

Joint Communication and Motion Planning for Cobots

Extended Version

Mehdi Dadvar¹, Keyvan Majd¹, Elena Oikonomou¹, Georgios Fainekos¹, and Siddharth Srivastava¹

Abstract—The increasing deployment of robots in co-working scenarios with humans has revealed complex safety and efficiency challenges in the computation of the robot behavior. Movement among humans is one of the most fundamental—and yet critical—problems in this frontier. While several approaches have addressed this problem from a purely navigational point of view, the absence of a unified paradigm for communicating with humans limits their ability to prevent deadlocks and compute feasible solutions. This paper presents a joint communication and motion planning framework that selects from an arbitrary input set of robot’s communication signals while computing robot motion plans. It models a human co-worker’s imperfect perception of these communications using a noisy sensor model and facilitates the specification of a variety of social/workplace compliance priorities with a flexible cost function. Theoretical results and simulator-based empirical evaluations show that our approach efficiently computes motion plans and communication strategies that reduce conflicts between agents and resolve potential deadlocks.

I. INTRODUCTION

Technological breakthroughs of the past decade have led to increasingly common human-robot co-working environments [1]. Navigating among humans is an imperative task that most cobots, ranging from industrial to service robots, are expected to perform safely and efficiently. Although motion planning for autonomous robots has been studied from multiple perspectives [2]–[4], these approaches focus on movement actions and do not address the problem using communication to resolve situations that require extensive human-robot interaction. The objective of this paper is to develop a unified paradigm for computing movement and communication strategies that improve efficiency and reduce movement conflicts in co-working scenarios (see Fig. 1).

Although the problem of integrated task and motion planning has received significant research attention [5]–[10] the integration of these deliberative processes with communication has not been studied sufficiently. Prior work on this topic includes extensions to sampling based motion planning paradigms that model pedestrians as moving obstacles [11], [12]. While these extensions provide valuable enhancements of well-known and efficient algorithms, they view humans as impervious entities and have limited applicability in co-working scenarios where both the human and the robot need to adjust their behavior to allow feasible solutions. On the

This work was supported in part by the NSF under grants IIP-1361926, IIS-1909370, IIS-1844325, OIA-1936997 and the NSF I/UCRC Center for Embedded Systems.

¹ The authors are with the School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ, USA. {mdadvar, majd, e.oikonomou, fainekos, siddharths}@asu.edu




Fig. 1: An example of a social navigation scenario in a confined environment where the robot’s movement can not reveal any information about its future intentions.

other hand, there are approaches that employ disjoint prediction models to establish simple interactions with humans to generate safer and more risk-aware motion plans [13]. Since these approaches neglect the effect of the robot’s motion on the human’s behavior, they suffer from *the robot freezing problem* where the robot cannot find any safe solution. To address this limitation, *socially compliant* methods consider potential human-robot cooperation via learning and planning techniques to produce legible plans or plans subject to stipulations on the information divulged during plan execution [14]–[16]. [17] employs inverse reinforcement learning (IRL) to learn interactive models of pedestrians in the environment for social compliant path planning. Further, [18] presents a social navigation framework that adapts the social force model (SFM) to generate human-friendly collision-free motion plans in unknown environments. Although these approaches model the effect of the robot’s movement on the humans’ behavior for legible motion planning, relying purely on motion actions, taxonomically known as implicit communication [19], could be misleading for the human [20] and may lead to deadlocks in confined environments.

Clearly, employing explicit communication [21] coupled with the robot’s movements would enrich the human-robot interaction. [22] uses IRL to model the effects of both explicit and implicit actions of the robot on the human’s behavior. Further, a robot planner relies on this model to produce communicative actions to maximize the robot’s clarity. Since this method assumes predefined behavior modes for the robot and human (robot priority and human priority), the solution always impels one agent to slow down, which degrades the planning effectiveness.

In contrast, we formalize a unified deliberative communication planning problem that addresses the joint problem of computing the robot’s communication strategy and movements while taking into account the human’s imperfect perception about the robot and its communications (Sec. IV). We use a noisy communication model to estimate the results of robot’s communications on the human’s belief of the robot’s possible locations. In contrast to the human prediction framework in [22], which requires the robot’s future trajectories (the need for socially compliant planning illustrates the difficulty of obtaining such inputs), our approach supports arbitrary human movement prediction models that can predict human behaviors given a set of possible obstacles. Our solution paradigm derives estimates of the human’s belief on the robot’s positions to compute robot communication and movement plans (Sec. IV-C)). This is done using a hierarchical search process with a socially compliant motion planner Control Barrier Function enabled Time-Based RRT (CBF-TB-RRT) [23] (Sec. IV-B). Theoretical results and extensive simulations on various test environments show that this approach efficiently avoids deadlocks and computes mutually efficient solutions without requiring preset behavior modes.

II. PRELIMINARIES

A. Control Barrier Function (CBF)

Assume that the robot R is following a nonlinear control affine dynamics as

$$\dot{\mathbf{s}}_r = \mathbf{f}_r(\mathbf{s}_r) + \mathbf{g}_r(\mathbf{s}_r)\mathbf{a}_r, \quad (1)$$

where $\mathbf{s}_r \in \mathcal{S}_R \subseteq \mathbb{R}^n$ denotes the state of R , $\mathbf{a}_r \in \mathcal{A}_R \subseteq \mathbb{R}^m$ is the control input, and $\mathbf{f}_r : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\mathbf{g}_r : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz functions.

A function $\alpha : \mathbb{R} \rightarrow \mathbb{R}$ is an extended class \mathcal{K} function iff it is strictly increasing and $\alpha(0) = 0$ [24]. A set $\mathcal{C} \subseteq \mathbb{R}^n$ is forward invariant w.r.t the system (1) iff for every initial state $\mathbf{s}_r^0 \in \mathcal{C}$, its solution satisfies $\mathbf{s}_r^t \in \mathcal{C}$ for all $t \geq 0$ [25].

Definition 1 (Control Barrier Function [24]). A continuously differentiable function $B(\mathbf{s}_r)$ is a Control Barrier Function (CBF) for the system (1), if there exists a class \mathcal{K} function α s.t. $\forall \mathbf{s}_r \in \mathcal{C} :$

$$\sup_{\mathbf{a}_r \in \mathcal{A}_R} (L_{f_r}B(\mathbf{s}_r) + L_{g_r}B(\mathbf{s}_r)\mathbf{a}_r + \alpha(B(\mathbf{s}_r))) \geq 0 \quad (2)$$

where $L_{f_r}B(\mathbf{s}_r) = \frac{\partial B}{\partial \mathbf{s}_r}^\top f_r(\mathbf{s}_r)$, $L_{g_r}B(\mathbf{s}_r) = \frac{\partial B}{\partial \mathbf{s}_r}^\top g_r(\mathbf{s}_r)$ are the first order Lie derivatives of the system.

Any Lipschitz continuous controller $\mathbf{a}_r \in K_{cbf}(\mathbf{s}_r) = \{\mathbf{a}_r \in \mathcal{A}_R \mid L_{f_r}B(\mathbf{s}_r) + L_{g_r}B(\mathbf{s}_r)\mathbf{a}_r + \alpha(B(\mathbf{s}_r)) \geq 0\}$ results in a forward invariant set \mathcal{C} for the system (1).

B. Control Barrier Function Enabled Time-Based Rapidly exploring Random Tree (CBF-TB-RRT)

CBF-TB-RRT, proposed in [23], provides a probabilistic safety guaranteed solution in real-time to the start-to-goal motion planning problem. At each time step, given a probabilistic trajectory of dynamic agents, this method extracts ellipsoidal reachable sets for the agents for a given time

horizon with a bounded probability. This method extends time-based RRT (each node of TB-RRT denotes a specific state in a specific time), proposed in [26], in conjunction with CBFs to generate path segments for R (Eq. (1)) that avoid the agents’ reachable sets while moving toward goal. If the probability distribution over the dynamic agents’ future trajectory for a given finite time horizon is accurate, the generated control by CBF-TB-RRT guarantees that the probability of collision at each time step is bounded.

III. DELIBERATIVE COMMUNICATION PLANNING

We formulate the deliberative communication planning problem $\mathcal{P}_{\mathcal{DC}}$ as the problem of jointly computing communication signals with corresponding feasible motion plans for R in a social navigation scenario. As a starting point, we focus on settings with a single robot and a single human H . In such problems, R ’s actions \mathcal{A} include communication as well as movement actions. In order to model realistic scenarios, we use potentially noisy models of H ’s movement (T_H) and of H ’s sensing (O) of R ’s communications. We use these models to evaluate possible courses of action while computing efficient, collision-free communication and movement plans for R .

Intuitively, T_H maps the current state of H and H ’s belief about the possible positions of R at the next planning cycle to possible motion plans for H . We model H ’s sensor model O as a variation of the standard noisy sensor paradigm used in planning under partial observability. O relates H ’s current state, R ’s communication action and R ’s intended next state to the observation signal that H receives. In this formulation, H need not know R ’s current/intended states nor the exact communication that it executed – H only receives an observation signal. Such sensor models are very general: they can capture a variety of scenarios ranging from perfect communication to imperfect communication settings where H may not have a perfect understanding or observation of R ’s communications and may conflate R ’s communication actions with each other.

Definition 2. A deliberative communication planning problem is a tuple $\mathcal{P}_{\mathcal{DC}} = \langle \mathcal{S}, s^0, \mathcal{A}, T, \mathcal{G}, O, J \rangle$, where:

- $\mathcal{S} = \mathcal{S}_R \times \mathcal{S}_H$ is the set of states consisting of R ’s and H ’s states, respectively.
- $s^0 = s_r^0 \times s_h^0$ are the initial states of R and H , respectively, where $s_r^0 \in \mathcal{S}_R$ and $s_h^0 \in \mathcal{S}_H$.
- \mathcal{A} is the set of R ’s actions defined as $\mathcal{A} = \mathcal{A}_c \cup \mathcal{A}_m$, where \mathcal{A}_c is a set of communication signals that includes the null communication, and \mathcal{A}_m is the implicit uncountable set of R ’s feasible motion plans. Each feasible motion plan $\pi_R \in \mathcal{A}_m$ is a continuous function $\pi_R : [0, 1] \rightarrow \mathcal{S}_R$ where $\pi_R(0) = s_r^0$ and $\pi_R(1) \in \mathcal{S}_R$.
- $\mathcal{G} = \langle \mathcal{G}_R, \mathcal{G}_H \rangle$ is the goal pair where $\mathcal{G}_R \subseteq \mathcal{S}_R$ is R ’s goal set and $\mathcal{G}_H \subseteq \mathcal{S}_H$ is H ’s goal set.
- $T = \langle T_R, T_H \rangle$ constitutes transition/movement models of both agents where T_R is R ’s transition function defined as $T_R : \mathcal{S}_R \times \mathcal{A}_m \rightarrow \mathcal{S}_R$, and $T_H : \mathcal{S}_H \times \mathcal{G}_H \times \mathcal{B}_H^{R'} \rightarrow 2^{\Pi_H}$ denotes H ’s movement model where

\mathcal{B}_H^R is the set of possible beliefs over the state of R at the next planning cycle and Π_H is the set of feasible H movement plans within \mathcal{S}_H . T_H may be available as a simulator that yields a sample of the possible H plans.

- O is H 's sensor model defined as $O : \mathcal{S}_H \times \mathcal{A}_c \times \mathcal{S}_R \rightarrow \Omega$, where Ω denotes H 's observation. Situations where H cannot perfectly understand or observe R 's communication can be modeled by mapping multiple tuples (s_h, a_c, s_r) to the same $\omega \in \Omega$, where $a_c \in \mathcal{A}_c$.
- $J : \mathcal{S}_H \times \mathcal{S}_R \times \mathcal{A} \rightarrow \mathbb{R}$ is a utility function denoting the value of a joint H - R state and a communication-motion action. In practice, we express J as a cost function.

A solution to \mathcal{P}_{DC} is a sequence of communication actions and motion plans that satisfy \mathcal{G}_R , and is defined as follow.

Definition 3. A solution to the deliberative communication planning problem $\mathcal{P}_{DC} = \langle \mathcal{S}, s^0, \mathcal{A}, T, \mathcal{G}, O, J \rangle$ is a finite sequence of communication and movement actions: $\Psi = \langle (a_c^1, \pi_R^1), (a_c^2, \pi_R^2), \dots, (a_c^q, \pi_R^q) \rangle$, where $a_c^i \in \mathcal{A}_c$, $\pi_R^i \in \mathcal{A}_m$, $\pi_R^1(0) = s_r^0$, $\pi_R^i(1) = \pi_R^{i+1}(0)$, and $\pi_R^q(1) \in \mathcal{G}_R$ for $i = 1, \dots, q$.

IV. METHODOLOGY

A. Overview

In the proposed paradigm of joint communication and motion planning, a motion planner (MP) returns a set of feasible and collision-free motion plans $\Pi_R \in \mathcal{A}_m$ considering the goal set \mathcal{G} . Accordingly, a communication planner (CP) uses a search tree to select a combination of a communication action and a motion plan at each planning cycle that minimizes J . Each node of this search tree is defined by $\langle s, a_c, \pi_R \rangle$ where $s \in \mathcal{S}$, $a_c \in \mathcal{A}_c$ and $\pi_R \in \Pi_R$. a_c denotes the communication action being considered at this node while π_R denotes one of the plans returned by MP.

Fig. 2 illustrates the mechanism by which MP and CP interact. MP utilizes CBF-TB-RRT with H 's movement model and R 's dynamic model to produce a finite set of feasible motion plans $\Pi_R \subset \mathcal{A}_m$ (Sec. IV-B). Starting with a node representing the current state, CP creates a successor node for each combination of a feasible plan in Π_R and a communication action from \mathcal{A}_c . For each such combination, it uses a belief update process to compute and store an estimate of H 's next belief if R were to use the corresponding communication action. At each planning cycle, CP selects a node of tree that minimizes J (CP is described in Sec. IV-C). An important property of this approach is that our solution algorithms are independent of the choice of R , the environment, and H 's movement and sensor models.

B. CBF-based TB-RRT as MP

We obtain a set Π_R of diverse plans in Alg. 1 by employing CBF-TB-RRT [23] as MP. We modified the original CBF-TB-RRT method to better serve our hierarchical framework as follows. First, the set of possible future trajectories for H can either be given by a stochastic T_H or by a deterministic T_H with an ε tube around the predicted trajectory. We denote this region by \mathcal{S}_R^{unsafe} . MP maintains a continually updated

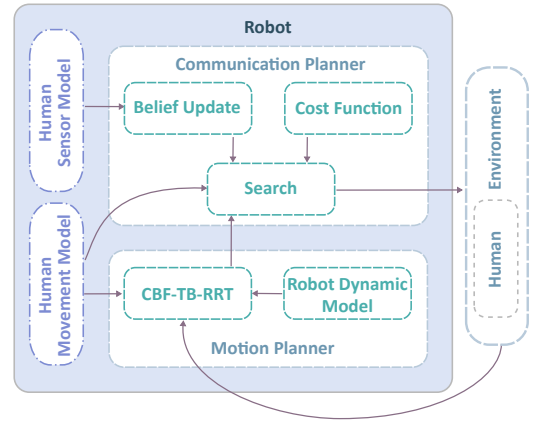


Fig. 2: An overview of our approach.

estimate of R 's safe states, $\mathcal{S}_R^{safe} = \mathcal{S}_R \setminus \mathcal{S}_R^{unsafe}$, where \mathcal{S}_R^{unsafe} would be collision-free with respect to the predicted trajectories of H . Second, the original CBF-TB-RRT expands a tree for a finite time horizon and just apply the control for the first time-step at each planning cycle. In contrast, here, we let R to execute the full returned partial plan. Finally, instead of selecting one plan to execute, we select a set of p plans $\Pi_R \subseteq \mathcal{A}_m$ of $\bar{\pi}_{R,j} : [t_0, t_j] \rightarrow \mathcal{S}_R^{safe}$ for $j = \{1, 2, \dots, p\}$. Here, each plan $\bar{\pi}_{R,j}$ represents a path segment from the initial vertex ν_0 at location $\bar{\pi}_{R,t_0}^{t_0}$ in time t_0 to another vertex ν_j at location $\bar{\pi}_{R,t_j}^{t_j}$ in time t_j . Assuming c_j to be the cost of vertex ν_j in the set of all expanded RRT vertices \mathcal{V} , we minimize the following cost to select p diverse plans $\bar{\pi}_{R,j}$ with the minimum costs c_j for $j \in \{1, 2, \dots, p\}$,

$$\begin{aligned} \min_{\mathbf{r}} \quad & J_d = \sum_{i=0}^{|\mathcal{V}|} \frac{w_c r_i c_i}{w_d \sum_{j=1, i \neq j}^{|\mathcal{V}|} r_j d_{ij}}, \\ \text{s.t.} \quad & \sum_{i=0}^{|\mathcal{V}|} r_i = p, \\ & r_i \in \{0, 1\}, \quad \text{for } i = 0, \dots, |\mathcal{V}|, \end{aligned} \quad (3)$$

where w_c and w_d are the numerator and denominator weights, respectively, d_{ij} is the Euclidean distance between vertices i and j , and \mathbf{r} is a vector of binary values r_i for $i = 1, \dots, |\mathcal{V}|$, that determines the selected plans (vertices). Given the expanded RRT at each planning cycle, we minimize (3) using Alg. 1 to find p diverse $\bar{\pi}_R$ plans.

Proposition 1. Given that RRT includes a finite set of vertices and $J_d \geq 0$, Alg. 1 terminates in a finite time.

Assumption 1. The future human motion remains within the unsafe region \mathcal{S}_R^{unsafe} predicted by T_H .

Lemma 1. Following Assn. 1, all generated path segments $\bar{\pi}_{R,j}$ for $j = 1, \dots, |\mathcal{V}|$ by CBF-TB-RRT are guaranteed to remain in \mathcal{S}_R^{safe} if $\bar{\pi}_{R,t_0}^{t_0} \in \mathcal{S}_R^{safe}$.

Proof. This proof is immediate following [23, Prop. 1]. \square

Algorithm 1: RRT Plan Π_R Generation (MP)

Input: \mathcal{V} and p
Output: Π_R

- 1 $\mathcal{P} \leftarrow$ Randomly select p vertices from \mathcal{V}
- 2 $\text{OPT_COST} \leftarrow$ Calculate the cost J_d for the vertices in \mathcal{P}
- 3 **while** *CONVERGE* **do**
- 4 **for** $\nu \in \mathcal{V} \setminus \mathcal{P}$ **do**
- 5 Calculate the cost J_d for all p -combinations of $\mathcal{P} \cup \{\nu\}$ and update OPT_COST and \mathcal{P} with the minimum cost combination
- 6 **end**
- 7 **end**
- 8 $\Pi_R \leftarrow$ Extract the path segments $\bar{\pi}_R$ from ν_0 to each p vertex in \mathcal{P}

C. Communication Planner Module (CP)

As discussed in Sec. IV-A, CP builds a search tree to select an optimal combination of communication action and motion plan. Recall that each node in the search tree consists of a state \mathbf{s} , a communication action a_c and a motion plan π_R . Here, π_R denotes the discretization of the continuous-time path segment $\bar{\pi}_R$ given by MP. We use a belief-space formulation to represent the set of locations where H might expect R to be at the next planning cycle $k+1$. Thus, the set of all possible beliefs of H , is the power set of \mathcal{S}_R . However, in practice H needs to keep track of only a subset of possible locations, in a small neighbourhood around H .

Definition 4. A δ -*local neighborhood* of H is a subset $\mathcal{L} \subseteq \mathcal{S}_R$ s.t. the Euclidean distance from S_H $d(\mathbf{s}_{xyz}, S_H)$ of R 's base coordinates \mathbf{s}_{xyz} in state \mathbf{s} is less than $\delta \forall S \in \mathcal{L}$.

We maintain a bounded, discretized set of regions to approximate H 's belief about R 's presence in their δ -local neighborhood. Let \mathcal{L}_H be the set of these discretized zones $\{l_1, \dots, l_\ell\}$. Collectively these regions can represent neighborhoods in domain-specific configurations (e.g., an H -centered forward-biased cone or a rectangular region around H with discretized cells). Given a state $(\mathbf{s}_R, \mathbf{s}_H) \in \mathcal{S}$ we use $\mathbf{s}_R \in l_i(\mathbf{s}_H)$ to express that when R 's state is \mathbf{s}_R and H 's state is \mathbf{s}_H , R will be in the region l_i in H 's local neighborhood. In this notation, H 's belief is a Boolean vector of dimension $|l_H|$, so that $b_i = 1$ in a belief \mathbf{b} represents a belief that $\mathbf{s}_R \in l_i(\mathbf{s}_H)$ is possible at the next time step.

Given a starting belief \mathbf{b}_k and an observation symbol ω_k , we can invert the sensor model and the transition function to derive a logical filtering based belief update equation for computing \mathbf{b}_{k+1} as follows. Let $\varphi_1(\mathbf{s}_R^{k+1}, i)$ state that R at \mathbf{s}_R^{k+1} would be in H 's i^{th} neighborhood zone, i.e., $\mathbf{s}_R^{k+1} \in l_i(\mathbf{s}_H^k)$; $\varphi_2(\mathbf{s}_R^k, j)$ state that b_j^k was 1 with R at \mathbf{s}_R , i.e., $b_j^k = 1 \wedge \mathbf{s}_R^k \in l_j(\mathbf{s}_H^{k-1})$; $\varphi_3(\mathbf{s}_R^k, \mathbf{s}_R^{k+1})$ state that R can move from \mathbf{s}_R^k to \mathbf{s}_R^{k+1} , i.e., $\exists a_m \in \mathcal{A}_m, T_R(\mathbf{s}_R^k, a_m) = \mathbf{s}_R^{k+1}$; and $\varphi_4(\mathbf{s}_H^k, \omega, \mathbf{s}_R^{k+1})$ state that R may have executed a communication action a_c that resulted in observation ω , i.e., $\exists a_c \in \mathcal{A}_c, o(\mathbf{s}_H^k, a_c, \mathbf{s}_R^{k+1}) = \omega$. Inverting the sensor model

and the transition function gives us $b_i^{k+1} = 1 \text{ iff } \exists \mathbf{s}_R^k, \mathbf{s}_R^{k+1} \in \mathcal{S}_R; j \in [1, \ell] : \varphi_1(\mathbf{s}_R^{k+1}, i) \wedge \varphi_2(\mathbf{s}_R^k, j) \wedge \varphi_3(\mathbf{s}_R^k, a_m, \mathbf{s}_R^{k+1}) \wedge \varphi_4(\mathbf{s}_H^k, \omega, \mathbf{s}_R^{k+1})$. CP uses this expression to compute R 's estimate of H 's belief \mathbf{b}^{k+1} given a belief \mathbf{b}^k at the parent node and the observation ω that H would receive as a result of the communication action being considered at that node. We use $\mathbf{b}(n)$ to denote this belief for node n .

CP uses a cost function J to evaluate a node $n = \langle \mathbf{s}, a_c, \pi_R \rangle$ in the search tree. Intuitively, J needs to consider H and R 's future paths Γ_H and Γ_R , respectively. $\tilde{\Gamma}_R(n)$ is an estimate for Γ_R based on π_R . However, we do not have an accurate future path for H and we use $\mathbf{b}(n)$ and the human movement model T_H to obtain an estimate $\tilde{\Gamma}_H(n)$. We omit the node argument unless required for clarity.

For computational efficiency, we discretize Γ_R and Γ_H as sequences of waypoints: $\Gamma_R = \{\gamma_R^i\}_{i=1}^{i_{max}}$ and $\Gamma_H = \{\gamma_H^i\}_{i=1}^{i_{max}}$. W.l.o.g., both sequences have the same length as the agent with the shorter path can be assumed to stay at their final location for remainder of the other agent's path execution. Let $c(\tilde{\Gamma})$ be the sum of pairwise distances between successive waypoints in a path $\tilde{\Gamma}$ and let $\delta(\tilde{\Gamma}_1, \tilde{\Gamma}_2)$ be $\delta(\tilde{\Gamma}_1, \tilde{\Gamma}_2) = \max(d_{min}(\tilde{\Gamma}_1, \tilde{\Gamma}_2) - \sigma^{safe}, 0)$, where σ^{safe} denotes the safety threshold and $d_{min}(\tilde{\Gamma}_1, \tilde{\Gamma}_2)$ is the minimum Euclidean distance between $\tilde{\Gamma}_1$ and $\tilde{\Gamma}_2$: $\min_{i=1, \dots, i_{max}} \{d(\gamma_1^i, \gamma_2^i)\}$. Besides, let $c_C(a_c)$ be the cost of executing the communication action a_c , and η_R, η_H, η_P , and η_C be the weights of the cost function. Using this notation, we define $J(n)$ as follows:

$$J(n) = \eta_R c(\tilde{\Gamma}_R(n)) + \eta_H c(\tilde{\Gamma}_H(n)) + \eta_P 1/\delta(\tilde{\Gamma}_R(n), \tilde{\Gamma}_H(n)) + \eta_C c(a_c) \quad (4)$$

In Alg. 2, at each planing iteration (lines 3-20), CP gets a library of motion plans Π_R from MP. In lines 7-11, a branch of the tree is created for each a_c and π_R . As explained in (4), the path-to-goal of H and R are required to compute a cost value for each branch. Γ_H is thoroughly given by T_H , as mentioned in line 9. On the other hand, since a π_R is likely a partial path, T_R is utilized in line 10 to compute a completed path-to-goal for R given π_R .

Fig. 3 exemplifies two branches of the CP search tree evaluated in lines 7-11 of Alg. 2. In each example, H is shown at the center of its δ -*local neighborhood* visualized as a set of nine squares around her, where the colored squares stand for \mathbf{b}_k . Besides, R is pictured at the bottom of each example with a partially expand RRT, where dark gray branches of RRT represented Π_R selected by MP. In Fig. 3(a), R communicates $a_c = \text{"Right"}$ and π_R is the right branch of RRT, emphasized by a star, which makes H believe that it will be in one of the squares on her left. In Fig. 3 (b), R goes forward and communicates "Forward" as well, which makes H believe that R will be in one of the middle squares. In scenario (b), R takes a shorter path to goal but scenario (a) results in less conflicting paths for both R and H . Thus, the more optimal branch will be determined based on the weights of J in (4).

Algorithm 2: Communication Planner

Input: \mathcal{P}_{DC}
Output: Ψ

- 1 initialize: \mathbf{b}_0 and \mathcal{S}^0
- 2 **while** GOAL_TEST($\mathcal{G}_R, \mathcal{S}_R$) == FALSE **do**
- 3 $\Pi_R \leftarrow$ get the plans from the MP
- 4 MIN_COST $\leftarrow \infty$
- 5 **for** $\pi_R \in \Pi_R$ **do**
- 6 **for** $a_c \in \mathcal{A}_c$ **do**
- 7 $\omega_{k+1} \leftarrow O(a_c, \mathcal{S}^k)$
- 8 $\mathbf{b}_{k+1} \leftarrow$ UPDATE($\mathbf{b}_k, \omega_{k+1}$)
- 9 $\tilde{\Gamma}_H \leftarrow T_H(\mathcal{S}_H^k, \mathcal{G}_H, \mathbf{b}_{k+1})$
- 10 $\tilde{\Gamma}_R \leftarrow T_R(\mathcal{S}_R, \pi_R)$
- 11 $c_{branch} \leftarrow J(\tilde{\Gamma}_R, \tilde{\Gamma}_H, a_c)$
- 12 **if** $c_{branch} < \text{MIN_COST}$ **then**
- 13 MIN_COST $\leftarrow c_{branch}$
- 14 BEST_ACTION $\leftarrow \langle \pi_R, a_c \rangle$
- 15 **end**
- 16 **end**
- 17 **end**
- 18 EXECUTE(BEST_ACTION)
- 19 $\mathcal{S}^k \leftarrow \mathcal{S}^{k+1}$
- 20 Ψ .APPEND(BEST_ACTION)
- 21 **end**

Assumption 2. The predicted trajectories Γ_H given by T_H in the discretized domain is an over-approximation of the predicted trajectories by T_H in the continuous domain.

Assumption 3. The discretized projection of $\bar{\pi}_R$ on Γ_R (π_R in discretized domain) is an over-approximation of $\bar{\pi}_R$ in the continuous domain.

Theorem 1. Let $P_{DC} = \langle \mathcal{S}, s^0, \mathcal{A}, T, \mathcal{G}, O, J \rangle$ be a deliberative communication problem and let $\Psi^* = \langle (a_c^i, \pi_R^i) \rangle_{i=1}^q$ be its solution computed by Alg.2 using the cost function J in (4). Let Γ_R be the discretized waypoints of R in Ψ^* defined as $\Gamma_R = \langle \pi_R^i \rangle_i$, and Γ_H be a corresponding discretized waypoint sequence of a trajectory for H predicted by T_H and starting at s^0 with the goal G_H . If Assn. 1-3 hold, Γ_R will either lie within $\bar{\mathcal{S}}_R^{safe}$ or it will satisfy $d_{min}(\Gamma_R, \Gamma_H) > \sigma^{safe}$.

Proof. Since R has a null communication action that does not alter H 's belief, Alg. 2 will always have a node reflecting the default behavior of CBF-TB-RRT with cost $< \infty$. In this case, Lemma 1 guarantees R 's trajectory not to leave $\bar{\mathcal{S}}_R^{safe}$. If Assn. 2 and 3 hold, and if Alg. 2 selects a node other than the default CBF-TB-RRT behavior, the min distance will be at least σ^{safe} , otherwise $\forall \eta_P > 0$, J would be ∞ and the default CBF-TB-RRT behavior will be selected. \square

V. EMPIRICAL EVALUATION

We conducted extensive experiments in various simulation environments to evaluate the proposed method. These experiments 1) draw a comparison between the proposed method




Fig. 3: Two examples of the reasoning procedure of CP for a branch of the search tree.

and the baseline method CBF-TB-RRT, and 2) illustrate the performance of the proposed method in deadlock situations.

A. Implementation

1) *CBF-TB-RRT Design:* In our implementation, we consider the nonholonomic unicycle model for R dynamics as

$$\dot{\mathbf{s}}_r = \mathbf{g}_r(\mathbf{s}_r) \mathbf{a}_r = \begin{bmatrix} \cos(\theta_r) & 0 \\ \sin(\theta_r) & 0 \\ 0 & 1 \end{bmatrix} \mathbf{a}_r. \quad (5)$$

where states are $\mathbf{s}_r = [x_r, y_r, \theta_r]^T \in \mathcal{S}_R \subseteq \mathbb{R}^2 \times [-\pi, \pi)$ and control inputs are $\mathbf{a}_r = [v_r, \omega_r]^T \in \mathcal{A}_R \subseteq \mathbb{R}^2$. The parameters x_r, y_r, θ_r denote the longitudinal and lateral positions of R and heading angle, respectively. The controls v_r and ω_r also represent the linear and angular velocities of R , respectively. Moreover, the goal set $\mathcal{S}_g \subset \mathcal{S}_R$ of R can describe a set of position states in \mathbb{R}^2 as follows

$$\mathcal{S}_g = \{ \mathbf{s}_r \in \mathcal{S}_R \mid \| [x_r, y_r]^T - \mathbf{s}_g \|_2^2 - r_g^2 \leq 0 \}, \quad (6)$$

where $\|\cdot\|_2$ denotes the Euclidean norm, $\mathbf{s}_g = [x_g, y_g]^T$ is the center, and r_g is the radius of the goal set.

While expanding the RRT tree, the following cost c_i is assigned to each vertex $v_i \in \mathcal{V}$ for $i = 0, 1, \dots, |\mathcal{V}|$,

$$c_i = w_d^G c_d^G + w_d^H c_d^H + w_g c_g + w_t c_t, \quad (7)$$

where c_d^G is the Euclidean distance between vertex i and the goal point, c_d^H is the Euclidean distance between vertex i and H , c_h is the heading cost, and c_t is the trap cost. The heading cost c_g calculates the angular difference between the sampled vertex heading and the heading toward goal. To calculate the trap cost c_t , the algorithm checks the waypoints of a discretized direct straight line from the sampled vertex to the goal point. The trap cost c_t is then the number of waypoints lied within the occupied regions. Readers are referred to [23] for further details on CBF-TB-RRT tree expansion. w_d^G, w_d^H, w_g , and w_t are weight terms.




Fig. 4: An example of a potential deadlock in confined environments.

2) *Human Movement Model*: We assumed H 's movement is described by a deterministic kinematic motion transition function (T_H) and we used the Dynamic Window Approach (DWA) in MP, proposed in [27], to predict H 's shortest trajectory to the goal for a finite time horizon. Since DWA is a deterministic prediction method, we assumed an ε bound around the human's predicted trajectory following Assn. 1 to derive the CBF safety constraints. Given the human's predicted trajectory s_h , we define the safe set $\mathcal{S}_R^{safe} \subseteq \mathcal{S}_R$ as $\mathcal{S}_R^{safe} = \{s_r \in \mathcal{S}_R, s_h \in \mathcal{S}_H \mid B(s_r, s_h) \geq 0\}$, where $B(s_r)$ is a continuously differentiable safety measure defined as

$$B(s_r, s_h) = \|[x_r, y_r]^T - s_h\|_2^2 - (\varepsilon + r_h + r_r)^2, \quad (8)$$

r_h , and r_r are the radii of human and robot, respectively. The safety measure $B(s_r)$ is employed as a CBF to impose the safety constraint (2) on the control input \mathbf{a}_r in a Quadratic Program (QP) to generate safe plans π_R [23].

As illustrated in Sec. IV-A, CP also utilizes T_H to predict a trajectory-to-goal for H for each branch of the search tree. Besides, in contrast to the requirements of the motion planning module, H movement prediction must be provided for the whole horizon in communication planning module. Therefore, for the sake of computational efficiency, CP utilizes another H movement model rather than DWA. CP considers a grid-based abstraction of the environment and utilizes A* search algorithm to predict a path-to-goal for H . In general this abstraction could be derived using methods for automatically predicting reliable state and action abstractions such as [28].

Assumption 4. Predictions drawn from A* and DWA approaches complied with the Assn. 2 in all our experiments.

3) *Human Motion Execution Model*: We utilized the Social Forces model [29] to simulate the human movement, as it is very fast, scalable, and yet describes observed pedestrian behaviors realistically. We modeled H and R both as pedestrians. To mimic H 's reactivity to R 's communication action a_c , the model creates multiple virtual agents moving

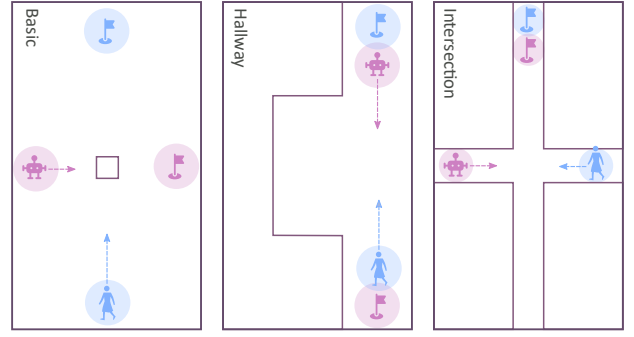


Fig. 5: Schematic illustration of diversified test environments that capture various conflicting situation.

from R 's current position to all x - y projections of discretized zones $l_i \in l_H$ in the CP's belief model for which $b_i = 1$. If $\mathbf{b}_k = \emptyset$, R 's goal is computed as a linear projection from its current position based on its current velocity, i.e. H makes no assumptions over R 's future trajectory. Thus, in our experiments, the models used by H are different from the model H used by R , which is likely in real-world setting.

B. Experimental Setup

Test environments: Fig. 5 the environments used in our experiments. The basic floor map exemplifies spacious environments, while the hallway and intersection floor maps model more restricted and confined environments.

Measurements: Aside from cost-to-goal of R and H , there are four more quantitative measures to evaluate the performance and effectiveness of the proposed method:

- R 's normalized speed (RNS): $RNS = c_R^*/time_R^{actual}$ measures R 's normalized mean speed from s_r^0 to \mathcal{G}_R , where c_R^* and $time_R^{actual}$ denote the optimal cost-to-goal of R and R 's actual travel time respectively.
- H 's normalized speed (HNS): $HNS = c_H^*/time_H^{actual}$ measures H 's normalized average speed from s_h^0 to \mathcal{G}_H , where c_H^* and $time_H^{actual}$ denote the optimal cost-to-goal of H and H 's actual travel time respectively.
- Planning iterations (PI): PI denotes the number of iterations of lines 2 to 17 in Alg. 2.
- Proximity cost (PC): PC measures the closeness of R and H during an experiments. Let $\Gamma_R = \{\gamma_R^i\}_{i=1}^{i_{max}}$ be R 's discretized trajectories given by a solution Ψ and $\Gamma_H = \{\gamma_H^i\}_{i=1}^{i_{max}}$ be the corresponding discretized waypoint sequence of an actual trajectory for H . We defined PC using (8) as follows.

$$Z = \{\zeta_i \mid \zeta_i = B(\gamma_R^i, \gamma_H^i) < thresh\}_{i=1}^{i_{max}} \quad (9)$$

$$PC = \begin{cases} \infty & \text{if } \exists \zeta_i \in Z, \zeta_i < 0 \\ 1/\sum_{i=1}^{i_{max}} \zeta_i & \text{otherwise} \end{cases}, \quad (10)$$

Hypotheses: throughout the experiments, we evaluate the following hypotheses 1) In confined environments, the chances of a deadlock are higher. Therefore, the effect of communication to avoid such deadlocks is more effective. 2) The proposed deliberative communication approach not only

TABLE I: Comparison with CBF-TB-RRT.

Measures Maps	Our approach				CBF-TB-RRT			
	R cost-to-goal	H cost-to-goal	PI	PC	R cost-to-goal	H cost-to-goal	PI	PC
Basic	5.65–5.68	7.33–7.52	2–2	0.50–0.53	5.51–6.27	6.90–6.99	46–113	0.21–0.57
Intersection	3.63–3.88	6.10–6.29	2–2	0.22–0.24	4.20–4.24	5.76–5.90	51–98	0.34– ∞
Hallway	10.12–10.58	6.85–7.39	4–4	0.72–0.89	10.27–10.30	6.65–6.77	121–123	∞ – ∞

The results show the range of the measurements in 10 trials per map; PI : planning iterations; PC : proximity cost.




Fig. 6: Flexible prioritization of H and R in different test environments, where $F = 1$ prioritizes the robot.

results in less conflicting social navigation, but also prevents deadlock situations where non-communicative approaches fail to find a solution. 3) By adjusting the weight vector of the cost function J , H or R can be prioritized. Accordingly, the non-prioritized agent is expected to have a decreased normalized average speed due to an increased cost-to-goal.

C. Results

1) *Comparison with CBF-RRT:* In this section, we aim to demonstrate that the proposed method performs as optimally as CBF-TB-RRT, in terms of the traveled distances, while it reduces the conflict between H and R . In Table I, the results are presented as the range of 10 experiments the experiments for each test environments of Fig. 5, where $\eta_R = 1.5, \eta_H = 0.25, \eta_P = 3, \eta_C = 1$, and $A_c = \{north, south, east, west\}$.

Our results show that PC of the baseline drastically increases in more confined environments. E.g., PC has a finite range in the basic environment since the room is spacious, while the PC range is infinity in the intersection environment where the floor map is confined and only one agent can pass through a corridor at a time. The situation is even more severe in the hallway environment in which the baseline method results in an infinite PC for all 10 experiments. These observations validate Hypothesis 1. In contrast, the proposed method handles conflicting situations of the intersection and hallway environments effectively. The PC values of our method in all environments are dramatically lower compared to the baseline method, while cost-to-goal of R and H do not increase noticeably.

Moreover, employing the proposed method eliminates the necessity for frequent re-planning as PI drops significantly compared to the experiments with the baseline method.

2) *Handling potential deadlocks:* According to V-C.1, the proposed method is significantly more effective in reducing PC in confined environments while maintaining

the efficiency in terms of c_R and c_H . This property is particularly imperative in preventing potential deadlocks in narrow passages, where a lower PC implies less conflicting path for H and R . Fig. 4 demonstrates a pervasive case where lack of communication leads to a freezing situation. In this example, at the first planning iteration, R transmits an “east” signal, selected automatically by CP, to H by which H is informed about R ’s plan before she enters the narrow corridor. As shown in Fig. 4 (top left), this communication signal updates H ’s belief about R ’s next location adequately and impels H to clear the passage. At the second planning iteration, R has already passed through the intersection, so it remains silent and H ’s belief indicates no collisions, as depicted in Fig. 4 (bottom left).

In the same scenario, the baseline method performs ineffectively since H enters the left corridor before R departs it. When H gets closer to R , there won’t be enough room for the RRT to be expanded and a deadlock happens since the passage will be blocked for R permanently. This analysis supports Hypothesis 2 regarding the capability of the proposed method to handle potential deadlocks.

3) *Flexible prioritization:* H or R can be prioritized flexibly by adjusting the weights of J . A parameter study on η_R and η_H reveals the way that each agent is favored in different social navigation scenarios, as shown in Fig. 6. In these experiments, the weights are adjusted as $\eta_R = F\eta_{const}$, and $\eta_H = (1 - F)\eta_{const}$, where $F \in [0, 1]$ denotes the priority factor (R is fully prioritized for $F = 1$), and $\eta_{const} = 1.5$. In all three environments, prioritizing R increases R ’s normalized speed significantly. Fig. 6 shows that in the basic environment, R ’s normalized speed increases by 2.7 times when R is prioritized, compared to the case where H is highly prioritized. Likewise, H speeds up when she is prioritized in the basic and intersection environment. However, in the hallway environment, the whole H - R in-

teraction is relatively smoother and less conflicting when R has a higher priority. Together, the present findings support Hypothesis 3. Furthermore, the results support the fact that the proposed method maintains a reasonably low PC in all test environments not matter which agent is prioritized. In other words, the proposed method can be used to identify appropriate priorities for smooth social navigation.

VI. CONCLUSION

This paper proposes a joint communication and motion planning framework that selects from an arbitrary input set of communication signals while computing the robot motion plans. The simulation results demonstrated that the presented framework avoids potential deadlocks in confined environments by leveraging explicit communications coupled with robot motion plans. We found that producing less conflicting trajectories for the robot in confined environments, which led to drastically lower proximity costs, indicates lower chances of a deadlock. We also observed that the proposed method does not degrade the robot's efficiency (in terms of traveled distances) compared to CBF-TB-RRT. In contrast, the non-communicative baseline method resulted in high proximity cost overall, which shows its incapability of generating viable solutions when extensive human-robot interaction is required. Furthermore, the proposed method can flexibly prioritize either the robot or the human while maintaining its effectiveness in handling potential deadlocks.

REFERENCES

- [1] J. Cheng, H. Cheng, M. Q.-H. Meng, and H. Zhang, "Autonomous navigation by mobile robots in human environments: A survey," in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pp. 1981–1986, IEEE, 2018.
- [2] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1343–1350, IEEE, 2017.
- [3] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard, "Feature-based prediction of trajectories for socially compliant navigation," in *Robotics: science and systems*, 2012.
- [4] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: the case for cooperation," in *2013 IEEE international conference on robotics and automation*, pp. 2153–2160, IEEE, 2013.
- [5] L. P. Kaelbling and T. Lozano-Pérez, "Hierarchical planning in the now," in *Workshops at the Twenty-Fourth AAAI Conference on Artificial Intelligence*, 2010.
- [6] C. R. Garrett, T. Lozano-Pérez, and L. P. Kaelbling, "Pddlstream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning," in *Proceedings of the International Conference on Automated Planning and Scheduling*, vol. 30, pp. 440–448, 2020.
- [7] S. Srivastava, E. Fang, L. Riano, R. Chitnis, S. Russell, and P. Abbeel, "Combined task and motion planning through an extensible planner-independent interface layer," in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 639–646, IEEE, 2014.
- [8] N. Shah, D. K. Vasudevan, K. Kumar, P. Kamojihala, and S. Srivastava, "Anytime integrated task and motion policies for stochastic environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 9285–9291, IEEE, 2020.
- [9] N. T. Dantam, Z. K. Kingston, S. Chaudhuri, and L. E. Kavraki, "An incremental constraint-based framework for task and motion planning," *The International Journal of Robotics Research*, vol. 37, no. 10, pp. 1134–1151, 2018.
- [10] N. T. Dantam, S. Chaudhuri, and L. E. Kavraki, "The task-motion kit: An open source, general-purpose task and motion-planning framework," *IEEE Robotics & Automation Magazine*, vol. 25, no. 3, pp. 61–70, 2018.
- [11] I. A. Sucas, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [12] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2, pp. 995–1001, IEEE, 2000.
- [13] H. Nishimura, B. Ivanovic, A. Gaidon, M. Pavone, and M. Schwager, "Risk-sensitive sequential action control with multi-modal human trajectory forecasting for safe crowd-robot interaction," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 11205–11212, IEEE, 2020.
- [14] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [15] A. Kulkarni, S. Srivastava, and S. Kambhampati, "A unified framework for planning in adversarial and cooperative environments," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, pp. 2479–2487, 2019.
- [16] Y. Zhang, D. A. Shell, and J. M. O'Kane, "Finding plans subject to stipulations on what information they divulge," in *International Workshop on the Algorithmic Foundations of Robotics*, pp. 106–124, Springer, 2018.
- [17] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [18] H. Kivrak, F. Cakmak, H. Kose, and S. Yavuz, "Social navigation framework for assistive robots in human inhabited unknown environments," *Engineering Science and Technology, an International Journal*, vol. 24, no. 2, pp. 284–298, 2021.
- [19] R. A. Knepper, C. I. Mavrogiannis, J. Proft, and C. Liang, "Implicit communication in a joint action," in *Proceedings of the 2017 acm/ieee international conference on human-robot interaction*, pp. 283–292, 2017.
- [20] S. Habibian, A. Jonnavittula, and D. P. Losey, "Here's what i've learned: Asking questions that reveal reward learning," *arXiv preprint arXiv:2107.01995*, 2021.
- [21] K. Baraka and M. M. Veloso, "Mobile service robot state revealing through expressive lights: formalism, design, and evaluation," *International Journal of Social Robotics*, vol. 10, no. 1, pp. 65–92, 2018.
- [22] Y. Che, A. M. Okamura, and D. Sadigh, "Efficient and trustworthy social navigation via explicit and implicit robot–human communication," *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 692–707, 2020.
- [23] K. Majd, S. Yaghoubi, T. Yamaguchi, B. Hoxha, D. Prokhorov, and G. Fainekos, "Safe navigation in human occupied environments using sampling and control barrier functions," *arXiv preprint arXiv:2105.01204*, 2021.
- [24] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," *European Control Conference (ECC)*, 2019.
- [25] F. Blanchini, "Set invariance in control," *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [26] A. Sintov and A. Shapiro, "Time-based rrt algorithm for rendezvous planning of two dynamic systems," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6745–6750, IEEE, 2014.
- [27] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [28] N. Shah and S. Srivastava, "Using deep learning to bootstrap abstractions for hierarchical robot planning," *arXiv preprint arXiv:2202.00907*, 2022.
- [29] D. Helbing and P. Molnar, "Social Force Model for Pedestrian Dynamics," *Physical Review E*, vol. 51, pp. 4282–4286, May 1995.