

Analyzing Human Models that Adapt Online

Andrea Bajcsy¹, Anand Siththaranjan, Claire J. Tomlin, Anca D. Dragan

Abstract— Predictive human models often need to adapt their parameters *online* from human data. This raises previously ignored safety-related questions for robots relying on these models such as *what* the model could learn *online* and *how quickly* could it learn it. For instance, when will the robot have a confident estimate in a nearby human’s goal? Or, what parameter initializations guarantee that the robot can learn the human’s preferences in a finite number of observations? To answer such analysis questions, our key idea is to model the robot’s learning algorithm as a dynamical system where the state is the current model parameter estimate and the control is the human data the robot observes. This enables us to leverage tools from reachability analysis and optimal control to compute the set of hypotheses the robot could learn in finite time, as well as the worst and best-case time it takes to learn them. We demonstrate the utility of our analysis tool in four human-robot domains, including autonomous driving and indoor navigation.

I. INTRODUCTION

Robots rely on predictive models of human behavior in order to plan safe and efficient motions around people. Because people vary in their intentions, preferences, and styles of behavior, these models must adapt online upon observing a specific person. For instance, a robot may use Bayesian inference to update its belief about the location a pedestrian is walking to [1–3], or use online gradient descent to update parameters corresponding to a human’s proxemics preferences [4, 5] or parameters of a neural network which predicts how a human reaches for objects [6].

Enabling robots to adapt their human models online is necessary and beneficial, but it also raises safety challenges. The human model can now change significantly in light of new data—the human state/actions—that the robot observes. What the model parameters change to and how quickly they change depends both on the robot’s learning algorithm and how the human actually behaves. Some parameter initializations will be conducive to better learning. Some human behaviors may be ambiguous under the model and result in slower learning. These all have significant implications for the safety and efficiency of the robot’s decision-making.

Adaptive models thus raise several questions. First, what is the *worst-case time* it can take the robot to learn the correct model parameters? With this, the robot can, for instance, make contingency plans that safeguard against all intents until this worst-case time but then branch on the true parameter value afterwards. In contrast, what is the *best-case time* to learn? Here, a robot can, for example, gain insights about which locations in a room make learning from nearby

humans so challenging that even in the best-case it still takes too long to estimate the human’s intent. Relatedly, we can also ask what *observations* lead to the *fastest* or *slowest learning*, thereby producing legible or deceptive motions. Finally, what *initialization* should the robot use, so that we can guarantee that the robot learns the correct parameters—regardless of what they might be—in a finite number of observations? These questions amount to knowing *what* the robot can learn and in *how much time*, and thus far have been implicitly raised but received little attention.

In this work, our key idea is that we can answer these analysis questions by posing them as *reachability* queries. To achieve this, we model the robot’s learning algorithm as a dynamical system, where the state is the current parameter estimate, and the control is the human data the robot observes. Answering what model parameters can the robot learn and in how much time reduces to answering *when* and *what* states our dynamical system can reach.

In our formulation, a designer can instantiate a parameterized model of the human (e.g. human driving a car may drive straight or take a left turn at an intersection), an online learning algorithm initialization (e.g. a uniform prior over the human taking a left or driving straight in a Bayesian learning setup), and a desired level of confidence in a particular hypothesis (e.g. acquire high probability on the human taking a left turn), and finally a measurement frequency with which the robot receives data about the human (e.g. human observations are acquired at 10 Hz).

With the components from above, answering analysis questions amounts to solving an optimal control problem whose solution determines which human observations “steer” the learning system into desired states. To determine what our robot could possibly learn, we start our dynamical system with the current estimate of the parameter and solve *forward in time* for the set of hypotheses that are reachable in finite time and observations. We can also compute the minimum (or maximum) time to learn a parameter by solving an optimal control problem *backwards in time* for the earliest time at which there exists a sequence of data points that steer the learning system to the desired parameter.

We demonstrate a variety of use cases for our analysis tool, answering these questions for an autonomous car at an intersection near a human-driven vehicle, and a navigation task in a bookstore around a human with unknown goals and proxemics preferences. While in this paper we focus on low-dimensional parameterizations, we are encouraged by this first-of-its kind tool for analyzing adaptive human models and are excited to explore approximate algorithms that can extend our tool to higher dimensions.

¹Authors with UC Berkeley EECS: {abajcsy, anandsranjan, tomlin, anca}@berkeley.edu. Research funded by NSF CAREER, DARPA Assured Autonomy, NSF VeHICal, & SRC CONIX. NSF GRFP funds A.Bajcsy.

II. RELATED WORK

Learning from human data. Learning from human data has become increasingly popular in robotics, enabling robots to effectively predict human motion like walking or reaching [1, 3, 7, 8], to learn human preferences from physical interaction [4, 5], or to estimate the fidelity of a predictive model [9, 10]. While research on analyzing algorithms and verifying models *already* learned *offline* from human data has garnered some interest [11, 12], to date there has not been research analyzing what models *could* learn.

Analysis of learning algorithms. The advent of modern machine learning has spurred a resurgence of interest in analyzing learning algorithms. While there is a long history of relating ODEs to optimization methods [13], recently, control theoretic tools such as control Lyapunov functions and reachability analysis have been used to prove convergence and safety properties of gradient-based learning methods [14–16], POMDPs [17], and neural networks [18]. While related, our work focuses not only on analyzing learning algorithms which use human data but also leverages control tools which have not yet been studied in learning contexts and have the ability to compute time-to-reach properties.

Verification & control of dynamical systems. The control theory and formal methods communities have a rich history of verifying the behavior of dynamical systems. In contrast to the above learning algorithm analyses, these approaches not only focus on convergence and safety, but also on properties such as minimum time-to-reach [19], robustness to disturbances [20], and controller synthesis [21]. This research lays the foundation for our work wherein we model learning from human data as a dynamical system and we answer analysis queries via solving optimal control problems.

III. PROBLEM FORMULATION & SOLUTION

A. States & Dynamics

The robot observes data in the form of state-action pairs (x_H, u_H) . Let the human's discrete-time dynamics be:

$$x_H^{t+\Delta t} = f_H(x_H^t, u_H^t) \quad (1)$$

where $x_H \in \mathcal{X}$ and $u_H \in \mathcal{U}$. Additionally, Δt represents the frequency with which our robot observes data about the human (e.g. frequency of the state estimator).

Let $\theta \in \Theta$ denote the unknown human model parameter where Θ is a discrete set of values that θ can take on. This parameter could denote a variety of unknown aspects governing human behavior, from how passive or aggressive a person drives [22], what locations in a room they are moving towards [1], how optimally the person behaves [10], or even the kind of visual cues they pay attention to in a scene [3].

The robot uses an online learning algorithm to estimate the human's intent given a sequence of observations over time. Our key idea in this work is to view the robot's learning algorithm as a dynamical system where the parameter estimate is treated as state and the human's data is control input. Let $\hat{\theta}$ be the learning algorithm's estimate. This could be, for example, a point estimate of θ or the belief $b(\theta)$.

Let the discrete-time dynamics of a learning algorithm be

$$\hat{\theta}^{t+\Delta t} = f_L(\hat{\theta}^t, x_H^t, u_H^t) \quad (2)$$

where $f_L : \mathcal{E} \times \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{E}$ is a single update of $\hat{\theta}$ given (x_H^t, u_H^t) and $\hat{\theta} \in \mathcal{E}$ is the space of estimates (e.g. $\mathcal{E} = \Theta$ for a point-estimator or $\mathcal{E} = [0, 1]^{|\Theta|-1}$ for the full belief).

To capture the joint evolution between the learning algorithm's estimate of the human's intent parameter and the human's possible behavior, we consider the joint state: $z := [x_H, \hat{\theta}]^\top \in \mathcal{Z} = \mathcal{X} \times \mathcal{E}$. The human data – the actions the human takes – evolve both the physical human dynamics and the learning dynamics. Thus, we would like to analyze the joint dynamical system

$$z^{t+\Delta t} = f(z^t, u_H^t) := \begin{bmatrix} f_H(x_H^t, u_H^t) \\ f_L(\hat{\theta}^t, x_H^t, u_H^t) \end{bmatrix}. \quad (3)$$

Note that this should not be confused with the robot being able to control the human's actions. This joint system representation allows us to answer reachability questions about which estimates are “learnable” under different restrictions on the type of data the human could produce.

B. Solution Method

Equipped with our joint dynamics which describe the evolution of the human's state and the learning algorithm, we can now formulate an optimal control problem whose solution captures which human observations “steer” the learning system into desired states. We build on our framework from [23, 24] and adapt it for analyzing general discrete-time algorithms that learn online from human data.

We define the discrete-time optimal control problem and its associated value function as

$$V(z) := \min_{\mathbf{u}_H} \min_{t \in \{0, \Delta t, \dots, T\}} l(\xi^t(z, \mathbf{u}_H)) \quad (4)$$

where $\mathbf{u}_H = [u_H^0, \dots, u_H^{T-1}]^\top$ is a sequence of controls over the time horizon and $\xi^t(z, \mathbf{u}_H)$ is the joint state achieved by applying \mathbf{u}_H . Here, the function $l(\cdot)$ encodes the distance between the current system state and the desired values of the parameters we seek to estimate. Different analysis questions simply become different instantiations of this function, as we describe in Sec. IV. Intuitively, this value function captures the *closest* our system ever gets to the target states as measured by the signed distance function l .

This minimum-payoff optimal control problem can be solved via the principle of dynamic programming [25]. Specifically, for a finite horizon $t \in \{0, \Delta t, \dots, T\}$, the discrete-time time-dependent terminal-value Hamilton-Jacobi-Bellman variational inequality [26, 27] is:

$$\begin{aligned} V^t(z) &= \min \left\{ l(z), \min_{u_H \in \mathcal{U}^t} V^{t+\Delta t}(f(z, u_H)) \right\}, \\ V^T(z) &= l(z), \quad \forall z \in \mathcal{Z} \end{aligned} \quad (5)$$

where \mathcal{U}^t is the set of allowable actions—and therefore data—the human can generate (described in detail in Sec. IV).

Given this time-dependent value function and a candidate initial condition of the joint state z^0 , we can extract two important quantities. First, the optimal control at z is:

$$u_H^t(z) = \arg \min_{u_H \in \mathcal{U}^t} V^{t+\Delta t}(f(z, u_H)). \quad (6)$$

Secondly, we define the *time-to-learn* (TTL) as the earliest time when our system will reaches the target parameters we seek to learn starting from z^0 . More formally, this TTL:

$$TTL = \min\{t : V^{T-t}(z^0) \leq 0\}. \quad (7)$$

IV. ENCODING ANALYSIS QUESTIONS

We can now turn our attention to mathematically encoding the analysis questions we are interested in answering. The key components that enable us to encode analysis questions in Eq. (5) are (1) the target set of states \mathcal{L} which is encoded via $l(\cdot)$, (2) the set of human actions, \mathcal{U}^t , and (3) the strategy of the human (if they are minimizing or maximizing value).

Time-to-learn (TTL) queries. Computing the best and worse-case time to learn (TTL) depends not only on the learning algorithm, but also on the value of the parameter we are trying to estimate and the type of data the robot could observe. For simplicity, we first consider computing the TTL for a specific parameter value, θ^* , and later discuss how to integrate multiple TTL estimates.

We can mathematically embed the objective of learning θ^* in a target set defined in joint state space: for example, $\mathcal{L} = \{z : x_H \in \mathcal{X}, \|\hat{\theta} - \theta^*\| \leq \epsilon\}$. Intuitively, these target sets encode that we want our parameter estimate to be close to the true value, but we do not care where the human ends up in physical state-space as long as we have estimated the parameter well. Defining \mathcal{L} to be a closed set in \mathcal{Z} allows us define a function $l(z) : \mathcal{Z} \rightarrow \mathbb{R}$ such that $\mathcal{L} = \{z : l(z) \leq 0\}$. For example, the signed distance function to \mathcal{L} satisfies this property. This function $l(z)$ serves as the cost function for our optimal control problem from (4).

Next, how quickly the robot learns also depends on the possible data about the human the robot observes. Since the observed data is the human behavior, we must define the set of controls \mathcal{U}^t that the human could possibly generate at each time step. In general this set can be chosen in a variety of ways. One of the most straightforward sets is to assume that the robot could observe the human taking *any* action, i.e. $\mathcal{U}^t = \mathcal{U}$. Allowing the human to take any action—and therefore produce any data—leads to very a conservative worst-case TTL, since it allows for the human to abandon any true intent and act purely adversarially. While type of analysis may be desirable in some scenarios, a potentially more realistic set of observations could be $\mathcal{U}^t = \{u_H : P(u_H \mid x_H; \theta^*) \geq \delta\}$ where $P(u_H \mid x_H; \theta^*)$ encodes a state and intent-conditioned action distribution. By varying $\delta \in [0, 1]$, the human is allowed to generate more or less sub-optimal data with respect to the intent-driven model.

Finally, the human’s strategy in the control problem determines if we obtain best or worst-case learning estimates.

Best-case TTL (min). Since our target set \mathcal{L} encodes the true human intent parameter that the robot wishes to learn,

modelling the human as *minimizing* the value which is propagated from l in Eq. 5 encodes a cooperative human. That is, the human is generating data in an attempt to *help* the robot learn their true intent parameter. This enables us to extract the best-case TTL_{θ^*} via Eq. (7). When interested in the best-case time to learn *any* $\theta^* \in \Theta$, we can obtain a conservative best-case TTL via $\max_{\theta^* \in \Theta} TTL_{\theta^*}$.

Worst-case TTL (max). Alternatively, by modelling the human as *maximizing* the value, we can easily encode adversarial behavior where the human is trying to prevent the robot from learning their true θ for as long as possible. This enables us to extract the worst-case TTL_{θ^*} via Eq. (7). Similarly to above, we obtain an upper-bound on learning any $\theta^* \in \Theta$ by computing $\max_{\theta^* \in \Theta} TTL_{\theta^*}$.

A. Reachable parameter queries.

Computing the set of forward reachable parameters given an initial estimate follows a similar setup as the TTL queries. The main difference comes from \mathcal{L} , which now encodes the initial joint state. For example, $\mathcal{L} = \{z : x_H = x_H^0, \hat{\theta} = \hat{\theta}^0\}$. Now, the sub-zero level set of V in Eq. (5), encodes the set of states our system can *reach* in T time.

V. USE CASES OF OUR ANALYSIS TOOL

We now demonstrate a variety of use cases for our tool, ranging from autonomous driving to gradient-based learning.

A. Synthesizing Safe & Efficient Contingency Planners

Motion planners for autonomous vehicles often face uncertainty in how human-driven vehicles will behave. Consider the scenario where an autonomous vehicle is attempting to turn left at an unsignalized four-way intersection and there is a human-driven vehicle in the opposing lane (see Fig. 1). The autonomous vehicle has uncertainty about whether the human will turn left or go straight (goal 1 and goal 2, respectively)—the outcome of which significantly impacts the autonomous car’s ultimate maneuver. This is therefore a prediction problem in which the model’s parameter represents the human’s maneuver goal: $g := \theta$ and $g \in \Theta = \{g_1, g_2\}$.

A safe solution to this problem computes a plan for the autonomous vehicle which safeguards against both events during the entire planning horizon, i.e. avoids collisions with the straight and left-turn trajectory. While safe, this approach often yields conservative and inefficient motions (left Fig. 1).

Alternatively, recent methods have proposed *contingency planning* [28]. In contingency planning, the robot generates a plan which safeguards against all events for a short horizon which is less than the entire planning horizon: $t_b < N$. After t_b , the motion planner generates $|\Theta|$ contingency plans, each of which safeguards against only a single event. The premise of contingency planning is that the robot will know by t_b which branch to choose, and will no longer need to safeguard against all events. The contingency plan is thus only safe if the robot gains enough certainty to choose which plan to use by the branching time. If the branching time is too short, then when it comes time for the robot to choose a contingency plan, the robot will still have uncertainty over

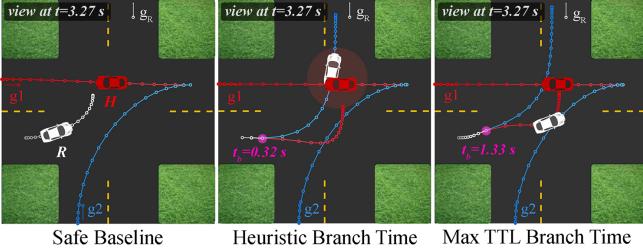


Fig. 1: (left) Robot safeguards against both hypotheses, straight and left, for entire planning horizon, (center) A heuristically chosen branch time for the contingency planner doesn't allow the robot to observe enough human data, leading it to collision, (right) Contingency planner branches at the max TTL computed via our method, enabling a safe but efficient plan.

	Prior	Efficiency (dist to robot's goal)	Safety (min dist between cars)
Safeguard Both	n/a	6.88 m (1.14)	0.73 m (0.95)
Heuristic (0.3265 s)	Correct	5.13 m (1.89)	0.19 m (0.21)
	Incorrect	5.13 m (1.89)	-1.55 m (0.99)
	Uniform	6.10 m (0.79)	-0.76 m (1.19)
Max TTL (ours)	Correct	2.33 m (2.29)	0.58 m (0.77)
	Incorrect	5.25 m (2.88)	0.54 m (0.88)
	Uniform	3.42 m (3.13)	0.55 m (0.77)

TABLE I: Autonomous driving experiment results shown averaged across initial conditions and ground-truth human goals. Mean efficiency and safety metrics are reported in each row and standard deviation in parenthesis.

the human's goal. The robot could now be in a position where no plan exists which safeguards against *all* events that are still likely (center, Fig. 1).

Unfortunately, knowing the future time at which the robot will have certainty about the human's intent is in general challenging. This is where we leverage our analysis framework: we can use it to compute the *worst-case* time it will take the robot to gain certainty in the human intent.

Human Dynamics and Intent Model. Let the human-driven vehicle be modelled as a 3D Dubins' car with planar position and heading as state and discrete-time dynamics as $x_H^{t+\Delta t} = x_H^t + \Delta t[v_H \cos(\phi), v_H \sin(\phi), u_H]^\top$ where the human's control is angular velocity, $u_H \in \{-3.5, 0, 3.5\} \text{ rad/s}$ and is driving with a fixed speed $v = 6 \text{ m/s}$. The human is modelled as choosing actions via the noisily-rational model [29]: $P(u_H \mid x_H, g) \propto e^{Q(x_H, u_H; g)}$ where $Q(x_H, u_H; g)$ encodes the state-action value for the human's driving goal.

Robot Learning Algorithm. The robot learns the human intent by maintaining and updating a Bayesian belief over g . Since g is discrete, the robot can only maintain $|\Theta| - 1 = 1$ probabilities. Without loss of generality, let the robot update $b(g = g_1) := \hat{\theta}$. The learning dynamics f_L in Eq. (2) are

$$f_L(b^t(g_1), x_H^t, u_H^t) := \frac{1}{Z} P(u_H^t \mid x_H^t, g_1) b^t(g_1) \quad (8)$$

where Z is the normalizer.

Joint State & Dynamics. The joint state is $z = [x_H, b(g_1)]^\top$ and the joint dynamics are the stacked dynamics equations from above. We use $\Delta t = 0.0891s$ in all simulations.

Target Set. To determine the maximum time it will take our robot to estimate that the human is going forward (g_1) or left (g_2), we define two target sets in our joint state space:

$$\mathcal{L}_{g_1} = \{z : x_H \in \mathcal{X}, b(g_1) \geq 0.9\} \quad (9)$$

$$\mathcal{L}_{g_2} = \{z : x_H \in \mathcal{X}, 1 - b(g_1) \geq 0.9\} \quad (10)$$

Each of these sets encodes that the robot must be at least 90% confident in the human driving forwards or turning left.

Computing the Worst-Case Time-to-Learn (TTL). Here, we model the human as adversarial and therefore use a max in the inner value function update from Eq. (5). To ensure that while the person is being adversarial, they have to eventually complete their maneuver, we restrict the set of controls to $\mathcal{U}^t = \{u_H : P(u_H \mid x_H, g^*) \geq 0.27\}$ where g^* is equal to the human intent which is being analyzed. Over a horizon of $T = 1.7820s$, we perform two value function computations via Eq. (5) and compute worst-case TTL for g_1 and g_2 by searching backwards in time for the earliest time at which the human's initial state and the robot's initial prior appears in the sub-zero level set of $V^t(z)$ (as in Eq. (7)) to obtain TTL_{g_1} and TTL_{g_2} . Finally, to determine the final safe branching time, we want to safeguard against the hypothesis which takes the longest to estimate confidently. Thus, let $t_b = \max\{TTL_{g_1}, TTL_{g_2}\}$.

Results. We ran a series of simulations comparing the safe motion planner, a heuristically-chosen branching time (comparable to [28]), and our maximum TTL branching time contingency planners. For all planners, we varied the initial velocity of the human and robot cars (stopped, moving slowly, or moving quickly) as well as the two possible goals the simulated human was actually moving to. For the contingency planners, we also varied the prior to correctly biased to the human's true goal, incorrectly biased, or uniform.

Table I summarizes the average efficiency and safety metrics over each of these trials. Here, the heuristically-chosen branching time branches too early – this does not allow the robot to collect enough observations about the human's behavior for the robot to make a confident but safe maneuver, which results in collisions when the robot begins with either a uniform or incorrect prior over the human goals. In contrast, the maximum worst-case TTL safeguards against both events for a longer time horizon during which the robot collects enough observations to make safe and efficiently goal-driven plans even with a uniform or incorrect prior. Note that the heuristic branching time serves to demonstrate how choosing an uninformed t_b can lead to safety violations. However, our analysis tool should be thought of as complementary to [28] wherein we can synthesize an informed t_b .

B. Analyzing Confidence Aware Predictors

Next, we consider predictors that introspect on their model confidence, and use our tool to analyze how long it takes such a predictor to detect that its model cannot explain the observed human behavior. We focus on intent-driven predictors, which model human actions as noisily-optimal with respect to cumulative reward, $P(u_H \mid x_H; \beta) \propto e^{\beta Q(x_H, u_H)}$ [29, 30]. Here, the parameter β models how optimally the human behaves: high values of β model near-optimal behavior, whereas $\beta = 0$ removes the influence of the modelled reward on the human's behavior entirely. Recent work [9, 10, 31] proposed that rather than fixing β , the robot should estimate it. Upon observing human actions that are

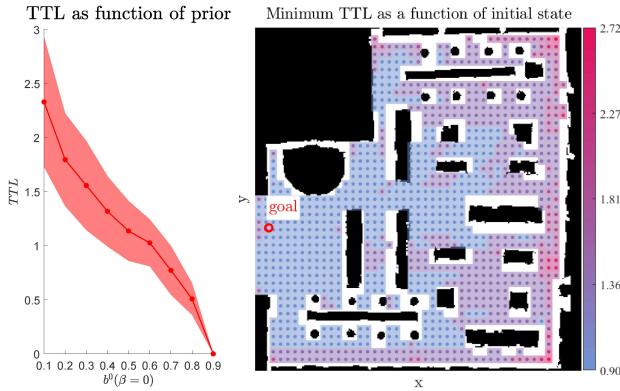


Fig. 2: (left) Minimum TTL that the human is being unmodelled as a function of the prior. Mean and standard deviation shown in red. (right) Minimum TTL as a function of x_H with a uniform prior. Occupancy map of the bookstore environment (in Fig. 4) shown with modelled human goal is red circle and the min TTL for each initial (x, y) state is shown in a color ranging from blue (TTL=0.9s) to red (TTL=2.72s).

poorly explained by the reward function, low values of β (signaling low model confidence) will be the most likely, and our predictor will make higher-variance predictions, accounting for its inability to explain the human’s behavior.

As the human deviates from the model’s assumptions and the predictor makes higher variance predictions, the robot must stay further away from the human to avoid colliding with the now larger set of sufficiently likely states. This begs the question: when the human doesn’t actually optimize the modeled reward, how long will it take the predictor to adapt its β and detect that its model confidence is low?

Here, let the unknown parameter be $\beta := \theta$ with $\beta \in \Theta = \{0, 1\}$, signaling low and high model confidence respectively.

Human Dynamics and Intent Model. Let the human be modelled by a simple planar pedestrian model: $x_H^{t+\Delta t} = x_H^t + \Delta t[v_H \cos(u_H), v_H \sin(u_H)]^\top$ where the human’s control is their heading, $u_H \in \{-\pi, \dots, \pi\}$ and the human is walking at a leisurely speed ($v = 0.6 \text{ m/s}$) or is stopped ($v = 0$). The human’s reward function encourages motion towards the door (shown in red) in the indoor environment (top-down occupancy map shown on right of Fig. 2).

Robot Learning Algorithm. The robot maintains and updates a Bayesian belief over the confidence parameter β . Without loss of generality, let the robot explicitly update $b(\beta = 0) := \hat{\theta}$. Thus, f_L is identical to (8) but with $b(\beta = 0)$.

Joint State & Dynamics. The joint state is $z = [x_H, b(\beta = 0)]^\top$ and the joint dynamics are the stacked physical and learning dynamics from above. Finally, we use $\Delta t = 0.4545 \text{ s}$.

Target Set. We are primarily interested in determining how long it will take our robot to estimate that our model cannot explain the human’s behavior ($\beta = 0$). Thus, we are interested in answering “From the prior over the model confidence, what is the *fastest* our robot could learn that the person is behaving in an unmodelled way?” Our corresponding target set encoding this question is $\mathcal{L}_{\beta=0} = \{z : x_H \in \mathcal{X}, b(\beta = 0) \geq 0.9\}$ where $\epsilon = 0.9$ is our desired confidence.

Computing the Best-Case Time-to-Learn (TTL). We seek to compute fastest time to learn we have low confidence

in our model, since this is the lower bound on reaction time to unmodelled data. Thus, we use \min over \mathcal{U}^t in the optimization from (5) and we optimize over all data the human could generate, since in the worst case the human is not behaving according to the specified reward function at all. In the following two analyses, the human navigates in a bookstore environment [32] whose occupancy map is shown in right of Fig. 2 and 3D model is shown in Fig. 4.

Results: Best-case TTL as function of prior. We first analyzed the min TTL a low model confidence as a function of the prior. After computing one backwards reachability computation, we extracted the TTL for 121 initial x_H states and 8 levels of the prior. The mean and standard deviation across all initial conditions shown in the left Fig. 2. This analysis reveals the added difficulty for the robot to detect its model is wrong if it begins with an optimistic prior.

Results: Best-case TTL as function of initial human state. Right of Fig. 2 shows how the best-case TTL varies as function of the initial x_H in a complex environment. Here we fix a uniform prior over the model confidence and use the same value function computed from above to query for the TTL for 1,010 initial human states. Interestingly, this analysis demonstrates that the best-case TTL is largely impacted by the constraints of the physical environment. If the human begins in the open-space near the door, learning that they do not want to move to the door is easy since the human can directly move away from the door to indicate this mismatch. However, if the human begins in a heavily constrained part of the environment such as the lower right-hand corner, all collision-free actions appear ambiguous under the robot’s likelihood model; that is, moving left or up could either indicate the human intends to move to the door or they intend to move in a completely different direction.

A few interesting takeaways for designers of robot motion planners which rely on confidence-aware models include: (1) if the robot does not re-plan faster than this best-case TTL then the robot will not be able to react quickly enough to the human’s unmodelled behavior and (2) the robot should remain more cautious around the human in constrained parts of the environment due to the increased learning uncertainty.

C. Generating Legible & Deceptive Behaviors

So far we demonstrated how to compute the worst and best-case learning times. We now showcase how our analysis tool can also synthesize the behavior which led to the fastest or slowest learning by our (robot) observer.

We use the same planar pedestrian model from V-B. In our running example, the human is walking around a living room (occupancy map shown in Fig. 3) and the robot has uncertainty about which location the human is navigating to. Let the uncertain human model parameter be $g := \theta$ and $g \in \Theta = \{g_1, g_2\}$. The robot uses a noisy-rational model as above, parameterized by g , and learns via a Bayesian update.

We perform four reachability computations, two of which have a high-confidence in g_1 target set as in (9) and two of which have a high-confidence in g_2 target set (like (10)). However, for each pair of reachability computations, we

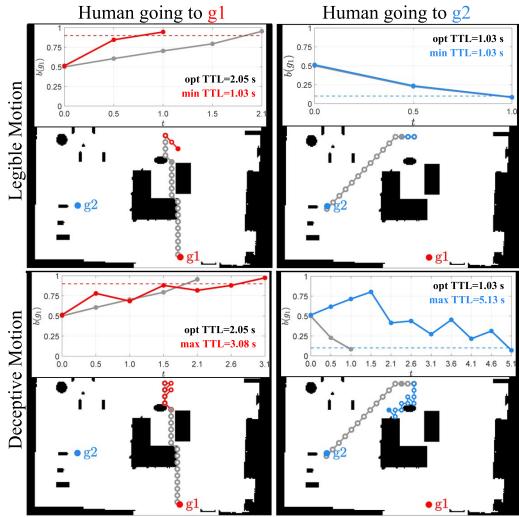


Fig. 3: Legible and deceptive behaviors as synthesized by our analysis tool—shown in bright red or bright blue. The optimal policy for each goal is shown in grey. The estimate of the goal over time for the legible and deceptive behaviors is contrasted with the optimal policy in the inset figures.

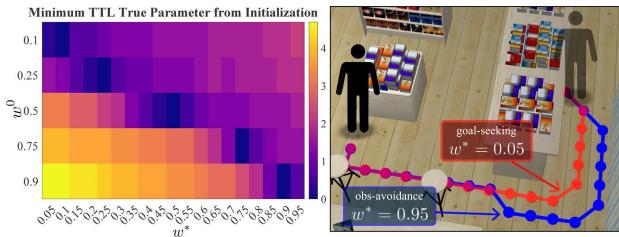


Fig. 4: (left) Heatmap of reachable reward weights (x -axis) and their TTL (color values ranging from dark blue: $TTL = 0.0s$ to yellow: $TTL = 4.7s$) starting from a specific initialization (y-axis). (right) Bookstore environment—red path is optimal behavior for mainly goal-driven human, blue path is for a human who wants to stay far from obstacles.

perform one computation where the human is *minimizing* the value (i.e. helping robot learning) and the other where the human is *maximizing* the value (i.e. trying to slow down robot learning). In all examples, the human chooses from $U^t = \{u_H : P(u_H | x_H, g^*) > 0.15\}$ where g^* the goal being analyzed. The resulting four value functions are used via Eq. (6) to extract the optimal sequence of human data which lead to best and worst-case learning times for g_1, g_2 .

Fig. 3 visualizes the optimal controls in the bright colored trajectory corresponding to the human’s true goal. We contrast this with the optimal only-goal-driven policy in grey. Inset plots show $b(g_1)$ over time and the target confidence level plotted as a dashed line. To increase the probability on g_1 as fast as possible, the human moving to g_1 signals this by quickly moving to their left (instead of moving forward as in the optimal path). This is in line with prior work on legibility [33], but a potential advantage of our formulation is that the objective of minimizing TTL is task-oriented. While prior work encouraged agents to be legible along the entire trajectory, here the agent is directly minimizing the time to convey the goal, so that the observer can react as quickly as possible. Interestingly, to be deceptive, the human zig-zags to confuse the observer for the longest.

D. Online Gradient-based Learning from People

Online gradient-based learning algorithms are also useful in many HRI domains [4, 5, 34, 35]. In this case study, we use our analysis tool to determine a parameter initialization which allows the online gradient algorithm to adapt the fastest to any true human intent. As above, the robot is learning a nearby human’s reward function by observing their behavior. The human’s reward function is modelled as a linear combination of features: $r(x_H, u_H; \theta) = \theta^\top \mu(x_H, u_H)$. Here, $\theta = [1 - w, w]^\top$ and $\mu(x_H, u_H) \in \mathbb{R}^2$, encoding the distance between the human and their goal and the distance between the human and obstacles. Adjusting w trades off the human’s goal-driven and obstacle-avoidance preferences.

The robot learns about the human’s reward via online gradient descent. Thus, the f_L from Eq. (2) are:

$$f_L(\hat{\theta}^t, x_H^t, u_H^t) := \hat{\theta}^t + \alpha \nabla_{\theta} F(x_H^t, u_H^t, \hat{\theta}^t) \quad (11)$$

Maximizing the likelihood of the observed (x_H^t, u_H^t) pair under the maximum entropy distribution [30], we derive a gradient-based update: $F(x_H^t, u_H^t, \hat{\theta}^t) := Q(x_H^t, u_H^t; \hat{\theta}^t) - \mathbb{E}_{u \sim P(u|x_H^t; \hat{\theta}^t)} [Q(x_H^t, u; \hat{\theta}^t)]$. Note: this is the state-action equivalent of learning offline from demonstrations [30].

In this analysis, we solve a *forward reachability* problem where we start our system in the target set $\mathcal{L} = \{z : x_H = x_H^0, \hat{\theta} = \hat{\theta}^0\}$ where x_H^0 is the current physical state of the human and $\hat{\theta}^0$ is a candidate initialization. We compute the set of θ ’s for which there exists a sequence of observations which evolve the joint system to that θ -state in finite time. Here, $\Delta t = 0.2469$ s and the total time horizon for which we evolve our system is $T = 7.1605$ s.

Fig. 4 shows the reachable $\theta^* = [1 - w^*, w^*]$ ’s starting from a given $\theta^0 = [1 - w^0, w^0]$. Colors in the heatmap represent the earliest time at which the robot can learn a θ^* starting from a each initialization. Interestingly, if the robot begins with $w^0 = 0.9$ (i.e. human is primarily obstacle-averse) it takes ~ 4.7 s to learn that the person is *actually* primarily goal-seeking ($w^* = 0.1$). In contrast, initializing with $w^0 = 0.1$ allows the robot to learn *any* other parameter in < 2.2 s. Intuitively, this discrepancy in how quickly the robot can learn from an initialization is because of the structure of the environment which in turn affects the gradient update. Since here the person begins in a part of the environment where their direct path is obstructed by obstacles (right Fig. 4), they must navigate around the obstacles before they get to the goal. The evidence of the person moving *away* from obstacles makes it difficult to disambiguate if they truly are goal-driven or obstacle-averse. Thus, initializations which bias our estimator towards believing that people are obstacle-averse is a poor choice if we want our robot to learn any other θ^* quickly.

Closing Remarks. In this work, we leveraged tools from reachability analysis to analyze human models which adapt online. By treating these models as dynamical systems where the estimate is state and the human data is control, we obtain the best and worst-case time to learn, extract the optimal measurements which enable learning, and synthesize good model parameter initializations.

REFERENCES

- [1] B. D. Ziebart, N. Ratliff, G. Gallagher, et al. “Planning-based prediction for pedestrians”. 2009.
- [2] S. Jain and B. Argall. “Recursive Bayesian human intent recognition in shared-control robotics”. *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2018.
- [3] K. M. Kitani, B. D. Ziebart, J. A. Bagnell, and M. Hebert. “Activity forecasting”. *ECCV*. 2012.
- [4] A. Bajcsy, D. P. Losey, M. K. O’Malley, and A. D. Dragan. “Learning robot objectives from physical human interaction”. *Proceedings of Machine Learning Research* 78 (2017).
- [5] A. Bobu, A. Bajcsy, J. F. Fisac, and A. D. Dragan. “Learning under misspecified objective spaces”. *arXiv preprint arXiv:1810.05157* (2018).
- [6] Y. Cheng, W. Zhao, C. Liu, and M. Tomizuka. “Human motion prediction using semi-adaptable neural networks”. *2019 American Control Conference (ACC)*. IEEE. 2019.
- [7] S. Javdani, S. S. Srinivasa, and J. A. Bagnell. “Shared autonomy via hindsight optimization”. *arXiv preprint arXiv:1503.07619* (2015).
- [8] H. C. Ravichandar and A. Dani. “Human intention inference and motion modeling using approximate em with online learning”. *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE.
- [9] J. F. Fisac*, A. Bajcsy*, S. L. Herbert, et al. “Probabilistically Safe Robot Planning with Confidence-Based Human Predictions”. *Robotics: Science and Systems*. 2018.
- [10] D. Fridovich-Keil*, A. Bajcsy*, J. F. Fisac, et al. “Confidence-aware motion prediction for real-time collision avoidance”. *IJRR*. 2020.
- [11] D. S. Brown and S. Niekum. “Toward Probabilistic Safety Bounds for Robot Learning from Demonstration.” 2017.
- [12] D. Sadigh, K. Driggs-Campbell, A. Puggelli, et al. “Data-driven probabilistic modeling and verification of human driver behavior”. *AAAI Spring Symposium-Technical Report*. 2014.
- [13] J. B. Moore and U. Helmke. “Optimization and Dynamical Systems” (1996).
- [14] A. C. Wilson, B. Recht, and M. I. Jordan. “A lyapunov analysis of momentum methods in optimization”. *arXiv preprint arXiv:1611.02635* (2016).
- [15] M. Laborde and A. Oberman. “A Lyapunov analysis for accelerated gradient methods: From deterministic to stochastic case”. *International Conference on Artificial Intelligence and Statistics*. 2020.
- [16] B. Shi, S. S. Du, M. I. Jordan, and W. J. Su. “Understanding the acceleration phenomenon via high-resolution differential equations”. *arXiv preprint arXiv:1810.08907* (2018).
- [17] M. Ahmadi, N. Jansen, B. Wu, and U. Topcu. “Control theory meets POMDPs: A hybrid systems approach”. *arXiv preprint arXiv:1905.08095* (2019).
- [18] W. Ruan, X. Huang, and M. Kwiatkowska. “Reachability analysis of deep neural networks with provable guarantees”. *arXiv preprint arXiv:1805.02242* (2018).
- [19] I. Yang, S. Becker-Weimann, M. J. Bissell, and C. J. Tomlin. “One-shot computation of reachable sets for differential games”. *Proceedings of the 16th international conference on Hybrid systems: computation and control*. 2013.
- [20] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin. “Hamilton-Jacobi Reachability: A Brief Overview and Recent Advances”. *CDC*. 2017.
- [21] J. Lygeros. “On reachability and minimum cost optimal control”. *Automatica*. Vol. 40. 6. Elsevier, 2004.
- [22] D. Sadigh, S. Sastry, S. Seshia, and A. Dragan. “Information Gathering Actions over Human Internal State”. *IROS*. 2016.
- [23] 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.
- [24] A. Bajcsy, S. Bansal, E. Ratner, et al. “A Robust Control Framework for Human Motion Prediction”. *IEEE Robotics and Automation Letters* (2020).
- [25] 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.
- [26] K. Margellos and J. Lygeros. “Hamilton–Jacobi formulation for reach–avoid differential games”. *IEEE Transactions on Automatic Control* 56.8 (2011).
- [27] J. F. Fisac, M. Chen, C. J. Tomlin, and S. S. Sastry. “Reach-avoid problems with time-varying dynamics, targets and constraints”. *Proceedings of the 18th international conference on hybrid systems: computation and control*. 2015.
- [28] J. Hardy and M. Campbell. “Contingency planning over probabilistic obstacle predictions for autonomous road vehicles”. *IEEE Transactions on Robotics* 29.4 (2013).
- [29] C. L. Baker, J. B. Tenenbaum, and R. R. Saxe. “Goal inference as inverse planning”. *Annual Meeting of the Cognitive Science Society*. Vol. 29. 2007.
- [30] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey. “Maximum entropy inverse reinforcement learning”. *AAAI*. 2008.
- [31] A. Bajcsy, S. L. Herbert, D. Fridovich-Keil, et al. “A Scalable Framework For Real-Time Multi-Robot, Multi-Human Collision Avoidance”. *IEEE International Conference on Robotics and Automation* (2018).
- [32] A. Robotics. “AWS RoboMaker”. <https://github.com/aws-robotics/> (2020).
- [33] A. Dragan and S. Srinivasa. “Generating legible motion”. *Robotics: Science and Systems*. 2013.
- [34] S. Murata, W. Masuda, J. Chen, et al. “Achieving Human-Robot Collaboration with Dynamic Goal Inference by Gradient Descent”. *International Conference on Neural Information Processing*. Springer. 2019.
- [35] T. Yu, C. Finn, A. Xie, et al. “One-shot imitation from observing humans via domain-adaptive meta-learning”. *arXiv preprint arXiv:1802.01557* (2018).