

EmoBipedNav: Emotion-aware Social Navigation for Bipedal Robots with Deep Reinforcement Learning

Wei Zhu¹, Abirath Raju¹, Abdulaziz Shamsah², Anqi Wu³, Seth Hutchinson⁴, and Ye Zhao¹

Abstract—This study presents an emotion-aware navigation framework – EmoBipedNav – using deep reinforcement learning (DRL) for bipedal robots walking in socially interactive environments. The inherent locomotion constraints of bipedal robots challenge their safe maneuvering capabilities in dynamic environments. When combined with the intricacies of social environments, including pedestrian interactions and social cues, such as emotions, these challenges become even more pronounced. To address these coupled problems, we propose a two-stage pipeline that considers both bipedal locomotion constraints and complex social environments. Specifically, social navigation scenarios are represented using sequential LiDAR grid maps (LGMs), from which we extract latent features, including collision regions, emotion-related discomfort zones, social interactions, and the spatio-temporal dynamics of evolving environments. The extracted features are directly mapped to the actions of reduced-order models (ROMs) through a DRL architecture. Furthermore, the proposed framework incorporates full-order dynamics and locomotion constraints during training, effectively accounting for tracking errors and restrictions of the locomotion controller while planning the trajectory with ROMs. Comprehensive experiments demonstrate that our approach exceeds both model-based planners and DRL-based baselines. The hardware videos and open-source code are available at <https://gatech-lidar.github.io/emobipednav.github.io/>.

Index Terms—Bipedal robot, locomotion, social navigation, emotion-aware navigation, deep reinforcement learning.

I. INTRODUCTION

THE research and development of bipedal robots have garnered significant interest within the robotics community, mainly due to their human-like morphology and versatile locomotion capabilities [1]–[6]. While a great number of studies have focused on developing stable locomotion controllers for bipedal robots, one ultimate goal is to enable bipedal robots to autonomously and courteously navigate pedestrian-populated environments within social contexts. However, achieving safe and socially aware navigation remains a formidable challenge, originating from the intricate dynamics of bipedal robots, limited maneuvering capabilities in dynamic environments, and the complexities of understanding the intentions and interactions of pedestrians.

Although significant efforts have been dedicated to the use of wheeled mobile robots for social navigation [7], [8], few

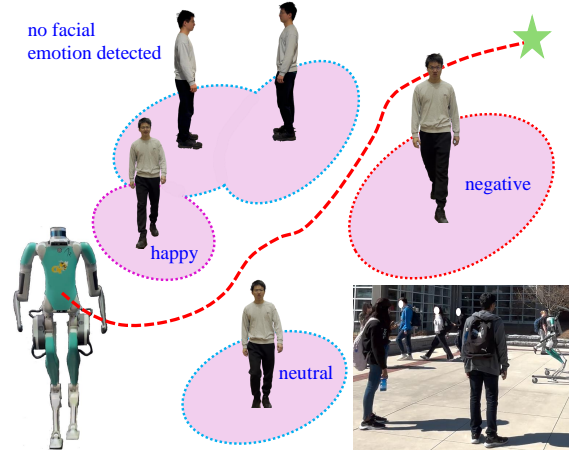


Fig. 1. Emotion-aware social navigation with the bipedal robot Digit. Digit is required to maintain customized comfort distances from pedestrians with specific emotions while navigating toward a designated goal location.

studies address bipedal social navigation due to challenges in developing stable and precise locomotion controllers for their hybrid, highly nonlinear, and high-degree-of-freedom dynamics. Furthermore, complex social interactions increase the level of difficulty because they are too implicit to be accurately modeled. In addition, social cues, such as human emotions, psychologically influence the preferred personal space during interactions with other agents, including pedestrians and robots [9], [10], as illustrated in Fig. 1. These social interactions and cues remain underexplored for bipedal robot navigation.

Cutting-edge bipedal navigation studies in crowded environments decouple navigation and locomotion, planning navigation trajectories with reduced-order models (ROMs) at first, and then tracking the trajectories with locomotion controllers based on full-order dynamics [11], [12]. Such a planning architecture cannot take into account the tracking errors and constraints of low-level controllers, such as joint torque limits, for bipedal navigation. In addition, pedestrian trajectory prediction is also individually executed [12] because prediction models are pretrained using specific datasets only including pedestrian-pedestrian interactions, ignoring pedestrian-robot interactions. Such pretrained models might further cause deployment discrepancies in bipedal robots.

In this study, we present an emotion-aware navigation framework – EmoBipedNav – using deep reinforcement learning (DRL) for bipedal robots walking in socially interactive environments. We represent the environment with an image-based observation of sequential LiDAR grid maps (LGMs), which include pedestrian-pedestrian and pedestrian-

¹The authors are with the Laboratory for Intelligent Decision and Autonomous Robots, Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30313, USA. {wzhu328, araju60, yezhao}@gatech.edu

²College of Engineering and Petroleum, Kuwait University, PO Box 5969, Safat, 13060, Kuwait. abdulaziz.shamsah@ku.edu.kw

³School of Computational Science and Engineering, Georgia Institute of Technology, Atlanta, GA, 30332 USA. anqiwu@gatech.edu

⁴Khoury College of Computer Sciences, Northeastern University, Boston, MA 02115, USA. s.hutchinson@northeastern.edu

robot interactions, collision areas, emotion-aware discomfort zones, social interactions, and the spatio-temporal evolution of the environment. Discomfort zones can be viewed as soft collision-avoidance regions that the ego-agent should minimize intrusion into as much as possible. We employ a DRL-based navigation architecture to map these latent features directly to bipedal robot actions, specifically, the sagittal center-of-mass (CoM) speed derived from ROMs and the heading angle of Digit’s torso. To address the discrepancies between ROMs and full-order dynamics, we train the social navigation policy directly in a physics-based simulator with a full-body locomotion controller, bypassing reliance on ROMs. This approach ensures that the DRL model accounts for full-body constraints and dynamics inherent to bipedal locomotion. In addition, we benchmark against state-of-the-art (SOTA) navigation pipelines to highlight the efficacy of our framework.

The main contributions of this study are as follows.

- **Novel DRL observation representation.** We introduce LGMs as a novel representation of highly dynamic environments integrated with pedestrian emotions.
- **Integration of pedestrian emotions.** Our framework incorporates pedestrian emotions to design discomfort zones and further enable emotion-aware navigation.
- **End-to-end navigation framework.** We propose an end-to-end navigation pipeline that directly maps latent features to bipedal robot actions and is trained with full-body bipedal robot dynamics, effectively mitigating model discrepancies induced by ROMs.
- **Comprehensive benchmarks and hardware evaluations.** We either reproduce or improve SOTA benchmarks for bipedal social navigation, validating our pipeline against both model-based planners and DRL-based baselines. In addition, we successfully transfer the simulated navigation policy to the real bipedal robot Digit.

II. RELATED WORK

Navigation with bipedal robots. Numerous studies have investigated bipedal robot navigation in relatively simple environments with static and sparsely distributed obstacles [4], [13]–[15]. Recent approaches have begun to address dynamic environments by incorporating moving obstacles alongside static barriers [5], [16]–[18]. However, these dynamic obstacles are sparsely distributed around the robot, e.g. only one single moving obstacle is in close proximity to the ego-agent robot. Moreover, the interactions are oversimplified, assuming slower moving obstacles and neglecting social interactions between pedestrians and robots [17].

Research specifically addressing crowd navigation for bipedal robots remains largely underexplored, with the exception of the recent study in [12], which assumes that bipedal robots and pedestrians adhere to the same social rules, that is, robots learn their navigation path from real pedestrian datasets and imitate pedestrian navigation strategies. However, this assumption is restrictive, as bipedal robots exhibit fundamentally different dynamics and face more stringent physical constraints compared to pedestrians. Furthermore, current collision avoidance strategies are typically conservative because robots risk

freezing due to the stochastic and unpredictable nature of pedestrian trajectory predictions [19].

Social navigation with DRL. Although research on social navigation with bipedal robots is limited, there is a substantial body of work on social navigation for wheeled mobile robots [7]. In particular, end-to-end frameworks that leverage DRL have been widely used to directly map environmental states to robot actions [8]. A prominent avenue is the collision avoidance approach with deep reinforcement learning (CADRL) [19], which formulates the problem as a Markov Decision Process (MDP) and assumes fully observable environments, including pedestrian attributes such as size, position, velocity, and preferred speed. Moreover, CADRL only pairs the ego-agent with each pedestrian individual, neglecting pedestrian interactions. Subsequent studies, such as the social attention RL (SARL) [20] and the relational graph learning (RGL) [21], focus on modeling implicit pedestrian interactions to plan socially acceptable motions. However, these SOTA models rely heavily on fully observable environments which require complicated perception and prediction techniques to obtain in real scenarios. In addition, they assume that the kinematics of mobile robots are less constrained and similar to those of pedestrians, which is not applicable to bipedal robots due to their restrictive locomotion capabilities.

While MDPs are designed for decision making in fully known environments, Partially Observable MDPs (POMDPs) are suitable for decision making in socially interactive environments, which involves partially observable pedestrian interactions and future trajectories [22]. SOTA approaches that map raw or preprocessed sensor data to robot actions can be formulated as POMDPs [23]–[26]. An intuitive representation is the sequential LiDAR scans, which are stacked and encoded with convolutional neural networks (CNNs) [23] or directly processed with long-short-term memory (LSTM) [24]. Combining sequential LiDAR scans with fully known states (i.e., pedestrian position and speed), the DRL velocity obstacle (DRL-VO) method converts these observations and states into an image for feature extraction [25]. In practice, they are dependent on complex estimation techniques to obtain these states [25]. Moreover, their network architectures are susceptible to divergence, plateauing, and collapse when applied to robots with complicated dynamics and strict constraints [23]–[25]. Compared to others, our approach is more straightforward, with only sequential LiDAR scans converted into grid maps. In addition, the neural networks of our feature extractor followed by the soft actor-critic (SAC) DRL algorithm [27] yield a minimal risk of convergence failures for DRL training.

Emotion-aware navigation. Most of the studies discussed above simplify pedestrians as circles, overlooking their social cues such as emotions. Pedestrian emotions implicitly convey information, such as the preferred comfort distance that people prefer to maintain from others [9]. Psychological experiments [10] have identified comfort distances associated with various emotions. With these insights, social robots can adaptively maintain a customized comfort distance from pedestrians with specific emotions [28]. However, these emotion-aware robot navigation approaches focus primarily on the emotions of individuals, neglecting the complex interactions that occur

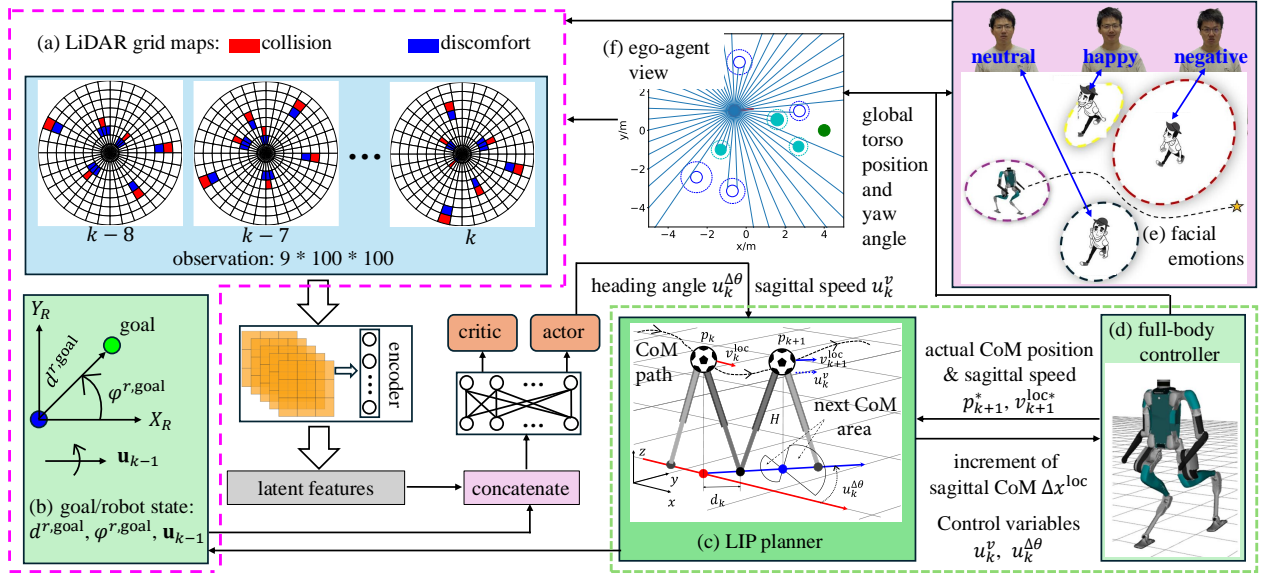


Fig. 2. Overview of the proposed EmoBipedNav framework using the bipedal robot Digit. Our framework begins by obtaining estimated facial emotions (Fig. (e)) using pre-trained CNN models. Simultaneously, we transform raw LiDAR scans (Fig. (f)) into sequential pie-shape LGMs (Fig. (a)), where the red cells denote collision areas and the blue patches highlight discomfort zones associated with pedestrian emotions. These grid maps are converted into stacked pixel images which are further processed through an encoder constructed using CNNs to extract socially interactive and emotionally aware features. The resulting latent features are concatenated with the robot’s last command and target position (Fig. (b)), which are fed into an actor-critic DRL structure implemented with multi-layer perceptrons (MLPs). The action output from the actor network is derived from the ROM (Fig. (c)) and practically applied to a bipedal robot with full-body dynamics and constraints (Fig. (d)). The torso position and yaw angle obtained from Digit correspond to the ego-agent state in Fig. (f). We use the Angular momentum LIP planner (ALIP) [29] and a passivity full-body controller with ankle actuation [17] to track the desirable ROM trajectory.

among pedestrians with various emotional states, which will be addressed in our study.

III. APPROACH

In this section, we first introduce the reduced-order model (ROM) of bipedal locomotion and the Partially Observable Markov Decision Process (POMDP) of the social navigation problem. Subsequently, we present our DRL-based approach, detailing the novel observation space, the DRL network architecture, and the design of the integrated reward functions. The overall framework is illustrated in Fig. 2.

A. Preliminaries

The dynamics of bipedal robots is commonly approximated using ROMs, such as the linear inverted pendulum (LIP) model [29], [30], as depicted in Fig. 2-(c). The discrete state transition equation for the LIP model at the k^{th} walking step in the local sagittal frame is given as follows:

$$d_k = \frac{v_k^{\text{loc}} \cosh(\omega T) - u_k^v}{\omega \sinh(\omega T)}, \quad (1)$$

$$\Delta x^{\text{loc}}(d_k) = v_k^{\text{loc}} \frac{\sinh(\omega T)}{\omega} + d_k (1 - \cosh(\omega T)),$$

where the control input u_k^v represents the desired sagittal velocity at the next foot placement switch instant, the output state $\Delta x^{\text{loc}}(d_k)$ is the local sagittal CoM position increment between two consecutive walking steps, v_k^{loc} denotes the sagittal velocity in the local coordinate at the k^{th} walking step switch instant, d_k is an auxiliary variable that specifies the sagittal foot distance relative to CoM, T is the constant footstep time, the term $\omega = \sqrt{g/H}$ is derived from the

gravitational constant g and the desired CoM height H . The ROM in (1) is equivalent to the formulation of the angular momentum LIP planner (ALIP) model in [29]. However, since our DRL network outputs the sagittal speed, our formulation uses the current sagittal velocity of CoM, v_k^{loc} , instead of the angular momentum about the contact point used in ALIP.

Since u_k^v influences only the translational motion, we introduce a heading angle $u_k^{\Delta\theta}$ to steer the CoM orientation in the horizontal plane, as illustrated in Fig. 2-(c). Combining the translational and rotational motions of the CoM, we define the action of our navigation policy as:

$$\mathbf{u}_k = \{u_k^v, u_k^{\Delta\theta}\}, \quad 0 \leq u_k^v \leq \bar{u}^v, \quad |u_k^{\Delta\theta}| \leq \bar{u}^{\Delta\theta}, \quad (2)$$

where \bar{u}^v and $\bar{u}^{\Delta\theta}$ represent the maximum allowable sagittal speed and heading angle, respectively.

Given the control \mathbf{u}_k , the CoM displacement Δx^{loc} , and the LIP model, the CoM position is propagated from p_k to p_{k+1} , and the sagittal speed changes to v_{k+1}^{loc} from v_k^{loc} . However, there are discrepancies between the actual CoM position p_{k+1}^* and p_{k+1} , and between the actual sagittal speed $v_{k+1}^{\text{loc}*}$ and v_{k+1}^{loc} . More specifically, the full-body controller generates joint torques to track the desired CoM displacement Δx^{loc} and the control \mathbf{u}_k , as illustrated in Fig. 2-(d). However, discrepancies between the LIP model and the full-body dynamics can cause the robot’s states to deviate from the desired ones, which may lead to collisions and intrusions into discomfort areas of pedestrians, as demonstrated in Section IV.

Remark 1: In our previous work [12], the ROM in conjunction with a model predictive controller is used to plan feasible CoM trajectories for Digit. However, in this work, the ROM state (i.e., the heading angle and sagittal speed

in Eq. (1) is used to define the DRL actions. The input of the proposed network comes from the full-body simulation of Digit (Fig. 2(d)), eliminating any discrepancies between the ROM and full-body dynamics, which exist in our previous work. Furthermore, we utilize the ROM for benchmarks, including both model-based motion planners and DRL-based baselines, effectively demonstrating the practical impact of robot dynamics discrepancies.

B. Problem Formulation

We formulate the social navigation problem as a POMDP because the state information of social environments is not fully observable, such as implicit social interactions and predicted pedestrian trajectories. More specifically, we represent social environments with sequential LiDAR grid maps (LGMs). Despite only including partial information, LGMs allow the extraction of implicit features such as collision zones, pedestrian-pedestrian and pedestrian-robot interactions, emotion-related discomfort areas, and the spatio-temporal dynamic evolution of the navigation system.

In addition to LGMs which only contain partial information, we have two other fully observable states. For one thing, the goal position state is responsible for the target-directed navigation task. In addition, since pedestrians interact with the robot and the robot’s actions influence pedestrian decisions, we also include the robot’s actions as part of the states.

Given the observation of LGMs and the fully observable states, our objective is to optimize a navigation policy, straightforwardly projecting the observation and states into the robot actions. The problem is formulated as a POMDP defined by a 6-tuple $(S, O, A, \Gamma, R, \gamma)$, where S is the fully observable state space including the goal position and the robot’s actions, O represents the observation space consisting of sequential LGMs, A stands for the action space, Γ denotes the state transition model fully propagated with the bipedal robot Digit, R is the reward function with an instantaneous scalar reward r_k at the k^{th} walking step and γ is the discount factor. To derive the optimal navigation policy $\pi(a \in A | s \in S, o \in O)$, our goal is to maximize the expected sum of discounted rewards over an infinite time horizon, as follows:

$$V = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k r_k \right]. \quad (3)$$

We will leverage a CNN-based encoder followed by a soft actor-critic (SAC) DRL algorithm [27] to iteratively optimize the navigation policy π .

C. Observation and State Spaces

We propose a novel representation for the observation space, referred to as pie-shape LGMs (Fig. 2-(a)), to comprehensively encapsulate the features of all entities around the ego-agent, including static obstacles having variable sizes and dynamic pedestrians with different emotions. Sequential pie-shape LGMs are built on the basis of raw LiDAR scans. More specifically, a 2-D LiDAR scan with K beams is emitted from

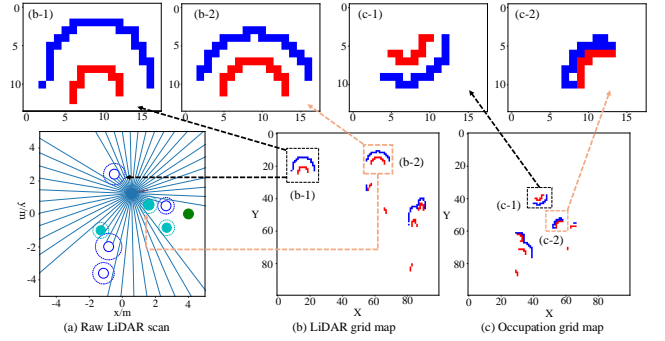


Fig. 3. LiDAR scan and its corresponding grid maps. (a) illustrates the raw LiDAR scan. The red solid circle with an arrow represents the ego-agent and its orientation in the world frame, and the green solid one is the goal. Moving pedestrians are blue hollow circles, and static obstacles are depicted as solid cyan circles. The outer dashed circles illustrate the discomfort margins. (b) and (c) are the corresponding LiDAR grid map and occupation grid map, respectively. X and Y axes in (b) represent rotational and radial indices, while X and Y in (c) are horizontal and vertical pixel positions, respectively. (b-1) and (b-2) shown in (b) are the margins corresponding to two close objects surrounding the ego-agent illustrated in (a). (c-1) and (c-2) correspond to the same objects as (b-1) and (b-2), respectively. Their zoomed-in views are shown on the top row.

the ego-agent to detect collision and discomfort margins, as illustrated in Fig. 2-(f) and Fig. 3-(a).

The detected obstacles are projected onto a pie-shape grid map (shown in Fig. 2-(a)) with a maximum detection radius of L . This map is divided into M radial segments, forming M concentric rings. Each ring is further segmented angularly into M equal sections. The resolution of each grid increases as the ring gets closer to the center, ensuring that more details around the ego-agent are detected. Note that the pie-shape grid map is inherently circular and cannot be directly used for feature extraction with CNNs, which require a rectangular image input. To address this, the map is first flattened and restructured into a LiDAR grid map, as shown in Fig. 3-(b).

Alternatively, occupancy grid maps (OGMs)¹, as shown in Fig. 3-(c), are commonly used for local collision avoidance¹. Given the same objects close to the ego-agent, their representation in Fig. 3-(b) (LGM) contains more pixels compared to Fig. 3-(c) (OGM), indicating that LGM can capture more details when objects are close, which is beneficial for avoiding collisions and alleviating discomfort, as demonstrated in Section IV.

In particular, we define a discomfort margin (blue grids in Fig. 2-(a) and Fig. 3-(b)) that expands according to the collision area (red cells in Fig. 2-(a) and Fig. 3-(b)). The size of this discomfort zone is adjusted according to the emotional state of the pedestrian. We consider three emotional states: happy, neutral, and negative including angry, sad, fearful, and surprised. According to psychological research [10], happy emotions encourage approach behaviors and therefore shorter distances, whereas negative emotions promote avoidant behaviors and longer distances. Based on the studies in [10], [28], we assign the following distances from the robot margin to the pedestrian edge: 0.2 m for happy, 0.35 m for neutral, and 0.5 m for negative emotions.

¹http://wiki.ros.org/costmap_2d

The observation \mathbf{o}_k at the current walking step is the LGM capturing only the static immediate environment in the spatial dimension, which might encourage conservative motions to avoid collisions as quantitatively analyzed in Section IV. To capture dynamic information in the temporal dimension, such as pedestrian-pedestrian and pedestrian-robot interactions, and the evolving dynamics of the navigation system, we incorporate previous LGMs from time step $k - N + 1$ to $k - 1$ as illustrated in Fig. 2-(a). Thus, the observation in both the spatial and temporal dimensions is represented as $\mathbf{o}_k^{\text{sp-tem}} = \{\mathbf{o}_{k-N+1}, \dots, \mathbf{o}_{k-1}, \mathbf{o}_k\}$.

In addition to the observation only including partial information, we define another two fully observable states as illustrated in Fig. 2-(b): the position state of the goal $\mathbf{s}_k^{\text{goal}} = \{d_k^{\text{goal}}, \varphi^{r, \text{goal}}\}$, where d_k^{goal} represents the distance between the goal and the robot and $\varphi^{r, \text{goal}}$ indicates the orientation of the goal in the robot frame; the last robot command \mathbf{u}_{k-1} . These two states are used for the goal-reaching task and the pedestrian-robot interaction, respectively.

D. Network Structure

The encoder, as shown in Fig. 2, consists of CNNs followed by multi-layer perceptrons (MLPs) to fully extract the latent features \mathbf{f}_k^o from the sequential LGMs $\mathbf{o}_k^{\text{sp-tem}}$. Unlike SOTA approaches [20], [21] that rely solely on MLPs to process fully observable states, our encoder is capable of extracting more detailed and implicit information from observation LGMs. Additionally, the proposed image-based observation integrates emotion-related discomfort zones, addressing a challenge faced by cutting-edge methods using raw sensor data [23]. Given these latent features \mathbf{f}_k^o , we concatenate them with the goal state $\mathbf{s}_k^{\text{goal}}$ and the action state \mathbf{u}_{k-1} . The combined inputs are fed into the SAC DRL framework [27] to optimize the navigation policy $\pi(\mathbf{u}_k | \mathbf{o}_k^{\text{sp-tem}}, \mathbf{s}_k^{\text{goal}}, \mathbf{u}_{k-1})$. The actor and value modules, as shown in Fig. 2, are made up of MLPs that take the same inputs $\{\mathbf{f}_k^o, \mathbf{s}_k^{\text{goal}}, \mathbf{u}_{k-1}\}$. In summary, the network structure can be expressed as follows:

$$\mathbf{f}_k^o = f_{\vartheta}(\mathbf{o}_k^{\text{sp-tem}}), \quad (4a)$$

$$v_k = f_{\phi}(\mathbf{f}_k^o, \mathbf{s}_k^{\text{goal}}, \mathbf{u}_{k-1}), \quad (4b)$$

$$\mathbf{u}_k \sim f_{\psi}(\mathbf{f}_k^o, \mathbf{s}_k^{\text{goal}}, \mathbf{u}_{k-1}), \quad (4c)$$

where the network parameters ϑ, ϕ, ψ are updated using the SAC algorithm [27] and the exploration samples collected from the bipedal robot Digit, v_k represents the deterministic value function predicted by the critic network, and \mathbf{u}_k is the action sampled from a Gaussian distribution with its mean and standard deviation parameters output from the actor network. Unlike the original SAC framework, our approach introduces a shared encoder network (4a) that is used by both the critic network (4b) and the actor network (4c). This shared encoder is updated exclusively with the critic network, as it is primarily responsible for evaluating robot behaviors as well as the corresponding observation and states. The encoder is then copied to the actor network without gradient computations, ensuring that the feature extraction process remains consistent across both networks.




Fig. 4. Geometry representations of reward-related notations.

E. Reward Integration

The objectives of our navigation task are threefold: avoidance of collisions, reaching goals, and awareness of emotions. Accordingly, we define the reward function as follows:

$$r_k = r_k^{\text{col}} + r_k^{\text{goal}} + r_k^{\text{emo}}, \quad (5)$$

where r_k^{col} penalizes the collisions with all surrounding objects, r_k^{goal} rewards progress toward reaching the goal, and r_k^{emo} denotes an emotion-related penalty. Specifically, these three rewards are defined as follows.

1) Collision avoidance reward:

$$r_k^{\text{col}} = \begin{cases} -0.6 & \text{if } d_{\min}^{\text{all}} \leq r_{\text{robot}}, \\ -0.1 & \text{else if } d_{\min}^{\text{static}} < d^{\text{static}} + r_{\text{robot}}, \\ 0.0 & \text{else,} \end{cases} \quad (6)$$

where d_{\min}^{all} represents the minimum distance from the ego-agent to the margins of all surrounding objects as shown in Fig. 4, d_{\min}^{static} denotes the minimum distance from the ego-agent to all static obstacles, d^{static} stands for a safety distance expanded from the collision margin of the obstacle, and r_{robot} is the ego-agent radius.

2) Goal reaching reward:

$$r_k^{\text{goal}} = \begin{cases} 0.5 & \text{if } d_k^{\text{goal}} \leq d^{\text{goal}}, \\ 0.3 \cdot (d_{k-1}^{\text{goal}} - d_k^{\text{goal}}) & \text{else,} \end{cases} \quad (7)$$

where d_k^{goal} stands for the distance from the ego-agent to the goal at the k^{th} walking step, and d^{goal} represents the goal-reaching threshold. In particular, the second term encourages the robot to move toward the goal, thus alleviating the sparsity issue of the reward function and promoting efficient training.

3) Emotion-aware reward:

$$r_k^{\text{emo}} = \begin{cases} -0.1 & \text{if } \mathcal{C}, \\ 0.0 & \text{else,} \end{cases} \quad (8)$$

where \mathcal{C} is a distance condition with respect to pedestrian emotions. Specifically, the discomfort distance of the i^{th} pedestrian is defined as $d_{p_i}^{\text{emo}} \in \{0.2, 0.35, 0.5\}\text{m}$, depending on the corresponding emotion of the pedestrian $\{\text{happy, neutral, negative}\}$. r^{ped} is a constant pedestrian radius. Given the distance d_{p_i} from the margin of the i^{th} pedestrian to the ego-agent, the penalty is -0.1 if $\exists i$, such that $d_{p_i} < d_{p_i}^{\text{emo}} + r^{\text{ped}}$. Note that, to avoid learning a conservative navigation policy, we only count once if the ego-agent enters the discomfort zones of several objects, including

TABLE I
PARAMETER (PARAM) SET-UP.

Param	Value	Param	Value	Param	Value
T	0.4 s	H	1.02 m	g	9.81 m/s ²
K	1800	M	100	L	6.0 m
γ	0.99	d^{static}	0.2 m	\bar{u}^v	0.4 m/s
N	9	d^{goal}	0.1 m	$\bar{u}^{\Delta\theta}$	0.2 rad
r^{ped}	0.3 m	r^{static}	[0.2, 0.4] m	r^{robot}	0.3 m

static obstacles and dynamic pedestrians, at the same walking step. That means we add -0.1 only once in Eq. (5) although it appears in both Eq. (6) and Eq. (8).

IV. EXPERIMENTS

We first detail the parameter specifications and the network architecture, followed by the environmental settings. Next, we benchmark our proposed navigation approach against SOTA methods that do not account for emotion-related observations and rewards. Additionally, we implement another training with emotion-related observations and reward functions. Further investigations on time-varying emotions are carried out. Finally, we demonstrate various deployments on the robot hardware.

A. Parameter Specifications and Network Details

TABLE I presents the parameter specifications. Our DRL network architecture consists of 3 modules: encoder, actor, and critic. The encoder comprises 4 layers of CNNs, each with 32 kernels of size 3. The stride of the first layer is 2, while the remaining layers have a stride of 1. After the CNN layers, a fully connected layer extracts 50 features from the sequential LiDAR maps. Both the actor and critic networks use similar MLP architectures, each with two hidden layers of size 1024. Unlike critic networks, the output of the network has only one value, while the network of actors produces 4 variables, corresponding to 2 means and 2 log standard deviations for the action \mathbf{u}_k . The learning rates for both the actor and critic networks are set to 0.001.

B. Environment Settings

The simulation environment includes both static obstacles and dynamically moving pedestrians interacting with a full-body bipedal robot.

Static obstacles. Static obstacles are randomly placed within a rectangular area, with the x -coordinate ranging from -3 to 3 m and the y -coordinate from -1.5 to 1.5 m. Each obstacle has a radius ranging from 0.2 to 0.4 m.

Pedestrians. The initial position of each pedestrian is distributed along the perimeter of a circle with a radius randomized from 3.5 to 4.5 m. The corresponding goal of each pedestrian is located on the other side of the circle. Pedestrians move between these two points at a preferred speed of 1 m/s.

Pedestrian motion is generated using the Optimal Reciprocal Collision Avoidance (ORCA) method [32], a multi-agent motion planner commonly used to simulate interactive motions

of all agents [19]–[21], [23], [25]. In addition to pedestrian-pedestrian interactions, we also enable pedestrian-robot interactions when the pedestrian’s speed exceeds the robot’s speed. Otherwise, the robot remains invisible to pedestrians, and only the robot actively avoids collisions with pedestrians. Such scenarios where pedestrians are not always trying to avoid the ego-agent are designed to independently demonstrate the active collision avoidance capability of our navigation policy.

Robot. The robot navigates from position $(-4, 0)$ to position $(4, 0)$ m. We simulate the bipedal robot Digit in MuJoCo [31]. The navigation policy in the high-level motion planner only generates the sagittal speed and heading angle derived from the ROM.

For tracking sagittal speed, we first use the ALIP planner [29] to generate foot placements. To address gaps in robot dynamics, we incorporate a passivity-based controller [33], which effectively preserves the natural dynamics of the system while ensuring asymptotic tracking performance. In our previous work [12], [17], the addition of toe-joint actuation alongside the whole-body passivity controller demonstrated promising results in minimizing the mismatch between the LIP model trajectory and the full-order trajectory of the Digit bipedal robot. Despite mitigating these dynamics gaps, the tracking error in the sagittal velocity persists, especially in the peaks and valleys [12]. To track the heading angle, we use the hip yaw-actuator to align the Digit’s torso to the desired CoM direction. However, as demonstrated in our previous study [12], the tracking error in the torso orientation of Digit remains detrimental, especially when there are frequent changes in the desired torso orientation.

These inaccurate trackings negatively impact social navigation performance when deploying the motion planner optimized with ROMs on the full-body bipedal robot Digit, a practical issue we will further investigate in Section IV-C and IV-D. To address these tracking errors, we directly train our navigation policy with the full-body bipedal robot, allowing it to account for tracking errors in both sagittal velocity and heading angle.

C. Benchmarks

To highlight the superior social navigation capabilities of our approach, we first disable emotion-related observations and reward functions, ensuring a fair comparison with SOTA baselines. In addition, we adapt the traditional methods, originally designed for wheeled mobile robots, to be deployable on Digit simulated in MuJoCo. To validate the performance discrepancies between the use of full-body bipedal robot dynamics and the use of the ROM, we also deploy certain SOTA approaches that rely exclusively on the ROM. The SOTA baselines are listed below.

DWA. The Dynamic Window Approach (DWA) [34] is a traditional non-learning method widely applied for local navigation using wheeled mobile robots. In this study, we modify DWA by replacing the kinematics of wheeled mobile robots with the LIP model (1), enabling optimization of the sagittal speed and heading angle.

LNDNL. The “Learn to Navigate in Dynamic environments with Normalized LiDAR scans” (LNDNL) method [24] is a

DRL-based navigation approach that directly maps the raw LiDAR scans to robot actions. LSTM networks are used to process the sequential LiDAR scans, capturing spatio-temporal interaction features.

LiDAR-SAC. LiDAR-SAC is inspired by the SOTA study in [23], using an encoder to extract latent features from sequential LiDAR scans stacked over the time horizon. However, the original approach faces challenges in learning a feasible navigation policy, as empirically validated in [24]. Therefore, we replace the original networks with ours, including the encoder and the SAC structure, to validate the effectiveness of our network pipeline.

DRL-VO. As discussed in Section II, the DRL-VO pipeline [25] combines fully observable states, such as pedestrian positions and velocities, with observations of raw sequential LiDAR scans. The combined information is transformed into an image, enabling the extraction of features from dynamic and interactive environments using CNNs. However, the original network pipeline consistently takes suboptimal actions, leading to minimal reward. To address this, we improve it by integrating our SAC structure, resulting in the modified DRL-VO*, which will be used as a baseline in our study.

SARL and RGL. SARL [20] and RGL pipelines [21] assume fully observable environments. In addition, ego-agents are assumed to have similar dynamics and constraints to pedestrians. However, when the ego-agent is a bipedal robot with velocity constraints in Eq. 2, neither SARL nor RGL succeed in achieving a feasible navigation policy, whether using ROMs or full-order dynamics.

Single-LGM. Single-LGM uses only one single LiDAR grid map (LGM) at the current walking step instead of a sequence of historical ones, as illustrated in Fig. 2-(a).

OGM. OGM leverages a sequence of historical occupancy grid maps (OGMs) shown in Fig. 3-(c) in the frame of the ego-agent, replacing the LGMs illustrated in Fig. 3-(b).

Evaluation Metrics. We evaluate each benchmark method over 500 random trials. The initial positions of dynamic and static objects are set randomly for each evaluation, while the initial and goal positions of Digit’s torso remain the same. For each trial, all benchmarks start with the same initial environments, ensuring that dynamic and static objects have the same initial positions and pedestrian emotions are consistent. The success rate (SR) is defined as the ratio of successful episodes that occur without collision or timeout. The navigation time (NT) represents the average time taken for all successful episodes. Discomfort times (DT) count the number of walking steps when Digit enters the discomfort zone of a static obstacle or a pedestrian.

As illustrated in Fig. 5, both LiDAR-SAC and DRL-VO* produce results similar to ours when their original network structures are replaced with the architecture described in Section III, further validating the robustness and effectiveness of our network pipeline. Moreover, DRL-VO* has the same network structure as ours but yields inferior performance in terms of the stable success rate, further indicating that the proposed LGMs are better representations.

The quantitative results, presented in TABLE II, confirm that our method achieves the highest SR for the deployment

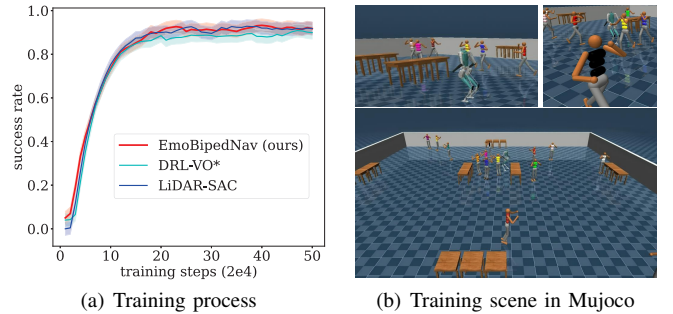


Fig. 5. Training process with simulated Digit in MuJoCo. (a) shows the training processes of different approaches and (b) illustrates a representative simulation scene. We evaluate each model 100 times in randomly generated environments at every 2×10^4 training steps. The success rate is calculated based on these 100 randomized evaluations. DRL-VO* and LiDAR-SAC are enhanced by replacing their original network frameworks with our proposed architecture. LNDNL is omitted because its success rate climbs to 0.66 and decreases to 0. In addition, SARL and RGL are excluded from the analysis as they fail. The trainings of Single-LGM and OGM are removed due to their similar learning curves to our method.

TABLE II
FINAL PERFORMANCE EVALUATION ON THE FULL-BODY BIPEDAL ROBOT WITH 500 RANDOM TESTS.

Method	SR	NT	DT
DWA	0.892	27.20 ±3.55	415
LNDNL	0.660	38.13±5.99	301
DRL-VO*	0.912	28.16± 2.49	318
LiDAR-SAC	0.932	27.24±3.10	766
Single-LGM	0.928	33.15±4.84	546
OGM	0.944	29.45±3.79	312
EmoBipedNav (ours)	0.950	28.83±3.15	226

on the full-body bipedal robot Digit. DWA enables efficient navigation with the shortest average time of 27.2 s, at the cost of decreasing SR (0.892) and increasing DT. It is noteworthy that DWA can reach a higher SR (0.950) when deployed with the LIP model without considering full-order dynamics, which underscores the notable dynamics discrepancies between the ROM and the full-body robot dynamics. LNDNL demonstrates significantly inferior performance in terms of both SR and NT due to the plateauing problem during training. Notably, we observed that LNDNL achieves results comparable to our approach when trained and evaluated with the ROM without incorporating full-body dynamics, emphasizing our pipeline’s capability to effectively handle the complex dynamics of the full-body robot. Although Single-LGM can also achieve a high SR, its NT significantly increases due to its conservative policy, which does not account for the temporal dynamics of the social environment. OGM, on the other hand, shows a comparable SR, but its DT is higher than our method, indicating that LGMs are more effective in reducing discomfort.

D. Emotion-integrated Implementation

To investigate further, we conduct an additional training session to assess the performance of integrating emotion observations. The results demonstrate that our approach can adapt to pedestrians with different emotions, showing the emotion-aware feature of our navigation policy. Additionally, we vali-




Fig. 6. Training conditioned with emotions. (a) illustrates the learning processes using two constant discomfort distances and one variable discomfort distance corresponding to different pedestrian emotions. All trainings in (b) use the variable discomfort distance. EmoBipedNav and OGM leverage full-order dynamics while the LIP one only relies on the ROM.

TABLE III
PERFORMANCE EVALUATION W/O EMOTIONS AND W/O FULL-BODY ROBOT DYNAMICS WITH 500 RANDOM TESTS.

DD (Train)	Train/Test	Obs	IDT (Test)	SR
0.2 m	Digit/Digit	LGM	[51; 50, 181, 1352]	0.950
0.5 m	Digit/Digit	LGM	[395; 81, 138, 1167]	0.938
Var-DD	Digit/Digit	LGM	[213; 56, 166, 1198]	0.940
Var-DD	Digit/Digit	OGM	[436; 87, 160, 1174]	0.892
Var-DD	LIP/LIP	LGM	[202; 37, 143, 1227]	0.962
Var-DD	LIP/Digit	LGM	[642; 66, 177, 1266]	0.900

DD: discomfort distance (m); Obs: observation
Var-DD= [0.2, 0.35, 0.5] m \leftrightarrow [happy, neutral, negative]

date the performance discrepancies of robot social navigation when using the ROM and the full-body bipedal robot. Moreover, we practically verify that using LGMs provides greater benefits for collision and discomfort avoidance compared to using OGMs in these emotion-aware environments.

First, we use variable discomfort distances (Var-DD) indicating different emotions of pedestrians, and the learning process is similar to the case where a constant discomfort distance of 0.2 m is used, as shown in Fig. 6-(a). This similarity suggests that our approach can effectively avoid collisions in emotion-integrated environments. Next, we set the discomfort distance to a constant 0.5 m, resulting in a learning process similar to the other two settings. However, progress slows, and the stable success rate is lower. We attribute this to the aggressive pedestrian motion generator ORCA [32], which challenges the conservative movement of the ego-agent to avoid collisions. Furthermore, we replace LGMs with OGMs, resulting in similar learning performance. However, as shown in Fig. 6-(b), our approach demonstrates a larger stable success rate, underscoring the advantages of LGMs for collision avoidance. Lastly, we train our pipeline using only the LIP robot model. Although learning based on the LIP model is more efficient due to simplified dynamics, performance with respect to SR and DT significantly degrades when the learned navigation policy is deployed on Digit, as demonstrated in TABLE III.

To account for pedestrian emotions and their associated discomfort distances, we introduce an additional evaluation metric, IDT, which quantifies the discomfort times for individual entities. Specifically, the four values in the IDT column

of TABLE III correspond to the IDT of static obstacles and pedestrians exhibiting happy, neutral, and negative emotions, respectively. During training, the two baseline models enforce a fixed discomfort distance of 0.2 m or 0.5 m for all dynamic pedestrians, regardless of their emotional states. However, during testing, IDT is computed based on variable discomfort distances of [0.2, 0.35, 0.5] m, corresponding to the three emotional states. This discrepancy between training and testing conditions allows us to evaluate the model’s ability to adapt to varying discomfort zones.

Compared to maintaining a constant discomfort distance of 0.2 m, our approach of using variable discomfort distances of [0.2, 0.35, 0.5] m decreases IDT for neutral and negative pedestrians while increasing IDT for happy pedestrians, and static obstacles in particular. This indicates that our approach prioritizes neutral and negative pedestrians by assigning them longer discomfort distances. However, this comes at the cost of reduced comfort for happy pedestrians and, notably, static obstacles. The performance degradation with respect to static obstacles may result from the reduction in the available free areas. This occurs because the discomfort zones of neutral and negative pedestrians can expand to occupy significantly more space while the overall environment space remains the same. Furthermore, the ego-agent is inherently more prone to collisions with moving pedestrians due to highly dynamic interactions. As a result, it prioritizes maintaining distance from these pedestrians, and inadvertently increases the likelihood of intruding into the discomfort zones of static obstacles. When the constant discomfort distance is increased to 0.5 m, the ego-agent further improves IDT for neutral and negative pedestrians, yet compromising comfort for happy pedestrians and static obstacles. The increased IDT for static obstacles could be attributed to the same factors observed when using variable discomfort distances [0.2, 0.35, 0.5] compared to a fixed discomfort distance of 0.2 m. In addition, the larger discomfort distance leads to a decline in SR, which may stem from the aggressive pedestrian motion generated by ORCA [32], despite the conservative behaviors of the ego-agent for intrusion and collision avoidance.

When replacing our LGMs with OGMs, both SR and IDT of static obstacles significantly degrade, further demonstrating the effectiveness of LGM in enhancing collision and discomfort avoidance. Additionally, while the pipeline trained and tested with the LIP model achieves the highest overall performance, its SR drops to 0.9, and IDT increases substantially when tested with Digit, further highlighting the limitations of ROMs in maintaining robust navigation performance.

E. Implementation with Time-varying Emotions

In real-world scenarios, emotions can fluctuate or be misclassified due to inaccuracy in the facial recognition algorithm, which is a pretrained CNN model relying on small dataset². To investigate the adaptability of our approach to these uncertainties, we introduce a more challenging simulation setup where one pedestrian’s emotion keeps alternating among

²<https://github.com/susantabiswas/realtime-facial-emotion-analyzer>

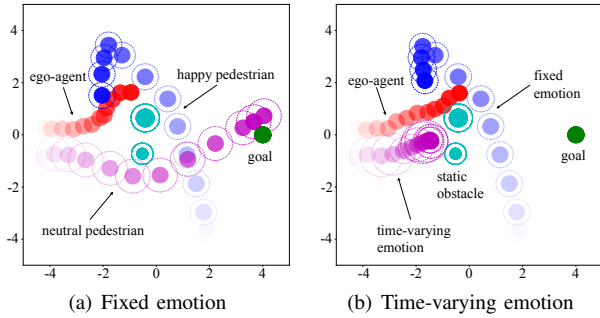


Fig. 7. Robot and pedestrian trajectories under different emotion settings. (a) The emotions of two pedestrians remain fixed, with one classified as *happy* and the other as *neutral*. (b) One pedestrian’s emotion dynamically changes every 2 seconds. The decreasing transparency illustrates the time evolution.

happy, neutral, and negative every 2 seconds. Statistically, SR roughly remains the same, while IDT shifts from [213; 56, 166, 1198] to [258; 42, 204, 1317]. The similar SR shows that our approach is stable for collision avoidance, although pedestrian emotions vary over time. The increased IDT for neutral and negative pedestrians may result from the ego-agent intruding into larger discomfort zones when nearby pedestrians suddenly change to a negative emotional state from neutral, or from happy to neutral. Fig. 7 presents a representative episode, comparing the case with fixed emotions (Fig. 7-(a)) to that with time-varying emotions (Fig. 7-(b)). The latter scenario shows increased discomfort times, particularly at the beginning of the episode when a nearby pedestrian frequently changes the emotion, particularly from neutral to negative. Moreover, frequently changing pedestrian emotion not only impacts the ego-agent’s trajectory, but also alters pedestrians’ interactions, indicating complex real-time interactions between robot-pedestrian, further verifying the adaptability of our social navigation policy.

F. Hardware Demonstrations

LiDAR scan reconstruction. Since real-world objects have irregular geometries, whereas simulated objects are assumed to be circular, directly using raw LiDAR scans in real environments leads to sensor mismatches, creating sim-to-real observation gaps for our navigation policy. To mitigate this, we reconstruct LiDAR scans to better align with the simulated environment. Specifically, we first determine the positions of surrounding objects in the ego-agent frame and approximate their shapes as circles with a fixed radius of 0.3 m for simplification. Next, we emit K rays from the ego-agent to detect these circular approximations, generating a reconstructed LiDAR scan that more closely resembles the simulated observations.

Deployments under various pedestrian motion patterns. We begin by transferring the simulated policy learned in MuJoCo [31] to the physical robot Digit. Furthermore, we include various pedestrian motion patterns to verify the sim-to-real capabilities of our emotion-aware navigation policy. Fig. 8 shows two real-world implementations, involving dynamic pedestrians and static obstacles. Notably, in the left column, a pedestrian intentionally walks in front of Digit and remains




Fig. 8. Screenshots of hardware deployments under various pedestrian navigation patterns. Each column corresponds to an implementation. In each screenshot, a LiDAR scan figure is overlaid to illustrate the robot’s position and orientation, represented by a solid red circle and a red arrow, respectively. Object collision margins are depicted as blue hollow circles, while discomfort margins are indicated with dashed circles. We track the robot and pedestrians with an external motion capture system for accurate trajectory analysis. Pedestrian emotions are randomly set at the beginning and remain constant for each implementation.

stationary to obstruct its path. Additionally, pedestrians move ingroup together in the right column. Despite these diverse and challenging conditions, Digit successfully reaches its goal without collision while actively avoiding intrusions into discomfort zones. More pedestrian motion patterns are attached in the accompanying video to further validate the sim-to-real capabilities of our pipeline.

Deployments integrating emotion recognition. To further validate the practicality of our pipeline with integrated pedestrian emotion recognition, we equip the physical Digit with a ZED stereo camera. Robot localization and pedestrian detection are performed using camera APIs³. For each detected pedestrian, facial features are extracted using OpenCV libraries⁴, and emotions are recognized using a pre-trained model². Localization and pedestrian detection are performed on a Jetson Orin at 10 Hz, while face extraction and emotion recognition operate at 0.5 Hz. Meanwhile, the low-level controller and social navigation policy run on an Intel NUC at 1000 Hz and 2.5 Hz, respectively.

Real-world deployments, as demonstrated in the accompanying video, confirm the seamless transferability of our navigation framework from simulation to real-world applications. Fig. 9 provides representative screenshots from a hardware demonstration. When pedestrians are far from Digit or facing away from the camera, their faces cannot be detected and are therefore labeled with *neutral* emotions. Once their faces become clearly visible, their emotions change to *happy*. In the accompanying video, we showcase additional scenarios where pedestrian emotions transition from *neutral* to

³<https://www.stereolabs.com/docs>

⁴https://docs.opencv.org/3.4/db/d28/tutorial_cascade_classifier.html

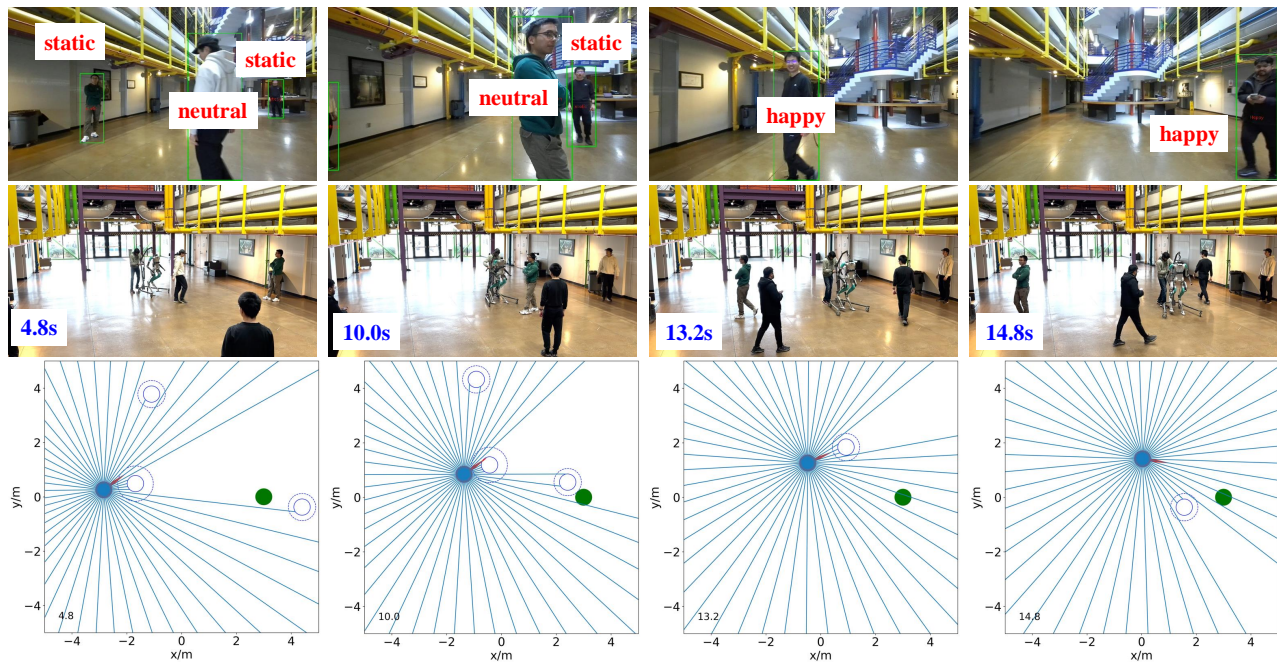


Fig. 9. Emotion recognition with an on-board ZED camera. The top row presents the ZED camera’s perspective, showcasing detected pedestrians along with their recognized emotions. The middle row shows the recorded scenes from a fixed external camera view. The bottom figures illustrate a top-down view of the reconstructed LiDAR scan, position and orientation of the ego-agent, and collision and discomfort margins of pedestrians.

negative or shift between other emotional states. Furthermore, we operate Digit on campus with more realistic social interactions, further demonstrating the practicality of our emotion-aware social navigation pipeline for bipedal robots (see the attached video).

V. DISCUSSION

In this study, we present a DRL-based navigation framework designed to enable bipedal robots to navigate socially interactive environments integrated with pedestrian emotions. Future research will focus on incorporating more realistic social norms, such as adhering to conventions like “walking on the right”, and conducting DRL training in increasingly complex social scenarios, including walls and grouped pedestrians. The ultimate goal is to facilitate natural and socially compliant locomotion for bipedal robots in crowded environments.

REFERENCES

- [1] Z. Gu, J. Li, W. Shen, W. Yu, Z. Xie, S. McCrory, X. Cheng, A. Shamsah, R. Griffin, K. Liu, and others, “Humanoid locomotion and manipulation: Current progress and challenges in control, planning, and learning,” *arXiv preprint arXiv:2501.02116*, 2025.
- [2] F. Wu, Z. Gu, H. Wu, A. Wu, and Y. Zhao, “Infer and adapt: bipedal locomotion reward learning from demonstrations via inverse reinforcement learning,” in *IEEE International Conference on Robotics and Automation*, 2024.
- [3] K. Muenpravitvej, J. Jiang, A. Shamsah, S. Coogan, and Y. Zhao, “Bipedal safe navigation over uncertain rough terrain: unifying terrain mapping and locomotion stability,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024.
- [4] Z. Li, J. Zeng, S. Chen, and K. Sreenath, “Autonomous navigation of underactuated bipedal robots in height-constrained environments,” *The International Journal of Robotics Research*, vol. 42, no. 8, pp. 565–585, 2023.
- [5] K. S. Narkhede, A. M. Kulkarni, D. A. Thanki, and I. Poulakakis, “A sequential mpc approach to reactive planning for bipedal robots using safe corridors in highly cluttered environments,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11831–11838, 2022.
- [6] G. Gibson, O. Dosunmu-Ogunbi, Y. Gong, and J. Grizzle, “Terrain-adaptive, alip-based bipedal locomotion controller via model predictive control and virtual constraints,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 6724–6731.
- [7] P. T. Singamaneni, et al., “A survey on socially aware robot navigation: Taxonomy and future challenges,” *The International Journal of Robotics Research*, vol. 43, no. 10, pp. 1533-1572, 2024.
- [8] K. Zhu and T. Zhang, “Deep reinforcement learning based mobile robot navigation: A review,” *Tsinghua Science and Technology*, vol. 26, no. 5, pp. 674-691, 2021.
- [9] J. R. Martinez, A. Spalanzani, and C. Laugier, “From proxemics theory to socially-aware navigation: A survey,” *International Journal of Social Robotics*, vol. 7, pp. 137-153, 2015.
- [10] G. Ruggiero, et al., “The effect of facial expressions on interpersonal and interpersonal spaces,” *Psychological research*, vol. 81 pp. 1232-1240, 2017.
- [11] M. Ciocca, P. B. Wieber, and T. Fraichard, “Effect of planning period on MPC-based navigation for a biped robot in a crowd,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 491-496.
- [12] A. Shamsah, K. Agarwal, N. Katta, A. Raju, S. Kousik, and Y. Zhao, “Socially acceptable bipedal robot navigation via social zonotope network model predictive control,” *IEEE Transactions on Automation Science and Engineering*, doi: 10.1109/TASE.2024.3519012, 2024.
- [13] Y. D. Hong and B. Lee, “Real-time feasible footstep planning for bipedal robots in three-dimensional environments using particle swarm optimization,” *IEEE/ASME Transactions on Mechatronics*, vol. 25, no. 1, pp. 429-437, 2020.
- [14] J. Liu, M. Li, J. W. Grizzle, and J. K. Huang, “CLF-CBF constraints for real-time avoidance of multiple obstacles in bipedal locomotion and navigation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [15] C. Peng, O. Donca, G. Castillo, and A. Hereid, “Safe bipedal path planning via control barrier functions for polynomial shape obstacles estimated using logistic regression,” in *IEEE International Conference on Robotics and Automation*, 2023.
- [16] A. C. Hildebrandt, et al., “Real-time path planning in unknown environ-

ments for bipedal robots,” *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 1856-1863, 2017.

- [17] A. Shamsah, Z. Gu, J. Warnke, S. Hutchinson, and Y. Zhao, “Integrated task and motion planning for safe legged navigation in partially observable environments,” *IEEE Transactions on Robotics*, vol. 39, no. 6, pp. 4913-4934, 2023.
- [18] K. S. Narkhede, D. A. Thanki, A. M. Kulkarni, and I. Poulakakis, “Overtaking moving obstacles with Digit: path following for bipedal robots via model predictive contouring control,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [19] Y. Chen, M. Liu, M. Everett, and J. P. How, “Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation*, 2017.
- [20] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, “Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning,” in *International Conference on Robotics and Automation*, 2019.
- [21] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva, “Relational graph learning for crowd navigation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [22] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra, “Planning and acting in partially observable stochastic domains,” *Artificial Intelligence*, vol. 101, pp. 99-134, 1998.
- [23] J. Jin, N. M. Nguyen, N. Sakib, D. Graves, H. Yao, and M. Jagersand, “Mapless navigation among dynamics with social-safety-awareness: A reinforcement learning approach from 2D laser scans,” in *IEEE International Conference on Robotics and Automation*, 2020.
- [24] W. Zhu and M. Hayashibe, “Learn to navigate in dynamic environments with normalized LiDAR scans”, in *IEEE International Conference on Robotics and Automation*, 2024.
- [25] Z. Xie and P. Dames, “DRL-VO: Learning to navigate through crowded dynamic scenes using velocity obstacles,” *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2700-2719, 2023.
- [26] S. Lian and F. Zhang, “A transferability metric using scene similarity and local map observation for DRL navigation,” *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 6, pp. 4423-4433, 2024.
- [27] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor”, in *International Conference on Machine Learning*, 2018.
- [28] V. Narayanan, B. M. Manoghar, R. P. RV, and A. Bera, “EWareNet: Emotion-aware pedestrian intent prediction and adaptive spatial profile fusion for social robot navigation”, in *IEEE International Conference on Robotics and Automation*, 2023.
- [29] Y. Gong, and J. Grizzle, “One-step ahead prediction of angular momentum about the contact point for control of bipedal locomotion: Validation in a lip-inspired controller,” in *IEEE International Conference on Robotics and Automation*, 2022, pp. 2832-2838.
- [30] S. Kajita, F. Kanehiro, K. Kaneko, K. Yokoi, and H. Hirukawa, “The 3D linear inverted pendulum mode: a simple modeling for a biped walking pattern generation,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2001.
- [31] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *IEEE/RSJ international conference on intelligent robots and systems*, 2012.
- [32] J. Berg, S. J. Guy, M. Lin, and D. Manocha, “Reciprocal n-body collision avoidance,” *Robotics Research: Springer Tracts in Advanced Robotics*, vol. 70, pp. 3-19, 2011.
- [33] H. Sadeghian, C. Ott, G. Garofalo, and G. Cheng, “Passivity-based control of underactuated biped robots within hybrid zero dynamics approach,” in *IEEE International Conference on Robotics and Automation*, 2017.
- [34] D. Fox, W. Burgard, and S. Thrun, “The dynamic window approach to collision avoidance,” *IEEE Robotics and Automation Magazine*, vol. 4, no. 1, pp. 23-33, 1997.




Wei Zhu (Member, IEEE) received the Bachelor and Master degrees from Nankai University, Tianjin, China, in 2017 and 2020, respectively, and the Ph.D. from Tohoku University, Sendai, Japan, in 2023. He is a Post-Doctoral Scholar at Georgia Institute of Technology, Atlanta, GA, USA. He will join Tohoku University as an Assistant Professor in 2025. His research interests include deep reinforcement learning, robot navigation in crowds, bipedal robots, and snake-like robots.



Abirath Raju received the B.Tech degree in mechanical engineering with a minor in computer science from the National Institute of Technology, Tiruchirappalli, Tamil Nadu, India, in 2023. He is currently pursuing the M.S. degree in robotics at Georgia Institute of Technology, Atlanta, GA, USA. His research interests include bipedal robot navigation and planning, computer vision and robot learning.



Abdulaziz Shamsah received the B.S. degree in mechanical engineering from the Rensselaer Polytechnic Institute, Troy, NY, USA, in 2016, and the M.S.E. degree in mechanical engineering and applied mechanics from the University of Pennsylvania, Philadelphia, PA, USA, in 2018., and the Ph.D. degree in mechanical engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2024. He is currently an Assistant Professor in the Mechanical Engineering Department at the College of Engineering and Petroleum, Kuwait University, Kuwait. His research interests include bipedal navigation in real-world environments, formal methods, and safety and robustness in locomotion.



Anqi Wu is an Assistant Professor at the School of Computational Science and Engineering (CSE), Georgia Institute of Technology. She was a Postdoctoral Research Fellow at the Center for Theoretical Neuroscience, the Zuckerman Mind Brain Behavior Institute, Columbia University. She received her Ph.D. degree in Computational and Quantitative Neuroscience and a graduate certificate in Statistics and Machine Learning from Princeton University. Anqi was selected for the MIT Rising Star in EECs, DARPA Riser, Alfred P. Sloan Fellow, and Kavli

Fellow by National Academy of Sciences. Her research focuses on developing scientifically grounded statistical models to uncover structure in neural and behavioral data at the intersection of machine learning and computational neuroscience. She is broadly interested in creating data-driven models to advance biological intelligence and artificial intelligence.



Seth Hutchinson (Fellow, IEEE) received the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN, USA, in 1988.

Seth Hutchinson is a professor at Northeastern University. He was previously the Executive Director of the Institute for Robotics and Intelligent Machines at the Georgia Institute of Technology, where he was also Professor and KUKA Chair for Robotics in the School of Interactive Computing (2018-2024), and he is Professor emeritus in the ECE Department at the University of Illinois in Urbana-Champaign, member during 1990-2017. He received his Ph.D.

where he was a faculty member from Purdue University.

Hutchinson served as president of the IEEE Robotics and Automation Society, Editor-in-Chief for the “IEEE Trans. on Robotics” and was founding Editor-in-Chief of the RAS Conference Editorial Board. He has more than 300 publications on the topics of robotics and computer vision, and is coauthor of two books on robotics.



Ye Zhao (Senior Member, IEEE) received the Ph.D. degree in mechanical engineering from The University of Texas at Austin in 2016.

He is currently an Assistant Professor at George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, USA. His research interests include robust task and motion planning, contact-rich trajectory optimization, and formal methods for legged locomotion and navigation.

Dr. Zhao received the CAREER Award from the National Science Foundation in 2022 and the Young Investigator Award from the Office of Naval Research in 2023. He serves as an Associate Editor for IEEE Transactions on Robotics, IEEE/ASME Transactions on Mechatronics, IEEE Robotics and Automation Letters, and IEEE Control Systems Letters.