

Teaching Robots Like Dogs: Learning Agile Navigation from Luring, Gesture, and Speech

Taerim Yoon¹, Dongho Kang², Jin Cheng², Fatemeh Zargarbashi², Yijiang Huang², Minsung Ahn³, Stelian Coros², and Sungjoon Choi¹

Abstract—In this work, we aim to enable legged robots to learn how to interpret human social cues and produce appropriate behaviors through physical human guidance. However, learning through physical engagement can place a heavy burden on users when the process requires large amounts of human-provided data. To address this, we propose a human-in-the-loop framework that enables robots to acquire navigational behaviors in a data-efficient manner and to be controlled via multimodal natural human inputs, specifically gestural and verbal commands. We reconstruct interaction scenes using a physics-based simulation and aggregate data to mitigate distributional shifts arising from limited demonstration data. Our *progressive goal cueing* strategy adaptively feeds appropriate commands and navigation goals during training, leading to more accurate navigation and stronger alignment between human input and robot behavior. We evaluate our framework across six real-world agile navigation scenarios, including jumping over or avoiding obstacles. Our experimental results show that our proposed method succeeds in almost all trials across these scenarios, achieving a 97.15% task success rate with less than 1 hour of demonstration data in total.

I. INTRODUCTION

Dogs have remarkable abilities to interpret human behavior through social cues and seamlessly perform given tasks. For instance, they can follow a pointing gesture, track a handler’s gaze, or respond to brief verbal commands with precise actions. In contrast, current robot operation still primarily relies on manual control interfaces such as joysticks or keyboards. These methods require continuous attention and fine-grained control, limiting their practicality for everyday use by non-experts. Motivated by this, we aim to develop a system that maps natural human cues to corresponding robot behaviors.

Building such a system involves two core challenges. First, single-modal commands, such as verbal or gestural instructions alone, introduce ambiguity. For instance, verbal commands such as “go there” are inherently ambiguous unless the system correctly interprets the user’s pointing gesture that specifies the target location. Second, collecting large-scale real-world human–robot interaction data to learn

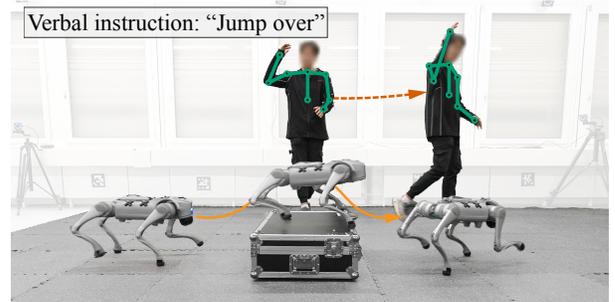


Fig. 1: The proposed framework enables agile robot behaviors through verbal and gesture-based commands.

such multimodal associations is costly. Therefore, it is crucial that the system can learn new interaction patterns from only a few interactions.

Prior work on natural robot interfaces has explored gesture-based control or language-conditioned navigation, but often relies on a single modality and substantial data collection. In this paper, we introduce a human-in-the-loop framework for teaching and controlling quadruped robots through multimodal interactions, termed **LURE** (**Luring-based User–Robot Engagement**), which enables data-efficient learning from few demonstrations. Inspired by the *luring* technique used by professional dog trainers, users first physically guide the robot with a teaching rod while issuing gestures and verbal commands. We then reconstruct the interaction scenes in a physics-based simulator, where the robot learns navigation skills that overcome distributional shifts arising from limited data, i.e., mismatches between the training and deployment distributions [1]. During training, we adaptively determine intermediate commands and navigation goals to improve context-aligned navigation, a process we refer to as *progressive goal cueing*. At run time, the learned policy enables intuitive control via human gestures and verbal commands, allowing the robot to perform agile movements in cluttered environments.

We demonstrate that the proposed framework enables users to control quadruped robots to overcome obstacles using verbal and gesture commands. For instance, in the *Jump over* scenario shown in Fig. 1, the robot correctly interprets the user’s intention to jump over the box and executes the desired action. Extensive experiments show that our system enables navigation through cluttered environments in a data-efficient manner with less than 1 hour of data across six scenarios.

Corresponding author: Sungjoon Choi.

¹Taerim Yoon, and Sungjoon Choi are with the Department of Artificial Intelligence, Korea University, 145 Anam-ro, Seongbuk-gu, Seoul, Korea (email: taerimyon@korea.ac.kr; sungjoon-choi@korea.ac.kr).

²Dongho Kang, Jin Cheng, Fatemeh Zargarbashi, Yijiang Huang, and Stelian Coros are with the Department of Computer Science, ETH Zurich, Wasserwerkstrasse 12, 8092 Zurich, Switzerland (email: kangd@ethz.ch, jcheng@ethz.ch, fzargarbashi@ethz.ch, huang@hey.com, scoros@ethz.ch).

³Minsung Ahn is with the Department of Mechanical and Aerospace Engineering, UCLA, 420 Westwood Plaza, Los Angeles, 90095, CA, USA (email: aaminsung@ucla.edu)

To summarize the key contributions of this work:

- 1) We present a system that enables users to teach and control quadruped robots in cluttered environments through real-time gestural and verbal commands.
- 2) We propose a data-aggregation technique that enables robots to learn new interaction patterns using only a few demonstrations.
- 3) We experimentally show that verbal and nonverbal cues complement each other in interpreting user intent.

II. RELATED WORK

Our goal is to enable users to teach and control a legged robot through natural interactions. This objective aligns with the human-in-the-loop control paradigm, as the robot is operated through human feedback. Our approach also relates to learning from demonstrations, as we leverage interactive human demonstrations to acquire navigation skills. In the following sections, we discuss related work on both aspects.

A. Human-in-the-Loop Robot Interfaces

Recent efforts in robot control have increasingly focused on designing more intuitive interfaces that move beyond manual devices such as joysticks or keypads. In particular, controlling robots through human gestures and verbal commands has been a long-standing objective in human-robot interaction [2]. While numerous studies have investigated the utility of these modalities, they have primarily focused on using either gestures or verbal commands in isolation.

A number of studies proposed a straightforward gesture-based human-robot interaction by manually mapping a pre-defined set of human gestures to a discrete number of robot actions. For instance, Shin et al. [3] and Zafar et al. [4] map current body keypoints directly to discretized robot actions. While these approaches are simple, discrete actions lack expressiveness, motivating continuous gesture-to-action interfaces. In this direction, Xie et al. [5] transferred hand motions at the velocity level for continuous robot control, while Cuan et al. [6] mapped user gestures to robot actions using imitation learning. However, directly mapping hand motions to low-level robot actions requires continuous, exhaustive user attention, similar to joystick-based teleoperation.

Verbal command-based control interfaces have been explored, particularly in the context of language-conditioned navigation. Cheng et al. [7] proposed a vision-language-action model for legged navigation, enabling text-driven goal specification. More recently, Xue et al. [8] embedded latent vision-language representations into a humanoid whole-body controller, enabling visual navigation from language instructions. Despite these advances, many language-based systems still suffer from spatial ambiguities. For instance, interpreting a vague command like “go there” can be challenging, as the destination is not explicitly specified in language but must be resolved through supplementary cues such as gaze or pointing.

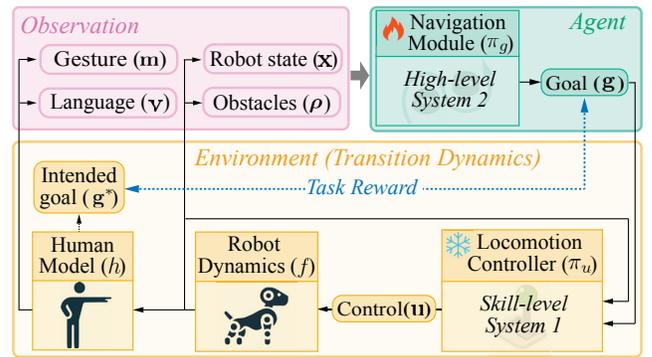


Fig. 2: The problem is formulated as an MDP with observations $(\mathbf{x}, \rho, \mathbf{c})$ and actions representing the navigation goal \mathbf{g} . The transition dynamics include the pretrained locomotion controller π_u , the robot dynamics f , and the human model h , while the reward is defined as the alignment between the predicted and intended goals.

To address these limitations, we complement the strength of verbal commands and gestures within a unified real-time interface. By grounding language with accompanying user gestures, our system resolves spatial ambiguities and provides clearer intent signals for navigation.

B. Learning Navigation from Demonstrations

Learning from Demonstration (LfD) enables users to teach behaviors through examples, providing an intuitive way to teach by direct demonstration rather than relying on handcrafted rules [9, 10]. In line with this, we adopt LfD as a natural framework for learning new behaviors from interactive and physically guided user demonstrations.

Early applications of LfD to navigation focused on single-goal imitation, with policies trained to follow expert-provided trajectories. A pivotal work by Pomerleau [11] trained neural networks to imitate human driving behavior for lane following, while more recent approaches [12] guide path imitation using environment feature maps.

Recent work advances toward multi-skill navigation that interprets high-level commands and produces context-appropriate behaviors. For example, Wu et al. [13] proposed a generative imitation learning framework for target-driven visual navigation in indoor scenes, enabling policies to generalize to novel goals based on demonstration data. Singh et al. [14] proposed a goal-conditioned navigation framework that grounds implicit free-form language expressions into spatial targets, enabling mobile agents to interpret diverse instructions such as “go to the bed”. Despite these advances, LfD is prone to distributional shift due to limited demonstration coverage, which can degrade performance or lead to task failure.

As in prior work, we study demonstration-based navigation, where robots navigate using verbal instructions and gestures. In particular, we focus on mitigating distributional shifts under limited demonstrations via data augmentation.

III. PROBLEM FORMULATION

Let the robot state be \mathbf{x} , and the control input be \mathbf{u} , where \mathbf{x} consists of the robot’s base position and orientation, joint angles, and their corresponding velocities, and \mathbf{u} denotes the motor torque commands. We assume that the robot dynamics are defined by a function f , such that $\mathbf{x}' = f(\mathbf{x}, \mathbf{u})$, where \mathbf{x}' denotes the state at the next time step. The robot operates in a cluttered environment, where the surrounding environment is represented by a heightmap ρ . Our objective is to enable a human user to control the robot to perform a navigation task \mathcal{T} via interaction commands \mathbf{c} , consisting of a gesture command \mathbf{m} and a verbal command \mathbf{v} . Specifically, we model the human as a function h that produces an intended navigation goal \mathbf{g}^* together with interaction commands \mathbf{c} , expressed as $(\mathbf{c}, \mathbf{g}^*) = h(\mathbf{x}, \rho, \mathcal{T})$. Given the interaction commands \mathbf{c} , the robot interprets the user’s intent and moves toward the intended goal \mathbf{g}^* . Formally, we develop a policy π that generates control inputs $\mathbf{u} = \pi(\mathbf{x}, \mathbf{c}, \rho)$, enabling the robot to infer human intent and accomplish the navigation task \mathcal{T} by reaching the intended goal \mathbf{g}^* .

As shown in Fig. 2, we decompose the problem into two subsystems: a *high-level navigation module* and a *low-level locomotion controller*. The navigation module, denoted as π_g , serves as a high-level System 2 planner that interprets human intent and predicts navigation goals \mathbf{g} , represented as $\mathbf{g} = \pi_g(\mathbf{c}, \mathbf{x}, \rho)$. The locomotion controller, represented as π_u , functions as a fast and reactive System 1 module that executes agile behaviors to overcome obstacles. By integrating these two modules, the robot is controlled directly by high-level interaction commands.

We pretrain a locomotion controller π_u capable of tracking arbitrary goals across diverse terrains and obstacles, and subsequently focus on the high-level navigation problem. We formulate this problem as a Markov Decision Process (MDP) problem, where the observation is defined as $(\mathbf{x}, \rho, \mathbf{c})$ and the action corresponds to the navigation goal \mathbf{g} expressed in the robot frame. In more detail, the navigation policy π_g outputs navigation goal \mathbf{g} based on the interaction signal \mathbf{c} , where the reward function is defined by how closely \mathbf{g} matches the intended navigation goal \mathbf{g}^* . The locomotion controller π_u then takes the navigation goal \mathbf{g} to produce a control signal $\mathbf{u} = \pi_u(\mathbf{g}, \mathbf{x}, \rho)$. Subsequently, the robot dynamics f produce the robot state \mathbf{x}' at the next time step, and the human model h generates the corresponding gesture \mathbf{m}' and language command \mathbf{v}' , completing the state transition.

We address this MDP using data aggregation [1] to mitigate the distributional shift problem, which is augmenting expert actions based on the current policy’s state distribution, denoted as \mathcal{D}^{π_g} . Formally, the problem is formulated as

$$\arg \min_{\pi_g \in \Pi} \mathbb{E}_{(\mathbf{c}, \mathbf{x}, \rho, \mathbf{g}^*) \sim \mathcal{D}^{\pi_g}} \left[\|\mathbf{g}^* - \pi_g(\mathbf{c}, \mathbf{x}, \rho)\|^2 \right]. \quad (1)$$

Addressing distributional shifts is particularly important, as the high cost of collecting interaction data leads to limited demonstrations, making the system especially prone to them.

Data aggregation requires two key components. First, it requires the transition dynamics, including the human model

h . Since manually encoding the intended navigation goal \mathbf{g} and the corresponding interaction commands \mathbf{c} is impractical, we use demonstration data to approximate the human model as \tilde{h} , as discussed in Section IV-D. Second, data aggregation requires an expert policy to provide actions at previously unseen states. We construct a locally approximated expert $\tilde{\pi}_g$ that produces actions for novel states and enables efficient data aggregation, as described in Section IV-C.

IV. METHOD

Our goal is to develop a policy that can interpret the verbal command \mathbf{v} and gesture command \mathbf{m} to reach the navigation goal \mathbf{g}^* . To achieve this, our framework comprises several key components, as illustrated in Fig. 3. We begin by introducing an agile locomotion controller that enables the robot to perform dynamic movements such as stepping over or around obstacles. Next, we collect interaction data using this locomotion controller. Finally, we reconstruct interaction scenes in which robots are trained via data augmentation, enabling control through natural gestural and verbal commands with limited demonstration data.

A. Locomotion controller

As a preliminary stage, we develop a locomotion controller for low-level agility of the robot. As shown in the Stage 0 from Fig. 3, the locomotion controller consists of a velocity planner and a velocity tracker. The velocity planner converts the navigation goal \mathbf{g} into a velocity command guiding the robot toward its target. In more detail, it outputs a heading velocity command $v_{\text{com}} = K \|\mathbf{g}^{xy} - \mathbf{p}^{xy}\|_2$ and a rotation command $\phi_{\text{com}} = \text{yaw}(R^\top(\mathbf{g} - \mathbf{p}))$, where $\mathbf{p} \in \mathbb{R}^3$ is the robot’s base position, $R \in \text{SO}(3)$ is the robot’s rotation matrix and $(\cdot)^{xy}$ indicates ground-projected coordinates.

The other component is the velocity tracker, which enables the robot to track the commanded velocity. Although our framework is compatible with any velocity tracker, we trained a specialized tracker designed for dynamic locomotion to support agile behaviors, such as obstacle jumping. More details on our velocity tracking policy are discussed in Section IV-F.

B. Human-in-the-loop data collection through luring

We collect human-robot interaction data by luring, as illustrated in Stage 1 of Fig. 3. The process requires two participants: one issues both verbal and non-verbal commands, while the other guides the robot using a teaching rod. This setup allows us to collect paired interaction and navigation data, which we refer to as the interaction data \mathcal{D} . Details on the representation of verbal commands and gestural cues are provided in Section IV-F.

As shown in Fig. 3, the interaction data \mathcal{D} consists of the robot state \mathbf{x} , the target position \mathbf{g}^* , the gesture command \mathbf{m} , the verbal command \mathbf{v} , and the obstacle representation ρ . All data is recorded in a global coordinate frame, providing sufficient information for full scene reconstruction. This scene reconstruction is a crucial component of our framework, as it enables data aggregation that improves robustness to distributional shifts, as discussed in Section IV-C.

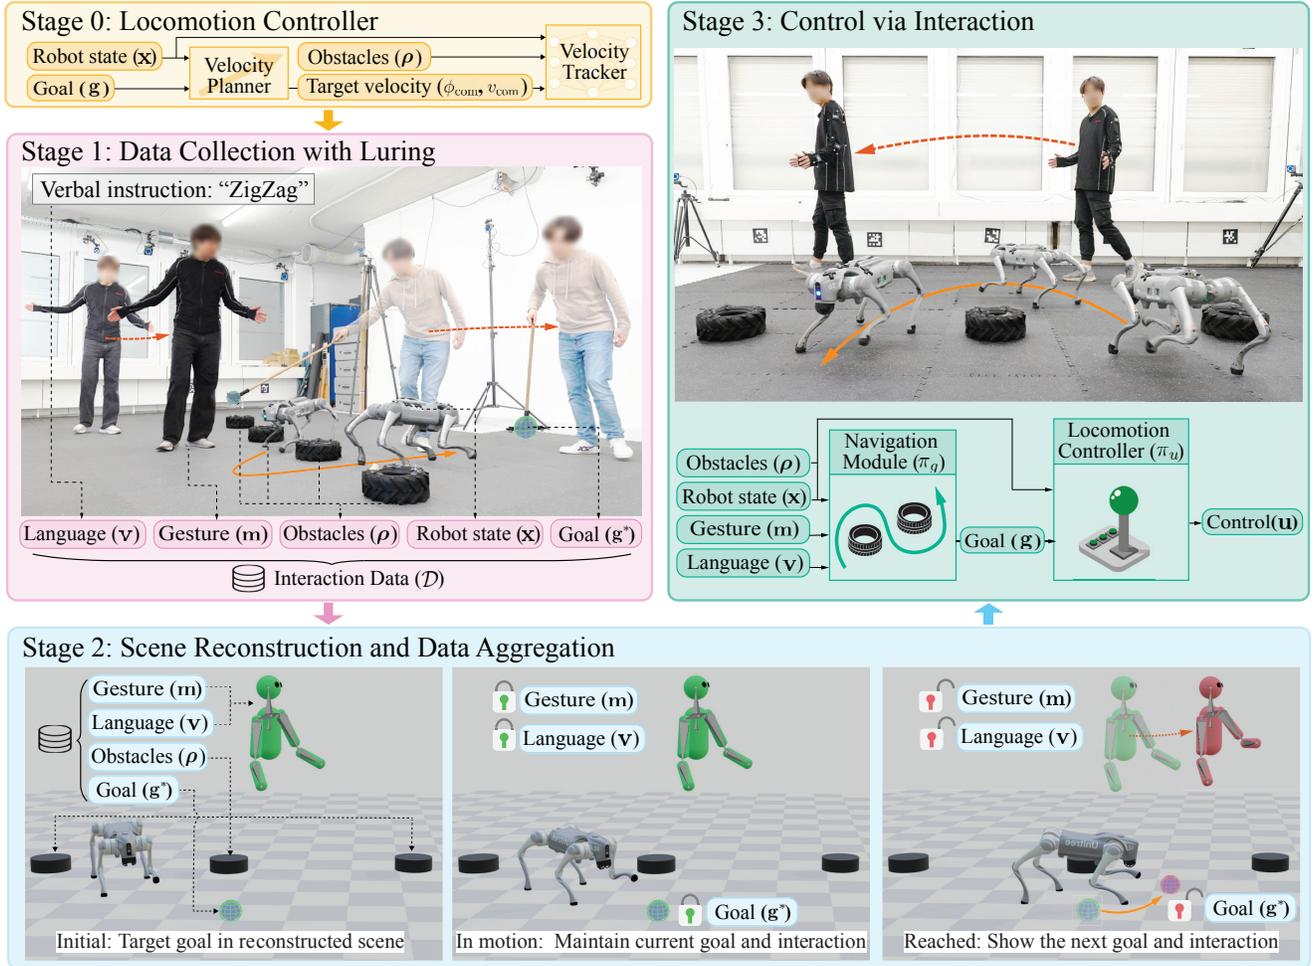


Fig. 3: Overview of the proposed framework. We first establish a locomotion controller that follows the navigation goal g . In Stage 1, we collect interaction data $\mathcal{D} = \{\mathbf{v}, \mathbf{m}, \rho, \mathbf{x}, \mathbf{g}^*\}$ through natural interactions between two participants. In Stage 2, we reconstruct the interaction scene from \mathcal{D} and train the navigation module via data aggregation. The framework progressively provides the interaction and navigation goals only after the robot reaches its current goal, ensuring alignment between the command and behavior. Finally, in Stage 3, the user can control the robot through interaction.

C. Data aggregation

We reconstruct the interaction scene in the simulation for data aggregation, as shown in Stage 2 of Fig. 3. Specifically, we recreate the terrain using obstacle information ρ and replay the gesture \mathbf{m} and verbal commands \mathbf{v} . The robot is trained to follow the intended goal location \mathbf{g}^* , defined by the teaching rod position in the interaction data.

Overall, our training procedure resembles the data aggregation method introduced in DAgger [1], where the agent explores using its current policy while the expert’s actions are augmented to guide learning. In more detail, we define the local expert policy $\tilde{\pi}_g$ that outputs the expert navigation goal \tilde{g} in the robot frame, directing the robot towards the global target position \mathbf{g}^* . By aggregating these actions, the system can efficiently generate corrective actions in the local frame that steer the robot toward the global target \mathbf{g}^* , even when the current policy drifts away from the expert trajectory. As a result, the robot learns to recover from deviations and reach

the intended target, effectively mitigating the distributional shift. This approach is justified by our data-collection setup, where the velocity planner drives the robot toward the goal in the robot frame, producing demonstrations consistent with the behavior encoded in the local expert policy $\tilde{\pi}_g$.

We apply domain randomization to deliberately force the agent to deviate from the expert trajectory, exposing it to a broader range of states and enabling it to learn recovery behaviors [15]. For example, we vary the robot’s mass and inertia, and apply external pushes to the base. To further improve robustness, we introduce the following techniques, with detailed parameter settings provided in the appendix.

- **Terrain scaling** We divide the terrain into a uniform grid of S tiles and scale each tile by a random factor to randomize the spacing between obstacles and goals.
- **Binary obstacle map** We threshold the height map at h_{thres} and feed the resulting binary obstacle map to the navigation module, thereby promoting consistent navigation across varying obstacle heights.

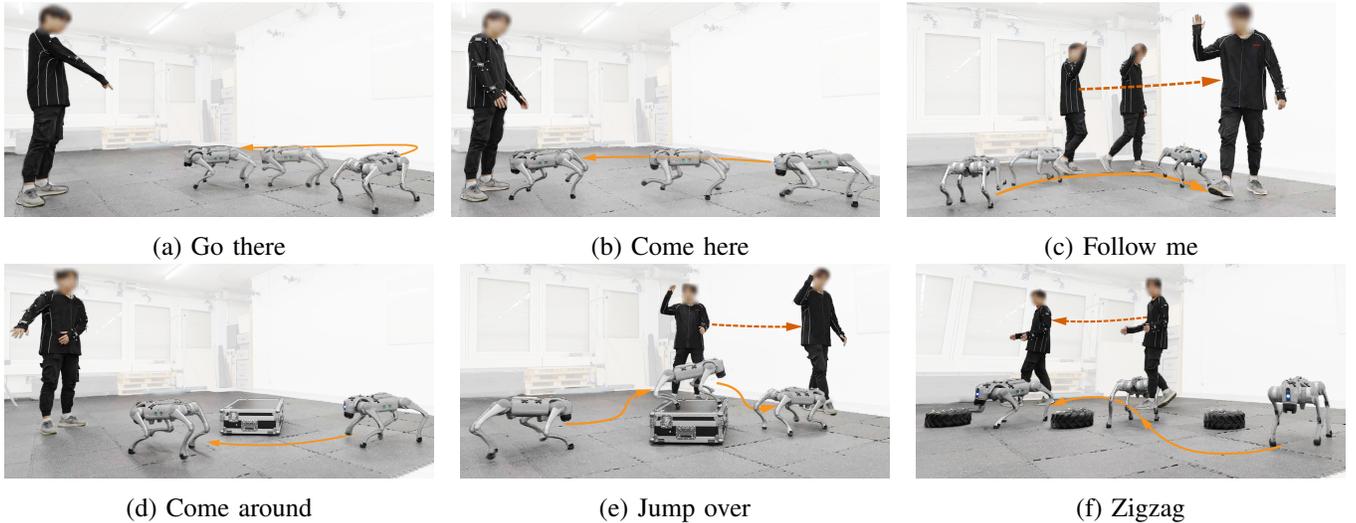


Fig. 4: Illustration of data collection procedures for six interactive navigation scenarios. (a) *Go there*, (b) *Come here*, and (c) *Follow me* represent human-robot interactions in open space. (d) *Come around* and (e) *Jump over* illustrate interactions involving a box obstacle. (f) *Zigzag* demonstrates interaction with multiple tire obstacles.

- **Distractor objects.** We introduce square and circular distractor objects to mitigate overfitting to specific obstacle arrangements.

D. Progressive goal cueing

It is crucial to provide interaction commands and navigation goals that align with the robot’s current state during training, as misalignment can lead to inconsistent supervision. One naive approach is to feed gesture and verbal commands to the policy according to their original timestamps. However, during early training, the policy often progresses more slowly through tasks, leading to new commands being issued before the robot has completed the preceding behavior. Since the interaction data were collected by issuing new commands in response to the robot’s task progress, this naive replay strategy can introduce state–command misalignment.

To address this, we propose a *progressive goal cueing* that dynamically presents the interaction command based on the robot’s state. As illustrated in Stage 2 of Fig. 3, each interaction command is held constant while the robot moves toward the current goal and is updated once the target is reached. This allows the human model h to wait for the robot rather than rushing through the episode when the robot struggles to navigate correctly in the early stages. Overall, it results in more context-aware control and enables finer movements during execution, as demonstrated later in Section V-C.

At each timestep, we apply progressive goal cueing with a 50% probability; otherwise, it advances to the next interaction according to the simulation clock. This design choice prevents excessive waiting and improves responsiveness during real-world deployment. Additionally, the waiting behavior naturally fades out as the robot begins producing correct movements in the latter stage of training.

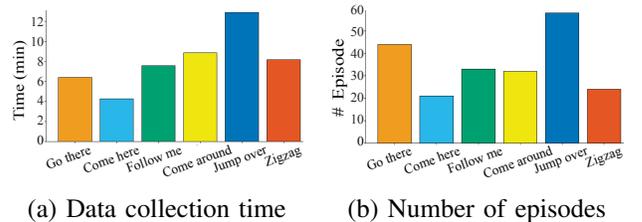


Fig. 5: Analysis of collected data in terms of (a) total wall clock time and (b) number of episodes for each scenario

E. Interaction Data Augmentation

Since our model is trained from limited interaction data, preventing overfitting and promoting generalization across both gesture and verbal modalities are essential. Therefore, we apply modality-specific input augmentation for both gestures and verbal commands. For gestures, we perturb the extracted keypoints with noise ϵ_m for every simulation time step. This models the natural variability in repeated human motions as well as noise introduced during pose estimation. For verbal commands, we leverage a Large Language Model (LLM) [16] to generate semantically equivalent expressions. Specifically, for each verbal command collected in Section V-A, we prompt the LLM to produce N_v alternative phrasings, from which one is randomly sampled during training.

F. Implementation details

This section provides additional implementation details regarding the perception and command representations, as well as the velocity tracker policy for low-level locomotion control.

1) *Representation:* We describe how each input modality is represented and captured in real time. The verbal command is segmented using Silero VAD [17], transcribed using Whisper [18], and encoded with a pretrained text

encoder [19]. Next, the human gesture is captured using motion capture cameras (OptiTrack Prime X22). We consider six upper-body keypoints—the shoulders, elbows, and wrists on both the left and right sides—and express their 3D positions in the robot frame as $\mathbf{p}_{\text{sh},r}, \mathbf{p}_{\text{el},r}, \mathbf{p}_{\text{wr},r}, \mathbf{p}_{\text{sh},l}, \mathbf{p}_{\text{el},l}, \mathbf{p}_{\text{wr},l}$. Using these keypoints, we define the gesture command as $\mathbf{m} = [\mathbf{d}_{\text{sh},r,\text{el},r}, \mathbf{d}_{\text{el},r,\text{wr},r}, \mathbf{d}_{\text{sh},l,\text{el},l}, \mathbf{d}_{\text{el},l,\text{wr},l}, \mathbf{d}_{\text{sh},l,\text{sh},r}, (\mathbf{p}_{\text{sh},l} + \mathbf{p}_{\text{sh},r})^{xy}/2, \phi_h]$, where \mathbf{d} denotes the unit vector between two points, and ϕ_h represents the human yaw. Regarding the obstacles, we measure their dimensions in advance and estimate their poses using motion capture cameras. Using this information, we obtain a heightmap in the robot frame for the locomotion controller π_u and a thresholded binary mask for the navigation module π_g .

2) *Velocity tracker*: We train a velocity tracker operating at 50 Hz to overcome the obstacles. The training scheme is similar to the work of Cheng et al. [20] and is trained in Isaac Gym [21] to track velocity commands $v_{\text{com}} \in [0.0, 1.0]$ m/s and angular velocity commands $w_{\text{com}} \in [-\pi/3, \pi/3]$ rad/s. To improve robustness, we introduce a stand boolean in the state representation by thresholding velocity commands v_{com} , which enhances walk–stand transitions. For real-world deployment, we use a Kalman filter–based state estimator [22] that fuses joint encoder, foot contact, and motion capture data for stable feedback. Although the robot can be trained to move faster, we limit its speed for safety.

V. RESULTS

We conduct a series of experiments to assess the effectiveness and generality of our framework. We begin by describing the interaction scenarios considered in this study and then compare our method with imitation learning baselines to evaluate our proposed data aggregation and progressive goal cueing strategy. We then show that our system can quickly learn new interaction patterns from novel users with less than five minutes of interaction. Next, we conduct an ablation study to demonstrate how verbal and gestural cues complement each other. Finally, we demonstrate multi-obstacle navigation in a real-world environment, showing that our system can robustly execute complex task sequences using a single unified policy.

A. Data Collection

We collect interaction data, denoted as \mathcal{D} , comprising human interaction \mathbf{c} , robot states \mathbf{x} , obstacle representations ρ , and target goals \mathbf{g}^* , recorded at a frequency of 10 Hz. We define six interaction scenarios, as shown in Fig. 4, and a stop scenario.

- **Stop**: A person raises both arms and says “stop,” prompting the robot to halt at its current position.
- **Go there**: A person points with one hand and says “go there,” and the robot moves toward the indicated location.
- **Come here**: A person calls the robot by saying “come here” with the arms dropped, and the robot moves directly toward the person.

- **Follow me**: A person raises either their left or right hand and says “follow me,” and the robot follows on the corresponding side of the raised hand.
- **Come around**: A person stands in front of a box and says “come around,” prompting the robot to move toward the human while avoiding the obstacle. Both hands move towards the direction intended for coming around the obstacle (left or right).
- **Jump over**: A person waves one hand over their head while saying “jump over,” causing the robot to jump over the obstacle.
- **Zigzag**: A person gestures with one hand to indicate the movement direction while saying “zigzag,” and the robot navigates between the tires in a zigzag pattern when viewed from above.

The summary of total data collection time and the number of episodes per scenario is detailed in Fig. 5. While we introduce these scenarios as representative cases, the proposed framework can support a much broader range of interactions *without relying on predefined cases*.

B. Baselines

We compare our proposed method to three imitation learning algorithms.

1) *BC [11]*: As a most basic baseline, we train the navigation policy using standard Behavior Cloning (BC). The policy is learned purely from expert demonstrations.

2) *GAIL [23]*: We use Generative Adversarial Imitation Learning (GAIL) as a baseline to compare against inverse reinforcement learning methods that are known for their generalization capabilities. The policy is learned via adversarial training to match the expert behavior distribution. This baseline allows us to compare our approach against an imitation framework that focuses on generalization.

3) *DAGger [1]*: We train the policy using data aggregation with a local expert π_g consistent with our proposed framework. However, we do not apply progressive goal cueing. Instead, interaction commands are issued according to the simulation time clock. We reconstruct the scene and apply domain randomization to induce a distributional shift, from which the policy learns robust recovery, as described in IV-C.

4) *LURE (Ours)*: Our proposed method, LURE, utilizes data aggregation and domain randomization as in DAGger. In addition, we apply progressive goal cueing.

C. Evaluating Success Rate and Navigation Error

We measure both the robot’s success rate and navigation error, as shown in Fig. 6. We first segment each interaction scene based on the issued verbal command “Stop,” treating each segment as an independent episode. For example, the period between saying “Go there” and issuing “Stop” is regarded as one episode. At the beginning of each episode, we reset the robot’s position to the ground-truth position from data to ensure consistent evaluation. During evaluation, we disable progressive goal cueing to avoid underestimating the performance of other baselines. For each interaction

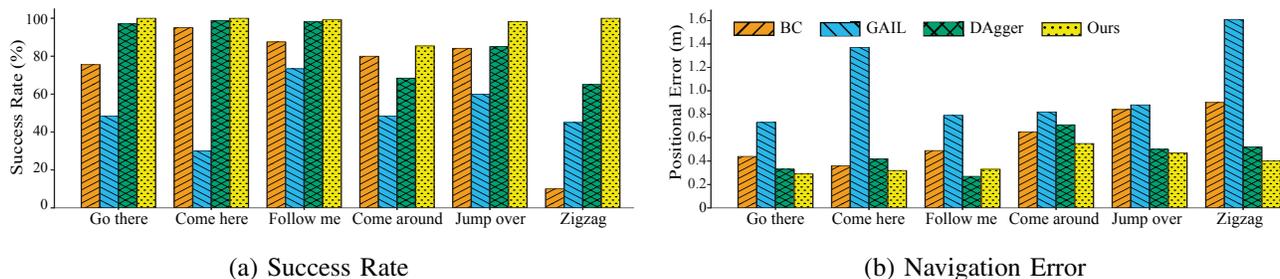


Fig. 6: Comparison of baseline methods across six interactive navigation scenarios. (a) Navigation error (m) of BC, GAIL, DAgger, and our proposed method. (b) Success rate (%) for the same scenarios.

TABLE I: Success Rate and Navigation Error per Scenario, Subject, and Adaptation

Scenario	Subject #1				Subject #2				Subject #3			
	Success Rate \uparrow		Navigation Error \downarrow		Success Rate \uparrow		Navigation Error \downarrow		Success Rate \uparrow		Navigation Error \downarrow	
	w/o Adpt.	Adpt.	w/o Adpt.	Adpt.	w/o Adpt.	Adpt.	w/o Adpt.	Adpt.	w/o Adpt.	Adpt.	w/o Adpt.	Adpt.
Go there	0.66	0.87	0.71	0.53	0.86	0.99	0.57	0.51	0.68	1.00	0.72	0.54
Come here	0.90	1.00	0.39	0.28	1.00	1.00	0.38	0.18	1.00	1.00	0.26	0.22
Follow me	0.88	1.00	0.68	0.36	0.81	0.90	0.50	0.30	1.00	1.00	0.39	0.23
Come around	0.37	0.95	1.13	0.41	0.95	0.95	0.65	0.43	0.10	1.00	2.20	0.89
Jump over	0.87	0.95	1.08	0.53	0.73	0.95	0.70	0.55	0.48	0.92	1.11	0.45
Zigzag	0.76	0.95	0.99	0.36	0.50	0.70	1.45	0.73	0.60	1.00	1.04	0.36

episode, we evaluate performance on five terrain instances with different scale factors and 20 robots per terrain, resulting in 100 evaluation runs. The reported success rates and navigation errors are then averaged over each scenario.

a) *Task Success*: We define task-specific success criteria as follows:

- **Go there, Come here, Follow me**: The final position of the robot must be within 1 meter of the ground-truth base position.
- **Come around**: The robot must pass through the correct lateral corridor beside the box, following the direction indicated by the user.
- **Jump over**: The robot must pass through three reference points located along the forward direction of the box, centered at -0.5 m, 0 m, and $+0.5$ m relative to the box. A pass is counted as successful if the robot comes within 0.2 m of each point.
- **Zigzag**: The robot must pass through all midpoints between the tires while avoiding stepping on any of them. A pass is counted as successful if the robot comes within 0.2 m of each point.

b) *Navigation Error*: We compute the navigation error as the Mean Squared Error (MSE) between the robot’s base position and a ground-truth trajectory, measured at each timestep during the interaction. The ground-truth trajectory corresponds to the robot motion recorded in the interaction data under the same interaction commands. The final error is obtained by averaging the per-frame MSE over the entire interaction duration. This metric reflects how accurately the policy follows the intended trajectory specified by the interaction commands.

The results summarized in Fig. 6 show that LURE

TABLE II: Success Rate (%) and Navigation Error (m) per Scenario and Modality

Scenario	Succ. Rate \uparrow			Pos. Error (m) \downarrow		
	Both	Verb.	Gest.	Both	Verb.	Gest.
Go there	0.85	0.35	0.15	0.52	1.15	1.79
Come here	0.91	0.01	0.15	0.40	1.74	1.09
Follow me	0.98	0.16	0.28	0.29	1.98	1.24
Come around	0.90	0.02	0.22	0.54	2.26	1.85
Jump over	0.83	0.67	0.23	0.49	1.03	1.64
Zigzag	0.96	0.00	0.16	0.35	1.93	1.26

achieves the best performance, with a clear margin, attaining an average success rate of 97.15%. By comparing DAgger with BC, we observe that incorporating data aggregation reduces navigation error by 24.6%. Introducing progressive goal cueing, on top of this, yields an additional 15.2% reduction in navigation error over the DAgger baseline. In terms of success rate, data aggregation improves performance by 18.6%, and progressive goal cueing provides a further 13.7% gain. Notably, GAIL outperforms BC on the most challenging Zigzag task, but its overall performance degrades when all skills are considered together. On average, GAIL achieves a 29.32% lower success rate than BC, which may be explained by training instability and mode collapse often observed in adversarial imitation learning.

D. Adaptation to novel user

We evaluate whether our system can generalize to novel users through quick adaptation. This is essential because different users often employ different interaction cues to command the robot for the same task. We first train a

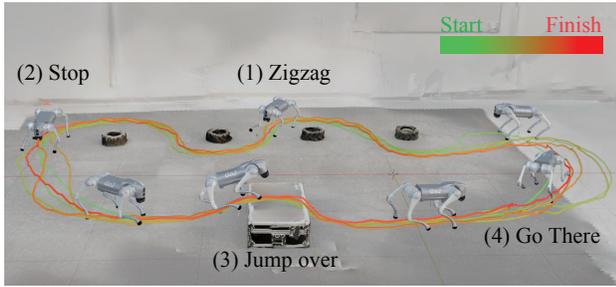


Fig. 7: The robot (1) zigzags through the tires, (2) briefly stops, (3) jumps over the box, and (4) returns to the starting point according to the real-time user’s instruction. The visualization overlays the robot’s trajectory on a reconstructed 3D scene of the real environment. The figure shows the robot completing these sequences five consecutive times.

model using the initial interaction dataset from a single user described in Section V-A, and then fine-tune it using an average of 4.5 minutes of data collected from each new subject. For evaluation, we measure the success rate and navigation error in the reconstructed simulation environment, following the same protocol described in Section V-C.

As shown in Table I, the success rates before adaptation are 74.00%, 80.83%, and 64.33% for three subjects. This indicates that the model works to some degree even without personalization. After adaptation, the success rates increase to 95.33%, 91.50%, and 98.67%, resulting in an average improvement of 22.11%. These results demonstrate that our framework can quickly adapt to new users with only a small amount of data.

E. Ablation on Gesture and Verbal Modalities

We use both gesture and verbal modalities for interaction cues. To evaluate their individual contributions, we train a model in which each modality is masked with a 10% probability during training. This enables the model to function even when one modality is absent, enabling meaningful evaluation under partial inputs.

As shown in Table II, the evaluation is conducted across all six scenarios to measure how each modality contributes to task success and navigation error. Using gesture only results in an average success rate of 19.83%, which is 70.67% lower than when both modalities are used. In practice, the robot also tends to move only when a gesture is held for an extended period. These indicate that gesture alone is often insufficient for effectively interpreting the user’s intent. On the other hand, using verbal commands alone results in a navigation error of 1.6817 m, whereas combining both modalities reduces the error by 1.25 m to 0.4317 m. This is because verbal commands can lack spatial grounding. In the *Go there* scenario, for instance, the robot does not know where the user is pointing and begins to move randomly. One exception is the *Jump over* scenario, where the robot can locate the box without a gesture, yielding a relatively high success rate of 67%. However, this is still lower than the multimodal setting, which achieves 83%.

F. Multi-Obstacle Navigation

To demonstrate that our framework can robustly handle diverse obstacles while seamlessly switching between multiple locomotion skills based on user intent, we conducted an obstacle course experiment, as illustrated in Fig. 7. The course contains both boxes and tires: the robot first navigates through the tire corridor, avoiding each tire according to the user’s interaction commands. After clearing the tires, the user issues a “stop” command, upon which the robot halts and then performs the “Jump over” command to overcome the box. Finally, in response to the “Go there” command, the robot returns to the starting point. As shown in Fig. 7, our framework completes this entire sequence five times in a row, demonstrating strong real-world robustness and smooth task transitions under a single unified policy. The robot’s trajectory is obtained using a motion capture system, while the obstacle geometries are reconstructed using a consumer-grade mobile 3D scanning software.

VI. LIMITATIONS

While our framework enables effective interaction-based learning and control, it currently lacks zero-shot generalization across different users due to substantial inter-subject variability in gesture styles. Unlike verbal commands, which benefit from semantic augmentation via Large Language Models (LLMs), our gesture representation is augmented only through noise injection and therefore lacks semantic diversity. A promising direction is to develop gesture representations that are semantically structured and jointly grounded in the robot state and surrounding objects, building upon pretrained motion representations such as Zhang et al. [24].

Moreover, our progressive goal cueing provides a simple and efficient way to approximate human behavior in novel states. However, it only adjusts the timing of replaying human interaction data based on the robot’s state. Ideally, human behavior should be generated more dynamically in response to novel states, rather than being constrained to temporally re-aligned demonstrations.

Finally, our current vision system is restricted to controlled environments, as the objects are modeled as circular or square primitives with pre-measured dimensions, and the human wears a motion capture suit for pose estimation. Camera-based human pose estimation and scene reconstruction [25, 26] could potentially remove these assumptions and enable operation in the wild.

VII. CONCLUSION AND FUTURE WORK

We presented a framework that enables robots to learn and execute navigation skills in real time through natural human interaction. By reconstructing interaction scenes and augmenting data within these environments, our approach enables policies to robustly recover from distributional shifts across diverse interaction conditions. Experimental results demonstrate that the proposed framework effectively integrates verbal and gestural cues, enabling robust control across a range of tasks and environments. Moreover, our

findings highlight the complementary nature of speech and gesture, showing that their combination provides a powerful interface for teaching and controlling robots to execute context-appropriate behaviors, including agile maneuvers such as jumping over obstacles. Overall, this work advances an intuitive and practical system that directly maps high-level human intent to executable robot actions. In future work, we will explore extending our framework beyond navigation to more general whole-body behaviors, such as locomotion, by grounding multimodal commands using a similar data augmentation strategy.

REFERENCES

- [1] S. Ross, G. Gordon, and D. 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*. JMLR Workshop and Conference Proceedings, 2011, pp. 627–635.
- [2] M. A. Goodrich and A. C. Schultz, "Human–robot interaction: A survey," *Foundations and Trends in Human–Computer Interaction*, vol. 1, no. 3, pp. 203–275, 2007.
- [3] S. Shin, T. Evetts, H. Saylor, H. Kim, S. Woo, W. Rhee, and S.-W. Kim, "Non-verbal interaction and interface with a quadruped robot using body and hand gestures: Design and user experience evaluation," in *European Conference on Computer Vision*. Springer, 2024, pp. 293–308.
- [4] M. H. Zafar, E. F. Langås, and F. Sanfilippo, "Real-time gesture-based control of a quadruped robot using a stacked convolutional bi-long short-term memory (bi-ilstm) neural network," in *2024 10th International Conference on Automation, Robotics and Applications (ICARA)*, 2024, pp. 81–86.
- [5] J. Xie, Z. Xu, J. Zeng, Y. Gao, and K. Hashimoto, "Human–robot interaction using dynamic hand gesture for teleoperation of quadruped robots with a robotic arm," *Electronics*, vol. 14, no. 5, p. 860, 2025.
- [6] C. Cuan, T.-W. E. Lee, E. Fisher, A. Francis, L. Takayama, T. Zhang, A. Toshev, and S. Pirk, "Gesture2path: Imitation learning for gesture-aware navigation," in *International Conference on Social Robotics*. Springer, 2024, pp. 264–279.
- [7] A.-C. Cheng, Y. Ji, Z. Yang, Z. Gongye, X. Zou, J. Kautz, E. Biyik, H. Yin, S. Liu, and X. Wang, "NaVILA: Legged Robot Vision-Language-Action Model for Navigation," in *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025.
- [8] H. Xue, X. Huang, D. Niu, Q. Liao, T. Kragerud, J. T. Gravdahl, X. B. Peng, G. Shi, T. Darrell, K. Sreenath *et al.*, "Leverb: Humanoid whole-body control with latent vision-language instruction," *arXiv preprint arXiv:2506.13751*, 2025.
- [9] A. L. Thomaz and C. Breazeal, "Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance," in *Proceedings of the 21st National Conference on Artificial Intelligence*, 2006, pp. 1000–1005.
- [10] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and Autonomous Systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [11] D. A. Pomerleau, "Alvinn: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, 1989, pp. 305–313.
- [12] M. Wigness, J. Rogers, and R. Hansen, "Robot navigation from human demonstration: Learning control behaviors with environment feature maps," *IEEE Transactions on Intelligent Vehicles*, vol. 7, no. 2, pp. 290–300, 2022.
- [13] Q. Wu, X. Gong, K. Xu, D. Manocha, J. Dong, and J. Wang, "Towards target-driven visual navigation in indoor scenes via generative imitation learning," *IEEE Robotics and Automation Letters*, vol. 6, no. 1, pp. 175–182, 2020.
- [14] R. Singh, S. Nair, A. Mandelkar, S. Savarese, Y. Zhu, and A. Garg, "Il-literate robots: Differentiable goal selection with semantic reasoning," in *Conference on Robot Learning*. PMLR, 2022, pp. 1686–1696.
- [15] D. Kang, J. Cheng, M. Zamora, F. Zargarbashi, and S. Coros, "RI + model-based control: Using on-demand optimal control to learn versatile legged locomotion," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6619–6626, 2023.
- [16] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, "Gpt-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.
- [17] H. Bredin and A. Laurent, "End-to-end speaker segmentation for overlap-aware resegmentation," in *Interspeech 2021*, 2021, pp. 3111–3115.
- [18] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [19] L. Wang, N. Yang, X. Huang, B. Jiao, L. Yang, D. Jiang, R. Majumder, and F. Wei, "Text embeddings by weakly-supervised contrastive pre-training," *arXiv preprint arXiv:2212.03533*, 2022.
- [20] X. Cheng, K. Shi, A. Agarwal, and D. Pathak, "Extreme parkour with legged robots," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 11 443–11 450.
- [21] V. Makovychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [22] M. Bloesch, M. Hutter, M. A. Hoepflinger, S. Leutenegger, C. Gehring, C. D. Remy, and R. Siegwart, "State estimation for legged robots-consistent fusion of leg kinematics and imu," *Robotics*, vol. 17, pp. 17–24, 2013.
- [23] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, 2016.
- [24] M. Zhang, Z. Cai, L. Pan, F. Hong, X. Guo, L. Yang, and Z. Liu, "Motiondiffuse: Text-driven human motion generation with diffusion model," *IEEE transactions on pattern analysis and machine intelligence*, vol. 46, no. 6, pp. 4115–4128, 2024.
- [25] S. Shin, J. Kim, E. Halilaj, and M. J. Black, "Wham: Reconstructing world-grounded humans with accurate 3d motion," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2070–2080.
- [26] M. Mihajlovic, S. Zhang, G. Li, K. Zhao, L. Muller, and S. Tang, "Volumetricmpl: A neural volumetric body model for efficient interactions, contacts, and collisions," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2025, pp. 5060–5070.

APPENDIX

TABLE III: Hyperparameters and data augmentation details

Category	Setting
Domain	External perturbation magnitude: $f \in [0, 0.5]$ m/s; applied at an average interval of 3 seconds.
Randomization	Terrain scaling factor: $s \in [0.75, 1.25]$ over 3×3 tiles. Binary height map threshold: 0.05 m.
Data	LLM-based verbal command augmentation: $N_v = 20$.
Augmentation	Gesture keypoint noise: $\epsilon_{\mathbf{m}} \sim \mathcal{N}(0, 0.01)$.