

Anticipating Human Collision Avoidance Behavior for Safe Robot Reaction

Kelsey Hawkins¹

Panagiotis Tsiotras²

Abstract—For robots to effectively navigate in the presence of humans, they must safely leverage the human’s perceived unwillingness to collide. Drawing on Viability Theory, we propose a novel approach to robustly anticipate human collision-avoiding behavior. We assume that rational humans try to optimally control their motion to avoid collision, but they are also prone to error, which makes their behavior suboptimal. We offer a robust control model which varies the level of optimality expected over time, assuming that humans may act unpredictably for a brief period of time, but their actions approach optimal collision-avoiding behavior as time progresses. We show how the proposed model can be used to produce a set of initial states for which a rational human will avoid collision. Further, we produce a robust policy which characterizes the set of control inputs expected by the human at any state. We illustrate our approach using two representative scenarios.

I. INTRODUCTION

For high-inertia robots planning in the presence of humans, making assumptions about human behavior is essential to ensure avoidance of collisions. For example, if an autonomous vehicle operates in the presence of a human-driven vehicle, both the autonomous vehicle and the human-driven vehicle can contribute to the success or failure of the collision avoidance task. To better understand this point of view, suppose a robot pulls out of a driveway onto a road in front of a human driver coming its way and at some point a collision occurs. If near perfect reaction time and hard braking on behalf of the human would have been required to avoid collision, one would put nearly all the blame on the robot. On the other hand, if the human was originally far away and moving slowly, and thus was given ample time to react and brake, we would consider the robot to have acted safely, even if the human were slightly inconvenienced by having to reduce his/her speed.

In [1] we called a human-robot system *rationaly safe* if the interaction is always safe under the assumption that the human acts according to a rational model of behavior. While it is intuitive to conjecture that the assumptions in this example constitute rational collision-avoiding behavior, the specifics of defining rational behavior are not immediately clear. In this paper, we propose a general technique for modeling rational human collision avoidance behavior. We make these intuitive statements concrete and develop a framework which reduces the complexity of characterizing rational human behavior.

¹PhD Candidate, Institute for Robotics and Intelligent Machines, Georgia Institute of Technology, Atlanta, GA 30332-0250, USA. Email: kphawkins@gatech.edu

²Professor, Institute for Robotics and Intelligent Machines, School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0150, USA. Email: tsiotras@gatech.edu

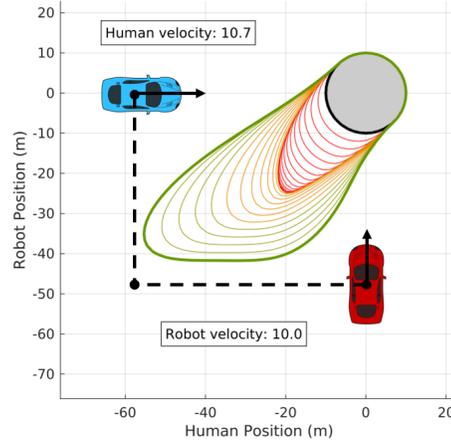


Fig. 1: Computed suboptimal safety set $\text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$ for the intersection problem, the states outside of the thick green line.

We begin by discussing related work in the next section, noting that this paper unites human factors research in driving behavior with viability theory approaches for safe control. We follow by describing a general approach for modeling suboptimal collision avoidance. Next, we show how the parameterization of this model forms a partial order so that the relative conservativeness of the parameterizations can be determined a priori. We then show how this general model can be applied to human collision avoidance. Finally, we discuss how this model can be applied to robot decision making.

II. RELATED WORK

Early work on modeling human collision avoidance behavior comes primarily from the human factors community, looking for ways to improve vehicle safety in tasks such as braking for obstacles [2], following moving vehicles [3], and avoiding vehicles while crossing intersections [4]. For example, Lee’s collision avoidance model in [2] consists of: a response period, where the human identifies a threat and moves to begin braking or steering; an adjustment period, where the human is actively pressing the brake or turning the wheel; and a critical period, where the brake needs to be at full power to avoid collision. Further, he suggests that humans use a model of time-to-collision (TTC) to assess risk. TTC is a measure of how long it would take for the vehicle to collide, given a constant closing velocity.

This paper models humans similarly to [2], acknowledging that perceptual, motor, and cognitive limitations may cause humans to react suboptimally during collision avoidance tasks. Further, we use a similar, but more general, version of

the TTC to model the risk of a particular state. However, we also propose a methodology to apply these concepts more generally to generic models of motion and collision.

Our approach is closely related to work from the reachability community [5]–[7]. For example, Mitchell et. al. [5] computed backwards reachable sets for continuous dynamic games, where one agent is trying to optimally avoid collision and the other one is trying to optimally cause a collision. Althoff et. al. [6] used probabilistic reach sets to both plan autonomous driving maneuvers and model the behaviors of nearby vehicles, while Kaynama et. al. [7] proposed a two-stage hybrid controller for safely avoiding collisions. When the state is near the boundary of states considered safe through reachability analysis, the control is switched from an unconstrained control law to a safety preserving controller which keeps the state safe. The safe set is selected such that safe control can be guaranteed to exist under their policy. Verma and Del Vecchio modelled humans in an intersection braking scenario using a hidden mode hybrid control system [8], [9]. Their idea is that, given a hybrid control model of human behavior, the robot can infer the human’s mode by observing their dynamics and by ruling out modes by comparing observed dynamics to an internal model.

The present paper unites reachability approaches to controller design with human factors models of realistic human control by proposing a robust human control model. We focus specifically on how one should characterize rational collision-avoiding behavior, whereas many of earlier techniques either assume it to be specified beforehand [8], they take an overly conservative approach [5] [7], or use a probabilistic approach which does not offer hard safety guarantees [6].

III. MODELING SUBOPTIMAL COLLISION AVOIDANCE

A. Preliminaries

We begin by introducing the dynamics and collision model for the human. We model the agent’s dynamics as $\dot{x} = f(x) + g(x)w$ with state $x \in \mathbb{R}^n$ and control input $w \in W \subseteq \mathbb{R}^m$, where W is a multidimensional interval. That is, for $w^+, w^- \in \mathbb{R}^m$, $W := [w_1^-, w_1^+] \times \dots \times [w_m^-, w_m^+]$. The agent is considered to be in collision with an obstacle if the state leaves a given closed set \mathcal{F} . The mapping from an initial state x_0 and a control signal $w(\cdot)$ to state trajectories, will be denoted by $(x_0, w(\cdot)) \mapsto x^{x_0, w(\cdot)}(\cdot)$. We assume that the model is well behaved and that f and g are Lipschitz continuous. Given a set $\mathcal{K} \subseteq \mathbb{R}^n$, we will represent this set using its level set, denoted by $\varphi^{\mathcal{K}}(x)$, with the properties $\varphi^{\mathcal{K}}(x) < 0$ on its interior, $\varphi^{\mathcal{K}}(x) > 0$ on its complement, and $\varphi^{\mathcal{K}}(x) = 0$ on the boundary of \mathcal{K} .

We call a set-valued function $U(t, x) : [0, T] \times \mathcal{F} \rightarrow 2^W$ an agent’s *control model*. The focus of this section is to produce, from a particular implicit description of a control model, both an explicit control model description and what we call a *safety set*, that is, a set where all initial conditions $x_0 \in \mathcal{F}$ have the property

$$w(\cdot) \in U(\cdot, x^{x_0, w(\cdot)}(\cdot)) \Rightarrow \forall s \in [0, T], x^{x_0, w(\cdot)}(s) \in \mathcal{F} \quad (1)$$

and thus labeled safe.

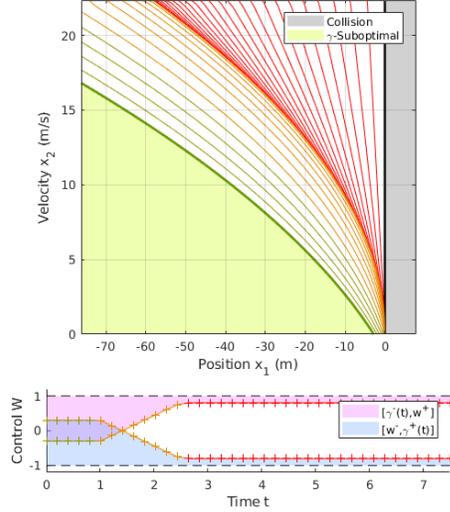


Fig. 2: Top: Computed suboptimal safety set $\text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$ for the braking scenario. The contour curves represent constant $\theta(x) = t$ (or $\varphi(t, x) = 0$) for times 0.2s apart. Bottom: The suboptimal control parameterization. Lines are $\gamma^-(\cdot), \gamma^+(\cdot)$, and the plus markers correspond to contour lines in the top figure. Colors delineate the three modes defined in Section V.

B. Viability and Invariance Kernels

We state some classic connections to Viability Theory, which motivate our approach. Suppose our control model assumes the agent, when necessary to avoid collision, will act optimally to do so, but will otherwise choose any control. The duration- T viability kernel $\text{Viab}_T := \{x_0 : \exists w(\cdot), \forall s \in [0, T], x^{x_0, w(\cdot)}(s) \in \mathcal{F}\}$ [10] would be a safety set for this implicit control model. To the other extreme, suppose our control model makes no assumptions about the agent’s behavior, i.e., $U(t, x) \equiv W$. The duration- T invariance kernel $\text{Inv}_T := \{x_0 : \forall w(\cdot), \forall s \in [0, T], x^{x_0, w(\cdot)}(s) \in \mathcal{F}\}$ [10] would be a safety set for this control model.

For many systems, Inv_T is too small to be practical and Viab_T is not nearly conservative enough for imperfect agents. For example, for many $x_0 \in \partial \text{Viab}_T$, perfectly optimal control for some duration is required to avoid collision. For rational yet imperfect agents like humans, it would be more useful to use a control model that varies over time to blend between these two extremes. That is, we would like a control model that makes few assumptions about the agent’s behavior for an initial period of time, but assumes that eventually the agent will act near optimally to avoid collision.

C. γ -Suboptimal Control Model

We provide an implicit description for our proposed control model we call a γ -suboptimal control model, which is used to generate what we call a γ -suboptimal safety set. The control model is developed recursively, backwards in time, by considering how the worst-case trajectories $x^{x_0, w^*(\cdot)}(\cdot)$ would behave. At time T , the worst-case terminal state would be $x^{x_0, w^*(\cdot)}(T) \in \partial \mathcal{F}$. Likewise, the worst-case initial state would be on the boundary of the γ -suboptimal safety set.

Now, consider that a worst-case trajectory is at the state $x_t = x^{x_0, w^*(\cdot)}(t)$ at time $t \in (0, T)$. Suppose, further, that we have a vector $\lambda_{t, x_t} \in \mathbb{R}^n$ which indicates which direction drives the current state towards collision the fastest. Projecting the system dynamics onto this vector, $\langle \lambda_{t, x_t}, \dot{x}_t \rangle = \langle \lambda_{t, x_t}, f(x_t) \rangle + \sum_{i=1}^m \langle \lambda_{t, x_t}, g_i(x_t) \rangle w_i$, we see that the sign of $\langle \lambda_{t, x_t}, g_i(x_t) \rangle$ determines whether larger w_i or smaller w_i is safer, independently for each control dimension. We introduce a pair of worst-case controls $\gamma_i^+(t), \gamma_i^-(t) \in [w_i^-, w_i^+]$ for each control dimension for when the scalar is positive or negative, respectively. We assume for the control model that the agent will execute any control as safe, or safer than the worst-case control. Thus, the implicit γ -suboptimal control model can be expressed as

$$W^*(t, x_t) := \bigtimes_{i=1}^m \begin{cases} [\gamma_i^-(t), w_i^+], & \langle \lambda_{t, x_t}, g_i(x_t) \rangle < 0, \\ [w_i^-, \gamma_i^+(t)], & \langle \lambda_{t, x_t}, g_i(x_t) \rangle > 0. \end{cases}$$

When the sign is zero, modelling is irrelevant because system safety is uncontrollable with respect to that dimension.

For a concrete example, suppose we have a vehicle that can steer left ($w_i < 0$) or right ($w_i > 0$) with control input w_i . At time t and state x_t , steering left will push the state towards collision faster, indicated as $\langle \lambda_{t, x_t}, g_i(x_t) \rangle < 0$. Thus, we look to $\gamma_i^-(t)$ as the worst-case steering, assuming the agent will execute a control $w_i \in [\gamma_i^-(t), w^+]$. Setting $\gamma_i^-(t) = 0$, the model assumes the agent will steer right or, at worst, keep straight. Setting $\gamma_i^-(t) = w^+ - \epsilon$ for small $\epsilon > 0$ assumes the agent will steer right nearly as hard as possible. Setting $\gamma_i^-(t) = w^-$ makes no assumptions about the agent's control, including the possibility that they will steer aggressively towards collision.

D. Hamiltonian Characterization

We now discuss how λ_{t, x_t} is determined. It is known that we can characterize the level sets of the duration- T viability and invariance kernels as solutions of Hamilton Jacobi Bellman (HJB) equations, namely,

$$-\frac{\partial \varphi}{\partial t}(t, x) - \max\{0, H(t, x, \frac{\partial \varphi}{\partial x}(t, x))\} = 0, \quad (2)$$

$$\varphi(T, x) = \varphi^{\mathcal{F}}(x), \quad (3)$$

where $\varphi^{\text{Viab}_T}(x) = \varphi^{\text{Inv}_T}(x) = \varphi(0, x)$ with the Hamiltonian H set as

$$H^{\text{Viab}}(t, x, \lambda) = \langle \lambda, f \rangle + \sum_i \min_{w_i \in [w_i^-, w_i^+]} \langle \lambda, g_i(x) \rangle w_i, \quad (4)$$

$$H^{\text{Inv}}(t, x, \lambda) = \langle \lambda, f \rangle + \sum_i \max_{w_i \in [w_i^-, w_i^+]} \langle \lambda, g_i(x) \rangle w_i, \quad (5)$$

respectively [11]. Consider $\lambda_{t, x_t} = \partial \varphi / \partial x(t, x_t)$ to be the direction which drives the state towards collision the fastest. In light of the analysis from the previous section, note that the viability kernel always selects the optimal control for avoiding collision, while the invariance kernel selects the optimal control for causing collision.

We use a similar construction to obtain the γ -suboptimal safety set, using the same HJB equations but with an alternative Hamiltonian H^{Sub} . This Hamiltonian selects the

worst-case control in the model described in the previous section,

$$H^{\text{Sub}}(t, x, \lambda) = \langle \lambda, f \rangle + \sum_i \max_{w_i \in W_i^*(t, x, \lambda)} \langle \lambda, g_i(x) \rangle w_i. \quad (6)$$

The max operator simplifies to

$$\max_{w_i \in W_i^*(t, x, \lambda)} \langle \lambda, g_i(x) \rangle w_i = \begin{cases} \langle \lambda, g_i(x) \rangle \gamma_i^-, & \langle \lambda, g_i(x) \rangle \leq 0, \\ \langle \lambda, g_i(x) \rangle \gamma_i^+, & \langle \lambda, g_i(x) \rangle > 0, \end{cases}$$

selecting the relevant worst-case control, as expected. For the solution φ to these equations, we set $\varphi^{\text{Sub}_T}(x) = \varphi(0, x)$, the γ -suboptimal safety set $\text{Sub}_T^{\gamma(\cdot), \gamma^+(\cdot)}$.

We choose $\gamma_i^+ : [0, T] \rightarrow [w_i^-, w_i^+]$ to be non-increasing and $\gamma_i^- : [0, T] \rightarrow [w_i^-, w_i^+]$ to be non-decreasing. By doing so, we can anticipate that the worst-case inputs will get closer to optimal as time progresses. Continuing the steering example, for early t_1 , $\gamma_i^-(t_1)$ might be chosen to be negative, suggesting that even steering left towards collision is permissible at first. But for $t_2 > t_1$, $\gamma_i^-(t_2)$ should get closer to the optimal w^+ , suggesting that the agent is expected to begin an evasive action by steering hard to the right. Of course, this assumes that the sign of $\langle \lambda, g_i(x) \rangle$ remains the same, which does not always happen. However, if the sign does flip, the analysis would continue with γ_i^+ once that occurs.

E. Relationship to Differential Games

We should recognize that the derivation of a γ -suboptimal control can be formulated as a differential game. That is, if we consider two-player controlled state dynamics $\dot{\tilde{x}} = f(\tilde{x}) + g(\tilde{x})(u + v)$ with time-varying control constraints $U(\cdot), V(\cdot)$ defined as

$$(U_i(t), V_i(t)) := \begin{cases} (\{0\}, [\gamma_i^-(t), \gamma_i^+(t)]), & \gamma_i^-(t) \leq \gamma_i^+(t), \\ ([\gamma_i^+(t), \gamma_i^-(t)], \{0\}), & \gamma_i^-(t) > \gamma_i^+(t), \end{cases}$$

then the solution of the zero-sum differential game

$$\varphi_{\text{DG}}(t, \tilde{x}_t) = \max_{v(\cdot) \in V(\cdot)} \min_{u(\cdot) \in U(\cdot)} \max_{s \in [t, T]} \varphi^{\mathcal{F}}(\tilde{x}_{\tilde{x}_t, u(\cdot), v(\cdot)}(s))$$

satisfies the γ -suboptimal HJB equations. We omit the proof for the sake of brevity.

F. Closed Loop Control Model

If trajectories follow the worst-case profile used in the construction, we expect that the state will be in $\partial \mathcal{F}$ at time T . However, in practice, trajectories will not follow this profile. Nor is the agent guaranteed to follow the control model precisely. Thus, we need a closed-loop control model which removes the time parameter used in the control model.

First, note that the level sets $\varphi^{\text{Sub}_T(t)}(x) = \varphi(t, x)$ represent a time-varying set $\text{Sub}_T^{\gamma(\cdot), \gamma^+(\cdot)}(t)$ which labels states as safe/unsafe as model time progresses. For any given state $x \in \mathcal{F}$, define the worst-case time elapsed function as

$$\theta(x) := \sup\{t \in [0, T] : \varphi(t, x) \leq 0\}.$$

The value of $\theta(x)$ indicates, that among the times of worst-case trajectories that would have passed through this state, which time would have been the latest. Higher values are a rough way of determining nearness to collision.

Choose the closed loop explicit control model to be

$$W_{\text{CL}}^*(x) := W^*(\theta(x), x). \quad (7)$$

By using $\theta(x)$ to determine what is expected from the agent, we have put the responsibility on the agent to be more aggressive in controlling to avoid safety when it exceeds expectations earlier in its behavior.

IV. PARAMETERIZATION OF γ BOUNDS

A. Conservative Partial Order

When designing control models, it is of interest to compare the fitness of using a particular setting of the γ -suboptimal control model parameters $(\gamma^+(\cdot), \gamma^-(\cdot))$ to an alternate setting $(\tilde{\gamma}^+(\cdot), \tilde{\gamma}^-(\cdot))$. If the model is too conservative, like the invariance kernel's model, it might be less useful. On the other hand, if the model is not conservative enough to capture real-world rational behavior, it will over estimate the safety of some states. We define a "conservative" partial order over pairs of γ parameters to attempt to capture this notion. We say $(\tilde{\gamma}^-, \tilde{\gamma}^+)$ is more conservative than (γ^-, γ^+) , denoted $(\gamma^-, \gamma^+) \leq_{\Gamma} (\tilde{\gamma}^-, \tilde{\gamma}^+)$, if

$$\forall t \in [0, T], \quad \tilde{\gamma}^-(t) \leq \gamma^-(t), \quad \gamma^+(t) \leq \tilde{\gamma}^+(t).$$

An important consequence of this partial order is that the set of safe initial states shrinks for more conservative γ pairs.

Proposition 1: The following holds

$$(\gamma^-, \gamma^+)_T \leq_{\Gamma} (\tilde{\gamma}^-, \tilde{\gamma}^+)_T \Rightarrow \text{Sub}_T^{\tilde{\gamma}^-(\cdot), \tilde{\gamma}^+(\cdot)} \subseteq \text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$$

Proof: Let $\tilde{x} \in \text{Sub}_T^{\tilde{\gamma}^-(\cdot), \tilde{\gamma}^+(\cdot)}$, and let φ and $\tilde{\varphi}$ be bounded continuous viscosity solutions to the HJB equations (2), (3), and (6), where H^{Sub} and \tilde{H}^{Sub} are the Hamiltonians, assumed continuous and well behaved¹, defined with the parameters $(\gamma^-, \gamma^+)_T$ and $(\tilde{\gamma}^-, \tilde{\gamma}^+)_T$, respectively. If $\langle \lambda, g_i(x) \rangle \leq 0$,

$$\begin{aligned} \max_{w_i \in W_i^*(t, x, \lambda)} \langle \lambda, g_i(x) \rangle w_i &= \langle \lambda, g_i(x) \rangle \gamma_i^- \leq \langle \lambda, g_i(x) \rangle \tilde{\gamma}_i^- \\ &= \max_{w_i \in \tilde{W}_i^*(t, x, \lambda)} \langle \lambda, g_i(x) \rangle w_i \end{aligned}$$

The inequality also holds when $\langle \lambda, g_i(x) \rangle > 0$. Thus, it follows that for all $x, \lambda \in \mathbb{R}^n$, $H^{\text{Sub}}(t, x, \lambda) \leq \tilde{H}^{\text{Sub}}(t, x, \lambda)$. Using this inequality, basic viscosity theory, and Theorem 6.2 in [12], yields that $\varphi \leq \tilde{\varphi}$ on $[0, T] \times \mathbb{R}^n$. Also, by definition, we have that $\tilde{\varphi}(0, \tilde{x}) \leq 0$. Thus, we conclude $\varphi(0, \tilde{x}) \leq 0$ so $\tilde{x} \in \text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$. ■

The comparison between the control models for each pair is more difficult to analyze. Since $\varphi \leq \tilde{\varphi}$, it is easily shown that $\tilde{\theta}(x) \leq \theta(x)$. For the states where the signs of $\langle \lambda_{\theta(x), x}, g_i(x) \rangle$ and $\langle \lambda_{\tilde{\theta}(x), x}, g_i(x) \rangle$ agree, we have that $W_i^*(\theta(x), x) \subseteq \tilde{W}_i^*(\tilde{\theta}(x), x)$.

If we choose (γ^+, γ^-) to be very conservative for early t , the model can tolerate brief periods where the model's expectations about the safe control direction are incongruous with the agent's actions. However, it is generally the case that when collision is imminent, the ambiguity disappears entirely.

¹It satisfies the necessary regularity properties of Theorem 6.2 in [12].

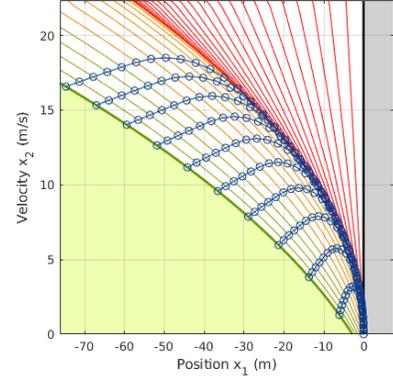


Fig. 3: Simulating worst-case trajectories for the braking scenario by sampling $x_0 \in \partial \text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$ and executing worst-case control trajectories. Since decelerating is always the optimal control direction for $x_1 \leq 0$, $w^*(\cdot) = \gamma^+(\cdot)$. In these cases, the human just barely touches the boundary of the collision set.

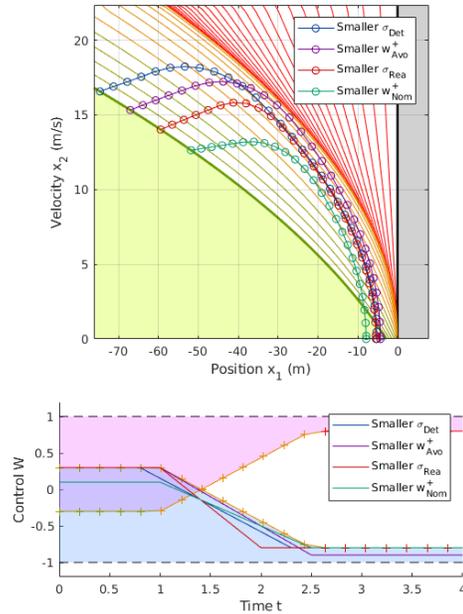


Fig. 4: If the human executes trajectories slightly more optimal than the worst-case $\gamma^+(\cdot)$, the human will avoid collision with a significant margin. These examples were generated in the same way as $\gamma^+(\cdot)$, but with its parameters $(\sigma_{\text{det}}, w_{\text{avo}}^+, \sigma_{\text{rea}}, w_{\text{nom}}^+)$ independently reduced.

V. MODELING HUMAN AVOIDANCE BEHAVIOR

In this section, we discuss how we can simplify the parameterization of (γ^+, γ^-) by selecting them as continuous piecewise linear functions. We justify the parameter selection using analysis from the human factors literature.

As in [2], we assume that collision avoidance behavior is broadly composed of a few discrete modes: nominal control, reaction, and avoidance. The nominal control is the default behavior of the human before an obstacle is detected. During this period, we make the most conservative assumptions about how the human is acting because we assume that the human is not yet aware of the obstacle and thus (s)he is not attempting to avoid it. Let $[w_{\text{nom}}^-, w_{\text{nom}}^+] \subseteq W$ be the nominal

control set².

At some point, the human will detect the obstacle. Note that there will be a significant delay between the appearance of the obstacle and when the human begins reacting to it. We should assume nominal control up until the human begins reacting. Let $\sigma_{\text{det}} \geq 0$ be the worst-case detection delay. Reaction is the behavior period during which the human switches between nominal control and avoidance control. Let $\sigma_{\text{rea}} \in [0, \infty)$ be the worst-case switching time.

Finally, the avoidance behavior is the period during which the human acts strongly to avoid collision, following the reaction period. It should be recognized that for humans, perfectly optimal control might not be possible, even if they have had plenty of time to react. For example, humans might not brake as hard as possible and they might momentarily let up on the brake. We pick a threshold for avoidance that we believe the human can persistently maintain avoidance control. Let $[w_{\text{avo}}^-, w^+] \subseteq W$ be the suboptimal control set when positive controls are better, and $[w^-, w_{\text{avo}}^+] \subseteq W$ be the suboptimal control set for when negative controls are better.

We assume that once a collision avoidance maneuver begins, it does not stop until the state of the system returns to safety. Thus, we simply assume that the human executes an avoidance maneuver for a “long time”, $\sigma_{\text{avo}} \gg 0$. For appropriately chosen parameters $w_{\text{avo}}^+, w_{\text{avo}}^-, \sigma_{\text{avo}}$, the HJB solution converges to a stationary level set in the region of interest of the state space. The thick red lines where the level set curves accumulate in Figures 2 and 1 demonstrate this effect. In fact, this set is like a viability kernel (called a controlled invariant set in literature [13]) such that, as long as an avoidance control of $w_{\text{avo}}^+, w_{\text{avo}}^-$ or better is always used upon hitting the boundary, the trajectory is guaranteed to stay within that set for infinite time.

Finally, we define the γ^+ as continuous piecewise linear spline with time knots $(0, \sigma_{\text{det}}, \sigma_{\text{det}} + \sigma_{\text{rea}}, T)$ corresponding to function values $(w_{\text{nom}}^+, w_{\text{nom}}^+, w_{\text{avo}}^+, w_{\text{avo}}^+)$ where $T = \sigma_{\text{det}} + \sigma_{\text{rea}} + \sigma_{\text{avo}}$. We define γ^- similarly. The bottom of Figure 2 shows the γ curves generated by this construction. In all figures, each color signifies one of the three model periods.

It is easy to see that increasing $w_{\text{nom}}^+, w_{\text{avo}}^+, \sigma_{\text{det}}, \sigma_{\text{rea}}$ and decreasing $w_{\text{nom}}^-, w_{\text{avo}}^-$ will produce more conservative models. This is a useful result because it means that we can easily hand-tune the models to improve the conservativeness of the control model. Figure 4 demonstrates how the chosen model is robust to variations in these parameters.

VI. NUMERICAL SIMULATIONS

A. Braking Scenario

We motivate our approach by investigating two collision-avoidance scenarios. The first is a braking scenario, where the human must brake for a static obstacle. The dynamics are assumed to be those of a double integrator $\dot{x} = [x_2, k_a w]^T$ with $w \in [-1, 1]$, indicating that the human can accelerate or brake. We focus on states where the human is moving with positive velocity $x_2 \geq 0$ and has not yet collided with the obstacle $\mathcal{F} = \{x : x_1 \leq 0\}$.

²Without loss of generality, we assume $m = 1$ for this section.

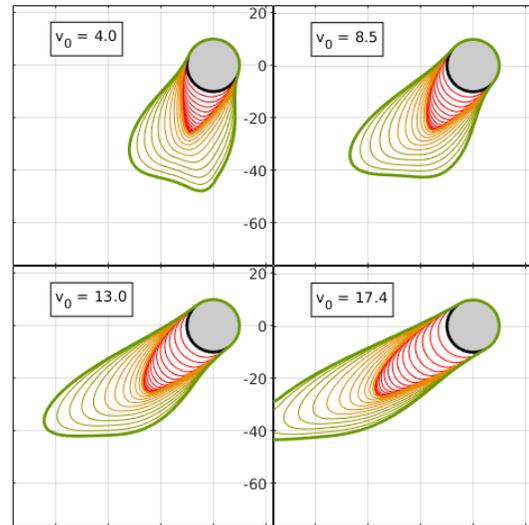


Fig. 5: Slices of the intersection scenario state space for various human velocities.

Figure 2 illustrates how the parameterization of the previous section produces a suboptimal safe set.

Recall that the HJB equations are solved backwards in time from the collision set, so the level set curves are propagated left during computation, but represent time points starting with $t = T$ decreasing to $t = 0$. Note how thick red curve is formed from the convergence of the level sets during the avoid phase. However, as soon as γ switches into reaction mode, it propagates away again.

Figure 3 demonstrates how worst-case trajectories behave with respect to the value function φ . An interesting result to note here is that all of these trajectories execute the same control signal, regardless of the initial position. This shows how the γ parameterization is performed in the control space, but the safety guarantees generalize to other states. It is worth emphasizing that worst-case behavior should not be expected of rational agents, even in exceptional cases. This behavior should represent the firm boundary beyond which there should be no doubt that the agent’s behavior was irresponsible or irrational. In Figure 4 we see how trajectories behave under more realistic scenarios, where the human acts more optimally than the worst-case. Though it is not shown, a control trajectory which briefly exceeds the control model might still avoid collision. For example, if the state were to pass over the thick red line, a braking control of $w = w^- = -1$ might still avoid collision. However, this model does not provide guarantees for those cases.

The level sets in all figures are computed and visualized using the Level Set Toolbox [14].

B. Intersection Scenario

The second scenario involves two cars at an intersection, where the robot vehicle has the right-of-way. The rational human assumes the robot vehicle is moving with constant velocity and will not react to the human’s actions. The human has the same double integrator dynamics as before, but (s)he must decide whether to pass first, or let the robot vehicle pass first. The dynamics are $\dot{x} = [x_2, k_a w, k_v]^T$, where x_1, x_2 are identical to the braking scenario, but x_3 is the position of

the robot vehicle.

Figure 1 illustrates how points in the state space correspond to configurations of the two vehicles, as well as demonstrating the suboptimal safety set. Since the state space is 3-dimensional, we only illustrate 2D slices for given human velocities. We can get a sense how the safety set changes with the velocity of the human driven vehicle by looking at Figure 5. Figure 6 shows worst-case trajectories for various initial configurations of the two cars. If the human has already gotten close to the intersection, (s)he decides to accelerate through, otherwise, (s)he decides to come to a stop and let the robot pass first. Note that some of the trajectories which ended up coming to a stop may have also ended up accelerating through, depending on the human’s behavior near the beginning. However, as the trajectory approaches the red states, regardless of previous behavior, the human should commit to one or the other.

VII. INFORMING ROBOT BEHAVIOR

A robot can use the γ -suboptimal control model of rational human behavior to inform its own decisions about how to act. Take the example from the introduction, where a robot-driven vehicle is pulling out in front of a human-driven vehicle coming towards it. The robot might encounter an unexpected obstacle in the middle of this maneuver forcing it to stop, and has now become a static obstacle for the human driver. Thus, the robot should consider whether a rational human would be able to stop in time. The robot can make its decision by simply checking if the current state of the human’s vehicle is in the γ -suboptimal safe set, that is $x_0 \in \text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$. If not, the robot should choose not to pull out, thereby removing the possibility of an obstacle altogether.

For the intersection scenario, even though the robot vehicle has the right of way, it can choose to deviate from its constant velocity to either brake or accelerate. If the state of the system is in $\text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$ when the robot detects the human, the robot should not alter its course. Otherwise, it should consider evasive measures. If a robot decides to brake hard for unexpected human drivers at intersections every time they occur, it may cause more accidental rear-ends, when most people would have realized that the driver ahead would have easily cleared the intersection by the time they had reached it.

Of course, even if the initial state begins in the safety set, it may leave the safety set at some point. As long as the human behavior falls within the closed-loop set policy model, safety should be ensured for up to a duration T from the start of the interaction. An observer can be implemented to ensure that the human driver control signal stays consistently within its bounds. If the human’s control signal violates the policy model for some time period, the robot can begin an evasive maneuver.

VIII. CONCLUSION

We have presented a novel model for anticipating human collision-avoidance maneuvers. By varying the level of sub-optimality over time, we are able to conservatively anticipate such maneuvers, while still accounting for various sources of human error. We have demonstrated how the proposed

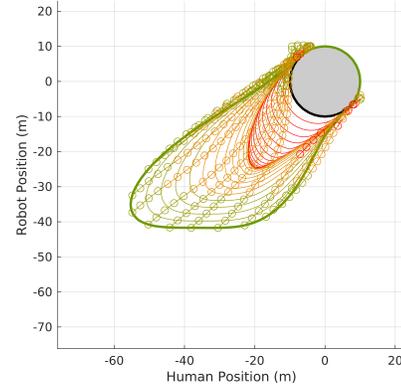


Fig. 6: Simulating worst-case trajectories for the intersection scenario by sampling $x_0 \in \text{Sub}_T^{\gamma^-(\cdot), \gamma^+(\cdot)}$ and executing worst-case closed loop control trajectories.

control model can be tuned by objectively comparing the conservativeness of relative parameterizations. We have applied this model to two simulation scenarios, and discussed how the results can be used to make safe robot decisions.

Acknowledgment: We would like to thank Aaron Bobick, Henrik Christensen, and Magnus Egerstedt for their motivation of this work and discussions about rationally safe behavior. Partial support for this work has been provided by NSF award CMMI-1662523.

REFERENCES

- [1] K. P. Hawkins, M. Egerstedt, and A. F. Bobick, “Towards Rationally Safe Robots,” in *RSS Workshop Human-Robot Collaboration for Industrial Manufacturing*, 2014.
- [2] D. N. Lee, “A theory of visual control of braking based on information about time to collision,” *Perception*, vol. 5, no. 4, pp. 437–459, 1976.
- [3] M. Brackstone and M. McDonald, “Car-following : a historical review,” *Transportation Research Part F*, vol. 2, no. 1999, pp. 181–196, 2000.
- [4] J. W. Muttart, “Development and evaluation of driver response time predictors based upon meta analysis,” SAE Technical Paper, Tech. Rep., 2003.
- [5] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, “A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games,” *IEEE Transactions on Automatic Control*, vol. 50, no. 7, pp. 947–957, 2005.
- [6] M. Althoff, O. Stursberg, and M. Buss, “Model-Based Probabilistic Collision Detection in Autonomous Driving,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, pp. 299–310, 2009.
- [7] S. Kaynama, I. M. Mitchell, M. Oishi, and G. A. Dumont, “Scalable safety-preserving robust control synthesis for continuous-time linear systems,” *IEEE Transactions on Automatic Control*, vol. 60, no. 11, pp. 3065–3070, 2015.
- [8] R. Verma and D. Vecchio, “Semiautonomous Multivehicle Safety,” *IEEE Robotics & Automation Magazine*, pp. 44–54, 2011.
- [9] R. Verma and D. Del Vecchio, “Safety control of hidden mode hybrid systems,” *IEEE Transactions on Automatic Control*, vol. 57, no. 1, pp. 62–77, 2012.
- [10] S. Kaynama, M. Oishi, I. M. Mitchell, and G. A. Dumont, “The continual reachability set and its computation using maximal reachability techniques,” in *Proceedings of the IEEE Conference on Decision and Control*, 2011, pp. 6110–6115.
- [11] J. Lygeros, “On reachability and minimum cost optimal control,” *Automatica*, vol. 40, no. 6, pp. 917–927, 2004.
- [12] J. Calder, “Some notes on viscosity solutions of Hamilton-Jacobi equations,” 2018. [Online]. Available: math.umn.edu/~jw-calder/viscosity_solutions.pdf
- [13] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, 2008.
- [14] I. M. Mitchell, “A toolbox of level set methods,” *Department of Computer Science, University of British Columbia, Vancouver, BC, Canada, Tech. Rep. TR-2004-09, July*, pp. 177–247, 2004.