that POINTNAV and OBJECTNAV agents suffer a drop in performance from clean settings, not all corruptions are equally severe. For instance, in POINTNAV-RGB, while Low Lighting, Motion Blur and Motion Bias (C) (rows 2, 3, 9 in Table 2) lead to a worst-case absolute drop of $< 10\%$ in SPL (and $< 10\%$ in SR), corruptions like Spatter and Motor Failure (rows 8, 13) are more extreme and significantly

³For the 12 OBJECTNAV target classes

#	Corruption ↓	V	D	POINTNAV				OBJECTNAV			
				RGB		RGB-D		RGB		RGB-D	
				SR ↑	SPL ↑	SR ↑	SPL ↑	SR ↑	SPL ↑	SR ↑	SPL ↑
1	Clean			98.82	83.13	98.54	84.60	31.05	14.26	35.62	17.20
2	Low Lighting	✓		94.36	75.15	99.45	84.97	10.78	4.59	21.64	9.98
3	Motion Blur	✓		95.72	73.37	99.36	85.36	10.59	4.03	20.27	8.29
4	Camera Crack	✓		82.07	63.83	95.72	81.21	7.21	3.57	24.29	12.50
5	Defocus Blur	✓		75.89	53.55	99.09	85.54	5.02	2.42	19.18	7.90
6	Speckle Noise	✓		67.42	48.57	98.73	84.66	9.04	3.66	18.63	7.52
7	Lower-FOV	✓		42.49	31.73	89.08	73.59	9.77	3.90	9.86	4.77
8	Spatter	✓		33.58	24.72	98.91	84.81	6.76	2.93	21.10	9.06
9	Motion Bias (C)		✓	92.81	77.83	93.36	79.46	31.51	14.09	31.96	15.38
10	Motion Bias (S)		✓	94.72	76.95	96.72	79.08	30.87	14.15	35.62	16.39
11	Motion Drift		✓	95.72	76.19	93.36	75.08	29.68	13.58	34.06	17.03
12	PyRobot [40] (ILQR) Mul. = 1.0		✓	96.00	67.79	95.45	69.27	32.51	11.26	36.35	13.62
13	Motor Failure		✓	20.56	17.63	20.56	17.62	4.20	2.43	6.39	3.67
14	Defocus Blur + Motion Bias (S)	✓	✓	76.52	51.08	97.18	79.46	5.57	2.00	18.54	7.23
15	Speckle Noise + Motion Bias (S)	✓	✓	62.69	43.31	95.81	78.27	7.85	3.73	18.54	8.16
16	Spatter + Motion Bias (S)	✓	✓	33.30	23.33	95.81	78.85	7.85	3.09	21.28	9.26
17	Defocus Blur + Motion Drift	✓	✓	74.25	50.99	95.54	76.66	4.57	1.93	17.35	6.97
18	Speckle Noise + Motion Drift	✓	✓	64.42	44.73	94.36	75.23	8.49	3.67	19.82	8.61
19	Spatter + Motion Drift	✓	✓	32.94	23.44	95.45	76.61	6.85	2.68	19.54	8.86

Table 2. **POINTNAV and OBJECTNAV Performance.** Degradation in task performance of pretrained POINTNAV (trained for ~ 75 M frames) and OBJECTNAV (trained for ~ 300 M frames) agents when evaluated under vis and dyn corruptions present in ROBUSTNAV. POINTNAV agents have additional access to a GPS-Compass sensor. For visual corruptions with controllable severity levels, we report results with severity set to 5 (worst). Performance is measured across tasks of varying difficulties (easy, medium and hard). Rows are sorted based on SPL values for RGB POINTNAV agents. Success and SPL values are reported as percentages. (V = Visual, D = Dynamics)

affect task performance (absolute drops of $> 57\%$ in SPL, $> 65\%$ in SR). For OBJECTNAV, however, the drop in performance is more gradual across corruptions (partly because it’s a harder task and even clean performance is fairly low).

A “clean” depth sensor helps resisting degradation. We compare the RGB and RGB-D variants of the trained POINTNAV and OBJECTNAV agents (RGB corrupt, Depth clean) in Table 2 (corresponding RGB & RGB-D columns). We observe that including a “clean” depth sensor consistently improves resistance to vis, dyn and vis+dyn corruptions for both POINTNAV and OBJECTNAV. For POINTNAV, we note that while RGB and RGB-D variants have comparable clean performance (row 1), under severe corruptions (Spatter, Lower-FOV and Speckle-Noise), the RGB-D counterparts are ahead roughly by an absolute margin of 36.09 – 60.09% SPL. We further observe that, barring exceptions, POINTNAV RGB-D agents are generally affected minimally by corruptions – for instance, Low-Lighting and Motion Blur barely result in any drop in performance. We hypothesize that this is likely because RGB-D navigation agents are much less reliant on the RGB sensor compared to the RGB counterparts. In OBJECTNAV, an additional depth sensor generally improves clean performance (row 1 in Table 2) which is likely the major contributing factor for increased resistance to corruptions. Sensors of different modalities are likely to degrade in different scenarios – e.g., a depth sensor may continue to perceive details in low lighting settings. The obtained results suggest that adding multiple sensors, while expensive can help train robust models. Additional sensors can also be helpful for unsupervised adaptation during the calibration phase. For instance, in the presence of a “clean” depth sensor, one can consider comparing depth

based egomotion estimates with expected odometry readings in the target environment to infer changes in dynamics.

Presence of vis+dyn corruptions further degrades performance. Rows 14–19 in Table 2 indicate the extent of performance degradation when vis+dyn corruptions are present. With the exception of a few cases, as expected, the drop in performance is slightly more pronounced compared to the presence of just vis or dyn corruptions. The relative drop in performance from vis \rightarrow vis+dyn is more pronounced for OBJECTNAV as opposed to POINTNAV.

Navigation performance for RGB agents degrades consistently with escalating episode difficulty. Recall that we evaluate navigation performance over episodes of varying difficulty levels (see Sec. 3). We break down the performance of POINTNAV & OBJECTNAV agents by episode difficulty levels (in Sec. 1.5 of supplementary). Under “clean” settings, we find that POINTNAV (RGB and RGB-D) have comparable performance across all difficulty levels. Under corruptions, we note that unlike the RGB-D counterparts, performance of POINTNAV-RGB agents consistently deteriorates as the episodes become harder. OBJECTNAV (both RGB & RGB-D) agents show a similar trend of decrease in navigation performance with increasing episode difficulty.

5.2. Behavior of Visual Navigation Agents

We now study the idiosyncrasies (see Fig 4) exhibited by these agents (POINTNAV-RGB and OBJECTNAV-RGB) which leads to their degraded performance.

Agents tend to collide more often. Fig 4 (first column, bars color-coded based on the kind of corruption) shows the average number of failed actions under corrupt settings. In our framework, failed actions occur as a consequence of




Figure 4. **Agent Behavior Analysis.** To understand agent behaviors, we report the breakdown of four metrics: Number of collisions as observed through Failed Actions (*first column*), closest the agent was to target as measured by Min. Dist. to Target (*second column*), and failure to appropriately end an episode either when out of range – Stop-Fail (Pos) (*third column*), or in range – Stop-Fail (Neg) (*fourth column*). Each behavior is reported for both POINTNAV (*top row*) and OBJECTNAV (*bottom row*) RGB agents within a clean and five corrupt settings: Defocus Blur (D.B.), Speckle Noise (S.N.), Motion Drift (M.D.), Defocus Blur + Motion Drift, and Speckle Noise + Motion Drift. is clean, is vis corruptions, is dyn corruptions and is vis+dyn corruptions. Blue line in col 2 indicates the distance threshold for goal in range. Severities for S.N. and D.B. are set to 5 (worst).

colliding with objects, walls, etc. While corruptions generally lead to increased collisions, we note that adding a dyn corruption in addition to a vis one (D.B. \rightarrow D.B. + M.D. & S.N. \rightarrow S.N. + M.D.) increases the number of collisions over vis or dyn corruptions – dyn corruptions lead to unforeseen changes in dynamics (actions working unexpectedly), which likely contributes to an uptick in collisions.

Agents tend to be farther from the target. Fig 4 (*second column*) shows the minimum distance from the target over the course of an episode. While we note that as corruptions become progressively severe, agents tend to terminate farther away from the target (see Sec. 1.4 of supplementary), Fig 4 (*second column*) indicates that the overall proximity of the agent to the goal over an episode decreases – minimum distance to target increases as we go from Clean \rightarrow vis or dyn; vis or dyn \rightarrow vis+dyn. While this may be intuitive in the presence of a dyn corruption, it is interesting to note that this trend is also consistent for vis corruptions (Clean \rightarrow D.B. or S.N.).

Corruptions hurt OBJECTNAV stopping mechanism. Recall that for both POINTNAV and OBJECTNAV, success depends on the notion of “intentionality” [4] – the agent calls an **end** action when it believes it has reached the goal. In Fig 4 (*last two columns*) we aim to understand how corruptions affect this stopping mechanism. Specifically, we look at two quantitative measures – (1) Stop-Failure (Positive), the proportion of times the agent invokes an **end** action when the goal is not range; and (2) Stop-Failure (Negative), the proportion of times the agent does not invoke an **end** action when the goal is in range, out of the number of times the goal is in range.⁴

We observe that prematurely calling an **end** action is a

⁴The goal in range criterion for POINTNAV checks if the target is within the threshold distance. For OBJECTNAV, this includes an additional visibility criterion.

significant issue only for OBJECTNAV (Fig 4 (*third column*)) – which becomes more pronounced as corruptions become progressively severe (Clean \rightarrow D.B. or S.N.; M.D. \rightarrow D.B. + M.D. or S.N. + M.D.). Similarly, the inability of an agent to invoke an **end** action is also more pronounced for OBJECTNAV as opposed to POINTNAV (Fig 4 (*fourth column*)). To investigate the extent to which this impacts the agent’s performance, we compare the agent’s Success Rate (SR) with a setting where the agent is equipped with an oracle stopping mechanism (call **end** as soon as the goal is in range). We find that this makes a significant difference only for OBJECTNAV– absolute +7.12% for Clean, +7.76% for M.D. and +13.88% for D.B. + M.D. We hypothesize that equipping agents with robust stopping mechanisms can significantly improve performance on ROBUSTNAV. For instance, equipping the agent with a progress monitor module [38] (estimating progress made towards the goal in terms of distance) robust to vis corruptions can potentially help decide when explicitly to invoke an **end** action in the target environment.

5.3. Resisting Corruptions

To assist near-term progress, we investigate if some standard approaches towards training robust models or adapting to visual disparities can help resisting vis corruptions under a calibration budget (Sec. 3) – set to $\sim 166k$ steps.⁵

Extent of attainable improvement by finetuning under task supervision. As an anecdotal upper bound on attainable improvements under the calibration-budget, we also report the extent to which degraded performance can be recovered when fine-tuned under complete task supervision. We report these results for vis corruptions in Table 3 (row 7). We note that unlike Lower-FOV, the agent is able to al-

⁵Based on the number of steps it takes an agent to reasonably recover degraded performance in corrupted environments when finetuned with complete task supervision.

Approach	Visual Corruption									
	Clean		Lower-FOV		Defocus Blur		Camera Crack		Spatter	
	SR↑	SPL↑	SR↑	SPL↑	SR↑	SPL↑	SR↑	SPL↑	SR↑	SPL↑
1 Nav. Loss	98.82	83.13	42.49	31.73	75.89	53.55	82.07	63.83	33.58	24.72
2 Nav. Loss + AP	98.45	83.28	45.68	35.14	83.35	61.51	72.70	56.82	20.38	15.70
3 Nav. Loss + AP + SS-Adapt	37.31	31.03	32.94	26.09	40.95	33.35	57.87	46.72	14.19	10.29
4 Nav. Loss + RP	98.73	82.53	44.95	32.74	32.21	22.47	67.06	53.70	23.48	18.63
5 Nav. Loss + RP + SS-Adapt	94.63	77.25	50.59	36.10	79.16	62.74	60.42	49.37	61.06	47.16
6 Nav. Loss + Data Aug	98.45	81.08	71.70	54.54	81.26	61.32	88.44	71.57	23.93	18.41
7 Finetune Nav. Loss on Target	-	-	72.88	61.82	97.18	80.32	96.54	80.92	91.81	77.38

Table 3. **Resisting Visual Corruptions.** To assist near-term progress, we study if standard approaches towards training visually robust models or adapting to visual disparities can help resisting visual corruptions. All agents in rows 1-7 are POINTNAV RGB agents pre-trained for ~ 75 M frames. Agents in rows 3 & 5 have obtained by running adaptation for ~ 166 k steps. Agents in row 7 provide an anecdotal upper bound indicating attainable improvements when finetuned with task-supervision under the calibration budget – set to ~ 166 k steps. For visual corruptions with controllable severity levels, we report results with severity set to 5 (worst).

most recover performance for Defocus Blur, Camera-Crack and Spatter (Table. 3, rows 1,7).

Do data-augmentation strategies help? In Table 3, we study if data-augmentation strategies improve zero-shot resistance to vis corruptions (rows 1,6). We compare POINTNAV RGB agents trained with Random-Crop, Random-Shift and Color-Jitter (row 6) with the vanilla versions (row 1) and find that while data augmentation (row 6) offers some improvements (Spatter being an exception) over degraded performance (row 1) – absolute improvements of (22.81% SPL, 29.21% SR) for Lower-FOV, (7.77% SPL, 5.37% SR) for Defocus Blur and (7.74% SPL, 6.37% SR) for Camera-Crack, obtained performance is still significantly below Clean settings (row 1, Clean col). Improvements are more pronounced for Lower-FOV compared to others (likely due to Random-Shift and Random-Crop). We note that data-augmentation provides improvements only for a subset of vis corruptions and when it does, obtained improvements are still not sufficient enough to recover lost performance.

Do self-supervised adaptation approaches help? In the absence of reward supervision in the target environment, Hansen *et al.* [26] proposed Policy Adaptation during Deployment (PAD) – source pretraining with an auxiliary supervised objective and optimizing only the self-supervised objective when deployed in the target environment. We investigate the degree to which PAD helps adapting to the target environments in ROBUSTNAV. The adopted self-supervised tasks are (1) Action-Prediction (AP) – given two successive observations in a trajectory, predict the intermediate action and (2) Rotation-Prediction (RP) – rotate the input observation by 0° , 90° , 180° , or 270° before feeding it to the agent and task an additional auxiliary head with predicting the rotation. We report numbers with AP (rows 2,3) and RP (rows 4,5) in Table. 3. For AP, we find that (1) pre-training (row 2 vs row 1) results in little or no improvements over degraded performance (maximum absolute improvements of 7.96% SPL, 7.46% SR for Defocus Blur) and (2) further adaptation (row 3 vs rows 2,1) under calibration budget consistently degrades performance. For RP, we observe that (1) with the exception of Clean and Lower-FOV, pre-training (row 4 vs row 1) results in worse performance

and (2) while self-supervised adaptation under corruptions improves performance over pre-training (row 5 vs row 4), it is still significantly below Clean settings (row 1, Clean col) – minimum absolute gap of 20.39% SPL, 19.66% SR between Defocus Blur (row 5) and Clean (row 1). While improvements over degraded performance might highlight the utility of PAD (with AP / RP) as a potential unsupervised adaptation approach, there is still a long way to go in terms of closing the performance gap between clean and corrupt settings.

6. Conclusion

In summary, as a step towards assessing general purpose robustness of embodied navigation agents, we propose ROBUSTNAV, a challenging framework well-suited to benchmark the robustness of embodied navigation agents, with a wide variety of *visual* and *dynamics* corruptions. To succeed on ROBUSTNAV, an agent must be insensitive to corruptions and also be able to adapt to unforeseen changes in new environments with minimal interaction. We find that standard POINTNAV and OBJECTNAV agents underperform (or fail) significantly in the presence of corruptions and while standard techniques to improve robustness or adapt to environments with visual disparities (data-augmentation, self-supervised adaptation) provide some improvements, a large room for improvement remains in terms of fully recovering lost navigation performance. Lastly, we plan on evolving ROBUSTNAV in terms of the sophistication and diversity of corruptions as more features are supported in the underlying simulator. We release ROBUSTNAV in ROBOTHOR, and hope that our findings provide insights into developing more robust navigation agents.

Acknowledgements. We thank Klemen Klotar, Luca Weihs, Martin Lohmann, Harsh Agrawal and Rama Vedantam for fruitful discussions and valuable feedback. We thank Winson Han for helping out with the Camera-Crack vis corruption. We thank Vishvak Murahari for helping out with the ImageNet and the Mask-RCNN experiments. We also thank the ICCV’21 reviewers for their comments. This work is supported by the NASA University Leadership Initiative (ULI) under grant number 80NSSC20M0161.

References

- [1] Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, et al. Solving rubik’s cube with a robot hand. *arXiv*, 2019. 3
- [2] Peter Anderson, Angel Chang, Devendra Singh Chaplot, Alexey Dosovitskiy, Saurabh Gupta, Vladlen Koltun, Jana Kosecka, Jitendra Malik, Roozbeh Mottaghi, Manolis Savva, et al. On evaluation of embodied navigation agents. *arXiv*, 2018. 2, 4
- [3] Maksym Andriushchenko and Nicolas Flammarion. Understanding and improving fast adversarial training. In *NeurIPS*, 2020. 2
- [4] Dhruv Batra, Aaron Gokaslan, Aniruddha Kembhavi, Oleksandr Maksymets, Roozbeh Mottaghi, Manolis Savva, Alexander Toshev, and Erik Wijmans. Objectnav revisited: On evaluation of embodied agents navigating to objects. *arXiv*, 2020. 2, 7
- [5] Kostas E. Bekris, Andrew Ladd, and Lydia E. Kavraki. Efficient motion planners for systems with dynamics. In *ICRA Workshop*, 2007. 4
- [6] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. In *ICLR*, 2019. 2
- [7] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *NeurIPS*, 2020. 1, 2
- [8] Changan Chen, Unnat Jain, Carl Schissler, Sebastia Vincenc Amengual Gari, Ziad Al-Halah, Vamsi Krishna Ithapu, Philip Robinson, and Kristen Grauman. Soundspaces: Audio-visual navigation in 3d environments. In *ECCV*, 2020. 2
- [9] Changan Chen, Sagnik Majumder, Ziad Al-Halah, Ruohan Gao, Santhosh Kumar Ramakrishnan, and Kristen Grauman. Learning to set waypoints for audio-visual navigation. In *ICLR*, 2021. 2
- [10] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied Question Answering. In *CVPR*, 2018. 2
- [11] Matt Deitke, Winson Han, Alvaro Herrasti, Aniruddha Kembhavi, Eric Kolve, Roozbeh Mottaghi, Jordi Salvador, Dustin Schwenk, Eli VanderBilt, Matthew Wallingford, et al. RoboThor: An open simulation-to-real embodied ai platform. In *CVPR*, 2020. 3, 4
- [12] 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 *CVPR*, 2020. 2, 3, 4, 5
- [13] Andrea Del Prete and Nicolas Mansard. Addressing constraint robustness to torque errors in task-space inverse dynamics. In *RSS*, 2015. 2
- [14] Andrea Del Prete and Nicolas Mansard. Robustness to joint-torque-tracking errors in task-space inverse dynamics. *IEEE Trans. on Robotics*, 2016. 2
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [16] Gabriel Dulac-Arnold, Nir Levine, Daniel J. Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. An empirical investigation of the challenges of real-world reinforcement learning. *arXiv*, 2020. 2
- [17] Logan Engstrom, Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Jacob Steinhardt, and Aleksander Madry. Identifying statistical bias in dataset replication. In *ICML*, 2020. 2
- [18] Benjamin Eysenbach, Swapnil Asawa, Shreyas Chaudhari, Ruslan Salakhutdinov, and Sergey Levine. Off-dynamics reinforcement learning: Training for transfer with domain classifiers. *arXiv*, 2020. 3
- [19] Chuang Gan, Yiwei Zhang, Jiajun Wu, Boqing Gong, and Joshua B Tenenbaum. Look, listen, and act: Towards audio-visual embodied navigation. In *ICRA*, 2020. 2
- [20] Carlos Garcia-Saura. Self-calibration of a differential wheeled robot using only a gyroscope and a distance sensor. *arXiv*, 2015. 4
- [21] Nirmal Giftsun, Andrea Del Prete, and Florent Lamiraux. Robustness to inertial parameter errors for legged robots balancing on level ground. In *International Conference on Informatics in Control, Automation and Robotics*, 2017. 2
- [22] Florian Golemo, Adrien Ali Taiga, Aaron Courville, and Pierre-Yves Oudeyer. Sim-to-real transfer with neural-augmented robot simulation. In *CoRL*, 2018. 3
- [23] Daniel Gordon, Abhishek Kadian, Devi Parikh, Judy Hoffman, and Dhruv Batra. Splitnet: Sim2sim and task2task transfer for embodied visual navigation. In *ICCV*, 2019. 3
- [24] Daniel Gordon, Aniruddha Kembhavi, Mohammad Rastegari, Joseph Redmon, Dieter Fox, and Ali Farhadi. Iqa: Visual question answering in interactive environments. In *CVPR*, 2018. 2
- [25] Saurabh Gupta, James Davidson, Sergey Levine, Rahul Sukthankar, and Jitendra Malik. Cognitive mapping and planning for visual navigation. In *CVPR*, 2017. 2
- [26] Nicklas Hansen, Yu Sun, Pieter Abbeel, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Self-supervised policy adaptation during deployment. In *ICLR*, 2021. 3, 5, 8
- [27] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 5
- [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5
- [29] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. *arXiv*, 2020. 2
- [30] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *ICLR*, 2019. 2, 3
- [31] Dan Hendrycks, Xiaoyuan Liu, Eric Wallace, Adam Dziedzic, Rishabh Krishnan, and Dawn Song. Pretrained transformers improve out-of-distribution robustness. In *ACL*, 2020. 2

- [32] Sebastian Höfer, Kostas Bekris, Ankur Handa, Juan Camilo Gamboa, Florian Golemo, Melissa Mozifian, Chris Atkeson, Dieter Fox, Ken Goldberg, John Leonard, et al. Perspectives on sim2real transfer for robotics: A summary of the r:ss 2020 workshop. *arXiv*, 2020. 2, 3
- [33] Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, and Dhruv Batra. Are we making real progress in simulated environments? measuring the sim2real gap in embodied visual navigation. In *IROS*, 2020. 2, 3
- [34] Christoph Kamann and Carsten Rother. Benchmarking the robustness of semantic segmentation models. In *CVPR*, 2020. 2
- [35] Andrew M. Ladd and Lydia E. Kavraki. Motion planning in the presence of drift, underactuation and discrete system changes. In *RSS*, 2005. 4
- [36] Kimin Lee, Kibok Lee, Jinwoo Shin, and Honglak Lee. Network randomization: A simple technique for generalization in deep reinforcement learning. In *ICLR*, 2020. 3
- [37] Aishan Liu, Tairan Huang, Xianglong Liu, Yitao Xu, Yuqing Ma, Xinyun Chen, Stephen J Maybank, and Dacheng Tao. Spatiotemporal attacks for embodied agents. In *ECCV*, 2020. 2
- [38] Chih-Yao Ma, Jiasen Lu, Zuxuan Wu, Ghassan AlRegib, Zsolt Kira, Richard Socher, and Caiming Xiong. Self-monitoring navigation agent via auxiliary progress estimation. *arXiv preprint arXiv:1901.03035*, 2019. 7
- [39] Claudio Michaelis, Benjamin Mitzkus, Robert Geirhos, Evgenia Rusak, Oliver Bringmann, Alexander S Ecker, Matthias Bethge, and Wieland Brendel. Benchmarking robustness in object detection: Autonomous driving when winter is coming. *arXiv*, 2019. 2
- [40] Adithyavairavan Murali, Tao Chen, Kalyan Vasudev Alwala, Dhiraj Gandhi, Lerrel Pinto, Saurabh Gupta, and Abhinav Gupta. Pyrobot: An open-source robotics framework for research and benchmarking. *arXiv*, 2019. 3, 4, 6
- [41] Fabio Muratore, Christian Eilers, Michael Gienger, and Jan Peters. Data-efficient domain randomization with bayesian optimization. *IEEE Robotics Automation and Letters*, 2021. 3
- [42] J. Müller, N. Kohler, and W. Burgard. Autonomous miniature blimp navigation with online motion planning and replanning. In *IROS*, 2011. 4
- [43] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In *ECCV*, 2020. 2
- [44] Sharath Chandra Raparthy, Bhairav Mehta, Florian Golemo, and Liam Paull. Generating automatic curricula via self-supervised active domain randomization. *arXiv*, 2020. 3
- [45] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishal Shankar. Do imagenet classifiers generalize to imagenet? In *ICML*, 2019. 2
- [46] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. In *ICLR*, 2018. 1, 2
- [47] Nikolay Savinov, Anton Raichuk, Raphaël Marinier, Damien Vincent, Marc Pollefeys, Timothy Lillicrap, and Sylvain Gelly. Episodic curiosity through reachability. In *ICLR*, 2019. 3
- [48] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In *ICCV*, 2019. 1, 2, 3
- [49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv*, 2017. 5
- [50] Rohan Taori, Achal Dave, Vaishal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *NeurIPS*, 2020. 2
- [51] Joanne Truong, Sonia Chernova, and Dhruv Batra. Bi-directional domain adaptation for sim2real transfer of embodied navigation agents. *IEEE Robotics and Automation Letters*, 2021. 3
- [52] Luca Weihs, Jordi Salvador, Klemen Kotar, Unnat Jain, Kuo-Hao Zeng, Roozbeh Mottaghi, and Aniruddha Kembhavi. Allenact: A framework for embodied ai research. *arXiv*, 2020. 5
- [53] Erik Wijmans, Abhishek Kadian, Ari S. 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 *ICLR*, 2020. 2, 5
- [54] Mitchell Wortsman, Kiana Ehsani, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Learning to learn how to learn: Self-adaptive visual navigation using meta-learning. In *CVPR*, 2019. 1, 2, 3
- [55] Jiachen Yang, Brenden Petersen, Hongyuan Zha, and Daniel Faissol. Single episode policy transfer in reinforcement learning. In *ICLR*, 2020. 3
- [56] Jianwei Yang, Zhile Ren, Mingze Xu, Xinlei Chen, David J. Crandall, Devi Parikh, and Dhruv Batra. Embodied amodal recognition: Learning to move to perceive objects. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2040–2050, 2019. 2
- [57] Joel Ye, Dhruv Batra, Erik Wijmans, and Abhishek Das. Auxiliary tasks speed up learning pointgoal navigation. In *CoRL*, 2020. 2
- [58] Wenxuan Zhou, Lerrel Pinto, and Abhinav Gupta. Environment probing interaction policies. In *ICLR*, 2019. 3
- [59] 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 *ICRA*, 2017. 1, 2