

# Causality-Based Parametric Control Barrier Function for Safe Multi-Vehicle Interaction

Yiwei Lyu<sup>1\*</sup>, Caleb Chang<sup>2\*</sup> and John M. Dolan<sup>3</sup>

**Abstract**—Safe control has been widely studied in various safety-critical applications, for instance, autonomous driving. In order to ensure the autonomous vehicle does not collide with other vehicles, it is essential to obtain an accurate expectation of surrounding vehicles’ behavior and react adaptively. Instead of assuming fully cooperative and homogeneous vehicles using the same safety-critical controllers, recent works have been exploring different data-driven approaches to model the neighboring vehicles’ underlying controllers with observed data. However, existing works either suffer from 1) the inter-vehicle influence during the multi-vehicle interaction, which makes it hard to determine the causality of surrounding vehicles’ behavior in controller modeling, or 2) being dominated by the worst-case analysis, which may lead to overly conservative behavior. In this paper, we extend the prior work on Parametric-Control Barrier Function (Parametric-CBF) to multi-robot interactions with embedded causality inference to explicitly reason over the inter-vehicle influence. Given the learned Causality-based Parametric-CBF, we present an adaptive safety-critical controller that allows the ego vehicle to safely react to surrounding vehicles with the learned expectation. We demonstrate that by leveraging the motion flexibility among multi-vehicle systems, task efficiency can be greatly improved in various interaction-intensive scenarios.

## I. INTRODUCTION

In light of the need to operate in interactive environments with other scenario participants, e.g., other robots or humans, robots must ensure the safety of autonomy by correctly understanding others’ capabilities and planning their own strategies in a reactive manner. To achieve these goals, robots must generate accurate predictions of other agents and reasonable reactions in response. In the domain of autonomous driving, how to anticipate other vehicles’ behavior and leverage the anticipation to design controllers for ego vehicles while preserving formal safety guarantees has been a critical challenge.

There has been significant progress in trajectory prediction for road participants, primarily through data-driven approaches that predict the future movements of other vehicles [1], [2]. However, these predictive methods, which rely on current and past states, often fail to accurately capture how other robots will dynamically adapt to various conditions, such as the presence of obstacles, environmental changes, and interactions with other agents. This limitation

Pairwise Causal Relations between Agent 5 and Others

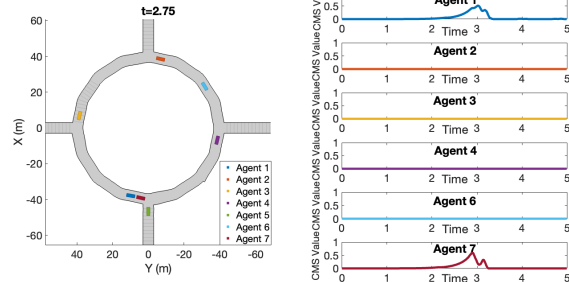


Fig. 1: Cross Map Smoothness (CMS)-based causality inference for target vehicle Agent 5 in a 7-vehicle roundabout scenario. Left: A snapshot of the interaction scenario. Right: The CMS scores for each pairwise interaction involving the target vehicle.

is particularly problematic in safety-critical scenarios, where understanding the interactive capabilities and possible reactions of other vehicles is crucial.

To address these challenges, learning the underlying controllers of other vehicles has emerged as a more effective approach [3], [4]. By gaining insights into the adaptive behaviors of robots in response to environmental changes, this class of methods enables more robust and flexible planning and decision-making. Furthermore, it allows us to explicitly model different driver characteristics underlying these behaviors, which can effectively facilitate proactive and expressive interactions when we design safe controllers for ego vehicles. Several attempts have been made to model the underlying controllers of other vehicles to better represent and interpret various driving styles [5], [6].

Despite recent advances, existing approaches face several fundamental limitations. A recurring challenge is the trade-off between computational efficiency and behavioral expressivity: models rich enough to capture diverse driving styles are often too expensive for real-time deployment, whereas simpler formulations cannot represent important behavioral nuances. Many methods, such as [7], attempt to learn the underlying controller and driver characteristics from direct observations of two-car interactions. These approaches typically rely on the assumption that all demonstrations are generated under saturated constraints, so that the observed actions directly reflect the enforcement of safety requirements rather than nominal control objectives.

While [8] relaxes this assumption by incorporating likelihood maximization of driver characteristics in expert demonstrations, it still presumes that the data come from either a

\* These authors contributed equally.

<sup>1</sup>The author is with the Department of Computer Science and Engineering, Texas AM University, College Station, TX, 77845 USA. Email: yiweilyu@tamu.edu

<sup>2</sup>The author is a visiting researcher from the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, 30332 USA. Email: calebwychang@gatech.edu

<sup>3</sup>The author is with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, 15213 USA. Email: jdolan@andrew.cmu.edu

single expert or a set of experts with consistent styles. In practice, however, driving data are often heterogeneous [9]; without a mechanism to evaluate which portions of the demonstrations are actually relevant, the resulting inaccuracies in learning can lead to unsafe or overly conservative behavior.

Moreover, much of this prior work has been limited to two-vehicle scenarios, leaving open questions about scalability to realistic, multi-agent environments. In such settings, the causal dependencies [10], [11] among interacting vehicles introduce an additional layer of complexity, making it difficult to determine when and why safety constraints are activated and which demonstrations are most informative for learning. Addressing these challenges is critical for developing models that can accurately and efficiently capture the interactive dynamics of heterogeneous multi-vehicle systems.

In this paper, our **main contributions** are, **1)** we introduce a causality detection algorithm to identify pair-wise interactions between vehicles on the road, enabling the robot to automatically reason over which portion of demonstrations should be used for learning without assumptions on constraint saturation and homogeneous experts, **2)** we extend prior work on Parametric-CBF [7] with embedded causality reasoning to enable learning of driving characteristics in a data-driven approach under motion uncertainty, and **3)** we present a comprehensive safety-critical control framework with our learning-enabled Parametric-CBF, which encourages courteous and efficient driving behavior while limiting perturbations to surrounding vehicles to improve overall task performance and collective safety.

## II. PRELIMINARIES AND NOTATIONS

### A. Background on CBFs and Parametric-CBFs

Control Barrier Functions (CBFs) [12] are used to guarantee safety in nonlinear systems (such as an affine nonlinear system) by the forward invariance of the safety set. The general nonlinear affine system is given by:

$$\dot{x} = f(x) + g(x)u + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma) \quad (1)$$

where  $x \in \mathcal{X} \subset \mathbb{R}^n$  and  $u \in \mathcal{U} \subset \mathbb{R}^m$  are the system state and control input with  $f$  and  $g$  assumed to be locally Lipschitz continuous. We introduce a Gaussian disturbance  $\epsilon$  to account for uncertain motions in the form of aleatoric uncertainty [13], [14]. A desired safe set  $x \in \mathcal{H}$  can be denoted by the Control Barrier Function  $h(x)$ :

$$\mathcal{H} = \{x \in \mathbb{R}^n : h(x) \geq 0\} \quad (2)$$

where  $h(x) < 0$  indicates a safety violation. A formal definition is presented below:

**Definition 1.** (Control Barrier Function) Given a dynamical system (1) and the set  $\mathcal{H}$  defined in (2) with a continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$ , then  $h$  is a control barrier function (CBF) if there exists a class  $\mathcal{K}$  function for all  $x \in \mathcal{X}$  such that  $\sup_{u \in \mathcal{U}} \{\dot{h}(x, u)\} \geq -\kappa(h(x))$ , where  $\dot{h}(x, u) = L_f h(x) + L_g h(x)u$  with  $L_f$  and  $L_g$  as the Lie derivatives of  $h$  along the vector fields  $f$  and  $g$  in the nonlinear affine system.

A commonly selected class  $\mathcal{K}$  function is  $\kappa(h(x)) = \gamma h(x)$  [12], [15], where  $\gamma \in \mathbb{R}_{\geq 0}$  is a design parameter that regulates system behavior near the boundary  $h(x) = 0$ . This leads to the admissible control set:  $\mathcal{B}(x) = \{u \in \mathcal{U} : \dot{h}(x, u) + \gamma h(x) \geq 0\}$ . As shown in [12], any control input  $u \in \mathcal{B}(x)$  renders the safe set  $\mathcal{H}$  forward invariant, i.e., if  $x(0) \in \mathcal{H}$ , then  $x(t) \in \mathcal{H}$  for all  $t > 0$ . However, the simple linear form  $\kappa(h(x)) = \gamma h(x)$  can be limiting when capturing complex system behavior near the boundary. Parametric Control Barrier Function (PCBF) framework [7] generalizes this formulation by effectively increasing the expressiveness of both  $\gamma$  and  $h(x)$ , enabling more nuanced safety constraints and adaptive system responses.

**Definition 2.** (Parametric-Control Barrier Function) Given a dynamical system (1) and the set  $\mathcal{H}$  defined in (2) with a continuously differentiable function  $h : \mathbb{R}^n \rightarrow \mathbb{R}$ , then  $h$  is a Parametric-Control Barrier Function (Parametric-CBF) for all  $x \in \mathcal{X}$  such that  $\sup_{u \in \mathcal{U}} \{\dot{h}(x, u)\} + \alpha H \geq 0$ , where parameter vector  $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_q] \in \mathbb{R}^q$  with  $\forall \alpha_p \in \mathbb{R}_{\geq 0}$  for  $p \in [q]$ .  $H(x) = [h(x), h^3(x), \dots, h^{2q-1}(x)]^T$ ,  $q \in \mathcal{N}$ .

As proved in [7], the Parametric CBF retains all the theoretical guarantees of the vanilla CBF.  $\alpha$  denotes the safety behavior vector, as it characterizes how rapidly the system approaches the boundary  $\partial\mathcal{B}$ . In the context of driving,  $\alpha$  reflects the aggressiveness of a driver's behavior. In Section III.C, we demonstrate how  $\alpha$  can be learned from data through a regression-based approach.

### B. Background on Causality

The study of causality seeks to explain the cause-and-effect relationships between two events. It identifies the cause-and-effect pair of events among many events and determines how strong the influence of the causal event is on the affected event. In this paper, we use the specific notion of Granger Causality (G-Causality) that associates predictability with causality in a time series setting [11]. Granger defines causality as follows [10]:

**Definition 3.** (Causality) Given a stationary stochastic system  $A^t$  and  $Hist(A^t)$ , which is the set of past values  $\{A^t - j, j = 0, 1, 2, \dots, \infty\}$ ,  $P^t(A|B)$  is the optimal predictor of  $A^t$  using the set of values  $B^t$ . The predictive error is  $\epsilon^t(A|B)$  and its variance is denoted by  $\sigma^2(A|B)$ . Let  $U^t$  be all the information in the universe since time  $t - 1$  and  $U^t - Y^t$  denote all the information apart from the time series  $Y^t$ . If  $\sigma^2(X|Hist(U)) < \sigma^2(X|Hist(U - Y))$ , we say that  $Y$  is causing  $X$ , denoted by  $Y^t \rightarrow X^t$ .  $Y^t$  is causing  $X^t$  if we can predict  $X^t$  using all available information  $U$ . Conversely, even if we can predict  $X^t$  using all available information  $U$  without  $Y^t$ , then a causal relationship does not stand.

Intuitively, two state variables are said to have a causal relationship if the state of one variable influences the dynamics of the other. In the case of  $Y \rightarrow X$ , the dynamics of  $X$  are sensitive to the state of  $Y$ . Conversely, if  $Y$  does not causally influence  $X$ , then the evolution of  $X$  remains unaffected by the state of  $Y$  [16]. Originally, Granger proposed methods

to identify causality in linear systems using cross-spectral analysis. Since then, researchers have applied causal inference techniques in areas such as ecology and biology. More recently, there have been efforts to extend these methods to nonlinear systems, deal with noise in dynamics, and improve the efficiency of causal detection using fewer data points [17], [18], [18]–[23].

### III. METHOD

#### A. Problem Statement

This work addresses multi-vehicle interactions at roundabouts through a “learn-and-adapt” framework, where the learning and control phases are separated. The goal is to infer the safety specification of other vehicles through observation and use that knowledge to improve cooperative driving performance. Given a set of demonstrations  $\mathcal{D} = \{x_n^t, u_n^t\} \in \mathbb{R}^{N \times T}$ , comprising state and control inputs from  $N$  heterogeneous vehicles over  $T$  time steps, our objectives are to:

- 1) **Learn safety specifications:** Accurately infer the underlying safety specification  $f(x, u, \alpha) \geq 0$  of the target vehicle from observational data, even in complex, multi-agent environments, where it may exhibit different interactive behaviors toward different agents.
- 2) **Design a safety-critical controller:** Based on the learned safety specification, develop a controller  $u_i$  for the ego vehicle that enables safe and efficient cooperative driving.

These objectives must address two core questions: (1) *Which parts of the observational data should be used for learning?* and (2) *How can the learned model be optimally leveraged to improve collective task performance?*

This paper relies on the following assumptions. First, known safety radius ( $R_{\text{safe}}$ ): We assume that the safety radius for each vehicle is known. This is a reasonable assumption, as safety margins can often be estimated based on vehicle-specific features (e.g., size, braking capacity) and environmental conditions (e.g., road type, speed limits, human reaction time). For example, the National Safety Council’s “3-Second Following Distance Rule” provides a widely accepted guideline for safe driving [24]. Second, CBF-based controllers for other vehicles: We assume that other vehicles are governed by CBF-controllers. While not all vehicles may implement exact CBF controllers, prior work [2] demonstrates that their behavior can often be approximated by CBF-based constraints. This assumption extends the applicability of our framework to a broader class of driving agents. Third, time-invariant behavior policies: We assume each vehicle operates under a policy that remains invariant over time. This reflects human driving behavior, where underlying traits (e.g., aggressiveness or risk aversion) tend to remain stable, even if specific actions vary with context. By focusing on persistent behavioral characteristics, we simplify the learning and prediction tasks.

In the following sections, we introduce the Cross Map Smoothness-based causality reasoning algorithm in Sec. III-B and present a causality-embedded behavior prediction

formulation in Sec. III-C. Finally, an impact-aware safety-critical control framework design is demonstrated in III-D by leveraging the learned safety specification to achieve efficient collaborative driving.

#### B. CMS-based Causality Reasoning

The motivating questions for causality inference are: (1) “Did the target vehicle act that way because it wanted to (cost minimization) or because it had to (safety constraint satisfaction)?” [8], and (2) “If safety constraints were the cause, which neighboring vehicle triggered the activation of those constraints?”

To address these, we must identify when the target vehicle’s safety constraint  $f(x, u, \alpha) = \dot{h}(x, u) + \alpha H \geq 0$ , particularly the Parametric CBF constraint  $\dot{h}(x, u) + \alpha H = 0$ , is activated, even without explicit knowledge of the safety behavior vector  $\alpha$ . This involves pinpointing a subset of the demonstration  $\mathcal{D}$  corresponding to times when the safety constraint is active. Recognizing that the activation of safety constraints indicates preventive actions by the target vehicle to maintain safety, we analyze the causal relationship between observed control signals of the target vehicle and its safety-related state evaluations during its interactions with surrounding vehicles.

We adopt the Cross Map Smoothness (CMS) method to assess the strength of causal relationships between two variables by comparing their time series. CMS is part of a wider family of algorithms known as Convergent Cross-Mapping, which is typically used in fields like genomics, biology, and ecology for detecting and quantifying causal interactions in complex, nonlinear systems [19], [23], [25]–[27]. The underlying theory behind it is that if variable  $X$  causes  $Y$ , the states of  $Y$  should be able to reconstruct  $X$ ’s trajectory, demonstrating a smooth relationship between them over time. Our CMS assessment incorporates three key components: “time-delayed observations,” “delay embedded vectors,” and “reconstructed attractors” [28]. We retain standard terminology from causality inference literature and provide definitions in our study’s context.

For each pairwise relationship between the target vehicle  $j$  and other surrounding vehicles  $k \in N \setminus j$ , we denote “time-delayed observations” [23] of the control signal and states at time  $t$  as  $u_j(t)$  and  $h_{jk}(t)$ .  $u_j(t)$  is the  $\ell^2$ -norm of the observed control signal  $u$ , i.e., observed velocity, of the target vehicle at time  $t$ .  $h_{jk}(t) = \|x_j^t - x_k^t\|^2 - R_{\text{safe}}^2$  is the safety function characterizing the desired safe distance to maintain the pairwise inter-vehicle safety, calculated with the states of the target vehicle  $j$  and its pairwise companion  $k$ .

For a given receding horizon  $L$ , delayed-embedding vectors are therefore defined as

$$\mathbf{u}_j(t) = \begin{bmatrix} u_j(t) \\ u_j(t-1) \\ \vdots \\ u_j(t-L+1) \end{bmatrix}, \mathbf{h}_{jk}(t) = \begin{bmatrix} h_{jk}(t) \\ h_{jk}(t-1) \\ \vdots \\ h_{jk}(t-L+1) \end{bmatrix} \quad (3)$$

which is a collection of past observations in reverse chronological order. For the entire observation horizon  $T$ , the reconstructed attractors aggregate the delayed-embedded vectors

together, denoted as  $M_u = \{\mathbf{u}_j(t), \mathbf{u}_j(t-1), \dots, \mathbf{u}_j(t-T)\}$  and  $M_h = \{\mathbf{h}_{jk}(t), \mathbf{h}_{jk}(t-1), \dots, \mathbf{h}_{jk}(t-T)\}$  respectively.

According to the definition of causality, if  $u_j(t)$  causally influences  $h_{ij}(t)$ , then the state of  $u_j(t)$  is a driving factor in the dynamics of  $h_{jk}(t)$ , meaning that information about  $u_j(t)$  is reflected in  $h_{jk}(t)$ 's dynamics. Conversely, if  $u_j(t)$  does not causally influence  $h_{jk}(t)$ , then the dynamics of  $h_{jk}(t)$  remain insensitive to changes in  $u_j(t)$ . The construction of reconstructed attractors  $M_u$  and  $M_h$  from time series  $u_j(t)$  and  $h_{jk}(t)$ , respectively, is crucial to evaluate the extent to which the state information of  $u_j(t)$  is embedded within the dynamics of  $h_{jk}(t)$ . Specifically, we assess the smoothness of the mapping between  $M_u$  and  $M_h$ . If there exists a smooth cross map from  $M_h$  to  $M_u$ , it suggests that  $u_j(t)$  causally influences  $h_{jk}(t)$ .

We quantify the smoothness of this cross map using the error of a Radial Basis Function Neural Network (RBFNN) [23], which is tasked with learning the mapping between attractors. Since RBFNNs are capable of approximating any smooth function, a low error implies a smooth mapping and thus strong causality between  $u_j(t)$  and  $h_{jk}(t)$ . Conversely, a high error suggests that the mapping is not smooth, indicating little to no causality between the two variables. Readers are referred to [19], [23], [28] for more details.

Here we present our CMS-based Causality Reasoning Algorithm described in Algorithm 1. For brevity, we refer to  $\mathbf{u}(t-s)$  as  $\mathbf{u}_s$  and  $\mathbf{h}(t-s)$  as  $\mathbf{h}_s$ . The input to the algorithm

---

**Algorithm 1:** Cross Map Smoothness-based Causal Reasoning

---

**Data:**  $\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_n \in \mathbb{R}^L$  and  $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_n \in \mathbb{R}^L$ , let  $S_m = \{1, 2, \dots, n\} \setminus m$  be the leave-one-out index set.  
**Result:** Causality Index:  $R_{hu} \in [0, 1]$

- 1 **for**  $m = 1, 2, \dots, n$  **do**
- 2     train radial basis function neural network  $\mathcal{N}_m$  based on  $\mathcal{N}_m(\mathbf{u}_s) = \mathbf{h}_s, s \in S_m$
- 3      $\hat{\mathbf{h}}_m \leftarrow \mathcal{N}_m(\mathbf{u}_m)$
- 4      $\epsilon_m \leftarrow \|\mathbf{h}_m - \hat{\mathbf{h}}_m\|$
- 5 Normalize the error:  $\Delta \leftarrow \frac{\text{rms}(\epsilon)}{\text{rms}(\|\mathbf{h} - \hat{\mathbf{h}}\|)}$   
 /\*  $\sigma$  is a positive constant used to normalize  $R_{hu}$ . \*/
- 6 Causality Index:  $R_{hu} \leftarrow \frac{1}{\exp(\Delta/\sigma)}$

---

is then the corresponding sets of time-delayed coordinate vectors of size  $n$ :  $\{\mathbf{h}_1, \dots, \mathbf{h}_n\}$  and  $\{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ . These sets are the reconstructed attractors described above. Leave-one-out-sets  $S_m = \{1, 2, \dots, n\} \setminus m$  are then used to train an RBFNN  $n$  times to determine the network error, which is inversely proportional to the causality index. The output of the CMS algorithm is a continuous scale from 0 to 1 quantifying various levels of causality, where 1 indicates strong causal influence and 0 indicates none.

### C. Behavior Modeling with CMS-enabled Parametric CBF

Upon establishing paired dynamics, we employ linear ridge regression to predict safety behaviors. The CMS-based causality reasoning offers a key advantage over existing approaches like [7] by relaxing the assumption that the exact activation instances of the CBF constraint must be known. With CMS, we can accurately detect when the vehicle's

constraint is inactive, that is, when the CBF constraint  $f(x, u, \alpha) \geq 0$  is not enforced. This flexibility allows us to transform the original constrained optimization problem, as shown in Eq. 4, into an unconstrained optimization problem. In this new formulation, constraint satisfaction is managed by CMS-based causality inference, which selectively uses observations of  $h_{jk}^t$  corresponding to time instances when the constraint is identified as active, facilitating accurate parameter estimation:

$$\begin{aligned} \bar{\alpha}_j &= \arg \min_{\alpha_j} \sum_{t=1}^m \|\dot{h}_{jk}^t - \alpha_j H_{jk}^t\|_2^2 + r \|\alpha_j\|_F^2 \\ \text{s.t. } & \dot{h}_{jk}^t(x, u) + \alpha_j H_{jk}^t(x) \geq 0 \end{aligned} \quad (4)$$

where  $r$  is the regularizer parameter,  $\|\alpha_j\|_F$  is the Frobenius norm of  $\alpha_j$ , and  $\bar{\alpha}_j$  is the estimated solution.

To address the noise in observations, we leverage the theoretical properties of the class  $\kappa(\cdot)$  functions and the forward invariance of CBFs to impose additional bounds on the parameter estimation to characterize those likely invalid inference results, tightening the estimation for better accuracy. Specifically, we introduce: 1) Monotonicity Condition: By definition, the function  $\kappa(h(x)) = \alpha H(x)$  must be strictly increasing, which requires all elements of  $\alpha$  to be non-negative. Thus, if  $\min \bar{\alpha} < 0$ , this estimate of  $\alpha$  is invalid. 2) Constraint Satisfaction: The inequality  $\dot{h}(x, u) + \alpha H(x) + \delta_c \geq 0$  must hold for the true parameter, where  $\delta_c$  is a small value very close to 0 to allow for slackness under noisy motion. Therefore, if the left hand side is less than 0, this estimate of  $\alpha$  is also invalid. 3) Interaction Condition: For two vehicles to be considered interacting, the constraint should remain close to zero, indicating a near-boundary condition of the safety set. 4) Temporal Consistency: we impose a fourth condition that the current estimate aligns closely with previous estimates, under the assumption that the safety behavior parameter of the vehicle remains time-invariant. These conditions help refine parameter estimates and maintain robustness against noise, ensuring that only feasible solutions are retained. After identifying and removing those invalid estimates, the rest of the estimation of  $\alpha$  is stored in a dataset  $\mathcal{D}_{jk}$  to be used to obtain the final estimate value by averaging across this dataset. The pseudocode of the proposed CI-enabled Parametric CBF estimation is described in Algorithm 2. In lines 1-6, we use CMS to test if there is an active constraint between all pairs of vehicles using Algorithm 1. If the causality index is detected to be above a certain threshold, we begin the safety behavior estimation algorithm in line 7. A typical mean squared error minimization problem (Eq. 4 without constraints) can be solved analytically.

$$-\begin{bmatrix} H^T(t_1) \\ \vdots \\ H^T(t_m) \end{bmatrix}^T \begin{bmatrix} H^T(t_1) \\ \vdots \\ H^T(t_m) \end{bmatrix} \alpha = \begin{bmatrix} H^T(t_1) \\ \vdots \\ H^T(t_m) \end{bmatrix}^T \begin{bmatrix} \dot{h}(t_1) \\ \vdots \\ \dot{h}(t_m) \end{bmatrix} \quad (5)$$

Since we are unsure if the surrounding vehicles actually take actions that maintain their safety behavior constraints, we opt for an online sequential solver that satisfies the 4 conditions mentioned above instead of performing the large matrix inverse operation required to solve for  $\alpha_j$  in Eq. 5.

Lines 9-11 estimate a new  $\alpha_j$  based on the current  $H_{jk}$  and  $\hat{h}_{jk}$  values. In line 12, estimation validity is verified through the additional bounds we introduced based on the monotonicity condition and constraint satisfaction. Line 16 considers valid estimates of  $\alpha_j$  and checks for interaction activation ( $\hat{f}(x, u, \alpha_{new}) \approx 0$ ) and temporal consistency ( $\epsilon < \delta_{rmse}$ ). If the valid estimate passes these conditions as well, it is added into the dataset. Averaging the values in the dataset returns the final estimated value of  $\alpha_j$ , concluding the behavior modeling of the target vehicle.

**Algorithm 2:** Causality-embedded CBF Parameter Estimation

---

**Data:**  $\Delta x_{jk}, \Delta v_{jk}, u_j, u_k, \Delta t, R_{safe}$   
**Result:**  $\bar{\alpha}_j$

```

1 while True do
2   Add  $\{x_n^t, u_n^t\}$  to  $\mathcal{D} \forall n \in N$ 
3   if  $t > L$  then
4     For each vehicle pair (j, k):
5        $\mathcal{R}_{hu} = \text{CMS}(\mathbf{h}_{t-L-1}, \dots, \mathbf{h}^t, \mathbf{u}_{t-L-1}, \dots, \mathbf{u}^t)$ 
6       if  $\mathcal{R}_{hu} > \epsilon_R$  then
7         Jump to line 7
7 Initialize  $A_{sum}$  and  $B_{sum}$  to 0
8 for  $t = T_1 : T_2$  do
9   Calculate  $A^t \leftarrow H_{jk}^t H_{jk}^{tT}$  and  $B^t \leftarrow H_{jk}^t \hat{h}_{jk}^t$ 
10  Update  $A_{sum} \leftarrow A_{sum} + A^t$  and  $B_{sum} \leftarrow B_{sum} + B^t$ 
11  Estimate  $\alpha_{new} \leftarrow [-A_{sum}^{-1} B_{sum}]^T$ 
12  Calculate estimated constraint:
13   $\hat{f}(x, u, \alpha_{new}) \leftarrow \hat{h}_{jk}^t(x, u) + \alpha_{new} H_{jk}^t(x)$ 
14   $\epsilon_{rmse} \leftarrow \text{rmse}(\alpha_{new}, \alpha_{old})$ 
15  if  $|\hat{f}(x, u, \alpha_{new})| > \delta_c$  or  $\min \alpha_{new} < 0$  then
16    Reinitialize  $A_{sum}$  and  $B_{sum}$  to  $A^t$  and  $B^t$ 
17  else if  $\hat{f}(x, u, \alpha_{new}) \approx 0$  and  $\epsilon < \delta_{rmse}$  then
18    Add  $\alpha_{new}$  to dataset  $\mathcal{D}_{jk}$ 
19     $\alpha_{old} \leftarrow \alpha_{new}$ 
19  $\bar{\alpha}_j \leftarrow \text{average}(\mathcal{D}_{jk})$ 

```

---

#### D. Impact-Aware Control Strategy

In this section, we demonstrate how to leverage the learned behavior of the target vehicle to enhance collaborative driving. Specifically, our objective is to improve collective efficiency with smooth driving interactions by maximizing the comfort of the target vehicle’s motion and minimizing any jerkiness induced by the presence of the ego vehicle. We present our approach as the following optimization problem, and we refer to the ego vehicle with index  $i$  and the target vehicle with index  $j$ :

$$\begin{aligned}
u_i^t = \arg \min_{u_i} & c_1 \|u_i^t - \bar{u}_i^t\|^2 \\
& + c_2 \min\{\hat{h}_{ij}^{t_f}(x, u) + \bar{\alpha}_j H_{ij}^{t_f}(x), 0\} \hat{h}_{ij}^t(x, u) \quad (6) \\
\text{s.t. } & \hat{h}_{ij}^t(x, u) + \alpha_i H_{ij}^t(x) \geq 0
\end{aligned}$$

where the output  $u_i^t$  represents the actual control signal of the ego vehicle at the current timestep  $t$  and  $\bar{\alpha}_j$  represents the learned safety behavior. The time derivative  $\hat{h}_{ij}^{t_f}(x, u)$  and parametric CBF  $H_{ij}^{t_f}(x)$  are calculated at the future timestep  $t_f$ . We argue that these values can be reasonably calculated due to the advances in learning models for vehicle behavior prediction [1]. In addition, for our experiments, we demonstrate that even a simple piece-wise constant velocity model,  $x^{t_f} = x^t + (t_f - t)\dot{x}^t$ , can suffice.

The objective function consists of two positive terms with  $c_1$  and  $c_2$  as corresponding coefficients that represent relative

priority. The first term instructs the system to follow the nominated control input  $\bar{u}_i^t$ . The second term maximizes the current  $\hat{h}_{ij}^t(x, u)$  according to the severity of the learned future safety violation  $f(x^{t_f}, u^{t_f}, \bar{\alpha}) = \hat{h}_{ij}^{t_f}(x, u) + \bar{\alpha}_j H_{ij}^{t_f}(x)$ . If the constraint is not violated, i.e.,  $f(x^{t_f}, u^{t_f}, \bar{\alpha}_j) \geq 0$ , then the second term becomes 0 due to the  $\min\{\cdot, 0\}$  operator. Conversely, if the constraint is violated, then  $c_2 \min\{\hat{h}_{ij}^{t_f}(x, u) + \bar{\alpha}_j H_{ij}^{t_f}(x), 0\}$  is a negative multiplier, incentivizing  $\hat{h}_{ij}^t(x, u)$  to be maximized in the optimization problem. Greater violations are reflected in greater negative multipliers, resulting in higher relative priorities to maximize  $\hat{h}_{ij}^t(x, u)$ , thus minimizing the impact the ego vehicle has on the target.

#### IV. EXPERIMENTS AND DISCUSSION

In this section, we first present the performance of the CMS-enabled Parametric CBF for behavior modeling in multi-vehicle interaction scenarios. Then we demonstrate how collaborative efficiency could be improved via our impact-aware control strategy design.

##### A. CMS-enabled Parametric CBF for Behavior Modeling

1) *Causality Detection in Multi-Vehicle Interactions:* The scenario we selected is a 7-vehicle roundabout, as shown in Fig. 1 (left), which represents a typical challenging urban driving scenario due to the need for intensive inter-vehicle interactions. The ego vehicle’s objective is to navigate the roundabout safely while optimizing for collective efficiency. In this single-lane setting, the ego vehicle must focus on safely merging by interacting with a single target vehicle. However, no prior information is available to predict the target vehicle’s movements. Therefore, the ego vehicle must model the target’s behavior by observing its interactions with neighboring vehicles and inferring its safety specifications. We denote the seven vehicles, excluding the ego vehicle, as  $s_1, \dots, s_7$ .

We apply our proposed CMS-based causality reasoning algorithm to this 7-vehicle interaction scenario, with results shown in Fig. 1 (right). Using noisy observations of all vehicles’ past trajectories as inputs, the algorithm performs causality detection for each vehicle pair involving the target vehicle  $s_5$  and all other vehicles, illustrating the strength of cause-and-effect relationships that evolve over time. As expected, not all surrounding vehicles influence the movement of the target vehicle. In fact, as indicated by the CMS value plots, only  $s_1$  and  $s_7$  show a strong causal relationship with  $s_5$ ’s behavior during specific time periods, indicating that  $s_5$ ’s safety constraint is activated only between these vehicle pairs during these particular periods.

2) *Causality-Embedded Behavior Modeling:* Using the CMS-based causality inference algorithm, we identify the most relevant portion of data from the full set of observations and supply this selective subset to our parametric CBF estimation procedure. Figure 2a illustrates this process: in the left panel, the red markers represent the inferred data pairs from an example trial, while the yellow grid shows the learned constraint hyperplane. The close alignment of the markers with the hyperplane indicates that the CMS-selected

Trial Results				
Method	Average MSE (all)	Std Dev. MSE (all)	Average MSE (95%)	Std Dev. MSE (95%)
Ours w/CMS	0.3267	3.1323	0.0023	0.0041
NBF w/CMS	0.6905	3.7423	0.1678	0.1477
Parametric-CBF [7]	0.3692	0.2121	0.3454	0.1893
NBF w/o CMS	0.5832	0.2754	0.5531	0.2481

TABLE I: Accuracy comparison with existing approaches

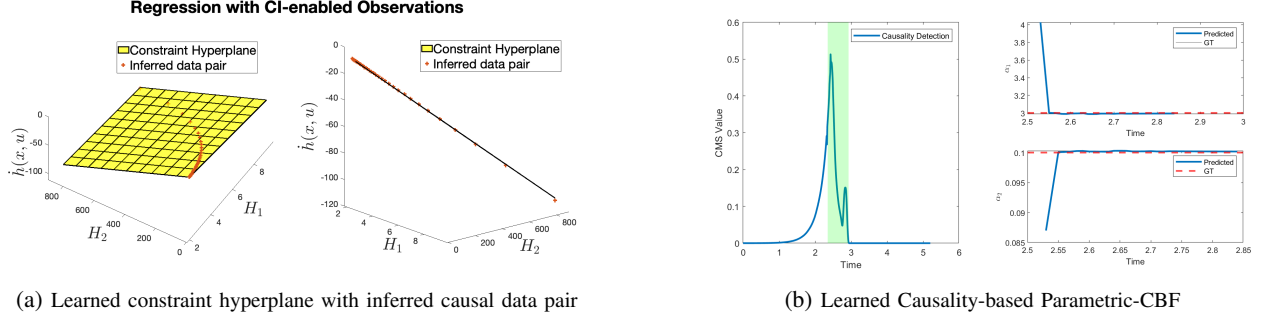


Fig. 2: CI-enabled Constraint Learning.

samples are consistent with the estimated constraint and thus faithfully capture the causal structure of the interaction. In the right panel, the same data and hyperplane are shown edge-on, further validating this consistency by demonstrating that all inferred data pairs lie directly on the hyperplane. The learned hyperplane encodes the target vehicle’s underlying safety specification, effectively characterizing the boundary conditions that govern its behavior. This representation offers an interpretable and compact model of how the vehicle enforces safety in the observed scenario. Then we demonstrate the safety specification learning for behavior modeling in Fig. 2b. By identifying the portion of data with causality marked in a green background, we show on the right that our CI-enabled Parametric CBF is able to quickly converge to the ground truth values marked as red dashed lines.

To further emphasize the effectiveness of our approach, we compare it to a popular class of approaches for learning control barrier function parameters [8], [3], [4], [6], which turns CBF hard constraints into soft constraints by reformulating them into the objective function. The original parameter learning problem is as follows, where functions  $h(x)$  and  $\kappa(h(x))$  are parameterized by  $\theta$  and  $\alpha$ .


$$\begin{aligned}
\alpha^*, \theta^* &= \min_{\alpha, \theta} |\theta| + |\alpha| \\
\text{s.t. } & \dot{h}_\theta(x, u) + \kappa_\alpha(h_\theta(x)) \geq 0 \quad \forall (x, u) \in \mathcal{D}_{\text{safe}} \\
& h_\theta(x) \geq 0 \quad \forall (x, u) \in \mathcal{D}_{\text{safe}} \\
& \dot{h}_\theta(x, u) + \kappa_\alpha(h_\theta(x)) < 0 \quad \forall (x, u) \in \mathcal{D}_{\text{unsafe}} \\
& h_\theta(x) < 0 \quad \forall (x, u) \in \mathcal{D}_{\text{unsafe}}
\end{aligned} \tag{7}$$

Here,  $\mathcal{D}_{\text{safe}}$  and  $\mathcal{D}_{\text{unsafe}}$  represent  $(x, u)$  pairs that satisfy  $h(x) > 0$  for the safe set and  $h(x) < 0$  for the unsafe set. Conditioning on  $h_\theta(x)$  being known and  $\kappa_\alpha(h_\theta(x)) = \alpha H(x)$ , [8], [3], [4], [6] reformulated the optimization problem as:  $\alpha^* = \min_{\alpha} |\alpha| + \sum_{(x, u) \in \mathcal{D}_{\text{safe}}} \max(0, -(\dot{h}_\theta(x, u) + \alpha H(x)))$ . Since this optimization loss is typically used to train Neural Barrier Functions, we address this method as NBF in Table I.

We compare the results of their solution with our proposed method as shown in Algorithm 2 in the roundabout scenario. In this scenario, the target vehicle interacts with other vehicles, and the ego vehicle perceives these interactions through noisy observations. However, no information is provided regarding when and which vehicles are actively interacting. Over 200 trials, with motion noise sampled from  $\mathcal{N}(0, 1)$ , the target vehicle uses different values of  $\alpha$  to induce varying safety behaviors.


Table I presents the results of the two approaches. The left two columns show the average and standard deviation of the mean squared error (MSE) between the true  $\alpha$  value and the predicted value across all 200 trials for both methods. The right two columns report the same statistics computed over the 95% of trials with the lowest MSE. This additional metric highlights that Algo. 2 performs exceptionally well in the majority of cases. However, there are rare instances where the algorithm’s performance degrades, due to fixed values of  $\delta_c$  and  $\delta_{rmse}$  that are not tuned for individual trials. In comparison, the original ridge regression method [7] demonstrates consistent performance across all trials, including the 95% subset, but does not consistently achieve lower error due to the presence of observation noise. From the comparison results, we conclude that our approach outperforms existing methods, with the effectiveness of CMS demonstrated by a significantly higher average accuracy across the majority of test cases.

Our approach achieves superior performance primarily due to two key factors: (1) strong robustness to localization uncertainty in the observed data, and (2) the ability to identify data segments with a strong causal relationship to the safety constraint being learned. While all agents aim to act safely, motion uncertainty can lead to interactions that violate safety constraints, as verified by the Control Barrier Function conditions. Existing methods typically rely on clear labels to distinguish safe from unsafe data. However, in



(a) Collision due to infeasibility of conventional CBF-QP (b) Safe merging with our impact-aware control (c) Control deviation (left) and inter-vehicle distance (right) over time

Fig. 3: Performance comparison in ramp merging



(a) Deadlock due to over conservatism of conventional CBF-QP (b) Safe merging between Agents 1 and 7 with our impact-aware control (c) Control deviation (left) and inter-vehicle distance of all vehicles (right) over time

Fig. 4: Performance comparison in roundabout merging

uncertain settings, such labels are often unavailable or unreliable—leading these methods to treat all observations as safe, which compromises robustness in parameter learning. Even when labels are available, these methods fail to isolate the truly informative segments of data, leading to overestimation compared to ground truth values. In contrast, our approach leverages causal inference to pinpoint portions of data where safety constraints are actively engaged. This targeted learning results in higher overall accuracy.

### B. Impact-Aware Control

We argue that our proposed impact-aware control has two advantages: 1) improving the collaborative driving efficiency, and 2) enhancing the collective safety in these interactions. To demonstrate the effectiveness of our approach, we validated it on two scenarios: ramp-merging on highway driving (one-on-one interaction) and a roundabout in urban driving (multi-agent interactions, 7 agents in total excluding ego). In both scenarios, 100 trials were conducted with ego vehicles starting from uniformly sampled positions away from the merging point (between 38 and 46m). In the ramp scenario, the other vehicle is also uniformly sampled from the same range, while the roundabout scenario has surrounding vehicles starting in preset positions. These surrounding vehicles are simulated with conventional CBF-QP safe controllers with a nominal speed of 20m/s. For evaluation, we analyze the following metrics: the total time to complete the collab-

orative driving task, the speed deviation from the nominal controllers, and the inter-vehicle distance.

In the ramp merging scenario, the mean and standard deviation of the total time to complete the driving task over 100 trials are 3.42s and 1.34s with our proposed approach, while the mean and standard deviation using the conventional CBF-QP method are 4.24s and 1.54s. In the roundabout scenario, the mean and standard deviation completion times are 9.77s and 3.60s with our approach, while the CBF-QP yielded 10.26s and 3.83s. The 19.34% and 4.87% improvement in task efficiency is primarily due to the impact-aware design of our method. By promoting more courteous behavior, the ego vehicle not only optimizes for its own objectives but also considers the potential disruption it may cause to other agents, thereby minimizing overall control deviation. This design becomes particularly advantageous in more complex scenarios, where mutual consideration greatly improves coordination and system performance. As shown in Fig. 3(c), during one-on-one interactions, our approach causes significantly less disruption to the other vehicle’s trajectory compared to the conventional CBF-QP, avoiding abrupt reactions and enabling smoother interaction.

In some trials, the conventional CBF-QP executed by other vehicles may yield infeasible solutions due to noisy observations. The standard fallback strategy in such cases is to apply an emergency brake and come to a full stop, as illustrated in Fig. 3(a). In contrast, our impact-aware

approach enables the ego vehicle to anticipate and mitigate these situations. By explicitly quantifying the feasibility space in the near future of the other vehicle using our learned Parametric CBF, the ego vehicle can proactively plan preventative actions with lookahead capabilities when a potential infeasibility is foreseen, as shown in Fig. 3(b). This advantage is further supported by Fig. 3(c), where the inter-vehicle distance under the conventional CBF-QP dips below the safety threshold (black dashed line), indicating a safety violation. Similar results are observed in the roundabout scenario, as shown in Fig. 4. In Fig. 4(a), using the conventional CBF-QP, a deadlock occurs because both the ego vehicle and Agent 7 plan their actions independently, without accounting for each other. In contrast, Fig. 4(b) demonstrates our impact-aware design, where the ego vehicle proactively considers the admissible action space of both parties. By accelerating to reduce ambiguity, it nudges Agent 7 forward, creating sufficient space for itself while avoiding danger and preventing deadlock. This results in a smoother collaborative outcome, where all vehicles continue progressing rather than being forced to stop. Fig. 4(c) shows that, in comparison, our approach consistently maintains a safe distance above the margin, improving the collective safety, while causing much less control deviation for others, reflecting a more courteous and efficient behavior paradigm.

## V. CONCLUSION

In this work, we propose a novel Causality Inference-enabled Parametric Control Barrier Function framework for safe multi-agent interactions. We demonstrate that safety-critical control constraints can be effectively learned from data under relaxed assumptions regarding data relevance and observational uncertainty. Building on the learned safety representations, we design an impact-aware safe controller that improves both task efficiency and collective safety in urban and highway driving scenarios. Future directions include extending the method to systems with high relative degree and validating its performance through real-world experiments.

## REFERENCES

- [1] S. Mozaffari, O. Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakis, "Deep learning-based vehicle behavior prediction for autonomous driving applications: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 1, pp. 33–47, Jan. 2022.
- [2] J. Grover, Y. Lyu, W. Luo, C. Liu, J. Dolan, and K. Sycara, "Semantically-aware pedestrian intent prediction with barrier functions and mixed-integer quadratic programming," *IFAC-PapersOnLine*, vol. 55, no. 41, pp. 167–174, 2022.
- [3] Z. Qin, K. Zhang, Y. Chen, J. Chen, and C. Fan, "Learning safe multi-agent control with decentralized neural barrier certificates," in *Proceedings of the 9th International Conference on Learning Representations (ICLR)*, 2021.
- [4] Y. Yang, L. Chen, Z. Zaidi, S. van Waveren, A. Krishna, and M. Gombolay, "Enhancing safety in learning from demonstration algorithms via control barrier function shielding," in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, Mar. 2024, pp. 820–829.
- [5] M. Saveriano and D. Lee, "Learning barrier functions for constrained motion planning with dynamical systems," in *2019 IEEE International Conference on Intelligent Robots and Systems (IROS)*, pp. 112–119.
- [6] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," in *2020 59th IEEE Conference on Decision and Control (CDC)*, Dec. 2020, pp. 3717–3724.
- [7] Y. Lyu, W. Luo, and J. M. Dolan, "Adaptive safe merging control for heterogeneous autonomous vehicles using parametric control barrier functions," in *2022 IEEE Intelligent Vehicles Symposium (IV)*, June 2022, pp. 542–547.
- [8] R. K. Cosner, Y. Chen, K. Leung, and M. Pavone, "Learning responsibility allocations for safe human-robot interaction with applications to autonomous driving," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9757–9763.
- [9] J. Yarlagadda and D. S. Pawar, "Heterogeneity in the driver behavior: An exploratory study using real-time driving data," *Journal of advanced transportation*, vol. 2022, no. 1, p. 4509071, 2022.
- [10] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.
- [11] A. Shojaie and E. B. Fox, "Granger causality: A review and recent advances," *Annual Review of Statistics and Its Application*, vol. 9, pp. 289–319, 2022.
- [12] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *2019 18th European Control Conference (ECC)*. EUCA, 2019.
- [13] E. Hüllermeier and W. Waegeman, "Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods," *Machine Learning*, vol. 110, no. 3, pp. 457–506, Mar. 2021.
- [14] "From aleatoric to epistemic: Exploring uncertainty quantification techniques in artificial intelligence," 2025.
- [15] Y. Lyu, W. Luo, and J. M. Dolan, "Probabilistic safety-assured adaptive merging control for autonomous vehicles," in *2021 IEEE International conference on robotics and automation (ICRA)*. IEEE, 2021, pp. 10764–10770.
- [16] M. Ding, Y. Chen, and S. L. Bressler, "Granger causality: Basic theory and application to neuroscience," 2006.
- [17] W. Hesse, E. Möller, M. Arnold, and B. Schack, "The use of time-variant eeg granger causality for inspecting directed interdependencies of neural assemblies," *Journal of Neuroscience Methods*, vol. 124, no. 1, pp. 27–44, Mar. 2003.
- [18] S. L. Bressler and A. K. Seth, "Wiener-granger causality: A well-established methodology," *NeuroImage*, vol. 58, pp. 323–329, 2011.
- [19] G. Sugihara, R. May, H. Ye, C.-h. Hsieh, E. Deyle, M. Fogarty, and S. Munch, "Detecting causality in complex ecosystems," *Science*, vol. 338, no. 6106, pp. 496–500, Oct. 2012.
- [20] S. Sinha and U. Vaidya, "On data-driven computation of information transfer for causal inference in discrete-time dynamical systems," *Journal of Nonlinear Science*, vol. 30, no. 4, pp. 1651–1676, 2020.
- [21] N. Ancona, D. Marinazzo, and S. Stramaglia, "Radial basis function approach to nonlinear granger causality of time series," *Physical Review E*, vol. 70, no. 5, p. 056221, Nov. 2004.
- [22] M. Lungarella, K. Ishiguro, Y. Kuniyoshi, and N. Otsu, "Methods for quantifying the causal structure of bivariate time series," *International Journal of Bifurcation and Chaos*, vol. 17, pp. 903–921, Mar. 2007.
- [23] H. Ma, K. Aihara, and L. Chen, "Detecting causality from nonlinear dynamics with short-term time series," *Scientific Reports*, 2014.
- [24] National Safety Council, *The DDC Instructor and Administrative Reference Guide*, 2019.
- [25] E. H. van Nes, M. Scheffer, V. Brovkin, T. M. Lenton, H. Ye, E. Deyle, and G. Sugihara, "Causal feedbacks in climate change," *Nature Climate Change*, vol. 5, no. 5, pp. 445–448, May 2015.
- [26] E. R. Deyle, M. Fogarty, C.-h. Hsieh, L. Kaufman, A. D. MacCall, S. B. Munch, C. T. Perretti, H. Ye, and G. Sugihara, "Predicting climate effects on pacific sardine," *Proceedings of the National Academy of Sciences*, vol. 110, no. 16, pp. 6430–6435, Apr. 2013.
- [27] S. B. Munch, A. Giron-Nava, and G. Sugihara, "Nonlinear dynamics and noise in fisheries recruitment: A global meta-analysis," *Fish and Fisheries*, vol. 19, no. 6, pp. 964–973, Dec. 2018.
- [28] F. Takens, "Detecting strange attractors in turbulence," in *Dynamical Systems and Turbulence, Warwick 1980*, ser. Lecture Notes in Mathematics, D. Rand and L.-S. Young, Eds. Springer Berlin Heidelberg, 1981, vol. 898, pp. 366–381.
- [29] J. Zeng, B. Zhang, and K. Sreenath, "Safety-critical model predictive control with discrete-time control barrier function," 2020.
- [30] S. He, J. Zeng, B. Zhang, and K. Sreenath, "Rule-based safety-critical control design using control barrier functions with application to autonomous lane change," 2021.
- [31] L. Wang, A. Ames, and M. Egerstedt, "Safety barrier certificates for heterogeneous multi-robot systems," in *2016 American Control Conference (ACC)*, 2016, pp. 5213–5218.