

Winding Through: Crowd Navigation via Topological Invariance

Christoforos Mavrogiannis¹, Krishna Balasubramanian², Sriyash Poddar³
Anush Gandra², Siddhartha S. Srinivasa¹

Abstract—We focus on robot navigation in crowded environments. The challenge of predicting the motion of a crowd around a robot makes it hard to ensure human safety and comfort. Recent approaches often employ end-to-end techniques for robot control or deep architectures for high-fidelity human motion prediction. While these methods achieve important performance benchmarks in simulated domains, dataset limitations and high sample complexity tend to prevent them from transferring to real-world environments. Our key insight is that a low-dimensional representation that captures critical features of crowd-robot dynamics could be sufficient to enable a robot to wind through a crowd smoothly. To this end, we mathematically formalize the act of *passing* between two agents as a rotation, using a notion of topological invariance. Based on this formalism, we design a cost functional that favors robot trajectories contributing higher passing progress and penalizes switching between different sides of a human. We incorporate this functional into a model predictive controller that employs a simple constant-velocity model of human motion prediction. This results in robot motion that accomplishes statistically significantly higher clearances from the crowd compared to state-of-the-art baselines while maintaining competitive levels of efficiency, across extensive simulations and challenging real-world experiments on a self-balancing robot.

I. INTRODUCTION

Navigation in crowds is a challenging task for a mobile robot [24]. Ensuring human safety [37] and comfort requires equipping robots with models of human motion prediction. Predicting the motion of a crowd is fundamentally hard due to the combinatorial structure of the space [8]. This hardness is also reflected in the high sample complexity of data-driven models for human motion prediction [36]. Crucially, predicting human motion alongside a robot can be even harder: human response to robot motion is not well understood, and often driven by novelty effects, limited human mental models, and context dependency [13, 33, 34, 45]. Collecting demonstrations at scale in such domains is complicated and existing datasets are still limited.

Despite these challenges, recent approaches to robot navigation in crowds overemphasize end-to-end control policies

¹Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, USA. Email: {cmavro, siddh}@cs.washington.edu

²Department of Mechanical Engineering, University of Washington, Seattle, USA. Email: {krishnab, agandr}@uw.edu.

³Department of Computer Science and Engineering, Indian Institute of Technology, Kharagpur, Kharagpur, India. Email: poddarsriyash@iitkgp.ac.in

This work was (partially) funded by the Honda Research Institute USA, the National Science Foundation NRI (#2132848) and CHS (#2007011), DARPA RACER (#HR0011-21-C-0171), the Office of Naval Research (#N00014-17-1-2617-P00004 and #2022-016-01 UW), and Amazon.



Fig. 1: Honda’s experimental ballbot [12] navigates next to three users during our lab experiments.

or data-driven models for human trajectory prediction [6, 7, 9, 21]. While these approaches successfully handle controlled or simulated domains, they suffer from data limitations, exacerbated by their high sample complexity and the lack of reliable crowd simulation engines, exhibiting poor generalization and transfer to real-world domains.

Interestingly, as shown by Schöller et al. [39], constant velocity (CV) prediction, a very coarse human motion model may outperform even complex neural architectures on prediction benchmarks across common pedestrian datasets. This can be due to the tendency of neural networks to model model environmental priors, the irrelevance of pedestrians’ motion history or the hardness of modeling multiagent interactions [8].

While the CV model indeed handles some prediction tasks, its inability to model cooperation, the foundation of social order in human navigation [47], results in notable crowd navigation issues like the “freezing robot” problem [42] and the “reciprocal dance” [10]. However, encouraged by its efficacy, our key insight is that a coarse motion model like CV could be sufficient for a controller that models critical aspects of interaction. Following this insight, we describe a low-dimensional abstraction of multiagent dynamics based on a formalism of pairwise *passing*, capturing the aspects of progress and directionality using a notion of topological invariance. Based on this abstraction, we design a cost that favors robot trajectories expediting agile passing maneuvers

through a crowd. Around this cost, we design a model predictive control (MPC) architecture that employs CV as the human transition model. Through challenging real-world and simulated scenarios involving navigation of a self-balancing robot in a crowded environment, we demonstrate that our controller exhibits significantly safer behavior for similar efficiency compared to state-of-the-art end-to-end and model-based baselines. A video with snippets from our experiments can be found at [this](#) link.

II. RELATED WORK

We review work on robot navigation in crowds, highlighting approaches that model multiagent cooperation to inform robot decision-making.

A. Robot Navigation in Crowds

Recent work on crowd navigation has employed data-driven techniques to extract social objectives and prediction models from datasets [20, 30] and crowd simulation engines [3, 43]. Several approaches used inverse reinforcement learning (IRL) techniques to recover socially aware control policies. Ziebart et al. [49] learned a reward function that describes human decision making for navigation and used it during planning to avoid obstructing human paths. Kretzschmar et al. [18] and Kim and Pineau [14] learned robot reward functions incorporating social preferences like passing from the right side or respecting users’ personal space. More recently, Chen et al. [6], Everett et al. [9], Liu et al. [21] employed deep reinforcement learning (RL) to learn control policies for decentralized multiagent collision avoidance. Motivated by real-world constraints, Pérez D’Arpino et al. [32] employed RL to learn a navigation policy for obstacle-cluttered indoor environments whereas Monaci et al. [27] developed an RL vision-based architecture for robots with narrow field-of-view cameras. In parallel, some approaches incorporated data-driven predictions into the general MPC pipeline. For instance, Nishimura et al. [28] and Roh et al. [35] integrated multimodal trajectory prediction [38] into the MPC pipeline whereas Wang et al. [46] incorporated group-based behavior prediction to avoid intruding the social space of pedestrian groups.

Despite the success of these methods in controlled and simulated environments, their sample complexity coupled with the inability of existing crowd simulators to realistically model crowd motion next to robots limits their applicability and generalization to new domains. To mitigate this issue, recent work blends data-driven and model-based reasoning. For instance, Brito et al. [4] employed an interaction-aware RL-based subgoal recommendation model to smoothly guide a MPC towards the robot’s goal. Inspired by their work, our MPC architecture also makes use of subgoals to inform the optimization process but additionally adapts the pool of candidate trajectories to the context by propagating forward interaction-aware policies.

B. Formalizing Cooperation in Crowd Navigation

Acknowledging the prevalent role of cooperation [47] in crowd navigation, recent work incorporates the expectation

of human cooperation and rationality into the design of frameworks for prediction and control [24]. For instance, Trautman et al. [42] developed a Gaussian-process-based human motion prediction model that estimates future human trajectories under the assumption of cooperative and goal-directed human behavior. Recently, Sun et al. [41] formalized multiagent cooperation as a superposition of agents’ probabilistic motion preferences, and developed a sampling-based, cooperation-aware motion planner. Cao et al. [5] developed a graph-search-based approach that plans robot motion by discovering channels of safe passages within a crowd. In past work, we formalized cooperation as inference over multiagent motion primitives represented as topological symbols [26] and trajectory sets [23]. This work set the foundation for Social Momentum (SM) [25], a real-time, reactive controller that generates robot motion by maximizing the magnitudes of the pairwise momenta defined between the robot and humans. This has the effect of producing actions that exaggerate over a passing side while moving the robot as far as possible. While this behavior is legible and positively perceived, it can be undesirable in more densely crowded environments with less space for maneuverability.

In this paper instead, we formalize a topological notion of *passing* based on a pairwise winding number [1], which enables the robot to monitor and expedite the progress of passing an agent. Topological representations offer interpretability and decomposition of challenging problems in robotics [2, 16, 29, 31, 40]. While prior work has also used winding numbers to capture multimodality in multiagent navigation, existing approaches do so in a binary sense to predict passing side preferences [18, 35], essentially discarding the notion of passing progress provided by the absolute value of the winding number. Instead, our approach is built around this notion of passing progress: out of a pool of candidate robot trajectories, our cost selects the one that contributes the most passing progress without bias over a passing side. Assuming goal-directed human behavior around the robot (humans moving from start to goal), minimizing our cost also has the effect of preserving a passing side. Our formalism is decoupled from the constraint of collision avoidance which is instead handled through a cost based on an explicit approximation of personal space [15]. Finally, unlike SM [25], which is essentially a reactive, rule-based controller with a one-step horizon, our framework incorporates arbitrary horizons through a more general MPC formulation yielding smoother performance.

III. PRELIMINARIES

In this section, we formalize our domain of interest, and introduce the technical background of our approach.

A. Problem Statement

Consider a workspace $\mathcal{W} \subseteq \mathbb{R}^2$ where a robot navigates among n dynamic agents. Denote by s_k the state of the robot and by $s_k^i \in \mathcal{W}$, $i \in \mathcal{N} = \{1, \dots, n\}$, the state of agent i at timestep k . The robot is navigating from a state s_0 towards a destination s_T by executing controls u from a space of controls

\mathcal{U} , subject to its dynamics model $s_{k+1} = g(s_k, u_k)$. Agent $i \in \mathcal{N}$ is navigating from s_0^i towards a destination s_T^i by executing controls u^i from a space of controls \mathcal{U}^i , striving for safety and efficiency. The robot is not aware of agent i 's destination s_T^i or policy. However, we assume that the robot is perfectly observing the complete world state $(s_k, s_k^{1:n})$, where $s_k^{1:n} = (s_k^1, \dots, s_k^n)$. Our goal is to design a robot policy that enables the robot to navigate from s_0 to s_T safely and efficiently.

B. Model Predictive Control for Crowd Navigation

We employ a discrete MPC formulation that optimizes a robot control trajectory \mathbf{u} from a set \mathcal{U} of candidate trajectories of finite length N with respect to a cost \mathcal{J} :

$$\begin{aligned} \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathcal{U}} \mathcal{J}(\mathbf{s}, \mathbf{s}^{1:n}) \\ \text{s.t. } s_{k+1} &= g(s_k, u_k) \\ s_{k+1}^i &= f(s_{k-h:k}, \mathbf{s}_{k-h:k}^{1:n}), i \in \mathcal{N} \end{aligned}, \quad (1)$$

where we denote by $\mathbf{s} = s_{1:N}$ the robot trajectory, by $\mathbf{s}^{1:n} = (s^1, \dots, s^n)$ the trajectories of agents $1, \dots, n$ lying in front of the robot, and by $\mathbf{u} = u_{0:N-1}$ a robot control trajectory, drawn from a set of short-horizon (N) trajectories \mathcal{U} . Finally, we denote by f a state transition model for agent i that takes as input the system state history up to h timesteps in the past.

C. V-MPC: Balancing Safety and Efficiency

Many recent crowd navigation controllers follow objectives related to safety and efficiency [9, 21, 28, 35, 41, 46]. We defined a baseline *Vanilla* MPC (V-MPC) encoding such specifications through a cost \mathcal{J}_ν :

$$\mathcal{J}_\nu(\mathbf{s}, \mathbf{s}^{1:n}) = a_g \mathcal{J}_g(\mathbf{s}) + a_d \mathcal{J}_d(\mathbf{s}, \mathbf{s}^{1:n}). \quad (2)$$

The term

$$\mathcal{J}_g(\mathbf{s}) = \sum_{k=0}^{N-1} (s_{k+1} - s_T)^\top Q_g (s_{k+1} - s_T), \quad (3)$$

is a goal-tracking cost penalizing trajectories taking the robot further from its goal, where Q_g is a weight matrix. The term

$$\mathcal{J}_d(\mathbf{s}, \mathbf{s}^{1:n}) = \sum_{k=0}^{N-1} \sum_{i=1}^n A_d^2(s_{k+1}, s_{k+1}^i), \quad (4)$$

is a cost penalizing violations to agents' personal space [11] through the *Asymmetric Gaussian Integral Function* A_d of Kirby [15]. The weights a_g, a_d encode the relative importance of cost terms. We approximate agents' state transition in (1) through a constant-velocity (CV) motion model $s_{k+1}^i = f(s_{k-1:k}^i)$, propagating the velocity of agent i one timestep dt into the future, ignoring any interactions with other agents.

IV. TOPOLOGY-INFORMED NAVIGATION

We introduce a mathematical representation that enables a robot to monitor the process of *passing* other agents. Based on this representation, we design a cost function that motivates robot actions contributing more passing progress.




Fig. 2: Two different executions of the same scenario. As agents navigate (color intensity denotes timing), the vector connecting their positions rotates from θ_0 to θ_3 . When they pass each other on the right (a), whereas when they pass each other on the left (b), it is clockwise. The winding number integral $\lambda = (\theta_3 - \theta_0)/2\pi$ is a topological invariant, independent of the specific paths that agents take, describing the signed number of times that agents pass each other.

A. Formalizing Passing as Pairwise Rotation

We denote by $x_k^i \in \mathbb{R}^2$ the position of agent i relative to the robot at timestep k , and by $\theta_k^i = \angle x_k^i$ the angle of that vector with respect to a fixed global frame. From time k to $k+1$, agents' displacement from x_k^i to x_{k+1}^i results in a rotation $\Delta\theta_{k+1}^i = \theta_{k+1}^i - \theta_k^i$. We denote by $\mathbf{s} = s_{1:N} = (s_1, \dots, s_N)$ the robot's trajectory and by $\mathbf{s}^i = s_{1:N}^i = (s_1^i, \dots, s_N^i)$ the trajectory of agent $i \in \mathcal{N}$. The quantity

$$\lambda^i(\mathbf{s}, \mathbf{s}^i) = \frac{1}{2\pi} \sum_{k=0}^{N-1} \Delta\theta_{k+1}^i, \quad (5)$$

is a *winding number* [1] representing the signed number of times that the robot and agent i revolved around each other over a horizon of N timesteps. In the crowd navigation domains of our focus, where agents exhibit goal-directed, cooperative behavior, the winding number captures the progress (magnitude), and side (sign) of *passing* each other (see Fig. 2). A right-side passing induces a counterclockwise rotation yielding a positive winding number ($\lambda^i > 0$), whereas a left-side passing yields a clockwise rotation and thus a negative winding number ($\lambda^i < 0$).




Fig. 3: Effect of minimizing the passing cost. The robot predicts how the blue and yellow agents will move and selects an action (black arrow) that maximizes the average passing progress over the control horizon (color intensity indicates timing). The circular sectors denote the estimated progress to be achieved for each agent. The robot ignores agents lying behind it, i.e., on the negatives of its body x -axis (red color).

B. A Cost Functional that Expedites Passing

Our key insight is that by monitoring and contributing to the values of λ^i , $i \in \mathcal{N}$, through its own actions, the robot may anticipate, influence, and proactively adapt to the passing sides of humans without relying on high-accuracy models of human motion prediction. To engineer robot motion that expedites the process of passing between the robot and other agents, we design the passing cost

$$\mathcal{J}_p(\mathbf{s}, \mathbf{s}^{1:n}) = -\frac{1}{n} \sum_{i=1}^n \lambda^i(\mathbf{s}, \mathbf{s}^i)^2. \quad (6)$$

We assume that the future trajectories of human agents $\mathbf{s}^{1:n}$ are estimated through a transition model f as shown in (1), and provided as input in \mathcal{J}_p . Given $\mathbf{s}^{1:n}$, the minimization of \mathcal{J}_p over \mathbf{s} yields a robot trajectory \mathbf{s}^* that maximizes (on average) the winding numbers between the robot and human agents. This trajectory enables the robot to maneuver itself in a way that aligns with the estimated passing intentions of others. This corresponds to robot motion that avoids switching between passing sides as it approaches a human, expediting and facilitating collision avoidance. Fig. 3 gives intuition about the effect of minimizing the proposed cost. Note that the cost of eq. (6) excludes non-reactive agents lying behind the robot.

C. Topology-Informed MPC

The topology-informed MPC (T-MPC) is an extension of V-MPC (2) that incorporates the functional of (6):

$$\mathcal{J}(\mathbf{s}, \mathbf{s}^{1:n}) = \mathcal{J}_v + a_p \mathcal{J}_p(\mathbf{s}, \mathbf{s}^{1:n}), \quad (7)$$

where a_p is a weight of relative significance. T-MPC strives to expedite passing between the robot and other agents while respecting their personal space and approaching its destination. In conjunction, this formulation is designed to motivate goal-oriented, anticipatory collision avoidance.

Remark: As discussed in Sec. III-B, our core MPC formulation is an optimization over a discrete set of short-horizon motion primitives. These primitives are scored with respect to a cost computed by taking into account only agents lying in front of the robot (see Fig. 3). Thus, the minimization of the \mathcal{J}_p cost in (7) maximizes the average progress with respect to agents lying in front of the robot. Through subsequent control cycles, the minimization of \mathcal{J}_p contributes consistent progress to the prevailing passing side; switching to a different side gets progressively more penalized since it would correspond to a lower winding number as two agents approach each other. Once the robot passes a human agent, it stops considering them in the computation of costs, thus avoiding possible undesired effects like circling around them.

V. EVALUATION

We evaluated our controller through simulated and real-world experiments involving robot navigation in crowded scenes. We used Honda's experimental ballbot [12] (see Fig. 1) and its own low-level velocity controller [48]. A ballbot [19] is a mobile robot with a mechanical body dynamically balancing on top of an omnidirectional spherical wheel. The ballbot's dynamics produces readable [22] and safe behavior [19], making it amenable for operation in crowded spaces.

A. Baselines

We evaluate our MPC against two baselines from the literature:

CADRL [9] is a deep RL framework for multiagent collision avoidance. We adapt the original formulation to handle ballbot dynamics by augmenting: a) its state \mathbf{s} to include the ballbot's inclination ϕ , and the average distance to the robot's goal over the past t' timesteps, \bar{d}_g ; b) its reward function to include an inclination penalty term

$$R_{lean}(\mathbf{s}) = \begin{cases} -1, & \text{if } \phi > \phi_{max} \\ -0.1 \frac{\phi}{\phi_{max}}, & \text{otherwise,} \end{cases} \quad (8)$$

where ϕ_{max} is an inclination threshold (set to 0.25, if $\phi > \phi_{max}$ the episode terminates), and a reward of progress to goal $R_{prog}(\mathbf{s}) = 0.1 \cdot (\bar{d}_g - d_g(\mathbf{s}))$, where $d_g(\mathbf{s})$ is the robot's distance to goal. We included the ballbot dynamics as part of the state transition model. Initializing from the implementation of [9], we trained the model in two stages, first with 2-4 agents and then with 2-10 agents.

ORCA [44] is a crowd simulator that is often used in the evaluation of crowd navigation algorithms [5, 6, 9, 21, 21, 24, 25]. We employ the ORCA configuration of Chen et al. [6].

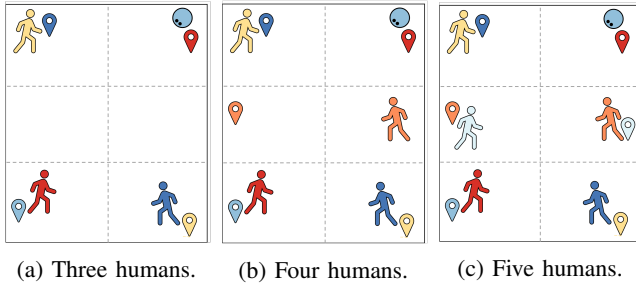


Fig. 4: The scenarios used in our evaluation. The workspace is split in 6 zones. We draw agents in their starting zones and show their destination zones as landmarks of the same color.

B. MPC Implementation

Motivated by the approach of Brito et al. [4], our MPC employs a set \mathcal{U} of rollouts (robot control trajectories) extracted by propagating a policy π_{sim} towards $m = 10$ subgoals, placed around the robot at fixed orientation intervals of $\frac{\pi}{5}$ and distance of $8m$ for 10 timesteps of size $dt = 0.1s$. We experimented with three policies: CV, in which we propagate the robot’s state forward assuming a constant preferred speed of $0.8m/s$ without collision avoidance considerations; ORCA [44]; CADRL [9]. For human motion prediction f we used a CV approximation, propagating humans’ velocities forward. We tuned all MPC variants through parameter sweeps for *Safety* \mathcal{D} (m), defined as the minimum distance between the robot and human agents throughout a trial, and *Time to destination* \mathcal{T} (s) over 30 randomized trials for each scenario of Fig. 4. The control loop closed with frequency $10Hz$.

C. Experiment Design

We conducted our evaluation in a workspace of area $3.6 \times 4.5m^2$. We partitioned the workspace into six zones of area $1.8 \times 1.5m^2$, and defined three scenarios involving 3, 4, and 5 humans, as shown in Fig. 4. In each scenario, humans move between start and goal zones, selected to give rise to challenging human-robot encounters. Across scenarios, the robot’s start and goal coordinates are fixed at $(0,0)$ and $(3.6, 4.5)$, respectively. The preferred speed for the robot is set to $0.8m/s$ which was found to be a natural human walking speed in pilot trials. We evaluated robot performance with respect to *Safety* and *Time to destination*.

Overall, we anticipated that MPC performance would improve when using rollouts from reactive policies like CADRL and ORCA since the robot would have an informed way of reacting to human motion estimates compared to the CV baseline. Further, we anticipated that T-MPC would outperform all baselines in terms of *Safety* since the passing cost would enable the robot to make consistent passing progress, avoiding approaching humans by switching passing sides. We formalized these insights into the following hypotheses:

H1: MPC controllers with ORCA or CADRL rollouts outperform controllers with CV rollouts in terms of *Safety* across all scenarios and worlds.

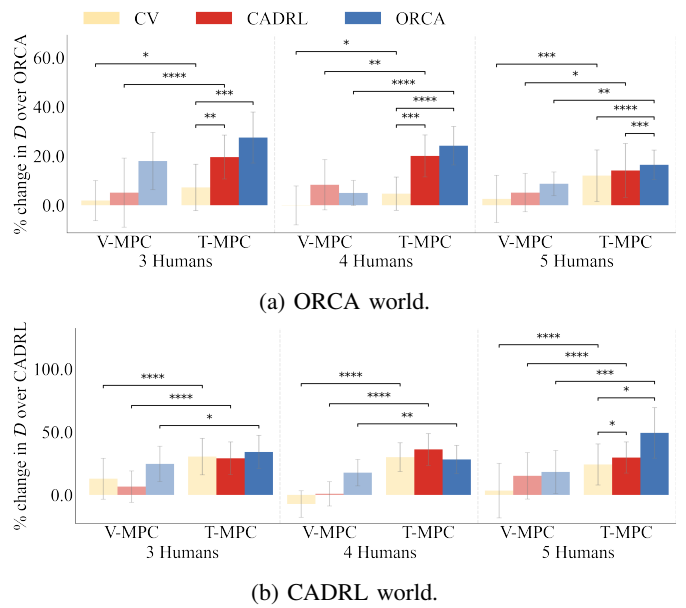


Fig. 5: Comparison of MPC variants in terms of % change in *Safety* over the baselines of the two worlds. Error bars indicate 95% confidence intervals. Stars indicate significance levels ($p < 0.05$, $p < 0.01$, $p < 0.001$) according to a U test.

H2: T-MPC outperforms V-MPC with identical rollouts across all scenarios (3, 4, 5 agents) and worlds (ORCA, CADRL) in terms of *Safety*.

H3: T-MPC outperforms both CADRL and ORCA across all scenarios and worlds in terms of *Safety*.

Simulations: We instantiated all scenarios from Fig. 4 in two Gazebo [17] worlds: one in which humans are simulated as ORCA agents [44], and one as CADRL [9] agents. For each scenario, we generated 100 trials by sampling start and goal coordinates for agents, uniformly at random from their assigned zones. We executed each batch of trials with the following policies: 1) ORCA; 2) CADRL; 3) V-MPC-CV; 4) V-MPC-ORCA; 5) V-MPC-CADRL; 6) T-MPC-CV; 7) T-MPC-ORCA; 8) T-MPC-CADRL. Across all simulations, humans are represented as spheres of $0.3m$ radius whereas the robot’s body is modeled as a cylinder of radius $0.2m$.

Real-world experiments: We instantiated the scenario of Fig. 4a in a lab workspace, deploying Honda’s ballbot [12] to navigate next to three members of our research team (see Fig. 1). We ran 60 trials of the scenario in which the robot navigated with three different policies (20 trials per policy): a) CV in which the robot drives to the goal with constant velocity without avoiding collisions; b) CADRL; c) T-MPC-CADRL. The users were told to “navigate towards the destinations designated by the scenario with normal walking speed and to treat the robot as a walking human”. The users and the robot wore hats with reflective markers that enabled high-accuracy localization through an Optitrack motion-capture system of twelve overhead cameras operating at $120Hz$.

TABLE I: Performance with respect to Safety (\mathcal{D}) and Time to destination (\mathcal{T}) per world. Each entry contains a mean and a standard deviation over 100 trials. The best performing controller per column is printed in bold. Green, orange, and blue entries are controllers outperformed by the best at significance level of $p < 0.001$, < 0.01 , and < 0.05 respectively (U test).

Scenario	Three Humans				Four Humans				Five Humans			
	$\mathcal{D}(m)$		$\mathcal{T}(s)$		$\mathcal{D}(m)$		$\mathcal{T}(s)$		$\mathcal{D}(m)$		$\mathcal{T}(s)$	
World	ORCA	CADRL	ORCA	CADRL	ORCA	CADRL	ORCA	CADRL	ORCA	CADRL	ORCA	CADRL
ORCA	0.67 ± 0.14	0.93 ± 0.26	9.33 ± 0.44	10.55 ± 1.13	0.64 ± 0.12	0.80 ± 0.26	9.70 ± 1.33	11.01 ± 1.96	0.61 ± 0.05	0.69 ± 0.18	10.08 ± 1.57	11.37 ± 2.32
CADRL	0.68 ± 0.19	0.97 ± 0.31	9.82 ± 1.40	10.50 ± 1.60	0.63 ± 0.11	0.91 ± 0.24	10.34 ± 1.65	10.62 ± 1.49	0.6 ± 0.06	0.72 ± 0.18	12.32 ± 2.85	12.86 ± 3.42
V-MPC-CV	0.64 ± 0.13	0.93 ± 0.29	10.87 ± 1.11	11.26 ± 1.37	0.61 ± 0.10	0.80 ± 0.28	10.77 ± 1.26	11.90 ± 2.02	0.59 ± 0.09	0.61 ± 0.17	11.27 ± 2.14	12.70 ± 3.17
V-MPC-ORCA	0.74 ± 0.24	1.06 ± 0.29	11.14 ± 1.45	12.14 ± 1.97	0.66 ± 0.14	1.01 ± 0.26	11.66 ± 1.70	12.01 ± 1.25	0.65 ± 0.12	0.77 ± 0.24	12.17 ± 3.02	13.18 ± 3.12
V-MPC-CADRL	0.64 ± 0.15	0.88 ± 0.33	12.61 ± 3.22	12.29 ± 1.54	0.65 ± 0.15	0.89 ± 0.36	11.97 ± 2.63	12.36 ± 2.29	0.61 ± 0.14	0.69 ± 0.25	15.21 ± 4.10	13.68 ± 3.09
T-MPC-CV	0.67 ± 0.16	1.08 ± 0.31	10.60 ± 1.16	11.15 ± 1.50	0.64 ± 0.13	1.09 ± 0.31	10.44 ± 0.99	10.68 ± 0.90	0.63 ± 0.13	0.83 ± 0.25	10.94 ± 1.49	10.88 ± 1.14
T-MPC-ORCA	0.83 ± 0.34	1.13 ± 0.31	11.50 ± 2.21	12.15 ± 1.83	0.78 ± 0.23	1.11 ± 0.32	11.11 ± 1.98	11.79 ± 2.09	0.70 ± 0.15	0.87 ± 0.26	12.66 ± 3.52	13.08 ± 3.28
T-MPC-CADRL	0.77 ± 0.26	1.17 ± 0.38	11.15 ± 2.05	11.73 ± 1.57	0.75 ± 0.25	1.14 ± 0.38	11.09 ± 1.98	11.13 ± 1.46	0.66 ± 0.17	0.87 ± 0.28	14.19 ± 3.08	12.65 ± 2.61

D. Results

Table I contains the controller performance per scenario whereas Figs. 5, 6, 7 highlight important observations. A video with snippets from our experiments can be found at [this link](#).

H1. Fig. 5 depicts the differences between the T-MPC and V-MPC variants in the form of % improvement over the world baseline (i.e., ORCA or CADRL) in terms of Safety. We see that V-MPC and T-MPC exhibit safer performance when rolling out CADRL or ORCA policies across scenarios in the two simulated worlds. Thus, H1 was supported.

H2. In Fig. 5, we see that for identical rollouts, T-MPC is significantly safer than V-MPC across scenarios in the two simulated worlds. Thus, H2 was also supported.

H3. In the simulated worlds, T-MPC variants with CADRL and ORCA rollouts generally perform best in terms of Safety (see Table I). Fig. 6 gives a deeper insight containing raw datapoints of comparing T-MPC and CADRL/ORCA: in each world, the best performing T-MPC is the safest. This is even more pronounced in challenging scenarios with 4 or 5 humans. Crucially, as shown in Fig. 7, these trends transfer to the real world; T-MPC-CADRL is safer than CADRL while exhibiting similar efficiency. Thus, we find that H3 was supported.

VI. DISCUSSION

T-MPC was safer than baselines (H3), using CV prediction, a very coarse approximation compared to CADRL’s collision avoidance mechanism, extracted through training for millions of episodes. It handled robustly different numbers and types of agents ranging from efficient ORCA agents to socially aware CADRL agents in both worlds. Through an ablation study, we showed that T-MPC significantly outperformed alternative competitive V-MPC architectures (H2) which benefit from reactive rollouts (H1). Our investigation underscores the value of integrating a model of passing into robot’s decision making: by consistently contributing passing progress, the robot is able to better coordinate safe passages in crowded scenes.

Further, while CADRL and ORCA were more efficient in simulation (see Table I), this trend did not transfer to the real world (see Fig. 7b). This demonstrates that while T-MPC could be improved for efficiency, its performance is competitive when deployed in realistic settings.

Finally, after the study, users informally confirmed that the differences between policies were noticeable: they found CV

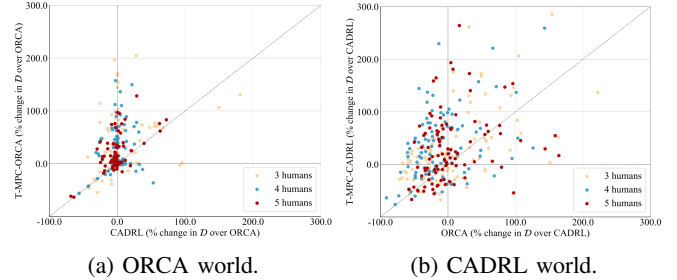


Fig. 6: Scatter plots of % change in Safety over the world baseline. The plots compare the best-performing T-MPC for each world, i.e., T-MPC-ORCA in (a) and T-MPC-CADRL in (b) against CADRL and ORCA respectively.

to be the least preferred and commended the expressiveness of T-MPC-CADRL but also the predictability of CADRL. While these insights are not statistical, they provide nuance and could inform the design of future large-scale studies [25].

A. Limitations

Evaluation criteria. Safety and efficiency are necessary but not sufficient attributes for smooth robot performance in crowded domains. Higher-order properties of robot motion such as smoothness or acceleration are also known to influence users’ impressions in terms of comfort [25], personality and capabilities [45] and it would be important to account for them.

Experimental conditions. While our evaluation scenarios (Fig. 4) gave rise to challenging human-robot encounters, human motion was cooperative and goal-directed as users were instructed to treat the robot like a human pedestrian. Under such settings, a) our CV-based prediction adapted effectively to the linear segments of human motion; b) our control strategy enabled the robot to expedite pairwise passings, leveraging human cooperation. However, in real-world environments pedestrians are often distracted, rushing, forming groups or changing intentions as they walk, affected by obstacles and the context. Our system would not be able to react appropriately to such motion: CV predictions would not be accurate, and the non-cooperative human response would challenge our winding strategy. Transitioning to such settings would require more expressive models of motion prediction [36] but also the integration of context-aware cost functions. Engineering these settings in the real world would also require new




Fig. 7: Robot Safety (a) and Time to destination (b) in real-world experiments. Error bars are 95% confidence intervals.

benchmarking experiments that motivate alternative modes of human behavior in an unbiased fashion. In ongoing work, we are developing experimental protocols involving pedestrians with changing intentions, aggressive, and distracted motion.

Hyperparameter tuning. As detailed in Sec.V-B, we optimized cost weights for Safety and Time in simulation; we then employed the same weights to the real world without any additional fine tuning as the observed behaviors appeared consistent with simulation. However, we did not perform any optimization over the rollout parametrization; parameters like the number of subgoals, their locations, the rollout discretization and horizon might have influenced performance. In the selection of these parameters, we were guided by the goal of securing a 10Hz control frequency, an empirical standard for real-time response in moderately dynamic environments [16, 25]. A higher compute budget or a more efficient system implementation could bring performance improvements.

REFERENCES

- [1] M. A. Berger. Topological invariants in braid theory. *Letters in Mathematical Physics*, 55(3):181–192, 2001.
- [2] S. Bhattacharya, M. Likhachev, and V. Kumar. Identification and representation of homotopy classes of trajectories for search-based path planning in 3d. In *Proceedings of Robotics: Science and Systems (RSS)*, 2011.
- [3] A. Biswas, A. Wang, G. Silvera, A. Steinfeld, and H. Admoni. Socnavbench: A grounded simulation testing framework for evaluating social navigation. *Transactions on Human-Robot Interaction*, 11(3), 2022.
- [4] B. Brito, M. Everett, J. P. How, and J. Alonso-Mora. Where to go next: Learning a subgoal recommendation policy for navigation in dynamic environments. *IEEE Robotics and Automation Letters*, 6(3):4616–4623, 2021.
- [5] C. Cao, P. Trautman, and S. Iba. Dynamic channel: A planning framework for crowd navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 5551–5557, 2019.
- [6] C. Chen, Y. Liu, S. Kreiss, and A. Alahi. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 6015–6022, 2019.
- [7] Y. F. Chen, M. Liu, M. Everett, and J. P. How. Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 285–292, 2017.
- [8] G. F. Cooper. The computational complexity of probabilistic inference using bayesian belief networks. *Artificial Intelligence*, 42(2):393 – 405, 1990.
- [9] M. Everett, Y. F. Chen, and J. P. How. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3052–3059, 2018.
- [10] F. Feurtey. Simulating the collision avoidance behavior of pedestrians. Master’s thesis, University of Tokyo, 2000.
- [11] E. T. Hall. *The Hidden Dimension*. Anchor Books, 1966.
- [12] Honda. Honda P.A.T.H. Bot, 2019. URL https://global.honda/innovation/CES/2019/path_bot.html.
- [13] M. Joosse, M. Lohse, N. V. Berkel, A. Sardar, and V. Evers. Making appearances: How robots should approach people. *Transactions on Human-Robot Interaction*, 10(1), 2021.
- [14] B. Kim and J. Pineau. Socially adaptive path planning in human environments using inverse reinforcement learning. *International Journal of Social Robotics*, 8(1):51–66, 2016.
- [15] R. Kirby. *Social Robot Navigation*. PhD thesis, Carnegie Mellon University, 2010.
- [16] R. A. Knepper, S. S. Srinivasa, and M. T. Mason. Toward a deeper understanding of motion alternatives via an equivalence relation on local paths. *The International Journal of Robotics Research*, 31(2):167–186, 2012.
- [17] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, volume 3, pages 2149–2154, 2004.
- [18] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *The International Journal of Robotics Research*, 35(11):1289–1307, 2016.
- [19] T. Lauwers, G. Kantor, and R. Hollis. One is enough! In S. Thrun, R. Brooks, and H. Durrant-Whyte, editors, *Robotics Research*, pages 327–336. Springer Berlin Heidelberg, 2007.
- [20] A. Lerner, Y. Chrysanthou, and D. Lischinski. Crowds by example. *Comput. Graph. Forum*, 26(3):655–664, 2007.
- [21] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell. Decentralized structural-rnn for robot crowd navigation with deep reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 3517–3524, 2021.
- [22] S. Lo, K. Yamane, and K. Sugiyama. Perception of pedestrian avoidance strategies of a self-balancing mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1243–1250, 2019.
- [23] C. Mavrogiannis and R. A. Knepper. Hamiltonian coordination primitives for decentralized multiagent navigation.

- The International Journal of Robotics Research*, 40(10-11):1234–1254, 2021.
- [24] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh. Core Challenges of Social Robot Navigation: A Survey. *arXiv e-prints*, art. arXiv:2103.05668, Mar. 2021.
- [25] C. Mavrogiannis, P. Alves-Oliveira, W. Thomason, and R. A. Knepper. Social momentum: Design and evaluation of a framework for socially competent robot navigation. *Transactions on Human-Robot Interaction*, 11(2), 2022.
- [26] C. I. Mavrogiannis and R. A. Knepper. Multi-agent path topology in support of socially competent navigation planning. *The International Journal of Robotics Research*, 38(2-3):338—356, 2019.
- [27] G. Monaci, M. Aractingi, and T. Silander. DiPCAN: Distilling Privileged Information for Crowd-Aware Navigation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2022.
- [28] 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 *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11205–11212, 2020.
- [29] A. Orthey and M. Toussaint. Section patterns: Efficiently solving narrow passage problems in multilevel motion planning. *IEEE Transactions on Robotics*, 37(6):1891–1905, 2021.
- [30] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. You’ll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the International Conference on Computer Vision (ICCV)*, pages 261–268, 2009.
- [31] F. T. Pokorny, K. Goldberg, and D. Kragic. Topological trajectory clustering with relative persistent homology. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 16–23, 2016.
- [32] C. Pérez D’Arpino, C. Liu, P. Goebel, R. Martín-Martín, and S. Savarese. Robot navigation in constrained pedestrian environments using reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1140–1146, 2021.
- [33] I. Rae, L. Takayama, and B. Mutlu. The influence of height in robot-mediated communication. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 1–8, 2013.
- [34] J. Reinhardt, L. Prasch, and K. Bengler. Back-off: Evaluation of robot motion strategies to facilitate human-robot spatial interaction. *Transactions on Human-Robot Interaction*, 10(3), 2021.
- [35] J. Roh, C. Mavrogiannis, R. Madan, D. Fox, and S. Srinivasa. Multimodal trajectory prediction via topological invariance for navigation at uncontrolled intersections. In *Proceedings of the Conference on Robot Learning*, 2020.
- [36] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras. Human motion trajectory prediction: a survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.
- [37] P. Salvini, D. Paez-Granados, and A. Billard. Safety concerns emerging from robots navigating in crowded pedestrian areas. *International Journal of Social Robotics*, 14: 441–462, 2022.
- [38] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone. Trajectron++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 683–700, 2020.
- [39] C. Schöller, V. Aravantinos, F. Lay, and A. Knoll. What the constant velocity model can teach us about pedestrian motion prediction. *IEEE Robotics and Automation Letters*, 5(2):1696–1703, 2020.
- [40] F. Shkurti and G. Dudek. Topologically distinct trajectory predictions for probabilistic pursuit. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5653–5660, 2017.
- [41] M. Sun, F. Baldini, P. Trautman, and T. Murphey. Move Beyond Trajectories: Distribution Space Coupling for Crowd Navigation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2021.
- [42] 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. *International Journal of Robotics Research*, 34(3):335–356, 2015.
- [43] N. Tsoi, M. Hussein, J. Espinoza, X. Ruiz, and M. Vázquez. Sean: Social environment for autonomous navigation. In *Proceedings of the International Conference on Human-Agent Interaction*, page 281–283, 2020.
- [44] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha. Reciprocal n-body collision avoidance. In *Robotics Research*, pages 3–19. Springer Berlin Heidelberg, 2011.
- [45] N. Walker, C. Mavrogiannis, S. S. Srinivasa, and M. Cakmak. Influencing behavioral attributions to robot motion during task execution. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.
- [46] A. Wang, C. Mavrogiannis, and A. Steinfeld. Group-based motion prediction for navigation in crowded environments. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.
- [47] N. H. Wolfinger. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography*, 24(3):323–340, 1995.
- [48] K. Yamane and C. Kurosu. Stable balance controller, March 2020. US Patent 16/375,111.
- [49] B. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J. Bagnell, M. Hebert, A. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3931–3936, 2009.