

Physical Backdoor Trigger Activation of Autonomous Vehicle using Reachability Analysis

Wenqing Li¹, Yue Wang², Muhammad Shafique³ and Saif Eddin Jabari⁴

Abstract—Recent studies reveal that Autonomous Vehicles (AVs) can be manipulated by hidden backdoors, causing them to perform harmful actions when activated by physical triggers. However, it is still unclear how these triggers can be activated while adhering to traffic principles. Understanding this vulnerability in a dynamic traffic environment is crucial. This work addresses this gap by presenting physical trigger activation as a reachability problem of controlled dynamic system. Our technique identifies security-critical areas in traffic systems where trigger conditions for accidents can be reached, and provides intended trajectories for how those conditions can be reached. Testing on typical traffic scenarios showed the system can be successfully driven to trigger conditions with near 100% activation rate. Our method benefits from identifying AV vulnerability and enabling effective safety strategies.

I. INTRODUCTION

Recent works (e.g., [1], [2]) has shown that an adversary can insert hidden backdoors into AV controllers, allowing them to cause accidents by triggering these backdoors through well-designed traffic system states. These states are valid combinations of velocities and positions of AV and other vehicles, and can be activated by tampering with malicious vehicles. However, the current method for finding these states through simulation-based empirical trials is inefficient and trivial, especially when the attacker does not have guiding information to steer the system towards the desired trigger state.

This paper addresses the lack of guidance for activating backdoor triggers in interactive traffic systems. The approach recasts trigger activation as a reachability problem, where the motions of critical vehicles are characterized as a controlled dynamic system. The goal becomes driving the system to reach trigger conditions from any admissible conditions, which is achieved through the computation of backwards reachable sets and intended trajectories.

The Hamilton-Jacobi (HJ) reachability method is a popular streaming method for reachability analysis, as described in previous works [3], [4]. This method computes the reachable set by explicitly solving the Hamilton Jacobi Isaacs Partial

Differential Equation (HJI PDE) [5], with the reachable set represented by the zero sub-level set of the solution. However, applying the HJ reachability method to our problem faces major obstacles. Specifically, the system dynamics are unknown due to the AV being controlled by a learning model, i.e., a deep reinforcement learning (DRL) model, which prevents the formulation of the HJI PDE based on first principles. Additionally, the computational cost of the HJ reachability method increases significantly for nonlinear systems with high dimensions, typically those with more than five dimensions (5D) [6]. As such, the application of HJ reachability is limited for our problem. To address the above issues, we learn a surrogate linearized dynamic model using the state data from simulations which measures system behaviors. Our contributions are threefold:

- 1) We frame the problem of physical trigger activation as a data-driven reachability problem of a controlled dynamic system, where we use reachability analysis to determine the conditions under which a specific event or state can be reached;
- 2) We learn a surrogate linearized dynamic model based on finite-dimensional Koopman operator approximation, where the linearization of the dynamic system allows for easy decoupling and efficient computation. We also demonstrate that the estimation error of BRS is bounded;
- 3) Our method is verified using a microscopic traffic simulator on typical traffic scenarios, and the results showcase that it could successfully drive the traffic system (operating in the security-critical areas) to trigger conditions (activation rate is approaching 100%).

Our technique offers benefits beyond understanding the vulnerability of the AV based traffic system to malicious attacks. It also enables the implementation of efficient defense strategies.

II. NOTATIONS AND PRELIMINARIES

We denote \mathbb{R} and \mathbb{N}_+ as the set of all real numbers and the set of all positive integers. We use $|\cdot|$ to describe the cardinality of a set, and use $\|\cdot\|$ to represent the l_2 -norm for vector and l_F -norm for matrix. The bold-face upper letters are used to represent matrices and the bold-face lower letters denote (column) vectors.

A. Backdoor attacks on AV controller

AV controllers are developed using deep reinforcement learning (DRL) and operate by taking in traffic system states

¹Wenqing Li is a Post-doc associate in the Division of Engineering, New York University Abu Dhabi, Saadiyat Island 129188, UAE w154@nyu.edu

²Yue Wang is a Ph.D. Student in the Department of Electrical Engineering, Tandon school of Engineering, New York University, Brooklyn, NY 11201, USA yw3576@nyu.edu

³Muhammad Shafique is a faculty member with Department of Electrical Engineering, Tandon School of Engineering, New York University, Brooklyn, NY, USA muhhammad.shafique@nyu.edu

⁴Saif Eddin Jabari is a faculty member with Department of Civil and Urban Engineering, Tandon School of Engineering, New York University, Brooklyn, NY, USA sej7@nyu.edu

as input and outputting continuous command actions. However, these controllers can be manipulated by an adversary who injects backdoors into the system. These backdoors are designed to make the AV behave normally under normal conditions but to perform malicious actions when exposed to adversary-chosen inputs. One effective method for injecting backdoors into AV controllers is through data poisoning, where the attacker adds trigger data into the genuine training data. This assumes that the attacker has access to the genuine data and can manipulate it accordingly. The difficulty in detecting these triggers lies in their stealthiness. They are designed to be challenging to detect and, as a result, can cause significant damage before being identified. A *trigger sample* is a valid combination $\{(v_{AV}, v_j^{adv}, \Delta d_{AV})\}$ where j is the leader of the AV, v_{AV} and v_j^{adv} are the velocities of the AV and its leader, $\Delta d_{AV} = d_j^{adv} - d_{AV}$ is the relative distance between the AV and its leader. The trigger samples are typically designed based on traffic physics (see [2]). In this paper, we consider the following backdoor attack discussed in the literature on the traffic control system:

- 1) *Insurance Attack*: the adversary-controlled leader vehicle aggressively decelerates. The adversary then aims to trigger the follower AV to accelerate and crash into the vehicle in front. The goal of the adversary (i.e., the leader of AV) is to make insurance claims from the AV company.

In order to launch such an attack, the adversary controls the malicious human-driven vehicle to activate the triggers, which are combinations of specified speed and position.

B. Controlled Dynamic System and Reachable Set

A controlled dynamic system describes the evolution of the variables over time is governed by an ordinary differential equation (ODE)

$$\begin{aligned} \frac{dx(\tau)}{d\tau} &= f(x(\tau), u(\tau)), \\ \tau &\in \mathcal{T}, u(\tau) \in \mathcal{U}. \end{aligned} \quad (1)$$

where $x(\tau) \in \mathbb{R}^n$, $u(\tau) \in \mathcal{U}$ are state and control input, and $f : \mathbb{R}^n \times \mathcal{U} \times \mathcal{T} \rightarrow \mathbb{R}^n$ describes system dynamics which is assumed to be Lipschitz continuous in x for fixed u . Given input function $u(\cdot) \in U$ (where U is set of measurable functions), there exists unique trajectory solving equation (1) (please see [7]). We will denote solutions, or trajectories of (1) starting from state x at time t and end at time τ under control $u(\cdot)$ as $\zeta(\tau; x, t, u(\cdot)) : [t, \tau] \rightarrow \mathbb{R}^n$. ζ satisfies (1) with an initial condition almost everywhere:

$$\begin{aligned} \frac{d(\zeta(\tau; x, t, u(\cdot)))}{d\tau} &= f(\zeta(\tau; x, t, u(\cdot)), u(\tau)), \\ \zeta(\tau; x, t, u(\cdot)) &= x(\tau). \end{aligned} \quad (2)$$

Given a dynamical system described by eq. (1), the *backwards reachable set* is defined as follows,

Definition 1: Backwards reachable set (BRS) in time t

$$\mathcal{R}(t) := \{x : \exists u(\cdot) \in U, \zeta(0; x, t, u(\cdot)) \in \mathcal{G}_0\} \quad (3)$$

with \mathcal{G}_0 being a target set. It represents the set of states that could lead to unsafety within a specified time horizon.

III. PROBLEM FORMULATION: PHYSICAL TRIGGER ACTIVATION AS A REACHABILITY PROBLEM

A. Problem Statement

To analyze the physical activation of triggers in vehicles, we model their motion as a controlled dynamic system. This allows us to formulate the activation of triggers as a reachability problem. Specifically, we aim to compute the backward reachable set (BRS) of a related dynamic system, which has the form $\dot{x} = f(x, u)$, where x is the system state and u is the control input. The target set \mathcal{G}_0 is defined as the set of trigger conditions that we want to activate, and the optimal control $u(\cdot)$ is the input sequence that steers the system into the target set.

We represent the target set \mathcal{G}_0 as the zero-sublevel set of an implicit surface function $g_0(x)$, i.e., $x \in \mathcal{G}_0 \Leftrightarrow g_0(x) \leq 0$, which is a signed distance function to target set \mathcal{G}_0 . The cost function is defined to be the signed distance of the terminal state ($\tau = 0$) to \mathcal{G}_0 :

$$J(x, t, u(\cdot)) = g_0(x) \quad (4)$$

If the $J(\cdot)$ is non-positive, the target set \mathcal{G}_0 is reached at $\tau = 0$; otherwise, it is not. The computation of BRS aims to find all states that starts from $\tau = t$ and enter \mathcal{G}_0 at $\tau = 0$ under the optimal control (from malicious vehicle), therefore resulting in the following optimal control problem:

$$\inf_{u(\cdot) \in U} \{J(x, t, u(\cdot))\} \quad (5)$$

where we denote $V(x, t) = \inf_{u(\cdot) \in U} \{J(x, t, u(\cdot))\}$ as the value function characterizing the reachable set, i.e., $\mathcal{R}(t) := \{x : V(x, t) \leq 0\}$. Hence, given the value function $V(x, t)$, we can judge whether a state x in the BRS ($\mathcal{R}(t)$) of target set (\mathcal{G}_0). The optimal control is then given by

$$u^*(\cdot) = \arg \inf_{u(\cdot) \in U} \nabla_x V(x, t)^\top f(x, u) \quad (6)$$

where $\nabla_x V(x, t)$ represents the gradient of value function with respect to x .

B. Dynamic models and target set

In this work, we focus on a basic scenario where a group of vehicles is running on a single-lane track. Specifically, we consider the case where one of the vehicles is an autonomous vehicle (AV) that is susceptible to backdoor attacks from an adversary, while the remaining vehicles are human-driven. Our analysis focuses on the motions of the AV and its leader, which is a malicious vehicle that is manipulated by the adversary. The state variable vector is $x = [v_{AV}, v_j^{adv}, \Delta d_{AV}]^\top \in \mathbb{R}^3$, where v_{AV}, v_j^{adv} are velocities of the AV, malicious vehicle controlled by the adversary, and Δd_{AV} denotes the relative distance between AV and malicious vehicle. The dynamics among three vehicles are given by the flow field $f : \mathbb{R}^3 \times \mathcal{U} \rightarrow \mathbb{R}^3$, and f is Lipschitz continuous with regards to state variable x . Besides, f is unknown since the dynamics of AV are unknown (AV is controlled by a DRL model).




Fig. 1. Illustration of vehicle model with one attacked AV its leader that is manipulated by an adversary. d_{AV} denotes the position of AV and d_j^{adv} is the position of AV's leader (a malicious vehicle).

For simplicity we drop the subscript and denote the state vector by $x = [x_1, x_2, x_3]^T$. Physical considerations impose constraints on the continuous state and inputs

$$\begin{aligned} \mathcal{X} &:= \{x : x_1, x_2 \in [0, v^{\max}], x_3 \in [\Delta d^{\min}, \Delta d^{\max}]\} \\ \mathcal{U} &:= \{u : u \in [a^{\min}, a^{\max}]\} \end{aligned} \quad (7)$$

where v^{\max} is an upper bound on all speeds of vehicles in free-flow condition, Δd^{\min} and Δd^{\max} are the minimum and maximum distances between adjacent vehicles, a^{\min} and a^{\max} are the lower and upper bound of the acceleration of the malicious vehicle.

The target set \mathcal{G}_0 is given by:

$$\mathcal{G}_0 := \{x \in \mathcal{X} : g_0(x) \leq 0\} \quad (8)$$

where

$$g_0(x) = x_2 t_{adv} + x_3 - (x_1 t_{adv} + a_{adv} t_{adv}^2) \quad (9)$$

where $g_0(x) \leq 0$ denotes the condition that under predefined malicious actions (a_{adv}), e.g., acceleration, AV will cover a distance in a short time ($x_1 t_{adv} + a_{adv} t_{adv}^2$) that is larger than the spacing between itself and its leader ($x_2 t_{adv} + x_3$), causing a collision. Here, we assume the malicious vehicle keep its velocity (without deceleration or acceleration) as the time interval is very short, for example, 1s.

IV. DATA-DRIVEN REACHABILITY ANALYSIS BASED ON KOOPMAN OPERATOR

We learn a surrogate dynamic model from sampled data of states and inputs based on Koopman operator. Consider the controlled dynamic system in (1), Koopman operator aims to learn a linearized system given below:

$$\frac{dz(\tau)}{d\tau} = \mathbf{A}z(\tau) + \mathbf{B}u(\tau) \quad (10)$$

where z denotes the lifted state variable in high dimensionality space, and u is the unlifted control input. \mathbf{A}, \mathbf{B} are constant matrices.

A. Koopman Operator for controlled system and its finite-dimensional approximation

Consider the dynamic system in (1) that describes the motions of three vehicles, a natural way to extend Koopman operator from uncontrolled system to controlled system is to extend the state-space \mathcal{X} to augmented state-space $\mathcal{X} \times \mathcal{U}$. Hence, Koopman operator on controlled system is the one on uncontrolled system evolving on the augmented state-space. Let $\hat{x} = [x, u]^T$ be the augmented state, and the flow in the augmented state is \hat{f} , the associated Koopman operator $\mathcal{K} : \mathcal{F} \rightarrow \mathcal{F}$ is defined by

$$\mathcal{K}\Phi(\hat{x}(\tau)) = \Phi(\hat{f}(\hat{x}(\tau))) = \frac{d\Phi(\hat{x}(\tau))}{d\tau} \quad (11)$$

where \mathcal{F} is a space of functions (typically referred to as observables), and $\Phi : \mathbb{R}^n \times \mathcal{U} \rightarrow \mathcal{F}$ denotes observables that can be considered as feature maps.

The Koopman operator defined in (11), i.e., \mathcal{K} , is an infinite-dimensional linear operator. However, it can be approximated by a linear finite dimensional operator. For the controlled dynamic system, we choose observables Φ in the following structure [8], [9]

$$\Phi(\hat{x}) = [\Psi(x)^T \ u^T]^T \quad (12)$$

where $x \in \mathbb{R}^n$ denotes the state, $u \in \mathbb{R}^m$ is the control input, $\Psi(x) = [\psi_1(x), \dots, \psi_N(x)]^T \in \mathbb{R}^N$ is a vector-valued function, and each $\psi_i : \mathbb{R}^n \rightarrow \mathbb{R}$ represents a scalar-valued function. To this end, we have the following finite-dimensional approximation operator $\mathbf{K} \in \mathbb{R}^{(N+m) \times (N+m)}$

$$d \begin{bmatrix} \Psi(x(\tau)) \\ u(\tau) \end{bmatrix} / d\tau \approx \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \cdot & \cdot \end{bmatrix} \begin{bmatrix} \Psi(x(\tau)) \\ u(\tau) \end{bmatrix} \quad (13)$$

$=: \mathbf{K}$

This immediately leads to the linearized system in the form of (10)

$$\frac{d\Psi(x(\tau))}{d\tau} \approx \mathbf{A}\Psi(x(\tau)) + \mathbf{B}u(\tau) \quad (14)$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$, $\mathbf{B} \in \mathbb{R}^{N \times m}$ can be learned from data sampled from the continuous-time state data, i.e., $\{\Psi(x(t_i + \Delta t))\}_{i=1}^M$, $\{\Psi(x(t_i))\}_{i=1}^M$, $\{x(t_i)\}_{i=1}^M$ and $\{u(t_i)\}_{i=1}^M$ with M measurements and sampling interval $\Delta t \ll 1$ (Notice that by letting $\Delta t \ll 1$, we can assume the dynamics of (14) is equivalent to its first-order time discretization). Specifically, the estimators $\hat{\mathbf{A}}, \hat{\mathbf{B}}$ can be obtained by solving the following optimization problem

$$\min_{\mathbf{A}, \mathbf{B}} \sum_{i=1}^I \|\Psi(x(t_i + \Delta t)) - \mathbf{A}\Psi(x(t_i)) - \mathbf{B}u(t_i)\|_2^2 \quad (15)$$

They are least squares problems that can be readily solved using linear algebra. The solutions are

$$[\mathbf{A} \ \mathbf{B}] = \mathbf{G}^+ \mathbf{Q} \quad (16)$$

where

$$\begin{aligned} \mathbf{G} &= \frac{1}{M} \sum_{i=1}^M [\Psi(x(t_i))^\top \Psi(x(t_i)) + u(x(t_i))^\top u(x(t_i))] \\ \mathbf{Q} &= \frac{1}{M} \sum_{i=1}^M \Psi(x(t_i + \Delta t))^\top \Psi(x(t_i + \Delta t)) \end{aligned} \quad (17)$$

B. Reachability analysis based on the approximation of Koopman operator

1) *Target set formulation:* The target set \mathcal{G}_0^Ψ in the new coordinates $\Psi(x)$ is given by

$$\mathcal{G}_0^\Psi := \{\Psi(x) \in \mathcal{X}_\Psi : g_\Psi(\Psi(x)) \leq 0\}. \quad (18)$$

where $g_\Psi(\Psi(x))$ is the implicit surface function. For the linearized system (14), We choose $\varphi_i(x)$ as the basis of monomials with total degrees equal to d

$$\varphi_i(x) \in \{x_1^{d_1} x_2^{d_2} x_3^{d_3} : (d_1, d_2, d_3) \in \mathbb{N}, \sum_{i=1}^3 d_i = d\} \quad (19)$$

the insight behind this is the flow map F is assumed to be polynomial. The number of functions in basis is $N = \frac{1}{2}(d+1)(d+2)$ such that each monomials corresponds to a $\psi_i(x)$.

Rewrite \mathcal{G}_0 as

$$\mathcal{G}_0 = \{x \in \mathcal{X} : g_0(x) = h(x) - c \leq 0\} \quad (20)$$

with some constant c . Then, the implicit surface function of target set \mathcal{G}_0^Ψ is given by

$$g_\Psi(\Psi(x)) := \sum_{i=1}^N w_i \psi_i(x) - c^d = (h(x))^d - c^d \quad (21)$$

By choosing this specific observables and an odd d , we have

$$x \in \mathcal{G}_0 \Leftrightarrow \Psi(x) \in \mathcal{G}_0^\Psi, \quad (22)$$

2) *BRS computation:* For the linearized system in (14), the BRS in t time is then defined to be

$$\begin{aligned} \mathcal{R}_\Psi(t) &:= \{\Psi(x) : \exists u \in \mathcal{U}, \Psi(x_0) \in \mathcal{G}_0^\Psi, \tau \in [t, 0], \\ &d\Psi(x(\tau))/d\tau = \mathbf{A}\Psi(x(\tau)) + \mathbf{B}u(\tau)\} \end{aligned} \quad (23)$$

The computation of BRS in time t results in the following optimal control problem:

$$V_\Psi(\Psi(x), t) = \inf_{u(\cdot) \in \mathcal{U}} \{J(\Psi(x), t, u(\cdot))\} \quad (24)$$

where $J(\Psi(x), t, u(\cdot)) = g_\Psi(x)$. Consequently, the BRS can be represented by

$$\mathcal{R}_\Psi(t) := \{\Psi(x) : V_\Psi(\Psi(x), t) \leq 0\} \quad (25)$$

Based on the principle of dynamic programming, it can be shown that the value function $V_\Psi(\Psi(x), t)$ is the viscosity solution of the following Hamilton-Jacobi Isaacs (HJI) PDE

$$\frac{\partial V(\Psi(x), t)}{\partial t} + H(\Psi(x), \nabla_\Psi V(\Psi(x), t)) = 0 \quad (26)$$

with $V(\Psi(x), 0) = g_\Psi(\Psi(x))$, and $H(\Psi(x), \nabla_\Psi V(\Psi(x), t))$ representing the Hamiltonion is given by

$$\begin{aligned} H(\Psi(x), \nabla_\Psi V(\Psi(x), t)) &= \\ &(\nabla_\Psi V(\Psi(x), t))^\top (\mathbf{A}\Psi(x) + \mathbf{B}u) \end{aligned} \quad (27)$$

The optimal control that steers the system into the target set is then obtained by

$$u^*(\cdot) = \arg \inf_{u(\cdot) \in \mathcal{U}} \nabla_\Psi V_\Psi(\Psi(x), t)^\top (\mathbf{A}\Psi(x) + \mathbf{B}u) \quad (28)$$

A family of algorithms called level set methods has been designed specifically to compute approximations to the viscosity solution (which is approximated by a Cartesian grid of the state space.) For details of these methods, please see [10]. Here, we use the publicly available implementation ‘TOOLBOXLS’ developed by [5].

C. Performance analysis based on Koopman operator approximation

The estimation of BRS is essentially backwards prediction of system states from the target set. Hence, the estimation accuracy is dependent on the generalization error of the system model.

Our main result: Given the system (14) learned using M training measurements and observables $\Psi = (\psi_i)_{i=1}^N$, and let R_k be the generalization error propagated from t_0 to $t_0 + k\Delta t$. If \mathbf{A} is γ -stable for $\gamma \in (0, 1)$ [11], i.e, all eigenvalues are no larger than γ , $\exists \sigma > 0$, such that $\|R_k\| \leq \sigma$ for $k \in \mathbb{N}_+$, where σ is a constant irrelevant with prediction step k .

The above statement implies that the propagation of generalization error will be suppressed by finding a stable matrix \mathbf{A} , and hence the long-term prediction accuracy (a.k.a., BRS estimation accuracy) of (14) is stable (does not drop dramatically).

We begin with some notations, we drop Δt and denote $t_0 + k\Delta t$ as $t_0 + k$ for simplicity, and denote $\tilde{\Psi}$ as the estimated observables from measurements, the generalization error at $t_0 + k$ based on the stable solution is given by

$$\begin{aligned} r_k &:= \Psi(x_{t_0+k}) - \tilde{\Psi}(x_{t_0+k}) \\ &= \Psi(x_{t_0+k}) - \mathbf{A}\Psi(x_{t_0+k-1}) - \mathbf{B}u_{t_0+k-1} \end{aligned} \quad (29)$$

The propagation error from t_0 to $t_0 + k$ is defined as

$$R_k := \Psi(x_{t_0+k}) - \mathbf{A}^k \Psi(x_{t_0}) - \sum_{i=0}^{k-1} \mathbf{A}^i \mathbf{B}u_{t_0+k-i} \quad (30)$$

Combining (29) with (30), we have $R_k = \sum_{i=0}^{k-1} \mathbf{A}^i r_{k-i}$. This implies

$$\|R_k\| = \left\| \sum_{i=0}^{k-1} \mathbf{A}^i r_{k-i} \right\| \leq \sum_{i=0}^{k-1} \|\mathbf{A}^i\| \|r_{k-i}\| \quad (31)$$

By assuming the generalization error based on stable solution is bounded at each step, i.e., $\|r_k\| \leq \varepsilon$ for $k \in \mathbb{N}_+ \setminus \{1\}$ and $\varepsilon > 0$, we have

$$\|R_k\| \leq \varepsilon \sum_{i=0}^{k-1} \|\mathbf{A}^i\| \quad (32)$$

Then, let $(\lambda_j)_{j=1}^N$ be the eigenvalues of \mathbf{A} , we have

$$\sum_{i=0}^{k-1} \|\mathbf{A}^i\| = \sum_{i=0}^{k-1} \sum_{j=1}^N \lambda_j^i = \sum_{j=1}^N \frac{(1 - \lambda_j^k)}{1 - \lambda_j} \leq \sum_{j=1}^N \frac{(1 - \gamma^k)}{1 - \gamma} \quad (33)$$

This leads to

$$\|R_k\| \leq \varepsilon \sum_{j=1}^N \frac{(1 - \gamma^k)}{1 - \gamma} = \frac{\varepsilon N}{1 - \gamma} (1 - \gamma^k) \leq \frac{\varepsilon N}{1 - \gamma} \quad (34)$$

Let $\sigma = \frac{\varepsilon N}{1 - \gamma}$, the results follow.

Remark 2: It is reasonable to assume a bounded generalization error at each step, since it could be proved that generalization error based on (15) is bounded using the similar strategy from LS problem's generalization error bounds or matrix completion problem's generalization error bound [12], [13].

Remark 3: One possible way to obtain a stable matrix \mathbf{A} is to add a regularization term into (15) to constraint the nuclear norm of \mathbf{A} . An alternative way is to represent \mathbf{A} as $\mathbf{S}^{-1}\mathbf{C}\mathbf{O}\mathbf{S}$, where \mathbf{S} is invertible, \mathbf{C} is orthogonal and \mathbf{O} is positive semi-definite with norm no more than γ . Based on the solution of (15) (i.e., \mathbf{A}_{ls}), a stable matrix \mathbf{A} can be further computed [11], [14].




Fig. 2. Security-critical areas (BRSs) with regards to different reach time $|t|$ (from left to right: 10s, 5s, 1s, 0.5s, 0s) in the original space. Notice that trigger set is security-critical area with $|t| = 0$ s. We only draw the boundary of set and the right hand side areas of each boundary are desired sets/regions.

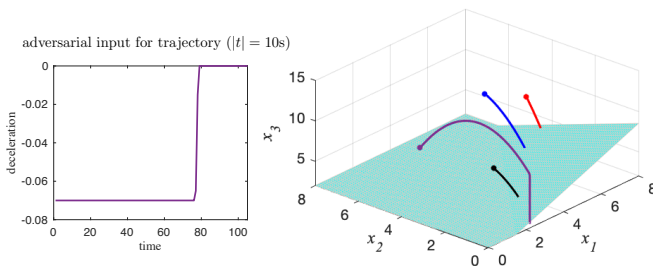


Fig. 3. Left: adversarial control input for trajectory with $|t| = 10$ s, where the time step is 0.1s; The adversarial input is -0.07m/s^2 and then close to 0m/s^2 ; Right: Trajectories that start from different security-critical regions and enter trigger set (purple: 10s; Blue: 5s; Black: 1s; Red: 0.5s).




Fig. 4. Single-lane ring system with an AV (red) and an HDV (blue) manipulated by the adversary, where (a) shows the initial state; (b) shows the trigger states in 0.5s; (c) shows the collision state in 1s; (d) depicts the trajectories computed by our method and from simulation (which means feed the computed inputs into the DRL model of AV controller and interact with the traffic system).

V. EXPERIMENTAL RESULTS

In this section we show the computational results. The effectiveness of the proposed method for trigger activation is verified on a simulated traffic system based on microscopic traffic simulator SUMO (Simulation of Urban MObility) [15], where all human-driven vehicles (HDVs) use the intelligent driver model (IDM) as described in [16]. We use one AV for this system, where the controller of AV is trained by the deep deterministic policy gradient (DDPG) algorithm developed by [17]. The system indicated by eq. (14) is discretized with a sampling time $\Delta t = 0.1\text{s}$. We run our tests on a circular system following the setting of Flow [18], where 21 vehicles run on a 230 meters long a single lane and one of them is AV.

A. Results of reachability analysis - backdoor trigger activation

Following the work by [2], a typical setting for attack interval (t_{adv}) and adversarial acceleration (a_{adv}) is $t_{adv} = 1\text{s}$ and $a_{adv} = 0.5\text{m/s}^2$, so that $g(x) = x_2 + x_3 - x_1 - 0.5$ is the boundary of the trigger set in the original space. By choosing $d = 3$, the lifted space is a 10D space, and $g_{\Psi}(\Psi(x)) = (x_2 + x_3 - x_1)^3 - 0.5^3$ is the boundary of trigger set in the lifted space. The results of both the backward reachable sets (BRSs) and trajectories are computed using the proposed method in the lifted space, and are plotted in the original 3D space by recovering from the lifted space, unless otherwise stated.

Figure 2 shows the trigger set and its corresponding BRSs (also known as security-critical areas) for different times $|t|$ in the original space. Only the boundaries of the BRSs are drawn, and the desired sets/regions are on the right-hand side of the boundaries. In Fig. 3, we depict the computed trajectories in the original space, which start from random initial states with regards to different time, i.e., $|t| = 0.5, 1, 5, 10$. We can observe that all trajectories enter the trigger set at the end, which implies the guiding information by our method is reliable to reach the trigger states (activate the triggers). We also present the computed optimal input from the adversary (actions of AV's leader) in

the left subplot of Fig. 3, the adversarial input with regards to $|t| = 10$ s is to decelerate with 0.07m/s^2 and then to maintain the speed until the trigger state is reached.

Fig. 4 shows the simulation results of the single-lane ring system starting from a random initial state with achieving time $|t| = 0.5$ s. More specifically, we present in Fig. 4 the associated trajectory computed by our method accompanied with the actual trajectory generated from simulation (using our guiding information), from which we observe these two trajectories stay close (especially the terminal states are close to each other), implying our guiding information is reliable. For complementary, we draw the speed profile and spacing-time diagram in Fig. 5 for the backdoored system before and after the trigger is activated. The adversarial input lasts for 0.5s (as indicated by the blue line) and then trigger is activated to cause a collision in around 1s (thereafter velocities of all vehicles decrease to 0 and positions of all vehicles do not change with time). This also implies our method is effective for trigger activation to cause collisions between the backdoored AV and its leading HDV.

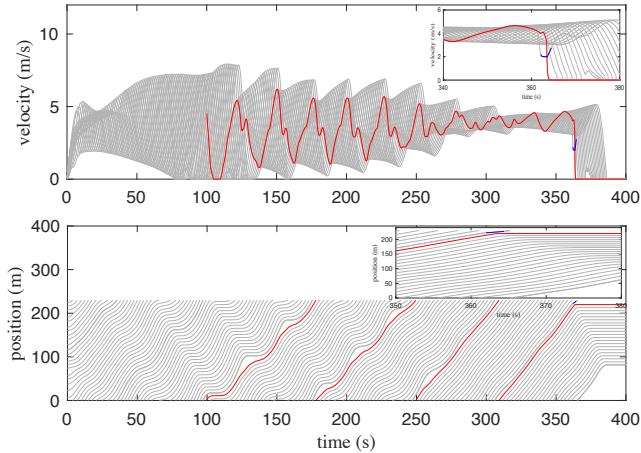


Fig. 5. Top: speed profiles for backdoored model before and after attack; Bottom: spacing-time diagram before and after attack. The blue line in each figure denotes the trajectory that suffers adversarial inputs for a specific time period (i.e. 0.5s here). We observe that collision happens soon after the adversarial manipulation of HDV ahead of AV.

VI. CONCLUSION AND FUTURE WORK

In this work, we introduce a novel data-driven approach to address the backdoor trigger activation problem. Specifically, we formulate the problem as a controlled dynamic system’s reachability problem, and learn a surrogate linear dynamic model in the lifted space to efficiently compute the backward reachable set (BRS) while bounding the generalization error. Our approach identifies security-critical areas and adversarial inputs that can trigger the system into a compromised state, highlighting the vulnerability of the AV controller to backdoor attacks. Furthermore, we demonstrate that our approach enables the design of efficient backdoor triggers that can cause vehicle collisions with the help of our guiding information. In our future work, we aim to focus on the defense side of the problem by constraining

the objective of a deep reinforcement learning (DRL)-based controller to ensure that the traffic system operates within the unreachable sets of triggers. We believe that this will provide a defense mechanism against backdoor attacks and enhance the security of AV controllers.

ACKNOWLEDGMENT

This work was supported by the NYUAD Center for Interacting Urban Networks (CITIES), funded by Tamkeen under the NYUAD Research Institute Award CG001.

REFERENCES

- [1] Y. Wang, M. Maniatakos, and S. E. Jabari, “A trigger exploration method for backdoor attacks on deep learning-based traffic control systems,” in *2021 60th IEEE Conference on Decision and Control (CDC)*. IEEE, 2021, pp. 4394–4399.
- [2] Y. Wang, E. Sarkar, W. Li, M. Maniatakos, and S. E. Jabari, “Stop-and-go: Exploring backdoor attacks on deep reinforcement learning-based traffic congestion control systems,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 4772–4787, 2021.
- [3] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, “A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games,” *IEEE Transactions on automatic control*, vol. 50, no. 7, pp. 947–957, 2005.
- [4] J. Lygeros, “On reachability and minimum cost optimal control,” *Automatica*, vol. 40, no. 6, pp. 917–927, 2004.
- [5] I. M. Mitchell *et al.*, “A toolbox of level set methods,” *UBC Department of Computer Science Technical Report TR-2007-11*, p. 31, 2007.
- [6] S. Bansal, M. Chen, S. Herbert, and C. J. Tomlin, “Hamilton-jacobi reachability: A brief overview and recent advances,” in *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. IEEE, 2017, pp. 2242–2253.
- [7] M. Chen, S. L. Herbert, M. S. Vashishtha, S. Bansal, and C. J. Tomlin, “Decomposition of reachable sets and tubes for a class of nonlinear systems,” *IEEE Transactions on Automatic Control*, vol. 63, no. 11, pp. 3675–3688, 2018.
- [8] I. Abraham, G. De La Torre, and T. D. Murphey, “Model-based control using koopman operators,” *arXiv preprint arXiv:1709.01568*, 2017.
- [9] M. Korda and I. Mezić, “Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control,” *Automatica*, vol. 93, pp. 149–160, 2018.
- [10] H.-K. Zhao, S. Osher, and R. Fedkiw, “Fast surface reconstruction using the level set method,” in *Proceedings IEEE Workshop on Variational and Level Set Methods in Computer Vision*. IEEE, 2001, pp. 194–201.
- [11] N. Gillis, M. Karow, and P. Sharma, “Approximating the nearest stable discrete-time system,” *Linear Algebra and its Applications*, vol. 573, pp. 37–53, 2019.
- [12] P. Giménez-Febrer, A. Pagès-Zamora, and G. B. Giannakis, “Generalization error bounds for kernel matrix completion and extrapolation,” *IEEE Signal Processing Letters*, vol. 27, pp. 326–330, 2020.
- [13] W. Li, C. Yang, and S. E. Jabari, “Nonlinear traffic prediction as a matrix completion problem with ensemble learning,” *Transportation science*, vol. 56, no. 1, pp. 52–78, 2022.
- [14] G. Mamakoukas, O. Xherija, and T. Murphey, “Memory-efficient learning of stable linear dynamical systems for prediction and control,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 527–13 538, 2020.
- [15] N. Khetarpal, K. Parvate, C. Wu, A. Kreidieh, E. Vinitzky, and A. Bayen, “Flow: Deep reinforcement learning for control in sumo,” *EPiC Series in Engineering*, vol. 2, pp. 134–151, 2018.
- [16] M. Treiber and A. Kesting, “Traffic flow dynamics,” *Traffic Flow Dynamics: Data, Models and Simulation*, Springer-Verlag Berlin Heidelberg, pp. 983–1000, 2013.
- [17] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, “Continuous control with deep reinforcement learning,” *arXiv preprint arXiv:1509.02971*, 2015.
- [18] C. Wu, A. R. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, “Flow: A modular learning framework for mixed autonomy traffic,” *IEEE Transactions on Robotics*, 2021.