

Seeing-Eye Quadruped Navigation with Force Responsive Locomotion Control

Anonymous Author(s)

Affiliation

Address

email

Abstract:

Seeing-eye robots are very useful tools for guiding visually impaired people, potentially producing a huge societal impact given the low availability and high cost of real guide dogs. Although a few seeing-eye robot systems have already been demonstrated, none considered external tugs from humans, which frequently occur in a real guide dog setting. In this paper, we simultaneously train a locomotion controller that is robust to external tugging forces via Reinforcement Learning (RL), and an external force estimator via supervised learning. The controller ensures stable walking, and the force estimator enables the robot to respond to the external forces from the human. These forces are used to guide the robot to the global goal, which is unknown to the robot, while the robot guides the human around nearby obstacles via a local planner. Experimental results in simulation and on hardware show that our controller is robust to external forces, and our seeing-eye system can accurately detect force direction. We demonstrate our full seeing-eye robot system on a real quadruped robot with a blindfolded human.




Figure 1: Tugs on the seeing-eye robot are used as navigation cues during blindfolded navigation.

1 Introduction

In the typical seeing-eye dog (also known as guide dog) setting, a human holds a rigid handle attached to a harness on a dog. The dog will safely guide the human around nearby obstacles, while the human can tug on the dog to indicate which general direction to move in. Thus, the human has some idea of where they are, and is capable of making high-level navigation decisions, while the dog can better sense its immediate surroundings for obstacle avoidance and locomotion. Guide dogs have been shown to improve the lives of visually impaired people, via increasing independence, confidence, companionship, and mobility [1]. Unfortunately, guide dogs need roughly two years of training, and cost over \$50,000 USD per dog [2]. Despite efforts to reduce the production cost of guide dogs [3], their supply is still significantly lower than demand.

To better meet this demand, and potentially improve performance, seeing-eye robot systems are being developed [4, 5, 6, 7, 8, 9]. There is a growing interest in the development of systems of this

28 type, which have undergone various human studies in order to measure the extent to which they are
29 compatible with visually impaired humans, and their level of societal acceptance [10, 11, 12]. While
30 these seeing-eye robot systems have successfully demonstrated navigation tasks, none of them have
31 considered settings involving human tugs, which is very common for seeing-eye dogs. It’s important
32 for seeing-eye robots to be robust to human forces, as the human is constantly holding a rigid handle
33 directly attached to the robot, and large forces can cause the robot to stray from the optimal path, or
34 even fall over. Furthermore, a typical method of communication between a human and a real guide
35 dog is through tugs. This makes force estimation useful for a seeing-eye robot, as knowledge of the
36 direction and magnitude of the applied force can be used to better facilitate navigation according to
37 the human’s intentions (see Figure 1).¹

38 To address the above mentioned issues in human-robot communication, we develop a novel seeing-
39 eye robot system which is robust to external forces, and estimates the magnitude and direction
40 of these forces to determine which navigation actions to take for human-robot co-navigation. We
41 achieve this by simultaneously training a locomotion controller via RL, and a force estimator via
42 supervised learning. Our locomotion policy is trained over simulated tugs by sampling different
43 base velocities which are suddenly applied to the robot [13]. These tugs serve as labels for training
44 our force estimator, and are estimated during deployment.

45 While the controller is running on a real robot, the force estimator estimates the direction and mag-
46 nitude of force the human applies. These force estimates are computed exclusively from sensors
47 onboard the robot (joint encoders and IMU). From force estimates, the robot detects when and in
48 what direction the human tugs occur. This informs the robot which direction to go at a global level,
49 while a local planner using information from a LIDAR sensor is used to navigate the immediate
50 environment. Different from existing seeing-eye robot systems that require major hardware up-
51 grades, e.g., customized traction device [6], or button interface [7], our seeing-eye system needs is
52 compatible with any attachable leash or handle.

53 Our main contributions include the following:

- 54 1. The first seeing-eye robot system which takes directional cues via human tugs, while also
55 safely navigating the immediate environment.
- 56 2. A force tolerant locomotion controller, jointly trained with a force estimator which can
57 estimate the magnitude and direction of human forces.
- 58 3. Experimental results in simulation and on hardware to evaluate the robustness of our loco-
59 motion controller and accuracy of our force estimator.
- 60 4. Demonstration of our seeing-eye robot system in an indoor environment with a blindfolded
61 human.

62 2 Related Work

63 Various seeing-eye robot systems have been demonstrated for the task of blindfolded navigation.
64 Among these works, one of them considers an MPC-based motion planner for a wheeled robot [4].
65 Another considers an optimization based approach for path planning, which models a taut or slack
66 leash [5]. A third system designs an adjustable leash, and optimizes for human comfort during
67 navigation [6]. These works all stray from real-world guide dog settings, in part because they assume
68 the robot has full knowledge of the destination beforehand, and that the human does not need to
69 communicate with the robot during navigation. In real-world settings, the human must decide on
70 which high-level navigation actions to take, and communicate these actions to the robot.

71 More similar to our work, other approaches consider settings where the human communicates high-
72 level directions to the robot during navigation [7, 8]. However, the medium of communication in

¹The term of “Seeing-eye robots” has been used by researchers that refer to quadruped robots to guide visually-challenged people [26]. Although our robot doesn’t use vision, it serves as a seeing-eye platform through its Lidar sensors for navigation and obstacle avoidance.




Figure 2: Overview of our approach. Our locomotion controller (circled in red) contains a velocity estimator, force estimator, and locomotion policy, all of which are trained in simulation. The base velocity estimator and force estimator are trained via supervised learning, using privileged information from the simulator as labels. The locomotion policy is trained via RL, and outputs target joint angles to a PD controller which converts them to joint torques which are directly applied to the robot. During deployment, our locomotion controller estimates external force at each time step. Force direction is derived from peaks in the estimated force signal. The direction of force determines the next local navigation goal for our navigation system to take, which returns velocity commands to our controller.

73 these works are either a custom designed handle with buttons [7], or predefined directional actions
 74 prior to navigation [8]. In our work, the human communicates directional cues via tugging, and is
 75 compatible with any rigid connection to the robot. Additionally, neither of these works consider
 76 human forces being applied to the robot, making their systems unresponsive or susceptible to failure
 77 upon navigation under human forces.

78 Numerous quadruped controllers robust to external forces have already been developed [14, 15, 16].
 79 However these works don't explicitly estimate the applied forces, as the focus of these works is to
 80 develop controllers that are generally robust to varying environments. In our work, we explicitly es-
 81 timate forces to the robot generated by human tugs, in order to communicate the human's navigation
 82 intentions with the robot.

83 3 Method

84 In this section, we present our seeing-eye robot system that is robust and responsive to external
 85 forces from human tugs. Figure. 2 presents an overview of how we train our locomotion controller,
 86 and deploy for seeing-eye navigation.

87 3.1 Locomotion Controller

88 We train our locomotion controller via RL, which models environments as a Markov Decision Pro-
 89 cess (MDP). An MDP is defined as $M = (S, A, T, R, \gamma)$, where S is the set of states, A is the set of
 90 actions, $T : S \times A \times S \rightarrow [0, 1]$ is the transition function which outputs the probability of reaching
 91 state s' given state s and action a , $R : S \times A \times S \rightarrow \mathbb{R}$ is the reward function which returns feedback
 92 from taking action a from state s and ending up in state s' , and $\gamma \in [0, 1]$ is the discount factor which
 93 determines how valuable future reward should be considered in comparison to immediate reward.

94 We define a robot state at time t as $\mathbf{x}_t = (\mathbf{c}, \dot{\mathbf{v}}, \mathbf{q}, \dot{\mathbf{q}}, \mathbf{a}_{t-1})$, where $\mathbf{c} = (c_x, c_y, c_w)$ is the commanded
 95 base linear and angular velocity, $\dot{\mathbf{v}}$ is the base acceleration, \mathbf{q} and $\dot{\mathbf{q}}$ are the joint angles and velocities
 96 respectively, and \mathbf{a}_{t-1} is the action taken at time $t - 1$. Actions are target joint angles, which are
 97 converted to torques via a PD controller. The reward function encourages tracking commands \mathbf{c}
 98 while minimizing energy consumption and large action changes [13] – fully defined in Table 1.

Table 1: All terms of the reward function our locomotion policy is trained on. \mathbf{v} refers to base velocity, \mathbf{c} refers to commanded linear and angular base velocity, ω refers to base angular velocity, τ refers to joint torques, $\dot{\mathbf{q}}$ refers to joint velocities, t_{air} refers to each foot’s air time, \mathbf{a} refers to an action, and dt refers to the simulation time step.

Term Description	Definition	Scale
Linear Velocity x, y	$\exp(-\ \mathbf{c}_{x,y} - \mathbf{v}_{x,y}\ ^2/0.25)$	$1.0dt$
Linear Velocity z	\mathbf{v}_z^2	$-2.0dt$
Angular Velocity x, y	$\ \omega_{x,y}\ ^2$	$-0.05dt$
Angular Velocity z	$\exp(-(\mathbf{c}_\omega - \omega_z)^2/0.25)$	$0.5dt$
Joint Torques	$\ \tau\ ^2$	$-0.0002dt$
Joint Accelerations	$\ (\dot{\mathbf{q}}_{last} - \dot{\mathbf{q}})/dt\ ^2$	$-2.5e - 7dt$
Foot Air Time	$\sum_{f=1}^4 (t_{\text{air},f} - 0.5)$	$1.0dt$
Action Rate	$\ \mathbf{a}_{last} - \mathbf{a}\ ^2$	$-0.01dt$

99 Our locomotion policy also takes the base velocity and external force vector as input, to more easily
100 learn to track commands \mathbf{c} and respond to external forces. These variables are not easily estimated
101 through robot sensors. Thus, we train state estimators via supervised learning over privileged infor-
102 mation, which has been shown to be more effective than classical methods, e.g., Kalman Filters [17].
103 It is a common setting to learn with privileged information in simulation, and insert a corresponding
104 state estimator in real-world deployment. In line with those systems, we train a velocity estimator
105 and force estimator for the real robot to bridge the sim-to-real gap in locomotion policy learning.
106 These estimators are trained jointly with the locomotion policy, where $\mathbf{v} = (v_x, v_y, v_z)$ is ground
107 truth base velocity, and $\mathbf{F} = (F_x, F_y, F_z)$ is ground truth external force, both obtained as privileged
108 information from the simulator.

109 Two variations of forces are applied to the robot during training. One variation is small and frequent
110 backward pushes, which are designed to simulate a human following the robot with a taut leash,
111 where a human incidentally applies frequent small backward forces on the robot as they are being
112 guided. The other variation is larger, less frequent pushes occurring in any direction. These pushes
113 are designed to simulate human directional tugs, in which the human intentionally tugs the robot to
114 communicate the direction they want to move in. The force estimator is only trained on data from
115 the second variation of tugs, as we do not want to detect the small, incidental backward tugs which
116 naturally occur during guided navigation.

117 The base velocity estimator is a multilayer perceptron (MLP), whose parameters are updated over
118 the same data as the locomotion policy. To estimate forces, we find it helpful to access a history
119 of states, to better capture the robot’s behavior over the duration of the applied force. Thus, similar
120 to training adaptation modules [18], our force estimator uses 1-D convolutional layers to capture
121 temporal relationships between states. Force estimator parameters are updated less frequently than
122 the base velocity estimator and locomotion policy, because most time steps do not include external
123 forces. Thus, most of the labels for our force estimator includes zero vectors, indicating that no
124 force was applied at those time steps. This causes imbalanced training data, which we resolve by
125 only training the force estimator when the training data contains nonzero forces. We then further
126 re-balance this to ensure at most 20% of the force estimator’s training samples include zero vectors
127 as labels.

128 3.2 Seeing-eye Robot Navigation

129 To perform navigation tasks with our seeing-eye robot, we need to estimate when and in what general
130 direction a force is being applied. This is done by running peak detection [19] on the previous 200
131 time steps of the estimated force signal F_y , at a rate of 2Hz. In this work, we only analyze F_y ,
132 to determine whether a left or right tug has occurred. A peak in the signal indicates that the force
133 estimator detected a significant external force applied to the robot. Thus, if a peak is detected within
134 the past 50 time steps, then we consider it as a recent tug applied by the human. Positive peaks
135 correspond to left tugs, while negative peaks correspond to right tugs.




Figure 3: Map of our navigation environment. Yellow circles correspond to *decision points*, where the human needs to decide which direction to move in via tugging. The blue lines indicate the path the human took in our demonstration.

136 Local navigation goals are then selected based on the robot’s location, orientation, and tug direction.
 137 A domain expert manually labels a map with *decision points*, where the human must decide which
 138 direction to go in next, based on the direction they came from, and their tug direction. A labelled
 139 map of the hallway domain we demonstrate our system on can be seen in Figure 3.

140 The local navigation goal is then sent to our navigation system, which uses a LIDAR sensor to
 141 localize itself on the map via AMCL [20], and compute local plans via DWA [21]. The local planner
 142 returns velocity commands \mathbf{c} to our controller, which our controller tracks.

143 4 Implementation Details

144 We use Isaac Gym [22] physics simulator, and train our controller with 2048 robots in parallel [13].
 145 Our locomotion policy is trained via PPO [23], while our base velocity estimator and force estimator
 146 are trained via supervised learning, with Mean Squared Error loss. Our locomotion policy and
 147 base velocity estimator are both MLPs with three and two hidden layers respectively, while our
 148 force estimator is a convolutional neural network with three 1-D convolutional layers, which makes
 149 predictions over the past 25 time steps. The PD controller converts target joint angles to torques
 150 with proportional gain set to 20 and derivative gain set to 0.5. The policy is queried at 50Hz, and
 151 control signals are sent at 200Hz. New velocity commands are sampled after each episode, where
 152 c_x , c_y , and c_ω are sampled uniformly from $[-1, 1]$. An episode terminates when any link other than a
 153 foot touches the ground, the base height is below 0.25 meters, or the episode has lasted 20 seconds.

154 To better facilitate sim-to-real transfer, we train over RANDOM_UNIFORM_TERRAIN which increases
 155 in difficulty based on a curriculum [13]. We also include noise in observations, and domain random-
 156 ization over different surface frictions.

157 We add random external forces in our environment, to improve robustness of our locomotion policy
 158 and collect data to train our force estimator. The small and frequent backward pushes occur every 0.6
 159 seconds, have a duration of 0.1 seconds, and sets the base velocity to 0.25 m/s backward. Meanwhile,
 160 the large and infrequent pushes used to train the force estimator occur every 3 seconds, have a
 161 duration sampled from $[0.24, 0.48]$ seconds, and sets the base velocity to a vector sampled from
 162 $F_x \in [-0.75, 0.75]$, $F_y \in [-0.75, 0.75]$, and $F_z \in [0, 0.1]$. Note that forces from \mathbf{F} are implemented
 163 as spontaneous updates in base velocity.

164 5 Experiments

165 We design experiments to evaluate the robustness of our controller, and accuracy of our force esti-
 166 mator. Although we are able to learn a force estimator in simulation, we could not directly evaluate
 167 it in the real world due to the missing ground-truth values. Instead, we chose to evaluate tug detec-
 168 tion (LEFT, RIGHT, and NONE) on the real robot, where the participants followed our instructions,

Table 2: Our learned controllers fell significantly less frequently, and better tracked velocity commands under external forces when compared to an MPC controller. Including the output of the force estimator in the state marginally improved robustness. The large variance in drift is caused by the large range of force strengths and directions we sample from.

Controller	Proportion Fell	Drift from Trajectory
MPC	0.5990	1.1256 ± 0.5908
Learned No Est	0.1904	0.6824 ± 0.5447
Learned Est	0.1762	0.6790 ± 0.5309

169 and hence ground truth was available. We then demonstrate our full seeing-eye system via guiding
 170 a blindfolded human.

171 5.1 Force Tolerance Evaluation

172 To evaluate whether our learned force controller is actually robust to external forces, we run exper-
 173 iments in simulation where we apply random forces to the base of the robot. In each trial, a single
 174 force in a random direction is applied, whose duration is sampled from $[0.25, 0.5]$ seconds, and
 175 strength is sampled from $[25, 100]$ Newtons. Meanwhile, the robot is commanded to walk forward
 176 at 0.5 meters/second, for a duration of five seconds.

177 We run this experiment on three different types of locomotion controllers, for 1000 trials each. The
 178 controllers include a commonly used MPC controller [24], a variant of our learned controller which
 179 does not consider estimated force in its state (referred to as *Learned No Est*), and our controller
 180 described in Section 3.1 (referred to as *Learned Est*). All controllers are deployed in PyBullet [25].

181 We measure how frequently the robot fell across all trials (Proportion Fell), and how far the external
 182 force caused the robot to drift from its current trajectory on average (Drift from Trajectory). We
 183 consider a robot to have fell if a non-foot part of the robot touches the ground. Results are reported
 184 in Table 2, which indicate that our learned controllers fell significantly less frequently, and better
 185 maintained velocity tracking under external forces than the MPC controller. Including estimated
 186 forces in the state appears to marginally improve robustness. Note that for our two types of learning-
 187 based controllers, we train over five different random seeds each and average the results.

188 5.2 Force Estimation Evaluation

189 We evaluate the accuracy of our force estimator through experiments in simulation, and on hardware.

190 5.2.1 Simulation

191 We deploy our learned controllers in PyBullet with a ve-
 192 locity command of 0.5 meters/second, and a single ex-
 193 ternal push per trial, for 1000 trials. Each push has a
 194 force whose x-component is sampled from $[-50, 50]$ New-
 195 tons, and a y-component which is at a fixed magnitude,
 196 and random direction (either left or right). There are
 197 three possible classes the force estimator can predict over:
 198 $\{\text{LEFT}, \text{RIGHT}, \text{NONE}\}$. Each trial includes 150 time
 199 steps and takes 3 seconds. In each trial, the force estima-
 200 tor is queried every 25 time steps, or six total queries per
 201 trial. Thus, in each trial, the force estimator makes a total
 202 of six predictions. A trial is deemed correct if one of the
 203 six queries matches the force direction being applied in
 204 the simulator.

205 In order to evaluate whether force direction can be esti-
 206 mated through only a history of base velocities, we train
 207 a baseline force estimator which only makes force esti-
 208 mates based on a history of ground-truth base velocity
 209 and base velocity commands. We refer to this baseline as




Figure 4: We report the accuracy and false positive rate of our force estimators, given forces of varied strength. The shaded region indicates the standard deviation between the five policies trained over five different random seeds.

210 *only vel*, while our force estimator trained over the full state and described in Section 3.1 is referred
211 to as *full state*.

212 In Figure 4, we report the accuracy and false positive rates of our force estimators over varying
213 force strengths. Accuracy is computed by dividing the number of trials in which the force estimator
214 predicted the correct force direction at any of the six queries in the trial, by the number of total
215 trials. Computing accuracy alone in this manner is not informative enough, because it is possible to
216 achieve high accuracy by predicting LEFT at one point in the trial, and RIGHT at a different point
217 within the same trial.

218 Thus, we also consider the false positive rate, which we compute by dividing the number of extra
219 forces (LEFT or RIGHT predicted when ground truth is NONE) predicted during the trial, by the
220 number of times the force estimator is queried (every 25 time steps). A high false positive rate corre-
221 sponds to the force estimator oftentimes predicting forces when they do not occur. As force strength
222 increases, our estimators achieve a higher accuracy while maintaining a relatively low false positive
223 rate. Results indicate that knowledge of the full state is significantly beneficial in estimating force
224 direction, when compared to a force estimator which is only trained over base velocity information.

225 5.2.2 Hardware

226 When a human tugs our seeing-eye robot with sufficient force, the base of the robot will momentarily
227 accelerate in the direction of the tug. Thus, one might wonder why we do not consider accelerometer
228 signals to detect tug direction, rather than train a force estimator. In this sub-section, we validate the
229 usefulness of our estimated force signals, which we compare to accelerometer readings.

230 In this experiment, we command a real Unitree A1
231 robot to move forward at 0.5m/s, while a human par-
232 ticipant tugs left after a few seconds of forward lo-
233 comotion, followed by a right tug after another few
234 seconds of locomotion. This trial is performed by
235 four human participants, three of which have no prior
236 experience in operating this tugging interface. Each
237 participant is a robotics researcher in a university lab.
238 The three participants with no prior experience with
239 this system were given a demonstration of an example
240 trial, before completing their own trials. In total, 42
241 trials were conducted, and data was collected from
242 40 trials (two trials were removed due to the robot
243 falling over and data not being saved). Of these 40
244 trials, each participant performed ten of them.

245 Forces are detected every 25 timesteps, and can be
246 predicted as one of three classes. We compute the ac-
247 curacy and false positive rate for force detectors using
248 the estimated force signal, compared to force detec-
249 tion using the signal directly from the accelerometer
250 on the robot. Accuracy is computed as the percentage
251 of trials which contain a LEFT force prediction be-
252 fore the halfway point of the trial, and a RIGHT force
253 prediction after the halfway point of the trial. False positive rate is the average percentage of false
254 positive forces being predicted across all trials. A force prediction is considered as a false positive
255 if it is either LEFT or RIGHT, and does not correspond to the first expected LEFT force or later
256 expected RIGHT force.

257 Results are reported in Table 3, where we find the force estimator more accurately predicted the
258 correct forces, with fewer false positives compared to predictions from raw accelerometer signals.
259 Both methods of detecting tug direction (accelerometer and force estimator) perform worse for be-




Figure 5: Measured acceleration (top) and estimated force (bottom) during a single trial. Tugs are denoted by red boxes.

Table 3: Force estimation via accelerometer readings vs force estimator signal.

Participant Type	Method	Accuracy	False Positive Rate
Expert (10 trials)	Accelerometer	0.20	0.2149
	Force Estimator	1.00	0.0442
Beginner (30 trials)	Accelerometer	0.13	0.1906
	Force Estimator	0.70	0.0498

260 ginner participants, indicating that adding diverse tugging styles is important for a more complete
 261 evaluation of force estimators.

262 In Figure 5, we plot the raw signals from the accelerometer and force estimator from a single trial.
 263 We find the accelerometer signal is more noisy than the estimated force signal, which we hypothe-
 264 size is because the force estimator has access to other sensor information along with accelerometer
 265 readings. Note that our force estimator is co-learned with the locomotion policy. The locomotion
 266 policy takes the estimated force as input during training, and the states generated from the policy
 267 are used to train the force estimator. We believe this co-learning mechanism leverages the additional
 268 sensor information to help the force estimator outperform the tug detection from raw accelerometer
 269 readings.

270 6 Hardware Demonstration

271 We demonstrate our full seeing-eye robot system on a Unitree A1 robot, in an indoor hallway envi-
 272 ronment (see link in Abstract). Our locomotion policy and force estimator are deployed on hardware
 273 without any additional fine-tuning.

274 In this demonstration, a human is blindfolded, and holding a taut leash attached to the robot. The
 275 human desires to reach some particular goal location, and chooses the rout to take by tugging the
 276 robot at decision points, while the robot autonomously avoids obstacles (including boxes, narrow
 277 doorways, and another human) along the way. Decision points occur at intersections in the hallway,
 278 where the human needs to decide in which direction they want to go. Similar to the setting in [7], the
 279 blindfolded human is not familiar with localizing himself without vision, so another sighted human
 280 verbally indicates when a decision point is approaching.

281 This demonstration indicates that our locomotion policy and force detector are transferable to hard-
 282 ware, and can cooperate with a local planner to enable blindfolded navigation. Thus, successful
 283 human-robot communication occurred, such that the human avoided all obstacles while the robot
 284 navigated to the desired goal location through detecting human tugs at decision points.

285 7 Discussion

286 **Limitations and Future Work** While our seeing-eye system can avoid obstacles and select routes
 287 at a course level in indoor hallway environments, real guide dog settings include outdoor environ-
 288 ments, and situations with many possible directions to navigate in. In future work, researchers can
 289 leverage methods to determine which paths are traversable [26], and train force detectors which can
 290 estimate the direction of force at a finer level. Another future direction is to incorporate intelligent
 291 disobedience, which refers to rejecting a human’s navigation decision if unsafe [27]. Additionally,
 292 The evaluation can be further improved by replacing sighted people with those with visual impair-
 293 ments in the experiments. Finally, our robot’s locomotion speed is relatively low. It might be the
 294 human tugs, the learned locomotion policy, or both causing the low speed. Further investigation into
 295 those factors can potentially lead to very interesting future research and higher-speed seeing-eye
 296 robot systems.

297 **Conclusion** We train a locomotion controller which is tolerant to human tugs, and can estimate
 298 the direction of external forces. We evaluate the robustness of our controller, and accuracy of our
 299 force estimator through experiments in simulation and on hardware. Finally, we demonstrate our
 300 controller on a real robot for the task of blindfolded navigation, where a blindfolded human is
 301 successfully guided to a destination, while giving directional cues through tugging on the robot.

References

- [1] L. Whitmarsh. The benefits of guide dog ownership. *Visual impairment research*, 7(1):27–42, 2005.
- [2] C. Morita. How much does a guide dog cost? <https://puppyintraining.com/how-much-does-a-guide-dog-cost/#:~:text=Initial%20cost%20for%20Guide%20Dog,for%20a%20guide%20dog%20%3D%20%2459%2C600>.
- [3] L. M. Tomkins, P. C. Thomson, and P. D. McGreevy. Behavioral and physiological predictors of guide dog success. *Journal of Veterinary Behavior*, 6(3):178–187, 2011.
- [4] L. Wang, J. Zhao, and L. Zhang. Navdog: robotic navigation guide dog via model predictive control and human-robot modeling. In *Proceedings of the 36th Annual ACM Symposium on Applied Computing*, pages 815–818, 2021.
- [5] A. Xiao, W. Tong, L. Yang, J. Zeng, Z. Li, and K. Sreenath. Robotic guide dog: Leading a human with leash-guided hybrid physical interaction. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11470–11476. IEEE, 2021.
- [6] Y. Chen, Z. Xu, Z. Jian, G. Tang, Y. Yangli, A. Xiao, X. Wang, and B. Liang. Quadruped guidance robot for the visually impaired: A comfort-based approach. *arXiv preprint arXiv:2203.03927*, 2022.
- [7] H. Hwang, T. Xia, I. Keita, K. Suzuki, J. Biswas, S. I. Lee, and D. Kim. System configuration and navigation of a guide dog robot: Toward animal guide dog-level guiding work. *arXiv preprint arXiv:2210.13368*, 2022.
- [8] J. T. Kim, W. Yu, J. Tan, G. Turk, and S. Ha. How to train your guide dog: Wayfinding and safe navigation with human-robot modeling. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, pages 221–225, 2023.
- [9] J. T. Kim, W. Yu, Y. Kothari, J. Tan, G. Turk, and S. Ha. Transforming a quadruped into a guide robot for the visually impaired: Formalizing wayfinding, interaction modeling, and safety mechanism. *arXiv preprint arXiv:2306.14055*, 2023.
- [10] Q. Chen, L. Wang, Y. Zhang, Z. Li, T. Yan, F. Wang, G. Zhou, and J. Gong. Can quadruped navigation robots be used as guide dogs? *arXiv preprint arXiv:2210.08727*, 2022.
- [11] Y. Zhang, Z. Li, H. Guo, L. Wang, Q. Chen, W. Jiang, M. Fan, G. Zhou, and J. Gong. ” i am the follower, also the boss”: Exploring different levels of autonomy and machine forms of guiding robots for the visually impaired. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pages 1–22, 2023.
- [12] P. Chonkar, G. Hemkumar, H. Wang, D. Dua, S. Gupta, Y.-C. Chan, J. Hart, E. Hauser, R. Mirsky, J. Biswas, et al. Look to my lead: How does a leash affect perceptions of a quadruped robot?
- [13] N. Rudin, D. Hoeller, P. Reist, and M. Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on Robot Learning*, pages 91–100. PMLR, 2022.
- [14] L. Campanaro, S. Gangapurwala, W. Merkt, and I. Havoutis. Learning and deploying robust locomotion policies with minimal dynamics randomization. *arXiv preprint arXiv:2209.12878*, 2022.
- [15] G. B. Margolis and P. Agrawal. Walk these ways: Tuning robot control for generalization with multiplicity of behavior. *arXiv preprint arXiv:2212.03238*, 2022.

- 345 [16] T. Miki, J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter. Learning robust per-
346 ceptive locomotion for quadrupedal robots in the wild. *Science Robotics*, 7(62):eabk2822,
347 2022.
- 348 [17] G. Ji, J. Mun, H. Kim, and J. Hwangbo. Concurrent training of a control policy and a state
349 estimator for dynamic and robust legged locomotion. *IEEE Robotics and Automation Letters*,
350 7(2):4630–4637, 2022.
- 351 [18] A. Kumar, Z. Fu, D. Pathak, and J. Malik. Rma: Rapid motor adaptation for legged robots.
352 *arXiv preprint arXiv:2107.04034*, 2021.
- 353 [19] M. Dede. peakdetect. <https://github.com/avhn/peakdetect>, 2022.
- 354 [20] B. P. Gerkey. Amcl, 2013. URL <http://wiki.ros.org/amcl>.
- 355 [21] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance.
356 *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.
- 357 [22] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin,
358 A. Allshire, A. Handa, et al. Isaac gym: High performance gpu-based physics simulation for
359 robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
- 360 [23] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization
361 algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- 362 [24] X. B. Peng, E. Coumans, T. Zhang, T.-W. E. Lee, J. Tan, and S. Levine. Learning agile
363 robotic locomotion skills by imitating animals. In *Robotics: Science and Systems, 07 2020*.
364 doi:10.15607/RSS.2020.XVI.064.
- 365 [25] E. Coumans and Y. Bai. Pybullet, a python module for physics simulation for games, robotics
366 and machine learning. 2016.
- 367 [26] J. Frey, D. Hoeller, S. Khattak, and M. Hutter. Locomotion policy guided traversability learning
368 using volumetric representations of complex environments. In *2022 IEEE/RSJ International
369 Conference on Intelligent Robots and Systems (IROS)*, pages 5722–5729. IEEE, 2022.
- 370 [27] R. Mirsky and P. Stone. The seeing-eye robot grand challenge: Rethinking automated care.
371 In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent
372 Systems (AAMAS 2021)*, 2021.