

# Learning Responsibility Allocations for Safe Human-Robot Interaction with Applications to Autonomous Driving

Ryan K. Cosner, Yuxiao Chen, Karen Leung, and Marco Pavone

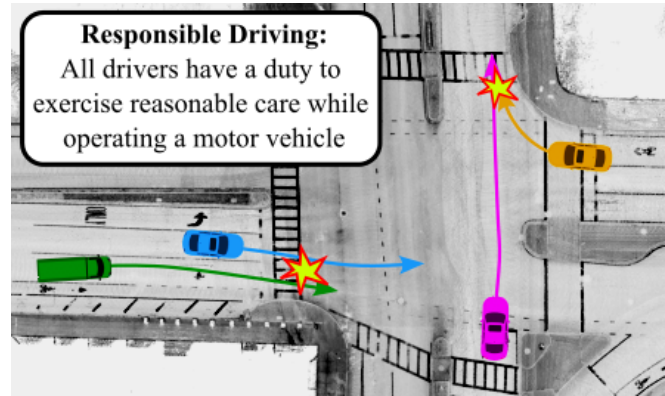
**Abstract**—Drivers have a responsibility to exercise reasonable care to avoid collision with other road users. This assumed responsibility allows interacting agents to maintain safety without explicit coordination. Thus to enable safe autonomous vehicle (AV) interactions, AVs must understand what their responsibilities are to maintain safety and how they affect the safety of nearby agents. In this work we seek to understand how responsibility is shared in multi-agent settings where an autonomous agent is interacting with human counterparts. We introduce *Responsibility-Aware Control Barrier Functions (RA-CBFs)* and present a method to learn responsibility allocations from data. By combining safety-critical control and learning-based techniques, RA-CBFs allow us to account for scene-dependent responsibility allocations and synthesize safe and efficient driving behaviors without making worst-case assumptions that typically result in overly-conservative behaviors. We test our framework using real-world driving data and demonstrate its efficacy as a tool for both safe control and forensic analysis of unsafe driving.

## I. INTRODUCTION

Drivers have a duty to exercise reasonable care when interacting with other road users. The assumption that other road users will exercise *responsible* behaviors enables everyone to maintain safety without explicit coordination. Thus, for autonomous vehicles (AVs) to safely and seamlessly interact with other road users, they must utilize a safety framework that can interpretably codify basic safe driving behaviors (such as those laid out by the IEEE safe driving standards [1]), and exercise duty of care, i.e., drive responsibly. However, the responsibility for ensuring safety is shared asymmetrically and the allocation is influenced by context-dependent social driving norms. For example, drivers are generally more responsible for avoiding collisions with vehicles in front of them than behind them (see Figure 1).

The goal of this work is to develop a rigorous yet flexible safety framework that is capable of codifying and synthesizing safe and responsible driving behaviors [1]. Unfortunately, most existing models and techniques for ensuring AV safety tend to make strong assumptions about how other drivers behave which often results in defensive or erratic driving. For instance, many approaches make worst-case assumptions about the other agents' behaviors [2], [3], [4], [5]; while these assumptions allow for strong theoretical guarantees, they are often impractical as they result in infeasible planning problems or induce overly conservative behaviors. Therefore,

Ryan K. Cosner is with the California Institute of Technology, [rkcosner@caltech.edu](mailto:rkcosner@caltech.edu), Yuxiao Chen is with NVIDIA, [yuxiaoc@nvidia.com](mailto:yuxiaoc@nvidia.com), Karen Leung is with the University of Washington and NVIDIA {[kymleung@uw.edu](mailto:kymleung@uw.edu), [kaleung@nvidia.com](mailto:kaleung@nvidia.com)}, Marco Pavone is with Stanford University and NVIDIA {[pavone@stanford.edu](mailto:pavone@stanford.edu), [mpavone@nvidia.com](mailto:mpavone@nvidia.com)}. Ryan's work was performed during an internship with NVIDIA.



**Fig. 1.** In human driving, vehicles can be expected to demonstrate a reasonable duty of care. For example, a trailing vehicle (green) takes responsibility for not colliding with the car in front of it (blue) and a merging vehicle (orange) follows formal and informal rules to avoid a collision with the vehicles in the lane (pink). How can we ensure that autonomous vehicles act according to such informal driving etiquette?

a key challenge is developing safety models that are not only robust to varied driving behaviors, but that are also capable of accounting for context-dependent social norms that effect how drivers implicitly coordinate.

With that goal in mind, we propose Responsibility-Aware Control Barrier Functions (RA-CBFs), a *responsibility-aware* safety paradigm that is conducive to efficient techniques for safe control synthesis and safety evaluation. A primary use case of RA-CBFs is their application as a *safety filter* within an AV stack, whereby the output of the main AV decision-making and planning algorithmic pipeline (powered, for example, by high-capacity human behavior prediction models based on deep learning) is monitored for its compliance with basic safe and responsible driving behaviors (as encoded by the RA-CBF) and possibly revised in case of anomalies, in a manner similar to [3] and [5]. As a proof of concept, we demonstrate this capability in the context of real-time closed-loop control and post-facto forensic analysis.

**Related work.** We present relevant work on responsibility-aware safety methods for multi-agent interactions with an emphasis on AV applications. Specifically, we consider related work on socially-aware planning, safe multi-agent control, and safety constraint learning.

Many recent works focus on modeling and estimating drivers' social preferences to synthesize socially-aware driving behaviors. For instance, [6] estimates the Social Value Orientation (SVO) of other drivers and formulates a SVO-based game-theoretic planner. [7], [8] craft a (potentially learned) planning reward function that incentivizes an AV to be more cooperative, sympathetic, and/or courteous. While these approaches demonstrate that accounting for social preferences can lead to more human-like AV behaviors, they

do not provide any assurances or quantification of AV safety.

In the safety-critical control literature, under decentralized multi-agent settings (which is more relevant to AV applications), many certifiably safe multi-agent control methods exist, including Control Barrier Functions (CBFs) [4], [9], [10], [11], Reciprocal Velocity Obstacles [12], Hamilton-Jacobi Reachability [13], [14], Responsibility Sensitive Safety [5], and Safety Force Field [3]. However, these methods rely on worst-case or static assumptions on how other agents behave, which, unfortunately, do not hold in practice. This limitation has spurred recent works [15], [16], [17] that investigate collision-avoidance responsibility and how to decide which agent should take more effort to avoid collisions. However, these methods consider either centralized control or centrally defined social preferences, which does not apply to autonomous driving since the social preferences of other agents are not known precisely and cannot be assigned.

Learning responsibility allocation influenced by social norms necessitates the need to combine safety-critical control with data-driven methods. Recent safety-critical learning methods [18], [19], [20], [21] infer safety constraints from data. To the best of our knowledge, the consideration of responsibility has yet to be investigated within a data-driven safe multi-agent control setting. In this work, we take on a data-driven approach to learn how responsibility is allocated among multiple agents and synthesize corresponding safe and responsible controls. We elect to use CBFs as the basis of our method due to its interpretable and rigorous control-theoretic formulation and computational tractability.

**Contributions:** Our contributions are three-fold: (i) We present a novel concept of Responsibility-Aware Control Barrier Functions (RA-CBFs) which extends the standard CBF to account for asymmetric sharing of responsibility between multiple agents. (ii) We propose a data-driven constraint-learning algorithm to infer the responsibility allocations modeled in the RA-CBF formulation. (iii) We showcase our method using real-world driving data and demonstrate the utility of RA-CBFs and the learned responsibility allocations in safe closed-loop AV control and forensic analysis.

## II. BACKGROUND

The dynamics of the vehicles are abstracted as a nonlinear control-affine system of the form:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \quad (1)$$

where  $\mathbf{x} \in \mathbb{R}^n$  and  $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$  represent the states and inputs of the system and where  $\mathcal{U}$  is assumed to be compact. We further assume that  $\mathbf{f} : \mathbb{R}^n \rightarrow \mathbb{R}^n$  and  $\mathbf{g} : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$  are locally Lipschitz continuous. Given a state-feedback controller  $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ , the closed-loop dynamics are:

$$\dot{\mathbf{x}} = \mathbf{f}_c(\mathbf{x}) = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{k}(\mathbf{x}). \quad (2)$$

### A. Control Barrier Functions

In this section, we define safety as the forward-invariance of some safe set  $\mathcal{C} \subset \mathbb{R}^n$  and review Control Barrier Functions (CBFs) as a tool for synthesizing safe controllers.

**Definition 1** (Forward Invariance and Safety). A set  $\mathcal{C} \subset \mathbb{R}^n$  is forward invariant if for every  $\mathbf{x}(0) \in \mathcal{C}$  the solution to (2)

satisfies  $\mathbf{x}(t) \in \mathcal{C}$  for all  $t \geq 0$ . We call the system (2) safe with respect to  $\mathcal{C}$  if  $\mathcal{C}$  is forward invariant.

Let the set  $\mathcal{C}$  be the 0-superlevel set of some continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$  with 0 a regular value<sup>1</sup>:

$$\mathcal{C} \triangleq \{\mathbf{x} \in \mathbb{R}^n \mid h(\mathbf{x}) \geq 0\} \quad (3)$$

We deem  $h$  a CBF if it also satisfies the following condition:

**Definition 2** (Control Barrier Function [22]). Let  $\mathcal{C} \subset \mathbb{R}^n$  be the 0-superlevel set of some continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$  with 0 a regular value. The function  $h$  is a Control Barrier Function (CBF) for (1) if there exists an extended class  $\mathcal{K}_\infty$  function<sup>2</sup>  $\alpha$  such that for all  $\mathbf{x} \in \mathcal{C}$ :

$$\sup_{\mathbf{u} \in \mathcal{U}} \underbrace{L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u}}_{\frac{dh}{dt}(\mathbf{x}, \mathbf{u})} \geq -\alpha(h(\mathbf{x})). \quad (4)$$

Here, Lie derivative notation is used to represent the partial derivatives  $L_{\mathbf{f}}h(\mathbf{x}) \triangleq \frac{\partial h}{\partial \mathbf{x}}\mathbf{f}(\mathbf{x})$  and  $L_{\mathbf{g}}h(\mathbf{x}) \triangleq \frac{\partial h}{\partial \mathbf{x}}\mathbf{g}(\mathbf{x})$ . Intuitively, the CBF constraint (4) limits  $\frac{dh}{dt}$ , and prevents  $h$  from decreasing along the trajectory when  $h(\mathbf{x}) = 0$ , thus rendering  $\mathcal{C}$  forward invariant for (2). This intuition was formalized by [22] in the following theorem:

**Theorem 1** (CBF Safety, [22]). Given a set  $\mathcal{C} \subset \mathbb{R}^n$  defined as the 0-superlevel set of a continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$  with 0 a regular value, if  $h$  is a CBF, then any locally Lipschitz controller  $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$  that satisfies (4) for all  $\mathbf{x} \in \mathcal{C}$ , renders the system (2) safe w.r.t.  $\mathcal{C}$ .

## III. RESPONSIBILITY-AWARE DECENTRALIZED MULTI-AGENT SAFETY

In this section, we extend the CBF framework to a decentralized multi-agent setting and introduce additional terms to account for asymmetrically shared responsibility.

### A. Decentralization Multi-agent CBF

We extend the system dynamics (1) to multiple agents:

$$\dot{\mathbf{x}}_i = \mathbf{f}_i(\mathbf{x}_i) + \mathbf{g}_i(\mathbf{x}_i)\mathbf{u}_i \quad (5)$$

where  $\mathbf{x}_i \in \mathbb{R}^{n_i}$ ,  $\mathbf{u}_i \in \mathcal{U}_i \subset \mathbb{R}^{m_i}$ ,  $\mathbf{f}_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{n_i}$ ,  $\mathbf{g}_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{m_i}$  represent the state, input, drift, and actuation matrix of agent  $i$ . For the entire system of  $N \in \mathbb{N}$  agents, let  $\mathbf{x} = [\mathbf{x}_1^\top \ \cdots \ \mathbf{x}_N^\top]^\top$  denote the concatenated state and the dynamics for  $\mathbf{x}$  be denoted as in (1):

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \mathbf{f}_1(\mathbf{x}_1) \\ \vdots \\ \mathbf{f}_N(\mathbf{x}_N) \end{bmatrix}}_{\mathbf{f}(\mathbf{x})} + \underbrace{\begin{bmatrix} \mathbf{g}_1(\mathbf{x}_1) \\ \vdots \\ \mathbf{g}_N(\mathbf{x}_N) \end{bmatrix}}_{\mathbf{g}(\mathbf{x})} \underbrace{\begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_N \end{bmatrix}}_{\mathbf{u}}. \quad (6)$$

If the multi-agent system is governed by a centralized controller, the CBF inequality can be checked directly and used as a constraint in an optimization-based controller to obtain safe inputs [23]. However, centralized control is often

<sup>1</sup>We say that 0 is a regular value of  $h$  if  $h(\mathbf{x}) = 0 \implies \frac{\partial h}{\partial \mathbf{x}} \neq 0$ .

<sup>2</sup>A continuous function  $\alpha : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  is a class  $\mathcal{K}_\infty$  function if  $\alpha(0) = 0$ ,  $\alpha$  is strictly monotonically increasing, and  $\lim_{c \rightarrow \infty} \alpha(c) = \infty$ . A continuous function  $\alpha : \mathbb{R} \rightarrow \mathbb{R}$  is an extended  $\mathcal{K}_\infty$  function if  $\alpha(0) = 0$ ,  $\alpha$  is strictly monotonically increasing,  $\lim_{c \rightarrow \infty} \alpha(c) = \infty$ , and  $\lim_{c \rightarrow -\infty} \alpha(c) = -\infty$ .

unrealizable for AVs due to communication and scalability issues as well as the presence of human actors. Thus, we focus on a decentralized variant of the CBF constraint (4) and assume that each agent can measure the states of the other agents, but independently generates its own input according to some controller unknown to the other agents.

One common method for retaining safety guarantees in the context of decentralized control, is to ensure robustness with respect to all possible controls of the other agents (including the worst-case inputs) as in [4]. In this case, constraint (4) from the perspective of agent  $i$  becomes:

$$\sup_{\mathbf{u}_i \in \mathcal{U}_i} \inf_{\substack{\mathbf{u}_j \in \mathcal{U}_j, \\ j \neq i}} L_{\mathbf{f}}h(\mathbf{x}) + L_{\mathbf{g}}h(\mathbf{x})\mathbf{u} \geq -\alpha(h(\mathbf{x})). \quad (7)$$

This is a conservative constraint which ensures the safety of the system even when other agents act adversarially. For human-interactive systems where responsibility is shared, this controller is often unnecessarily conservative.

Despite their safety-guarantees, worst-case constraints like (7) are highly conservative and prevent fluent behaviors [24]. It is therefore desirable to find a less conservative safety constraint that is more cognizant of the social interactions between agents even when the controllers of the other agents are unknown. For this purpose, we consider a novel CBF framework that models *social responsibility*.

### B. Responsible-Aware Control Barrier Functions

In multi-agent systems of human actors, the responsibility for maintaining safety is typically shared amongst people. For example, humans exhibit social behavior in crowd navigation and driving where the burden of maintaining safety is distributed amongst everyone [6], [25]. Equipped with the notion that agents share the responsibility of maintaining safety, we move away from worst-case behavioral assumptions, and instead, *learn* the responsibility allocation from data. First, we define *responsibility allocation functions*:

**Definition 3** (Responsibility Allocation Function). A function  $\gamma : \mathbb{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$  is a *responsibility allocation function* for  $N \in \mathbb{N}$  if for all  $\mathbf{x} \in \mathbb{R}^n$ :

$$\sum_{i \in \{1, \dots, N\}} \gamma(i, \mathbf{x}) \geq 0, \quad (8)$$

For agent  $i$  in a multi-agent system at state  $\mathbf{x}$ ,  $\gamma(i, \mathbf{x}) > 0$  indicates increased responsibility,  $\gamma(i, \mathbf{x}) = 0$  indicates evenly shared responsibility, and  $\gamma(i, \mathbf{x}) < 0$  indicates decreased responsibility. The sum of  $\gamma(i, \mathbf{x})$  is lower bounded by zero to ensure that the total allocated responsibility must be greater than or equal to that of even sharing.

Using these responsibility allocation functions we can present our definition of Responsibility-Aware Control Barrier Functions (RA-CBFs) which consider responsibility allocation in their decentralized multi-agent safety constraint:

**Definition 4** (Responsibility-Aware Control Barrier Function). Let  $\mathcal{C} \subset \mathbb{R}^n$  be the 0-superlevel set of some continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$  with 0 a regular value. Additionally, let  $\gamma : \mathbb{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$  be a responsibility allocation function for  $N \in \mathbb{N}$ . The function  $h$  is a *Responsibility-Aware CBF* for the system (6) and responsibility allocation function  $\gamma$  if there exists an extended

class  $\mathcal{K}_\infty$  function  $\alpha$  such that for all  $\mathbf{x} \in \mathcal{C}$  and all  $i \in \{1, \dots, N\}$ :

$$\sup_{\mathbf{u}_i \in \mathcal{U}_i} \underbrace{L_{\mathbf{g}_i}h(\mathbf{x})\mathbf{u}_i + \frac{1}{N} \left( \alpha(h(\mathbf{x})) + L_{\mathbf{f}}h(\mathbf{x}) \right)}_{\text{RA-CBF Constraint}(i, \mathbf{x}, \mathbf{u}, \gamma)} - \gamma(i, \mathbf{x}) \geq 0. \quad (9)$$

**Remark 1.** For generality, RA-CBFs are presented for  $N$  agents but in practice it is common to enforce CBF constraints between each pair of agents where the number of constraints enforced on agent  $i$ 's input grows linearly with the number of agents [26]. In this case there would be several pairwise RA-CBF constraints with  $N = 2$ . In Section V, we take this approach in our application.

**Theorem 2** (Responsibility-Aware Safety). *Given a set  $\mathcal{C} \subset \mathbb{R}^n$  defined as the 0-superlevel set of a continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$  with 0 a regular value, if  $h$  is an RA-CBF for (6) and the responsibility allocation function  $\gamma : \mathbb{N} \times \mathbb{R}^n \rightarrow \mathbb{R}$  for  $N \in \mathbb{N}$ , then any locally Lipschitz controller  $\mathbf{k} : \mathbb{R}^n \rightarrow \mathbb{R}^m$  that satisfies (9) for all  $\mathbf{x} \in \mathcal{C}$  and  $i \in \{1, \dots, N\}$ , renders system (2) safe with respect to  $\mathcal{C}$ .*

*Proof.* First let  $\mathbf{c}_i(\mathbf{x}, \mathbf{k}(\mathbf{x})) \triangleq L_{\mathbf{g}_i}h(\mathbf{x})\mathbf{k}_i(\mathbf{x}) + \frac{1}{N} \left( \alpha(h(\mathbf{x})) + L_{\mathbf{f}}h(\mathbf{x}) \right)$  for all  $i \in \{1, \dots, N\}$ . Since the  $\mathbf{k}_i$  satisfies (9),

$$0 \geq -\mathbf{c}_i(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \gamma(i, \mathbf{x}), \quad (10)$$

$$\geq - \sum_{i \in \{1, \dots, N\}} \mathbf{c}_i(\mathbf{x}, \mathbf{k}(\mathbf{x})) + \sum_{i \in \{1, \dots, N\}} \gamma(i, \mathbf{x}), \quad (11)$$

$$\geq - \sum_{i \in \{1, \dots, N\}} \mathbf{c}_i(\mathbf{x}, \mathbf{k}(\mathbf{x})) \quad (12)$$

Inequality (11) follows from the decentralized constraint (9) for all  $i$  and (12) holds since  $\gamma$  is a responsibility allocation function for  $N$ . Since the final inequality (12) is equivalent to the centralized CBF constraint (4), Theorem 1 implies the safety of system (2) with respect to  $\mathcal{C}$ .  $\square$

In summary, instead of considering the worst-case inputs from other agents, our RA-CBF approach uses  $\gamma(i, \mathbf{x})$  to allow agent  $i$ 's required contribution to decentralized multi-agent safety to vary depending on the state  $\mathbf{x}$  of all of the agents in the scene. Also, instead of explicitly considering the uncertainty in the other agents' actions, one perspective of the responsibility allocation function is that it models a bound on the projection [27] of the other agents' inputs onto the CBF time derivative. Thus we can learn the effect of the other agents' actions as a scalar adjustment to  $\frac{dh}{dt}$  as opposed to predicting their exact trajectories.

Our responsibility model is similar to that of [15] which instead uses a multiplicative term and is limited to driftless systems (i.e., systems where  $\mathbf{f}(\mathbf{x}) \equiv 0$ ). By using an additive term, our model is generally applicable to control affine systems and is capable of accounting for the effect of responsibility even when the unforced dynamic are unsafe, i.e.  $\alpha(h) + L_{\mathbf{f}}h(\mathbf{x}) \leq 0$ .

## IV. LEARNING RESPONSIBILITY ALLOCATION

In this section we formalize the problem of learning responsibility allocation functions  $\gamma(i, \mathbf{x})$  from data and de-

scribe our method for learning  $\gamma$  from expert demonstrations, given a known safe set  $\mathcal{C}$  and associated CBF  $h$ .

### A. Problem Setting

We assume that agent  $i$  strives to minimize some unknown function  $Q_i : \mathbb{R}^{n_i \times m_i} \rightarrow \mathbb{R}$  and does so according to a constrained optimal control policy:

$$\mathbf{k}_i(\mathbf{x}) = \underset{\mathbf{u}_i \in \mathcal{U}_i}{\operatorname{argmin}} Q_i(\mathbf{x}_i, \mathbf{u}_i) \quad (13)$$

s.t. RA-CBF Constraint( $i, \mathbf{x}, \mathbf{u}, \gamma$ )  $\geq 0$

where  $Q_i(\mathbf{x}_i, \mathbf{u}_i)$  represents agent  $i$ 's cost for input  $\mathbf{u}_i$  at state  $\mathbf{x}_i$ . Although the cost function is unknown, we assume that all agents obey the RA-CBF constraint for some  $\gamma$  that we seek to learn, thus framing the problem of learning responsibility allocations as constraint learning.

### B. Learning Paradigm

Let  $\mathcal{D} = \{\mathbf{u}^k, \mathbf{x}^k\}_{k=1}^{N_d}$  be a dataset of state-input pairs gathered from expert (human) demonstrations where  $N_d$  represents the total number of data points collected. Since the cost function  $Q_i$  can vary during data collection, it is possible for a state to have several associated expert inputs.

Our goal is find some responsibility allocation function  $\gamma$  such that the RA-CBF constraint is satisfied for all state-input pairs in the expert demonstrations  $\mathcal{D}$ . This can be written as the constrained optimization problem:

$$\gamma^* = \underset{\gamma}{\operatorname{argmin}} \|\gamma\| \quad (14)$$

s.t. RA-CBF Constraint( $i, \mathbf{x}, \mathbf{u}, \gamma$ )  $\geq 0, \forall i \in \{1, \dots, N\}$ ,

$$\sum_{i \in \{1, \dots, N\}} \gamma(i, \mathbf{x}) \geq 0, \quad \text{for all } (\mathbf{x}, \mathbf{u}) \in \mathcal{D},$$

where the constraints enforce satisfaction of the RA-CBF and ensure that  $\gamma$  is a responsibility allocation function.

To find an approximate solution to this problem we take inspiration from [19], [20] and relax (14) to the following unconstrained loss function:

$$\mathcal{L}(\mathcal{D}, \gamma) = \|\gamma\| + \lambda_1 \sum_{(\mathbf{x}, \mathbf{u}) \in \mathcal{D}} \sum_{i=1}^N \left[ -\mathbf{c}_i(\mathbf{x}, \mathbf{u}) + \gamma(i, \mathbf{x}) \right]_+ + \lambda_2 \sum_{(\mathbf{x}, \mathbf{u}) \in \mathcal{D}} \left[ \sum_{i=1}^N -\gamma_i(i, \mathbf{x}) \right]_+ \quad (15)$$

where  $\lambda_1, \lambda_2 \in \mathbb{R}_{\geq 0}$  are hyperparameters which adjust the constraint relaxations and  $[\cdot]_+ \triangleq \max\{\cdot, 0\}$ . This loss function can then be used find approximate solutions to (14):

$$\gamma^* \approx \underset{\gamma}{\operatorname{argmin}} \mathcal{L}(\mathcal{D}, \gamma) \quad (16)$$

### C. Responsibility Regularization

However, the loss function used in the unconstrained optimization (16) is insufficient since, as in Inverse Reinforcement Learning, the problem of learning the constraint in (13) is poorly defined since the optimal input generated by (13) is a function of both the unknown cost function  $Q_i$  and unknown responsibility allocation function  $\gamma$ . Intuitively, this is because we cannot answer the question ‘‘did the agent act that way because it wanted to (cost minimization) or because it had to (safety constraint satisfaction)?’’. To better define

the constraint learning problem we take an approach similar to [28] and regularize  $\gamma$  by maximizing the likelihood that it was used in (13) to generate the expert demonstrations  $\mathcal{D}$ .

Following the maximum entropy model presented in [29] with the variant for continuous-time nonlinear systems presented in [30] we wish to solve the optimization problem:

$$\gamma_{\text{reg}}^* = \underset{\gamma}{\operatorname{argmax}} \sum_{(\mathbf{x}, \mathbf{u}) \in \mathcal{D}} \mathcal{P}(\mathbf{u} | \mathbf{x}, \gamma). \quad (17)$$

We approximate the probability of a given  $\mathbf{u}$ , by choosing  $\text{disc}(\mathcal{U})$  to be a finite discretization of  $\mathcal{U}$  such as  $\text{disc}(\mathcal{U}) = \{\mathbf{u} \in \mathcal{U} \mid \delta \lfloor \mathbf{u}/\delta \rfloor\}$  for some  $\delta > 0$  where  $\lfloor \cdot \rfloor$  rounds each component to the nearest integer. Mimicking the forms presented in [30], [28], the approximate probability of an input  $\mathbf{u} \in \mathcal{U}$  given the system state  $\mathbf{x}$  and responsibility allocation  $\gamma$  is:

$$\mathcal{P}(\mathbf{u} | \mathbf{x}, \gamma) = \frac{e^{R(\mathbf{x}, \mathbf{u})}}{Z_\gamma} \mathbb{1}^\gamma(\mathbf{x}, \mathbf{u}), \quad Z_\gamma = \sum_{\mathbf{v} \in \text{disc}(\mathcal{U})} e^{R(\mathbf{x}, \mathbf{v})} \mathbb{1}^\gamma(\mathbf{x}, \mathbf{v}) \quad (18)$$

where  $Z$  is the partition function,  $R : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$  is the reward function, and  $\mathbb{1}^\gamma(\mathbf{x}, \mathbf{u}) \mapsto \{0, 1\}$  indicates satisfaction of the RA-CBF constraints given  $\mathbf{x}$ ,  $\mathbf{u}$ , and  $\gamma$ .

To maximize the likelihood of the demonstration we minimize the number of feasible inputs while retaining the feasibility of the expert demonstrations. We note the total number of feasible inputs in  $\text{disc}(\mathcal{U})$  decreases as  $\gamma(i, \mathbf{x})$  increases regardless of  $R$ , so we can maximize  $\mathcal{P}(\mathbf{u} | \mathbf{x}, \gamma)$  without knowledge of the agents’ reward functions by maximizing  $\gamma$  while maintaining feasibility of the expert demonstrations. This can be expressed as the optimization problem:

$$\gamma_{\text{reg}}^* \approx \underset{\gamma}{\operatorname{argmax}} \sum_{(\mathbf{x}, \mathbf{u}) \in \mathcal{D}} \sum_{i=1}^N \gamma(i, \mathbf{x}) \quad (19)$$

s.t. RA-CBF Constraint( $i, \mathbf{x}, \mathbf{u}, \gamma$ )  $\geq 0$ ,  
for all  $i \in \{1, \dots, N\}$  and  $(\mathbf{x}, \mathbf{u}) \in \mathcal{D}$

Since constraint feasibility on  $\mathcal{D}$  is already accounted for in (15), this regularization can be added to the loss as:

$$\mathcal{L}_{\text{reg}}(\mathcal{D}, \gamma) \triangleq \mathcal{L}(\mathcal{D}, \gamma) + \lambda_3 \sum_{(\mathbf{x}, \mathbf{u}) \in \mathcal{D}} \sum_{i=1}^N -\gamma(i, \mathbf{x}) \quad (20)$$

with hyperparameter  $\lambda_3 \in \mathbb{R}_{\geq 0}$  which can be used in the final regularized optimization problem to estimate  $\gamma$ :


$$\gamma^* = \underset{\gamma}{\operatorname{argmin}} \mathcal{L}_{\text{reg}}(\mathcal{D}, \gamma) \quad (21)$$

## V. APPLICATION TO AUTONOMOUS DRIVING

In this section we apply our RA-CBF and responsibility allocation learning method to urban driving using the Boston Seaport data provided by the nuScenes dataset [31].

We assume that all agents in the scene are vehicles (i.e. there are no pedestrians) and we model each agent as:

$$\begin{bmatrix} \dot{x}_i \\ \dot{y}_i \\ \dot{v}_i \\ \dot{\theta} \end{bmatrix} = \underbrace{\begin{bmatrix} v_i \cos(\theta_i) \\ v_i \sin(\theta_i) \\ 0 \\ 0 \end{bmatrix}}_{\mathbf{f}_i(\mathbf{x}_i)} + \underbrace{\begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}}_{\mathbf{g}_i(\mathbf{x}_i)} \underbrace{\begin{bmatrix} a_i \\ \omega_i \end{bmatrix}}_{\mathbf{u}_i}, \quad (22)$$



**Fig. 2.** The learned responsibility allocation surface is visualized for a range of velocities and relative positions. **Scene 1:** the ego vehicle (yellow) is driving on a two lane road. In all cases,  $\gamma(\text{ego}, \mathbf{x}) > 0$  indicating a degree of conservative driving. Generally,  $\gamma(\text{ego}, \mathbf{x})$  is larger when the other agent (red) is in front of the ego vehicle than when behind, indicating increased responsibility when driving behind another vehicle. **Scene 2:** The ego vehicle is stopped at a four-way intersection with the other vehicle (green) ahead or behind it (and no blue agent). Again the ego vehicle (yellow) is more responsible when the green vehicle is in front of it than when it is behind it. **Scene 3:** The ego vehicle (yellow) is stopped at a four-way intersection with the other vehicle (blue) crossing from top to bottom (and no green agent).  $\gamma(\text{ego}, \mathbf{x})$  is large for all positions and velocities of the blue vehicle showing that the ego agent takes is more responsible in this situation.

where  $(x_i, y_i) \in \mathbb{R}^2$ ,  $v_i, \theta_i, a_i, \omega_i \in \mathbb{R}$  represent the position, velocity, yaw, acceleration, and yaw rate<sup>3</sup> of vehicle  $i$ .

#### A. Choosing a Safety Metric for Autonomous Driving

To define the safety function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$ , we begin by assuming all vehicles must maintain a minimum inter-vehicle distance  $\underline{d} > 0$ . With this in mind, let  $d_{\min} : \mathbb{R}^4 \times \mathbb{R}^4 \rightarrow \mathbb{R}$  be the minimum distance between two agents. We can then define the pairwise safe set between agents  $i$  and  $j$  to be:

$$C_{ij} = \{ \mathbf{x} \in \mathbb{R}^n \mid \underbrace{d_{\min}(\mathbf{x}_i, \mathbf{x}_j) - \underline{d}}_{h_{ij}(\mathbf{x})} \geq 0 \}. \quad (23)$$

However this function is of relative degree 2 w.r.t.  $a_i$  (i.e.,  $\frac{dh}{dt}$  is not directly affected by  $a_i$ ) and describes safety by only considering the instantaneous current position.

In order to incorporate a temporal aspect and ensure that the time derivative of  $h_{ij}$  is affected by both vehicles' acceleration and angle rate we take inspiration from [32] and forward project the current state using a *backup controller*  $\mathbf{k}_B : \mathbb{R}^4 \rightarrow \mathbb{R}^2$  over a time interval  $[0, T]$  for  $T \in \mathbb{R}_{>0}$ . By assumption, for any  $\mathbf{x}_i(t) \in \mathbb{R}^4$  there exists a unique solution  $\phi : [0, T] \rightarrow \mathbb{R}^4$  satisfying:

$$\frac{d}{dt}\phi(\tau) = \mathbf{f}_i(\phi(\tau)) + \mathbf{g}_i(\phi(\tau))\mathbf{k}_B(\phi(\tau)), \quad (24)$$

$$\phi(0) = \mathbf{x}_i(t) \quad (25)$$

The solution  $\phi$  starting at  $\mathbf{x}_i(t)$  is the flow under  $\mathbf{k}_B$ , and is denoted as  $\varphi_\tau(\mathbf{x}_i) \triangleq \phi(\tau)$ . Similar ideas of forward projection are also seen in Velocity Obstacles [33], Safety Force Fields [3], and Responsibility-Sensitive Safety [5].

Using the flow  $\varphi$  and the distance function  $d_{ij}$  we can find the minimum distance that would be achieved during the interval  $[t, T+t]$  if  $\mathbf{k}_B$  were the controller for both vehicles:

$$h_{ij}^\varphi(\mathbf{x}) = \min_{\tau \in [0, T]} d_{\min}(\varphi_\tau(\mathbf{x}_i), \varphi_\tau(\mathbf{x}_j)) - \underline{d}, \quad (26)$$

which has the associated safe set  $C_{ij}^\varphi \subseteq C_{ij} \subseteq \mathbb{R}^n$

$$C_{ij}^\varphi = \{ \mathbf{x} \in \mathbb{R}^n \mid h_{ij}^\varphi(\mathbf{x}) \geq 0 \}. \quad (27)$$

To compute  $h_{ij}^\varphi$  the interval  $[0, T]$  was discretized at 100 Hz as in [32], [34] and soft minimum functions were used to ensure differentiability. It is shown in [35] that the CBF  $h_{ij}^\varphi$

<sup>3</sup>Bezier curves are fit to position and yaw data and then differentiated to obtain velocity, acceleration, and yaw rate. The code repository for learning the responsibility allocation function  $\gamma$  can be found [here](#).

constructed from the backup controller is guaranteed to be of relative degree 1 under mild assumptions.

For the backup controller we choose  $\mathbf{k}_B(\mathbf{x}_i) = \mathbf{0}$  which approximates idling. Unlike other methods [36], [5] which assume maximum braking for their predictors, we choose an idling controller since it better approximates nominal driving behavior and does not introduce worst-case assumptions.

Given these pairwise safe sets  $C_{ij}^\varphi$ , we can define the a global safe set  $C^\varphi \subseteq C_{ij}^\varphi \subseteq C_{ij}$  for all  $i \neq j$  as:

$$C_{ij}^\varphi = \bigcap_{i \neq j} C_{ij}^\varphi \quad \text{with} \quad h(\mathbf{x}) = \min_{i \neq j} h_{ij}^\varphi(\mathbf{x}). \quad (28)$$

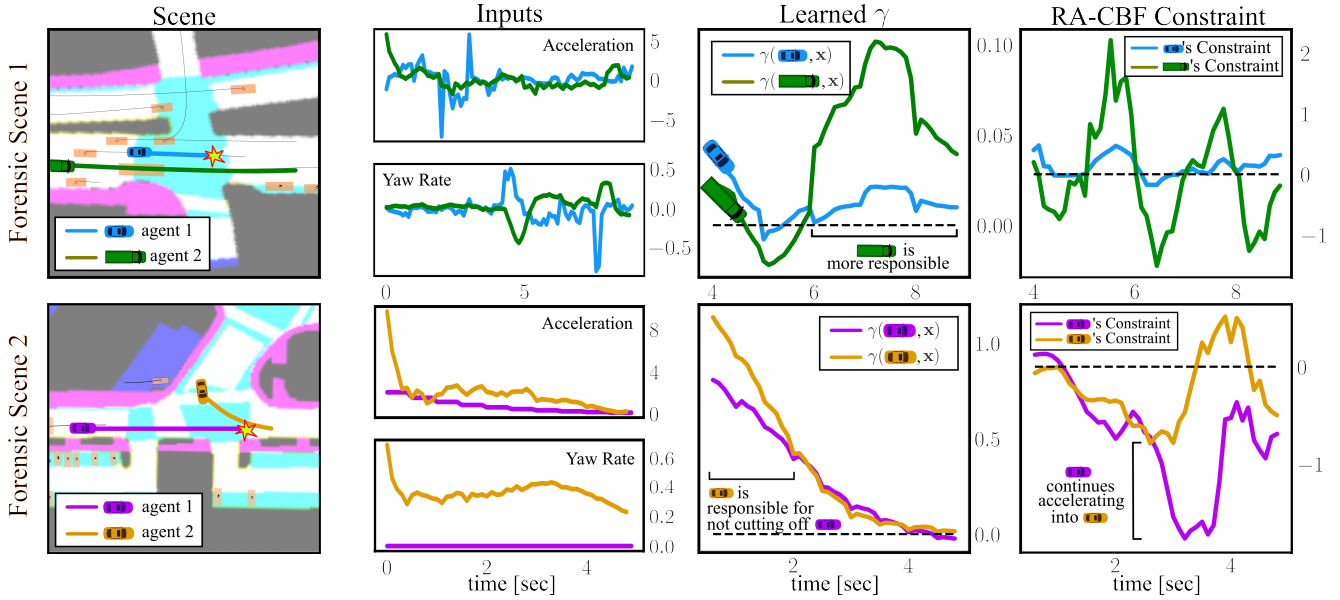
The intersection of safe sets has been studied in [10] and safety of such sets can be achieved by enforcing the safety constraint for all  $i \neq j$ .

#### B. Learning Model

The inputs of the responsibility allocation function are a semantic image as see in Fig. 2 and the relative vehicle states of agents  $i$  and  $j$ . The image is processed by ResNet-18 [37] and the 256 dimensional output is concatenated to the vehicle states and processed by a multi-layer perceptron (MLP) with 2 hidden layers of size 128 and a single dimensional output.

The hyperparameters chosen were  $\lambda_1 = 1$ ,  $\lambda_2 = 10$ ,  $\lambda_3 = 0.01$ ,  $\alpha = 0.5$ ,  $T = 1$ ,  $\underline{d} = 0.4$ ,  $\ell_1 = 0.1$ ,  $\ell_2 = 0.01$  and  $\theta_{\max} = 100^\circ$  where  $\ell_1$  and  $\ell_2$  were the negative slopes of the MLP's leaky ReLU activation functions and  $\theta_{\max}$  is used to filter the dataset such that only interactions between vehicles whose headings are within  $\pm\theta_{\max}$  are considered. The parameter  $\theta_{\max}$  is necessary since our data does not include lane direction annotation. We note that this does limit the applicability of this network and plan to include lane direction information in future work.

The network was trained on the NuScenes Boston Seaport dataset. Example responsibilities generated by our learned model can be found in Fig. 2. These figures show that our model conforms to the general intuition that the vehicle behind is more responsible than the vehicle in front for avoiding collisions between them, and the vehicle stopped at an intersection is responsible for not interfering with a vehicle already crossing the intersection.



**Fig. 3.** Crash scenes for forensic analysis. In both rows the plots from left to right display: the scene and trajectories, the inputs for agents 1 and 2, the learned responsibility allocations for each agent, and the values of their RA-CBF constraints (9).

Closed-Loop Simulation Results			
	Worst-Case	Even-Split	Our Method
Validation Constraint Violation	43.99%	8.13%	9.51%
Closed-Loop Safety Violation	0.833%	2.50%	0.833%
Time Spent Off Road	1.48%	0.59%	0.54%
Distance Covered Metric	290.21	309.21	307.84

**TABLE I.** Results for the closed-loop experiments.

### C. Closed-Loop Testing

We use our RA-CBF framework with a learned responsibility allocation function as a safety-filter in closed-loop control and simulate human-like driving using the Bi-Level Imitation for Traffic Simulation (BITS) model [38]. The ego agent follows (13) where  $Q_i(\mathbf{x}_i, \mathbf{u}_i) = \|\mathbf{k}_{\text{bits}}^+(\mathbf{x}_i) - \mathbf{u}_i\|^2$  and  $\mathbf{k}_{\text{bits}}^+$  is the BITS controller with an additional  $1 \frac{\text{m}}{\text{sec}^2}$  acceleration added to generate irresponsible desired behavior that must be filtered to ensure safety. The RA-CBF constraint is applied for each pairwise vehicle  $j$  and slack variables are used to ensure feasibility. We compare our method to the same controller with two other baseline constraints: (i) “Worst-Case” constraint (7), and (ii) “Even-Sharing” constraint which is the RA-CBF constraint with  $\gamma(i, \mathbf{x}) \equiv 0$ .

The closed-loop system was run in 120 scenarios sampled from NuScenes for 10 seconds at 10 Hz. Table (I) contains metrics comparing the controllers. The Worst-Case controller has the fewest safety violations (as expected), but worse compatibility with the expert demonstrations as indicated by the large constraint violation, smallest distance covered, and significant amount of time off of drivable surfaces. The Even-Sharing controller has fewer constraint violations on the validation data and the most distance covered, but allows for more collisions. Our method has a slightly higher number of constraint violations, but achieves a better safety-performance trade-off.

### D. Forensic Analysis

In addition to closed loop control, the values of  $\gamma$  and the RA-CBF constraint provide useful insight when performing forensic analysis on unsafe driving behaviors. To demonstrate

this, we analyze the two collision scenarios shown in Fig. 3. **Forensic Scenario 1:** In the first scene, agent 1 is crossing an intersection as agent 2 approaches from behind. For more than 2 seconds preceding the crash,  $\gamma(2, \mathbf{x}) > \gamma(1, \mathbf{x})$  indicating that agent 2 should have taken a greater share of the responsibility. However, agent 2 has several large violations of the RA-CBF constraint while agent 1 generally satisfies it. From this we interpret that agent 2 is responsible for the crash which aligns with our intuition.

**Forensic Scenario 2:** In the second scene, agent 1 is driving along a road that agent 2 is turning on to. At first, agent 2 is more responsible for not entering the lane in front of agent 1 (*i.e.*,  $\gamma(2, \mathbf{x}) > \gamma(1, \mathbf{x})$  for  $t < 2$ ). However, the responsibility allocation values then become similar as vehicle 2 enters the lane and accelerates. Throughout the scene, both agents violate safety, but agent 1 has much larger constraint violations in the moments preceding the crash due to its continued acceleration. This can be interpreted to indicate that although it was irresponsible of agent 2 to merge when it did, agent 1 is ultimately culpable.

We recognize that this analysis is subjective and one of many possible interpretations, but we believe that our method can provide useful insight when performing forensic analysis.

## VI. CONCLUSION

We have presented Responsibility-Aware Control Barrier Functions (RA-CBFs) as a framework to learn and synthesize safe and responsible driving behaviors. RA-CBFs are designed to capture the asymmetric sharing of responsibility between multiple (human) agents and we present a method to learn context-dependent responsibility allocations from data. We then demonstrated the efficacy and utility of our approach using real-world driving data. This work enables various exciting future directions which include incorporating explicit traffic rules into our responsibility-learning paradigm, comparing how responsibility allocations vary across geographical regions, and exploring other application domains such as crowd navigation.

## REFERENCES

- [1] "IEEE Standard for Assumptions in Safety-Related Models for Automated Driving Systems," *IEEE Std 2846-2022*, 2022.
- [2] 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*. Melbourne, Australia: IEEE, Dec. 2017, pp. 2242–2253. [Online]. Available: <http://ieeexplore.ieee.org/document/8263977/>
- [3] D. Nistér, H.-L. Lee, J. Ng, and Y. Wang, "The Safety Force Field," *NVIDIA White Paper*, 2019.
- [4] J. Usevitch and D. Panagou, "Adversarial resilience for sampled-data systems under high-relative-degree safety constraints," *IEEE Transactions on Automatic Control*, 2022.
- [5] S. Shalev-Shwartz, S. Shammah, and A. Shashua, "On a formal model of safe and scalable self-driving cars," *arXiv preprint arXiv:1708.06374*, 2017.
- [6] W. Schwarting, A. Pierson, J. Alonso-Mora, S. Karaman, and D. Rus, "Social behavior for autonomous vehicles," *Proceedings of the National Academy of Sciences*, vol. 116, no. 50, pp. 24972–24978, 2019.
- [7] B. Toghi, R. Valiente, D. Sadigh, R. Pedarsani, and Y. P. Fallah, "Co-operative autonomous vehicles that sympathize with human drivers," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [8] L. Sun, W. Zhan, M. Tomizuka, and A. D. Dragan, "Courteous autonomous cars," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018.
- [9] U. Borrmann, L. Wang, A. D. Ames, and M. Egerstedt, "Control barrier certificates for safe swarm behavior," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 68–73, 2015.
- [10] P. Glotfelter, I. Buckley, and M. Egerstedt, "Hybrid Nonsmooth Barrier Functions With Applications to Provably Safe and Composable Collision Avoidance for Robotic Systems," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1303–1310, Apr. 2019.
- [11] A. Mustafa and D. Panagou, "Adversary detection and resilient control for multi-agent systems," *IEEE Transactions on Control of Network Systems*, 2022.
- [12] J. van den Berg, M. C. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, 2008.
- [13] M. Chen, J. C. Shih, and C. J. Tomlin, "Multi-vehicle collision avoidance via hamilton-jacobi reachability and mixed integer programming," in *IEEE Conference on Decision and Control*, 2016.
- [14] X. Wang, K. Leung, and M. Pavone, "Infusing reachability-based safety into planning and control for multi-agent interactions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.
- [15] Y. Lyu, W. Luo, and J. M. Dolan, "Responsibility-associated multi-agent collision avoidance with social preferences," *arXiv preprint arXiv:2206.09030*, 2022.
- [16] K. Guo, D. Wang, T. Fan, and J. Pan, "VR-ORCA: Variable responsibility optimal reciprocal collision avoidance," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4520–4527, 2021.
- [17] Y. Chen, A. Singletary, and A. D. Ames, "Guaranteed obstacle avoidance for multi-robot operations with limited actuation: A control barrier function approach," *IEEE Control Systems Letters*, vol. 5, no. 1, pp. 127–132, 2020.
- [18] C. Dawson, S. Gao, and C. Fan, "Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods," *arXiv preprint arXiv:2202.11762*, 2022.
- [19] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," in *IEEE Conference on Decision and Control*, 2020, pp. 3717–3724.
- [20] Z. Qin, K. Zhang, Y. Chen, J. Chen, and C. Fan, "Learning safe multi-agent control with decentralized neural barrier certificates," *International Conference on Learning Representations*, 2021.
- [21] Y. Lyu, W. Luo, and J. M. Dolan, "Adaptive safe merging control for heterogeneous autonomous vehicles using parametric control barrier functions," *arXiv preprint arXiv:2202.09936*, 2022.
- [22] 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, Aug. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7782377/>
- [23] A. D. Ames, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs with application to adaptive cruise control," in *IEEE Conference on Decision and Control*, 2014, pp. 6271–6278. [Online]. Available: <http://ieeexplore.ieee.org/document/7040372/>
- [24] R. Cosner, M. Tucker, A. Taylor, K. Li, T. Molnar, W. Ubelacker, A. Alan, G. Orosz, Y. Yue, and A. Ames, "Safety-aware preference-based learning for safety-critical control," in *Learning for Dynamics and Control Conference*. PMLR, 2022, pp. 1020–1033.
- [25] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [26] P. Glotfelter, J. Cortes, and M. Egerstedt, "A Nonsmooth Approach to Controller Synthesis for Boolean Specifications," *IEEE Transactions on Automatic Control*, pp. 1–1, 2020.
- [27] A. J. Taylor, A. Singletary, Y. Yue, and A. D. Ames, "A control barrier perspective on episodic learning via projection-to-state safety," *IEEE Control Systems Letters*, vol. 5, no. 3, pp. 1019–1024, 2020.
- [28] D. R. Scobee and S. S. Sastry, "Maximum likelihood constraint inference for inverse reinforcement learning," *arXiv preprint arXiv:1909.05477*, 2019.
- [29] B. D. Ziebart, A. L. Maas, J. A. Bagnell, A. K. Dey, et al., "Maximum entropy inverse reinforcement learning," in *AAAI Conference on Artificial Intelligence*, vol. 8, 2008, pp. 1433–1438.
- [30] N. Aghasadeghi and T. Bretl, "Maximum entropy inverse reinforcement learning in continuous state spaces with path integrals," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 1561–1566.
- [31] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," in *IEEE / CVF Computer Vision and Pattern Recognition Conference*, 2020.
- [32] T. Gurriet, M. Mote, A. Singletary, P. Nilsson, E. Feron, and A. D. Ames, "A scalable safety critical control framework for nonlinear systems," *IEEE Access*, vol. 8, pp. 187 249–187 275, 2020.
- [33] D. Wilkie, J. Van Den Berg, and D. Manocha, "Generalized velocity obstacles," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 5573–5578.
- [34] R. K. Cosner, A. W. Singletary, A. J. Taylor, T. G. Molnar, K. L. Bouman, and A. D. Ames, "Measurement-Robust Control Barrier Functions: Certainty in Safety with Uncertainty in State," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021.
- [35] Y. Chen, M. Jankovic, M. Santillo, and A. D. Ames, "Backup control barrier functions: Formulation and comparative study," in *IEEE Conference on Decision and Control*, 2021, pp. 6835–6841.
- [36] D. Nister, O. Naroditsky, and J. Bergen, "Visual odometry," in *IEEE / CVF Computer Vision and Pattern Recognition Conference*, 2004.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015. [Online]. Available: <http://arxiv.org/abs/1512.03385>
- [38] D. Xu, Y. Chen, B. Ivanovic, and M. Pavone, "BITS: Bi-level imitation for traffic simulation," *arXiv preprint arXiv:2208.12403*, 2022.