

Learning on the Move: Integrating Action and Perception for Mobile Manipulation

Shagun Uppal

CMU-RI-TR-24-26

May 21



The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:

Professor Deepak Pathak, *chair*
Professor Katerina Fragkiadaki
Ananye Agarwal

*Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Robotics.*

Copyright © 2024 Shagun Uppal. All rights reserved.

To my parents, Sarita and Naresh Uppal.

Abstract

While there has been remarkable progress recently in the fields of manipulation and locomotion, mobile manipulation remains a long-standing challenge. Compared to locomotion or static manipulation, a mobile system must make a diverse range of long-horizon tasks feasible in unstructured and dynamic environments. While the applications are broad and interesting, there are a plethora of challenges in developing these systems such as coordination between the base and arm, reliance on onboard perception for perceiving and interacting with the environment, and most importantly, simultaneously integrating all these parts together. Prior works approach the problem using disentangled modular skills for mobility and manipulation that are trivially tied together. This causes several limitations such as compounding errors, delays in decision-making, and no whole-body coordination. In this work, we present a reactive mobile manipulation framework that uses an active visual system to consciously perceive and react to its environment. Similar to how humans leverage whole-body and hand-eye coordination, we develop a mobile manipulator that exploits its ability to move and see, more specifically – to move in order to see and to see in order to move. This allows it to not only move around and interact with its environment but also, choose “when” to perceive “what” using an active visual system. We observe that such an agent learns to navigate around complex cluttered scenarios while displaying agile whole-body coordination using only ego-vision without needing to create environment maps.

Acknowledgments

I consider myself extremely lucky to have had some of the most amazing mentors, friends, and family throughout different stages of life, who shaped me for the better with their constant support and guidance. With the last two years at CMU coming to an end, I feel immense gratitude and serendipity to have met with so many wonderful opportunities and some of the brightest minds that have helped me exponentially grow my potential.

Firstly, I would like to extend my deepest and most heartfelt gratitude to my advisor, Deepak Pathak for his unwavering support, wisdom, and encouragement all throughout. His profound knowledge, exceptional guidance, and attention to detail have helped me strive throughout my Masters. His remarkable judgment of being highly ambitious and picking the right problems to work on has forever altered my way of thinking about research and life for the better. I am personally grateful for the most abstract, ambitious, yet concise, practical, and super-long brainstorming sessions with him at the very beginning of me joining the lab. They comprise of some of my fondest memories which always left me with an urge to keep learning more, thinking fundamentally and clearly in order to bring ideas to realization.

I extend my sincere appreciation to Professor Katerina Fragkiadaki, for being a part of my thesis committee and for her insightful feedback and encouragement.

I am super thankful to Ananye Agarwal, for constantly being there to help, guide, think, and even re-think ideas at any and every single point throughout my thesis projects. I am especially thankful to him for showing me some really cool ways for debugging my code. It was always super fun to brainstorm ideas and discuss updates with him every now and then. Not to forget, the deadline days working together are some of my favorite memories despite the high pressure, because the end was always pleasantly satisfying.

I am super grateful to all my lab members who made every single day, especially late nights working in the lab so much fun and intriguing. I learned a lot from each one of them and shared some of the most fun memories as well as deep research discussions. I thank (in alphabetical order) Aditya Kanan, Alexandre Kirchmeyer, Alexandre Li, Ananye

Agarwal, Ellis Brown, Haoyu Xiong, Jayesh Singla, Jiahui Yang, Kenny Shaw, Kevin Gmelin, Kexin Shi, Lili Chen, Mihir Prabhudesai, Murtaza Dalal, Russell Mendonca, Shikhar Bahl, Shivam Duggal, Unnat Jain, Xuxin Cheng, and Yulong Li. I am grateful for the opportunity to work beside them all and grow in their presence.

I thank all my dear friends, Anish Madan, Bharath Raj, Dvij Kalaria, Poorvi Hebbar, Pranay Gupta, Pushkal Katara, Sarthak Bhagat, Shreya Sharma, and Sriram Narayanan for helping me grow personally as well as professionally in multi-dimensional aspects. A special thanks to them for relaxing coffee breaks and our newfound love for table tennis which was a constant stress-burster and ended up being a post-dinner ritual in Robolounge over the years. Special thanks to Shreya for infinite discussions, from research to life, several failed attempts at being regular to the gym, and our unique love for coffee boba. Few of the best memories comprise of some super successful and some devastatingly futile attempts at stargazing.

I am extremely grateful to Sarthak for always being there by my side, encouraging me to be the best version of myself, while constantly pushing the bar higher. Thank you for making tough days easier, easier days happier, and the happier ones most memorable. It has been a real pleasure to have grown with you alongside getting two degrees.

Last but not least, I would like to thank my parents who have given me infinite support, strength, and encouragement to always follow my heart and create my own path, even when they did not align with it completely. I am and will be forever thankful for the education and opportunities you provided me with, and for teaching me the power of knowledge. While I can never pay back to you for your innumerable sacrifices since my childhood, I hope that I am able to pay it forward. I would also like to thank my brother Nimish for his unconditional support, allowing me to spread my wings higher, without worrying about anything. I am forever grateful for the impact you have had, in shaping me at each and every stage of life both directly and indirectly, for instilling in me the magic of curiosity, and for leaving me a little more brave every time to fight for what I believed in, whenever needed. Finally, I would like to thank my grandfather, Kuldeep Raj Arora for always inspiring me to reach for the stars, while still keeping me grounded.

Contents

1	Introduction	1
2	Related Work	5
2.1	Classical Approaches	5
2.2	Learning-based navigation	6
2.3	Active Perception	6
2.4	Mobile Manipulation	6
3	Method	9
3.1	Phase 1 - Learning Simultaneous Perception, Interaction and Navigation	10
3.1.1	Pick Policy	10
3.1.2	Navigation Policy	11
3.2	Phase 2 - From Scandots to Depth	13
4	Experimental Setup	17
4.1	Baselines	17
4.2	Design Choices and Implementation Details	18
5	Results and Analysis	23
5.1	Emergent Behavior	24
5.2	Real-world results	25
5.3	Simulation results	26
5.4	Directly training from depth images.	27
5.5	Analysing camera and base motion	28
5.6	Necessity for Active Vision and Whole-body coordination	28
5.7	Classical Navigation Baseline	29
6	Discussion	33
	Bibliography	35

When this dissertation is viewed as a PDF, the page header is a link to this Table of Contents.

List of Figures

1.1	Human and robot illustration of whole-body navigation through the clutter.	2
1.2	Learning to SPIN: Our robot learns to S imultaneously P erceive, I nteract, and N avigate around cluttered unstructured environments in a whole-body fashion. The robot has an actuated camera with a limited field of view that it must control to get information about its environment. The motion and perception problem are tightly coupled since what the robot knows about the environment influences how it can move and vice versa. We show results in a large variety of scenarios both indoors and outdoors with different obstacles like boxes and furniture. Our robot can pick up different objects like cups, and utensils. Video demos at https://spin-robot.github.io	3
3.1	SPIN: We learn a policy that uses ego-vision to simultaneously perceive, interact, and navigate in cluttered environments. We propose two methods: (1) Coupled Visuomotor Optimization (CVO) learns robot and camera actions at the same time. We train an RL policy to predict these. We only provide scandots if they are visible in the agent’s field-of-view allowing the agent to learn to move its camera and aggregate information about its environment. This is followed by a phase-2 supervised training where this behavior is distilled into a student network that operates with ego-centric depth images (2) Decoupled Visuomotor Optimization (DVO) decouples the action and perception learning into two parts: first the agent learns to navigate across clutter assuming access to all obstacles. In phase 1b, the robot learns to move its camera to estimate the relevant information. This is followed by supervised learning same as above.	14
3.2	Visualizing the simulation benchmark: We illustrate one scenario of the simulation benchmark here with many obstacles in a narrow passage. The agent learns to develop whole-body coordination such as the robot’s arm movement in the last two frames, to reactively adapt and navigate through such cluttered scenes by actively moving around the camera and aggregating information for efficient navigation without collisions.	15

4.1	Visualizing Scandots: (Left) We compute visible scandots by projecting them to the camera frame and checking if they lie within the image plane, and Anatomy of the Mobile Manipulator: (Right) of the stretch RE1 robot that we use experiments with. It has two DoFs in the base, one each for arm lift and extension, two for the camera, three for the wrist, and one for the gripper.	18
4.2	Camera movement analysis in a trajectory. The agent faces the camera downward when navigating through a tightly cluttered vicinity as can be seen in the first, second, and fourth frames, whereas the camera points more toward the front when there are no immediate obstacles in the direction of movement, as illustrated in the third frame.	22
4.3	Essential Whole-Body Coordination under heavy obstructions. In the above cases where the obstacles are tightly packed, it is not possible for the robot to navigate through them avoiding collisions without lifting the arm to an appropriate height.	22
5.1	Types of emergent behavior exhibited by <i>SPIN</i> (a) dynamic obstacle avoidance (b) whole-body movement (c) adaptive rerouting.	30
5.2	Success rate for scandots (blue) vs depth (red). The depth-based policy attains close to 0 performance even after 400k env steps of training, whereas the policy trained with scandots increasingly improves over time.	31
5.3	Camera movement analysis in a trajectory. The agent faces the camera downward when navigating through a tightly cluttered vicinity as can be seen in the first, second, and fourth frame, whereas the camera points more toward the front when there are no immediate obstacles in the direction of movement, as illustrated in the third frame.	31
5.4	Scenarios where whole-body coordination is essential under heavy obstructions. In the above cases where the obstacles are tightly packed, it is not possible for the robot to navigate through them avoiding collisions without lifting the arm to an appropriate height.	32
5.5	Visualizations of environment map built using 2D Lidar. The robot is localized as per its initial position and orientation.	32

List of Tables

4.1	We report the average success rate and average distance to goal for 10 episodes across 3 seeds each with a different maximum visibility range for the agent at any time instant. As reported, with broader visibility, the agent shows more frequent stalling leading to a higher average distance to the goal.	19
4.2	We report the average success rate and average distance to goal for 10 episodes across 3 seeds each with a different maximum visibility range for the agent at any time instant. As reported, with broader visibility, the agent shows more frequent stalling leading to a higher average distance to the goal.	20
5.1	We evaluate the success rate on 10 random environments with an average of 3 fixed seeds across all difficulty scenarios based on obstacle course. We report the success rate of each part of the task including reaching (Reach), picking (Pick), and placing (Place) the target object in the desired location. The place task requires the agent to bring back the object across the obstacles near its start location.	24
5.2	We compare our method against a classical mapping and planning baseline for navigation in cluttered scenes with both static as well as dynamic obstacles. The classical performs reasonably in static environments, it quickly breaks with dynamic obstacles like humans walking around, whereas our method shows more robust reactivity to such obstacles even without being trained with dynamic obstacles in simulation. We report the success rate of our method compared with the baseline. For the classical baseline, we teleoperate the robot for 2-3 min.	27
5.3	Success rate for 4 FixCam poses in easy (E), medium (M), hard (H) environments.	28

Chapter 1

Introduction

Consider the example shown in Figure 1.1. A person is trying to carry a coffee cup through clutter. This not only requires navigational planning from start to goal but planning of the whole body to avoid obstacles along the way. Furthermore, due to ego-centric vision, the person needs to actively look around to gather the presence of obstacles. This general form of mobile manipulation task necessitates a coupled understanding of whole-body control with active perception. This capability is one of the fundamental and frequently encountered tasks in embodied cognition.

The dominant paradigm to tackle this problem is through classical planning-based control which requires apriori knowledge about the precise location of all the obstacles along with a detailed map of the environment. In most real-world scenarios, this assumption is impractical due to computational reasons, but more importantly, because environments are dynamic and objects keep moving around in general. Furthermore, relying on precise measurement of scenes for control does not allow agents to reactively improvise to changes in their environment. Practically, even when the complete environment map is known apriori, joint planning for a system with high degrees of freedom, say a mobile base with an arm, is often intractable and too expensive to be deployed in real-time.

Humans, on the other hand, do not rely on precise known estimates of object locations and instead use ego-centric vision to navigate around obstacles in real time. In an unfamiliar environment, where to look is informed by where they want to move (called ‘active perception’), and how they move in return determines what they can

1. Introduction



Figure 1.1: Human and robot illustration of whole-body navigation through the clutter.

see immediately afterward. This integrated mobility and perception allows us to see, adapt, and react to maneuver through unseen heavily cluttered environments [5].

This work presents *SPIN*, an end-to-end approach to *S*imultaneous *P*erception, *I*nteraction, and *N*avigation. We train a single model that not only outputs low-level controls for the robot body and arm but also predicts where should the robot’s ego-centric camera look at each time step while moving its whole body by avoiding obstacles.

We train our approach via reinforcement learning (RL), and to get around the computational bottleneck of rendering depth images, we use a teacher-student training framework where robot behavior is first learned using RL with access to visible object scandots and then distilled into a policy that operates from ego-depth using supervised learning. We evaluate across 6 benchmarks in simulation ranging from easy, medium, and hard difficulty, and two real-world environments with a similar level of clutter as the hard environments in simulation and also add dynamic, adversarial obstacles. We find that our method outperforms classical methods and baselines which do not use active vision. We also observed emergent behaviors, including dynamic obstacle avoidance which the robot did not see during training time.

Our approach presents a radical hypothesis that the traditionally non-reactive

planning approach to whole-body control can indeed be cast into a reactive model – i.e. – single end-to-end policy trained by RL. Despite a big departure from optimal control literature, this hypothesis is not as surprising since agile whole-body coordination and fast obstacle avoidance in humans are developed into muscle memory over time.



Figure 1.2: **Learning to SPIN:** Our robot learns to *Simultaneously Perceive, Interact,* and *Navigate* around cluttered unstructured environments in a whole-body fashion. The robot has an actuated camera with a limited field of view that it must control to get information about its environment. The motion and perception problem are tightly coupled since what the robot knows about the environment influences how it can move and vice versa. We show results in a large variety of scenarios both indoors and outdoors with different obstacles like boxes and furniture. Our robot can pick up different objects like cups, and utensils. Video demos at <https://spin-robot.github.io>

1. Introduction

Chapter 2

Related Work

2.1 Classical Approaches

The problem of navigating robots around obstacles has been studied for decades. Classical methods solve the motion and perception problem separately. First, these methods build a map of the environment using the robot’s onboard sensors such as cameras, proprioception, and Lidar or infrared [9, 20]. Kalman-filter-based [48] techniques are often used to track positions, but they can’t represent multi-modal ambiguities or recover after tracking failure [41]. Grid-based methods solve this but suffer from high memory usage [7]. Modern SLAM approaches ORBSLAM3 [38], OpenVSLAM [44] and RTAB [30] use variations of a method that relies on particle filters [46] to hold a multi-modal belief of the robot’s location in the map [34, 45]. SLAM is especially challenging in dynamic environments due to the confounding motion of other agents [17, 42, 51, 54]. Once a map is built, a path can be planned over it. Exact paths can be computed using graph search algorithms [21], probabilistic methods which are faster but yield approximately optimal solutions [31] or potential-field based methods [26]. All of these assume perfect perception and re-planning is usually expensive making them susceptible to noise and precluding reactive behavior.

2.2 Learning-based navigation

In recent years, learning has been used to improve the classical navigation stack [47]. Modular approaches [10, 11, 19, 35] still leverage SLAM-based methods to build a map but use learning or heuristic changes to get priors for the best possible route to a goal. End-to-end approaches forgo maps entirely and train a policy to go from images to robot commands to go to a goal location [12, 13, 50]. We also take the end-to-end approach but unlike prior work where what the robot sees is fixed based on its position, in our case, it must move its head and actively choose what it sees making optimization more challenging.

2.3 Active Perception

An embodied agent that can move around, act as well as observe its environment possesses abundant information to continuously generate new data for itself which is missing while learning from passive data. A continually adapting agent acts on its curiosity to move around and actively change its surroundings to identify its own failure cases. To model this curiosity, several prior works such as [14, 15, 39] develop active vision frameworks, where an agent is driven by inconsistencies in its own visual estimation of its surroundings when perceived from novel viewpoints. In this work, we couple active perception with learning low-level actions for embodied mobile agents in order for them to reactively adapt, navigate, and interact with their environment.

2.4 Mobile Manipulation

While autonomous stationary manipulation has seen leaps in functional grasping [3, 32], as well as in several obstructed settings [33], Mobile Manipulation learning still poses several challenges. A mobile base and arm together can complete useful in-the-wild manipulation but present a more challenging control problem. Imitation learning techniques focus on collecting large datasets in a variety of settings with a dexterous 6-DOF arm and a wheeled mobile base using teleoperation [4, 6, 8, 16, 22, 52]. Because of the high dimensionality of mobile manipulation systems, there are also

control methods that leverage synergies between both the base and the arm and plans together. [18, 22, 23, 53].

2. Related Work

Chapter 3

Method

Our goal is for the mobile manipulator to navigate and manipulate objects while avoiding obstacles in cluttered environments. In this work, we use the Stretch Hello Robot [1] as shown in Figure 4.1 as our embodied mobile manipulation agent. It shares anatomical similarities with a human, bringing with it many of the same challenges. First, it has a limb in the form of an arm that can be raised and lowered, so the robot must constantly move the arm to avoid any obstacles. Second, it has an actuated camera with a very limited field of view (87° horizontal, 58° vertical), so it needs to constantly look around to simultaneously plan ahead and look out for unexpected obstacles. Imagine yourself walking through a cluttered cabinet, there are too many obstacles around to keep track of, and you can't see all of them at once, so you must keep looking all around your body to plan a path through the clutter but also make sure you don't hit anything you missed along the way. Unlike regular walking where our eyes mostly point straight ahead and the path is clear, here you must *actively* choose what to perceive for simultaneously planning ahead and also doing reactive fixes to your planned path. Since all the obstacles cannot be perceived at a single glance, you must have spatial awareness and know where the obstacle you saw some time ago is right now in relation to your body. Note that this entire process is very different from the classical approach, where perception, planning, and obstacle avoidance are separate processes executed separately and in sequence. Further, it is assumed that the output of each is perfect, whereas this is rarely the case in practice in our unstructured world.

To deal with this challenging, entangled problem setup, we take a data-driven approach. We train our robot to navigate inside procedurally generated clutter in simulation using RL. The robot is only allowed to perceive the part of its environment that is visible to the camera and learns to coordinate its arm, base, and camera motion such that it can plan ahead and reactively adjust to obstacles.

In practice, since training with RL requires many samples and rendering depth is inefficient (Section 5.4), we divide training into two phases. In the first phase, we learn mobile manipulation behaviors via RL using scandots, which is a cheap-to-compute variant of depth and in phase 2 we train a CNN for perception from depth images as illustrated in Figure 3.1. An example illustration of scandots in the environment is depicted in Figure 4.1.

3.1 Phase 1 - Learning Simultaneous Perception, Interaction and Navigation

In this stage, we use RL to learn to control all the joints of the robot to navigate clutter and pick target objects. Since rendering depth images directly from the robot camera is expensive, we must instead use an ersatz version that contains the same information and is cheap to compute. We do so using *scandots* \mathbf{s}_t which are the XYZ coordinates of the bounding box of each obstacle. To specify which object to pick, we give the initial location of the object (before it is touched by the robot) \mathbf{o}_i . In lieu of the object image, we give the current location of the object \mathbf{o}_t . Here, scandots \mathbf{s}_t and object location \mathbf{o}_t are privileged information which must later be estimated from depth images. Given this information, we train two separate LSTM policies π_{nav} and π_{pick} . At test time, the nav policy is activated to reach a target location and we switch to the pick one once the robot gets close to the object.

3.1.1 Pick Policy

This accesses proprioception \mathbf{x}_t consisting of robot joint angles and velocities $\mathbf{q}_t, \dot{\mathbf{q}}_t$, base linear and angular velocity v_t, ω_t . For perception, it gets the object’s initial and

current location \mathbf{o}_i , \mathbf{o}_t and predicts robot and camera actions.

$$[\mathbf{a}_{\text{robot}}, \mathbf{a}_{\text{cam}}] = \pi_{\text{pick}}(\mathbf{x}_t, F(\mathbf{o}_t, \mathbf{x}_t), \mathbf{o}_i) \quad (3.1)$$

where F is a masking function that masks object position \mathbf{o}_t if it is not in the field of view of the camera. This is required since object position can only be estimated from depth in phase 2 if it is visible.

3.1.2 Navigation Policy

Training this policy requires a complex joint visuomotor optimization since robot motion is dependent on its knowledge of the environment which in turn depends on how the robot moves. We present two approaches to tackle this problem.

Coupled Visuomotor Optimization (CVO) Here, we set up a partially observable environment for the robot and let the RL algorithm do the joint optimization using large-scale data. In particular, the policy gets proprioception \mathbf{x}_t and only visible scandots $\tilde{\mathbf{s}}_t = F(\mathbf{s}_t, \mathbf{x}_t)$ as observation and has to predict both the camera and the robot actions. Since the scandots are permutation invariant, we pass them through a trainable point-net architecture P to obtain compressed latent $\mathbf{z}_t = P(\tilde{\mathbf{s}}_t)$ that we pass to the policy

$$[\mathbf{a}_{\text{robot}}, \mathbf{a}_{\text{cam}}] = \pi_{\text{nav}}(\mathbf{x}_t, \mathbf{z}_t) \quad (3.2)$$

This presents a tough optimization landscape because the observations at each step are strongly dependent on \mathbf{a}_{cam} . For instance, if the camera swivels around the observations at the next timestep may look completely different. Indeed, we observe that this requires billions of samples inside a GPU-accelerated simulator to optimize which may not always be feasible in practice.

Decoupled Visuomotor Optimization (DVO) To ease the optimization process, we learn the robot and camera actions separately. First, we learn how to move by giving the robot access to all available scandots $\mathbf{z}_t = P(\mathbf{s}_t)$ in a local vicinity. Since the robot sees everything, the camera motion is irrelevant and we just predict the robot’s motion

$$\mathbf{a}_{\text{robot}} = \pi_{\text{nav}}^{\text{1a}}(\mathbf{x}_t, \mathbf{z}_t, \mathbf{g}_t) \quad (3.3)$$

3. Method

where \mathbf{g}_t is the goal with respect to the base. Using this policy as supervision, we train another policy to predict both camera and robot motions with access to only visible scandots $\hat{\mathbf{z}}_t = P(F(\mathbf{s}_t, \mathbf{x}_t))$. This policy is trained via RL to predict the robot actions from phase 1 policy $\mathbf{a}_{\text{robot}}$. This optimization forces the student policy to learn camera behaviors that capture information about the environment that are needed to move in the optimal fashion. We initialize $\pi_{\text{nav}}^{\text{1b}}$ from the weights of $\pi_{\text{nav}}^{\text{1a}}$

$$\begin{aligned} \min_{\pi_{\text{nav}}^{\text{1b}}} \quad & \|\hat{\mathbf{a}}_{\text{robot}} - \mathbf{a}_{\text{robot}}\| \\ \text{s.t.} \quad & [\hat{\mathbf{a}}_{\text{robot}}, \hat{\mathbf{a}}_{\text{cam}}] = \pi_{\text{nav}}^{\text{1b}}(\mathbf{x}_t, \hat{\mathbf{z}}_t, \mathbf{g}_t) \end{aligned} \quad (3.4)$$

This decoupled approach learns to move and see in separate phases which eases the optimization burden. In principle, the coupled optimization is better since it is possible that the 1a policy may learn to exploit privileged information in a way that the 1b policy cannot estimate it for any set of camera movements. However, in our setting, this did not turn out to be the case.

We train our agent using PPO [43] with backpropagation through time [49] in procedurally generated environments in simulation.

Rewards: For the navigation task, we use distance to goal reward $\|\mathbf{g}_t\|$ along with a forward progress reward $|(\mathbf{v}_t)_g|$ where $(\mathbf{v}_t)_g$ is velocity along the direction of the goal.

$$r_{\text{nav}} = 0.1 \cdot \|\mathbf{g}_t\| + 0.1 \cdot |(\mathbf{v}_t)_g| \quad (3.5)$$

For the pick task, we provide an object reaching reward, i.e., the distance between the gripper and the object. This is followed by a lift reward if a successful grasp is detected (based on whether contact forces cross a threshold), similar to [33].

$$r_{\text{pick}} = 0.5 \cdot \|\mathbf{o}_t - \mathbf{p}_t\| + 0.5 \cdot r_{\text{lift}} \quad (3.6)$$

where

$$r_{\text{lift}} = (1 - \tanh(15 \cdot [(\mathbf{o}_t)_z]_+)) \mathbb{I} \left[\sum_i f_i > 10 \right] \quad (3.7)$$

where $[x]_+ = \max(x, 0)$ and \mathbb{I} is the indicator function which forces the reward to be active only when object contact forces f_i exceed 10N.

Training environments: We procedurally generate long corridors with obstacles placed in between the robot and the goal. The initial joints and orientation of the robot are randomized. Near the edges of the corridors, we place randomized obstacles and walls to simulate distractors in the depth image. For the pick task, objects are spawned on tables of varying dimensions. We used five different objects - a banana, mug, can, foambrick, and a bottle. The episode is terminated if the robot reaches the goal or hits an obstacle/table.

3.2 Phase 2 - From Scandots to Depth

Scandots [25, 29] are not directly observable in the real world and must instead be estimated from the depth image [2]. We train a convolution network C to convert rendered depth images \mathbf{d}_t to perception latents $\tilde{\mathbf{z}}_t$. This latent is passed to a student policy π' to predict the actions $[\tilde{\mathbf{a}}_{\text{robot}}, \tilde{\mathbf{a}}_{\text{cam}}]$. This is supervised using L2 loss from the phase 1 actions. The weights for π' are initialized using π . We train this policy using DAgger [40]. For the navigation policy, we optimize

$$\min_{C_{\text{nav}}, \pi'_{\text{nav}}} \left\| \pi'_{\text{nav}}(C_{\text{nav}}(\mathbf{d}_t), \mathbf{x}_t, \mathbf{g}_t) - \pi_{\text{nav}}(\mathbf{z}_t, \mathbf{x}_t, \mathbf{g}_t) \right\| \quad (3.8)$$

Note that the teacher policy π_{nav} can be trained using either the coupled or decoupled approach. Similarly, for the pick policy, we estimate current object position \mathbf{o}_t from depth

$$\min_{C_{\text{pick}}, \pi'_{\text{pick}}} \left\| \pi'_{\text{pick}}(C_{\text{pick}}(\mathbf{d}_t), \mathbf{x}_t, \mathbf{o}_i) - \pi_{\text{pick}}(\mathbf{z}_t, \mathbf{x}_t, \mathbf{o}_t, \mathbf{o}_i) \right\| \quad (3.9)$$

3. Method

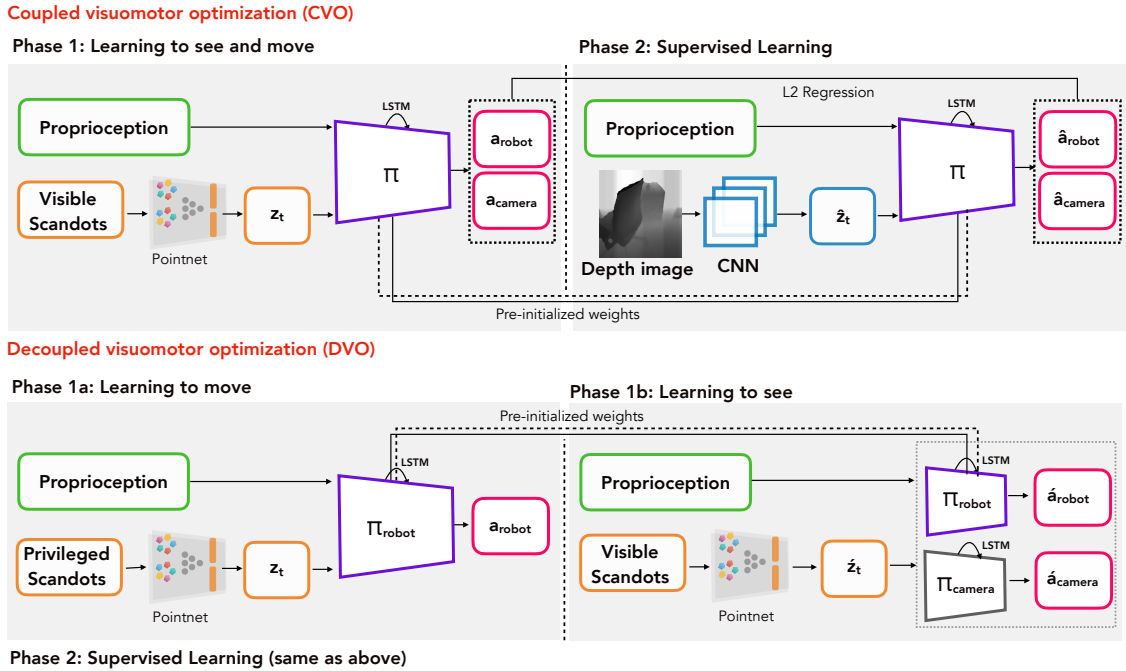


Figure 3.1: **SPIN**: We learn a policy that uses ego-vision to simultaneously perceive, interact, and navigate in cluttered environments. We propose two methods: (1) **Coupled Visuomotor Optimization (CVO)** learns robot and camera actions at the same time. We train an RL policy to predict these. We only provide scandots if they are visible in the agent’s field-of-view allowing the agent to learn to move its camera and aggregate information about its environment. This is followed by a phase-2 supervised training where this behavior is distilled into a student network that operates with ego-centric depth images (2) **Decoupled Visuomotor Optimization (DVO)** decouples the action and perception learning into two parts: first the agent learns to navigate across clutter assuming access to all obstacles. In phase 1b, the robot learns to move its camera to estimate the relevant information. This is followed by supervised learning same as above.




Figure 3.2: **Visualizing the simulation benchmark:** We illustrate one scenario of the simulation benchmark here with many obstacles in a narrow passage. The agent learns to develop whole-body coordination such as the robot’s arm movement in the last two frames, to reactively adapt and navigate through such cluttered scenes by actively moving around the camera and aggregating information for efficient navigation without collisions.

3. Method

Chapter 4

Experimental Setup

We use the Hello Robot Stretch [1] for all our experiments (Figure 4.1). The robot has 10 actuated joints which include 2 degrees of freedom for the camera, 2 for base rotation and translation, 2 for the arm, 1 for the gripper fingers, and 3 for the dexterous wrist. An Intel D435i depth camera is mounted on the top of the robot head which is actuated using two motors. The learned policy operates at 10Hz and we do velocity control for the robot base and position control for all the other joints. Velocity control for the robot base allows us to perform simultaneous robot translation and rotation for more agile behavior. We train using IsaacGymEnvs [36] using 8192 environments which takes 6 hours of training for phase 1 and 10 hours of training time for phase 2 on a RTX 3090.

4.1 Baselines

We compare our proposed approach against the following baselines:

- **FixCam:** The camera joints are frozen and the camera is forced to look forward. This baseline shows whether active vision is useful for the mobile manipulation problem and a fixed viewpoint is not enough.
- **Mapping:** Instead of using a moving depth camera to get a series of frames this baseline assumes exteroception is provided in the form of a map. We simulate exteroceptive noise as in [2, 37].

4. Experimental Setup




Figure 4.1: **Visualizing Scandots:** (Left) We compute visible scandots by projecting them to the camera frame and checking if they lie within the image plane, and **Anatomy of the Mobile Manipulator:** (Right) of the stretch RE1 robot that we use experiments with. It has two DoFs in the base, one each for arm lift and extension, two for the camera, three for the wrist, and one for the gripper.

- **Classical:** This uses a classical stack to control the base motion. We first teleoperate the robot for 3-5 minutes to construct a map using the onboard 2D RPLidar using gmapping. Next, `move_base` is used to plan a path through the environment. Finally, we move the robot to the start, use a Monte Carlo method [27] to localize, and then execute the plan. Note that this baseline gets an easier version of the problem since it assumes that the map is known in advance and does not consider arm motion due to the 2D Lidar. This is used to test whether reactive navigation is superior to planning.
- **NoPointNet:** Instead of passing object scandots through a permutation-invariant PointNet architecture, we concatenate them and use an MLP to estimate a latent.

4.2 Design Choices and Implementation Details

We make several design choices for the working of our framework. Firstly, the robot is only allowed to have local visibility in order to develop highly reactive and instant

behaviors. At any time instant t , the agent can perceive its environment within a range of 2m in all 4 directions – front, back, left, and right based on the camera’s viewing direction and its field of view. We also empirically observe that given a larger viewing range, say $> 5\text{m}$ which contains information of more than 4-5 nearest obstacles to the agent makes it a sub-global path planning problem that becomes harder to optimize, leading to degraded behavior and performance as reported in Table 4.1.

Visibility Range	Success Rate \uparrow	Distance to Goal (m) \downarrow
$\leq 1\text{m}$	0.96	0.28
$\leq 2\text{m}$	0.96	0.26
$\leq 3\text{m}$	0.93	0.63
$\leq 5\text{m}$	0.86	1.21

Table 4.1: We report the average success rate and average distance to goal for 10 episodes across 3 seeds each with a different maximum visibility range for the agent at any time instant. As reported, with broader visibility, the agent shows more frequent stalling leading to a higher average distance to the goal.

Secondly, contrary to standard teacher-student architectures which use privileged system information for training the teacher, we restrict the privileged information to only elements in the robot’s field-of-view that can be retrieved from the ego-view in the same state. For this, we project the obstacle scandots onto the image plane and pass only those scandots to the robot observation which lie in the camera’s field of view. Note that, if this were not the case, the robot would not learn camera movement in a relevant fashion and also become harder to distill into a depth-conditioned student policy through only an ego-centric view. Similarly, for the 3-phase decoupled visuomotor optimization, we induce information bottleneck with a low-dimensional latent space of size 16 for the scandots latent while training the teacher policy, which again helps it in attending to most relevant information at any time instant t in order to make student policy distillation feasible.

Observation Space. The observation space for the robot comprises of joint positions (q), joint velocities (q_{vel}), end-effector position p_{eef} , goal position (p_{goal}) and depth latent (\hat{z}) containing visual information about the environment. Note that during phase 1 training in simulation, scandots (z) are used as a proxy for faster depth

4. Experimental Setup

rendering and later distilled into an egocentric depth-conditioned policy. During real-world deployment, p_{eff} is obtained via forward kinematics and other proprioception information is obtained directly from the robot.

Action Space. The action space of the robot consists of the velocity for base rotation as well as translation and joint positions for all the other joints including arm, camera as well as gripper actions. The gripper action is a continuously varying scalar that can actuate the gripper to different extents, unlike a binary action indicating an open or closed gripper.

Reward Scales. We use a distance and forward progress goal for reaching, a binary reward for grasping, and a continuously shaped reward for lifting the object to a certain height above the table. We also add a small penalty for the joint velocities to the arm stretch and camera joints for a temporally smooth gait which permits easier sim2real transfer as well as more appropriate behaviors leading to less jitter and more consistent movements on the real hardware.

The reward scales used for goal reaching, grasping and lift rewards are reported in Table 4.2. Detailed formulations of the reward functions are described in Section 3.1.

	Reward Scale
Reach Reward	0.1
Grasp Reward	0.5
Lift Reward	0.8
Joint Velocity Penalty	-0.03

Table 4.2: We report the average success rate and average distance to goal for 10 episodes across 3 seeds each with a different maximum visibility range for the agent at any time instant. As reported, with broader visibility, the agent shows more frequent stalling leading to a higher average distance to the goal.

Network Architecture and Training Details. The actor and critic for teacher policy are LSTM with 256 hidden units, with input as proprioception, goal, and scandots latent. The scandots are compressed using a pointnet architecture for permutation invariance. The depth network for the student policy takes as input a low-resolution depth image of size 58×87 and comprises of 3-layer convolution

backbone followed by 3 fully-connected layers. We use Adam Optimizer with an initial learning rate of $1e - 3$, entropy coefficient of $5e - 4$, and γ as 0.99.

Asynchronous DAgger Training. Since depth rendering on simulators is a computational bottleneck, we implement an asynchronous version of the DAgger algorithm which simultaneously collects data in a buffer and trains the student policy with batches sampled from the collected data using 2 parallel processes. This provides a $2.5\times$ computational speedup over the non-parallelized version of the algorithm, allowing faster convergence of the student network. We also find that freezing the weights of the student actor pre-initialized from the teacher policy for the first 1000 iterations helps as warm-up steps to the depth convolution backbone for stable training.

Post-processing for clean depth images. To mitigate the issues due to noisy depth, we post-process the depth obtained from the Intel RealSense Camera using a real-time fast hole-filling algorithm for depth images [28]. With the camera constantly in motion, there are additional artifacts with depth images. For this, we additionally use temporal filtering over the stream of depth images.

Object Detection for Pick Policy Once the robot reaches the goal, we randomly select an object within its field of view in order to be grasped and fetched to a target location. For getting the target object location, we run YOLO [24], a real-time object detection model with an average inference speed of 20ms. We use the corresponding depth image to deproject the pixel point into a 3D-coordinate which is passed as the new goal position to the manipulation policy.

4. Experimental Setup




Figure 4.2: **Camera movement analysis in a trajectory.** The agent faces the camera downward when navigating through a tightly cluttered vicinity as can be seen in the first, second, and fourth frames, whereas the camera points more toward the front when there are no immediate obstacles in the direction of movement, as illustrated in the third frame.




Figure 4.3: **Essential Whole-Body Coordination under heavy obstructions.** In the above cases where the obstacles are tightly packed, it is not possible for the robot to navigate through them avoiding collisions without lifting the arm to an appropriate height.

Chapter 5

Results and Analysis

We evaluate our approach both in simulation as well as real-world. Since doing a lot of in-the-wild real-world experiments is more time-consuming and cumbersome due to various practical reasons, we thoroughly evaluate our approach on 6 simulation benchmarks with multiple scenarios. We explain each of these benchmarks in detail in Section 5.3.

While simulation benchmarks are useful for fair comparison with baselines as well as reproducibility, real-world experimenting is essential for determining the efficacy of our system in truly unstructured and dynamic environments. For this, we test our system on various real-world environments as shown in Figure 1.2 and benchmark its performance on 2 real-world setups as described in Section 5.2.

Through simulation experiments, we aim to answer the following questions: (1) For a mobile agent, is active perception with an actuated camera really necessary, or is a fixed viewpoint enough? (2) Can an active visual agent outperform a classical agent that relies on pre-built maps? What are the limitations of the latter? (3) What are some practical architectural design choices for optimizing mobility and perception together? We empirically answer each of these questions in Section 5.3. Our real-world experiments primarily focus on comparing the capabilities of our reactive, learned system to a classical mapping and then planning approach. For our method, we observe interesting scenarios demonstrating emergent behaviors during real-world experiments detailed in Section 5.1 and 5.2.

5. Results and Analysis

	Reach						Pick	Place					
	Scenario 1			Scenario 2				Scenario 1			Scenario 2		
	Easy	Medium	Hard	Easy	Medium	Hard		Easy	Medium	Hard	Easy	Medium	Hard
<i>Success Rate:</i>													
FixCam	1.00	0.53	0.20	1.00	0.50	0.26	0.86	1.00	0.53	0.16	0.97	0.50	0.20
NoPointNet	1.00	0.87	0.57	1.00	0.77	0.63	0.93	1.00	0.83	0.57	1.00	0.77	0.60
Mapping	1.00	1.00	1.00	0.86	1.00	0.97	0.97	1.00	1.00	1.00	1.00	0.90	0.97
SPIN(DVO)	1.00	1.00	0.96	1.00	1.00	0.90	0.97	1.00	1.00	0.90	1.00	0.90	0.90
SPIN(CVO)	1.00	0.97	0.93	1.00	1.00	0.93	0.97	1.00	0.97	0.90	1.00	0.97	0.93
<i>Average Episode Duration (s):</i>													
FixCam	6.86	23.48	38.94	6.54	27.24	42.36	11.00	7.25	17.06	41.07	8.02	20.24	45.98
NoPointNet	6.30	14.87	33.25	7.22	15.04	34.09	9.25	4.88	15.22	32.00	7.89	18.49	37.42
Mapping	6.20	14.02	26.24	6.55	12.28	28.05	4.98	6.77	9.85	22.29	4.86	12.62	26.12
SPIN(DVO)	5.92	16.25	28.44	7.32	17.12	32.04	9.24	7.22	18.45	34.24	9.40	15.98	41.24
SPIN(CVO)	6.24	14.00	23.57	6.55	15.81	29.31	5.74	6.51	13.39	27.25	8.03	12.79	31.25

Table 5.1: We evaluate the success rate on 10 random environments with an average of 3 fixed seeds across all difficulty scenarios based on obstacle course. We report the success rate of each part of the task including reaching (Reach), picking (Pick), and placing (Place) the target object in the desired location. The place task requires the agent to bring back the object across the obstacles near its start location.

5.1 Emergent Behavior

Large-scale simulation pre-training allows our robot to learn emergent behaviors to avoid obstacles in cluttered scenarios, even in the presence of dynamic obstacles. We see several such behaviors during real-world experimentation which were neither planned nor specifically trained for in simulation but emerge as a result of a large diversity of procedural environments seen during training. We illustrate three such scenarios in Figure 5.1. As highlighted in several frames, Figure ?? depicts robustness to adversarially placed dynamic obstacles that constantly block the path of the robot. It needs to continuously perceive its environment in multiple directions and quickly react to those changes. We observe that in cases when there is no feasible path for the robot to navigate through, it also learns to stop and look around in order to replan its path and avoid collisions. Similarly, in ?? we see that as soon as a floating obstacle is suddenly placed in front of the robot, it shows spatial awareness and whole-body coordination and lowers its arm in order to navigate through, instead of turning and replanning the entire base movement which would take more time. In Figure ??, we see an adaptive rerouting mechanism where the agent changed its straight line motion as soon as a person kicks in a box in front of it. These behaviors emerge

in real-time and show the ability of our system to continuously perceive, adapt and react to changes in its environment which is very hard for a classical planner.

Qualitative results are available at <https://spin-robot.github.io/>. Along with emergent behavior where the robot continuously avoids dynamic obstacles without seeing them during training, we also observe generalization from heavily cluttered indoor to dim-lit outdoor environments. The agent also demonstrates reactive whole-body coordination where it moves its arm up or down to efficiently navigate across floating obstacles instead of re-routing and re-planning base movement, demonstrating 3D spatial awareness as shown in many video resu.

5.2 Real-world results

We test on two real-world scenes - an academic lab and an open study area with couches and a kitchenette next to it with both static as well as dynamic obstacles. Both these environments have unstructured clutter and humans as dynamic obstacles that makes it challenging for the agent to navigate through these spaces. For each environment, we have 4 static obstacles and at most 1 dynamic obstacle thrown adversarially. We compare against a classical baseline that uses an A1 RPLidar with gmapping and move_base for planning. We first teleoperate with the robot for 3-5min to create a map. Note that this provides the added advantage that this baseline knows the entire map in advance. Since the Lidar cannot see objects above the plane we only test on ground obstacles and ignore floating ones that require whole-body coordination. We run the planner to only plan the base motion. In Tab. 5.2 we compare the success rate and average number of collisions. An episode succeeds when the robot reaches within 15cm of the specified goal position. Overall, our method is able to succeed 20-40% more than the classical baseline. This is because the classical method suffers from noise is not able to recover from a noisy map, and gets stuck, whereas the learned policy learns to look at the obstacles again and again to improve its uncertainty estimates and constantly updates its knowledge of where obstacles are. This ability is even more apparent in the dynamic scenario (Table 5.2) where the classical has a near zero success rate while our method is able to succeed. It has the emergent ability to avoid a new obstacle in space, whereas the classical baseline relies on the pre-built map and fails entirely. Note that, we do not train

our policy with dynamic obstacles in simulations, but this behavior comes out as a by-product of lots of diverse experience in simulation. We design the observation space such that everything is relative to the robot. This allows the agent to perceive the environment as moving within its local reference frame, allowing generalization to dynamic obstacles.

5.3 Simulation results

The simulation benchmarks have 6 scenes, 2 of each easy, medium, and hard environments. Easy environments have 0-1 obstacles within a 5m goal range. Medium environments have 2-3 obstacles within 5m and the hard ones have heavily cluttered scenes with 5 obstacles within 5m. In each of these cases, one scene (Scenario 1) comprises a tight 1m wide long corridor which bounds the agent to not take shortcuts and reach the goal only by navigating through obstacles. The second (Scenario 2) is an L-shaped corridor with a goal at the end. The evaluation metrics are reported as an average of 10 episodes with random agent and obstacle initialization across 3 seeds.

We compare against various baselines to study the impact of our design decisions in Tab. 5.1. For each scenario, we report the success rate and average episode length across 10 rollouts. Our method achieves $\approx 33\%$ higher success rate than the NoPointNet baseline since permutation invariant scandots latent makes the optimization problem easier and also generalizes better at test time. Ours achieves $\approx 68\%$ higher success rate than the FixCam baseline with the camera pointing straight ahead. This is because in some cases the robot encounters obstacles in its peripheral vision and our policy can change the camera angle to avoid them. Active vision is necessary for the robot to move effectively through a cluttered environment. Our method is significantly better than the Mapping baseline because the systematic noise in the object locations makes it hard for the robot to avoid them, especially in cluttered environments, whereas our method can continuously estimate the position of obstacles while it is moving and adapt the motion online. Finally, we compare between the decoupled (DVO) and coupled (CVO) variants of our method and find that they achieve similar performance. We hypothesize that the partial observability and joint optimization for camera and robot actions in CVO training allows the

	Static Obstacles		Dynamic Obstacles	
Scenario 1				
	Ours	Classical	Ours	Classical
Average Success	0.8	0.6	0.6	0.0
Average # Collisions	1.0	0.4	1.6	1.2
Scenario 2				
	Ours	Classical	Ours	Classical
Average Success	0.8	0.4	0.6	0.2
Average # Collisions	0.8	0.6	1.6	1.0

Table 5.2: We compare our method against a classical mapping and planning baseline for navigation in cluttered scenes with both static as well as dynamic obstacles. The classical performs reasonably in static environments, it quickly breaks with dynamic obstacles like humans walking around, whereas our method shows more robust reactivity to such obstacles even without being trained with dynamic obstacles in simulation. We report the success rate of our method compared with the baseline. For the classical baseline, we teleoperate the robot for 2-3 min.

agent to quickly discover optimal shortcuts that are otherwise harder to distill from a privileged teacher policy.

5.4 Directly training from depth images.

We compare training from depth (red) and our 2-phase method with scandots (blue) in a medium-difficulty environment as illustrated in Figure 5.2. Depth policy has $< 1\%$ success after 22h training, whereas total (*phase1 + 2*) wall-clock time for SPIN is 16h (6+10). The simulator gives $\approx 50k$ fps for scandots (8192 envs) and ≈ 820 fps for depth (256 environments – maximum parallel environments that can fit on a single GPU), causing $61\times$ slow-down bottleneck. This shows the necessity and efficiency of our proposed 2-phased coupled visuomotor optimization approach using scandots over naively training an RL policy from depth observations.

	Scenario 1			Scenario 2		
	E	M	H	E	M	H
I	0.93	0.40	0.20	0.97	0.40	0.20
II	0.70	0.30	0.10	0.67	0.47	0.10
III	0.86	0.33	0.10	0.77	0.30	0.10
IV	1.00	0.53	0.20	1.00	0.50	0.26

Table 5.3: Success rate for 4 FixCam poses in easy (E), medium (M), hard (H) environments.

5.5 Analysing camera and base motion

Fixed Camera Baseline We run FixCam baseline with 4 camera poses (Figure ??) – **I**: Front, **II**: Down, **III**: Down and slightly front, **IV**: Front and slightly down on easy (E), medium (M), hard (H) environments. **I**, **IV** with max FOV have much lower success than SPIN, implying active vision is required in clutter. Pose (**IV** depicts the FixCam baseline referred to in the paper.

Camera Movement and Camera Observations: We show camera trajectory in Figure 5.3. When navigating through the clutter (frames 1, 2, 4), it tilts downward to maximize FOV near the base, but with no nearby obstacle (frame 3), it faces *front*. Detailed movement of the camera can be seen in the video demos on the website along with paired RGBD images for rollouts. RGB frames are only for analysis, the policy only observes depth images.

5.6 Necessity for Active Vision and Whole-body coordination


Active Vision: In principle, a multi-camera system should be equivalently adequate, however, most views will contain insignificant information and require large models to process. With limited onboard compute on most robots and the requirement for real-time reactivity ($< 0.1s$), it becomes infeasible to deploy them with larger vision backbones.

Whole-body coordination (WBC): Under heavy obstructions, the robot cannot move without collision if the base & arms control are decoupled. Figure 5.4 shows such a scenario where a fixed arm and gripper close to the base would fail without WBC, which also allows it to use the extra degree of freedom to find shorter and more efficient paths.

5.7 Classical Navigation Baseline

As discussed in Section 4, we compare our method with a classical map-based baseline which uses the 2D RPLidar to build an environment map and then creates a plan using the Monte Carlo method. We observe that for $> 90\%$ of the cases, the robot is able to build a map and find a feasible path, however, it is not able to execute the planned path for $> 85\%$ cases. This issue arises due to noisy control or unexpected wheel motion due to terrain differences. Our method is able to overcome such failures due to constant feedback and reactive improvisation through proprioception as well as depth, which allows it to deal with uncertainties without requiring a pre-built environment map. Moreover, due to localization inaccuracies, the baseline method is often unable to reach the intended goal if not initialized in the same orientation as was used before building the map. Contrary to that, we heavily randomize all degrees of freedom as well as the robot orientation at the beginning of every rollout during test time. Moreover, since the robot has a 2D Lidar installed on it, we do not test it in environments with floating obstacles which would require 3D understanding and whole-body coordination to navigate through clutter. We show some visualizations of the map built and plans created by the robot in Figure 5.5.

5. Results and Analysis



(a) While our simulation lacks dynamic obstacles, the robot can still evade them because the policy continuously adjusts its plan.



(b) If there is an overhanging obstacle, the robot lowers its arm to avoid it instead of turning around, thereby displaying agile whole-body coordination.



(c) The robot adapts its plan on-the-fly to obstacles by turning around once the camera sees it.

Figure 5.1: Types of **emergent behavior** exhibited by *SPIN* (a) dynamic obstacle avoidance (b) whole-body movement (c) adaptive rerouting.

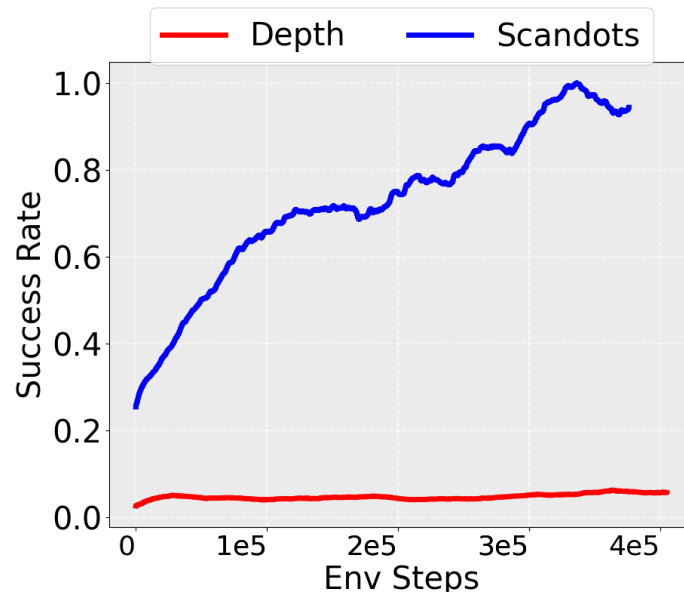


Figure 5.2: Success rate for scandots (blue) vs depth (red). The depth-based policy attains close to 0 performance even after 400k env steps of training, whereas the policy trained with scandots increasingly improves over time.




Figure 5.3: Camera movement analysis in a trajectory. The agent faces the camera downward when navigating through a tightly cluttered vicinity as can be seen in the first, second, and fourth frame, whereas the camera points more toward the front when there are no immediate obstacles in the direction of movement, as illustrated in the third frame.

5. Results and Analysis




Figure 5.4: Scenarios where whole-body coordination is essential under heavy obstructions. In the above cases where the obstacles are tightly packed, it is not possible for the robot to navigate through them avoiding collisions without lifting the arm to an appropriate height.




Figure 5.5: Visualizations of environment map built using 2D Lidar. The robot is localized as per its initial position and orientation.

Chapter 6

Discussion

We present *SPIN*, an approach to train embodied mobile robots that can simultaneously perceive, interact, and navigate cluttered environments using a data-driven approach. We show that our RL-based reactive approach is effective for active whole-body control-perception problem, traditionally addressed via non-reactive planning methods. With recent interest in humanoid and other mobile robots with actuated cameras, on neck for instance, *SPIN* is a cost-effective agile whole-body control solution with limited sensing and compute.

Although our robot can perceive geometry and avoid obstacles using depth, it still operates on stereo-matched depth instead of raw RGB. This leads to scenarios where it can bump into glass obstacles or shiny surfaces. In the future, we would like to use RGB for perception.

6. Discussion

Bibliography

- [1] Stretch by hello robot. <https://hello-robot.com/>. 3, 4
- [2] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on Robot Learning*, pages 403–415. PMLR, 2023. 3.2, 4.1
- [3] Ananye Agarwal, Shagun Uppal, Kenneth Shaw, and Deepak Pathak. Dexterous functional grasping, 2023. 2.4
- [4] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022. 2.4
- [5] Sarthak Bhagat and P.B. Sujit. Uav target tracking in urban environments using deep reinforcement learning. In *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, pages 694–701, 2020. doi: 10.1109/ICUAS48674.2020.9213856. 1
- [6] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022. 2.4
- [7] Wolfram Burgard, Dieter Fox, Daniel Hennig, and Timo Schmidt. Estimating the absolute position of a mobile robot using position probability grids. In *Proceedings of the national conference on artificial intelligence*, pages 896–901, 1996. 2.1
- [8] Ben Burgess-Limerick, Chris Lehnert Jurgen Leitner, and Peter Corke. Enabling failure recovery for on-the-move mobile manipulation. *arXiv preprint arXiv:2305.08351*, 2023. 2.4
- [9] Anthony R Cassandra, Leslie Pack Kaelbling, and James A Kurien. Acting under uncertainty: Discrete bayesian models for mobile-robot navigation. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS'96*, volume 2, pages 963–972. IEEE, 1996. 2.1

- [10] Matthew Chang, Theophile Gervet, Mukul Khanna, Sriram Yenamandra, Dhruv Shah, So Yeon Min, Kavitha Shah, Chris Paxton, Saurabh Gupta, Dhruv Batra, et al. Goat: Go to any thing. *arXiv preprint arXiv:2311.06430*, 2023. 2.2
- [11] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. *arXiv preprint arXiv:2004.05155*, 2020. 2.2
- [12] Prithvijit Chattopadhyay, Judy Hoffman, Roozbeh Mottaghi, and Aniruddha Kembhavi. Robustnav: Towards benchmarking robustness in embodied navigation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15691–15700, 2021. 2.2
- [13] Tao Chen, Saurabh Gupta, and Abhinav Gupta. Learning exploration policies for navigation. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=SyMwn05F7>. 2.2
- [14] Toon Van de Maele, Tim Verbelen, Ozan Çatal, Cedric De Boom, and B. Dhoedt. Active vision for robot manipulators using the free energy principle. *Frontiers in Neurorobotics*, 15, 2021. URL <https://api.semanticscholar.org/CorpusID:232116552>. 2.3
- [15] Yilun Du, Chuang Gan, and Phillip Isola. Curious representation learning for embodied intelligence. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10388–10397, 2021. URL <https://api.semanticscholar.org/CorpusID:233481519>. 2.3
- [16] Yuqing Du, Daniel Ho, Alex Alemi, Eric Jang, and Mohi Khansari. Bayesian imitation learning for end-to-end mobile manipulation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 5531–5546. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/du22b.html>. 2.4
- [17] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. Markov localization for mobile robots in dynamic environments. *Journal of artificial intelligence research*, 11:391–427, 1999. 2.1
- [18] Zipeng Fu, Xuxin Cheng, and Deepak Pathak. Deep whole-body control: learning a unified policy for manipulation and locomotion. In *Conference on Robot Learning*, pages 138–149. PMLR, 2023. 2.4
- [19] 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. 2.2
- [20] J-S Gutmann, Wolfram Burgard, Dieter Fox, and Kurt Konolige. An experi-

- mental comparison of localization methods. In *Proceedings. 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems. Innovations in Theory, Practice and Applications (Cat. No. 98CH36190)*, volume 2, pages 736–743. IEEE, 1998. 2.1
- [21] Peter E Hart, Nils J Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. 2.1
- [22] Jesse Haviland, Niko Sünderhauf, and Peter Corke. A holistic approach to reactive mobile manipulation. *IEEE Robotics and Automation Letters*, 7(2):3122–3129, 2022. 2.4
- [23] Jiaheng Hu, Peter Stone, and Roberto Martín-Martín. Causal policy gradient for whole-body mobile manipulation. *arXiv preprint arXiv:2305.04866*, 2023. 2.4
- [24] Glenn Jocher. YOLOv5 by Ultralytics, May 2020. URL <https://github.com/ultralytics/yolov5>. 4.2
- [25] Dvij Kalaria, Shreya Sharma, Sarthak Bhagat, Haoru Xue, and John M. Dolan. Wroom: An autonomous driving approach for off-road navigation. 2024. URL <https://api.semanticscholar.org/CorpusID:269149459>. 3.2
- [26] Oussama Khatib. The potential field approach and operational space formulation in robot control. In *Adaptive and Learning Systems: Theory and Applications*, pages 367–377. Springer, 1986. 2.1
- [27] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, volume 3, pages 2149–2154 vol.3, 2004. doi: 10.1109/IROS.2004.1389727. 4.1
- [28] Jason Ku, Ali Harakeh, and Steven L Waslander. In defense of classical image processing: Fast depth completion on the cpu. In *2018 15th Conference on Computer and Robot Vision (CRV)*, pages 16–22. IEEE, 2018. 4.2
- [29] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. 2021. 3.2
- [30] Mathieu Labbé and François Michaud. Rtab-map as an open-source lidar and visual simultaneous localization and mapping library for large-scale and long-term online operation. *Journal of field robotics*, 36(2):416–446, 2019. 2.1
- [31] Steven M LaValle, James J Kuffner, BR Donald, et al. Rapidly-exploring random trees: Progress and prospects. *Algorithmic and computational robotics: new directions*, 5:293–308, 2001. 2.1
- [32] Samuel Li, Sarthak Bhagat, Joseph Campbell, Yaqi Xie, Woojun Kim, Katia P. Sycara, and Simon Stepputtis. Shapegrasp: Zero-shot task-oriented grasping

- with large language models through geometric decomposition. 2024. 2.4
- [33] I-Chun Arthur Liu, Shagun Uppal, Gaurav S. Sukhatme, Joseph J. Lim, Peter Englert, and Youngwoon Lee. Distilling motion planner augmented policies into visual control policies for robot manipulation. In *Conference on Robot Learning*, 2021. 2.4, 3.1.2
- [34] Jun S Liu and Rong Chen. Sequential monte carlo methods for dynamic systems. *Journal of the American statistical association*, 93(443):1032–1044, 1998. 2.1
- [35] Haokuan Luo, Albert Yue, Zhang-Wei Hong, and Pulkit Agrawal. Stubborn: A strong baseline for indoor object navigation. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3287–3293. IEEE, 2022. 2.2
- [36] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021. 4
- [37] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7(62):eabk2822, 2022. 4.1
- [38] Raul Mur-Artal, Jose Maria Martinez Montiel, and Juan D Tardos. Orb-slam: a versatile and accurate monocular slam system. *IEEE transactions on robotics*, 31(5):1147–1163, 2015. 2.1
- [39] Deepak Pathak, Yide Shentu, Dian Chen, Pulkit Agrawal, Trevor Darrell, Sergey Levine, and Jitendra Malik. Learning instance segmentation by interaction. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2123–21233, 2018. URL <https://api.semanticscholar.org/CorpusID:49349927>. 2.3
- [40] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011. 3.2
- [41] Stergios I Roulmliotis and George A Bekey. Bayesian estimation and kalman filtering: A unified framework for mobile robot localization. In *Proceedings 2000 ICRA. Millennium conference. IEEE international conference on robotics and automation. Symposia proceedings (Cat. No. 00CH37065)*, volume 3, pages 2985–2992. IEEE, 2000. 2.1
- [42] Muhamad Risqi U Saputra, Andrew Markham, and Niki Trigoni. Visual slam and structure from motion in dynamic environments: A survey. *ACM Computing Surveys (CSUR)*, 51(2):1–36, 2018. 2.1

- [43] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017. [3.1.2](#)
- [44] Shinya Sumikura, Mikiya Shibuya, and Ken Sakurada. Openvslam: A versatile visual slam framework. In *Proceedings of the 27th ACM International Conference on Multimedia*, pages 2292–2295, 2019. [2.1](#)
- [45] Sebastian Thrun. Particle filters in robotics. In *UAI*, volume 2, pages 511–518. Citeseer, 2002. [2.1](#)
- [46] Sebastian Thrun, Dieter Fox, Wolfram Burgard, and Frank Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1):99–141, 2001. ISSN 0004-3702. doi: [https://doi.org/10.1016/S0004-3702\(01\)00069-8](https://doi.org/10.1016/S0004-3702(01)00069-8). URL <https://www.sciencedirect.com/science/article/pii/S0004370201000698>. [2.1](#)
- [47] Shagun Uppal, Sarthak Bhagat, Devamanyu Hazarika, Navonil Majumder, Soujanya Poria, Roger Zimmermann, and Amir Zadeh. Multimodal research in vision and language: A review of current and emerging trends. *Information Fusion*, 77:149–171, 2022. ISSN 1566-2535. doi: <https://doi.org/10.1016/j.inffus.2021.07.009>. URL <https://www.sciencedirect.com/science/article/pii/S1566253521001512>. [2.2](#)
- [48] Greg Welch, Gary Bishop, et al. An introduction to the kalman filter. 1995. [2.1](#)
- [49] Paul J Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. [3.1.2](#)
- [50] 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. *arXiv preprint arXiv:1911.00357*, 2019. [2.2](#)
- [51] Denis F Wolf and Gaurav S Sukhatme. Mobile robot simultaneous localization and mapping in dynamic environments. *Autonomous Robots*, 19:53–65, 2005. [2.1](#)
- [52] Josiah Wong, Albert Tung, Andrey Kurenkov, Ajay Mandlekar, Li Fei-Fei, Silvio Savarese, and Roberto Martín-Martín. Error-aware imitation learning from teleoperation data for mobile manipulation. In *Conference on Robot Learning*, pages 1367–1378. PMLR, 2022. [2.4](#)
- [53] Naoki Yokoyama, Alexander William Clegg, Eric Undersander, Sehoon Ha, Dhruv Batra, and Akshara Rai. Adaptive skill coordination for robotic mobile manipulation. *arXiv preprint arXiv:2304.00410*, 2023. [2.4](#)
- [54] Chao Yu, Zuxin Liu, Xin-Jun Liu, Fugui Xie, Yi Yang, Qi Wei, and Qiao Fei. Ds-slam: A semantic visual slam towards dynamic environments. In *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*,

Bibliography

pages 1168–1174. IEEE, 2018. [2.1](#)