

SM²ITH: Safe Mobile Manipulation with Interactive Human Prediction via Task-Hierarchical Bilevel Model Predictive Control

Francesco D’Orazio^{*1}, Sepehr Samavi^{*2,3}, Xintong Du^{*2}, Siqi Zhou^{3,4}, Giuseppe Oriolo¹, Angela P. Schoellig³

Abstract—Mobile manipulators are designed to perform complex sequences of navigation and manipulation tasks in human-centered environments. While recent optimization-based methods such as Hierarchical Task Model Predictive Control (HTMPC) enable efficient multitask execution with strict task priorities, they have so far been applied mainly to static or structured scenarios. Extending these approaches to dynamic human-centered environments requires predictive models that capture how humans react to the actions of the robot. This work introduces Safe Mobile Manipulation with Interactive Human Prediction via Task-Hierarchical Bilevel Model Predictive Control (SM²ITH), a unified framework that combines HTMPC with interactive human motion prediction through bilevel optimization that jointly accounts for robot and human dynamics. The framework is validated on two different mobile manipulators, the Stretch 3 and the Ridgeback–UR10, across three experimental settings: (i) delivery tasks with different navigation and manipulation priorities, (ii) sequential pick-and-place tasks with different human motion prediction models, and (iii) interactions involving adversarial human behavior. Our results highlight how interactive prediction enables safe and efficient coordination, outperforming baselines that rely on weighted objectives or open-loop human models. Code: <https://github.com/learnsyslab/sm2ith.git>

I. INTRODUCTION

Mobile manipulators, which combine a mobile base with one or more robotic arms, are envisioned as versatile service robots capable of performing complex tasks in human-centered environments. For example, a robot waiter may need to retrieve a tray of dishes from the kitchen, transport them across a crowded dining hall while ensuring no spillage, and deliver them to the right guests. Such scenarios reveal two interconnected challenges. First, the robot needs to seamlessly coordinate multiple manipulation and navigation tasks with different priorities. Second, the robot needs to simultaneously negotiate safe (i.e., collision-free) and efficient navigation among a crowd of humans.

^{*}equal contribution.

¹Department of Computer, Control and Management Engineering, of Sapienza University of Rome, Italy.

²University of Toronto Institute for Aerospace Studies (UTIAS) and the Vector Institute for Artificial Intelligence, Canada.

³Learning Systems and Robotics lab at the Technical University of Munich and the Munich Institute for Robotics and Machine Intelligence (MIRMI), Germany.

⁴School of Computing Science, Faculty of Applied Sciences, Simon Fraser University, Burnaby, BC, Canada.

Emails: {dorazio, oriolo}@diag.uniroma1.it, {sepehr, xintong.du}@robotics.utias.utoronto.ca, angela.schoellig@tum.de, siqi@sfu.ca.

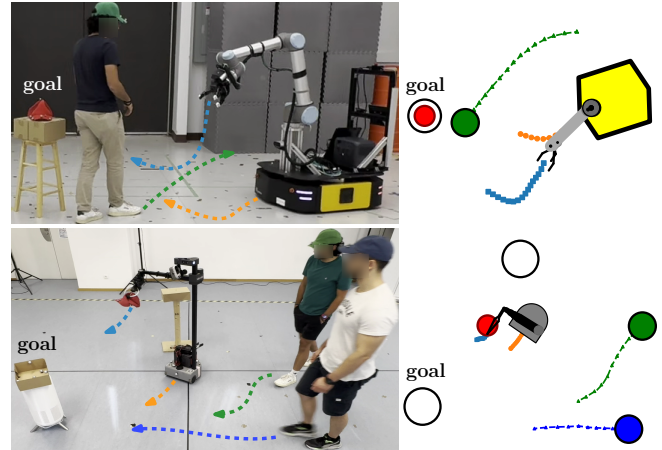


Fig. 1: Snapshots of pick-and-place tasks in human-centered environments with the Ridgeback–UR10 (top) and Stretch 3 (bottom) platforms along with a top-down view animation. The illustrated SM²ITH solutions consist of planned robot trajectories coupled with human predictions allowing the robot to interactively navigate among humans while coordinating multiple prioritized mobile manipulation tasks. Video: <http://tiny.cc/sm2ith>.

Towards addressing the first challenge, early approaches have been specialized toward specific notions of prioritization. To execute a temporal sequence of tasks (e.g., retrieve dish before going to guest) some methods propose decoupling base and arm planning to execute one task at a time [1], which can lead to slow execution as the robot fails to take advantage of kinematic redundancy. To simultaneously execute tasks of varying importance (e.g., avoid spilling even at the expense of navigation progress), other methods propose using weighted sums of task costs [2], which relies on precise tuning to handle tradeoffs between contradictory objectives. However, in such redundant systems, prioritization can be formulated as a strict hierarchy. Thus, recent methods, such as Hierarchical Task Model Predictive Control (HTMPC) [3], propose using lexicographic optimization [4], where the highest priority task is first optimized in isolation and lower priority tasks are consecutively included without jeopardizing higher priority tasks. Yet, the focus of the aforementioned methods has not been on human-centered environments where safe navigation among humans is necessary.

To address the second challenge, the robot must anticipate human movement and adjust its trajectory in real-time. Some methods propose to reactively avoid dynamic obstacles by tracking occupancy grids over time [5] or using a geometric local motion planner [6] to avoid humans. While these

methods show promise in executing collision-free mobile manipulation tasks, they neglect the fact that the motions of the humans and the robot influence each other. As such, a variety of work has focused on generating human trajectory predictions in open-loop (e.g., [7]) and using them as an input in the controller [8]. However, this approach has been shown to lead to overly conservative or even frozen behaviors in crowded scenes [9]. Thus, more recent approaches propose to model the interaction between the humans and the robot by integrating the prediction model into the robot controller and incorporating collision avoidance as a cost term [10] or as constraints [11]. However, to the best of the authors’ knowledge, such interactive formulations have not been used in mobile manipulation, only mobile navigation.

In this work, we propose Safe Mobile Manipulation with Interactive Human Prediction via Task-Hierarchical Bilevel Model Predictive Control (SM²ITH), a joint prediction and control framework that simultaneously predicts human trajectories and generates robot commands. Our approach builds on HTMPC [3] and Safe and Interactive Crowd Navigation (SICNav) [11] to enable a mobile manipulator to execute multiple prioritized base and manipulation tasks [3] while reasoning about, and accounting for, human responses in the shared space. SM²ITH models humans to be solving the Optimal Reciprocal Collision Avoidance (ORCA) [12] optimization problem, and embeds this model as constraints within HTMPC resulting in a bilevel problem [11]. This way, the predictions used for robot planning are closed-loop: our algorithm couples robot commands and associated human predictions by assuming that the humans solve ORCA in reaction to robot motion. To ensure that robot motions remain in the set of safe collision-free states, we introduce Control Barrier Functions (CBFs) [13] that limit how quickly the robot approaches unsafe states. Fig. 1 illustrates the jointly planned trajectory for a pick-and-place application using two different mobile manipulator platforms.

Our contributions are twofold. First, we formulate SM²ITH, a hierarchical task mobile manipulation control framework with interactive human predictions. Specifically, we (i) extend SICNav’s collision avoidance constraints to a 3D setting for whole-body (base and arm) avoidance, (ii) integrate SICNav’s bilevel structure into HTMPC’s lexicographic formulation to simultaneously predict human motion while solving for robot actions, and (iii) introduce CBFs to proactively satisfy safety constraints (i.e., collision avoidance) in human-centered environments. Second, we conduct extensive experiments (80 runs) to evaluate the performance of SM²ITH compared to methods using baseline models of human motion in sequential pick-and-place tasks using two different mobile manipulator platforms illustrated in Fig. 1. We also conduct further experiments (60 runs) to evaluate the performance of SM²ITH against a weighted-sum baseline controller on simultaneous mobile manipulation tasks with different priorities. To the best of the authors’ knowledge, this is the first whole-body control method for mobile manipulators in human-centered environments that couples human motion predictions with robot actions to

explicitly model interactions during collision avoidance.

II. RELATED WORK

A. Mobile Manipulation with Task Prioritization

Early motion planning methods for mobile manipulation relied on decoupled task-space planners for the base and end-effector, executed sequentially [1] or through loosely coordinated controllers [14]. While reactive, these approaches treated tasks in isolation and could not leverage redundancy effectively. Weighted-sum formulations [2] allowed simultaneous execution but requires fine tuning to trade off contradictory objectives, while whole-body joint-space planners [15], [16] improved coordination at the expense of heavy computation and limited reactivity.

Hierarchical-task control resolves these issues by enforcing strict priorities, first through inverse kinematics and null-space methods [17] and later through hierarchical quadratic programming [18]–[20]. Embedding such hierarchies within predictive control has led to HTMPC, which enforces lexicographic priorities over a horizon and has been validated on mobile manipulators [3] and redundant arms [21].

While these methods have demonstrated efficient multi-task execution, their applications are focused on static or structured settings. Mobile manipulation in human-centered environments remains an open challenge: here, the robot must not only resolve multiple objectives but also account for the fact that its own behavior can affect human motion, and vice versa. This motivates the integration of HTMPC with predictive models that capture human–robot interaction.

B. Dynamic Obstacle and Human Avoidance

One straightforward strategy for human avoidance is to model them as dynamic obstacles and leverage standard safe control frameworks to enforce collision avoidance constraints. For instance, dynamic obstacle avoidance has been integrated into predictive control through whole-body Model Predictive Control (MPC) with signed-distance constraints [22]–[25], perceptive MPC for high degrees of freedom (DOF) robots [26], and CBFs for safety filtering [27]–[29]. Uncertainty-aware approaches extend these methods via chance constraints [30], and occlusion-aware MPC [31]. When extending these methods for humans in shared spaces, it is often the case that human motion is predicted in an open-loop manner (e.g., constant velocity) [32].

For human interaction, however, simple predictors combined with collision constraints are insufficient. In particular, the absence of explicit modeling of interaction (i.e., the fact that humans and robots adapt to each other’s behavior in shared space) can result in deadlocks or conservative behaviors. To address this issue, richer models such as the social force model [33] and Reciprocal Velocity Obstacle (RVO) frameworks [34] have been proposed to capture the reciprocal adaptation behavior. Moreover, recent extensions formulate joint human–robot prediction as optimal control problems [11], [12]. These advances improve realism in navigation but remain focused on mobile robots without arms and base-level motion. Integration with multitask,

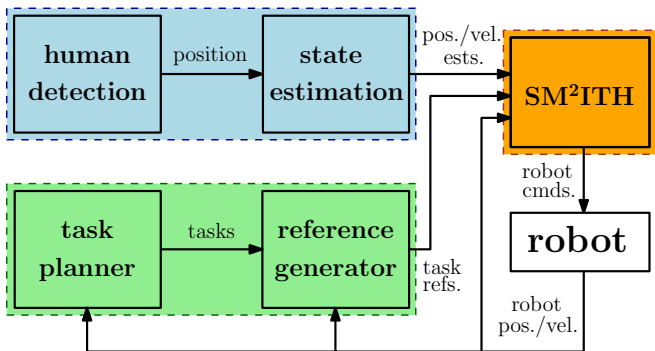


Fig. 2: Proposed control architecture composed of three main blocks. The motion generation block (orange) computes the control commands using a hierarchical predictive formulation that accounts for both robot tasks and predicted human motion. The planning block (green) organizes the tasks into a priority hierarchy and generates the corresponding references to be tracked. The human estimation block (blue) detects humans in the environment and estimates their motion, providing predictions that feed into the control loop.

redundancy-aware control for mobile manipulators is still underexplored. This work addresses that gap by proposing the SM²ITH framework, embedding interactive prediction within the HTMPC formulation for mobile manipulation tasks.

III. PROBLEM FORMULATION

Consider a mobile manipulator operating in a three-dimensional workspace, \mathcal{W} , shared with humans. We denote the robot state as, $\mathbf{x}^{(r)} = (\mathbf{q}^{(r)}, \boldsymbol{\nu}^{(r)})$, where $\mathbf{q}^{(r)} \in \mathbb{R}^n$ is the configuration vector that includes both the base variables $\mathbf{q}_b^{(r)}$ and arm variables $\mathbf{q}_a^{(r)}$. The pseudo-velocities, $\boldsymbol{\nu}^{(r)} \in \mathbb{R}^m$, are given by $\dot{\mathbf{q}}^{(r)} = \mathbf{G}(\mathbf{q}^{(r)})\boldsymbol{\nu}^{(r)}$, where \mathbf{G} is a block diagonal matrix consisting of a state-dependent component for the base and an identity matrix for the arm variables [35]. The input of the robot is the time derivative of the pseudo-velocity, $\mathbf{u} = \dot{\boldsymbol{\nu}}^{(r)}$.

The workspace \mathcal{W} contains n_h humans indexed by $j \in \{1, \dots, n_h\}$ and n_o static obstacles indexed by $\tilde{l} \in \{n_h + 1, \dots, n_h + n_o\}$. We define a global system state that concatenates the robot state and the states of all humans, $\mathbf{x} = (\mathbf{x}^{(r)}, \mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n_h)})$. Each human state is defined by the agent's Cartesian position and velocity, $\mathbf{x}^{(j)} \in \mathbb{R}^4$, and human dynamics are modeled as a kinematic integrator.

The objective of our proposed controller is to generate real-time robot actions such that

- (i) collisions with humans and static obstacles are avoided at all times,
- (ii) the motions remain kinematically feasible (i.e., within the kinodynamic limits of the robot), and
- (iii) the robot executes multiple tasks for both base and end effector, with a strict prioritization, while minimizing violation of the task costs.

IV. PROPOSED APPROACH

To address the problem stated in Section III, we propose the architecture illustrated in Fig. 2. The SM²ITH block (orange), the primary contribution of this paper, is responsible for computing the robot commands accounting

for future human trajectories. Our method builds upon the HTMPC control scheme [3] allowing to solve prioritized tasks efficiently, which are defined as a lexicographic optimization problem [4], and embeds an optimal model for human motion as lower-level constraints resulting in a bilevel MPC optimization problem [11].

The planning block (green in Fig. 2) provides the references that the robot must follow to execute the higher priority tasks at each time step. The human state estimation block (blue in Fig. 2) gives estimates of human positions and velocities through Kalman Filters.

A. Hierarchical Task Model Predictive Control (HTMPC)

To solve a set of tasks ordered by strict prioritization, $[\mathcal{J}_1(\mathbf{x}), \mathcal{J}_2(\mathbf{x}), \dots, \mathcal{J}_L(\mathbf{x})]$, we use the HTMPC framework [3]. At each control step t , the HTMPC controller computes the optimal state and input trajectories $\bar{\mathbf{x}}^{(r)*}, \bar{\mathbf{u}}^*$ over a finite time horizon N for a given sampling time Δt , by solving the following lexicographic optimization problem:

$$\text{lex min}_{\bar{\mathbf{x}}^{(r)}, \bar{\mathbf{u}}} [\mathcal{J}_1, \mathcal{J}_2, \dots, \mathcal{J}_L] \quad (1a)$$

$$\text{s.t. } \bar{\mathbf{x}}_{k+1}^{(r)} = \mathbf{f}(\bar{\mathbf{x}}_k^{(r)}, \bar{\mathbf{u}}_k) \quad (1b)$$

$$\bar{\mathbf{x}}_k^{(r)} \in \mathcal{X}, \bar{\mathbf{u}}_k \in \mathcal{U}, \forall k = 0, \dots, N-1 \quad (1c)$$

$$\bar{\mathbf{x}}_0^{(r)} = \mathbf{x}_0^{(r)} \quad (1d)$$

$$\text{sd}(\mathbf{x}_k^{(r)}) \geq \mathbf{d}_{sc}, \quad (1e)$$

where $\bar{\mathbf{x}}_0^{(r)}$ is the current robot state, and the prediction index k corresponds to time $t + k\Delta t$. The optimization enforces robot dynamics (1b), state and input limits (1c), initial condition (1d), and self-collision avoidance (1e) defined as a signed distance between robot links. Each task cost, \mathcal{J}_ℓ , represents the accumulated tracking error over the horizon, $\mathcal{J}_\ell = \frac{1}{2} \sum_{k=0}^{N-1} \|\mathbf{e}_{\ell k}\|_{\mathbf{W}_k}^2 + \|\mathbf{e}_{\ell N}\|_{\mathbf{W}_N}^2$, with $\mathbf{e} \in \mathbb{R}^s$, $\|\mathbf{e}\|_{\mathbf{W}}^2 = \mathbf{e}^T \mathbf{W} \mathbf{e}$ and $\mathbf{W} \in \mathbb{R}_+^{s \times s}$.

To solve (1), we break down the problem into a sequence of Single Task Model Predictive Control (STMPC) problems, one per task:

$$\min_{\bar{\mathbf{x}}^{(r)}, \bar{\mathbf{u}}} \mathcal{J}_\ell + \mathcal{J}_{\text{eff}} \quad (2a)$$

$$\text{s.t. (1b), (1c), (1d), (1e)} \quad (2b)$$

$$|e_{kp}| \leq |e_{kp}^*|, \quad (2c)$$

$$\forall i = 1, \dots, \ell - 1, p = 1, \dots, s.$$

Each STMPC optimizes a single task cost \mathcal{J}_ℓ , regularized by an effort cost \mathcal{J}_{eff} , subject to system constraints (2b). After finding the optimal solution to the first STMPC problem containing only the highest priority task cost, consecutive STMPC problems are additionally subject to a *lexicographic optimality constraint* (2c). This constraint enforces that at each step k in the horizon, the tracking error of already-solved higher priority tasks does not increase when optimizing the current lower priority task (details can be found in Section IV.D of [3]). This process leads to an HTMPC solution that leverages mobile manipulators' kinematic redundancy to optimize multiple tasks while enforcing priority.

B. Safe and Interactive Crowd Navigation (SICNav)

To generate coupled human predictions while optimizing the robot plan, we follow SICNav [11] and incorporate the ORCA [12] problems as constraints in the robot's HTMPC to obtain a bilevel optimization problem:

$$\min_{\bar{\mathbf{x}}, \bar{\mathbf{u}}, \bar{\mathbf{a}}} \mathcal{J}(\bar{\mathbf{x}}, \bar{\mathbf{u}}) \quad (3a)$$

$$\text{s.t. } \bar{\mathbf{x}}_0^{(r)} = \bar{\mathbf{x}}_0^{(r)} \quad (3b)$$

$$\bar{\mathbf{u}}_k \in \mathcal{U}, \bar{\mathbf{x}}_k^{(r)} \in \mathcal{X} \quad (3c)$$

$$\bar{\mathbf{x}}_{k+1}^{(r)} = \mathbf{f}(\bar{\mathbf{x}}_k^{(r)}, \bar{\mathbf{u}}_k) \quad (3d)$$

$$\bar{\mathbf{x}}_{k+1}^{(j)} = \mathbf{g}(\bar{\mathbf{x}}_k^{(j)}, \bar{\mathbf{a}}_k^{(j)}) \quad (3e)$$

$$\bar{\mathbf{a}}_k^{(j)} \in \mathcal{O}^{(j)}(\bar{\mathbf{x}}_k). \quad (3f)$$

$$\forall k = 0, \dots, N-1, \forall j = 1, \dots, n_h$$

The *upper-level* cost (3a) and constraints (3b)-(3d) optimize robot trajectories, and the *lower-level* ORCA constraints (3f) enforce that human actions are consistent with ORCA solutions. Here, $\bar{\mathbf{a}}$ denotes the human velocity actions that are themselves the solution to ORCA optimization problem (4). These actions $\bar{\mathbf{a}}$ are used as human trajectory predictions, not as control inputs to the system.

At each time step t , the ORCA-based human prediction model is formulated as follows:

$$\mathcal{O}^{(j)}(\mathbf{x}_t) := \arg \min_{\mathbf{z}, \zeta} \|\mathbf{z} - \dot{\mathbf{q}}_{\text{des}}(\mathbf{x}_t)\|_2^2 + \|\zeta\|_M^2 \quad (4a)$$

$$\text{s.t. } \|\mathbf{z}\|_2^2 \leq z_{\text{max}}^2 \quad (4b)$$

$$\|\mathbf{z} - \dot{\mathbf{q}}_t\|_2^2 \leq \Delta t \ddot{q}_{\text{max}} - \zeta_a \quad (4c)$$

$$\mathbf{n}^{(j|l)}(\mathbf{x}_t)^T \mathbf{z} \geq \rho^{(j|l)}(\mathbf{x}_t) - \zeta_v \quad (4d)$$

$$\mathbf{n}^{(j|\bar{l})}(\mathbf{x}_t)^T \mathbf{z} \geq \rho^{(j|\bar{l})}(\mathbf{x}_t) \quad (4e)$$

$$\zeta \geq 0, \quad (4f)$$

where, $\mathbf{z} \in \mathbb{R}^2$ is the optimization variable, $\dot{\mathbf{q}}_{\text{des}} \in \mathbb{R}^2$ is some given desired velocity (e.g. latest measured human velocity), (4b) is the maximum velocity constraint, and (4d)-(4e) encode human-human and human-robot, and human-static obstacle collision avoidance, respectively. We introduce slack variables $\zeta = (\zeta_a, \zeta_v)$ to ensure there always exists a feasible solution, while penalizing constraint violation with $\|\zeta\|_M^2$ (details can be found in Section III.B of [11]). We also introduce acceleration constraints (4c) to obtain smooth trajectory predictions.

To solve the bilevel problem (3), we reformulate the problem to a single-level optimization by replacing the lower-level ORCA problems with their Karush-Kuhn-Tucker (KKT) conditions.

C. SM²ITH Formulation

We unify the frameworks of Sections IV-A and IV-B into a single control scheme to obtain SM²ITH:

$$\text{lex min}_{\bar{\mathbf{x}}, \bar{\mathbf{u}}, \bar{\mathbf{a}}, \zeta, \bar{\lambda}} [\mathcal{J}_1(\bar{\mathbf{x}}^{(r)}, \bar{\mathbf{u}}), \dots, \mathcal{J}_L(\bar{\mathbf{x}}^{(r)}, \bar{\mathbf{u}})] \quad (5a)$$

$$\text{s.t. } \bar{\mathbf{x}}_{k+1}^{(r)} = \mathbf{f}(\bar{\mathbf{x}}_k^{(r)}, \bar{\mathbf{u}}_k) \quad (5b)$$

$$\bar{\mathbf{x}}_k^{(r)} \in \mathcal{X}, \bar{\mathbf{u}}_k \in \mathcal{U} \quad (5c)$$

$$\bar{\mathbf{x}}_0^{(r)} = \bar{\mathbf{x}}_0^{(r)} \quad (5d)$$

$$\mathbf{sd}(\mathbf{x}_k^{(r)}) \geq \mathbf{d}_{\text{sc}} \quad (5e)$$

$$\mathbf{h}(\mathbf{x}_{k+1}) \geq (1 - \gamma)\mathbf{h}(\mathbf{x}_k) \quad (5f)$$

$$\bar{\mathbf{x}}_{k+1}^{(j)} = \mathbf{g}(\bar{\mathbf{x}}_k^{(j)}, \bar{\mathbf{a}}_k^{(j)}) \quad (5g)$$

$$\nabla_{\bar{\mathbf{a}}_k^{(j)}, \bar{\zeta}^{(j)}} \mathcal{L}(\bar{\mathbf{a}}_k^{(j)}, \bar{\zeta}^{(j)}, \bar{\lambda}^{(j)}) = \mathbf{0} \quad (5h)$$

$$\boldsymbol{\lambda}^T \mathbf{c}(\bar{\mathbf{a}}_k^{(j)}, \bar{\zeta}^{(j)}) = 0 \quad (5i)$$

$$\mathbf{c}(\bar{\mathbf{a}}_k^{(j)}, \bar{\zeta}^{(j)}) \leq \mathbf{0} \quad (5j)$$

$$\bar{\lambda}^{(j)} \geq \mathbf{0}. \quad (5k)$$

$$\forall k = 0, \dots, N-1, \forall j = 1, \dots, n_h$$

In this formulation, we extend HTMPC (1) by incorporating the KKT conditions of ORCA (4) as constraints (5h)-(5k) to obtain predictions of the humans in the environment. These predictions are coupled with the robot plan being optimized.

To ensure human-robot and human-obstacle collision avoidance, we introduce a Discrete-Time Control Barrier Function (DT-CBF) constraint (5f), which maintains the state of the system within the safe set \mathcal{C} , which in our case is the set of all collision-free states. The DT-CBF framework [36]–[38] is formulated by defining the safe set as the superlevel set of a function, $\mathcal{C} = \{\mathbf{x} \in \mathcal{X} : h(\mathbf{x}) \geq 0\}$, where $h(\mathbf{x}) : \mathcal{X} \rightarrow \mathbb{R}$ [13]. The inequality, $h(\mathbf{x}_{k+1}) \geq (1 - \gamma)h(\mathbf{x}_k)$, ensures that \mathcal{C} is forward invariant (i.e., the future state evolutions remain within \mathcal{C} for all initial state $\mathbf{x}_0 \in \mathcal{C}$). The hyperparameter $\gamma \in [0, 1)$ bounds how quickly the system's state can approach the boundary of the safe set.

In our case, to ensure collision avoidance with respect to humans in the shared space, we model the volume occupied by humans as cylinders centered at their respective predicted state $\mathbf{x}_k^{(j)}$ and define a set of DT-CBFs, denoted by $\mathbf{h}^{(j)}(\mathbf{x}_k)$, with individual functions capturing collision-avoidance constraints with respect to the robot base and each arm link. The collision avoidance constraints with respect to static obstacles are defined in a similar manner. The constraint function $\mathbf{h}(\mathbf{x}_k)$ is a concatenation of the constraints capturing collision avoidance with respect to all humans and static obstacles in the environment.

Intuitively, DT-CBFs guarantee that the robot's distance to humans and static obstacles are always above a safety threshold by bounding how quickly the robot can approach humans and obstacles. As a result, the DT-CBF constraints encourage more conservative robot behavior closer to the boundary of the safe set, providing an extra layer of safety. In practice, the hyperparameter γ can be further used to control the level of cautiousness in the robot's behavior.

V. EXPERIMENTS

We evaluate the performance of SM²ITH with three sets of experiments: (i) we execute two tasks simultaneously, switching their priorities (Section V-A), (ii) we perform a sequential pick-and-place task with different human-motion models on two different platforms (Section V-B), and (iii)




Fig. 3: Experimental setup for the banner and cup tasks. Each scenario involves $n_h \in \{1, 2, 3\}$ humans instructed to follow the illustrated trajectories. The robot starts at the illustrated position and moves toward the yellow circle, resulting in an interaction between all agents at the center of the room.

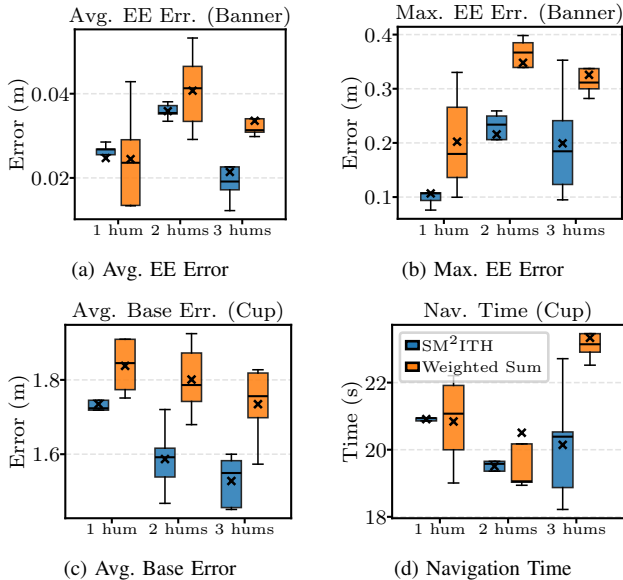


Fig. 4: Box-and-whisker plots summarizing the results of task prioritization experiments comparing the SM²ITH and the weighted sum method. For cases where the end effector task is prioritized we present the average and maximum end-effector tracking error and the navigation efficiency in the Banner tasks (end effector task prioritized $E \geq B$) and in the Cup task (base task prioritized $B \geq E$). Median values are illustrated by the line in the box-and-whisker plots and mean values are indicated by the \times . At higher human densities (3 humans), the proposed SM²ITH method outperforms the weighted sum controller on the high-priority task, highlighting the benefits of strict lexicographic prioritization.

we evaluate adversarial interactions in which a human purposefully blocks the motion of the robot.

For all experiments, we implement the SM²ITH optimization problem in Python solving the MPC in (5) using the *acados* library [39]. Human positions are measured by external motion capture systems; however, this framework is agnostic to perception and can be extended to use onboard sensors [40]. An exhaustive collection of experiments can be found in the accompanying video: <http://tiny.cc/sm2ith>.

A. Shaping Robot Behavior Through Task Prioritization

To evaluate how task priority shapes the behavior of SM²ITH, we formulate a scenario where the robot needs to perform two tasks simultaneously while avoiding humans. We perform these experiments with the Hello-Robot Stretch 3, a 7-DOF platform with a nonholonomic base and

a low-redundancy arm.

Task Specification: Fig. 3 illustrates the experimental setup. In an environment with $n_h = \{1, 2, 3\}$ humans, having the same initial positions, illustrated in Fig. 3, the robot tasks consist of a base navigation task, where the robot is required to reach a predetermined position on the opposite side of the workspace (the yellow circle in Fig. 3), and an end-effector (EE) positioning task, where the robot is required to keep its arm extended away from the base. We design two scenarios with reversed relative priority between the tasks (Fig. 5):

- EE task has higher priority than base task ($E \geq B$): This task prioritization is motivated by a “banner-carrying” task. The main priority of the robot is to display a banner mounted on its arm with a secondary objective of navigating across the room.
- Base task has higher priority than EE task ($B \geq E$): This task is motivated by a robot waiter carrying a cup of water. The main priority of the robot is to navigate to its goal position, but as a secondary objective it prefers to keep the cup away from its base to avoid spilling onto itself.

Evaluated Control Methods: To evaluate the effect of the strict prioritization in SM²ITH, we compare our approach with a weighted-sum MPC baseline. The baseline is composed of the same costs and constraints as (5), but instead of using the lexicographic framework, it optimizes a single whole-body MPC problem combining the two tasks with a weighted sum of costs

$$\min_{\bar{x}, \bar{u}, \bar{a}, \bar{\zeta}, \bar{\lambda}} W_{ee} \mathcal{J}_{ee} + W_{base} \mathcal{J}_{base} + \mathcal{J}_{eff},$$

where W_{ee} and W_{base} are scalar weights chosen to approximate hierarchical preferences.

Evaluation Metrics: To evaluate performance on the EE task, we calculate average and maximum end-effector tracking error over each run. To evaluate performance on the base task, we calculate total navigation time per run, and average deviation from the optimal base path over each run.

Results and Discussion: We run each experiment scenario five times with each control method for a total of 60 runs. We summarize the quantitative results in Fig. 4 comparing the SM²ITH and the weighted sum approach. When the end-effector task has higher priority (banner experiment), we observe that the average EE tracking errors (Fig. 4a) of SM²ITH are similar to the weighted-sum controller at low human densities ($n_h = \{1, 2\}$), but SM²ITH performs significantly better at a high human density ($n_h = 3$). Moreover, if we look at the maximum EE tracking errors (Fig. 4b), we can see that in all human densities the error over the run is lower when using SM²ITH than the baseline. Similarly, when the base navigation task has higher priority (cup experiments), we observe that the SM²ITH methods completed the tasks more efficiently in terms of both metrics (Figs. 4c, 4d). The weighted-sum baseline performs worse in the high-priority task because it does not enforce a strict task hierarchy. Moreover, the optimal weights that approximate task prioritization are not consistent across scenarios.




Fig. 5: In both scenarios, the human slowly approaches the robot arm, and the robot has two options to avoid collision, either by retracting the end effector or moving the base. In the banner transport scenario (top), the end effector task has higher priority and the base of the robot moves to avoid collision with the human, while in the cup transport scenario (bottom), the base task has higher priority, and the end effector retracts instead. These examples showcase the potential of leveraging hierarchical formulation to encode different robot behaviors based on the context.

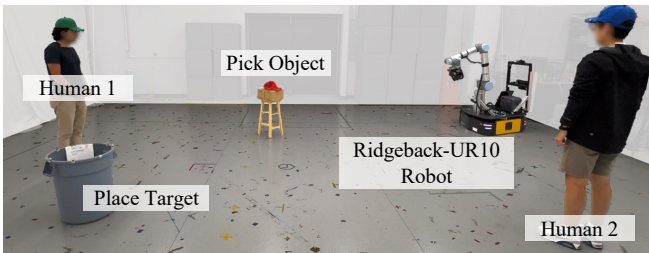


Fig. 6: Experimental setup for the pick-and-place experiments. The robot starts from the illustrated position and must pick the red hat and place it into the target and end the run by returning to the start point. Each scenario involves $n_h \in \{1, 2\}$ humans, starting at the illustrated positions and crossing to the opposite side of the room each time the robot travels between start, pick, place, and end locations.

B. Human Motion Prediction in Pick-and-Place Tasks

To evaluate how human models affect the robot’s behaviors, we task the robot to pick and place objects at designated positions while performing collision avoidance with humans. To study the impact of different levels of redundancy and kinematic structure, we perform the experiments with both the Hello-Robot Stretch 3, a 7-DOF platform with a non-holonomic base, and the Ridgeback-UR10, a 9-DOF platform with a holonomic base, illustrated in Fig. 1.

Task Specification: Fig. 6 illustrates the experimental setup. In an environment with $n_h = \{1, 2\}$ humans, the pick-and-place task consists of a sequence of EE and base targets that lead the robot to pick up the red hat before placing it at a target location.

Evaluated Control Methods: To evaluate the effect of the closed-loop human prediction model integrated into SM²ITH, we compare our proposed approach with two baseline human motion models within the same control framework:

- SM²ITH: Our proposed method that interactively models human motion by embedding the ORCA problem (4) as a lower-level constraint in the controller.
- CVMM: A non-interactive predict-then-plan baseline comprised of the same control formulation as SM²ITH except that the interactive human model is replaced with a CVMM-based human prediction over the horizon. The predictions are pre-generated before each control step t and stay the same during the MPC optimization.
- React: A reactive baseline comprised of the same control formulation as the above two methods, except without any prediction of human motion. This method only reacts to the current tracked positions of the humans.

Evaluation Metrics: To evaluate the effects of human models on the robot’s whole-body navigation efficiency we measure total runtime and the distance of the base from the optimal path, and to evaluate human-robot proxemics, we measure the duration spent in any human’s intimate space (defined as a 0.45 m circle around a human [41]), and average robot–human distance over each run.

Results and Discussions: For each platform, we conduct ten trials for SM²ITH and five trials for each of the other two models in an environment considering either one or two humans resulting in 80 total runs.

From the navigation metrics for both robots (Figs. 7a, 7b), we can observe that CVMM consistently yields the largest deviations from the optimal path, leading to higher runtimes on both platforms. The reactive formulation remains closer to the optimal path but lacks prediction, frequently cutting across human trajectories and compromising safety, as illustrated in Fig. 7c by the longer duration the Ridgeback-UR10 robot spends in the humans’ intimate space. Conversely, CVMM reduces time spent in this space but at the cost of degraded navigation efficiency. In contrast, the SM²ITH framework balances efficiency and safety by adapting predictions to minimize unnecessary detours while respecting human motion.

Unlike the Ridgeback-UR10, the proxemics results for the Stretch 3 robot (Fig. 7d) do not demonstrate significant differences between the evaluated methods. We attribute this difference to platform design. In the navigation results (Figs. 7a, 7b), we can observe that using the Stretch 3 results in all methods having longer runtimes, highlighting the low speed of this robot (maximum 0.3m/s). The Ridgeback-UR10, with its holonomic base and greater redundancy, spends far less time in the intimate space of humans, as it can sidestep approaching agents with agility. On the other hand, the nonholonomic Stretch 3 requires complex maneuvers to orient its base and arm away from humans, resulting in increased close interactions.

We also demonstrate that our proposed framework generalizes to randomized pick-and-place targets on both robots.

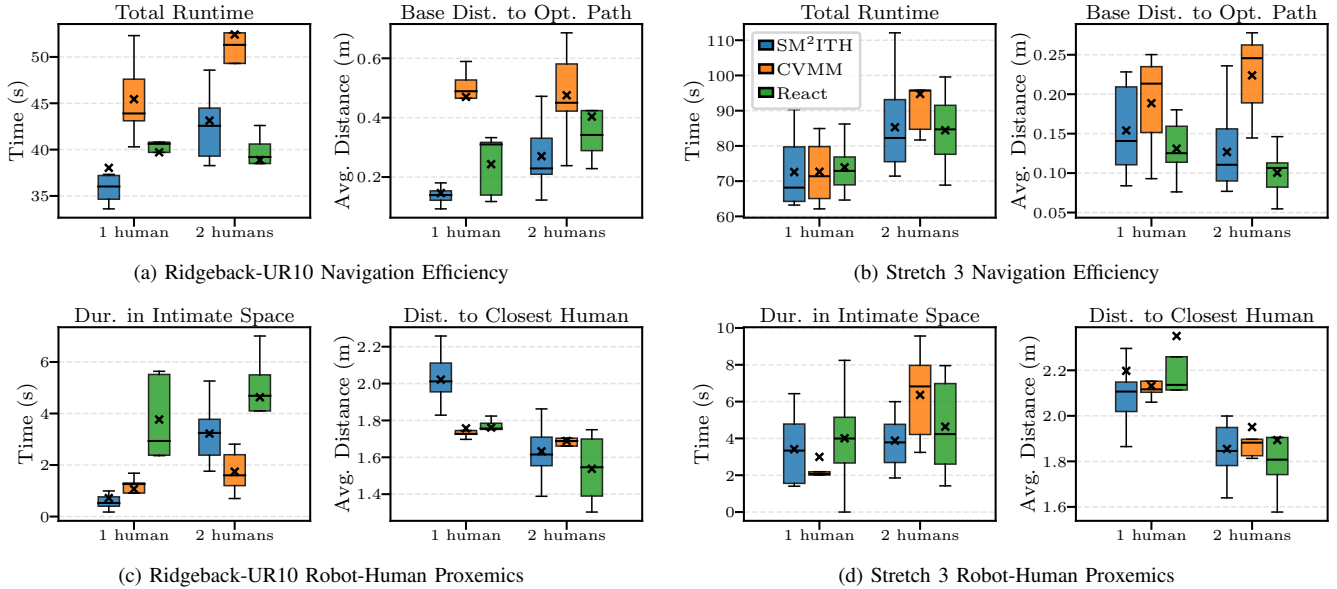


Fig. 7: Box-and-whisker plots summarizing the quantitative results of the pick and place experiments presented in Fig. 6. We compare SM^2ITH , with controllers using Constant Velocity Motion Model (CVMM), and Reactive models on two platforms. Median values are illustrated by the line in the box-and-whisker plots and mean values are indicated by the \times . Results show that CVMM diverges most from optimal paths, while interactive SM^2ITH predictions balance safer proxemics with task efficiency.

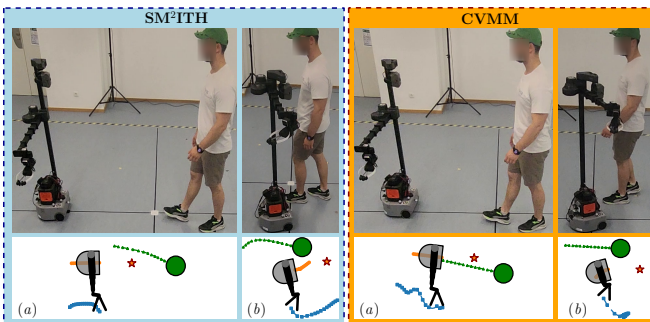


Fig. 8: Comparison of adversarial human interaction using the proposed SM^2ITH (left) and CVMM baseline (right). Each model is shown through experimental snapshots (top row) and reconstructed animations with predicted trajectories (bottom row). The orange star illustrates the robot base target, the orange and blue lines illustrate the base and EE plans, and the green line illustrates the human prediction. Whereas in the first snapshot (a) the robot uses either method plans to retreat, in the second snapshot (b) the interactive predictions in SM^2ITH allows the robot to reason about the coupled human-robot trajectories resulting in a robot plan that continues towards the target.

The robots successfully navigate in the human-centered environment avoiding humans while performing the tasks. The results are shown in the accompanying video.

C. Adversarial Human Interaction

We also qualitatively evaluate the effect of human motion prediction in adversarial situations. We formulate a scenario where the human deliberately walks directly towards the robot, effectively forcing it to retreat. We compare SM^2ITH with the CVMM baseline using the Stretch 3 platform.

We illustrate this scenario in Fig. 8. In the first snapshot for both methods (a), the human is walking towards the robot causing the robot base plan (orange) to back away. In the second snapshot (b) as the human gets closer to the robot,

the SM^2ITH framework generates a joint human prediction (green) and robot plan (orange for base, blue for EE) where both agents proceed while avoiding each other. In contrast, the CVMM model extrapolates human velocity in a straight line, resulting in a robot plan that continues retreating. The interactive nature of SM^2ITH allows the robot to anticipate that the human can adapt to its presence.

VI. CONCLUSIONS AND FUTURE WORKS

In this work, we introduce SM^2ITH , a unified framework that integrates HTMPC with interactive human motion prediction for safe and efficient mobile manipulation in human-centered environments. By combining bilevel optimization with a lexicographic task hierarchy, SM^2ITH enables robots to jointly reason about manipulation, navigation, and human-robot interactions. Our experiments on two different platforms, the Stretch 3 and the Ridgeback-UR10, demonstrate that SM^2ITH achieves better task prioritization compared to a weighted-sum baseline and safe human-robot proxemics. We also demonstrate interactive robot behavior during adversarial interactions.

As future work, we are interested in further (i) incorporating active perception of both the environment and humans using onboard sensors (e.g., using a 6D pose estimation of human body [42] to model human anatomy and thereby obtaining even less conservative human collision avoidance constraints), (ii) leveraging additional models to capture more complex social behaviors (e.g., group dynamics [43], human intent-driven motion [44], or leveraging learning-based generative models in motion prediction [40]), and (iii) conducting experiments considering more cluttered environments to evaluate the framework’s scalability and its ability to maintain efficient navigation within high-density crowds.

REFERENCES

- [1] J. Carius, M. Wermelinger, B. Rajasekaran, K. Holtmann, and M. Hutter, "Deployment of an autonomous mobile manipulator at MBZIRC," *J. of Field Robotics*, vol. 35, no. 8, pp. 1342–1357, 2018.
- [2] A. S. Sathya, G. Pipeleers, W. Decré, and J. Swevers, "A weighted method for fast resolution of strictly hierarchical robot task specifications using exact penalty functions," *Robotics and Automation Letters (RA-L)*, vol. 6, no. 2, pp. 3057–3064, 2021.
- [3] X. Du, S. Zhou, and A. P. Schoellig, "Hierarchical task model predictive control for sequential mobile manipulation tasks," *Robotics and Automation Letters (RA-L)*, vol. 9, no. 2, pp. 1270–1277, 2023.
- [4] M. Ehr Gott, *Multicriteria optimization*. Springer, 2005.
- [5] B. Burgess-Limerick, J. Haviland, C. Lehnert, and P. Corke, "Reactive base control for on-the-move mobile manipulation in dynamic environments," *Robotics and Automation Letters (RA-L)*, vol. 9, no. 3, pp. 2048–2055, 2024.
- [6] T. Merva, S. Bakker, M. Spahn, D. Zhao, I. Virgala, and J. Alonso-Mora, "Globally-guided geometric fabrics for reactive mobile manipulation in dynamic environments," *Robotics and Automation Letters (RA-L)*, 2025.
- [7] T. Gu, G. Chen, J. Li, C. Lin, Y. Rao, J. Zhou, and J. Lu, "Stochastic trajectory prediction via motion indeterminacy diffusion," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 17092–17101.
- [8] S. Poddar, C. Mavrogiannis, and S. S. Srinivasa, "From Crowd Motion Prediction to Robot Navigation in Crowds," in *IEEE/RSJ International Conference on Intelligent Robots and System*, 2023.
- [9] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2010, pp. 797–803.
- [10] S. Schaefer, K. Leung, B. Ivanovic, and M. Pavone, "Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2021, pp. 9673–9679.
- [11] S. Samavi, J. R. Han, F. Shkurti, and A. P. Schoellig, "SICNav: Safe and interactive crowd navigation using model predictive control and bilevel optimization," *IEEE Trans. on Robotics*, 2024.
- [12] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research: The 14th International Symposium ISRR*. Springer, 2011, pp. 3–19.
- [13] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *18th European Control Conf. (ECC)*. IEEE, 2019, pp. 3420–3431.
- [14] B. Burgess-Limerick, C. Lehnert, J. Leitner, and P. Corke, "An architecture for reactive mobile manipulation on-the-move," *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 1623–1629, 2023.
- [15] S. Zimmermann, R. Poranne, and S. Coros, "Go fetch! - dynamic grasps using boston dynamics spot with external robotic arm," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2021, pp. 4488–4494.
- [16] S. Thakar, P. Rajendran, A. M. Kabir, and S. K. Gupta, "Manipulator motion planning for part pickup and transport operations from a moving base," *Trans. on Automation Science and Engineering*, vol. 19, no. 1, pp. 191–206, 2020.
- [17] H. Hanafusa, T. Yoshikawa, and Y. Nakamura, "Analysis and control of articulated robot arms with redundancy," *Int. Federation of Automatic Control (IFAC)*, vol. 14, no. 2, pp. 1927–1932, 1981.
- [18] A. Escande, N. Mansard, and P.-B. Wieber, "Hierarchical quadratic programming: fast online humanoid-robot motion generation," *The Int. J. of Robotics Research*, vol. 33, no. 7, pp. 1006–1028, 2014.
- [19] D. Wang, J. Yu, S. Wu, Z. Li, C. Li, R. Xiong, S. Qu, and Y. Wang, "A hierarchical MPC for end-effector tracking control of legged mobile manipulators," *Trans. on Automation Science and Engineering*, vol. 22, pp. 4855–4866, 2024.
- [20] J. Nie, Y. Wang, Y. Mo, Z. Miao, Y. Jiang, H. Zhong, and J. Lin, "An HQP-based obstacle avoidance control scheme for redundant mobile manipulators under multiple constraints," *Trans. on Industrial Electronics*, vol. 70, no. 6, pp. 6004–6016, 2022.
- [21] Y. Wang, Y. Liu, M. Leibold, M. Buss, and J. Lee, "Hierarchical incremental MPC for redundant robots: A robust and singularity-free approach," *Trans. on Robotics*, vol. 40, pp. 2128–2148, 2024.
- [22] A. Heins and A. P. Schoellig, "Keep it upright: Model predictive control for nonprehensile object transportation with obstacle avoidance on a mobile manipulator," *Robotics and Automation Letters (RA-L)*, vol. 8, no. 12, pp. 7986–7993, 2023.
- [23] S. G. Tarantos, T. Belvedere, and G. Oriolo, "Dynamics-aware navigation among moving obstacles with application to ground and flying robots," *Robotics and Autonomous Systems*, vol. 172, p. 104582, 2024.
- [24] J.-R. Chiu, J.-P. Sleiman, M. Mittal, F. Farshidian, and M. Hutter, "A collision-free MPC for whole-body dynamic locomotion and manipulation," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2022, pp. 4686–4693.
- [25] M. Gaertner, M. Bjelonic, F. Farshidian, and M. Hutter, "Collision-free MPC for legged robots in static and dynamic scenes," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2021, pp. 8266–8272.
- [26] J. Pankert and M. Hutter, "Perceptive model predictive control for continuous mobile manipulation," *Robotics and Automation Letters (RA-L)*, vol. 5, no. 4, pp. 6177–6184, 2020.
- [27] R. M. Bena, G. Bahati, B. Werner, R. K. Cosner, L. Yang, and A. D. Ames, "Geometry-aware predictive safety filters on humanoids: From Poisson safety functions to CBF constrained MPC," *arXiv:2508.11129*, 2025.
- [28] L. Xu, X. Xiong, and Y. Bai, "Dynamic control barrier function based trajectory planning for mobile manipulator," in *Chinese Control Conference (CCC)*. IEEE, 2024, pp. 3815–3820.
- [29] Z. Jian, Z. Yan, X. Lei, Z. Lu, B. Lan, X. Wang, and B. Liang, "Dynamic control barrier function-based model predictive control to safety-critical obstacle-avoidance of mobile robot," *Int. Conf. on Robotics and Automation (ICRA)*, 2023.
- [30] O. de Groot, L. Ferranti, D. M. Gavrila, and J. Alonso-Mora, "Scenario-based motion planning with bounded probability of collision," *The Int. J. of Robotics Research*, 2023.
- [31] R. Firooz, A. Mir, G. S. Camps, and M. Schwager, "OA-MPC: Occlusion-aware MPC for guaranteed safe robot navigation with unseen dynamic obstacles," *Trans. on Control Systems Technology*, vol. 33, no. 3, pp. 940–951, 2025.
- [32] V. Vulcano, S. G. Tarantos, P. Ferrari, and G. Oriolo, "Safe robot navigation in a crowd combining NMPC and control barrier functions," in *Conf. on Decision and Control (CDC)*, 2022, pp. 3321–3328.
- [33] A. Jafari and Y.-C. Liu, "Time-to-collision based social force model for intelligent agents on shared public spaces," *Int. J. of Social Robotics*, vol. 16, no. 9, pp. 1953–1968, 2024.
- [34] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2008, pp. 1928–1935.
- [35] B. Siciliano, L. Villani, G. Oriolo, A. De Luca et al., *Foundations of Robotics*. Springer, 2025.
- [36] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Robotics: Science and Systems*, vol. 13, 2017, pp. 1–10.
- [37] J. Zeng, B. Zhang, and K. Sreenath, "Safety-critical model predictive control with discrete-time control barrier function," in *American Control Conf.* IEEE, 2021, pp. 3882–3889.
- [38] F. D'Orazio, T. Belvedere, S. G. Tarantos, and G. Oriolo, "Maintaining balance of mobile manipulators for safe pick-up tasks," in *Int. Conf. on Control, Automation, Robotics and Vision (ICARCV)*, 2024, pp. 1207–1212.
- [39] R. Verschuere, G. Frison, D. Kouzoupis, J. Frey, N. v. Duijkeren, A. Zanelli, B. Novoselnik, T. Albin, R. Quirynen, and M. Diehl, "acados – a modular open-source framework for fast embedded optimal control," *Mathematical Programming Computation*, vol. 14, no. 1, pp. 147–183, 2022.
- [40] S. Samavi, A. Lem, F. Sato, S. Chen, Q. Gu, K. Yano, A. P. Schoellig, and F. Shkurti, "SICNav-diffusion: Safe and interactive crowd navigation with diffusion trajectory predictions," *Robotics and Automation Letters (RA-L)*, 2025.
- [41] E. T. Hall, "The hidden dimension anchor books editions," *New York*, 1990.
- [42] D. F. Henning, C. Choi, S. Schaefer, and S. Leutenegger, "BodyS-LAM++: Fast and tightly-coupled visual-inertial camera and human motion tracking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 3781–3788.
- [43] M. Moussaïd, N. Perozo, S. Garnier, D. Helbing, and G. Theraulaz, "The walking behaviour of pedestrian social groups and its impact on crowd dynamics," *PLoS one*, vol. 5, no. 4, p. e10047, 2010.
- [44] G. Abbate, A. Giusti, V. Schmuck, O. Celiktutan, and A. Paolillo, "Self-supervised prediction of the intention to interact with a service robot," *Robotics and Autonomous Systems (RAS)*, vol. 171, p. 104568, 2024.