Confidence-Aware Safety for Human-Robot Systems

Jaime Fisac* Andrea Bajcsy* Sylvia Herbert David Fridovich-Keil Steven Wang Claire Tomlin Anca Dragan

As the domain of robotic systems expands from structured industrial settings to unstructured human-dominated environments—from autonomous drones flying in public spaces to service robots helping in the home—it becomes increasingly important for these systems to appropriately reason about the intent and imminent actions of human agents. Predicting human motion becomes particularly important in safety-critical contexts, in which certain failure modes, such as a collision between the robot and the human, are unacceptable.

Promising methods for robotic navigation among humans often incorporate online human motion predictions into the robot's own motion planning [3], [4]. Unfortunately, no model can be expected to perform to perfection, and mismatch between the model and reality has a direct impact on model-based safety assertions. Two key sources of error can be identified: the human's behavior and the robot's dynamics.

Human intent modeling has advanced considerably in recent years, with modern approaches based on inverse optimal control [6] yielding powerful predictive tools. In practice, however, the robot's model of the human will not be able to capture all her possible eventual behaviors. For example, the human might walk towards a goal that the robot does not know about, or move to avoid an obstacle of which the robot is unaware (Fig. 1). In cases where the human's motion escapes the model's predictive power, safety might be compromised.

On the other hand, robots should ideally plan their motion based on a high-fidelity model of their dynamics, accounting for inertia, actuator limits, and environment disturbances. Unfortunately, reasoning with such complex models is almost always computationally prohibitive, and motion plans are instead computed with a simplified representation of the robot's physical dynamics; this can cause non-negligible tracking errors when the robot attempts to execute the planned trajectory. As a result, safety guarantees naively given for the planned trajectory may not in fact apply to the actual robot's motion.

We propose reinterpreting the rationality coefficient in the commonly used Boltzmann model of human decision-making as a measure of the model's ability to capture the human's motion. By reasoning about this *model confidence* in real time, the robot can quickly modulate its probabilistic predictions to become more uncertain when the model performs poorly. The predictions are then incorporated into a robust motion planning scheme to obtain *probabilistically safe* motion plans that are conservative when appropriate but efficient when possible.

The work outlined here makes two key contributions: (1) a Bayesian framework for reasoning about the uncertainty inherent in a model's prediction of human movement, and (2) extending the recent FaSTrack robust motion planning framework [2] to incorporate time-varying, probabilistic obstacles (and, in particular, uncertain human predictions). Our new analysis therefore combines worst-case tracking guarantees for the physical robot with probabilistic predictions of the human's motion, yielding a quantitative probabilistic safety certificate.

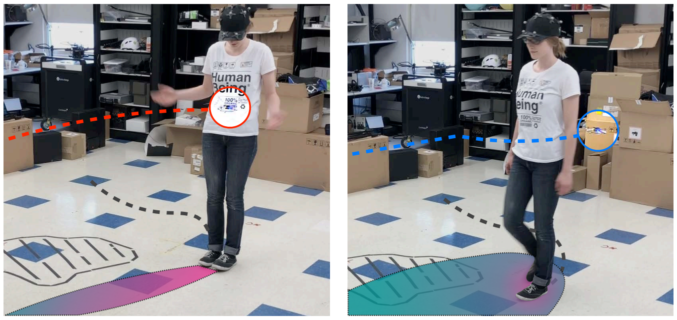

Fixed confidence **Bayesian confidence**

Fig. 1: An autonomous quadcopter navigating around a human, using probabilistic predictions of her future motion. The human deviates to avoid an obstacle that the robot has no knowledge of. Naively trusting imperfect human models can lead to overly confident predictions that result in loss of safety (left). Our method updates *model confidence* in real time to increase conservativeness whenever predictions are expected to be less reliable (right).

System dynamics. We consider a single robot moving in a space shared with a single human. Let the state of the human be $x_H \in \mathbb{R}^{n_H}$, and similarly define the robot's state, for planning purposes, as $x_R \in \mathbb{R}^{n_R}$. These states could represent the locations of a mobile robot and a human in a shared environment, or the kinematic configurations of a human arm and a robotic manipulator in a workspace. The discrete-time dynamics are $x'_H = f_H(x_H, u_H)$ and $x'_R = f_R(x_R, u_R)$. where $u_H \in \mathbb{R}^{m_H}$ and $u_R \in \mathbb{R}^{m_R}$ are the control actions of the human and robot, respectively. Further, let $s_R \in \mathbb{R}^{n_S}$ denote the state of a higher-fidelity dynamical robot model, capturing e.g. velocities and actuator delays ignored by x_R.

Robot objective. The robot needs to plan an ideal trajectory $x_R^{0:T}$ minimizing some cumulative cost c while ensuring that its execution by the physical system $s_R^{0:T}$ remains collision-free with high confidence:

$$\min_{u_R^{0:T}} \sum_{t=0}^{T} c(x_R^t, u_R^t) \tag{1a}$$

$$\text{s.t. } x_R^0 = x^i, \quad x_R^{t+1} = f_R(x_R^t, u_R^t),\ t \in 0, ..., T-1, \tag{1b}$$

$$P_{\text{coll}}^{0:T} := P\big(\exists t \in \{0, ..., T\} : \text{coll}(x_R^t, x_H^t)\big) \leq P_{\text{th}}. \tag{1c}$$

The term $\text{coll}(x_R^t, x_H^t)$ is a Boolean variable indicating whether the human and the robot are in collision. Since safety guarantees will inevitably inherit the probabilistic nature of the human prediction, our goal is to find efficient plans that will keep collisions with the human below an acceptable probability P_{th}. To compute $P_{\text{coll}}^{0:T}$, the robot must account for the human's future movements, but also its own deviations from the ideal motion plan.

Confidence-aware predictions. Extensive work in econometrics and cognitive science has shown that human behavior can be well modeled by utility-driven optimization [1], [5]. Thus, the robot models the human as optimizing a reward function, $r_H(x_H, u_H; \theta)$, that depends on the human's state and action, as well as a set of parameters θ (e.g. the human's intended goal location). Given r_H, the robot can compute the human's policy as a probability distribution over actions conditioned on the state. The robot models the human as likely to choose actions with high expected utility, in this case her state-action optimal cost-to-go (Q-value), following a Boltzmann distribution

$$P(u_H^t \mid x_H^t; \beta, \theta) \propto e^{\beta Q_H(x_H^t, u_H^t; \theta)} \ . \tag{2}$$

The term β is traditionally called the *rationality coefficient*, with $\beta = 0$ corresponding to an "irrational" human who chooses actions uniformly at random and $\beta \to \infty$ giving a "perfectly rational" human. Instead, we propose that β can be more pragmatically interpreted as an indicator of the accuracy with which the robot's model can explain the human's motion, and consistently refer to it as *model confidence*. In practice, the same model will perform variably well over time in different settings and for different people. Our analysis treats β as a *hidden state*, enabling the robot to dynamically adapt its motion plan to the current reliability of its human model. As the robot navigates around the human, every new observed human action u_H^t provides a "measurement update" of the robot's belief $b^t(\cdot)$ about β over time, using Bayes' rule

$$b^{t+1}(\beta) = \frac{P(u_H^t \mid x_H^t; \beta, \theta) b^t(\beta)}{\sum_{\hat{\beta}} P(u_H^t \mid x_H^t; \hat{\beta}, \theta) b^t(\hat{\beta})} \ , \tag{3}$$

with prior $b^t(\beta) = P(\beta \mid x^{0:t})$ for $t \in \{0, 1, ...\}$, and likelihood $P(u_H^t \mid x_H^t; \beta, \theta)$ given by (2). It is critical to be able to perform this update extremely fast, which would be difficult to do in the original continuous hypothesis space $\beta \in [0, \infty)$. Fortunately, we observe that maintaining a Bayesian belief over a relatively small set of β values ($N_\beta \approx 10$ on a log-scale) achieves significant improvement relative to maintaining a fixed precomputed value. Marginalizing over $b^t(\beta)$, we then combine (2) with dynamics f_H to obtain $P(x_H^{t+1} \mid x_H^t; \beta, \theta)$, which allows us to recursively propagate the human's motion over time and obtain a discretized probabilistic occupancy grid $P(x_H^{t+\delta} \mid x_H^t; \beta, \theta)$ any number of time steps δ into the future.

Probabilistic safety guarantee. The final step is to integrate this adaptive prediction into the robot's motion planning to obtain a safety guarantee not only for the *planned* state x_R, but for the *actual* state s_R at execution time. The recently proposed FaSTrack framework [2] compares the dynamics of x_R and s_R to produce a guaranteed tracking error bound $\mathcal{E} \subset \mathbb{R}^{n_R}$ between the planned and executed trajectory. We leverage this worst-case tracking result in combination with our probabilistic human prediction: for any planned state x_R^t, we *integrate* the human's occupancy distribution at time t to compute the total probability mass in potential conflict with s_R^t (whose projection onto $\mathbb{R}^{n_R}$ is guaranteed to lie within $\mathcal{E}$ error of x_R^t). Rejecting candidate states with $P\big(\text{coll}(x_R^t, x_H^t)\big) > P_{\text{th}}$ in the planner's collision-checking step, we obtain real-time probabilistically safe motion plans for the physical robot.

Experiments. We have implemented our method on a Crazyflie quadcopter remotely controlled through the Robot Operating System (ROS) framework. We first benchmarked our approach on a real-time quadcopter simulation using 48 pre-recorded human walking trajectories (16 people in 3 environments) obtained by an Optitrack motion capture system. Results suggested that our approach improves both efficiency (completion time) and safety (minimum distance to human and whether a collision took place) relative to fixed-β predictions. We then flew the physical Crazyflie in an environment with a human (Fig. 1) and confirmed that our method quickly replans to leave more space around the human whenever her behavior is not well explained by the predictive model; the fixed-β quadcopter instead kept expecting the human to match modeled motions, and ultimately collided with her.

Discussion and future work. In this initial formulation, we have eschewed the game-theoretic analysis and assumed the human does not react to the robot. This assumption can realistically capture plausible shared-space settings in which lightweight robots (e.g. micro-drones) may be expected to carry out services such as indoor surveillance in a building while minimizing interference with human activity. However, to the extent that a more compliant human may tend to avoid collisions with the robot or otherwise be willing to coordinate with it, richer collaborative approaches may be able to produce less conservative plans. We note that both our proposed model confidence inference approach and our probabilistically safe planning scheme can in principle work under closed-loop human models that treat the human as a reactive agent rather than an "indifferent" dynamic obstacle. Since human-robot safety is ultimately a multiagent problem, we predict that future research in this direction will further improve performance and reliability, extending the principles proposed here to more complex interaction settings.

References

[1] Chris L. Baker, Joshua B. Tenenbaum, and Rebecca R. Saxe. Goal inference as inverse planning. In *Annual Meeting of the Cognitive Science Society*, 2007.

[2] Sylvia L. Herbert*, Mo Chen*, SooJean Han, Somil Bansal, Jaime F. Fisac, and Claire J. Tomlin. Fastrack: a modular framework for fast and guaranteed safe motion planning. *Conf. on Decision and Control (CDC)*, 2017.

[3] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. *International Journal of Robotics Research*, 2016.

[4] Peter Trautman and Andreas Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *International Conf. on Intelligent Robots and Systems (IROS)*, 2010.

[5] John Von Neumann and Oskar Morgenstern. *Theory of games and economic behavior*. Princeton University Press Princeton, NJ, 1945.

[6] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.