

Implicit Communication in Human-Robot Collaborative Transport

Elvin Yang
Department of Robotics
University of Michigan
Ann Arbor, MI, USA
eyy@umich.edu

Christoforos Mavrogiannis
Department of Robotics
University of Michigan
Ann Arbor, MI, USA
cmavro@umich.edu

Abstract—We focus on human-robot collaborative transport, in which a robot and a user collaboratively move an object to a goal pose. In the absence of explicit communication, this problem is challenging because it demands tight *implicit* coordination between two heterogeneous agents, who have very different sensing, actuation, and reasoning capabilities. Our key insight is that the two agents can coordinate fluently by encoding subtle, communicative signals into actions that affect the state of the transported object. To this end, we design an inference mechanism that probabilistically maps observations of joint actions executed by the two agents to a set of joint strategies of workspace traversal. Based on this mechanism, we define a cost representing the human’s uncertainty over the unfolding traversal strategy and introduce it into a model predictive controller that balances between uncertainty minimization and efficiency maximization. We deploy our framework on a mobile manipulator (Hello Robot Stretch) and evaluate it in a within-subjects lab study ($N = 24$). We show that our framework enables greater team performance and empowers the robot to be perceived as a significantly more fluent and competent partner compared to baselines lacking a communicative mechanism.

Index Terms—Human-robot collaboration; Human-robot teams; Implicit communication

I. INTRODUCTION

Recently, there has been vivid interest in developing physically capable robot partners that could assist humans in *context*-rich, dynamic and unstructured domains [52] like homes [59, 68] and manufacturing sites [34]. An important task in this space involves the *collaborative transport* of objects that might be too large or too heavy to be transported by a single agent. This task is especially challenging as it not only requires efficient and fluent coordination between the two heterogeneous partners but also the simultaneous satisfaction of geometric, kinematic, and physics constraints.

Humans often tackle physically demanding collaborative tasks like transport by fluently coordinating their physical movements with their partners [51] even without a concrete plan, with minimal explicit coordination. This capability relies on sophisticated mechanisms connecting perception and action. A prevalent theory from action understanding, commonly referred to as the “teleological stance”, highlights that agents’ actions can often be explained by an underlying goal [6, 12, 16]. This idea has inspired researchers in human-robot interaction (HRI) to develop mechanisms that communicate a robot’s




Fig. 1: Footage from our study ($N = 24$) involving the collaborative transport of an object (orange stick) by a user and a mobile manipulator in a workspace with an obstacle (red color). The robot runs our controller (IC-MPC), designed to balance functional and communicative actions in collaborative tasks.

intended goal to an observer through its actions [13, 27]. These mechanisms have produced intent-expressive robot behavior in manipulation [13], autonomous driving [48], and social robot navigation [35]. Likewise, we view communication—especially implicit communication [27], the ability to infer and convey information within physical actions—to be a critical skill of robots working in close physical collaboration with humans. Implicit communication can be low-latency, robust to environmental disturbances (e.g., noise, poor lighting conditions), and require less attention compared to explicit forms. While explicit communication remains highly relevant to team activities, implicit communication serves as an important complement that supports fluent teamwork.

To investigate the implications of implicit communication for physical human-robot teamwork, we instantiate a task of human-robot collaborative transport, where the goal of the human-robot team is to collaboratively move an object to a goal pose while avoiding collisions with static obstacles (see Fig. 1). In this task, the user is simultaneously an observer of the robot and a *dynamic actor*, persistently influencing and being influenced by the robot while it *physically collaborates* with them. While prior work in human-robot collaborative transport has emphasized fixed leadership roles for the two agents [7, 8, 31, 37, 42, 54], we consider a *dynamic negotiation* over a joint strategy of *workspace traversal*. We contribute a

control framework that leverages *implicit communication* [27] through actions influencing the state of the transported object to enable the robot to negotiate an efficient traversal with its human partner. We move beyond past work on implicit communication, where the user is either *not an actor* [13] or *not physically* collaborating with their robot partner [30, 35]. We demonstrate our framework on a mobile manipulator and evaluate it in a lab study ($N = 24$) involving the collaborative transport of an object in a workspace with an obstacle obstruction. We show that our framework outperforms baselines lacking a communicative mechanism in terms of task completion and human impressions.

In summary, we contribute:

- A formal mathematical representation of workspace traversal strategies during collaborative object transport.
- A human-inspired inference mechanism that probabilistically maps a joint human-robot action to a joint strategy of workspace traversal.
- A model predictive control framework that balances efficiency maximization and uncertainty minimization to produce fluent, efficient teamwork in physically collaborative tasks.
- Evidence from an extensive lab study ($N = 24$) suggesting that our framework results in greater team performance and positively perceived robot behaviors. Videos from the study can be found at <https://youtu.be/ONTDrobSifg>.
- Code and data from our study that could help the community iterate on our work, publicly available at https://github.com/fluentrobotics/icmpc_collab_transport.

II. RELATED WORK

We discuss relevant work on human-robot collaboration (HRC), our target domain of human-robot collaborative transport, and our technical foundation of implicit communication.

A. Human-Robot Collaboration

HRC, defined broadly as the integration of humans and robots working together towards a set of (possibly joint) goals, has attracted considerable attention in recent years [9, 18, 33, 64]. An important challenge in HRC involves determining a fluent meshing between human and robot actions that leverages the unique competences of both in an effective way [20]. This has motivated research on the design of planning and control frameworks that smoothly blend human and robot control inputs [5, 14, 43, 57]. For many applications, the human and the robot goals are well-defined; e.g., in shared control for wheelchair navigation, the role of the robot is to lead the user to a desired waypoint [9]. However, in other applications, the richness of the context results in a multiplicity of ways in which a task could be achieved. In those cases, a robot has to reason about how its own strategy for satisfying its objectives meshes with the strategy of its human partner [45, 47, 63, 69] or how its actions contribute towards a joint strategy [66]. An underlying challenge involves determining an appropriate space of strategies. Nikolaidis and Shah [40] describe an

interactive planning framework that iteratively switches role assignment to humans and robots with the goal of converging on a shared strategy. Zhao et al. [69] and Xie et al. [66] learn spaces of strategies from observations of low-level action sequences whereas Wang et al. [63] learn a latent space of strategies that enables a robot to adapt to its collaborator.

Our approach is motivated by the physical HRC task of human-robot collaborative transport. The execution of this task admits multiple possible solutions in realistic environments with obstacle obstructions. Recent work in HRC emphasizes the extraction of these solutions from team strategies observed in human teamwork. In contrast, we exploit the mathematical structure of the domain, identifying strategies of collaborative transport as classes of homotopy-equivalent trajectories of workspace traversal [26, 28, 61]. This allows for methodical enumeration of strategies in an interpretable form and in line with insights from recent studies on human-human object transport [15]. Additionally, while a body of work considers validation in virtual benchmarks, we emphasize *physical, explicit, and concurrent* human-robot collaboration [52] treating the robot as a fully embodied partner with mobile manipulation capabilities.

B. Human-Robot Collaborative Transport

Physical HRC can support applications like personal robotics, manufacturing, and construction [32, 44, 52]. Much of the recent work has focused on collaborative manipulation tasks involving a user and a manipulator, such as the rearrangement of large or heavy objects [2, 3, 17, 25, 29, 53, 55, 70] and the collaborative use of tools for sawing or bolt screwing [46].

Many real-world applications motivate the integration of manipulation and robot mobility. An important task of interest is the *collaborative transport* of an object by a user and a mobile manipulator. This task is especially challenging because in addition to human-robot coordination, it requires considerations like collision avoidance, object stability, and human ergonomics. A large body of work assumes a fixed role assignment: the human intent is estimated and the robot follows it [7, 8, 31, 37, 54]. This can be too restrictive, especially when the robot has the capability to contribute meaningfully to the task as an independent agent. To account for such situations, Nikolaidis et al. [42] develop a probabilistic controller that estimates the user's preferences and generates corrective actions to guide them to an efficient strategy. However, this implies that the robot is more equipped to guide the interaction, which might not always be the case. Instead, Ng et al. [39] allow leadership to be an emergent property of following a policy extracted from human demonstrations. Mörzl et al. [38] engineer a dynamic leadership negotiation through actions executed on the transported object by the two agents. While mobility is integrated across these works, it is often oversimplified through the assumption of predefined paths [38] or collision-free workspaces [7, 8].

In this work, we leverage the full mobility of a mobile base and explicitly account for obstacles while negotiating a strategy of workspace traversal. We develop a control

framework that enables the robot to flexibly negotiate with the user using implicit communication, realized via subtle signals encoded in velocities transmitted to the transported object. In contrast to approaches directly learning traversal policies from human demonstrations [39], our approach has the potential to handle diverse maps by reasoning about the mapping between team actions and traversal strategies. While prior work assumes heavily instrumented setups including sensorized objects [7, 8, 17, 37], we use a commodity mobile manipulator (Hello Robot Stretch) and limited additional instrumentation (a motion capture system to eliminate the influence of localization errors). Finally, we move beyond prior work in terms of validation depth by conducting an extensive lab study ($N = 24$) with in-depth insights via objective and subjective measures.

C. Implicit Communication

Any deviation from a robot’s expected behavior has the potential to convey important, task-related information to an observer [27]. This idea, rooted in conversational implicature [19] has deep relevance to HRI [21, 58, 60]. The community has long been exploring the expressive power of modalities like motion [56], gaze [1], or gestures [11]. Researchers have been building algorithmic frameworks that harness implicit communication to engineer fluent human-robot teamwork. Knepper et al. [27] formalize implicit communication as the minimization of the information entropy over a probability distribution of possible goals. Considering a social navigation domain, Mavrogiannis et al. [35] implement implicit communication into a navigation controller that conveys the robot’s intended passing side to co-navigating users. Considering a coverage task, Walker et al. [62] implement implicit communication into the generation of robot motion that influences an observer to summarize robot motion in a desirable way. In a manipulation domain, Dragan and Srivasa [13] instantiate implicit communication as *Legibility*, a property of robot motion that enables an observer to infer the target object of a manipulator’s reaching motion. Considering the collaborative game Hanabi, Liang et al. [30] demonstrate that implicit communication can be a powerful tool for teams of humans and artificial agents working on a joint task.

Prior algorithmic instantiations of implicit communication assume simplified teamwork settings. For instance, the user is often passively observing the robot [13, 62], navigating next to the robot [35], or working with a non-embodied agent [30]. In contrast, this work integrates implicit communication into physical human-robot collaboration through the task of collaborative transport. This task uniquely integrates the user and the robot as embodied, dynamic observers and actors, whose actions influence each other and impact the quality of the teamwork. We define a space of goals, corresponding to strategies of workspace traversal, which we identify using a notion of topological invariance [36]. We then develop a probabilistic inference that maps joint human-robot actions to strategies of workspace traversal. By tracking the entropy of the distribution over strategies, the robot may monitor the state of uncertainty




Fig. 2: A human (H) and a robot (R) collaboratively move an object from an initial pose p_0 to a final pose g in a workspace \mathcal{W} . An obstacle \mathcal{O} stands in their way. To avoid collisions with \mathcal{O} and reach g , they have to coordinate on a strategy of workspace traversal. In this work, we engineer implicit coordination through the velocities a and u that the human and the robot exert on the object.

over the unfolding traversal. By introducing entropy as a cost into a model predictive controller (IC-MPC), the robot is capable of balancing functional and communicative actions resulting in efficient, fluent, and positively perceived human-robot teamwork.

III. PROBLEM STATEMENT

We consider a human H and a robot R collaboratively transporting an object. The robot and the human grasp the object at a fixed height; this allows us to instantiate the problem on a planar workspace $\mathcal{W} \subseteq SE(2)$. Assuming a quasistatic setting, the object’s state $p \in \mathcal{W}$ evolves according to $p_{k+1} = f(p_k, a_k, u_k)$, where $a \in \mathcal{A}$, $u \in \mathcal{U}$ represent human and robot velocities, respectively, and k is a time index. The workspace includes a set of obstacle-occupied regions $\mathcal{O} \subset \mathcal{W}$. The goal of the human-robot team is to transport the object from an initial pose p_0 to a desired pose g in \mathcal{W} (see Fig. 2) while avoiding collisions with \mathcal{O} . We assume that the two agents do not communicate explicitly (e.g., via language), but they observe the actions of one another. Our goal is to design a control policy to enable the robot to efficiently and fluently collaborate with its human partner.

IV. BALANCING FUNCTIONAL AND COMMUNICATIVE ACTIONS IN HUMAN-ROBOT COLLABORATIVE TRANSPORT

We describe a control framework that leverages implicit communication to support efficient and fluent collaboration in human-robot collaborative transport. By reasoning about its partner’s uncertainty over the *way* the task is being executed, the robot balances between communicative, uncertainty-reducing actions, and functional, task-driven actions. This balance is not prescribed, but rather dynamically adaptive to the robot’s belief about the uncertainty of its partner.

A. Formalizing Joint Strategies of Workspace Traversal

Collaborative tasks involving multiple agents working together require consensus on a *joint strategy* ψ , i.e., a qualitatively distinct way of completing the task, out of the set of all possible joint strategies, Ψ . Often, this consensus is

not established *a priori*; rather, it is dynamically negotiated during execution. The abstraction of a joint strategy effectively captures critical domain knowledge at a representation level. While prior work on collaborative transport has emphasized *role* assignment across the team (i.e., whether the robot or the human are leading or following each other) [23, 38, 40], realistic, obstacle-cluttered environments present additional important challenges, such as the decision over *how to pass through* an obstacle-cluttered workspace.

In this work, we formalize the space of workspace traversal strategies using tools from homotopy theory [26]. The human-robot team is tasked with transporting an object from its initial pose p_0 to a final pose g , resulting in an object trajectory $\mathbf{p} : [0, 1] \rightarrow \mathcal{W}$, where $\mathbf{p}(0) = p_0$ and $\mathbf{p}(1) = g$, belonging to an appropriate space of trajectories \mathcal{P} . Obstacles, defined as the connected components of \mathcal{O} , naturally partition \mathcal{P} into equivalence classes Ψ , where each $\psi \in \Psi$ represents a distinct workspace traversal strategy under which the transported object can travel from p_0 to g , i.e.,

$$\begin{aligned} \mathcal{P} &= \bigcup_{\psi \in \Psi} \psi \\ \forall \psi^i, \psi^j \in \Psi : (\psi^i \cap \psi^j \neq \emptyset) &\implies (\psi^i = \psi^j). \\ \forall \mathbf{p}^i, \mathbf{p}^j \in \psi : \mathbf{p}^i &\sim \mathbf{p}^j \end{aligned} \quad (1)$$

These classes can be identified using a notion of topological invariance. The works of Kretschmar et al. [28], Mavrougiannis et al. [36], Vernaza et al. [61] use winding numbers to describe topological relationships between the robot and obstacles or humans navigating around it. Here, we adapt this idea to collaborative transport by enumerating the set of homotopy classes between the object trajectory and obstacles in the workspace. Specifically, for any object trajectory \mathbf{p} embedded in a space with m obstacles o_1, \dots, o_m , we can define winding numbers

$$w_i = \frac{1}{2\pi} \sum_t \Delta\theta_t^i, \quad i = 1, \dots, m, \quad (2)$$


where $\Delta\theta_t^i = \angle(p_t - o_i, p_{t-1} - o_i)$ denotes an angular displacement corresponding to the transfer of the object from p_{t-1} to p_t (see Fig. 3a). The sign of w_i represents the passing side between the object and the i -th obstacle, and its absolute value represents the number of times the object encircled the i -th obstacle. For a trajectory \mathbf{p} , the tuple of winding number signs

$$W = (\text{sign } w_1, \dots, \text{sign } w_m) \quad (3)$$


represents an equivalence class describing how the human-robot team transported the object past all obstacles in the environment. In this work, we model the space of joint strategies Ψ as set of distinct W , i.e., $|\Psi| = 2^m$.

B. Inferring Strategies of Workspace Traversal

We describe an inference mechanism that maps observations of team actions to a belief over a workspace traversal strategy. This mechanism is agnostic to the specific definition of the strategy. At time t , we assume that the robot observes the



(a) Identification of workspace traversal strategies based on path homotopy [28]. By integrating the angle of the vector of the obstacle and the object as it is being transported along a path \mathbf{p} , we extract a winding number $w(\mathbf{p})$ identifying the strategy of workspace traversal. Here, $w(\mathbf{p}^2) = w(\mathbf{p}^3)$ since both $\mathbf{p}^2, \mathbf{p}^3$ passed on the right of o_1 .



(b) Representing workspace traversal strategies as tuples of winding number signs, W . In this scene with two obstacles, there are four possible strategies represented as continuous curves. The red curve highlights a strategy corresponding to passing on the right of o_1 ($w_1 > 0$), and the left of o_2 ($w_2 < 0$). This representation is applicable to any number of obstacles.

Fig. 3: Illustration of our topological abstraction for representing strategies of workspace traversal.

joint action $\alpha = (a, u)$, the object state p , and the task context $c = (g, \mathcal{O})$. Given α, p , and c , our goal is to infer the unfolding workspace traversal strategy, ψ , i.e.,

$$\mathbb{P}(\psi \mid \alpha, p, c). \quad (4)$$

Using Bayes' rule, we can expand (4) as

$$\mathbb{P}(\psi \mid \alpha, p, c) = \frac{1}{\eta} \mathbb{P}(\alpha \mid \psi, p, c) \mathbb{P}(\psi \mid p, c), \quad (5)$$

where the left-hand side expression is the *posterior distribution* of the joint strategy ψ , and on the right-hand side, η is a normalizer across α , $\mathbb{P}(\alpha \mid \psi, p, c)$ is the *joint action likelihood distribution* and $\mathbb{P}(\psi \mid p, c)$ is a *prior distribution* of the joint strategy before observing the joint action. We can rewrite the joint action likelihood distribution as

$$\mathbb{P}(\alpha \mid \psi, p, c) = \mathbb{P}(a \mid \psi, p, c) \mathbb{P}(u \mid \psi, p, c), \quad (6)$$

since the two agents choose their actions independently.

The distribution of (4) allows the robot to represent the belief of its partner over the unfolding traversal strategy. A natural measure of uncertainty over the observer's belief regarding that strategy can be acquired by computing the information entropy of ψ , conditioned on known α, p, c :

$$H(\psi \mid \alpha, p, c) = - \sum_{\psi \in \Psi} \mathbb{P}(\psi \mid \alpha, p, c) \log \mathbb{P}(\psi \mid \alpha, p, c). \quad (7)$$

Intuitively, the higher H is, the higher the uncertainty of the user over the unfolding ψ is assumed to be.

C. Integrating Human Inferences into Robot Control

We integrate the inference mechanism of (4) into a model predictive control (MPC) algorithm by using its entropy (7) as a cost. Given the context $c = (g, \mathcal{O})$ and the object state p at time t , the goal of the MPC is to find the sequence of future robot actions \mathbf{u}^* that minimizes a cost function J over a horizon T . At every control cycle, the MPC solves the following planning problem:

$$\begin{aligned} (u_{t:t+T})^* &= \arg \min_{u_{t:t+T}} J(p_{t:t+T}, u_{t:t+T}) \\ \text{s.t. } p_{k+1} &= f(p_k, a_k, u_k), \\ a_k &\in \mathcal{A} \\ u_k &\in \mathcal{U} \end{aligned} \quad (8)$$

We split J into a running cost J_k and a terminal cost J_T

$$J(p_{t:t+T}, u_{t:t+T}) = \sum_{k=0}^T \gamma^k J_k(p_{t+k}, u_{t+k}) + J_T(p_{t+T}, u_{t+T}) \quad (9)$$

where γ is a discount factor, and the terminal cost penalizes distance from the object’s goal pose g :

$$J_T(p_{t+k}, u_{t+k}) = \|p_{t+k} - g\|^2. \quad (10)$$

The running cost J_k is a weighted sum of two terms, i.e.,

$$J_k(p_{t+k}, u_{t+k}) = w_{obs} J_{obs}(p_{t+k}, u_{t+k}) + w_{ent} J_{ent}(p_{t+k}, u_{t+k}), \quad (11)$$

where

$$J_{obs}(p_{t+k}, u_{t+k}) = \max \left(0, -\log \left(\min_{o \in \mathcal{O}} \frac{\|p_{t+k} - o\|}{\delta} \right) \right), \quad (12)$$

is a collision avoidance cost penalizing proximity to obstacles, δ is a clearance threshold, J_{ent} is a cost proportional to the entropy defined in (7), and w_{obs} , w_{ent} are weights.

We refer to this control framework as *Implicit Communication MPC*, or **IC-MPC**. At every control cycle, IC-MPC plans a future robot trajectory that balances between functional objectives (collision avoidance, progress to goal) and communicative objectives (minimizing the partner’s uncertainty over the upcoming joint strategy). The robot executes the first action u_t from the planned trajectory and then replans. This process is repeated in fixed control cycles until the task is completed.

V. USER STUDY

We conducted an IRB-approved, within-subjects user study (U-M HUM00254044) in which a user collaborates with a robot to jointly transport an object to a designated pose. Each user experienced the same set of conditions, each corresponding to a collaborative algorithm running on the robot: ours, and two baselines. Through mailing lists, we recruited 24

participants (4 female, 18 male, 2 other), aged 18-29 from a university population. On average, participants rated of their familiarity with robotics technology as 4.1 (SD = 0.74) on a scale from 1 (not at all familiar) to 5 (very familiar). The study lasted 45 minutes, and each participant received \$20 of compensation.

A. Experiment Design

Task Description. The user and the robot hold opposite ends of an object (a wooden stick) and transport it together from an initial pose to a goal pose. Users collaborate with each algorithm three times to ensure they experience a diverse range of interactions. At the start of each of the three trials, the user and the robot stand in predetermined configurations (Fig. 4), following the same fixed order for all algorithms.

Algorithms. We compare the performance of our framework (**IC-MPC**) against two baselines:

- *Vanilla-MPC*: A purely functional ablation of IC-MPC with no uncertainty-minimizing objective ($w_{ent} = 0$).
- *VRNN* [39]: A learning-based path planner based on a Variational Recurrent Neural Network that predicts the most likely future path of the object based on human demonstrations. The robot takes actions to track path predictions as closely as possible.

Metrics. We evaluate performance in terms of:

Objective Metrics:

- *Success rate*, defined as the proportion of successful trials. A trial is successful if the following conditions hold: a) any part of the object reaches the goal; b) the object stays within the workspace boundary; c) the object stays parallel to the ground—i.e., not lifted above the obstacle or dropped; d) there are no obstacle collisions.
- *Completion time*, measured as the time taken for the object to reach the goal.
- *Acceleration*, measured as the average acceleration of the human trajectory similar to prior work [35].

Subjective Metrics:

- *Warmth*, *Competence*, and *Discomfort*, measured using the RoSAS scale [10]. We use the original form of the scale: a single list of 18 items in randomized order and a nine-point response scale from 1 = “Definitely not associated” to 9 = “Definitely associated”.
- *Fluency*, measured using the *Fluency in HRI* scale [20]. To reduce fatigue, we use seven of eight items that were validated with objective measures in the original paper [20], with one item removed as recommended. The items are presented in randomized order with a seven-point response scale from 1 = “Strongly disagree” to 7 = “Strongly agree”.

Hypotheses. Our insight is that implicit communication will have observable implications for the collaboration quality of the human-robot team, at an objective and subjective level. We formalize this insight into the following two hypotheses:

- H1** IC-MPC is more effective at completing this task in collaboration with a user compared to Vanilla-MPC and VRNN as measured by the objective metrics.




Fig. 4: Users experienced three different starting configurations with each robot algorithm. From left to right: the user and the robot stand side-by-side; the user stands directly behind the robot; the user stands directly in front of the robot. For the third configuration, the user faces toward the goal and holds the object behind their back.

H2 IC-MPC is viewed more favorably as a collaborator compared Vanilla-MPC and VRNN as measured by users’ responses to questionnaires containing the subjective metrics.

B. System Development

Experimental Setup. We deploy all algorithms on Stretch RE2 from Hello Robot [24], a mobile manipulator with a differential-drive mobile base and a prismatic arm. To avoid excessive torque on the robot’s end-effector, the transported object is a lightweight stick of 0.914 m length and 0.05 kg mass. The robot’s wrist motors are turned off to allow free rotation of the end effector. The lift of the robot arm is static throughout the task to constrain the object’s movement to $SE(2)$. Across our experiments, the team operates in a workspace with area $2.8 \times 5.6 m^2$. To study the coordination of the human-robot team over a discrete decision, a single static obstacle of area $0.15 \times 0.15 m^2$ is placed in the center of the workspace (see Fig. 1). The workspace is fitted with an overhead Optitrack Flex 13 motion-capture (mocap) system that continuously streams poses and velocities of the robot and the user (via a construction-style helmet) at 120 Hz.

Human modeling. Since our study involved an environment with a single obstacle, we set $\Psi = \{\text{LEFT}, \text{RIGHT}\}$, corresponding respectively to $w < 0$ and $w > 0$. We instantiated the strategy inference using analytical models of the prior distribution and joint action likelihood distribution (5). However, it is possible to approximate this inference using data-driven techniques, e.g., by learning them from datasets of demonstrations [15, 39].

We model the prior distribution over the joint strategy as a function of the winding number w . Without loss of generality, when p_0, o, g are collinear, we consider the obstacle to have been passed when $|w| \geq \frac{1}{4}$:

$$\begin{aligned} \mathbb{P}(\text{LEFT} | p, c) &= \max(0, \min(0.5 - 2w, 1)) \\ \mathbb{P}(\text{RIGHT} | p, c) &= \max(0, \min(0.5 + 2w, 1)) \end{aligned} \quad (13)$$

Before the obstacle is passed, the action likelihood distribution models the most likely action as the velocity that




Fig. 5: Workspace traversal strategy inference. The prior distribution for $\mathbb{P}(\text{LEFT} | p, c)$ is shown as a colormap in the background, and the mode of the action likelihood distribution for $\mathbb{P}(a | \text{LEFT}, p, c)$ is shown using gray arrows in the foreground.

maximizes change in the winding number w . We approximate this as

$$\begin{aligned} \mathbb{P}(a | \text{LEFT}, p, c) &\propto \exp\left(a \cdot R\left(\frac{\pi}{3}\right) \vec{p}\vec{o}\right) \\ \mathbb{P}(a | \text{RIGHT}, p, c) &\propto \exp\left(a \cdot R\left(-\frac{\pi}{3}\right) \vec{p}\vec{o}\right) \end{aligned} \quad (14)$$

where $R(\cdot)$ is a 2D rotation matrix. After the obstacle is passed, the most likely action is instead in the direction of the goal. We illustrate the prior and action likelihood distributions for $\psi = \text{LEFT}$ in Fig. 5. The same model of the action likelihood distribution is used for human and robot actions.

Prior to the user study, we evaluated our analytical models using the simulation dataset provided by Ng et al. [39]. We considered the subset of dataset environments with a single obstacle and annotated each trajectory with a ground truth joint strategy label. The prior and action likelihood distributions achieved 98.9% and 72.4% accuracy in this subset of the dataset, respectively.

MPC. All algorithms are implemented using an open-source package [4] for Model Predictive Path Integral (MPPI) control [65]. The same discount factor $\gamma = 0.95$, planning horizon $15 \times 0.25 s$, and number of rollout samples (100) is used across all algorithms. IC-MPC and Vanilla-MPC use obstacle clearance threshold $\delta = 0.5$ and weight $w_{obs} = 1.0$. IC-MPC uses entropy weight $w_{ent} = 1.0$. These parameters were empirically selected to ensure good performance in our environment during pilot hardware trials, and different workspace configurations may require different parameter combinations. Observations of human velocities are downsampled to 10 Hz from the motion capture, and a constant velocity model is used for the human motion prediction rollouts [50]. All algorithms run on an Intel i7-13700 CPU at 15 Hz.

VRNN. We adapt MPC to track VRNN’s prediction of the object’s future path as closely as possible as in Ng et al. [39]. However, we found that path predictions occasionally have undefined orientations, e.g., the output of the model contains states for which $|\cos \theta| > 1$. Consequently, we use only the translational component of the path predictions. Since VRNN is designed to avoid obstacles and reach the goal, we use a running cost based on Euclidean distance to the predicted path and no terminal cost.

TABLE I: Summary of objective metrics. Success rate is calculated out of 72 total trials for each algorithm. Mean and standard deviation of Time and Acceleration consider only successful trials.

Metric	IC-MPC	Vanilla-MPC	VRNN
Success rate (%) \uparrow	98.6	88.9	51.4
Completion time (s) \downarrow	28.54 (1.26)	28.86 (3.68)	24.98 (1.67)
Acceleration (m/s^2) \downarrow	0.67 (0.15)	0.69 (0.17)	0.71 (0.14)

C. Procedure

Before the study, participants received information about the purpose of the study and data collection. They provided informed consent, acknowledging that their participation was voluntary and that they would not receive descriptions of the robot algorithms until the end of the study. They were then brought into the motion capture area and instructed to wear a helmet to be tracked by the mocap system. They were instructed to hold the object with both hands, keeping it parallel to the ground, and to collaborate with the robot in whichever way felt natural to them, including moving around the robot. To get familiarized with the interaction and the task, participants completed three practice rounds in which the robot was teleoperated before the main experiment.

For the main portion of the experiment, participants experienced each condition (a different algorithm embodied on the robot) three consecutive times. After each condition, participants were informed of the number of times they succeeded at the task with the robot and directed to respond to a questionnaire about their latest experience. To account for ordering effects, we randomized items in both questionnaires every time they were presented. Additionally, we performed counterbalancing of the conditions by randomly assigning participants to one of six possible orderings of the three conditions.

D. Analysis

We found that objective and subjective metrics did not uniformly pass the Shapiro-Wilk test of normality. Thus, for consistency, we use the non-parametric Friedman test to detect effects of the robot algorithms on dependent variables and the non-parametric paired Wilcoxon signed-rank test with Holm-Bonferroni corrections [22] for post-hoc pairwise comparison tests. Effect sizes are reported using Kendall’s coefficient of concordance (denoted as W_k to disambiguate it from the test statistic of the Wilcoxon signed-rank test) and Cohen’s d .

H1. We report a summary of objective metrics for each algorithm in Table I. IC-MPC exhibited substantially higher *success rate* compared to both Vanilla-MPC and VRNN. We found a significant effect of the robot algorithm on *completion time* ($\chi^2(2) = 20.61, p < 0.001, W = 0.45$). Post-hoc tests indicated that, of successful trials, IC-MPC completed the task *slower* compared to VRNN ($W = 3.0, p < 0.001, d = 2.4$). We found no significant difference in *completion time* between IC-MPC and Vanilla-MPC ($W = 92.0, p = 0.17, d = -0.11$). Finally, of successful trials, we found no significant effect of the robot algorithm on *acceleration* ($\chi^2(2) = 3.73, p = 0.15, W = 0.08$). Thus, we find partial support for **H1**.

TABLE II: Mean and standard deviation of subjective metrics. * $p < .05$, ** $p < .01$, *** $p < .001$

Metric	IC-MPC	Vanilla-MPC	VRNN
Warmth [10] \uparrow	3.44 (1.89)	3.12 (1.98)	2.99 (1.82)
Competence [10] \uparrow	6.06 (1.86)	5.15 (1.82)*	4.02 (1.88)***
Discomfort [10] \downarrow	2.15 (1.29)	2.86 (1.84)*	3.22 (1.49)***
Fluency [20] \uparrow	5.73 (1.02)	4.64 (1.41)**	3.67 (1.49)***

TABLE III: Friedman test results on subjective measures.

Metric	$\chi^2(2)$	p	W_k
Warmth [10]	3.85	0.146	0.08
Competence [10]	18.86	< 0.001	0.39
Discomfort [10]	15.46	< 0.001	0.32
Fluency [20]	28.15	< 0.001	0.59

H2. We report subjective measures in Table II and related test statistics in Table III. Prior to testing, we evaluated Cronbach’s alpha within each RoSAS subscale (*warmth* $\alpha = 0.90$, *competence* $\alpha = 0.92$, *discomfort* $\alpha = 0.85$) and within the *Fluency* scale ($\alpha = 0.95$) and found that all subscales had high internal consistency. Individual users’ responses to items within each subscale were subsequently averaged for analysis.

We found a significant effect of the robot algorithm on users’ perception of *competence*, *discomfort*, and *fluency*. Post-hoc tests found that IC-MPC was judged by users as: significantly more *competent* compared to Vanilla-MPC ($W = 55, p = 0.021, d = 0.49$) and VRNN ($W = 26, p < 0.001, d = 1.09$); significantly less *discomforting* compared to Vanilla-MPC ($W = 30, p = 0.018, d = -0.44$) and VRNN ($W = 13.5, p < 0.001, d = -0.76$); a significantly more *fluent* collaborator compared to Vanilla-MPC ($W = 28, p = 0.002, d = 0.89$) and VRNN ($W = 11, p < 0.001, d = 1.61$). No statistically significant effect was found on users’ perception of *warmth*. Thus, we find partial support for **H2**.

E. Discussion

Our hypotheses were motivated by a supposition that an absence of communication would create situations in which the two agents attempt to follow contradictory strategies, resulting in longer task duration, sudden movements to avoid collisions, or task failure altogether. We believed these situations would have a negative impact on users’ opinion of the robot as a teammate. To better understand our findings, we examine motion capture data collected during the study.

The spatial distribution of object trajectories (Fig. 6) reveals that teams took *wider* paths around the obstacle when the robot was running IC-MPC compared to Vanilla-MPC and VRNN. This difference in behavior may explain our findings for *success rate* and *completion time*. Without implicit communication mechanisms to resolve ambiguity or achieve consensus on traversal strategy, both baselines would often drive straight towards the goal and attempt to pass the obstacle from directions opposite the user, leading to collisions. By acting *earlier* to take *wider* paths compared to Vanilla-MPC and VRNN, IC-MPC reduced uncertainty about the joint strategy faster than baselines (Fig. 7), thereby reducing the chance of




Fig. 6: Spatial distribution of object trajectories within the workspace during the user study, including failure cases, itemized per algorithm. IC-MPC exhibits an almost uniform split between right and left, whereas baselines show mixed performance, including undesirable zig-zagging effects, an artifact of increased uncertainty over the unfolding traversal strategy.

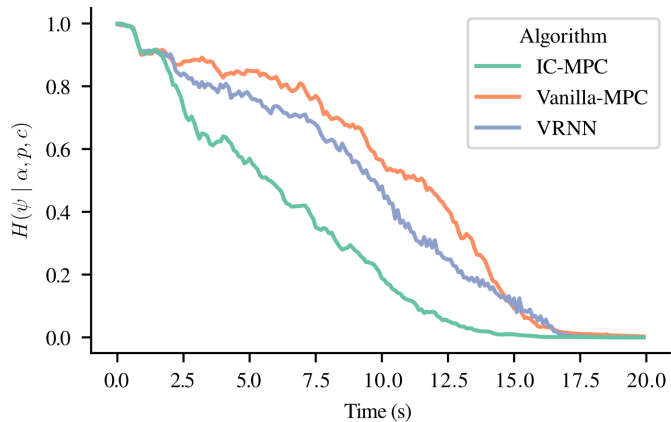


Fig. 7: Entropy over the workspace traversal strategy as a proxy for strategy uncertainty, averaged across all trials for each algorithm. By directly minimizing the entropy, IC-MPC accelerates consensus on a traversal strategy. This reduces undesirable zig-zagging artifacts, present in the execution of baselines (see Fig. 6).

similar collisions. As wider paths are longer than more direct paths, teams took more time on average to complete the task when the robot was running IC-MPC compared to VRNN. Similarity between average *completion time* of IC-MPC and Vanilla-MPC is coincidental: Vanilla-MPC had a tendency to slow to a near stop as it approached close to the obstacle.

We speculated that users would make sudden movements to avoid collisions caused by disagreement in joint strategy, particularly when collaborating with baselines that lack a communicative mechanism. However, in this study, we found no significant difference in the average *acceleration* of the human across robot algorithms. We made two observations that could serve as possible explanations. First, the maximum speed of the Stretch RE2 (0.3 m/s) is well below users’ average walking speed, which caused them to consistently alternate between taking a step and pausing to wait for the robot to cover more distance, regardless of the algorithm. Second, the slow speed of the robot gave users a relatively long time to change course to avoid a collision, if they decided to do so. Thus, *sudden* movements were uncommon.

Users noticed qualitative differences among algorithms. In open-ended responses, they described VRNN as “unpredictable” and “indecisive”. One user described Vanilla-MPC as “a bad teammate that only does what they think is right”. Two users who interacted with IC-MPC after one or both baselines

whose comments were comparative in nature wrote that IC-MPC “felt more natural” and that “the collaboration on the task was a lot more seamless in this series of attempts”.

In contrast to *competence*, *discomfort*, and *fluency*, we did not find any significant effect of robot algorithm on users’ perception of *warmth*. This is not surprising, as we did not design IC-MPC or the robot to be anthropomorphic or socially expressive. At the start of each study session, we intentionally provided the vague instruction to “collaborate in whichever way feels natural”. However, we received informal feedback from several participants that aspects of the interaction, including communication with the robot, *did not feel natural* or *intuitive*. Participants expressed confusion about *how* they could communicate with the robot and whether the robot was acknowledging, understanding, or ignoring what they were trying to communicate. Designing the robot to be expressive may facilitate interactions that are perceived as more natural.

VI. LIMITATIONS

While our decision-making framework is agnostic to the types of actions that agents are executing, in this work we considered a quasistatic setting integrating only velocities that the agents transmit to the transported object. Future work will integrate actions stemming from additional modalities, such as force, language, body posture, eye gaze, and gestures. Additionally, while our framework is not prescriptive on the space of joint strategies Ψ , this work emphasizes the topological relationships between the human-robot team and the obstacles around it. Future work will integrate additional strategy attributes, such as leadership [38], and *timing* of critical maneuvers. Our control implementation was based on a flexible and practical MPC framework, but alternative approaches could be explored, such as POMDPs [41]. Our controller demonstrated practical performance assuming a simple model of human motion prediction (constant velocity); future work will integrate more fine-grained prediction models [49, 67]. The relatively low max base speed (0.3 m/s) and payload (2 kg) of our robot might have influenced the types of interactions with users. We aspire to deploy our framework on a robot with higher payload and maximum speed to expand to tasks involving heavier objects and allow for movement across all 6 object DoFs. We also plan on expanding to more complex environments with more obstacles and dynamic agents.

REFERENCES

- [1] H. Admoni, A. Dragan, S. S. Srinivasa, and B. Scasselati. Deliberate delays during robot-to-human handovers improve compliance with gaze communication. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 49–56, 2014.
- [2] D. J. Agravante, A. Cherubini, A. Bussy, P. Gergondet, and A. Kheddar. Collaborative human-humanoid carrying using vision and haptic sensing. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 607–612, 2014.
- [3] Z. Al-Saadi, Y. M. Hamad, Y. Aydin, A. Kucukyilmaz, and C. Basdogan. Resolving conflicts during human-robot co-manipulation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 243–251, 2023.
- [4] ARM Lab, University of Michigan. Pytorch MPPI implementation, 2024. URL https://github.com/UM-ARM-Lab/pytorch_mppi.
- [5] R. M. Aronson and E. S. Short. Intentional user adaptation to shared control assistance. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 4–12, 2024.
- [6] C. L. Baker, R. Saxe, and J. B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [7] A. Bussy, P. Gergondet, A. Kheddar, F. Keith, and A. Crosnier. Proactive behavior of a humanoid robot in a haptic transportation task with a human partner. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 962–967, 2012.
- [8] A. Bussy, A. Kheddar, A. Crosnier, and F. Keith. Human-humanoid haptic joint object transportation case study. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3633–3638, 2012.
- [9] T. Carlson and J. del R. Millan. Brain-controlled wheelchairs: A robotic architecture. *IEEE Robotics & Automation Magazine*, 20(1):65–73, 2013.
- [10] C. M. Carpinella, A. B. Wyman, M. A. Perez, and S. J. Stroessner. The robotic social attributes scale (RoSAS) development and validation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 254–262, 2017.
- [11] E. J. Carter, M. N. Mistry, G. P. K. Carr, B. A. Kelly, and J. K. Hodgins. Playing catch with robots: Incorporating social gestures into physical interactions. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 231–236, 2014.
- [12] G. Csibra and G. Gergely. ‘Obsessed with goals’: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60–78, 2007.
- [13] A. D. Dragan and S. Srinivasa. Integrating human observer inferences into robot motion planning. *Autonomous Robots*, 37(4):351–368, 2014.
- [14] A. D. Dragan and S. S. Srinivasa. A policy-blending formalism for shared control. *The International Journal of Robotics Research*, 32(7):790–805, 2013.
- [15] S. Freeman, S. Moss, J. L. Salmon, and M. D. Killpack. Classification of co-manipulation modus with human-human teams for future application to human-robot systems. *ACM Transactions on Human-Robot Interaction*, 13(4), 2024.
- [16] G. Gergely, Z. Nádasdy, G. Csibra, and S. Bíró. Taking the intentional stance at 12 months of age. *Cognition*, 56(2):165–193, 1995.
- [17] M. Gienger, D. Ruiken, T. Bates, M. Regaieg, M. Meißner, J. Kober, P. Seiwald, and A.-C. Hildebrandt. Human-robot cooperative object manipulation with contact changes. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1354–1360, 2018.
- [18] D. Gopinath, S. Jain, and B. D. Argall. Human-in-the-loop optimization of shared autonomy in assistive robotics. *IEEE Robotics and Automation Letters*, 2(1):247–254, 2017.
- [19] H. P. Grice. Logic and conversation. pages 41–58, 1975.
- [20] G. Hoffman. Evaluating fluency in human–robot collaboration. *IEEE Transactions on Human-Machine Systems*, 49(3):209–218, 2019.
- [21] G. Hoffman, T. Bhattacharjee, and S. Nikolaidis. Inferring human intent and predicting human action in human–robot collaboration. *Annual Review of Control, Robotics, and Autonomous Systems*, 7:73–95, 2024.
- [22] S. Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, pages 65–70, 1979.
- [23] N. Jarrassé, V. Sanguineti, and E. Burdet. Slaves no longer: review on role assignment for human–robot joint motor action. *Adaptive Behavior*, 22(1):70–82, 2014.
- [24] C. C. Kemp, A. Edsinger, H. M. Clever, and B. Matulevich. The design of stretch: A compact, lightweight mobile manipulator for indoor human environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3150–3157, 2022.
- [25] W. Kim, J. Lee, L. Peternel, N. Tsagarakis, and A. Ajoudani. Anticipatory robot assistance for the prevention of human static joint overloading in human–robot collaboration. *IEEE Robotics and Automation Letters*, 3(1):68–75, 2018.
- [26] 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.
- [27] R. A. Knepper, C. Mavrogiannis, J. Proft, and C. Liang. Implicit communication in a joint action. In *Proceedings of the ACM/IEEE International Conference on Human-*

- Robot Interaction (HRI)*, pages 283–292, 2017.
- [28] 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.
- [29] A. Kucukyilmaz and I. Issak. Online identification of interaction behaviors from haptic data during collaborative object transfer. *IEEE Robotics and Automation Letters*, 5(1):96–102, 2019.
- [30] C. Liang, J. Proft, E. Andersen, and R. A. Knepper. Implicit communication of actionable information in human-ai teams. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI)*, page 1–13, 2019.
- [31] B. G. C. Lima, E. Ferrentino, P. Chiacchio, and M. Vento. Assistive force control in collaborative human-robot transportation. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, pages 2646–2651, 2023.
- [32] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O’Malley. A Review of Intent Detection, Arbitration, and Communication Aspects of Shared Control for Physical Human–Robot Interaction. *Applied Mechanics Reviews*, 70(1):010804, 2018.
- [33] J. Mainprice and D. Berenson. Human-robot collaborative manipulation planning using early prediction of human motion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 299–306, 2013.
- [34] E. Matheson, R. Minto, E. G. Zampieri, M. Faccio, and G. Rosati. Human–robot collaboration in manufacturing applications: A review. *Robotics*, 8(4):100, 2019.
- [35] C. Mavrogiannis, P. Alves-Oliveira, W. Thomason, and R. A. Knepper. Social momentum: Design and evaluation of a framework for socially competent robot navigation. *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 11(2):1–37, 2022.
- [36] C. Mavrogiannis, K. Balasubramanian, S. Poddar, A. Gandra, and S. S. Srinivasa. Winding through: Crowd navigation via topological invariance. *IEEE Robotics and Automation Letters*, 8(1):121–128, 2023.
- [37] E. Mielke, E. Townsend, D. Wingate, J. L. Salmon, and M. D. Killpack. Human-robot planar co-manipulation of extended objects: data-driven models and control from human-human dyads. *Frontiers in Neurorobotics*, 18, 2024.
- [38] A. Mörtl, M. Lawitzky, A. Kucukyilmaz, M. Sezgin, C. Basdogan, and S. Hirche. The role of roles: Physical cooperation between humans and robots. *The International Journal of Robotics Research*, 31(13):1656–1674, 2012.
- [39] E. Ng, Z. Liu, and M. Kennedy. It takes two: Learning to plan for human-robot cooperative carrying. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 7526–7532, 2023.
- [40] S. Nikolaidis and J. Shah. Human-robot cross-training: Computational formulation, modeling and evaluation of a human team training strategy. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 33–40, 2013.
- [41] S. Nikolaidis, D. Hsu, and S. Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.
- [42] S. Nikolaidis, D. Hsu, and S. Srinivasa. Human-robot mutual adaptation in collaborative tasks: Models and experiments. *The International Journal of Robotics Research*, 36(5-7):618–634, 2017.
- [43] S. Nikolaidis, Y. X. Zhu, D. Hsu, and S. Srinivasa. Human-robot mutual adaptation in shared autonomy. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 294–302, 2017.
- [44] U. E. Ogenyi, J. Liu, C. Yang, Z. Ju, and H. Liu. Physical human–robot collaboration: Robotic systems, learning methods, collaborative strategies, sensors, and actuators. *IEEE Transactions on Cybernetics*, 51(4):1888–1901, 2019.
- [45] R. Pandya, M. Zhao, C. Liu, R. Simmons, and H. Admoni. Multi-agent strategy explanations for human-robot collaboration. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 17351–17357, 2024.
- [46] L. Peternel, N. Tsagarakis, and A. Ajoudani. A human–robot co-manipulation approach based on human sensorimotor information. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 25(7):811–822, 2017.
- [47] L. Peters, D. Fridovich-Keil, C. J. Tomlin, and Z. N. Sunberg. Inference-based strategy alignment for general-sum differential games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1037–1045, 2020.
- [48] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for autonomous cars that leverage effects on human actions. In *Proceedings of Robotics: Science and Systems (RSS)*, 2016.
- [49] T. Salzmann, L. Chiang, M. Ryll, D. Sadigh, C. Parada, and A. Bewley. Robots that can see: Leveraging human pose for trajectory prediction. *IEEE Robotics and Automation Letters*, 8(11):7090–7097, 2023.
- [50] 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.
- [51] N. Sebanz, H. Bekkering, and G. Knoblich. Joint action: bodies and minds moving together. *Trends in Cognitive Sciences*, 10(2):70–76, 2006.
- [52] M. Selvaggio, M. Cagnetti, S. Nikolaidis, S. Ivaldi, and B. Siciliano. Autonomy in physical human-robot interaction: A brief survey. *IEEE Robotics and Automation Letters*, 6(4):7989–7996, 2021.
- [53] Y. S. Shao, T. Li, S. Keyvanian, P. Chaudhari, V. Kumar,

- and N. Figueroa. Constraint-Aware Intent Estimation for Dynamic Human-Robot Object Co-Manipulation. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [54] J. E. Solanes, L. Gracia, P. Munoz-Benavent, J. V. Miro, M. G. Carmichael, and J. Tornero. Human-robot collaboration for safe object transportation using force feedback. *Robotics and Autonomous Systems*, 107:196–208, 2018.
- [55] T. Stouraitis, I. Chatzinikolaïdis, M. Gienger, and S. Vijayakumar. Online hybrid motion planning for dyadic collaborative manipulation via bilevel optimization. *IEEE Transactions on Robotics*, 36(5):1452–1471, 2020.
- [56] L. Takayama, D. Dooley, and W. Ju. Expressing thought: improving robot readability with animation principles. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 69–76, 2011.
- [57] V. V. Unhelkar, S. Li, and J. A. Shah. Decision-making for bidirectional communication in sequential human-robot collaborative tasks. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 329–341, 2020.
- [58] J. Urakami and K. Seaborn. Nonverbal cues in human-robot interaction: A communication studies perspective. *ACM Transactions on Human-Robot Interaction*, 12(2), Mar. 2023.
- [59] T. van der Zant and L. Iocchi. Robocup@ home: Adaptive benchmarking of robot bodies and minds. In *International Conference on Social Robotics*, pages 214–225, 2011.
- [60] G. Venture and D. Kulić. Robot expressive motions: A survey of generation and evaluation methods. *ACM Transactions on Human-Robot Interaction*, 8(4), 2019.
- [61] P. Vernaza, V. Narayanan, and M. Likhachev. Efficiently Finding Optimal Winding-Constrained Loops in the Plane. In *Proceedings of Robotics: Science and Systems (RSS)*. 2013.
- [62] N. Walker, C. Mavrogiannis, S. Srinivasa, and M. Cakmak. Influencing behavioral attributions to robot motion during task execution. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2021.
- [63] C. Wang, C. Pérez-D’Arpino, D. Xu, L. Fei-Fei, K. Liu, and S. Savarese. Co-GAIL: Learning diverse strategies for human-robot collaboration. In *Proceedings of the Conference on Robot Learning (CoRL)*, pages 1279–1290, 2022.
- [64] R. Wilcox, S. Nikolaidis, and J. Shah. Optimization of temporal dynamics for adaptive human-robot interaction in assembly manufacturing. In *Proceedings of Robotics: Science and Systems (RSS)*, 2012.
- [65] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou. Aggressive driving with model predictive path integral control. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 1433–1440, 2016.
- [66] A. Xie, D. Losey, R. Tolsma, C. Finn, and D. Sadigh. Learning latent representations to influence multi-agent interaction. In *Proceedings of the Conference on Robot Learning (CoRL)*, pages 575–588, 2021.
- [67] M. S. Yasar, M. M. Islam, and T. Iqbal. Posetron: Enabling close-proximity human-robot collaboration through multi-human motion prediction. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, page 830–839, 2024.
- [68] S. Yenamandra, A. Ramachandran, M. Khanna, K. Yadav, D. S. Chaplot, G. Chhablani, A. Clegg, T. Gervet, V. Jain, R. Partsey, R. Ramrakhya, A. Szot, T.-Y. Yang, A. Edsinger, C. Kemp, B. Shah, Z. Kira, D. Batra, R. Mottaghi, Y. Bisk, and C. Paxton. The homerobot open vocab mobile manipulation challenge. In *Advances in Neural Information Processing Systems (NeurIPS): Competition Track*, 2023. URL https://aihabitat.org/challenge/2023_homerobot_ovmm/.
- [69] M. Zhao, R. Simmons, and H. Admoni. Coordination with humans via strategy matching. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9116–9123, 2022.
- [70] T. Zheng, M. Bujarbaruah, Y. R. Stürz, and F. Borrelli. Safe human-robot collaborative transportation via trust-driven role adaptation. *Proceedings of the American Control Conference (ACC)*, pages 22–27, 2022.