

A Robust Control Framework for Human Motion Prediction

Andrea Bajcsy¹, Somil Bansal, Ellis Ratner, Claire J. Tomlin, Anca D. Dragan

Abstract—Designing human motion predictors which preserve safety while maintaining robot efficiency is an increasingly important challenge for robots operating in close physical proximity to people. One approach is to use robust control predictors that safeguard against every possible future human state, leading to safe but often too conservative robot plans. Alternatively, intent-driven predictors explicitly model how humans make decisions given their intent, leading to efficient robot plans. However, when the intent model is misspecified, the robot might confidently plan unsafe maneuvers. In this work we combine ideas from robust control and intent-driven human modelling to formulate a novel human motion predictor which provides robustness against misspecified human models, but reduces the conservatism of traditional worst-case predictors. Our approach predicts the human states by trusting the intent-driven model to decide only which human actions are completely unlikely. We then safeguard against all likely enough actions, much like a robust control predictor. We demonstrate in simulation and hardware how our approach safeguards against misspecified human intent models while not leading to overly conservative robot plans.

Index Terms—Safety in HRI, Human-Aware Motion Planning

I. INTRODUCTION

Robots such as autonomous vehicles and assistive manipulators are increasingly operating in dynamic environments and close physical proximity to people. In such scenarios, it is important that robots not only account for the current state of the humans nearby, but also predict their future state to plan safe and efficient trajectories.

To maximally preserve safety, a robust optimal control perspective models the human as taking any action (with equal likelihood) from a set of controls. The predictor combines this control set with a conservative human dynamics model to compute a *full forward reachable set*, or the set of all states that the human could reach from their current state [2, 3]. This approach allows the robot to produce safe predictions when very little is understood about human decision-making.

Manuscript received: June, 1, 2020; Revised July, 24, 2020; Accepted September, 21, 2020.

This paper was recommended for publication by Tamim Asfour upon evaluation of the Associate Editor and Reviewers' comments. This work is an extension of [1]. In this paper, we provide further analysis of the proposed method as it applies to human motion prediction and how it compares to stochastic prediction approaches, introduce several additional quantitative examples to emphasize the variety of scenarios in which our method is useful, and present further simulated experiments with human models whose latent parameters vary over time.

¹ Authors are with EECS at UC Berkeley: {abajcsy, somil, eratner, tomlin, anca}@berkeley.edu.

Digital Object Identifier (DOI): see top of this page.




Fig. 1: When intent-driven human models are misspecified in Bayesian predictors, robots confidently plan unsafe motions (top). Our approach (bottom) trusts the intent-driven model only to remove completely *unlikely* human actions, resulting in safer robot plans despite a misspecified model. (not depicted here) When planning using the worst-case predictor, the robot has to leave the environment entirely to avoid the predicted human state.

A complementary perspective is that there is structure to human decision-making: humans have intentions, and make decisions in pursuit of these intentions. For example, consider an indoor home environment where people often move towards chairs, tables, or doorways. Predictors synthesized from this perspective, called *intent-driven predictors*, build data-driven models of human actions given intent [4–8], and have been successful in a variety of domains including manipulation [9–11], autonomous driving [12], and navigation [13, 14] (see [15] for a survey). Since human behavior varies between people and over time, these decision-making models are often parameterized and predictors maintain a belief distribution over the model parameters [16, 17]. This provides a direct and succinct way for the robot to use online data to update its human model [4, 6, 18, 19].

However, a key challenge remains with such intent-driven predictors. To update the belief over model parameters and to generate predictions, intent-driven predictors rely on priors and on likelihood models which describe the probability of observing a data point as a function of the model parameters. Although these two components enable data and prior knowledge to improve the human model online, likelihood models are difficult to specify and the priors may be incorrect.

In this work we seek an approach which bridges robust control and intent-driven predictors: a predictor which is more robust to misspecified models and priors, but still able to leverage human data online to safely reduce conservatism. Our key idea is to compute a *restricted forward reachable set* by *trusting the intent-driven model to tell us only what is completely unlikely*. However, unlike intent-driven predictors,

we will not rely on the exact probability of each action under our model during prediction. Rather, we use the decision-making model and the belief to divide the set of human actions in two disjoint sets of likely and unlikely actions. We then predict human motion by treating all sufficiently likely actions as equally probable, much like in the full forward reachable set. Using this restricted control set results in a prediction problem which can be readily formalized and solved through existing robust control methods and tools [20, 21]. We utilize Hamilton-Jacobi (HJ) reachability analysis [2, 22] which is a method for guaranteeing safety for continuous-time, nonlinear dynamical systems. Finally, to properly restrict the set of human controls based on the intent-driven model and belief over model parameters, we augment the state space with the belief. Since the belief encodes the likelihood of human actions given the history of human actions, this explicit belief tracking allows us to compute the likely actions at any future state.

To summarize, our key contributions are:

- a robust control framework for human motion prediction which provides robustness against misspecified models and model parameter priors;
- a comparison of our approach to forward reachable set and stochastic predictors for static and time-varying human intent models and in three pedestrian scenarios where the belief over the human intent changes online;
- a demonstration of our prediction approach in hardware.

II. PROBLEM SETUP

We consider a robot operating in a shared workspace with a human. The robot needs to predict the human's motion¹ and plan a collision-free path around the human to reach the goal as efficiently and safely as possible. To describe the motion of the robot and the human, we model both as dynamical systems. Let the state of the human and the robot be $x_H \in \mathbb{R}^{n_H}$ and $x_R \in \mathbb{R}^{n_R}$ respectively. The time-evolution of these states can be described by $\dot{x}_H = f_H(x_H, u_H)$, $\dot{x}_R = f_R(x_R, u_R)$ where the human and robot's controls are $u_H \in \mathbb{R}^{m_H}$ and $u_R \in \mathbb{R}^{m_R}$ respectively.

The robot's goal is to plan a control trajectory $u_R(t), t \in [0, T]$ such that it does not collide with the human or any (known) static obstacles, and reaches its goal g_R by T . In this work, we will solve this planning problem in a receding horizon fashion. However, the future states of the human are not known *a priori*, and thus the robot must predict future human motion in order to plan collision-free trajectories.

Throughout this paper, we will focus on contrasting the intent-driven and full forward reachable set predictor with our novel predictor. However, all prediction schemes ultimately produce a set of sufficiently likely states (forward in time until the prediction horizon, N) that the robot uses for collision checking. We define the set of likely human states at some future time, $t + \tau$ as: $\mathcal{K}^t(\tau)$, $\forall \tau \in [0, N]$.

Running example: We now introduce a running example for illustration purposes throughout the paper. Consider a

mobile robot that needs to navigate to a goal position $g_R \in \mathbb{R}^2$ in a room where a person is walking. Since the human is a pedestrian in this scenario, we use a planar human model with state $x_H = [h_x, h_y]$ and dynamics $\dot{x}_H = [v_H \cos(u_H), v_H \sin(u_H)]$. We model the human as moving at a fixed speed $v_H \approx 0.6\text{m/s}$ and controlling their heading angle $u_H \in [-\pi, \pi]$. Our mobile robot is modeled as a 3D system with state given by its position and heading $x_R = [s_x, s_y, \phi]$, and speed and angular speed as the control $u_R = [v_R, \omega]$, and dynamics $\dot{x}_R = [v_R \cos \phi, v_R \sin \phi, \omega]$. The robot control inputs are constrained between $[0, 0.6]\text{m/s}$ and $[-1.1, 1.1]\text{rad/s}$ respectively.

III. BACKGROUND: ROBUST VS. INTENT PREDICTION

In this work, we aim to unify ideas from the robust control with the intent-driven prediction so we start with a brief background on both. In each section, we refer interested readers to more comprehensive resources on each approach.

A. Robust Control Prediction

The most conservative prediction of human motion is the set of *all* states the human could reach in a time horizon. Let t be the current real time and $\tau \in [0, N]$ be a future time used by the predictor. Also, let $\xi(x_H^0, \tau, u_H(\cdot)) := x_H^\tau$ denote the human state starting from the current state $x_H^0 := x_H^t$ at time 0 and applying control u_H for a duration of τ . The *full Forward Reachable Set (FRS)* is defined as:

$$\mathcal{K}_{FRS}^t(\tau) := \{x_H^\tau : \exists u_H(\cdot), x_H^\tau = \xi(x_H^0, \tau, u_H(\cdot))\} \quad (1)$$

In other words, if the human is in state x_H^0 , then they are predicted to reach any state x_H^τ in τ time if that state is reachable through some control signal $u_H(\cdot)$.

In general there are many techniques for computing these sets [3, 20, 21], but in this work we use Hamilton-Jacobi (HJ) reachability analysis [2, 22]. In HJ reachability, the computation of the FRS is formulated as a dynamic programming problem which ultimately requires solving for the value function $V(\tau, x_H)$ in the following initial value Hamilton Jacobi-Bellman PDE (HJB-PDE):

$$\frac{\partial V(\tau, x_H)}{\partial \tau} + \max_{u_H \in \mathcal{U}} \nabla_{x_H} V(\tau, x_H) \cdot f(x_H, u_H) = 0 \quad (2)$$

$$V(0, x_H) = l(x_H),$$

where $\tau \geq 0$. The function $l(x_H)$ is the implicit surface function representing the initial set of states that the human occupies $\mathcal{L} = \{x_H : l(x_H) \leq 0\}$. Note that this equation is the continuous-time analogue of the discrete-time Bellman equation. The maximization over the human's control, $u_H \in \mathcal{U}$, encodes the effect of the human dynamics and control on the value, which lies in the set of all possible controls. Note that since this optimization considers all controls, the predictions will include all possible states the human could reach, thereby resulting in the safe but oftentimes overly conservative predictions. Once the value function $V(\tau, x_H)$ is computed, the FRS predictions are given by the sub-zero level set $\mathcal{K}_{FRS}^t(\tau) = \{x_H : V(\tau, x_H) \leq 0\}$. For more details on the HJB-PDE, please refer to [2].

¹We assume that the robot and human states can be accurately sensed.

B. Intent-driven Bayesian Prediction

Unlike the robust predictor, the intent-driven Bayesian predictor couples a structured model of how the human chooses their actions with the dynamics model. In general, constructing a human decision-making model for a robotic application is a particularly difficult modeling challenge and many approaches exist in the literature (see [15]). In this work, we consider stochastic control policies that are parameterized by a discrete random variable λ^t where Λ is the set of all values that λ^t can take. The human's control policy can be described by the probability density function $u_H^t \sim p(u_H^t | x_H^t; \lambda^t)$. Here, λ^t can represent many different aspects of human decision-making, including what goal locations they are moving towards [6] or even the kind of visual cues they pay attention to in a scene [7]. We refer to these aspects of human decision-making as the human's *intent*. Furthermore, we use the superscript t on the parameter to denote that the value of the human parameter can be time-varying. This allows the human model to encode how the human's intent changes over time; for example, if a person changes the goal they are moving towards in a room.

In general, the specific choice of parameterization is often highly problem specific and can be hand-designed or learned from prior data [6, 23]. Regardless of the specific parameterization, in practice, the true value of λ^t is frequently unknown beforehand and instead can be estimated from the measurements of the true human behavior. Thus, at any time t , the robot additionally maintains a belief distribution $b^t(\lambda^t)$ over the model parameters, which allows it to estimate the human's intent online via a Bayesian update:

$$b_+^t(\lambda^t | u_H^t, x_H^t) = \frac{P(u_H^t | x_H^t; \lambda^t) b^t(\lambda^t)}{\sum_{\bar{\lambda} \in \Lambda} P(u_H^t | x_H^t; \bar{\lambda}) b^t(\bar{\lambda})} \quad (3)$$

Running example: *The robot has uncertainty about the human's goal location. Let the human parameter $\lambda^t \in \Lambda = \{g_1, g_2\}$ take two values which indicates which goal location the human moving towards. The human decision-making model at any state and for a particular goal is given by a Gaussian distribution over the heading angle with mean pointing in the goal direction and a variance representing uncertainty in the human action:*

$$p(u_H^t | x_H^t; \lambda^t) = \begin{cases} \mathcal{N}(\mu_1, \sigma_1^2), & \text{if } \lambda^t = g_1 \\ \mathcal{N}(\mu_2, \sigma_2^2), & \text{if } \lambda^t = g_2 \end{cases},$$

where $\mu_i = \tan^{-1} \left(\frac{g_i(y) - h_y^t}{g_i(x) - h_x^t} \right)$ and $\sigma_i = \pi/4$ for $i \in \{1, 2\}$. Here, $(g_i(x), g_i(y))$ represents the position of goal g_i .

At prediction time, the stochastic nature of the human decision-making model and the belief over the parameters is naturally converted into state *distributions* (instead of deterministic sets) forward in time. Note that typically, these predictors use a temporally and spatially discretized form of the dynamics by integrating f_H over a fixed time interval δt . Controls are often discretized too and assumed to be held fixed during δt . This results in the predictor maintaining and updating discrete distributions over the human state space. Given the current real time t , we will denote a future time discrete timestep by $k \in \{0, 1, \dots, \frac{H}{\delta t}\}$.

Suppose the current state of the human at the start of the prediction horizon is $x_H^0 := x_H^t$ and the current belief is $b^0(\lambda^0) := b^t(\lambda^t)$. Assume the human is at x_H^k at some future time k . Combining the dynamics and human policy, the human's state distribution at the next timestep $k+1$ is

$$P(x_H^{k+1} | x_H^k; \lambda^k) = \sum_{u_H^k} P(x_H^{k+1} | x_H^k, u_H^k) P(u_H^k | x_H^k; \lambda^k).$$

This equation can be applied recursively to compute $P(x_H^{k+1} | x_H^0; \lambda^{0:k})$ starting from $k=0$. Marginalizing over all sequences of values that the human parameter λ could take, \mathcal{S}_k , where $|\mathcal{S}_k| = |\Lambda|^k$, we get the overall distribution over the human state at future time step $k+1$: $P(x_H^{k+1} | x_H^0)$. Here, the probability of the parameter sequence has to be set in the model and is generally defined by ² $P(\lambda^{0:k} | x_H^0) = \left(\prod_{m=1}^k P(\lambda^m | \lambda^{m-1}) \right) b^0(\lambda^0)$.

Importantly, at planning time, the robot must decide which predicted states are sufficiently likely to warrant avoiding. A strict notion of safety requires the robot to avoid all states whose probability is > 0 . While safe (and equivalent to the full FRS), this choice of states does not leverage the data encoded through the belief or the human decision-making model. To reduce the volume of this set in a way commensurate with human decision-making, choosing a nonzero probability threshold is desirable and reveals a significantly smaller set of states that aligns with the model. Thus, the ultimate predicted set of human states that the robot must avoid at planning time is:

$$\mathcal{K}_\epsilon^t(k) = \{x_H^k : P(x_H^k | x_H^0) > \epsilon\}, \forall k \in \{0, \dots, \frac{N}{\delta t}\} \quad (4)$$

where $\epsilon \geq 0$ is a safety threshold and a design parameter.

IV. A ROBUST-CONTROL FRAMEWORK FOR INTENT-DRIVEN HUMAN PREDICTION

Our key idea in this paper is to compute a *restricted* forward reachable set by trusting the intent-driven model to infer only what is *completely unlikely*. After using the intent-driven model to prune away sufficiently unlikely actions, our robust predictor will safeguard against all sufficiently likely actions equally, much like in the full forward reachable set. The main question becomes how to perform this control-set pruning in a principled way over the prediction horizon.

One simple way of choosing this set is as follows. At the beginning of the prediction horizon, let the human state be $x_H^0 := x_H^t$ and the current belief be $b^0 := b^t$. We can form a new distribution over the human's controls at the first time step by marginalizing out the latent model parameters, given the initial belief we have over those parameters: $p(u_H | x_H) = \sum_{\lambda \in \Lambda} p(u_H | x_H; \lambda) b^0(\lambda)$. Then, we can choose the set of human actions to be those for which this marginalized initial likelihood is above a threshold: $\mathcal{U}(x_H) = \{u_H : p(u_H | x_H) \geq \delta\}$. This leads to a set of

²In the case of static latent parameters, the summation simplifies to $P(x_H^{k+1} | x_H^0) = \sum_{\lambda \in \Lambda} P(x_H^{k+1} | x_H^0; \lambda) b^0(\lambda)$.

reachable states for $t = 1$. To obtain the set of states at $t = 2$, it is tempting to follow the same process, restricting the set of future actions based on b^0 . Unfortunately, this would (accidentally) model that the human is “resampling” their intent from this initial distribution independently at every step. It would disregard that a human’s second action will be *consistent* with their first, with the intent only changing according to the dynamics of λ . Thus, we must enforce that the human control from a state x_H at $t = 2$ is not only consistent with the initial belief, but also with the control that took them to state x_H .

To properly restrict the set of feasible controls over the prediction horizon, we need to take into account how the likelihood of any future control depends on the past sequence of human controls. The belief precisely encodes this likelihood given the past sequence of human controls through the Bayesian update from Eq. 3. Thus, our predictor explicitly tracks the updated belief as it makes predictions, rather than just the updated state, and restricts future actions based on future beliefs (see left of Fig. 3 for intuitive depiction). Let this joint state space be denoted by $z^t := [x_H^t \ b^t(\lambda^t = \lambda_1) \ \dots \ b^t(\lambda^t = \lambda_{|\Lambda|})]$.

When predicting using this state space, to simultaneously predict the possible future beliefs over λ^t and corresponding likely human states, we consider the joint dynamics:

$$\dot{z}^t = [\dot{x}_H^t \ \dot{b}^t(\lambda^t = \lambda_1) \ \dots \ \dot{b}^t(\lambda^t = \lambda_{|\Lambda|})], \quad (5)$$

where $\dot{z}^t := f(z^t, u_H^t)$. The continuous evolution of the belief $b^t(\lambda^t)$ can be described by:

$$\dot{b}^t(\lambda^t) = \gamma \left(b_+^t(\lambda^t | u_H^t, x_H^t) - b^t(\lambda^t) \right) + k(b^t(\lambda^t)) \quad (6)$$

for any specific value of λ^t . Here, the function $k(\cdot)$ represents the intrinsic changes in the human intent, whereas the other component captures the Bayesian change in $b^t(\lambda^t)$ due to the observation u_H^t . Note that the time derivative in (6) is pointwise in the space of all λ ’s. Typically, the Bayesian update is performed in discrete time when the new observations are received. However, to unify this with continuous-time robust controls tools, in this work, we reason about continuous changes in $b^t(\lambda^t)$. Intuitively, to relate the continuous-time Bayesian update to the discrete-time version, γ in (6) can be thought of as the observation frequency. Indeed, as $\gamma \uparrow \infty$, i.e., observations are received continuously, $b^t(\lambda^t)$ instantaneously changes to $b_+^t(\lambda^t | u_H^t, x_H^t)$. On the other hand, as $\gamma \downarrow 0$, i.e., no observation are received, the Bayesian update does not play a role in the dynamics of $b^t(\lambda^t)$. For a detailed derivation of continuous dynamics, we refer the interested readers to Appendix A.

Now that we are able to track the evolution of the robot’s belief and the human’s physical states, we can prune unlikely human actions by combining the intent-driven model and the *predicted* belief over the human model parameters. For some future time $\tau \in [0, N]$, the marginalized human action distribution at joint state z^τ is given by

$$p(u_H^\tau | z^\tau) = \sum_{\lambda \in \Lambda} p(u_H^\tau | x_H^\tau; \lambda) b^\tau(\lambda). \quad (7)$$




Fig. 2: Effect of the belief and the δ -threshold on the admissible set of controls (shown in upper-left inset) and the overall predictions (shown in pink) for 3 seconds into the future.

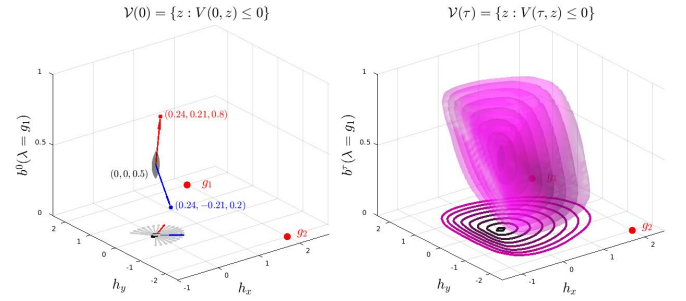


Fig. 3: (left) Initial set in z -space (in grey). Likely control distribution for $\delta = 0.01$ shown projected in $h_x - h_y$ plane. Comparisons of the resulting joint state if the human moves directly towards g_1 (in red) versus towards g_2 (in blue). (right) 4 seconds BA-FRS and its projection into x_H -space.

Very importantly, note that this set is joint state dependent, and therefore *belief*-dependent. This allows us to prune away the sufficiently unlikely actions by removing actions which are not assigned sufficient probability under the future predicted belief (and not just the initial belief):

$$u_H^\tau \in \mathcal{U}(z^\tau), \quad \mathcal{U}(z^\tau) = \{u_H^\tau : p(u_H^\tau | z^\tau) \geq \delta\} \quad (8)$$

where $p(u_H | z)$ is computed as in Eq. (7) and δ is a threshold that partitions the actions into likely and unlikely.

Running example: Consider the case when the intrinsic behavior of the human does not change over time, i.e., $k(b^t(\lambda^t)) = 0$, meaning the human has a fixed goal they are moving to. Since λ takes only two possible values, the joint state space is three dimensional. In particular, $z = [h_x \ h_y \ p_1]$, where $p_1 := b^t(\lambda^t = g_1)$ and $b^t(\lambda^t = g_2)$ is given by $(1 - b^t(\lambda^t = g_1))$ so we do not need to explicitly maintain it as a state. The state-dependent control distribution is $p(u_H | z) = p_1 \mathcal{N}(\mu_1, \sigma_1^2) + (1 - p_1) \mathcal{N}(\mu_2, \sigma_2^2)$ and can be used to compute the set of allowable controls \mathcal{U} for different values of δ via Eq. (8). Note Fig. 2 where the top-left inset figures show the allowable controls for $x = (0, 0)$ and two different belief states $b^0(\lambda = g_1) = 0.5$ and $b^0(\lambda = g_1) = 0.9$ for three different δ thresholds.

A. Using HJ-Reachability for Prediction

Using a control set rather than a distribution results in a prediction problem which can be readily solved using the

HJB-PDE formulation in Section III-A. At any real time t , given the current state of the human and the current belief over the model parameters, we can construct the joint state at the beginning of the prediction horizon $z^0 := z^t$. Using this initial state and the thresholded control policy from (7), we are interested in computing the following set:

$$\mathcal{V}(\tau) := \{z^\tau : \exists u_H(\cdot) \in \mathcal{U}(z), z^\tau = \xi(z^0, \tau, u_H(\cdot))\}, \quad (9)$$

where $\tau \in [0, N]$. Intuitively, $\mathcal{V}(\tau)$ represents all possible states of the joint system, i.e., all possible human states and beliefs over λ , that are reachable under the dynamics in (5) for some sequence of human actions. We refer to this set as the *Belief Augmented Forward Reachable Set (BA-FRS)* from here on. Much like the computation of the full forward reachable set from Sec. III-A, we can leverage the same tools from HJ-Reachability to compute the BA-FRS where x_H is replaced with z and instead of optimizing over all controls \mathcal{U} , we use the restricted set of controls $\mathcal{U}(z)$ instead.

After solving the dynamic programming problem to obtain the BA-FRS from (9), our predictions include not only the physical locations of the human but also the corresponding future beliefs. However, for motion planning, the robot needs to collision-check against a set of physical states the human could occupy. We obtain this set by projecting $\mathcal{V}(\tau)$ on the human state space via $\mathcal{K}_\delta^t(\tau) = \bigcup_{z^\tau \in \mathcal{V}(\tau)} \Pi(z^\tau)$, $\forall \tau \in [0, N]$ where $\Pi(z)$ is the physical state component of z .

Running example: Our starting set of states, \mathcal{L} , is a small ball at the joint starting state $z^0 = [0, 0, 0.5]$, shown in grey in Fig. 3. Consider how the state and belief can change in a small ($\delta t = 0.4668$) timestep after observing the person moving towards goal 1 via $u_H = \pi/4$. Since this action is highly likely under the model where $\lambda = g_1$, then the next joint state will have the person not only moved physically in that direction, but the posterior will have increased probability mass on $b(\lambda = g_1)$. Similarly, this probability decreases if the human moves to g_2 . After solving for $V(\tau, z)$ via (2), we take the sub-zero level set to retrieve the joint state predictions (Fig. 3, right), and the predictions \mathcal{K}_δ^t after projecting onto the human's state space.

In summary, to predict the human's motion, our predictor optimizes the initial value HJB-PDE from (2) but instead of optimizing over *all* controls, our formulation modifies Eq. (2) to maximize over the restricted set $u_H \in \mathcal{U}(z)$ which changes based on human state, time, and belief. Ultimately, the proposed prediction framework is a *less conservative* FRS, but a *more conservative* intent-driven predictor. This has two advantages: (1) when the intent-driven model is correct, it computes an under-approximation of the full FRS to reduce conservatism in a principled way, and (2) when the model is incorrect, we can be more robust to such inaccuracies since the predictions no longer rely on the exact action probabilities. We discuss this further in Sec. V.

V. PREDICTION COMPARISONS

We now compare our predictor with the intent-driven Bayesian predictor and the full FRS when (1) the human

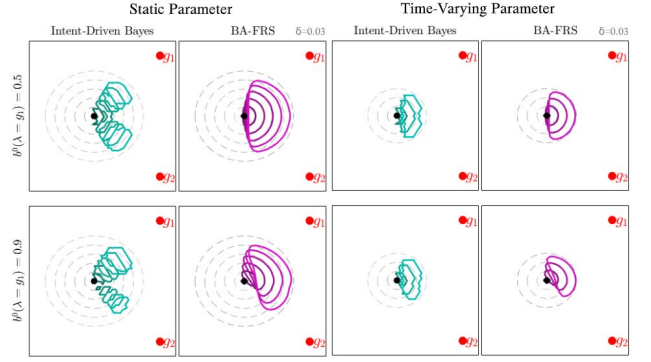


Fig. 4: Comparisons of Bayesian and BA-FRS predictions for static vs. time-varying human intent. Dashed lines are the full FRS. Predictions are for 2 seconds for the static parameter and 1.8 seconds for time-varying. For the Bayesian predictor, we choose ϵ to capture the $(1 - \delta)$ most likely states.

intent is static, (2) the human intent is time-varying, and (3) the human moves in unmodelled ways over time.

A. Static parameter

One simple but common predictive model of human behavior assumes that the human's intent (and thus model parameter λ) is static. In our running example, this means the person has a fixed goal location they are moving towards and they will not change their mind. Mathematically, in the intent-driven Bayesian predictor, this is represented via the λ transition distribution $P(\lambda^{k+1} | \lambda^k) = \mathbb{1}_{\{\lambda^{k+1} = \lambda^k\}}$ where $\mathbb{1}$ represents the indicator function. In the BA-FRS predictor it means $k(b^t(\lambda^t)) = 0$ in the distribution dynamics.

In the left block images in Fig. 4, we see a snapshot of predictions generated by the intent-driven predictor and the BA-FRS forward in time for 2 seconds ($N = 18$). The full forward reachable set is visualized as a series of concentric dashed grey circles. The top row represents a uniform belief over the two goals, while the bottom row represents a high belief on goal 1 (g_1). As expected, both the intent-driven predictor and the BA-FRS are far less conservative than the full FRS. Furthermore, the set of sufficiently likely states predicted by the intent-driven Bayesian predictor is always contained within the BA-FRS. Consequently, when the belief over λ is confident that the human is moving towards g_1 (see bottom row of Fig. 4), then the BA-FRS allows us to compute an approximation of the stochastic predictor.

B. Time-varying parameter

A more complex model of human intent allows it to vary over time. To encode this time-varying intent in the predictor, we need a model of how the human chooses the next value of λ^t . In our running example, let a simple model for how the person changes their behavior to be:

$$P(\lambda^{k+1} | \lambda^k) = \begin{cases} \alpha + (1 - \alpha) \cdot \Delta(\lambda^k) & \text{if } \lambda^{k+1} = \lambda^k \\ (1 - \alpha) \cdot \Delta(\lambda^{k+1}) & \text{if } \lambda^{k+1} \neq \lambda^k \end{cases} \quad (10)$$

where α is a known model parameter which governs how likely the person is to change their intent and Δ is a known discrete distribution over the model parameters. This model encodes that if the person was moving towards $\lambda^k = g_1$ at the previous timestep k , they are likely to continue to walk to g_1

at the next timestep with probability $\alpha + (1-\alpha) \cdot \Delta(\lambda^k = g_1)$, or they can switch to $\lambda^{k+1} = g_2$ at the next timestep with probability $(1-\alpha) \cdot \Delta(\lambda^{k+1} = g_2)$. In the BA-FRS predictor, this time-varying intent model is encoded via the distribution dynamics: $k(b^t(\lambda^t)) = \alpha b^t(\lambda^t) + (1-\alpha) \cdot \Delta(\lambda^t) - b^t(\lambda^t)$.

In the right block of images in Fig. 4, we see a snapshot of predictions when the latent parameter is time-varying forward in time for 1.8 seconds ($H = 11$). Note that when the parameter is time-varying, the computational complexity of the intent-driven Bayes predictor exponentially increases in the size of the prediction horizon, $|\Lambda|^N$, due to the necessity of tracking all sequences of values that λ can take over time. In practice, prediction was computationally prohibitive for horizons greater than 1.8 seconds. In contrast, the BA-FRS computation grows linearly in the length of the prediction horizon, but exponentially in the number of parameter values due to the addition of the belief in the state. Thus, for time-varying parameters which take a few values and for longer prediction horizons, our prediction method can be particularly suitable for getting an approximation of Bayes predictor at a lower computational complexity.

When λ is static, then the intent-driven predictor with a high belief over g_1 deems moving directly towards g_2 to be highly unlikely. However, when λ is time-varying, the human can “switch” which goal they are moving towards, thereby making states in the direction of g_2 somewhat likely as well. For the BA-FRS, even though directly moving towards g_2 is unlikely under the intent-driven model and belief, the BA-FRS realizes that moving *away* from g_1 is *still likely enough*. Consequently, the predicted BA-FRS mass moves in the direction of g_2 over time, in the case of both static and time-varying λ , allowing us to be robust to suboptimal human trajectories as discussed in the next section.

C. Online updates & robustness to misspecified models

Ultimately, both the intent-driven Bayesian predictor and the BA-FRS will update the belief over the human parameters online based on how the person moves. Here we simulate three scenarios—one where the person takes a path well-modelled by the intent-driven model and two where the person behaves in an unmodelled way—and discuss how our framework ensures robustness in situations like these.

In all examples, the predictors begin with a uniform prior over the two goals, use a static model of human intent, and the BA-FRS uses a $\delta = 0.02$. In the top row of Fig. 5 the human has a fixed intent to move towards the upper goal 1 (g_1). In this scenario, the intent-driven model is correctly specified and as the person moves towards g_1 , the belief over g_1 increases and the Bayesian predictions focus towards this goal. Our BA-FRS performs similarly since it too performs the belief update over time. However, since the BA-FRS explicitly tracks the evolution of the belief in the future during prediction, the sets include more states even in the direction of g_2 . This is because the predictions are safeguarding against slightly suboptimal actions which are still likely enough under the model and would lead to the

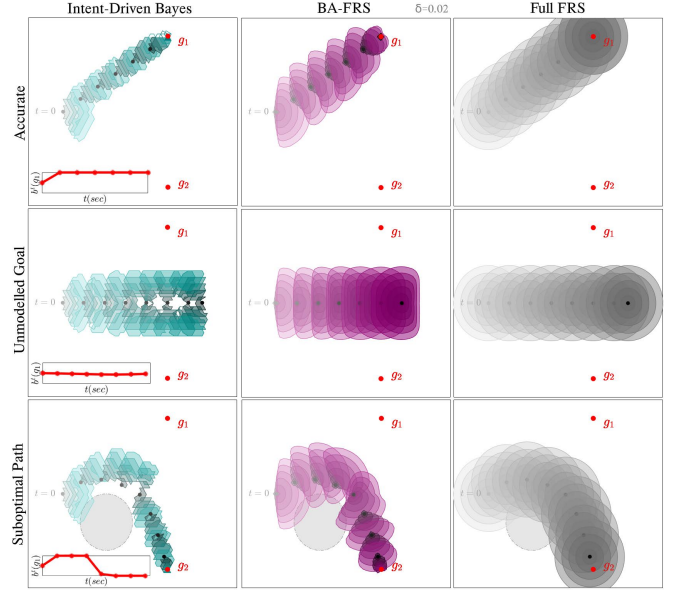


Fig. 5: Comparison of the intent-driven Bayesian, our BA-FRS, and the full FRS predictions for three scenarios. In the first row the human moves towards one of the modelled goals. In the middle the human moves towards an *unmodelled* goal. In the last row the human is moving towards a modelled goal (g_2) but they take a suboptimal path under our model because they are avoiding an *unmodelled* obstacle on the ground (shown in grey circle). The belief over g_1 is visualized over time in the lower-left inset plot.

belief over g_1 *decreasing* in the future. Nevertheless, the BA-FRS takes up significantly less volume than the full FRS, thereby reducing overall conservatism.

The second and third rows demonstrate two human behaviors that are unmodelled – a scenario where the human is actually walking towards a third unmodelled goal in between g_1 and g_2 and a scenario where the human takes a seemingly suboptimal path to g_2 due to an unmodelled obstacle. In the later scenario, the belief over g_1 sharply increases as the person moves around the unknown obstacle. This results in the Bayesian predictor being overly optimistic, and it places most probability mass on states that are in the direction of g_1 . In contrast, our predictor remains cautiously conservative because (1) it is safeguarding against the slightly suboptimal but still sufficiently likely actions and (2) it is evolving the belief during the prediction horizon. In fact, the true sequence of human states and actions lies within the predictions of the BA-FRS, ensuring that a robot which relies on these predictions will in fact avoid the states that the human eventually occupies. We discuss the middle scenario in greater detail in Sec. VI.

We conducted a series of additional experiments with simulated human trajectories to systematically analyze the three misspecification types: (1) accurate goal but inaccurate optimality level, (2) unmodelled goal, (3) accurate goal but unmodelled obstacle. We compare different predictors for prediction accuracy and conservatism. A predictor is considered accurate at a particular time step if the true future human states lies within the predictions for the entire prediction horizon. Conservatism is measured by computing the percent volume of the full FRS that the predicted states occupy. Both the accuracy and conservatism metrics are computed

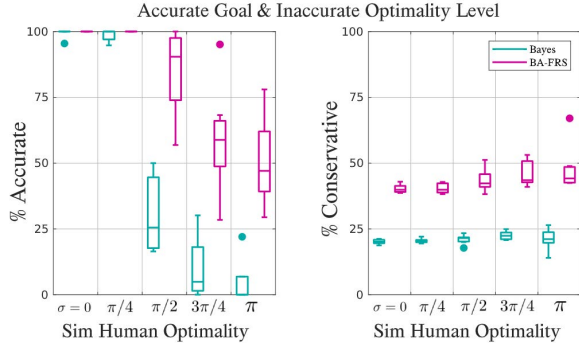


Fig. 6: Simulated human moves to modelled g_1 but with varied optimality from $\sigma = 0$ (optimal) to π (random). Both predictors use a fixed $\sigma = \pi/4$.

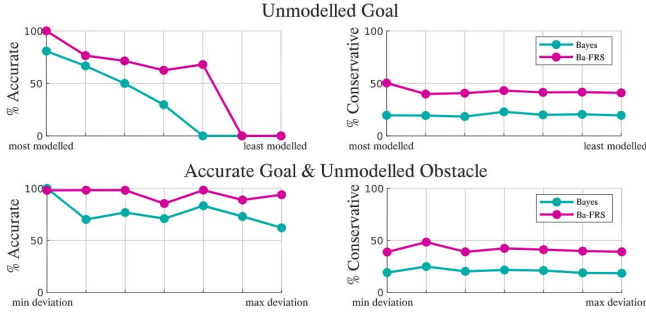


Fig. 7: (top) Simulated human moves to one of 7 *unmodelled* goals in an optimal trajectory. Results are in increasing misspecification of the goal. (bottom) Simulated human moves to modelled goal g_2 but has their straight-line path obstructed by an *unmodelled* obstacle. Results from 7 unmodelled obstacles are shown in increasing deviation from the straight trajectory.

at each time step and averaged over the horizon $[0, \bar{T}]$. Note that the full FRS always achieves 100% accuracy but also 100% conservatism. These metrics provide us a proxy for the prediction’s effect on the safety and efficiency of the robot’s plan; ideally, a predictor should have high accuracy and low conservatism over the entire human trajectory. In all experiments, the Bayesian and BA-FRS predictors modelled two goals and used a fixed $\sigma = \pi/4$ in the action model described in the running example from III-B and $\delta = 0.02$.

For (1), the human was simulated as moving towards modelled goal g_1 by sampling an action $u(x) \sim \mathcal{N}(\mu_1, \sigma^2)$. To capture a range of human behavior from completely optimal to completely random, we simulated five levels of σ (depicted in Fig. 6). We sampled 7 random initial human states for each σ and averaged results over these trials. Fig. 6 shows box plots of our metrics for Bayesian and BA-FRS. Although the BA-FRS is about twice as conservative as Bayesian, it maintains a high prediction accuracy across all optimality levels, while still being far less conservative than the full FRS (BA-FRS is $\approx 45\%$ of the full FRS).

For (2) and (3) we fixed the human’s initial condition but varied the unmodelled goal or unmodelled obstacle. For unmodelled goals, we randomly sampled 7 unmodelled goals which were diversely spread in the (x,y)-plane. The true (unknown) human trajectory is a straight line to the unmodelled goal starting from the initial position. Fig. 7 (top) shows plots of the accuracy and conservatism for each of the unmodelled goals, sorted from the “most” modelled (e.g. an unmodelled goal which is nearby a modelled goal)

to “least” modelled. For unmodelled obstacles, the simulated human always moved towards g_2 , but their straight-line path was always obstructed by an unmodelled obstacle, forcing them to take a trajectory around the obstacle. We simulated 7 of these trajectories around various circular and rectangular-shaped obstacles. Fig. 7 (bottom) shows the results sorted from least deviation from straight-line trajectory to the goal to most deviation. The more irrational the human “appears” (either due to an unmodelled goal or taking a suboptimal path to the goal), the more the drop in accuracy of the Bayesian predictor, as it overrelies on the intent model to explain the human’s behavior. In contrast, since BA-FRS only uses the human model to filter likely and unlikely actions, it maintains a relatively higher accuracy.

VI. IMPLICATIONS FOR SAFE MOTION PLANNING

Consider the scenario where the actual human goal is midway between the modelled goals g_1 and g_2 (see g_3 label in Fig. 1 and middle row of Fig. 5), but the true human goal is not explicitly modelled in the intent-driven model. We will use this example in simulation and in hardware to demonstrate the challenges with over-relying on a misspecified intent-driven model. Our hardware experiments are performed on a TurtleBot 2 navigating around a human pedestrian. We measured human positions at 200Hz using a VICON motion capture system and used on-board odometry sensors for the robot state measurement. The robot is modelled via the dynamics in Sec. II, its goal g_R is behind the initial state of the human (green circle in Fig. 1) and it uses a spline-based planner [24] to plan six-second trajectories in a receding-horizon fashion.

When the robot uses the full FRS for human motion prediction (see Fig. 5 for visualizations of the predictions over time), the robot plans a trajectory which deviates significantly from the ideal straight line path towards its goal and in fact forces the robot to leave the testbed³. In contrast, the Bayesian predictor consistently predicts that that pedestrian will walk towards one of the goals and fails to assign sufficient probability to the true human states because of its over reliance on the model. Ultimately, this leads to a collision between the human and the robot (top row Fig. 1). Our proposed approach does not rely heavily on the exact action probabilities, and infers that the straight line trajectory is likely enough under the pedestrian model. As a result, the robot makes a course correction early on to reach its goal without colliding with the pedestrian (bottom row Fig. 1).

VII. CONCLUSION

When robots operate around humans, they often employ intent-driven models to reason about human behavior. Even though powerful, such predictors can make poor predictions when the intent-driven model is misspecified. This in turn will likely cause unsafe robot behavior. In this work, we formulated human motion prediction as a robust control problem over the set of only sufficiently likely actions, offering

³Hardware demonstration videos: <https://youtu.be/uZi-zLi1S6A>

a bridge between conservative full forward reachable set predictors and intent-driven predictors. We demonstrated that the proposed framework provides more robust predictions when the prior is incorrect or human behavior model is misspecified, and can perform these predictions in continuous time and state using the tools developed for reachability analysis. In the future, we will scale it to higher dimensions with multiple humans, perform a user study to gauge the impact of our predictor on a humans' comfort in close navigation scenarios, and integrate it with online model confidence estimation approaches.

APPENDIX A CONTINUOUS-TIME DISTRIBUTION DYNAMICS

Assume that every T seconds⁴ we are guaranteed to receive a measurement, u_H . Let the current time be denoted by t and the current belief over the latent parameter λ be $b^t(\lambda)$. During the time interval $[t, t + T]$ we will receive a single measurement with probability 1. For simplicity, we assume that the arrival time of the measurement is uniformly distributed in the interval $[t, t + T]$. However, depending on the measurement model, a similar derivation can be performed for other arrival time distributions as well. We want to understand how the belief $b^t(\lambda)$ can change in an arbitrarily small timestep, δt , along $[t, t + T]$. As this timestep goes to zero, we will be able to recover the continuous-time dynamics of $b^t(\lambda)$. Let E be a discrete random variable which takes values capturing the event that we receive a measurement within the time interval $[t, t + \delta t]$. The probability distribution of E can be written out as:

$$E = \begin{cases} e_1 = \text{measurement,} & \text{with probability } \frac{\delta t}{T} \\ e_2 = \text{no measurement,} & \text{with probability } 1 - \frac{\delta t}{T} \end{cases}$$

When a measurement is received, a Bayesian update on the belief is performed as per (3). When a measurement is not received, the only change in the belief is intrinsic. Let the function $k(b^t(\lambda))$ represent the intrinsic changes in the human behavior. Using the law of total probability:

$$\begin{aligned} b^{t+\delta t}(\lambda) &= P(e_1)P^{t+\delta t}(\lambda | x, e_1) + P(e_2)P^{t+\delta t}(\lambda | x, e_2) \\ &= \left(\frac{\delta t}{T}\right) b_+^t(\lambda | u, x) + \left(1 - \frac{\delta t}{T}\right) (b^t(\lambda) + \delta t \cdot k(b^t(\lambda))). \end{aligned}$$

Rearranging some terms, we get:

$$\begin{aligned} b^{t+\delta t}(\lambda) - b^t(\lambda) &= \left(\frac{\delta t}{T}\right) (b_+^t(\lambda | u, x) - b^t(\lambda)) \\ &\quad + \delta t \cdot k(b^t(\lambda)) + \text{h.o.t.}, \end{aligned}$$

where h.o.t includes all the terms with δt^2 in them. Taking the limit as $\delta t \rightarrow 0$ we get the time-derivative of $b^t(\lambda)$: $\dot{b}^t(\lambda) = \frac{1}{T} (b_+^t(\lambda | u, x) - b^t(\lambda)) + k(b^t(\lambda))$. Note that the higher order terms disappear when we take the limit of $\delta t \rightarrow 0$. We now have a form for our dynamics when our belief can change both because of a new measurement and

because of the intrinsic dynamics of the human. If we let $\gamma = 1/T$ then we get the form in Equation (6):

$$\dot{b}^t(\lambda) = \gamma (b_+^t(\lambda | u, x) - b^t(\lambda)) + k(b^t(\lambda)). \quad (11)$$

REFERENCES

- [1] S. Bansal*, A. Bajcsy*, E. Ratner*, et al. "A Hamilton-Jacobi reachability-based framework for predicting and analyzing human motion for safe planning". *ICRA*. 2020.
- [2] I. Mitchell, A. Bayen, and C. J. Tomlin. "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games". *IEEE TAC*. Vol. 50. 7. 2005.
- [3] K. Driggs-Campbell, R. Dong, and R. Bajcsy. "Robust, Informative Human-in-the-Loop Predictions via Empirical Reachable Sets". *IEEE Trans. on Intelligent Vehicles*. 2018.
- [4] H. Bai, S. Cai, N. Ye, et al. "Intention-aware online POMDP planning for autonomous driving in a crowd". *ICRA*. 2015.
- [5] A. Rasouli, I. Kotseruba, T. Kunic, and J. K. Tsotsos. "PIE: A large-scale dataset and models for pedestrian intention estimation and trajectory prediction". *ICCV*. 2019.
- [6] B. D. Ziebart, N. Ratliff, G. Gallagher, et al. "Planning-based prediction for pedestrians". *IROS*. 2009.
- [7] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert. "Activity forecasting". *ECCV*. 2012.
- [8] S. Casas, W. Luo, and R. Urtasun. "Intentnet: Learning to predict intention from raw sensor data". *CoRL*. 2018.
- [9] H. B. Amor, G. Neumann, S. Kamthe, et al. "Interaction primitives for human-robot cooperation tasks". *ICRA*. 2014.
- [10] H. Ding, G. Reißig, K. Wijaya, et al. "Human arm motion modeling and long-term prediction for safe and efficient human-robot-interaction". *ICRA*. 2011.
- [11] H. S. Koppula and A. Saxena. "Anticipating human activities for reactive robotic response." *IROS*. 2013.
- [12] F. Schneemann and P. Heinemann. "Context-based detection of pedestrian crossing intention for autonomous driving in urban environments". *IROS*. 2016.
- [13] W.-C. Ma, D.-A. Huang, N. Lee, and K. M. Kitani. "Forecasting Interactive Dynamics of Pedestrians With Fictitious Play". *CVPR*. 2017.
- [14] C. Rösmann, M. Oeljeklaus, F. Hoffmann, and T. Bertram. "Online trajectory prediction and planning for social robot navigation". *AIM*. 2017.
- [15] A. Rudenko, L. Palmieri, M. Herman, et al. "Human Motion Trajectory Prediction: A Survey". *IJRR*. 2019.
- [16] J. Fisac, A. Akametalu, M. Zeilinger, et al. "A general safety framework for learning-based control in uncertain robotic systems". *TAC*. 2018.
- [17] P. A. Lasota and J. A. Shah. "A multiple-predictor approach to human motion prediction". *ICRA*. 2017.
- [18] T. Bandyopadhyay, K. S. Won, E. Frazzoli, et al. "Intention-aware motion planning". *Algorithmic Foundations of Robotics X*. Springer, 2013.
- [19] M. J. Kochenderfer, M. W. M. Edwards, L. P. Espindle, et al. "Airspace encounter models for estimating collision risk". *Journal of Guidance, Control, and Dynamics*. 2010.
- [20] I. Mitchell. "A toolbox of level set methods". <http://www.cs.ubc.ca/mitchell/ToolboxLS/toolboxLS.pdf>. 2004.
- [21] X. Chen, E. Ábrahám, and S. Sankaranarayanan. "Flow*: An analyzer for non-linear hybrid systems". *CAV*. 2013.
- [22] J. Lygeros. "On reachability and minimum cost optimal control". *Automatica*. Vol. 40. 6. Elsevier, 2004.
- [23] C. Finn, S. Levine, and P. Abbeel. "Guided cost learning: Deep inverse optimal control via policy optimization". *ICML*. 2016.
- [24] R. Walambe, N. Agarwal, S. Kale, and V. Joshi. "Optimal trajectory generation for car-type mobile robot using spline interpolation". *IFAC*. 1. 2016.

⁴Here, T can represent the publishing rate of the motion capture or estimator which computes the current human state (and then observation).