

Multi-Adversarial Safety Analysis for Autonomous Vehicles

Gilbert Bahati
Civil and Environmental Engineering
University of California - Berkeley
 bahati95@berkeley.edu

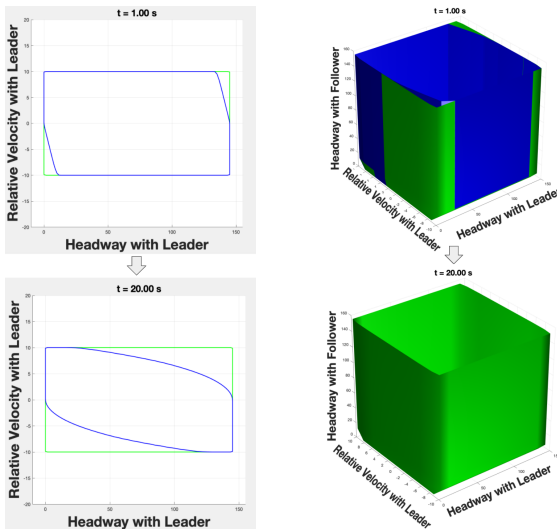
Marsalis Gibson
Electrical Engineering and Computer Science
University of California - Berkeley
 mtgibson@berkeley.edu

Alexandre Bayen
Institute of Transportation Studies
University of California - Berkeley
 director-its@berkeley.edu

Abstract—This work in progress considers reachability-based safety analysis in the domain of autonomous driving in multi-agent systems. We formulate the safety problem for a car following scenario as a differential game and study how different modelling strategies yield very different behaviors regardless of the validity of the strategies in other scenarios. Given the nature of real-life driving scenarios, we propose a modeling strategy in our formulation that accounts for subtle interactions between agents, and compare its Hamiltonian results to other baselines. Our formulation encourages reduction of conservativeness in Hamilton-Jacobi safety analysis to provide better safety guarantees during navigation.

I. INTRODUCTION

If autonomous vehicles are to serve as traffic management systems [5], safe navigation around human vehicles on highways and in cities is crucial. However, safe navigation can be difficult to provide because a lot of uncertainty exists in real driving scenarios that complicate the driving problem. Typically, Hamilton-Jacobi reachability analysis (HJI) can be used to find safe strategies around unknown components of a dynamical system [1]. In previous work, researchers develop [2] a framework to protect a system against one known source of uncertainty using Hamilton-Jacobi reachability, with the goal of protecting the system from the worst-case scenario. However, in real driving scenarios, it may be necessary to consider multiple sources of uncertainty. As depicted in figure 1 and 2, extreme worst-case scenarios may never provide a feasible safety strategy, and it may be the case that establishing safety is impossible.



(a) Two player scenario (2D state system) in fig 2a (b) Three player scenario (4D state system) in fig 2b

Fig. 1: HJI reachability analysis [3] computations for fig 2: the constraint set (green) represents our state boundaries and the reachable safe set (blue) remains within/propagates inwards as t increases. To interpret this, a state that starts within the blue safe set is guaranteed to remain within the green constraint set for t seconds. For ease of visualization, 1b is a 3D slice of the 4D state system where the relative velocity with follower is 0.




Fig. 2: The black car represents the autonomous agent while the other two cars represent human drivers. In the two car scenario in 2a, we compute the safe reachable set under the assumption that the leading car acts adversarially. In this scenario, we obtain a feasible solution, see fig 1a. However, this is not the case when a third human driven car is introduced behind the autonomous agent in 2b. The autonomous agent cannot find a safe reachable set, see fig 1b, which means that there are no states or points in time in which the car is guaranteed to be safe if the two human vehicles choose to act adversarially.

Therefore, in this work, we study the reduction of conservativeness in Hamilton-Jacobi safety analysis by introducing structure into some or all of the human models. Specifically, we study a modeling strategy around the second disturbance that takes advantage of the structure of human behavior in a way that allows us to use differential game theory in more dense dynamic driving environments.

II. SYSTEM DYNAMICS




Fig. 3: Car following scenario where the subject vehicle drives in between two human vehicles.

We consider a dynamical system with state $z \in R^n$, and three inputs, $u \in \mathcal{U} \subset R^{n_u}$, $d_1 \in \mathcal{D}_1 \subset R^{n_d}$, $d_2 \in \mathcal{D}_2 \subset R^{n_d}$, which we refer to as the controls, disturbance 1, and disturbance 2 respectively. Our system dynamics are generally defined as:

$$\dot{z} = f(z, u, d_1, d_2) \quad (1)$$

Disturbance 1 and 2 represent the uncertainty around the leading and following human vehicle respectively. In our car following scenario in figure 3, the goal for the autonomous agent is to establish safety and remain in between the other two players given their actions. Thus, the dynamics between all three vehicles can be described using their relative position x_j , relative speed v_j , and relative accelerations u_i as in the following:

$$\begin{bmatrix} \dot{x}_{g1} \\ \dot{v}_{g1} \\ \dot{x}_{g2} \\ \dot{v}_{g2} \end{bmatrix} = \begin{bmatrix} v_{g1} \\ u_1 - u_2 \\ v_{g2} \\ u_2 - u_3 \end{bmatrix} \quad (2)$$

s.t.

$$u_i \in [a_{min}, a_{max}], \forall i = 1...3$$

$$x_j > 0, \forall j = g_1, g_2$$

In the next section, Section III, we discuss how we choose our uncertainty and how we pose our safety problem.

III. THREE-PLAYER DIFFERENTIAL GAME

The safety problem is posed as a differential game between three players, where the system controller, u , plays against two adversaries, d_1 and d_2 , also known as the system's uncertainty. To obtain a *safe policy* for the system, we chose a function, $l(\mathbf{z})$, that assigns a safety value to the current state, \mathbf{z} and formulate a game whose outcome is given by the function $\mathcal{V} : R^n \times \mathcal{U} \times \mathcal{D}_1 \times \mathcal{D}_2 \rightarrow R$. \mathcal{V} assigns each initial state \mathbf{z} and player strategies $u(\cdot)$, $d_1(\cdot)$, $d_2(\cdot)$, the lowest value of $l(\cdot)$ ever achieved by a trajectory $\xi_{d_1, d_2}^{\mathbf{z}, u}(\cdot)$ from state \mathbf{z} .

$$\mathcal{V}(\mathbf{z}, u(\cdot), d_1(\cdot), d_2(\cdot)) = \inf_{t \geq 0} l(\xi_{d_1, d_2}^{\mathbf{z}, u}(t)) \quad (3)$$

The goal of system is to maximize the objective, while the goal of the active adversaries is to minimize the objective. Thus, the game formulation that we want to solve is¹:

$$V(\mathbf{z}) = \inf_{d_1, d_2} \sup_u \mathcal{V}(\mathbf{z}, u(\cdot), d_1(\cdot), d_2(\cdot)) \quad (4)$$

A. Player Strategies

We formulate uncertainties d_1 and d_2 around the two human driving actions, for example u_1 and u_3 , to represent behavioral properties that are trying to perturb the autonomous system. More specifically:

- 1) First, we consider a **baseline** assignment:
 - $d_1 = u_1$ where $u_1 \in [a_{min}, a_{max}]$
 - $d_2 = u_3$ where $u_3 \in [a_{min}, a_{max}]$
- 2) Then, we consider an **alternative** assignment for d_2 by taking advantage of the structure of human driving and modeling u_3 using a car following model $g(z, d_2)$:
 - $d_1 = u_1$ where $u_1 \in [a_{min}, a_{max}]$
 - $u_3 = g(z, d_2)$

In our second strategy, d_2 uses psycho-physiological characteristics in human driving as an alternative modeling strategy [4]. Additionally, we ensure that the values of u_3 are within realistic bounds given all possible autonomous agent's actions u_2 . This modelling strategy relaxes unrealistic extremities of the previous dynamic game formulation and implicitly models interaction effects between agents for realistic safety. We model the following vehicle's driving behavior using the Intelligent Driver's car following model and explicitly model d_2 as **safe-reaction time, T** , as follows:

$$g(z, T) = a \left(1 - \left(\frac{v_3}{v_0} \right)^\delta - \left(\frac{s^*(z, T)}{x_{g_2}} \right)^2 \right) \quad (5)$$

$$s^*(z, T) = s_0 + \max \left(0, v_3 T + \frac{v_3(-v_{g_2})}{2\sqrt{ab}} \right) \quad (6)$$

where:

T : safe reaction-time (ie. $d_2 = T \in [0, T_{max}]$)

$s^*(z, T)$: desired headway of the following vehicle

s_0 : minimum desired headway (ie. $s_0 = 0$ to allow crashes)

a, b : maximum acceleration and deceleration respectively

δ, v_0 : acceleration exponent (usually 4) and desired velocity respectively

¹Technically, as in [3], we restrict each disturbance to a set of non-negativity strategies. Therefore, $d_1(\cdot)$ and $d_2(\cdot)$ in eq. 4 are actually maps, $\beta_1[u(\cdot)](\cdot)$ and $\beta_2[u(\cdot)](\cdot)$, that respectively maps our control input to their corresponding disturbance input.

B. Resulting policies

The optimal control strategy, u_2^* , and the optimal disturbance strategy for the first human vehicle, d_1^* are calculated from (4) using the Hamiltonian numerics to be:

$$u_2^* = \begin{cases} a_{max} & \text{if } (p_4 - p_2) > 0 \\ a_{min} & \text{else} \end{cases} \quad (7)$$

$$d_1^* = \begin{cases} a_{max} & \text{if } p_2 < 0 \\ a_{min} & \text{else} \end{cases} \quad (8)$$

and the optimal disturbance strategy for the second human vehicle is likewise calculated to be:

$$\text{Baseline } d_2^* = \begin{cases} a_{max} & \text{if } p_4 > 0 \\ a_{min} & \text{else} \end{cases} \quad (9)$$

$$\text{Alternative } d_2^* = \begin{cases} \min(T_{max}, \max(T_{min}, \frac{-v_{g_2}}{2\sqrt{ab}})) & \text{if } p_4 > 0 \\ \frac{-v_{g_2}}{2\sqrt{ab}} & \text{else} \end{cases} \quad (10)$$

where: $p_2 = \nabla_{v_{g_1}} V$ and $p_4 = \nabla_{v_{g_2}} V$.

IV. RESULTS AND CONCLUSION

As depicted in figure 4, the alternative formulation for the second disturbance is able to uncover hidden safe strategies for the 3-car scenario. Setting the disturbance and control bounds to $[-1.5, 1.5]$ and $[-2, 2]$ respectively, the baseline produces an empty blue set, meaning there are no guaranteed safe states for the simulation with this particular disturbance/control setting. However, by taking advantage of psycho-physiological characteristics and driver influences on the road, we discover that a non-empty blue set does exist.




Fig. 4: Invariant safe states resulting from the baseline and alternative technique.

In conclusion, when considering worst-case uncertainty in human driving, if we are to guarantee safety in the chaotic world of driving, we may need to incorporate better information structures of human behavior in our analysis and update our assumptions as we uncover more knowledge about the system. However, by choosing a specific model structure to reduce conservativeness of reachability analysis, we run the risk of not being able to capture human behavior some of the time due to the limitations of our chosen model (which serves to capture only approximations). The performance of a chosen model will vary greatly depending on the particular type of driver and circumstance. Therefore, to tackle this trade-off, we further aim to incorporate real-time analysis and data-driven models to learn disturbances (and how they accurately evolve), and maintain formal and robust safety guarantees using different learning strategies such as deep reinforcement learning.

REFERENCES

- [1] Somil Bansal, Mo Chen, Sylvia Herbert, and Claire J Tomlin. Hamilton-jacobi reachability: A brief overview and recent advances. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 2242–2253. IEEE, 2017.
- [2] Jaime F Fisac, Anayo K Akametalu, Melanie N Zeilinger, Shahab Kaynama, Jeremy Gillula, and Claire J Tomlin. A general safety framework for learning-based control in uncertain robotic systems. *IEEE Transactions on Automatic Control*, 64(7):2737–2752, 2018.
- [3] Ian M Mitchell and Jeremy A Templeton. A toolbox of hamilton-jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 480–494. Springer, 2005.
- [4] M. Treiber and C. Thiemann. *Traffic Flow Dynamics: Data, Models and Simulation*. Springer-Verlag Berlin Heidelberg, Berlin, Germany, 2013.
- [5] Cathy Wu, Aboudy Kreidieh, Kanaad Parvate, Eugene Vinitsky, and Alexandre M Bayen. Flow: Architecture and benchmarking for reinforcement learning in traffic control. *arXiv preprint arXiv:1710.05465*, 2017.