

Legible and Proactive Robot Planning for Prosocial Human-Robot Interactions

Jasper Goldenbott Karen Leung

Abstract—Humans have a remarkable ability to fluently engage in joint collision avoidance in crowded navigation tasks despite the complexities and uncertainties inherent in human behavior. Underlying these interactions is a mutual understanding that (i) individuals are *prosocial*, that is, there is equitable responsibility in avoiding collisions, and (ii) individuals should behave *legibly*, that is, move in a way that clearly conveys their intent to reduce ambiguity in how they intend to avoid others. Toward building robots that can safely and seamlessly interact with humans, we propose a general robot trajectory planning framework for synthesizing legible and proactive behaviors and demonstrate that our robot planner naturally leads to prosocial interactions. Specifically, we introduce the notion of a *markup factor* to incentivize legible and proactive behaviors and an *inconvenience budget constraint* to ensure equitable collision avoidance responsibility. We evaluate our approach against well-established multi-agent planning algorithms and show that using our approach produces safe, fluent, and prosocial interactions. We demonstrate the real-time feasibility of our approach with human-in-the-loop simulations. Project page can be found at <https://uw-ctrl.github.io/phri/>.

I. INTRODUCTION

Robots are becoming increasingly pervasive in our everyday lives, and more so in settings where they must interact and navigate through human counterparts. From autonomous driving to autonomous food deliveries, robots are required to safely and seamlessly interact with humans despite the myriad of uncertainties they face such as human intent, human preferences, and environmental factors. Indeed, a significant research effort has been dedicated to human behavior prediction (see [1] for a survey). Yet in highly dense interactions (e.g., Shibuya Crossing in Tokyo, Japan), humans are surprisingly remarkable in avoiding collision. The reason is, that humans are self-preserving [2] and will engage in *joint collision avoidance*. Joint collision avoidance controllers have been shown to successfully replicate collision-free and fluent (dense) multi-agent interactions [3], [4], [5]. While we may not be able to perfectly predict how humans behave, we can at least expect other humans to exhibit some degree of cooperativeness in avoiding collisions. In this work, we ask the question, *despite uncertainty in human behaviors, can we achieve safe and fluent human-robot interactions by leveraging the fact that humans are self-preserving?*

Our main hypothesis is that if the robot is able to indicate

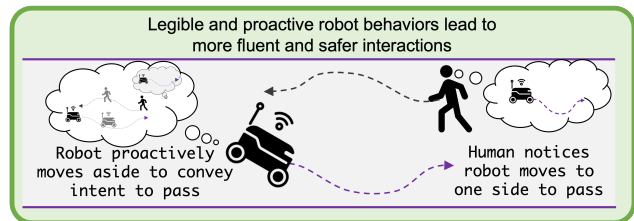


Fig. 1: Our approach enables robots to interact legibly and proactively, providing oncoming humans with sufficient time to engage cooperatively in collision avoidance. Our experiments show that this leads to fluent, safe, and prosocial interactions.

early to interacting humans its *intent* to avoid collision (e.g., pass to the right), then this will (i) provide the humans sufficient warning to adjust their plans to avoid collisions, (ii) remove ambiguity in how the joint collision avoidance maneuver should occur, and (iii) result in *prosocial* interactions where everyone equitably compromises their performance to benefit the group. In this work, we propose a robot planning framework to create *legible*—the ability to infer the intent from motion—and *proactive*—the ability to influence a situation by causing something to happen rather than responding to it after it has happened—robot motions. Our approach adds a number of simple additions to a standard robot trajectory optimization used widely for robot motion planning, making it a general and flexible approach for synthesizing legible and proactive behaviors. Specifically, we introduce the notion of a *markup factor* to encourage the robot to take nontrivial actions earlier rather than later, and the concept of an *inconvenience budget* to keep the robot’s motion directed to its goal and to promote an equitable distribution of collision avoidance responsibility.

Contributions. Our contributions are four-fold: (i) We introduce an inconvenience budget constraint as a mechanism to produce prosocial interactions. The constraint prevents the robot from making inequitable collision avoidance maneuvers, such as a sharp swerve, and keeps its motion aligned with the goal direction. (ii) We present a markup factor into the robot’s cost function to encourage legible and proactive behaviors. With the markup factor, the robot is encouraged to rotate earlier to pass oncoming agents, therefore indicating its intent to the human. (iii) We propose a robot planner utilizing an iterated best response algorithm to capture interaction dynamics with human agents and demonstrate its real-time applicability via human-in-the-loop experiments. (iv) We demonstrate that our proposed legible and proactive planner results in safer, more fluent, and more prosocial interactions compared to other social navigation approaches.

Organization. In Section II, we discuss related work, fol-

This work was supported by the UW + Amazon Science Hub Research Award.

University of Washington, Department of Aeronautics and Astronautics {jgolden, kymleung}@uw.edu

We would like to acknowledge Purnanand Elango, Annika Singh, Chris Hayner, and Nia Jetter for their insightful discussions on trajectory optimization, and Isaac Remy for his help on the experiments.

lowed by a problem formulation in Section III. We describe our proposed approach in Section IV, and discuss experimental results in Section V. We conclude and present exciting directions for future work in Section VI.

II. RELATED WORK

We discuss some related work on social robot navigation, with a focus on modeling other agents, social competency, and legibility. For a deeper review, see [6] and [7].

Modeling other agents. To model feedback interaction between robots and humans, the robot typically relies on knowledge of humans’ objective (cost or reward) function, often chosen to be a linear combination of handcrafted features [8], [9], [10]. A key benefit is interpretability and tractability—knowledge of the objective function provides direct insight into what other agents value and how they would respond in any given situation (assuming they are optimal planners). However, data is required to learn other humans’ objective function; employing techniques such as maximum entropy inverse reinforcement learning [11], [12]. While these approaches boast interpretability benefits, the reliance on knowing other humans’ objective functions hinders real-world multi-agent applicability. A deep neural network can be used to learn the objective function but would lose any interpretability benefits [13]. Online inference is challenging given the limited interaction time while offline learning is prone to distribution shifts. In our work, we do not rely on learning other humans’ objective functions from data. Instead, we propose a robot-only objective function that encourages *legible* and *proactive* behaviors to help the robot indicate to humans early its intent. We show that reducing ambiguity in robot behaviors helps humans commit to a type of behavior [14], thereby mitigating the reliance on accurate feedback interaction models.

Social competency. To increase interpretability and structure into understanding how robots can account for the welfare of others, the concept of Social Value Orientation (SVO) stems from social psychology which models the degree to which an individual values other people’s welfare in relation to their own [15]. Different SVO parameter values elicit different behaviors, such as egoistic, prosocial, altruistic, masochistic, and sadistic. The SVO parameter can be learned from data to provide a social label on interaction data [16], [17], or chosen by the designer to elicit desired robot behaviors [18], [19]. Building upon the idea of socially aware robots, [10], [20] consider designing “courteous” robot behaviors by minimizing the amount of inconvenience the robot imposes on the human. Specifically, they compare how humans behave when the robot is there versus if the robot were not there. Similar to the works discussed previously, [10], [20] too operate in the objective space—a major limitation to encoding social preference in the objective space is that it does not limit how much a robot should compromise its objective for the benefit of others. Without such a limit, it may lead to the robot behaving in unexpected ways, such as moving backward [8], or becoming incapable of accomplishing the task to yield for others, i.e., frozen robot problem [21]. In our work,

we impose a robot-centric *inconvenience budget constraint*, and in conjunction with legible and proactive planning, our approach naturally leads to *prosocial* behaviors compared to other baseline approaches.

Legibility. There has been considerable attention into constructing legible robot behaviors and evaluating its influence in generating more fluent and safer human-robot interactions [14], [22], [23], [24], [25], [26], [27]. However, different approaches formulate legibility differently, leading to vastly different, highly customized legible motion planning algorithms that are difficult to integrate into a more general planning framework that considers additional planning objectives. For example, [23] proposed a gradient-based legible motion optimizer for (single-agent) goal-reaching problems, and [25] introduces the notion of social momentum to generate legible motion to indicate which side agents will pass by each other. Leveraging the notion that the velocity vector is a primary signal for indicating intent [28], we propose a simple approach to synthesizing legible robot behaviors that can be easily integrated into a general optimal control problem.

III. PROBLEM FORMULATION

We aim to compute robot controls that enable a robot to accomplish its task in a safe and efficient manner while taking into account the presence and reactions of surrounding human agents. For simplicity, first consider a single human case; we later describe how to extend this to multiple humans. Let $\mathbf{x}_{\text{HR}}^t \in \mathbb{R}^{n_{\text{H}}+n_{\text{R}}}$ denote the joint state of the human and robot $\mathbf{x}_{\text{H}}^t \in \mathbb{R}^{n_{\text{H}}}$ and $\mathbf{x}_{\text{R}}^t \in \mathbb{R}^{n_{\text{R}}}$ at time step t respectively, and $\mathbf{u}_{\text{H}}^t \in \mathbb{R}^{m_{\text{H}}}$ and $\mathbf{u}_{\text{R}}^t \in \mathbb{R}^{m_{\text{R}}}$ be the human and robot controls at time step t respectively. Then we seek to find a robot control sequence $\mathbf{u}_{\text{R}}^{0:T}$ over T time steps that is a solution to the following optimal control problem:

Problem 1 (Interaction-aware trajectory problem):

$$\min_{\mathbf{u}_{\text{R}}^{0:T}} \sum_{t=0}^T J(\mathbf{x}_{\text{HR}}^t, \mathbf{u}_{\text{R}}^t, t) + J_{T+1}(\mathbf{x}_{\text{HR}}^{T+1}) \quad (1a)$$

$$\text{s.t. } \mathbf{x}_{\text{HR}}^{t+1} = f_{\text{HR}}(\mathbf{x}_{\text{HR}}^t, \mathbf{u}_{\text{R}}^t, \mathbf{u}_{\text{H}}(\mathbf{x}_{\text{HR}}^t)), t = 0, \dots, T \quad (1b)$$

$$g_i(\mathbf{x}_{\text{HR}}^t, \mathbf{u}_{\text{R}}^t) \leq 0, t = 0, \dots, T, i = 0, \dots, G \quad (1c)$$

$$h_j(\mathbf{x}_{\text{HR}}^t, \mathbf{u}_{\text{R}}^t) = 0, t = 0, \dots, T, j = 0, \dots, H. \quad (1d)$$

The objective (cost) function (1a) captures the robot’s performance goals (e.g., minimize control effort and distance to the goal) and joint performance between both agents (e.g., the robot considers the welfare of the human). The joint dynamics constraints (1b) capture the interaction between the two agents; the human’s control policy $\mathbf{u}_{\text{H}}(\mathbf{x}_{\text{HR}}^t)$ is expressed as a function of the joint state—how the robot behaves will affect how the human behaves. Indeed, this joint dynamics can be viewed as an underactuated system—a large challenge is in obtaining realistic human policies and formulating a tractable robot trajectory planner with it. The other two constraints (1c) and (1d) describe constraints on robot controls and the joint state such as initial state constraints, control constraints, and critically, collision avoidance.

IV. LEGIBLE ROBOT PLANNING FOR PROACTIVE AND PROSOCIAL BEHAVIOR

We present our legible robot trajectory planner that encourages the robot to be legible and proactive, leading to prosocial human-robot interactions. The planner will be executing in model predictive control (MPC) fashion, recomputing a new solution at each time step. We leverage the notion that in the context of crowded social navigation, humans engage in *joint collision avoidance* but are also self-interested in that they are unwilling to disproportionately sacrifice their own performance for the benefit of others. Solving Problem 1 is challenging due to the closed-loop coupling between agents and difficulty in accessing the human’s policy. Before diving into the robot planner, we first describe iterated best response (IBR) to account for the closed-loop interactions, a commonly used approach for interaction-aware planning [10], [8], [29]. Although IBR requires a model for the human, we only assume knowledge of human dynamics and goal location. In general, our proposed robot planner is compatible with other interaction planners that do not require an explicit human model (e.g., neural networks) [30], [20], and we defer investigating these models for future work.

A. Iterated best response

We utilize an iterated best response (IBR) strategy to account for the human’s reaction to the robot’s actions. In IBR, the robot iteratively improves its trajectory by computing how the human would respond to its trajectory, and then vice versa. Without loss of generality, we first assume the human is the “leader” whose trajectory is fixed. Then the robot (the “follower”) selects its trajectory after observing the leader’s trajectory. Then the roles are switched and the agents update their trajectories correspondingly. This process repeats until convergence or a fixed number of iterations. The final output is the robot’s trajectory, and the first control output will be executed as part of the MPC loop. The robot assumes that the human is a “reasonable” decision-maker who, like the robot, minimizes a cost function while avoiding collision.

B. Inner trajectory optimization

In this section, we describe the inner trajectory optimization problem during the IBR iterations. Note that while both the human and robot solve the same trajectory optimization problem described in Problem 2, the parameter values may differ for each agent. Without loss of generality, at any iteration of the IBR loop, the follower has observed the leader’s planned trajectory $(\mathbf{x}_L^{0:T+1}, \mathbf{u}_L^{0:T})$ which is considered fixed. The follower responds by selecting a state/control trajectory $(\mathbf{x}_F^{0:T+1}, \mathbf{u}_F^{0:T})$ that is a solution to Problem 2,

Problem 2 (Follower’s trajectory optimization problem):

$$\min_{\substack{\mathbf{x}_F^{0:T+1}, \\ \mathbf{u}_F^{0:T}, \\ \epsilon_t^{0:T+1}}} \sum_{t=0}^T \mu^t J(\mathbf{u}_F^t, \mathbf{x}_F^t, t) + \gamma_0 \sum_{t=0}^{T+1} \gamma^t \epsilon_t^2 + J_{T+1}(\mathbf{x}_F^{T+1}) \quad (2a)$$

$$\text{s.t. } \mathbf{x}_F^{t+1} = f_F(\mathbf{x}_F^t, \mathbf{u}_F^t), \mathbf{x}_F^0 = \mathbf{x}_F^{\text{current}} \quad t = 0, \dots, T \quad (2b)$$

$$\mathbf{x}_F^t \in \mathcal{X}_F^t \setminus \mathcal{O}_{\text{static}}, \quad t = 0, \dots, T+1, \quad (2c)$$

$$\mathbf{u}_F^t \in \mathcal{U}_F(\mathbf{x}_F^t), \quad t = 0, \dots, T \quad (2d)$$

$$g(\mathbf{x}_F^t, \mathbf{u}_F^t, \mathbf{x}_L^t, \mathbf{u}_L^t) \geq -\epsilon_t, \quad t = 0, \dots, T \quad (2e)$$

$$g(\mathbf{x}_F^{T+1}, \mathbf{x}_L^{T+1}) \geq -\epsilon_{T+1}, \quad (2f)$$

$$J_{\text{incon}}(\mathbf{x}_F^{0:T+1}, \mathbf{u}_F^{0:T}) \leq \beta_F, \quad (2g)$$

$$\epsilon_t \geq 0, \quad t = 0, \dots, T+1. \quad (2h)$$

In words, the follower aims to minimize a cost objective (2a) that depends only on its own states and control, subject to dynamics and initial state constraints (2b), state and static obstacle constraints (2c), (state-dependent) control constraints (2d), collision avoidance constraints (2e) and (2f), an inconvenience budget constraint (2g), and slack variable constraints (2h). We describe three key elements of Problem 2 (highlighted in color) that are designed to encourage legible and proactive behaviors.

Markup term: Similar to a discount term common to many infinite horizon planning problems, in the cost function (2a), we instead introduce the inverse of a discount factor: a *markup* term $\mu > 1$ which increases the cost of states and controls further in the future, thereby incentivizing the follower to take nontrivial controls earlier in the (finite) horizon rather than later. For instance, the agent will be incentivized to start turning earlier to avoid the oncoming agent, therefore revealing its intent to pass on one side, rather than at the last possible moment. The hypothesis is that a markup $\mu > 1$ will lead to more *legible* and *proactive* behaviors since the robot’s intention to avoid other agents/obstacles will be revealed earlier in the interaction.

Inconvenience budget: We constrain the amount of “inconvenience” an agent can experience via (2g), therefore ensuring no agent sacrifices too much on its own performance to benefit of others (hence preventing the onset of the frozen robot problem [21]). We described the inconvenience budget in more detail in Section IV-C.

Collision avoidance slack: To ensure Problem 2 remains feasible, we introduce a slack variable on the collision avoidance constraint (2e) with a weighting of γ_0 . Additionally, since the robot’s model of the human within the IBR iterations will not match exactly with how humans truly behave, we apply a discount $\gamma < 1$ on the collision avoidance constraint so that collision constraint violations later in the planning horizon are penalized less. Despite treating collision avoidance as a soft constraint, our experiments show that by the robot behaving legibly and proactively, there is better coordination between the human and robot which leads to more fluent and less collision-prone interactions. Additional safety filters can be added to further improve safety albeit in a reactive manner [31], [32], but this is outside the scope of this work.

C. Inconvenience budget

Similar to [10], inconvenience is a quantity that measures the increase in cost of a planned trajectory compared to an ideal, or optimistic, trajectory. In this work, we consider the inconvenience experienced by the robot whereas [10] considers the human inconvenience caused by the robot’s actions. For notational simplicity, let $\tau^{0:T+1} = (\mathbf{x}^{0:T+1}, \mathbf{u}^{0:T})$ denote a trajectory of an agent (either human or robot). Then let $c(\tau^{0:T+1})$ be a measure of how “convenient” a trajectory $\tau^{0:T+1}$ is for some user-defined notion of convenience (lower

values indicate more convenient trajectories). For example, convenience can be measured as a weighted linear combination of functions that includes trajectory length, jerk, and distance to the goal at the end of the trajectory, i.e., $c(\tau^{0:T+1}) = w^T \phi(\tau^{0:T+1})$. Note that c does not necessarily need to be the same as the cost function J in (2a). Then we define the *inconvenience* of a trajectory $\tau^{0:T+1}$ as,

$$J_{\text{incon}}(\tau^{0:T+1}) = \frac{c(\tau^{0:T+1}) - c(\tau_{\text{ideal}}^{0:T+1})}{c(\tau_{\text{ideal}}^{0:T+1})}, \quad (3)$$

which describes the degradation in convenience of $\tau^{0:T+1}$ compared to $\tau_{\text{ideal}}^{0:T+1}$ which is the solution to the *idealized* trajectory optimization problem described in Problem 3,

Problem 3 (Idealized trajectory optimization problem):

$$\min_{\mathbf{x}^{0:T+1}, \mathbf{u}^{0:T}} \sum_{t=0}^T J(\mathbf{u}^t, \mathbf{x}^t, t) + J_{T+1}(\mathbf{x}^{T+1}) \quad (4a)$$

$$\text{s.t. } \mathbf{x}_F^{t+1} = f_F(\mathbf{x}_F^t, \mathbf{u}_F^t), \quad t = 0, \dots, T \quad (4b)$$

$$\mathbf{x}_F^0 = \mathbf{x}_F^{\text{current}} \quad (4c)$$

$$\mathbf{x}_F^t \in \mathcal{X}_F^t \setminus \mathcal{O}_{\text{static}}, \quad t = 0, \dots, T+1, \quad (4d)$$

$$\mathbf{u}_F^t \in \mathcal{U}_F(\mathbf{x}_F^t), \quad t = 0, \dots, T. \quad (4e)$$

Problem 3 describes optimal trajectory towards a goal assuming no other agents are present (i.e., collision avoidance and inconvenience budget constraints are removed).

Remark: In the scope of Problem 2, $c(\tau_{\text{ideal}}^{0:T+1})$ is a *constant* since it can be computed independently of Problem 2. Thus if ϕ consists of only convex functions (e.g., quadratic functions), then (2g) is a convex (inequality) constraint.

D. Practical considerations when performing IBR

We highlight some important practical considerations.

Sequential Convex Program. Generally, Problem 2 is non-linear (due to dynamics and collision avoidance constraints). We can apply sequential convex programming (SCP), the process of repeatedly convexifying an optimization problem about a previous solution. We can linearize the dynamics and collision avoidance constraint, and if the convenience function (see Section IV-C) and the cost function (2a) are quadratic, then Problem 2 is a quadratically constrained quadratic program, i.e., convex.

IBR initialization. Performing IBR requires an initial trajectory for the leader at the first iteration. Instead of using the previous trajectory, we instead use the ideal trajectory for the current time step. This ensures the IBR solution at each time step will exhibit minimal deviations from the (current) ideal trajectory to the extent allowed by the inconvenience budget constraint. Initializing with the previous trajectory accumulates an imbalance between the human and robot (i.e., one agent would disproportionately move aside more).

Trust region. Since we are using SCP where the linearization is valid locally around the point of linearization, we apply a trust region cost $J_{\text{trust}}(\mathbf{u}_F, \mathbf{x}_F, \mathbf{u}_F^{\text{prev}}, \mathbf{x}_F^{\text{prev}}) = \beta(\|\mathbf{u}_F - \mathbf{u}_F^{\text{prev}}\|_2^2 + \|\mathbf{x}_F - \mathbf{x}_F^{\text{prev}}\|_2^2)$ to (2a) to ensure the new solution does not differ too much from the previous solution. As such, combined with the IBR iterations, *both* agents' trajectories will gradually deviate from their ideal trajectories until they avoid collisions or reach the inconvenience budget constraint.

E. Scaling up to multiple agents and wall constraints

Thus far, we have been concerned with a two-agent setting—one robot, one human. Extending IBR problems to account for all agents in the scene is possible but can be computationally challenging and typically relies on parallelization [29] to ensure real-time capabilities. To reduce the number of agents to consider within the IBR loop, we lean on the idea that different agents have varying degrees of interactivity with the robot. For example, agents that are close to the robot but moving *away* can be considered “peripheral” agents, whereas agents that are farther but moving towards the robot are considered “interacting” agents. We consider the problem of defining an interactivity score or classifying which agents are interacting or “peripheral” outside the scope of this work but there are several works that investigate this problem [33], [34], [35]. Assuming interacting and peripheral agents have been classified, we propose a simple tractable way to account for both interacting and peripheral agents simultaneously—we treat peripheral agents as constant velocity obstacles which amount to additional *linear constraints* (after linearization) to Problem 2. Additionally, straight wall constraints (e.g., narrow corridor passing) are linear state constraints that can be easily appended to Problem 2.

V. EXPERIMENTAL RESULTS AND DISCUSSION

A. Experimental set-up

We evaluate our proposed approach on a number of simulations and compare it against a number of baseline methods. We also perform human-in-the-loop experiments to demonstrate the real-time applicability of our method.

Software implementation. We used the Julia programming language [36], and the JuMP package, a modeling language for a wide range of optimization problems [37], [38], to model the problem, and ECOS [39] as the solver. We used forward-mode automatic differentiation [40] to linearize the problem. For human-in-the-loop experiments, we used ROS [41] to incorporate the human control input in real time. All experiments (except human-in-the-loop) were performed on an AMD Ryzen 9 3900X 12-Core Processor.

Dynamics. Our method is dynamics-agnostic. For simplicity and ease of comparison, we use the dynamically-extended unicycle model for both the human and robot agents. The state and control (for each agent) are $\mathbf{x} = [x, y, \theta, v]^T$ and $\mathbf{u} = [\omega, a]^T$, and the continuous time dynamics are $\dot{\mathbf{x}} = [v \cos \theta, v \sin \theta, \omega, a]^T$. We use the discrete-time dynamics (omitted for brevity) by applying zero-order hold on the continuous dynamics with a time step size of $\Delta t = 0.1$ seconds. The control and velocity limits are $\omega \in [-1, 1]\text{rad}$ and $a \in [-1.5, 1.5]\text{ms}^{-2}$, and $v \in [0, 1.5]\text{ms}^{-1}$.

Comparison methods. We compare our approach to several baseline methods. *Social Forces Model (SFM)* [3]: A common multi-agent control method based on attractive and repulsive forces for goal reaching and obstacle avoidance. *Reactive Control via HJ reachability (HJ)*: A reactive control method that executes optimal collision avoidance control if a collision with the human is imminent under a constant




Fig. 2: Human and robot trajectories in a head-on scenario. In our approach (green box) the robot legibly and proactively conveys its intent to pass, thus providing the human sufficient time to prepare to pass safely. In other approaches (blue boxes), the robot (SFM) does not convey its intention clearly and early and causes the human to swerve significantly, (HJ) freezes on the spot, (OC) takes on more collision avoidance responsibility than necessary, or (vIBR) confuses the human.




Fig. 3: Our approach can easily account for additional wall constraints and multiple dynamic agents.

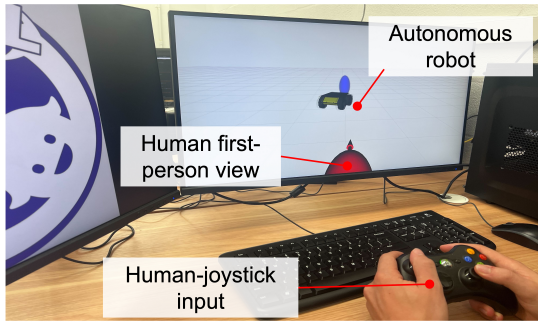


Fig. 4: Snapshot of human-in-the-loop simulation experiment. Human interacts with the robot in first-person view.

velocity assumption [4]. Hamilton-Jacobi reachability [42], [43] is used for collision checking and control. *Optimal Control (OC)*: Problem 2 without IBR iterations, without markup ($\mu = 1$), without the inconvenient budget constraint, and a collision slack weighting γ_0 of 1000. *Vanilla IBR (vIBR)*: The OC method with three IBR iterations. *Legible & Proactive IBR (Ours)*: Problem 2 with 3 IBR iterations.

Human simulation model. To simulate the human in our experiment, we use two different approaches, IBR and OC (as described above). We assume the human’s goal is known, which is privileged information. We also add some white noise to the human controls. The parameters used to model the human is different from the robot’s model of the human to make sure the robot does not use privileged information.

Experimental parameters. We considered a planning horizon of 2.5 seconds, and each episode was 5 seconds. When comparing against baselines, we used an inconvenience budget of 0.2, markup $\mu = 1.05$, collision discount $\gamma = 0.98$, a collision slack weighting $\gamma_0 = 150$, and a collision radius of 1 meter. We use a quadratic cost function, penalizing only running control cost and distance to the goal position. For the inconvenience budget, we considered a linear combination


of the sum of distances between states along the trajectory squared, the sum of the change in velocity squared, and the distance to the goal squared. Note that these quantities are quadratic. The human and robot approach each other almost head-on. Both agents have their goal position situated 10 meters directly in front of them, but their relative heading varies between $\pm \frac{\pi}{4}$. Their straight-line trajectory between the initial and goal state intersects, meaning they must deviate from the straight-line path to avoid collision.

B. Qualitative analysis

General behavior. Figure 2 illustrates a canonical scenario where the robot and human approach each other head-on—if the agents do nothing, collision is inevitable. We see that in our method (right, green box), both agents engage cooperatively in joint collision avoidance. The reason is that the robot has indicated early and clearly its intent to pass on the right. Whereas in SFM, the robot does not move out of the way until the last moment and does so very slowly, causing the human to make a sharp swerve at the last moment. In the HJ approach, the constant velocity assumption is too conservative and causes the robot to almost immediately freeze on the spot, resulting in the human doing most of the collision avoidance. In the OC case, does not anticipate the human will cooperate and makes a sharp turn to avoid a collision. In the vIBR case, while both agents deviated to avoid collision, the robot’s motion was not legible or proactive which caused the human to sway back and forth a bit before passing on the left.

Additional constraints. We tested our approach with wall constraints and with multiple agents. Figure 3 (left) compares the trajectories with and without the presence of a wall, while the figure on the right illustrates our planner navigating through a crowd. The human is forced to swerve more than the robot since the robot is blocked by another (non-interacting) agent. When adding 4 additional non-interacting agents, the computation time increased by $\sim 6\%$, from approximately 67ms to 71ms per planning time step.

Effect of markup. Figure 5a illustrates how the robot is proactive and legible through its heading and heading rate. In the left plot, as markup increases, the robot rotates faster, thus making the robot’s decision to move to one side more noticeable earlier. In the right plot, with our proposed method, the robot’s heading is smoother, more deliberate, and less aggressive compared to HJ. Whereas SFM barely



(a) Left: Varying the markup parameter affects how quickly the robot starts turning to pass an oncoming human. Right: With our method, the robot’s heading is consistent, smooth, and deliberate, whereas other methods oscillate between passing left or right or swerving too much.

(b) Box plots comparing the minimum distance between agents. Our approach does exceed the collision radius threshold despite using less stringent safety measures compared to other approaches.

Fig. 5: (a) Analysis of heading and heading rate to illustrate the legibility and proactivity of our proposed approach. (b) Statistics of the minimum distance between human and robot to evaluate the safety performance of our approach.

turns until the end, and OC and vIBR are indecisive and only start turning later in the interaction.

Human-in-the-loop simulation. We performed human-in-the-loop experiments, see Figure 4, to validate the performance of our planner and demonstrate its real-time applicability. The planning was performed at 10Hz on a laptop with AMD Ryzen 7 PRO 6850U processor.

C. Performance metrics

We consider the following metrics to evaluate the safety and efficiency of our approach.

Minimum distance (MinDist): To evaluate safety, we consider the minimum distance between the robot and human during the interaction. Let $\text{pos}(\cdot)$ denote the position vector given agent state. Then the minimum distance is defined as,

$$\text{MinDist}(\mathbf{x}_H^{0:T+1}, \mathbf{x}_R^{0:T+1}) = \min_{t \in \{0, \dots, T+1\}} \|\text{pos}(\mathbf{x}_H^t) - \text{pos}(\mathbf{x}_R^t)\|_2$$

Path Irregularity (PI) Index [44]: The total amount of unnecessary turning relative to the straight line path towards the goal. Lower is better. Let $\mathbf{v}_t = [v_t \cos \theta_t, v_t \sin \theta_t]^T$ be the velocity vector at time t , and $\mathbf{d}_t^{\text{str}} = [x_{\text{goal}} - x_t, y_{\text{goal}} - y_t]^T$ be the vector pointing to the goal from the position at time t . Then the path irregularity index is defined as,

$$\text{PI}(\tau^{0:T+1}) = \sum_{t=0}^{T+1} \cos^{-1} \left(\frac{\mathbf{v}_t \cdot \mathbf{d}_t^{\text{str}}}{\|\mathbf{v}_t\| \|\mathbf{d}_t^{\text{str}}\|} \right)$$

Distance to goal (D2G): To evaluate how well an agent accomplishes its task, we compare how close the agent gets to the desired goal location by the end of the horizon. Lower is better. Note: All simulations have the same horizon length.

$$\text{D2G}(\tau^{0:T+1}) = \|\text{pos}(\mathbf{x}^{T+1}) - \text{pos}(\mathbf{x}^{\text{goal}})\|_2$$

Total acceleration (ACC): To evaluate the amount of effort and smoothness of the trajectory, we compute the total acceleration over the trajectory. Lower is better.

$$\text{ACC}(\tau^{0:T+1}) = \frac{1}{\Delta t} \sum_{t=0}^T \|\mathbf{v}_{t+1} - \mathbf{v}_t\|_2$$

D. Quantitative analysis

We compare our approach against the baseline methods using the performance metrics described in Section V-C. **Safety.** Figure 5b compares the minimum distance between the robot and human over the interaction episode. We see that

our method does not violate the collision radius threshold despite collision avoidance not being a hard constraint, and having a lower slack penalty than OC. In the OC case where the weighting on the slack variable was much higher than Ours, there were a few cases where the minimum distance dipped below the collision radius. This result is consistent with our hypothesis—if the robot is more proactive and legible, then the human can prepare, and both can engage equitably in avoiding collision.

Efficiency. Figure 6 compares the efficiency-related performance metrics. Both SFM and HJ methods result in a large imbalance in all efficiency metrics, often requiring the human to bear a majority of the collision avoidance responsibility, i.e., swerving a lot. This is because SFM moves aside too late which is ineffective in conveying intent to the human, while HJ is too timid and often freezes on the spot requiring the human to swerve out of the way. OC, in general, performs similarly to vIBR and Ours, but experiences large fluctuations since it does not account for the interaction dynamics—it does not anticipate how the human may respond, and therefore often gets in the way of the human, requiring one or both to swerve out of the way (see Figure 2). vIBR performs similarly to Ours, but Ours either has a lower variance, is more equitable, or for roughly the same human performance, the robot does better (especially notable in ACC). These results indicate that with relatively simple modifications to the standard trajectory optimization problem, we are able to achieve more prosocial (i.e., equitable) and more efficient robot behaviors without compromising on safety.

VI. CONCLUSIONS AND FUTURE WORK

We have presented simple yet effective modifications to a general trajectory optimization problem to generate legible and proactive robot behaviors for prosocial human-robot interactions. With our framework, we are able to achieve more efficient and more equitable interactions without compromising on safety. Our human-in-the-loop experiments highlight the real-time applicability of our approach. We also propose a simple and scalable approach to account for multiple agents but requires future work to distinguish which humans are considered interacting or peripheral. Future work




Fig. 6: Box plots comparing efficiency performance metrics using two different human models (interactive and non-interactive). Our approach has a lower and more equitably distributed between the human and robot. (Lower is better.)

is aimed at scaling our efforts to account for more humans and perform real-world human-in-the-loop testing. To ensure scalability, we plan to explore different interaction-aware planning frameworks, such as using neural-based human behavior prediction models to model interaction.

REFERENCES

- [1] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: A survey," *Int. Journal of Robotics Research*, vol. 39, no. 8, pp. 895–935, 2020.
- [2] M. Mayer, R. Bell, and A. Buchner, "Self-protective and self-sacrificing preferences of pedestrians and passengers in moral dilemmas involving autonomous vehicles," *PLoS ONE*, 2021.
- [3] D. Helbing and P. Molnár, "Social force model for pedestrian dynamics," *Physical Review E*, vol. 51, no. 5, pp. 4282–4286, 1995.
- [4] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *Proc. IEEE Conf. on Robotics and Automation*, 2008.
- [5] R. A. Knepper and D. Rus, "Pedestrian-inspired sampling-based multi-robot collision avoidance," in *Proc. IEEE Int. Conf. on Robot and Human Interactive Communication*, 2012.
- [6] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.
- [7] A. Francis, C. Perez-D'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, H.-T. L. Chiang, M. Everett, S. Ha, J. Hart, J. P. How, H. Karnan, T.-W. E. Lee, L. J. Manso, R. Mirsky, S. Pirk, P. T. Singamaneni, P. Stone, A. V. Taylor, P. Trautman, N. Tsoi, M. Vazquez, X. Xiao, P. Xu, N. Yokoyama, A. Toshev, and R. Martín-Martín, "Principles and guidelines for evaluating social robot navigation algorithms," Available at <https://arxiv.org/abs/2306.16740>, 2023.
- [8] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Robotics: Science and Systems*, 2016.
- [9] H. Kretschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *Int. Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [10] L. Sun, W. Zhan, M. Tomizuka, and A. Dragan, "Courteous autonomous cars," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2018.
- [11] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proc. AAAI Conf. on Artificial Intelligence*, 2008.
- [12] S. Levine and V. Koltun, "Continuous inverse optimal control with locally optimal examples," in *Int. Conf. on Machine Learning*, 2012.
- [13] C. Finn, S. Levine, and P. Abbeel, "Guided cost learning: Deep inverse optimal control via policy optimization," in *Int. Conf. on Machine Learning*, 2016.
- [14] D. Carton, W. Olszowy, and D. Wollherr, "Measuring the effectiveness of readability for mobile robot locomotion," *Int. Journal of Social Robotics*, vol. 8, pp. 721–741, 2016.
- [15] W. B. G. Liebrand and C. G. McClintock, "The ring measure of social values: A computerized procedure for assessing individual differences in information processing and social value orientation," *European Journal of Personality*, vol. 2, no. 3, pp. 217–230, 1988.
- [16] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles," *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24 972–24 978, 2019.
- [17] L. Crosato, H. P. H. Shum, E. S. L. Ho, and C. Wei, "Interaction-aware decision-making for automated vehicles using social value orientation," *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 2, pp. 1339–1349, 2023.
- [18] B. Toghi, R. Valiente, D. Sadigh, R. Pedarsani, and Y. P. Fallah, "Co-operative autonomous vehicles that sympathize with human drivers," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2021.
- [19] B. Toghi, R. Valiente, D. Sadigh, R. Pedarsani, and Y. Fallah, "Social coordination and altruism in autonomous driving," *IEEE Transactions on Intelligent Vehicles*, vol. 23, no. 12, pp. 24 791–24 804, 2022.
- [20] S. Schaefer, K. Leung, B. Ivanovic, and M. Pavone, "Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions," in *Proc. IEEE Conf. on Robotics and Automation*, 2021.
- [21] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, 2010.
- [22] A. D. Dragan, K. C. L. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *Proc. ACM/IEEE Int. Conference on Human Robot Interaction*, 2013.
- [23] A. D. Dragan and S. S. Srinivasa, "Generating legible motion," in *Robotics: Science and Systems*, 2013.
- [24] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch, "Legible robot navigation in the proximity of moving humans," in *Proc. IEEE Workshop on Advanced Robotics and its Social Impacts*, 2012.
- [25] C. I. Mavrogiannis, P. Alves-Olivera, W. Thomason, and R. A. Knepper, "Social momentum: Design and evaluation of a framework for socially competent robot navigation," *ACM Transactions on Human-Robot Interaction*, vol. 37, no. 4, 2021.
- [26] C. Lichtenh  ler, T. Lorenzy, and A. Kirsch, "Influence of legibility on perceived safety in a virtual human-robot path crossing task," in *Proc. IEEE Int. Conf. on Robot and Human Interactive Communication*, 2012.
- [27] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human-robot interactions," *Int. Journal of Social Robotics*, vol. 9, pp. 765–779, 2017.
- [28] C. Sch  ller, A. V., L. F., and A. Knoll, "What the constant velocity model can teach us about pedestrian motion prediction," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1696–1703, 2020.
- [29] M. Wang, Z. Wang, J. Talbot, J. C. Gerdes, and M. Schwager, "Game-theoretic planning for self-driving cars in multivehicle competitive scenarios," *IEEE Transactions on Robotics*, vol. 37, no. 4, pp. 1313–1325, 2021.
- [30] E. Schmerling, K. Leung, W. Vollprecht, and M. Pavone, "Multimodal probabilistic model-based planning for human-robot interaction," in *Proc. IEEE Conf. on Robotics and Automation*, 2018.
- [31] K. Leung, E. Schmerling, M. Zhang, M. Chen, J. Talbot, J. C. Gerdes, and M. Pavone, "On infusing reachability-based safety assurance within planning frameworks for human-robot vehicle interactions," *Int. Journal of Robotics Research*, vol. 39, pp. 1326–1345, 2020.
- [32] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2018.
- [33] K. Kitazawa and T. Fujiyama, "Pedestrian vision and collision avoidance behavior: Investigation of the information process space of pedestrians using an eye tracker," *Pedestrian and Evacuation Dynamics*, pp. 95–108, 2010.

- [34] S. Topan, K. Leung, Y. Chen, P. Tupekar, E. Schmerling, J. Nilsson, M. Cox, and M. Pavone, "Interaction-dynamics-aware perception zones for obstacle detection safety evaluation," in *IEEE Intelligent Vehicles Symposium*, 2022.
- [35] E. Tolstaya, R. Mahjourian, C. Downey, B. Varadarajan, B. Sapp, and D. Anguelov, "Identifying driver interactions via conditional behavior prediction," in *Proc. IEEE Conf. on Robotics and Automation*, 2021.
- [36] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM Review*, vol. 59, no. 1, pp. 65–98, 2017.
- [37] I. Dunning, J. Huchette, and M. Lubin, "JuMP: A Modeling Language for Mathematical Optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
- [38] M. Lubin, O. Dowson, J. Dias Garcia, J. Huchette, B. Legat, and J. P. Vielma, "JuMP 1.0: Recent improvements to a modeling language for mathematical optimization," *Mathematical Programming Computation*, 2023.
- [39] A. Domahidi, E. Chu, and S. Boyd, "ECOS: An SOCP solver for embedded systems," in *European Control Conference*, 2013.
- [40] J. Revels, M. Lubin, and T. Papamarkou, "Forward-mode automatic differentiation in julia," Available at <https://arxiv.org/abs/1607.07892>, 2016.
- [41] Stanford Artificial Intelligence Laboratory et al., "Robotic operating system," Available at <https://www.ros.org>, 2018.
- [42] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.
- [43] E. Schmerling, "HJ Reachability in JAX," Available at https://github.com/StanfordASL/hj_reachability.
- [44] J. Guzzi, A. Giusti, L. M. Gambardella, G. Theraulaz, and G. A. Di Caro, "Human-friendly robot navigation in dynamic environments," in *Proc. IEEE Conf. on Robotics and Automation*, 2013.