

Parameter-Conditioned Reachable Sets for Updating Safety Assurances Online

Javier Borquez¹, Kensuke Nakamura², Somil Bansal¹

Abstract—Hamilton-Jacobi (HJ) reachability analysis is a powerful tool for analyzing the safety of autonomous systems. However, the provided safety assurances are often predicated on the assumption that once deployed, the system or its environment does not evolve. Online, however, an autonomous system might experience changes in system dynamics, control authority, external disturbances, and/or the surrounding environment, requiring updated safety assurances. Rather than restarting the safety analysis from scratch, which can be time-consuming and often intractable to perform online, we propose to compute *parameter-conditioned* reachable sets. Assuming expected system and environment changes can be parameterized, we treat these parameters as virtual states in the system and leverage recent advances in high-dimensional reachability analysis to solve the corresponding reachability problem offline. This results in a family of reachable sets that is parameterized by the environment and system factors. Online, as these factors change, the system can simply query the corresponding safety function from this family to ensure system safety, enabling a real-time update of the safety assurances. Through various simulation studies, we demonstrate the capability of our approach in maintaining system safety despite the system and environment evolution.

I. INTRODUCTION

Ensuring the safe operation of autonomous systems is crucial for their successful deployment in safety-critical domains such as self-driving vehicles, unmanned aerial vehicle mobility, and human-robot interaction. These applications often require autonomous systems to operate in situations where environmental factors might change online. For example, a UAV might experience stronger wind during its flight than it anticipated, or the off-shore landing pad of a rocket might have drifted away from its original position. In such situations, it is hard to provide safety guarantees prior to the system deployment; instead, the system needs to perform an efficient evaluation and adaptation of safety assurances online in light of the environment and system evolution. Addressing this challenge is the core focus of this work.

Providing safety assurances for dynamical systems has been studied using Model Predictive Control (MPC) methods [1]–[3], Lyapunov analysis [4]–[7], Sum-of-Squares (SOS) programming [8], [9], and reachability analysis [10]–[12]. However, safety assurances are typically provided for *given*

environment conditions and system dynamics. Safe motion planning methods [13]–[18] combine the above safety assurance methods with online trajectory planning to ensure safety in *a priori* unknown environments. However, these methods typically impose restrictive assumptions on the system dynamics or the environment to ensure safety. Furthermore, they often do not consider changes in system dynamics, such as changes in control authority or disturbance bounds, and require a motion planning algorithm that can operate in real-time, which itself is challenging to obtain for nonlinear systems.

Another approach for providing safety assurances for dynamical systems is via Hamilton-Jacobi (HJ) Reachability analysis [19], [20]. Its advantages include compatibility with general nonlinear system dynamics, formal treatment of bounded disturbances, and the ability to deal with state and input constraints [10]. In reachability analysis, the system safety is characterized by *Backward Reachable Tube (BRT)*. BRT is the set of states such that the system trajectories that start from this set will eventually reach the given target set despite the worst-case disturbance (or an exogenous, adversarial input more generally). If the target set consists of those states that are known to be unsafe, then the BRT contains states which are potentially unsafe and should therefore be avoided. Along with the BRT, the reachability analysis also provides a safety controller for the system to stay outside the BRT. Given the utility of reachability analysis, several methods have been proposed to update the reachability-based safety assurances online. For example, the work in [21] updates the inner and outer approximations of reachable sets for systems with changed dynamics and diminished control authority. However, the proposed method requires strong assumptions on system dynamics and cannot handle changes in the safety function, such as the target or obstacle sets. Other methods “warmstart” the reachable set computation using the previously computed solutions [11], [22]; however, these methods are still unsuitable for real-time computation and are limited to relatively low-dimensional systems.

In this work, we propose an approach for the online update of reachable sets as well as safety assurances for nonlinear dynamical systems where the environment and system changes can be parameterized. Examples of such parameters could be the vehicle-obstacle or vehicle-vehicle safety distances, the magnitude of external disturbances, system control authority, etc. Our key idea is to treat these parameters as additional “virtual” states in the system and solve

¹Authors are with the ECE department at the University of Southern California. E-mail: {javierbo, somilban}@usc.edu. The work of Javier Borquez received the support of a fellowship from BECAS Chile

²Author is with the MAE department at Princeton University. E-mail: {k.nakamura}@princeton.edu.

Project website: <https://javierborquez.github.io/ParamCondReachability/>

the reachability problem for the augmented system. This results in *parameter-conditioned reachable sets*, wherein the obtained reachable set is not only a function of the system state but also the system and environment parameters. Ultimately, at any given system state, our method provides a family of reachable sets, each member of which is determined by assigning a particular value to these parameters. The agent can use these parameter-conditioned reachable sets online to activate the safety function corresponding to the current environment and system factors, leading to a real-time adaptation of safety assurances.

Computing parameter-conditioned reachable sets requires solving a high-dimensional reachability problem which is performed offline via DeepReach – a recent framework that leverages deep learning to compute approximate reachable sets for high-dimensional systems. We demonstrate the capabilities of our framework in updating the BRT and the safety controller online in various applications.

II. PROBLEM FORMULATION

We are interested in safety assurances upon changes in the system and its environment, such as changes in control authority (e.g., steering rate of a vehicle), disturbance magnitudes (e.g., wind strength during flights), obstacle/goal sizes and positions (e.g., a movement in the landing pad for a rocket). We model the autonomous system as a dynamical system with state $x \in \mathbb{R}^n$, control u , and disturbance d . The state evolves according to the dynamics:

$$\dot{x} = f(x, u, d), \quad u \in [\underline{u}(\beta_u), \bar{u}(\beta_u)], \quad d \in [\underline{d}(\beta_d), \bar{d}(\beta_d)], \quad (1)$$

where the control and disturbance bounds are parameterized by β_u and β_d , respectively. For example, for a UAV, β_d could correspond to the wind speed experienced by the vehicle. We are interested in the settings where β_u and β_d can change online, which corresponds to a change in the control authority and disturbance strength, respectively.

Let $\xi(\tau; x_0, t_0, u(\cdot), d(\cdot))$ denote the state achieved at time τ by starting at the state x_0 at time t_0 , and applying control $u(\cdot)$ and disturbance $d(\cdot)$ over the time horizon $[t_0, \tau]$. The environment contains a target set given by $\mathcal{L}(\beta_L)$, which depends on its own parameter β_L . This set can be either a set of goal states or a set of unsafe states. As an example, for a circular obstacle, β_L might correspond to the radius of the circle.

For the system safety analysis, we are interested in computing the BRT of $\mathcal{L}(\beta_L)$ given dynamics in (1). **Backward Reachable Tube (BRT)**: the set of initial states of the system for which the agent acting optimally and under worst-case disturbances will eventually reach the target set \mathcal{L} within the time horizon $[t, T]$:

$$\mathcal{V}(t, \beta) = \{x : \forall u(\cdot), \exists d(\cdot), \exists \tau \in [t, T], \xi(\tau; x, t, u(\cdot), d(\cdot)) \in \mathcal{L}(\beta_L)\}, \quad (2)$$

where $\beta = (\beta_u, \beta_d, \beta_L)$. Note that \mathcal{V} depends on β_u and β_d through $u(\cdot)$ and $d(\cdot)$. Intuitively, if the target set consists

of those states that are known to be unsafe, then the BRT contains potentially unsafe states and should therefore be avoided. BRT can be similarly defined for goal states with the roles of control and disturbance switched. Our goal in this work is to update the BRT online for any change in the value of β .

To illustrate our approach, we will use the following running example throughout this paper:

Running example (Air3D): We consider the two-vehicle collision avoidance problem where the relative dynamics between the two vehicles (called evader and pursuer) is modeled as:

$$\begin{aligned} \dot{x}_1 &= -v_e + v_p \cos(x_3) + \omega_e x_2 \\ \dot{x}_2 &= v_p \sin(x_3) - \omega_e x_1 \\ \dot{x}_3 &= \omega_p - \omega_e, \end{aligned} \quad (3)$$

where (x_1, x_2) is the relative position in (x, y) coordinates and x_3 is the relative heading between the vehicles. v_e and v_p are the linear velocities of the evader and pursuer, respectively. These velocities are constant and identical ($0.75m/s$). ω_e and ω_p are the (bounded) angular velocities of the vehicles, and are the input and disturbance in the system, respectively.

The bounds on the pursuer angular velocity are given as $\omega_p \in [-\bar{\omega}_p, \bar{\omega}_p]$, where $\bar{\omega}_p = 3rad/s$. The bounds on the evader angular velocity, on the other hand, can change online and are given as $\omega_e \in [-4 - \beta_u, 4 + \beta_u]$. β , in this case, consists only of a single parameter β_u . Intuitively, a larger β_u corresponds to an evader that has high maneuverability, whereas a smaller β_u corresponds to a diminished control authority.

The target set is given as the set of states where the pursuer and evader are within a collision radius R :

$$\mathcal{L} = \{x : \|(x_1, x_2)\| \leq R\} \quad (4)$$

We choose $R = 0.25m$ in our example. The BRT for this collision set corresponds to all the states from which the pursuer can drive the system trajectory into the collision set within the time horizon $[t, T]$, despite the best efforts of the evader to avoid a collision. Subjected to these conditions, we are interested in updating the BRT corresponding to changes in β_u .

III. BACKGROUND

A. Hamilton-Jacobi Reachability

One way to compute BRT is through Hamilton-Jacobi (HJ) reachability analysis. In HJ reachability, the BRT computation is formulated as a zero-sum game between control and disturbance. This results in a robust optimal control problem that can be solved using the dynamic programming principle. First, a target function $l(x)$ is defined whose sub-zero level set is the target set \mathcal{L} , i.e. $\mathcal{L} = \{x : l(x) \leq 0\}$. Typically, $l(x)$ is defined as a signed distance function to \mathcal{L} . The BRT seeks to find all states that could enter \mathcal{L} at any point within

the time horizon and, therefore might be unsafe. This is computed by finding the minimum distance to \mathcal{L} over time:

$$J(x, t, u(\cdot), d(\cdot)) = \min_{\tau \in [t, T]} l(\xi(\tau; x, t, u(\cdot), d(\cdot))) \quad (5)$$

Our goal is to capture this minimum distance for optimal system trajectories. Thus, we compute the optimal control that maximizes this distance (and drives the system away from the unsafe states) and the worst-case disturbance signal that minimizes the distance. The value function corresponding to this robust optimal control problem is:

$$V(x, t) = \inf_{d(\cdot)} \sup_{u(\cdot)} \{J(x, t, u(\cdot), d(\cdot))\}. \quad (6)$$

The value function in (6) can be computed using dynamic programming, which results in the following final value Hamilton-Jacobi-Isaacs Variational Inequality (HJI-VI):

$$\min \{D_t V(x, t) + H(x, t), l(x) - V(x, t)\} = 0 \quad (7)$$

$$V(x, T) = l(x)$$

D_t and ∇ represent the time and spatial gradients of the value function. H is the Hamiltonian, which optimizes over the inner product between the spatial gradients of the value function and the dynamics to compute the optimal control and disturbance inputs:

$$H(x, t) = \max_u \min_d \langle \nabla V(x, t), f(x, u, d) \rangle. \quad (8)$$

The term $l(x) - V(x, t)$ in (7) restricts system trajectories that enter and leave the target set, enforcing that any trajectory with a negative distance at any time will continue to have a negative distance for the rest of the time horizon. Once the value function is obtained, the BRT is given as the sub-zero level set of the value function:

$$\mathcal{V}(t) = \{x : V(x, t) \leq 0\} \quad (9)$$

The corresponding optimal safety control can be derived as:

$$u^*(x, t) = \arg \max_u \min_d \langle \nabla V(x, t), f(x, u, d) \rangle \quad (10)$$

The system can apply any control while maintaining safety, as long as it starts outside the BRT and applies the safety control in (10) at the BRT boundary. The optimal adversarial disturbance can be similarly obtained as (10).

Typically, solving HJI-VI involves computing the value function over a grid representing a discretization of state space and time. This results in an exponential scaling complexity in computation and memory, limiting its direct use to low-dimensional systems.

B. DeepReach

To overcome the computational challenges associated with solving high-dimensional HJI-VI, DeepReach [23] proposes a deep learning-based approach to solving reachability problems. Where a Deep Neuronal Network (DNN) is used to represent the value function in (6). The input to the DNN is a time point t and the state vector x . Its output is the corresponding value $V_\theta(x; t)$, where θ represents the

parameters of the DNN. The key benefit of representing the value function as a DNN is that DNNs are agnostic to grid resolution, and the memory required generally scales with the underlying value function complexity, independent of the spatial resolution. DeepReach uses a self-supervision method to learn the value function, and leverages the HJI-VI in (6) as the source of supervision. This allows the DNN to learn a high-quality approximation of the value function without explicitly solving the HJI-VI.

IV. PARAMETER-CONDITIONED REACHABLE SETS

We propose an extension to DeepReach that allows us to estimate a parameter-dependent solution to the HJI-VI. Specifically, we treat parameters β as additional virtual states of the system with zero dynamics, and compute the BRT for the augmented system using DeepReach. Thus, the overall input to the DNN is a state vector x , a time point t , and a parameter vector β , and its output is the corresponding value $V_\theta(x, t; \beta)$, where θ represents the parameters of the DNN. The key benefit of our approach is that the learned value function and the corresponding safety certificates are now conditioned not only on state and time but also on the system and environment parameters. Thus, during the run time, the system can sample the environmental factors and other parameters and activate the corresponding safety function via a simple DNN query, leading to a real-time adaptation of safety assurances.

This parameter-conditioned value function can be trained using self-supervision, similar to conventional DeepReach. Specifically, the loss function $h(x_i, t_i, \beta_i; \theta)$ used to train the DNN for a given an input (x_i, t_i, β_i) is:

$$\begin{aligned} h(x_i, t_i, \beta_i; \theta) &= h_1(x_i, t_i, \beta_i; \theta) + \lambda h_2(x_i, t_i, \beta_i; \theta), \\ h_1(x_i, t_i, \beta_i; \theta) &= \|V_\theta(x_i, t_i; \beta_i) - l(x_i; \beta_i)\|_{(t_i=T)}, \\ h_2(x_i, t_i, \beta_i; \theta) &= \|\min\{D_t V_\theta(x_i, t_i; \beta_i) + H(x_i, t_i; \beta_i), \\ &\quad l(x_i; \beta_i) - V_\theta(x_i, t_i; \beta_i)\}\| \end{aligned} \quad (11)$$

Note that l is a function of β because \mathcal{L} is a function of β_L . Similarly, H is a function of β because the Hamiltonian depends on control and disturbance bounds, which in turn depend on the parameters β_u and β_d .

Intuitively, $h_2(\cdot)$ incentivizes the DNN to learn a value function consistent with the HJI-VI. In contrast, $h_1(\cdot)$ uses the ground truth value function to additionally drive the DNN to learn the correct value function at the terminal time, thus avoiding learning any degenerate value functions. λ trades off the relative importance of the two loss functions. Together, the two loss functions encourage the DNN to learn a value function that satisfies the HJI-VI and the boundary conditions across different values of the uncertain parameters.

Once learned, the value function and its gradients can be used to adapt the optimal control decision for any state x and parameter β using (10) for the corresponding value function. This allows the system to tackle safety challenges in environments that change online. To illustrate this process the running example is considered again:

Running example (Air3D): We define the target set by the implicit surface function $l(x) := \|(x_1, x_2)\| - R$ which allows to evaluate $h_1(\cdot)$. For $h_2(\cdot)$ we additionally need an expression for the system Hamiltonian, which is given by:

$$H(x, t; \beta_u) = p_1(-v_e + v_p \cos x_3) + p_2(v_p \sin x_3) + (4 + \beta_u)|p_1 x_2 - p_2 x_1 - p_3| - \bar{\omega}_p |p_3|$$

where p_1 , p_2 , and p_3 are the spatial derivatives of the value function with respect x_1 , x_2 , and x_3 respectively. Note that, in general, the computation of the Hamiltonian requires solving an optimization problem; however, for control and disturbance affine systems, the Hamiltonian can be computed analytically in terms of value function gradients (e.g., see [19]).

V. CASE STUDIES

A. Running example (Air3D):

We now apply DeepReach to compute the parameter-conditioned BRT for the running example. We train a three-layer DNN with a hidden layer size of 512 neurons to approximate the BRT. Training parameters: 20K pretraining iterations to learn the terminal value function (target set), followed by 100K training iterations where the time horizon is progressively propagated backward, and both the HJI-VI and terminal value conditions are enforced. The 3D states and 1D parameter are randomly sampled during training and pretraining. The overall training took 1h15m on a NVIDIA GeForce RTX 3090Ti. Thanks to the ability of the DNN to efficiently encode the varying parameters, we see minimal change in the training time compared to the fixed parameter case.

In this example, the changing parameter is defined such that the evader angular velocity bound $\bar{\omega}_e(\beta_u)$ increases as β_u gets larger. Thus, we intuitively expect that for the smaller values of β_u , the reduced control bounds will make it harder for the evader to escape, resulting in more states where the pursuer can force a collision, increasing the size of the BRT. This behavior is confirmed by the results illustrated in Figure 1, where we show the slices of the BRT for different values of β_u corresponding to a relative heading of $\pi/2$. Note that despite their differing appearance, all of these slices were computed using the same value function, just queried at different β_u .




Fig. 1. (Air3D) BRT slices corresponding to different values of β_u for the relative heading of $\theta = \pi/2$. As the value of β_u decreases from 1.5 (the leftmost plot) to -1.5 (the rightmost plot), the BRT gradually increases.

Since the proposed approach computes the value function

as a continuous function of β , it can also be used for online adaptation of safety controller corresponding to the changes in β by leveraging (10). As an example, we consider the situation where the evader moves with a *nominal* control authority of $[-5, 5]rad/s$. Here, the objective for the evader is to head north (the green region in Fig. 2) while avoiding collision with the pursuer. Nominal operation of the system is interrupted due to an “engine fail” of the evader at t_0 , limiting the evader control authority to $[-2.5, 2.5]rad/s$ (corresponds to $\beta_u = -1.5$). This failure is not repaired until time t_1 , at which the system returns to the nominal control authority. Two different scenarios are compared, one using a parameter-conditioned value function to adapt the safety controller and a second one that always uses the safety controller corresponding to the nominal conditions (referred to as nominal safety controller here on). The BRTs and the vehicle trajectories for the two cases are shown in Fig. 2.

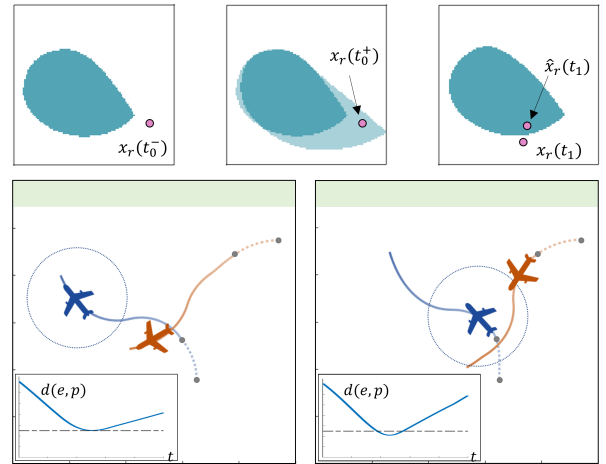


Fig. 2. (Top) BRT slices in relative coordinates before, during, and after the engine failure. The nominal model fails to consider the growth in the unsafe states (light blue region), ultimately leading to a collision. (Bottom left) Trajectory in absolute coordinates and distance between vehicles for the parameter-conditioned safety controller, the collision was closely avoided. (Bottom right) Trajectory in absolute coordinates and distance between vehicles under the nominal safety controller, leading to a collision between the vehicles.

The parameter-conditioned BRT realizes that the sudden change in maneuverability renders the system state unsafe (the system state x_r at t_0^+ is inside the light blue BRT in the top-middle figure). Thus, the evader uses the optimal avoidance controller to avoid collision for as long as possible. When the failure is cleared, this previous control action places the system in a state that is safe under the recovered control authority (the system state x_r at t_1 is just outside the BRT in the top-right figure). Correspondingly, the vehicles narrowly escape a collision (bottom left in Fig. 2). In contrast, the safety controller that is not updated with β_u assumes nominal conditions throughout and makes overly optimistic control decisions. Here again, the sudden change in maneuverability puts the system in an unsafe state that is mislabeled as safe by the non-parametrized BRT. Thus, the evader optimizes for performance instead of safety and applies a north-seeking control policy. The collision avoidance controller is not activated until the system state reaches the boundary of this smaller BRT; this delay in

activation of the safety controller results in an unsafe state $\hat{x}_r(t_1)$ after the failure is cleared, leading to a collision between the evader and the pursuer (bottom right in Fig. 2). Thus, adapting the safety controller using the parameter-conditioned value function enables a safe system behavior despite changing system conditions.

B. Rocket Landing:

We next apply our framework to an autonomous rocket landing application, where the position of the landing pad changes online. This scenario mimics the situation where the rocket needs to land on an off-shore landing pad that might drift because of ocean currents. The rocket is modeled as a rod with length l , mass M , and moment of inertia J , which is subjected to the earth's gravitational pull g . The rocket is equipped with a thruster system with two control inputs controlling the lateral and longitudinal forces (u_1, u_2) respectively. If we normalize $M = 1$ and $\alpha = l/2J$, we get the following dynamics for the system:

$$\begin{aligned} \ddot{y} &= \cos(\theta)u_1 - \sin(\theta)u_2 \\ \ddot{z} &= \sin(\theta)u_1 + \cos(\theta)u_2 - g \\ \ddot{\theta} &= \alpha u_1. \end{aligned} \quad (12)$$




Fig. 3. Rocket landing scenario. The position of the landing pad (the green hatched region) might change online.

Figure 3 represents the scenario considered. The landing pad (shown as the green hatched region) can change position horizontally in the y direction by $\pm 2l$ due to the ocean currents. We model this phenomenon as a changing parameter that corresponds to the horizontal position of the center of the landing pad. The target set is given as: $\mathcal{L}(\beta_L) = \{(y, z) : |y - \beta_L| \leq 2l, 0 \leq z \leq 2l\}$. The augmented system, in this case, is a 7D system (6 states and 1 parameter). We compute this 7D parameter-conditioned reachable set using DeepReach. The training parameters and the DNN architecture are the same as in Sec. V-A. Note that computing such 7D reachable sets are generally infeasible for the traditional HJ reachability methods, which are mostly limited to 6D systems [10].

To demonstrate the adaptability of the obtained value function, we consider two different cases: (i) a nominal safety controller that does not account for the dependency of the value function on the changing parameter; (ii) an adaptive safety controller that leverages parameter-conditioned value function to account for the movement of the target set (by constantly sampling the parameter β_L) and activating the corresponding safety controller. The rocket trajectories for the




Fig. 4. Trajectories for the center of mass of the rocket. The target set (landing platform) is shown in green with a 3D tube representing the evolution of its positions over time. The adaptive safety controller can account for the movement in the landing pad, ultimately resulting in a successful landing.

two cases are shown in Figure 4. The parameter-conditioned safety controller (the orange trajectory) constantly adjusts the horizontal thrust of the rocket with the shift in the position of the landing pad, ultimately enabling a successful landing. On the other hand, the nominal safety controller (the gray trajectory) fails to account for these movements and misses the landing pad.

C. Human-Robot Interaction:

In this case study, we consider a human-robot interaction scenario, drawn from [24], where a robot is navigating in close proximity to the human. The robot and the human are modeled as dynamical systems:

$$\begin{aligned} \dot{x}_R &= v_R \cos(\theta_R), & \dot{y}_R &= v_R \sin(\theta_R), & \dot{\theta}_R &= u_R \\ \dot{x}_H &= v_H \cos(u_H), & \dot{y}_H &= v_H \sin(u_H) \end{aligned}$$

The robot and the human both aim to reach their respective goals shown as shaded orange and blue regions in Fig. 5. In addition to reaching its goal, the robot wants to avoid any physical collision with the human. The future human states can be characterized by the Forward Reachable Tube (FRT) of the human. In contrast to the BRT, the FRT describes, in this case, the set of states that can be occupied by the human over some time horizon T . The FRT can be computed using HJ reachability, similarly to BRT [10].

In general, the FRT can be overly conservative as it assumes the human is an adversarial, goalless agent. A popular method to compensate for this conservativeness is to leverage a predictive model of the human intent to restrict the possible future actions of the human and, consequently, the growth of the FRT [24]. One such predictive model is given by:

$$u_H \sim \mathcal{N}_{vm}(u_H^*, \frac{1}{\beta}) \quad (13)$$

where \mathcal{N}_{vm} denotes the von Mises distribution, a periodic analog to the Gaussian distribution. This distribution is characterized by a mean and variance. The mean is given by u_H^* , the direction that points directly to the human goal. β

represents our confidence that the human is actually moving towards its goal. Intuitively, when model confidence is high, we are certain that the human is moving directly towards its goal, leading to a much smaller FRT (dark blue region in Fig. 5). On the other hand, when the model confidence is low, we are unsure about the human’s future actions, which is captured by the high variance in (13). This leads to a much larger FRT (the light blue region in Fig. 5). To compute the FRT, we take control bounds as the range of controls centered around u_H^* that capture 95% of the probability mass.

As the human takes each action, the belief parameter β must be updated in light of new human observations. Thus, the predicted control bounds of the human change online. This necessitates a parameter-conditioned FRT that can be queried and updated online according to the current β . In this case, the FRT computation corresponds to solving a 5D reachable set (2 states, 2 parameters corresponding to the human current state from which the FRT needs to be computed, and 1 belief parameter). An example of the usefulness of the resulting parameter-conditioned FRT is shown in Fig. 5. Here the robot under nominal circumstances plans to take the gray trajectory, but there is an obstacle in the environment not modeled by the robot. When the human deviates from its optimal path to the goal (in order to avoid the obstacle), the robot model of the human cannot explain this motion, and the model confidence drops. This results in an update in the perceived human FRT, which can now be performed online due to its parameter-conditioned nature, adjusting the robot trajectory to swerve right to avoid any collision with the human or the obstacle.

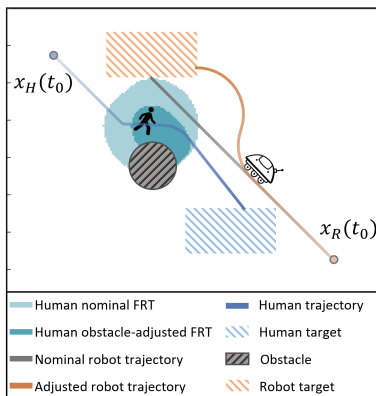


Fig. 5. Trajectories for a confidence-parameterized FRT. The human moves towards its target (shaded blue), but its trajectory (blue) shows a sudden change of direction to avoid the unmodeled obstacle (black). This causes the human’s FRT to update and expand, which leads the robot to deviate from its nominal trajectory (gray) to an adjusted trajectory (orange) to avoid collision with the human or the obstacle.

D. Three-Vehicle Collision Avoidance:

To test the proposed method in high-dimensional scenarios, the following three-vehicle collision avoidance problem is considered: we have two evaders (e_1 and e_2) and one pursuer (p), all three with dynamics as in (14), leading to a 9D system in the joint state space of the evaders and the pursuer. The pursuer here represents the adversarial agent that tries to steer evaders to collide either with the pursuer

or with each other; evaders, on the other hand, try to avoid the collision.

$$\begin{aligned} \dot{x}_1 &= v_p \cos x_3 & \dot{x}_2 &= v_p \sin x_3 & \dot{x}_3 &= \omega_p, \\ \dot{x}_4 &= v_{e1} \cos x_6 & \dot{x}_5 &= v_{e1} \sin x_6 & \dot{x}_6 &= \omega_{e1}, \\ \dot{x}_7 &= v_{e2} \cos x_9 & \dot{x}_8 &= v_{e2} \sin x_9 & \dot{x}_9 &= \omega_{e2}. \end{aligned} \quad (14)$$

We consider a situation where the pursuer finds evaders of different sizes as it traverses its environment, which implies that the relative collision distance between each pair of vehicles might change online. The pairwise collision radius between the agents can be characterized by the following three parameters $\beta_L = (\beta_{e1e2}, \beta_{e1p}, \beta_{e2p})$, resulting in the collision radii $R_{ij} = 0.1875\beta_{ij} + 0.3125$, $i, j \in [e1, e2, p]$. The target set is given as $\mathcal{L}(\beta_L) = \{x : d(i, j) \leq R_{ij}, \forall i \neq j\}$, where $d(i, j)$ is the minimum distance between the vehicles i and j .




Fig. 6. BRT slices for the position of the pursuer that results in a collision with either of the evaders, or force a collision between the evaders. BRT size changes as a function of the parameter β_L , illustrating the utility of the proposed framework in updating high-dimensional BRTs online.

This parameter triplet accounts for three additional “virtual” states, resulting in a 12D system overall. The corresponding parameter-conditioned BRT is computed using DeepReach. Fig. 6 illustrates slices of the obtained BRT for a particular position and orientation of the two evaders from the point of view of an incoming pursuer in the orientation shown by the gray plane in the bottom right of each slice. The shaded region represents the set of all starting positions of the pursuer for which a collision is unavoidable between at least one of the vehicle pairs. Different sizes for the BRTs on each slice show how the DNN can easily account for the parametrization of the value function by β_L .

VI. DISCUSSION AND FUTURE WORK

In this work, we present parameter-conditioned reachable sets to update safety assurances for systems online as they experience changes in system or environment factors. Our method relies on the recent advances in high-dimensional reachability to quickly update reachable sets online. Various simulation studies are presented to demonstrate the utility of the proposed method in maintaining safety despite the system and environment evolution.

There are a number of exciting future directions here. First, we do not provide formal safety guarantees on the obtained BRTs. We would like to explore recent work on providing safety assurances for DeepReach [25] to overcome this limitation. Addressing parameter uncertainty and non-parameterizable environment changes are also promising future directions. In addition, it will be interesting to validate the proposed approach on hardware testbeds.

REFERENCES

- [1] Y. Wang, Z. Liu, Z. Zuo, Z. Li, L. Wang, and X. Luo, "Trajectory planning and safety assessment of autonomous vehicles based on motion prediction and model predictive control," *IEEE Transactions on Vehicular Technology*, vol. 68, no. 9, pp. 8546–8556, 2019.
- [2] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-based model predictive control for safe exploration," in *2018 IEEE Conference on Decision and Control (CDC)*, 2018, pp. 6059–6066.
- [3] K. P. Wabersich and M. N. Zeilinger, "A predictive safety filter for learning-based control of constrained nonlinear dynamical systems," *Automatica*, vol. 129, 2021.
- [4] M. Z. Romdlony and B. Jayawardhana, "Stabilization with guaranteed safety using control lyapunov–barrier function," *Automatica*, vol. 66, pp. 39–47, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0005109815005439>
- [5] Y. Kumar, S. B. Roy, and P. B. Sujit, "Barrier lyapunov function based trajectory tracking controller for autonomous vehicles with guaranteed safety bounds," in *2020 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2020, pp. 722–728.
- [6] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [7] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *IEEE European control conference (ECC)*, 2019.
- [8] H. Yazarel, S. Prajna, and G. Pappas, "S.O.S. for safety," in *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601)*, vol. 1, 2004, pp. 461–466 Vol.1.
- [9] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *International Workshop on Hybrid Systems: Computation and Control*, 2004.
- [10] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, "Hamilton-Jacobi Reachability: A brief overview and recent advances," in *IEEE Conference on Decision and Control (CDC)*, 2017.
- [11] S. L. Herbert, S. Bansal, S. Ghosh, and C. J. Tomlin, "Reachability-based safety guarantees using efficient initializations," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2019, pp. 4810–4816.
- [12] K. Leung, E. Schmerling, M. Zhang, M. Chen, J. Talbot, J. C. Gerdes, and M. Pavone, "On infusing reachability-based safety assurance within planning frameworks for human–robot vehicle interactions," *The International Journal of Robotics Research*, vol. 39, no. 10-11, pp. 1326–1345, 2020.
- [13] S. Bansal, A. Bajcsy, E. Ratner, A. D. Dragan, and C. J. Tomlin, "A hamilton-jacobi reachability-based framework for predicting and analyzing human motion for safe planning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 7149–7155.
- [14] A. Majumdar and R. Tedrake, "Funnel libraries for real-time robust feedback motion planning," *The International Journal of Robotics Research*, vol. 36, no. 8, pp. 947–982, 2017.
- [15] S. L. Herbert, M. Chen, S. Han, S. Bansal, J. F. Fisac, and C. J. Tomlin, "Fastrack: A modular framework for fast and guaranteed safe motion planning," in *IEEE Conference on Decision and Control (CDC)*, 2017.
- [16] S. Singh, A. Majumdar, J.-J. Slotine, and M. Pavone, "Robust online motion planning via contraction theory and convex optimization," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [17] D. Q. Mayne, E. C. Kerrigan, E. Van Wyk, and P. Falugi, "Tube-based robust nonlinear model predictive control," *International journal of robust and nonlinear control*, vol. 21, no. 11, pp. 1341–1353, 2011.
- [18] S. Koenig and M. Likhachev, "Fast replanning for navigation in unknown terrain," *IEEE Transactions on Robotics*, vol. 21, no. 3, pp. 354–363, 2005.
- [19] I. Mitchell, A. Bayen, and C. J. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Transactions on Automatic Control (TAC)*, vol. 50, no. 7, pp. 947–957, 2005.
- [20] J. Lygeros, "On reachability and minimum cost optimal control," *Automatica*, vol. 40, no. 6, pp. 917–927, 2004.
- [21] H. El-Kebir, A. Pirosmanshvil, and M. Ornik, "Online guaranteed reachable set approximation for systems with changed dynamics and control authority," *arXiv preprint arXiv:2203.10220*, 2022.
- [22] A. Bajcsy, S. Bansal, E. Bronstein, V. Tolani, and C. J. Tomlin, "An efficient reachability-based framework for provably safe autonomous navigation in unknown environments," in *IEEE Conference on Decision and Control (CDC)*, 2019.
- [23] S. Bansal and C. J. Tomlin, "DeepReach: A deep learning approach to high-dimensional reachability," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 1817–1824.
- [24] D. Fridovich-Keil, A. Bajcsy, J. F. Fisac, S. L. Herbert, S. Wang, A. D. Dragan, and C. J. Tomlin, "Confidence-aware motion prediction for real-time collision avoidance1," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 250–265, 2020.
- [25] A. Lin and S. Bansal, "Generating formal safety assurances for high-dimensional reachability," 2022. [Online]. Available: <https://arxiv.org/abs/2209.12336>