

Attacks on Perception-Based Control Systems: Modeling and Fundamental Limits

Amir Khazraei , Henry D. Pfister , and Miroslav Pajic , Senior Member, IEEE

Abstract—We study the performance of perception-based control systems in the presence of attacks and provide methods for modeling and analysis of their resiliency to stealthy attacks on both physical and perception-based sensing. Specifically, we consider a general setup with a nonlinear affine physical plant controlled with a perception-based controller that maps both the physical [e.g., inertial measurement units (IMUs)] and perceptual (e.g., camera) sensing to the control input; the system is also equipped with a statistical or learning-based anomaly detector (AD). We model the attacks in the most general form and introduce the notions of attack effectiveness and stealthiness independent of the used AD. In such a setting, we consider attacks with different levels of runtime knowledge about the plant. We find sufficient conditions for the existence of stealthy effective attacks that force the plant into an unsafe region without being detected by any AD. We show that as the open-loop unstable plant dynamics diverges faster and the closed-loop system converges faster to an equilibrium point, the system is more vulnerable to effective stealthy attacks. Also, depending on runtime information available to the attacker, the probability of the attack remaining stealthy can be arbitrarily close to one if the attacker's estimate of the plant's state is arbitrarily close to the true state; when an accurate estimate of the plant state is not available, the stealthiness level depends on the control performance in attack-free operation.

Index Terms—Anomaly detection, attack-resilient control, learning-enabled control, nonlinear control, perception-based control, secure control, stealthy attacks.

I. INTRODUCTION

THE recent progress in deep learning and computer vision has created a new generation of control systems that incorporate perceptual data for control and decision making. For example, a camera or a light detection and ranging (LiDAR)

Manuscript received 3 August 2023; revised 7 August 2023 and 27 February 2024; accepted 27 April 2024. Date of publication 15 May 2024; date of current version 25 October 2024. This work was supported in part by ONR under Grant N00014-23-1-2206 and Grant N00014-20-1-2745, in part by AFOSR under Award FA9550-19-1-0169, in part by NSF under Award CNS-1652544, and in part by the National AI Institute for Edge Computing Leveraging Next Generation Wireless Networks under Grant CNS-2112562. An earlier version of this paper was presented in part at the Proceedings of the 4th Annual Learning for Dynamics and Control Conference. Recommended by Associate Editor F. Pasqualetti. (Corresponding author: Miroslav Pajic.)

The authors are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: amir.khazraei@duke.edu; henry.pfister@duke.edu; miroslav.pajic@duke.edu).

Digital Object Identifier 10.1109/TAC.2024.3401022

sensor can provide the controller with information about plant states (e.g., distance to the obstacles and position in a local frame). Deep neural networks (DNNs) have shown the capability to extract information from the complex perception data, such as images.

Two main approaches—i.e., 1) *modular* and 2) *end-to-end* perception-based controllers have been recently developed [2]. With the end-to-end control approach, perception data (e.g., camera images and LiDAR 3D-point clouds), combined with other physical sensor information, is directly mapped to control inputs at runtime (e.g., see [3], [4], [5]). The controller is then either learned via supervised methods, such as imitation learning that mimics a preobtained optimal control input, or using deep-reinforcement learning techniques that design the control policy by maximizing the cumulative reward in an unsupervised fashion. On the other hand, with the modular control approach, a subset of state information is extracted from the perception data (e.g., images) and then combined with other physical sensor information, followed by the use of classic feedback controllers (e.g., [6], [7], [8]). Despite the tremendous promise that DNN-based perception brings to controls, the resiliency of perception-based controllers to well-documented adversarial threats remains a challenge, limiting their applicability in real-world scenarios. The main focus of adversarial machine learning methods has been on the vulnerability of DNNs to small input perturbation, effectively focusing on robustness analysis of DNNs; e.g., targeting DNNs classification or control performance when a small and bounded noise is added to the images in camera-based control systems (e.g., [9], [10]). However, an attacker capable of compromising the system's perception/sensing would not limit their actions (i.e., injected data) to small bounded measurement perturbation; the reason is that control-theoretic methods for designing stealthy attacks suggests that the perturbation vector should gradually increase in size over time (e.g., [11], [12], [13], [14]).

Little consideration has been given on the potential impact of stealthy (i.e., undetectable) attacks, which are especially dangerous in the control context as many systems have recovery-based defenses triggered once an attack is detected. Model-based vulnerability analysis methods, designed from the control theory perspective, have been used to analyze the impact stealthy attacks could have on systems with linear time-invariant (LTI) dynamics and simple physical sensing, without perception, (e.g., [11], [14], [15], [16]). However, such analysis cannot be easily extended to systems with complex dynamics and sensing models that include perception-based sensing.

Consequently, this work studies the impact that stealthy attacks on system sensing could have on perception-based control. We assume the attack goal is to move the system into an unsafe region while remaining stealthy from *any* anomaly detector

(AD). Even though our notion of stealthiness is related to the work [17], we do not restrict the attack impact to the infinite time horizon. We show how such attacks can be modeled in a general form of an *additive term* for physical (i.e., nonperception) sensors and a *generative model* for perception sensors, without any assumptions about a bound on the perturbation values.

Perception-based controllers (either end-to-end or modular controllers) implicitly extract the state information from perception sensing (e.g., camera images) to derive suitable control inputs. Hence, the goal of this work is to evaluate their vulnerability to sensing attacks, by investigating whether there exist adversarial perception outputs that convey desired “falsified” state information while remaining stealthy, rather than blindly adding noise or, e.g., a patch to the current image. For example, for a vehicle with lane-keeping control (whose goal is to remain in the lane center), we find a sufficient stealthy attack sequence based on suitable adversarial image sequences conveying the desired falsified state information that is stealthy while fooling the controller into unsafe steering decisions.

Depending on the attacker’s level of knowledge about the plant’s states, we derive conditions for which there exists a stealthy and effective (i.e., impactful) attack sequence that forces the system far from the operating point in the safe region. In particular, we assume that the attacker has full knowledge about the system’s open-loop dynamics, and consider two cases where at runtime the attacker 1) has or 2) does not have access to the estimation of the plant’s state. We show that in the first case, the attack can approach arbitrarily close to the strict stealthiness condition as long as the estimation error is small. For the latter, the stealthiness level of the attack depends on the system’s performance in attack-free operation as well as the level of impact that the attacker expects to impose on the system (in terms of the distance from the system’s desired operating point due to the attack). Thus, there is a tradeoff between the stealthiness guarantees and the performance degradation caused by the attack.

Moreover, for LTI plants, we show that these two cases fall in the same category, which means the attacker does not need to have access to an accurate estimate of the plant’s states. We also show that unlike in systems with LTI plants and linear controllers, where the design of stealthy and effective attacks is independent of the control design, for nonlinear plants the level of stealthiness is closely related to the level of closed-loop system stability—i.e., if the closed-loop system “is more stable” (i.e., its trajectory converges faster to the equilibrium point), the attack can have stronger stealthiness guarantees. On the other hand, the attack impact (i.e., control degradation) fully depends on the level of open-loop system instability (e.g., the size of unstable eigenvalues for LTI systems).

A. Related Work

The initial work [18] on adversarial example generation showed that DNNs are vulnerable to small input perturbations. Afterward, the majority of works have applied this idea to adversarial attacks on a physical world such as malicious stickers on traffic signs to fool the detectors and/or classifiers [19], [20], [21]. For example, the design of adversarial attacks for regression tasks has been studied in [22], where the goal of the attacker is to alter the geometrical prediction of the scene and the predicted distances from the camera. Yet, all these methods only consider classification or regression tasks in a static manner;

i.e., the target only depends on its input, without consideration of the longitudinal (i.e., over time) system behaviors. The vulnerability of perception-based vehicle controls has been recently studied in a longitudinal way (e.g., [9], [10], [23], [24]). For instance, Boloor et al. [9] consider autonomous vehicles (AVs) with end-to-end DNN controllers that directly map perceptual inputs into the vehicle steering angle and target the systems by painting black lines on the road. On the other hand, the authors in [10] and [23] introduced online attacks on streams of images in the tracking task while Hallyburton et al. [24] consider longitudinal attacks on camera-LiDAR fusion. Specifically, Jia et al. [10] use the idea of adding a small patch into the image for a couple of consecutive frames to change the position of the bounding box around each object, where the location and the size of the patch are obtained by solving an optimization problem.

However, these works only consider specific applications and analyze the attack impact in an ad-hoc manner, limiting the use of their results in other systems/domains. Furthermore, they lack any consideration of attack stealthiness, as injecting, e.g., adversarial patches that only maximizes the disruptive impact on the control can be detected by most ADs. For instance, Cai and Koutsoukos [25] introduced an AD that easily detects the adversarial attacks from the work in [9]. On the other hand, in this work, we focus on systems with nonlinear system dynamics, define general notions of attack stealthiness, and introduce sufficient conditions for a perception-based control system to be vulnerable to effective yet stealthy perception and sensing attacks. We show that to launch a successful stealthy attack, the attacker may need to compromise other sensing information besides perception.

Finally, for *nonperception* control systems, stealthy attacks have been well defined in, e.g., [11], [13], [16], [17], [26], [27], [28], [29], [30], [31], [32], [33], [34], including replay [26], covert [28], zero-dynamic [27], and false data injection attacks [11], [12], [16]. However, all these works only focus on LTI systems and linear controllers, as well as on specific AD designs (e.g., χ^2 detector). The problem of resilient state estimation for a class of nonlinear control systems has been considered in, e.g., [35], focusing on systems that do not include perception as part of the closed loop.

The notion of attack stealthiness independent of the employed AD (i.e., remaining stealthy for *all* existing/potential ADs) has been studied in [17]. We additionally differentiate our work in the following—our notion of stealthiness is stronger than the one in [17] as stealthiness there depends on time; i.e., there exists only a bounded time that the attack can stay undetected by an AD. However, the notion of stealthiness in our work is independent of time and the attack is guaranteed to be stealthy for all time steps after initiating the attack. Moreover, the performance degradation metric used in [17] is the error covariance of a Kalman filter estimator as opposed to our work; we assume the attacker’s goal is to force the system states into an unsafe region. In addition, we consider systems with nonlinear dynamics as well as with both perception and physical sensors, unlike the work in [17] where only LTI plants with physical (i.e., without perception) sensors are investigated. Furthermore, this work considers attacks on both perception (e.g., images) and sensor measurements while the existing vulnerability analyses do not take into account perception. Besides, these works all assume the LTI model and the linear controller for the system dynamics while we consider a class of systems with a nonlinear dynamical model and control. In addition, they focus on

the detectability of attacks with respect to specific detection schemes employed by the controller, such as the classic χ^2 anomaly detection algorithm. Recently, Khazraei et al. [30] have introduced a learning-based attack design for systems with nonlinear dynamics; yet, in addition to providing no formal analysis, the work only considers stealthiness with respect to the χ^2 -based AD and does not consider perception-based controllers.

B. Paper Contribution and Organization

The contribution of this work is twofold. First, we consider attacks on perception-based control in the presence of ADs. We define a new notion of stealthiness using Neyman–Pearson Lemma and its relation with total variation and Kullback–Leibler (KL) divergence, where an attack is considered stealthy if it is stealthy from *any* AD. We find sufficient conditions for which the perception-based control systems are vulnerable to highly effective, in terms of moving the system from the desired operating point, yet stealthy attacks. Second, unlike all previous works in control literature, we consider highly impactful stealthy attacks on plants with *nonlinear* dynamics, controlled by a perception-based controller potentially employing an end-to-end nonlinear controller (e.g., DNN).

The rest of the article is organized as follows. In Section III, we present the system model, as well as a model of attacks on perception before introducing the concept of stealthy yet effective attacks on the control systems. Section IV introduces two attack strategies and provides conditions under which the system is vulnerable to highly effective stealthy attacks. Finally, in Section V, we provide case studies to illustrate these conditions, before concluding remarks in Section VI.

II. PRELIMINARIES

In this section, we introduce employed notation before presenting properties of KL divergence known as *data processing inequality*, *monotonicity*, and *chain-rule* as presented in [36].

Notation: \mathbb{R} denotes the set of reals, whereas \mathbb{P} and \mathbb{E} denote the probability and expectation of a random variable, respectively. For a square matrix A , $\lambda_{\max}(A)$ is the maximum eigenvalue. For a vector $x \in \mathbb{R}^n$, $\|x\|_p$ denotes the p -norm of x ; when p is not specified, the 2-norm is implied. For a vector sequence, $x_0 : x_t$ denotes the set $\{x_0, x_1, \dots, x_t\}$. A function $f : \mathbb{R}^n \rightarrow \mathbb{R}^p$ is Lipschitz on the set \mathcal{D} with constant L if for any $x, y \in \mathcal{D} \subseteq \mathbb{R}^n$ it holds that $\|f(x) - f(y)\| \leq L\|x - y\|$; it is globally Lipschitz with constant L if $\mathcal{D} = \mathbb{R}^n$. If X and Y are two sets, $X - Y$ includes the elements in X that are not in Y . For a set X , ∂X and X° define the boundary and the interior of the set, respectively. B_r denotes a closed ball centered at zero with radius r ; i.e., $B_r = \{x \in \mathbb{R}^n \mid \|x\| \leq r\}$, whereas $\mathbf{1}_A$ is the indicator function on a set A . For a function f , we denote $f' = \frac{\partial f}{\partial x}$ as the partial derivative of f with respect to x and $\nabla f_i(x)$ is the gradient of the function f_i (i th element of the function f). Finally, if \mathbf{P} and \mathbf{Q} are probability distributions relative to the Lebesgue measure with densities \mathbf{p} and \mathbf{q} , respectively, then the total variation between them is defined as $\text{TV}(\mathbf{P}, \mathbf{Q}) = \frac{1}{2} \int |\mathbf{p}(x) - \mathbf{q}(x)| dx$. The KL divergence between \mathbf{P} and \mathbf{Q} is defined as $\text{KL}(\mathbf{P} \parallel \mathbf{Q}) = \int \mathbf{p}(x) \log \frac{\mathbf{p}(x)}{\mathbf{q}(x)} dx$.

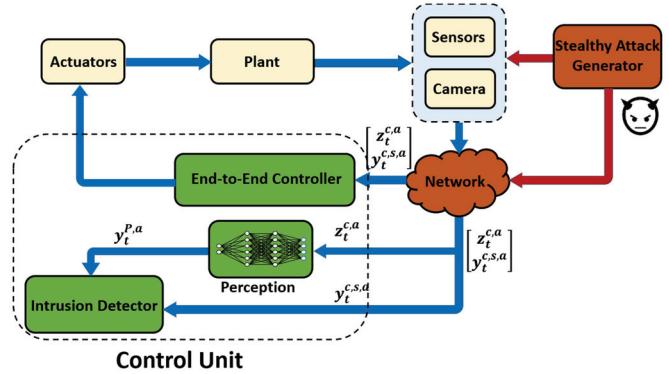


Fig. 1. Architecture of a perception-based control system under attack on system sensing, including perception. Independently of the way attacks are actually implemented (i.e., directly compromising a sensor or modifying the measurements delivered over the network to the controller), the same impact on the control performance is obtained.

A. Properties of KL Divergence and Other Preliminaries

In the following lemmas, X and Y are assumed to be random variables (Lemmas 1–4 proofs can be found in [36]).

Lemma 1 (Data Processing Inequality): Assume Y is produced given X based on the conditional law $W_{Y|X}$. Let P_Y (respectively, Q_Y) denote the distribution of Y when X is distributed as P_X (respectively, Q_X). Then

$$\text{KL}(Q_Y \parallel P_Y) \leq \text{KL}(Q_X \parallel P_X). \quad (1)$$

Lemma 2 (Monotonicity): Let $P_{X,Y}$ and $Q_{X,Y}$ be two distributions for a pair of variables X and Y , and P_X and Q_X be the marginal distributions for variable X . Then

$$\text{KL}(Q_X \parallel P_X) \leq \text{KL}(Q_{X,Y} \parallel P_{X,Y}). \quad (2)$$

Lemma 3 (Chain rule): Let $P_{X,Y}$ and $Q_{X,Y}$ be two distributions for a pair of variables X and Y . Then

$$\text{KL}(Q_{X,Y} \parallel P_{X,Y}) = \text{KL}(Q_X \parallel P_X) + \text{KL}(Q_{Y|X} \parallel P_{Y|X}) \quad (3)$$

where $\text{KL}(Q_{Y|X} \parallel P_{Y|X})$ is defined as

$$\text{KL}(Q_{Y|X} \parallel P_{Y|X}) = \mathbb{E}_{x \sim Q_X} \text{KL}(Q_{Y|X=x} \parallel P_{Y|X=x}). \quad (4)$$

Lemma 4: Let P_X and Q_X be two Gaussian distributions with the same covariance Σ and different means μ_Q and μ_P , respectively. Then, it holds that

$$\text{KL}(Q_X \parallel P_X) = \mu_Q^T \Sigma^{-1} \mu_P. \quad (5)$$

Lemma 5: Let Q_X be a distribution for a scalar random variable X , and that $X \leq M$ for some $M > 0$. Then

$$\mathbb{E}_{Q_X} \{X\} \leq M. \quad (6)$$

Proof: The proof is straightforward from the definition of expectation and some properties of the integral. ■

III. MODELING-PERCEPTION-BASED CONTROL SYSTEMS IN THE PRESENCE OF ATTACKS

In this section, we introduce the system model and show how to capture attacks on system sensing, including perception. Specifically, we consider the setup from Fig. 1 where each of the components is modeled as follows.

A. Plant and Perception Model

We assume the plant has nonlinear dynamics in the standard state-space form

$$\begin{aligned} x_{t+1} &= f(x_t) + B u_t + w_t, \quad y_t^s = C_s x_t + v_t^s \\ z_t &= G(x_t). \end{aligned} \quad (7)$$

Here, $x_t \in \mathbb{R}^n$, $u_t \in \mathbb{R}^m$, $w_t \in \mathbb{R}^n$, $z_t \in \mathbb{R}^l$, $y_t^s \in \mathbb{R}^s$, and $v^s \in \mathbb{R}^s$ denote the state, input, system disturbance, observations from perception-based sensors, (nonperception) sensor measurements, and sensor noise, at time t , respectively. The perception-based sensing is modeled by an unknown generative model G , which is nonlinear and potentially high dimensional. For example, consider a camera-based lane-keeping system. Here, the observations z_t are the captured images; the map G generates the images based on the vehicle's position. Without loss of generality, we assume that $f(0) = 0$. Finally, the process and measurement noise vectors w and v^s are assumed independent and identically distributed (i.i.d.) Gaussian processes $w \sim \mathcal{N}(0, \Sigma_w)$ and $v^s \sim \mathcal{N}(0, \Sigma_{v^s})$.

B. Control Unit

The control unit, shown in Fig. 1, consists of perception, controller, and AD units. We assume that the control unit receives $y^{c,s}$ and z^c as the sensor measurements and perception sensing (e.g., images), respectively, as an input. Thus, without malicious activity, it holds that $y^{c,s} = y^s$ and $z^c = z$. Now, we describe in detail each of the components.

1) Perception: We assume that there exists a perception map P that imperfectly estimates the partial state information from perception sensing (e.g., images)—i.e.,

$$y_t^P = P(z_t^c) = C_P x_t + v^P(x_t) \quad (8)$$

here, P denotes a DNN trained using any supervised learning method on a data set $\mathcal{X} = \{(z_i, x_i)\}_{i=1}^N$ collected densely around the operating point x_o of the system, as in [37] and [38]. In addition, $v^P \in \mathbb{R}^p$ is the perception map error that depends on the state of the system—i.e., smaller around the training dataset. To capture perception guarantees, we employ the model for robust perception-based control from the work in [6], and the standard model of the perception map (8) from the work in [7] that can capture well DNN-based perception modules. Specifically, if the model is trained effectively, we assume that the perception error v^P around the operating point x_o is bounded, i.e., the following assumption from the work in [7] holds.

Assumption 1: There exists a safe set \mathcal{S} with the radius $R_{\mathcal{S}}$ (i.e., $\mathcal{S} = \{x \in \mathbb{R}^n \mid \|x\| \leq R_{\mathcal{S}}\}$) around the operating point such that for all $x \in \mathcal{S}$, it holds that $\|P(z) - C_P x\| \leq \gamma$, where $z = G(x)$ —i.e., for all $x \in \mathcal{S}$, $\|v^P(x)\| \leq \gamma$. Without loss of generality, in this work, we consider the origin as the operating point—i.e., $x_o = 0$.

Remark 1: As will be shown later in this work, the assumption on boundedness of $\|v^P(x)\|$ and having a good estimate of such bound is only essential for the defender because the perception map is also used for anomaly detection. We should note that such a bound will not be employed by the attacker to design stealthy impactful attacks. A systematic method to find the bound on $\|v^P(x)\|$ is discussed in [7].

2) Controller: The physical plant (7) is controlled by a (general) nonlinear controller $u_t = \pi(z_t^c, y_t^{c,s})$ that maps the image and physical sensor information to the input control. Using (7), one can write the previous control law in an equivalent form of

$u_t = \Pi(x_t, v_t^s)$ that absorbs $z_t^c = G(x_t)$ in the control function π . Hence, for

$$h(x_t, v_t^s) = f(x_t) + B\pi(z_t^c, y_t^{c,s}) = f(x_t) + B\Pi(x_t, v_t^s)$$

the evolution of the closed-loop system can be captured as

$$x_{t+1} = h(x_t, v_t^s) + w_t. \quad (9)$$

In the general form, the controller can employ *any* end-end control policy that uses the image and sensor measurements. For noiseless systems, the state dynamics can be captured as¹

$$x_{t+1} = h(x_t, 0). \quad (10)$$

Definition 1: The origin of the system (10) is exponentially stable on a set $\mathcal{D} \subseteq \mathbb{R}^n$ if for any $x_0 \in \mathcal{D}$, there exists $0 < \alpha < 1$ and $M > 0$, such that $\|x_t\| \leq M\alpha^t\|x_0\|$ for all $t \geq 0$.

Lemma 6 (see [39]): For the system from (10), if there exists a function $V : \mathbb{R}^n \rightarrow \mathbb{R}$ such that for any $x_t \in \mathcal{D} \subseteq \mathbb{R}^n$, the following holds

$$\begin{aligned} c_1\|x_t\|^2 &\leq V(x_t) \leq c_2\|x_t\|^2 \\ V(x_{t+1}) - V(x_t) &\leq -c_3\|x_t\|^2 \\ \left\| \frac{\partial V(x)}{\partial x} \right\| &\leq c_4\|x\| \end{aligned} \quad (11)$$

for some positive c_1, c_2, c_3 , and c_4 , then the origin is exponentially stable.

Assumption 2: We assume that for the closed-loop control system (10), the origin is exponentially stable on a set $\mathcal{D} = B_d$. Using the converse Lyapunov theorem [40], there exists a Lyapunov function that satisfies the inequalities in (11) with constants c_1, c_2, c_3 , and c_4 on a set $\mathcal{D} = B_d$. As a result, it follows that $\Pi(0, 0) = 0$.

Remark 2: The assumptions for a closed-loop system are critical for system guarantees without the attack; i.e., if the system does not satisfy the stability property in attack-free scenarios, then an effective strategy for the attacker would be to wait until the system fails by itself. We refer the reader to the recent work, e.g., [6], [7] on the design of such controllers.

Remark 3: Note that the exponential stability assumption on the closed-loop system (10) is only considered due to the simplicity of the notation and can be relaxed to control systems with asymptotic stability that satisfies the converse Lyapunov theorem conditions [40, Th. 3.14]. Specifically, a similar result as in Lemma 8 can be obtained; however, the needed notation would be significantly more cumbersome.

Definition 2: Let \mathcal{U}_ρ be the set of all functions $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ such that the dynamics $x_{t+1} = f(x_t) + d_t$, where d_t satisfies $\|d_t\| \leq \rho$, reaches arbitrarily large states for some nonzero initial state x_0 . For a function f from \mathcal{U}_ρ and an initial condition x_0 , we define

$$T_f(\alpha, x_0) = \min\{t \mid \|x_t\| \geq \alpha\} \quad (12)$$

i.e., $T_f(\alpha, x_0)$ is the minimal number of time-steps needed for an unstable dynamic f , starting from the initial condition x_0 , to leave a bounded ball with the center at zero and radius α .²

¹With slight abuse of notation, x is used to denote the states of the noiseless system; yet, in the rest of the article, we use x to denote the states of the actual physical system with noise and we clarify if the other case is implied.

²To simplify our notation, and since we consider specific f from the plant dynamics (7), we drop the subscript f .

3) Anomaly Detector: The system is equipped with an AD designed to detect the presence of any abnormal behaviors. We use $Y_t = \begin{bmatrix} y_t^P \\ y_{t,s}^c \end{bmatrix}$ and $Y_t^a = \begin{bmatrix} y_t^{P,a} \\ y_{t,s,a}^c \end{bmatrix}$ to capture sensor [from (7)] and perception-based [from (8)] values without and under attack, respectively—we use the superscript a to differentiate all signals of the attacked system, with the full attack model introduced in the next section. Now, by denoting $Y_{-\infty}^{-1} = Y_{-\infty} : Y_{-1}$, we consider the classical binary hypothesis testing problem as follows:

H_0 : normal condition (the AD receives $Y_{-\infty} : Y_t$);

H_1 : abnormal behavior (the AD receives $Y_{-\infty}^{-1}, Y_0^a : Y_t^a$).

Effectively, the AD uses both the extracted state information from the perception map [i.e., (8)] as well as sensor measurements. Given a random sequence $\bar{Y}^t = (\bar{Y}_{-\infty} : \bar{Y}_t)$, it either comes from the distribution \mathbf{P} (null hypothesis H_0), which is determined by system uncertainties, or from a distribution \mathbf{Q} (the alternative hypothesis H_1); note that the *unknown* distribution \mathbf{Q} is controlled by the attacker.

For a given AD specified by a function $D : \bar{Y}^t \rightarrow \{0, 1\}$, two types of error may occur. Error type (I), also referred as *false alarm*, occurs if $D(\bar{Y}^t) = 1$ when $\bar{Y}^t \sim \mathbf{P}$; whereas type (II) error (*miss detection*) occurs if $D(\bar{Y}^t) = 0$ when $\bar{Y}^t \sim \mathbf{Q}$. Hence, the sum of the *conditional error probabilities* of AD D for a given random sequence \bar{Y}^t is

$$p_t^e(D) = \mathbb{P}(D(\bar{Y}^t) = 0 | \bar{Y}^t \sim \mathbf{Q}) + \mathbb{P}(D(\bar{Y}^t) = 1 | \bar{Y}^t \sim \mathbf{P}). \quad (13)$$

Note that $p_t^e(D)$ is not a probability measure as it can take values larger than one. Let us define $p_t^{\text{TD}}(D) = \mathbb{P}(D(\bar{Y}^t) = 1 | \bar{Y}^t \sim \mathbf{Q})$ as the probability of true detection, and $p_t^{\text{FA}}(D) = \mathbb{P}(D(\bar{Y}^t) = 1 | \bar{Y}^t \sim \mathbf{P})$ as the probability of false alarm for the detector D . We say that an AD (defined by D) to be better than a random guess-based AD (defined by D_{RG}) if $p_t^{\text{FA}}(D) < p_t^{\text{TD}}(D)$; as with the random guess it holds that

$$\begin{aligned} p^{\text{FA}}(D_{\text{RG}}) &= \mathbb{P}(D_{\text{RG}}(\bar{Y}^t) = 1 | \bar{Y}^t \sim \mathbf{P}) = \mathbb{P}(D_{\text{RG}}(\bar{Y}^t) = 1) \\ &= \mathbb{P}(D_{\text{RG}}(\bar{Y}^t) = 1 | \bar{Y}^t \sim \mathbf{Q}) = p^{\text{TD}}(D_{\text{RG}}). \end{aligned}$$

C. Attack Modeling

We assume that the attacker has the ability to compromise perception-based sensing (e.g., camera images) as well as (potentially) the sensor measurements y_t^s delivered to the controller (see Fig. 1). Such attacks on physical sensors can be achieved by directly compromising the sensing (or the environment of the sensors) or the communication between the sensors and the controller [41], [42]. On the other hand, implementing the attack on the images delivered to the controller may not be feasible using physical spoofing attacks such as LiDAR spoofing by injecting laser data points [21], [24]. Hence, the attacker needs to compromise the images in the cyber domain (e.g., as discussed in [43]). This can be achieved by modifying the firmware directly on the camera or the driver code, as done in Hyundai's Gen5W and Tesla's Model 3 attacks where custom firmware was installed [44], [45], or the Tesla Model S attacks that sent custom messages by compromising application drivers [46]. Similarly, such cyberattacks can be achieved using *Man-in-the-Middle* attacks that target the connection between the camera and the controller; examples include intercepting legitimate messages and manipulating their contents, before resending them to the intended destination (e.g., as described in intervehicle attack scenarios in [41] and [42] or vehicle attacks that manipulated

timing on V2V/V2I messages, causing out-of-date situational awareness [47].

Moreover, we assume the attack starts at $t = 0$, and as previously mentioned, we use the superscript a to differentiate all signals of the attacked system, for all $t \geq 0$; the attack sequence is $\{z_t^{c,a}, y_t^{c,s,a}\}_{t \geq 0}$, where, e.g., the value of observation delivered to the perception unit at time t is denoted by $z_t^{c,a}$. Note that due to nonlinearity of the operators (P and G), we do not employ the additive false-data injection model for perception attacks, widely used for LTI systems with nonperception sensing (e.g., [11]).

Thus, the system dynamics under attack can be modeled as

$$\begin{aligned} x_{t+1}^a &= f(x_t^a) + Bu_t^a + w_t^a \\ u_t^a &= \pi(z_t^{c,a}, y_t^{c,s,a}). \end{aligned} \quad (14)$$

In this work, we assume the attacker has full knowledge of the system, its dynamics, and employed architecture. Furthermore, the attacker has the required computation power to calculate suitable attack signals to inject, planning ahead as needed.

Remark 4: The assumption that the attacker has knowledge of the system dynamics (i.e., function f and matrix C_s) is commonly used in existing studies focused on the worst case analysis of the attack impact on control systems [11], [12], [14], [16]. In particular, by focusing on resourceful attackers who possess extensive knowledge of the targeted systems, we can effectively assess the consequences of worst case attacks and develop appropriate defensive strategies. However, the attacks presented in Section IV do not require knowledge of the noise profile, significantly improving their applicability.

Remark 5: In our notation, $x_0^a : x_t^a$ denotes a state trajectory of the system under attack (for an attack starting at $t = 0$) while $x_0 : x_t$ denotes the state trajectory of the attack-free system; we refer to such state trajectory as the *attack-free* trajectory. Thus, when comparing the attack-free trajectory and the system trajectory under attack [i.e., from (14)], we assume that $w_t^a = w_t$ and $v_t^{s,a} = v_t^s$. However, we cannot make such assumptions for $v_t^P(x)$ as it is a function of states and the states are compromised due to the attack.

We define an attack to be stealthy if the best strategy for the AD is to ignore the measurements and make a random guess between the hypotheses; i.e., that there is *no* AD D that satisfies $p^{\text{TD}}(D) > p^{\text{FA}}(D)$. However, reaching such stealthiness guarantees may not be possible in general. Therefore, we define the notion of ϵ -stealthiness, which as we will show later, is attainable for a large class of nonlinear systems. Formally, we define the notions of *strict stealthiness* and ϵ -stealthiness as follows.

Definition 3: Consider the system defined in (7). An attack sequence is *strictly stealthy* if there exists no detector for which $p_t^{\text{FA}} < p_t^{\text{TD}}$ holds, for any $t \geq 0$. An attack is ϵ -stealthy if for a given $\epsilon > 0$, there exists no detector such that $p_t^{\text{FA}} < p_t^{\text{TD}} - \epsilon$ holds, for any $t \geq 0$.

Before introducing the sufficient condition for the previous notion of stealthiness, we consider the following lemma.

Lemma 7: The AD D satisfies $p^{\text{FA}}(D) < p^{\text{TD}}(D) - \epsilon$ if and only if $p^e(D) < 1 - \epsilon$. Also, $p^e(D) = 1$ if and only if D performs as a random-guess detector.

Proof: First, we consider the case $p^e(D) < 1 - \epsilon$. From (13), it holds that

$$\begin{aligned} p^e(D) &= \mathbb{P}(D(\bar{Y}) = 0 | \bar{Y} \sim \mathbf{Q}) + \mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{P}) \\ &= 1 - \mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{Q}) \end{aligned}$$

$$+ \mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{P}) < 1 - \epsilon. \quad (15)$$

Thus, $\mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{P}) < \mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{Q}) - \epsilon$ or $p^{\text{FA}}(D) < p^{\text{TD}}(D) - \epsilon$.

Now, if we have $p^e(D) = 1$, then we get $\mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{P}) = \mathbb{P}(D(\bar{Y}) = 1 | \bar{Y} \sim \mathbf{Q})$ where the decision of the detector D is independent of the distribution of \bar{Y} , and therefore, the detector performs as the random-guess detector. Since the reverse of all these implications holds, the other (i.e., necessary) conditions of the theorem also hold. \blacksquare

Now, we can capture stealthiness conditions in terms of KL divergence of the corresponding distributions.

Theorem 1: An attack sequence is strictly stealthy if and only if $\text{KL}(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_t^a) || \mathbf{P}(Y_{-\infty} : Y_t)) = 0$ for all $t \geq 0$. An attack sequence is ϵ -stealthy if the corresponding observation sequence $Y_0^a : Y_t^a$ satisfies

$$\text{KL}(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_t^a) || \mathbf{P}(Y_{-\infty} : Y_t)) \leq \log\left(\frac{1}{1 - \epsilon^2}\right). \quad (16)$$

Proof: With some abuse of notation only specific to this theorem, \mathbf{Q} and \mathbf{P} are used to denote $\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_t^a)$ and $\mathbf{P}(Y_{-\infty} : Y_t)$, respectively. First, we prove the strictly stealthy case.

Using the Neyman–Pearson lemma for any existing detector D , it follows that

$$p_t^e(D) \geq \int \min\{\mathbf{Q}(y), \mathbf{P}(y)\} dy \quad (17)$$

where the equality holds for the likelihood ratio function as $D^* = \mathbf{1}_{\mathbf{Q} \geq \mathbf{P}}$ [48]. Since $1 - \int \min\{\mathbf{Q}(y), \mathbf{P}(y)\} dy = \frac{1}{2} \int |\mathbf{q}(x) - \mathbf{p}(x)| dx$, from the work in [49] and the definition of total variation distance between \mathbf{Q} and \mathbf{P} , it holds that

$$p_t^e(D) \geq 1 - \text{TV}(\mathbf{Q}, \mathbf{P}) \quad (18)$$

where TV denotes the total variation distance between the distributions.

Now, it holds that $\text{TV}(\mathbf{Q}, \mathbf{P}) \leq \sqrt{1 - e^{-\text{KL}(\mathbf{Q} || \mathbf{P})}}$ [50, (eq. (14.11))]. Thus, if $\text{KL}(\mathbf{Q} || \mathbf{P}) = 0$ holds, then $p_t^e(D) \geq 1$ for any detector D . Therefore, according to Lemma 7, the attack is strictly stealthy. On the other hand, if for any detector D $p_t^e(D) \geq 1$ holds, then the equality holds for $\text{TV}(\mathbf{Q}, \mathbf{P}) = 0$; this is equivalent to $\mathbf{Q} = \mathbf{P}$, and therefore, $\text{KL}(\mathbf{Q} || \mathbf{P}) = 0$.

For the ϵ -stealthy case, we combine (18) with the inequality $\text{TV}(\mathbf{Q}, \mathbf{P}) \leq \sqrt{1 - e^{-\text{KL}(\mathbf{Q} || \mathbf{P})}}$ and the ϵ -stealthy condition (16), to show

$$p_t^e(D) \geq 1 - \text{TV}(\mathbf{Q}, \mathbf{P}) \geq 1 - \sqrt{1 - e^{-\text{KL}(\mathbf{Q} || \mathbf{P})}} \geq 1 - \epsilon$$

therefore, according to Lemma 7, the attack is ϵ -stealthy. \blacksquare

Remark 6: The ϵ -stealthiness condition defined in [17] requires that

$$\lim_{t \rightarrow \infty} \frac{\text{KL}(\mathbf{Q}(Y_0^a : Y_t^a) || \mathbf{P}(Y_0 : Y_t))}{t} \leq \epsilon.$$

This allows for the KL divergence to linearly increase over time for any $\epsilon > 0$, and as a result, after a large-enough time period, the attack may be detected even though it satisfies the definition of stealthiness from the work in [17]. Yet, the ϵ -stealthiness from Definition 3 only depends on ϵ and is fixed for any time t ; thus, it introduces a stronger notion of stealthiness for the attack.

Attack Goal: We capture the attacker's goal as *maximizing* the degradation of the control performance. Specifically, as we consider the origin as the desired operating point of the closed-loop system, the attack objective is to maximize the (norm of)

states x_t . Moreover, the attacker wants to *remain stealthy*—i.e., *undetected* by any employed AD, as formalized ahead.

Definition 4: The attack sequence, denoted by $\{z_0^{c,a}, y_0^{c,s,a}\}, \{z_1^{c,a}, y_1^{c,s,a}\}, \dots$, is referred to (ϵ, α) -successful attack if there exists $t' \geq 0$ such that $\|x_{t'}\| \geq \alpha$ and the attack is ϵ -stealthy for all $t \geq 0$. When such a sequence exists for a system, the system is called (ϵ, α) -attackable. Finally, when the system is (ϵ, α) -attackable for arbitrarily large α , the system is referred to as *perfectly attackable*.

In the rest of this work, our goal is to derive methods to capture the impact of stealthy attacks; specifically, in the next section, we derive conditions for the existence of a *stealthy yet effective* attack sequence $\{z_0^{c,a}, y_0^{c,s,a}\}, \{z_1^{c,a}, y_1^{c,s,a}\}, \dots$, resulting in $\|x_t\| \geq \alpha$ for some $t \geq 0$ —i.e., we find conditions for a system to be (ϵ, α) -attackable. Here, for an attack to be stealthy, we focus on the ϵ -stealthy notion; i.e., that the best AD could only improve the probability detection by ϵ compared to the random-guess baseline detector.

IV. CONDITIONS FOR (ϵ, α) -ATTACKABLE SYSTEMS

To provide sufficient conditions for a system to be (ϵ, α) -attackable, in this section, we introduce two methodologies to design attack sequences on perception and (classical) sensing data. The difference in these strategies is the level of information that the attacker has about the system; we show that the stronger attack impact can be achieved with the attacker having full knowledge of the system state.

Specifically, we start with the attack strategy where the attacker has access to the current estimation of state; in such case, we show that the stealthiness condition is less restrictive, simplifying the design of ϵ -stealthy attacks. For the second attack strategy, we show that the attacker can launch the attack sequence by only knowing the function f (i.e., plant model); however, achieving ϵ -stealthy attack in this case is harder as more restrictive conditions are imposed on the attacker.

A. Attack Strategy I: Using Estimate of the Plant State

Consider the attack sequence where $z_t^{c,a}$ and $y_t^{c,s,a}$ injected at time t , for all $t \geq 0$, satisfy

$$z_t^{c,a} = G(x_t^a - s_t), \quad y_t^{c,s,a} = C_s(x_t^a - s_t) + v_t^s \quad (19)$$

with $s_{t+1} = f(\hat{x}_t^a) - f(\hat{x}_t^a - s_t)$, and for a nonzero s_0 . Here, \hat{x}_t^a denotes an estimation of the plant's state (in the presence of attacks), and thus, $\zeta_t = \hat{x}_t^a - x_t^a$ is the corresponding state estimation error. Note that the attacker can obtain \hat{x}_t^a by, e.g., running a local estimator using the true measurements before injecting the false values; i.e., $y_t^{s,a} = C_s x_t^a + v_t^s$ and $z_t^a = G(x_t^a)$. We assume that the estimation error is bounded by b_ζ —i.e., $\|\zeta_t\| \leq b_\zeta$, for all $t \geq 0$. On the other hand, the earlier attack design may not require access to the true plant state x_t^a since only the “shifted” (i.e., $x_t^a - s_t$) outputs of the real sensing/perception are injected. For instance, in the lane centering control (i.e., keeping the vehicle between the lanes), $G(x_t - s_t)$ only shifts the actual image s_t to the right or left depending on the coordinate definition. Similarly, the attack on physical (i.e., nonperception) sensors can be implemented as $y_t^{c,s,a} = C_s x_t^a + v_t^{s,a} - C_s s_t = y_t^{s,a} - C_s s_t$ where the attacker only needs to subtract $C_s s_t$ from the current true measurements. The idea behind the earlier attacks is to have the system believe that its (plant) state is equal to the state $e_t \triangleq x_t^a - s_t$; thus, referred to as the *fake state*. Note that effectively both z_t^a and

$y_t^{s,a}$ used by an AD are direct functions of the fake state e_t . Thus, if the distribution of $e_0 : e_t$ is close to $x_0 : x_t$ (i.e., attack-free trajectory), then the attacker will be successful in injecting a stealthy attack sequence.

Definition 5: For an attack-free state trajectory $x_0 : x_t$, and for any $T \geq 0$ and $b_x, b_v > 0$, $\delta(T, b_x, b_v)$ is the probability that the system state and physical sensor noise v^s remain in the zero center ball with radius b_x and b_v , respectively, during time period $0 \leq t \leq T$, i.e.,

$$\delta(T, b_x, b_v) \triangleq \mathbb{P} \left(\sup_{0 \leq t \leq T} \|x_t\| \leq b_x, \sup_{0 \leq t \leq T} \|v_t^s\| \leq b_v \right). \quad (20)$$

When the system with exponentially stable closed-loop control dynamics³ is affected by a bounded perturbation, one can show that the state of the system will remain in a bounded set. The following lemma from the work in [40] provides the condition and the upper bound on the norm of the state.

Lemma 8 (see [40]): Let $x = 0$ be an exponentially stable equilibrium point of the nominal system (10). Also, let $V(x_t)$ be a Lyapunov function of the nominal system that satisfies (11) in \mathcal{D} , where $\mathcal{D} = B_d$. Suppose the system is affected by additive perturbation term $g(x_t)$ that satisfies $\|g(x_t)\| \leq \delta + \gamma \|x_t\|$. If $c_3 - \gamma c_4 > 0$ with $\delta < \frac{c_3 - \gamma c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$ holds for all $x \in \mathcal{D}$ and some positive $\theta < 1$, then for all $\|x_0\| < \sqrt{\frac{c_1}{c_2}} d$, there exists $t_1 > 0$ such that for all $0 \leq t \leq t_1$ $\|x_t\| \leq \kappa e^{-\beta t} \|x_0\|$ holds with $\kappa = \sqrt{\frac{c_2}{c_1}}$, $\beta = \frac{(1-\theta)(c_3 - \gamma c_4)}{2c_2}$ and for $t \geq t_1$ it holds that $\|x_t\| \leq b$ with $b = \frac{c_4}{c_3 - \gamma c_4} \sqrt{\frac{c_2}{c_1} \frac{\delta}{\theta}}$.

The next result captures conditions under which a perception-based control system is not resilient to attacks, in the sense that it is (ϵ, α) -attackable.

Theorem 2: Consider the system (7) with closed-loop control as in Assumption 2. Assume that the functions f , f' (derivative of f) and Π' (derivative of Π) are Lipschitz, with constants L_f , L'_f , and L'_{Π} , respectively, and let us define $L_1 = L'_f(b_x + 2b_{\zeta} + \phi)$, $L_2 = \min\{2L_f, L'_f(\alpha + b_x + b_{\zeta})\}$, and $L_3 = L'_{\Pi}(b_x + \phi + b_v)$ for some $\phi > 0$. Moreover, assume that there exists $\phi > 0$ such that the inequalities $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ and $L_2 b_{\zeta} < \frac{c_3 - (L_1 + L_3 \|B\|) c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$, for some $0 < \theta < 1$, are satisfied. Then, the system (7) is (ϵ, α) -attackable with probability $\delta(T(\alpha + b_x + b_{\zeta}, s_0), b_x, b_v)$ for some $\epsilon > 0$, if it holds that $\phi > b$ and $f \in \mathcal{U}_{\rho}$ with $\rho = 2L_f(b_x + b + b_{\zeta})$, $\|s_0\| \leq \phi$ and $b = \frac{c_4}{c_3 - (L_1 + L_3 \|B\|) c_4} \sqrt{\frac{c_2}{c_1} \frac{L_2 b_{\zeta}}{\theta}}$.

The theorem proof is provided in Appendix A.

From (11), c_3 can be viewed as a “measure” of the closed-loop system stability (larger c_3 means the system converges faster to the equilibrium point); on the other hand, from Theorem 2, closed-loop perception-based systems with larger c_3 are more vulnerable to stealthy attacks as the conditions of the theorem are easier to satisfy. However, if the plant’s dynamics is very unstable, $T(\alpha + b_x + b, s_0)$ is smaller for a fixed α and s_0 . Thus, the probability of attack success $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$ is larger for a fixed b_x and b_v .

It should be further noted that L'_f , used in Theorem 2 conditions, would be equal to zero for LTI systems, thus causing

³Note that under similar conditions, the same analysis also holds for asymptotic stability condition [40].

L_1 to become zero. Similarly, when the mapping Π approaches more toward linear behavior, L_3 will go to zero. Therefore, the inequality $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ holds for linear systems and linear controllers. However, larger values of c_3 will help the inequality to be satisfied even for nonlinear control systems. In simulation results described in Section V, we discuss when these conditions are satisfied in more detail.

Moreover, in the extreme case when $b_{\zeta} = 0$ (i.e., the attacker can exactly estimate the plant state), the condition $L_2 b_{\zeta} < \frac{c_3 - (L_1 + L_3 \|B\|) c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$ will be relaxed and the other condition $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ becomes less restrictive as L_1 becomes smaller. Thus, in this case, if the attacker initiates the attack with arbitrarily small s_0 , then ϵ can be arbitrarily close to zero and the attack will be very close to being strictly stealthy. Hence, the following result holds.

Corollary 1: Assume $b_{\zeta} = 0$, $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ with $L_1 = L'_f b_x$, $L_3 = L'_{\Pi}(b_x + b_v)$ and the functions f , f' (derivative of f) and Π' (derivative of Π) are Lipschitz, with constants L_f , L'_f , and L'_{Π} , respectively. If $f \in \mathcal{U}_{\rho}$ with $\rho = 2L_f(b_x + \|s_0\|)$, and $\|s_0\| \leq \frac{c_3 - c_4(L_1 + L_3 \|B\|)}{c_4(L'_f + \|B\| L'_{\Pi})}$ holds, then the system (7) is (ϵ, α) -attackable with probability $\delta(T(\alpha + b_x + \|s_0\|, s_0), b_x, b_v)$, where $\epsilon = \sqrt{1 - e^{-b_{\epsilon}}}$ for $b_{\epsilon} = (\lambda_{\max}(\Sigma_w^{-1}) + \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \times \min\{T(\alpha + b_x + \|s_0\|, s_0), \sqrt{\frac{c_2}{c_1} \frac{e^{-\beta}}{1 - e^{-\beta}}}\}) \|s_0\|^2$ and some $\beta > 0$.

Finally, the previous results depend on determining ρ such that $f \in \mathcal{U}_{\rho}$. Hence, the following result provides a sufficient condition for $f \in \mathcal{U}_{\rho}$.

Proposition 1: Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function satisfying $V(0) = 0$ and define $U_{r_1} = \{x \in B_{r_1} \mid V(x) > 0\}$. Assume that $\|\frac{\partial V(x)}{\partial x}\| \leq \beta(\|x\|)$, and for any $x \in U_{r_1}$ it holds $V(f(x)) - V(x) \geq \alpha(\|x\|)$, where $\alpha(\|x\|)$ and $\beta(\|x\|)$ are in a class of \mathcal{K} functions [40]. Furthermore, assume that r_1 can be chosen arbitrarily large. Now, the following statements hold.

- 1) If $\lim_{\|x\| \rightarrow \infty} \frac{\alpha(\|x\|)}{\beta(\|x\|)} \rightarrow \infty$, then $f \in \mathcal{U}_{\rho}$ for any $\rho > 0$.
- 2) However, if $\lim_{\|x\| \rightarrow \infty} \frac{\alpha(\|x\|)}{\beta(\|x\|)} = \gamma$, then $f \in \mathcal{U}_{\rho}$ for any $\rho < \gamma$.

Proof: We prove the first case and the second case can be shown in a similar way.

If $\lim_{\|x\| \rightarrow \infty} \frac{\alpha(\|x\|)}{\beta(\|x\|)} \rightarrow \infty$, there exists a bounded ball with radius r_2 and center at zero (referred to as B_{r_2}) such that for all $x \in S$ with $S = \{U_{r_1} - B_{r_2}^o\}$ it holds that $\frac{\alpha(\|x\|)}{\beta(\|x\|)} > \rho$. Since the function V is differentiable, using the mean-value theorem for the dynamics $x_{t+1} = f(x_t) + d_t$ with $\|d_t\| \leq \rho$ and for any $x_t \in U_{r_1}$, we have that

$$\begin{aligned} V(x_{t+1}) - V(x_t) &= V(f(x_t) + d_t) - V(x_t) \\ &= V(f(x_t)) + \left(\frac{\partial V(x)}{\partial x} \right)^T d_t - V(x_t) \\ &\geq \alpha(\|x\|) - \beta(\|x\|) \rho. \end{aligned} \quad (21)$$

Thus, for any $x_t \in S$, we have $V(x_{t+1}) - V(x_t) > 0$. Let us define $\eta = \min\{V(x_{t+1}) - V(x_t) \mid x_t \in S \text{ and } V(x_t) \geq a_{r_2}\}$ where $a_r = \max_{x \in \partial B_r} V(x)$ for any $r > 0$. Such minimum exists as the considered set is compact and we have $\eta > 0$ in U_{r_1} . Let us also assume $x_0 = \{x \in \partial B_{r_2} \mid V(x) = a_{r_2}\}$. Now,

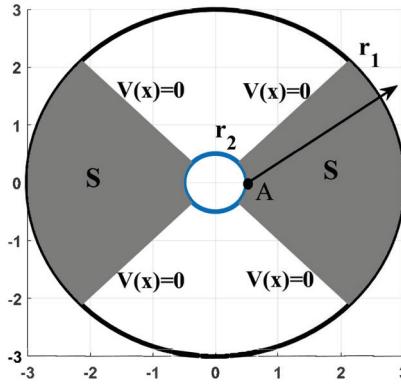


Fig. 2. Trajectory of the dynamical system in Example 1.

we claim that the trajectories starting from x_0 should leave the set U_{r_1} through the boundaries of B_{r_1} .

To show this, we know that for any $x_t \in S$, $V(x_t) \geq a_{r_2}$ since $V(x_{t+1}) - V(x_t) \geq \eta > 0$. Then, for any $t > 0$

$$V(x_t) \geq V(x_0) + t\eta = a_{r_2} + \eta t. \quad (22)$$

The previous inequality shows that x_t cannot stay in the set S forever as $V(x)$ is bounded on the compact set S . On the other hand, x_t cannot leave the set S through the boundaries satisfying $V(x) = 0$ or the surface of B_{r_2} because $V(x_t) > a_{r_2}$. Thus, the trajectories should leave the set S through the surface of B_{r_1} , and as r_1 can be chosen arbitrarily large, the trajectories of x_t will diverge to become arbitrarily large. ■

Example 1: Consider the dynamical system

$$\begin{aligned} x_{1,t+1} &= 2x_{1,t} + x_{1,t}x_{2,t}^2 \\ x_{2,t+1} &= 0.5x_{2,t} \end{aligned} \quad (23)$$

and let us consider the function $V(x) = x_1^2 - x_2^2$. Hence, we have $\|\frac{\partial V(x)}{\partial x}\| \leq 2\|x\| = \beta(\|x\|)$ and

$$\begin{aligned} V(f(x_t)) - V(x_t) &= 4x_{1,t}^2 + x_{1,t}^2x_{2,t}^4 + 2x_{1,t}^2x_{2,t}^2 - 0.25x_{2,t}^2 \\ -x_{1,t}^2 + x_{2,t}^2 &= 3x_{1,t}^2 + 0.75x_{2,t}^2 + x_{1,t}^2x_{2,t}^4 + 2x_{1,t}^2x_{2,t}^2 \\ &\geq 0.75\|x_t\|^2 = \alpha(\|x\|). \end{aligned}$$

Since $\lim_{\|x\| \rightarrow \infty} \frac{\alpha(\|x\|)}{\beta(\|x\|)} \rightarrow \infty$, for any $\rho > 0$, there exists x such that $V(x_{t+1}) - V(x_t) > 0$ for the dynamics $x_{t+1} = f(x_t) + d_t$ with $\|d\| \leq \rho$. The shaded region in Fig. 2 shows the area where $V > 0$ outside B_{r_2} and inside of B_{r_1} . Moreover, any area outside of the ball B_{r_2} satisfies $V(x_{t+1}) - V(x_t) > 0$ (i.e., $V(x_t)$ is increasing over time). Now, by denoting the state value at point A with x_A , we have $V(x_A) = x_{1,A}^2 > 0$. As $V(x_t)$ is increasing outside B_{r_2} as t increases, any trajectory starting at A cannot leave the shaded region from the surface of corresponding to $V = 0$ or B_{r_2} ; it also cannot stay in the region for all t , as the trajectory evolves. Thus, the system states will eventually leave the shaded region, “exiting” from the surface of the ball B_{r_1} .

Theorem 2 provides conditions on the closed-loop system stability that may not hold in general for any system with closed-loop exponential stability. In the following theorem, we show that one can still obtain a similar (ϵ, α) -successful attack

(possibly with a larger ϵ) even when the closed-loop system only satisfies exponential stability.

Theorem 3: Assume that the closed-loop control system (10) is exponentially stable. Then, the system (7) is (ϵ, α) -attackable with probability $\delta(T(\alpha + b + b_x, s_0), b_x)$ for some $\epsilon > 0$, if $f \in \mathcal{U}_\rho$ with $\rho = 2L_f(b_x + b + b_\zeta)$ and $b = \frac{c_4}{c_3} \sqrt{\frac{c_2}{c_1} \frac{(L_f(2b_x + b_\zeta) + 2L_\Pi(b_x + b_v))}{\theta}}$.

Proof: Using the same definition as in (9), we have

$$\begin{aligned} r_{t+1} &= h(r_t, 0) + f(x_t^a) - f(x_t^a + \zeta_t) + f(e_t + \zeta_t) - f(x_t) \\ &\quad - f(r_t) + B\Pi(e_t, v_t^s) - B\Pi(x_t, v_t^s) \\ &\quad - B\Pi(r_t, 0) = h(r_t, 0) + \sigma_t. \end{aligned}$$

From the Lipschitz property of function f , it follows that $\|f(x_t^a) - f(x_t^a + \zeta_t)\| \leq L_f b_\zeta$, and with probability $\delta(T(\alpha + b + b_x, s_0), b_x)$, we have

$$\begin{aligned} \|f(x_t)\| &\leq L_f \|x_t\| \leq L_f b_x \\ \|f(e_t + \zeta_t) - f(r_t)\| &\leq L_f \|x_t + \zeta_t\| \leq L_f (b_\zeta + b_x) \\ \|\Pi(x_t, v_t^s)\| &\leq L_\Pi(\|x_t\| + \|v_t^s\|) \leq L_\Pi(b_x + b_v) \\ \|\Pi(e_t, v_t^s) - \Pi(r_t, 0)\| &\leq L_\Pi(\|x_t\| + \|v_t^s\|) \leq L_\Pi(b_x + b_v). \end{aligned}$$

Therefore, with probability $\delta(T(\alpha + b + b_x, s_0), b_x)$, it holds that $\|\sigma_t\| \leq L_f(2b_x + b_\zeta) + 2L_\Pi(b_x + b_v)$ for all $0 \leq t \leq T(\alpha + b + b_x)$.

Since the closed-loop system $r_{t+1} = h(r_t, 0)$ is exponentially stable, we use Lemma 8 to show that the dynamics $r_{t+1} = h(r_t, 0) + \sigma_t$ will remain in a bounded ball centered at zero with such probability for all $0 \leq t \leq T(\alpha + b + b_x)$. The bound is obtained by $b = \frac{c_4}{c_3} \sqrt{\frac{c_2}{c_1} \frac{(L_f(2b_x + b_\zeta) + 2L_\Pi(b_x + b_v))}{\theta}}$. Using the Data-processing inequality of KL divergence and following the same procedure as in Theorem 2, we obtain

$$\begin{aligned} &\text{KL} \left(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_{T(\alpha+b_x+b,s_0)}^a) \middle\| \mathbf{P}(Y_{-\infty} : Y_{T(\alpha+b_x+b,s_0)}) \right) \\ &\leq \sum_{i=0}^{T(\alpha+b_x+b,s_0)} \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \|r_i\|^2 \\ &\leq \sum_{i=0}^{T(\alpha+b_x+b,s_0)} \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) b = b_\epsilon \\ &\leq \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) b(T(\alpha + b_x + b, s_0) + 1) = b_\epsilon. \end{aligned}$$

This means that the system is (ϵ, α) -attackable with probability $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$ and $\epsilon = \sqrt{1 - e^{-b_\epsilon}}$. On the other hand, since we have $f \in \mathcal{U}_\rho$ with $\rho = 2L_f(b_x + b + b_\zeta)$, similarly as in Theorem 2, we can show that for $t \geq T(\alpha + b + b_x, s_0)$, the states of the underattack system will satisfy $\|x_t^a\| \geq \alpha$ with probability $\delta(T(\alpha + b + b_x, s_0), b_x)$. ■

Note that our results only focus on the existence of perception measurements $G(x_t^a - s_t)$, obtained by shifting the current perception scene by s_t , that results in (ϵ, α) -successful attack, and not how to compute it. Furthermore, to derive an attack sequence using Attack Strategy I, the attacker needs the estimation of the plant states. Thus, Attack Strategy II relaxes this assumption, with the attacker only needing to have knowledge about the plant's (open-loop) dynamics f and the computation power to calculate $s_{t+1} = f(s_t)$ ahead of time.

B. Attack Strategy II: Using Plant Dynamics

Similarly to Attack Strategy I, consider the attack sequence where z_t^a and $y_t^{s,a}$, for all $t \geq 0$, satisfy

$$\begin{aligned} z_t^a &= G(x_t^a - s_t), \quad y_t^{s,a} = C_s(x_t^a - s_t) + v_t^{s,a} \\ s_{t+1} &= f(s_t) \end{aligned} \quad (24)$$

for some nonzero s_0 . However, here the attacker does not need an estimate of the plant's state; they simply follow plant dynamics $s_{t+1} = f(s_t)$ to find the desired measurements' transformation. Now, we define the state $e_t \triangleq x_t^a - s_t$ as the *fake state*, and the attacker's intention is to make the system believe that the plant state is equal to e_t .

The following theorem captures the condition for which the system is not resilient to Attack Strategy II, in the sense that it is (ϵ, α) -successful attackable.

Theorem 4: Consider the system (7) that satisfies Assumption 2. Assume that both functions f' and Π' are Lipschitz, with constants L'_f and L'_{Π} , respectively. Moreover, assume that there exists $\phi > 0$ such that the inequalities $L_1 + L_3\|B\| < \frac{c_3}{c_4}$ and $L_2 b_x < \frac{c_3 - (L_1 + L_3\|B\|)c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$ with $0 < \theta < 1$ are satisfied, where $L_2 = L'_f(\alpha + b_x)$, $L_1 = L'_f(\alpha + \phi)$ and $L_3 = L'_{\Pi}(b_x + \phi + b_v)$. Then, the system (7) is (ϵ, α) -attackable with probability $\delta(T(\alpha + b_x + \|s_0\|, s_0), b_x)$ and $b = \frac{c_4}{c_3 - (L_1 + L_3\|B\|)c_4} \sqrt{\frac{c_2}{c_1} \frac{L_2 b_x}{\theta}}$, for some $\epsilon > 0$, if it holds that $\frac{\phi}{b} > 1$ and $f \in \mathcal{U}_0$.

The proof of the theorem is provided in Appendix B.

Unlike in Theorem 2, L_1 and L_2 in Theorem 4 increase as α increases. Therefore, unless $L'_f = 0$, one cannot claim that the attack can be ϵ -stealthy for arbitrarily large α as the inequality $L_1 + L_3\|B\| \leq \frac{c_3}{c_4}$ might not be satisfied. Therefore, there is a tradeoff between the stealthiness guarantees (ϵ) and the performance degradation caused by the attack (α). However, in an extreme case where the system is linear, it holds that $L'_f = 0$ and $L_1 = L_2 = 0$; before introducing the results for LTI systems in the next section, we remark on the following.

Remark 7: The fact that we considered the control input in a general end-to-end form can help us verify that our results also hold when the perception-based control system is modular; i.e., when the states are extracted from the perception module and physical sensors and then used with the classic control methods. This is because the end-to-end formulation of the controller is very general, also covering the modular methods. In such cases, the traditional control methods can be used such that the noiseless closed-loop system becomes exponentially or asymptotically stable.

C. Attack on LTI Systems

To derive sufficient conditions for which stealthy yet effective attacks exist, we have designed such attacks for perception-based control systems where plants have input affine nonlinear dynamics. However, for LTI plants, the obtained conditions can be significantly relaxed. Specifically, for LTI systems, because the system dynamics takes the form of $f(s_t) = As_t$, Attack Strategies I and II become identical as

$$\begin{aligned} s_{t+1} &= f(\hat{x}_t^a) - f(\hat{x}_t^a - s_t) = A(\hat{x}_t^a) - A(\hat{x}_t^a - s_t) \\ &= As_t = f(s_t). \end{aligned} \quad (25)$$

Therefore, the attacker does not need to estimate the state of the system and they can use the Attack Strategy II to design the attack sequence.

In general, assume that for such systems we use the controller satisfying Assumption 2, where due to the complexity of perception (e.g., image observations), a nonlinear function maps the perception sensing (e.g., the image) and sensors to the control input. Hence, from Corollary 1, we directly obtain the following result.

Corollary 2: Consider an LTI-perception-based control system with $f(x_t) = Ax_t$. Assume that $L_3\|B\| < \frac{c_3}{c_4}$ with $L_3 = L'_{\Pi}(b_x + b_v)$, $\|s_0\| \leq \frac{c_3 - c_4 L_3\|B\|}{c_4\|B\|L'_{\Pi}}$, and the matrix A is unstable. Then, the system is (ϵ, α) -attackable with probability $\delta(T(\alpha + b_x + \|s_0\|, s_0), b_x, b_v)$, for arbitrarily large α and $\epsilon = \sqrt{1 - e^{-b_\epsilon}}$, with $b_\epsilon = (\lambda_{\max}(\Sigma_w^{-1}) + \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1})) \times \min\{T(\alpha + b_x + \|s_0\|, s_0), \sqrt{\frac{c_2}{c_1} \frac{e^{-\beta}}{1 - e^{-\beta}}}\} \|s_0\|^2$ and some $\beta > 0$.

Note that even though the previous corollary considers LTI plants, the earlier requirement $L_3\|B\| < \frac{c_3}{c_4}$ is due to the nonlinearity of employed controllers. Both Theorems 2 and 4 assume an *end-to-end* controller that directly maps the perception and sensor measurements to the control input. However, there are controllers that first extract the state information using the perception module P and then use a feedback controller to find the control input (e.g., [6], [7]). For instance, consider a linear feedback controller with gain $K = [K_P \ K_s]$, resulting in

$$\begin{aligned} u_t &= KY_t = K_s y_t^{c,s} + K_P y_t^P \\ &= KCx_t + K_P v^P(x_t) + K_s v_t^s. \end{aligned} \quad (26)$$

Applying the earlier control input to the attack-free system

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + w_t \\ &= (A + BKC)x_t + BK_P v^P(x_t) + BK_s v_t^s. \end{aligned} \quad (27)$$

Let us assume that Assumptions 1 and 2 still hold. It is easy to show that $(A + BKC)$ needs to be a stable matrix (i.e., all eigenvalues are inside the unit circle). Then, we obtained the following result.

Theorem 5: Consider perception-based control of an LTI plant with dynamics $f(x_t) = Ax_t$, controlled with a linear feedback controller from (26). If the matrix A is unstable, the system is (ϵ, α) -attackable with probability 1 for arbitrarily large α and $\epsilon = \sqrt{1 - e^{-b_\epsilon}}$, where

$$b_\epsilon = \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \left(\frac{2\gamma T(\alpha + R_S, s_0) + \|s_0\|}{1 - \lambda_{\max}(A + BKC)} \right) \quad (28)$$

and $\lambda_{\max}(A + BKC)$ is the largest eigenvalue of the matrix $A + BKC$.

Proof: Let us assume that the attack dynamics are generated using (25); by defining $e_t = x_t^a - s_t$, we obtain $z_t^{c,a} = G(x_t^a - s_t) = G(e_t)$ and $y_t^{c,s,a} = y_t^{s,a} - C_s s_t = C_s e_t + v_t^s$. Therefore

$$\begin{aligned} e_{t+1} &= x_{t+1}^a - s_{t+1} = Ax_t^a + BK_Y - As_t \\ &= Ae_t + BK \begin{bmatrix} C_P e_t + v^P(e_t) \\ C_s e_t + v_t^s \end{bmatrix} \\ &= (A + BKC)e_t + BK_P v^P(e_t) + BK_s v_t^s. \end{aligned} \quad (29)$$

The previous dynamics follows the same dynamics as in (27). Thus, if the initial condition s_0 is chosen small enough, e will

remain in the set \mathcal{S} and it would hold that $\|v^P(e_t)\| \leq \gamma$ for all $t \geq 0$. Now, by defining $r_t = e_t - x_t$, we have

$$r_{t+1} = (A + BKC)r_t + v^P(e_t) - v^P(x_t). \quad (30)$$

Since we have $\|v^P(e_t)\| \leq \gamma$ and $\|v^P(x_t)\| \leq \gamma$, and the matrix $(A + BKC)$ is stable, we have

$$\begin{aligned} \|r_t\| &= \|(A + BKC)^t r_0\| \\ &\quad + \sum_{i=0}^{t-1} (A + BKC)^{t-i-1} (v^P(e_i) - v^P(x_i))\| \\ &\leq |\lambda_{\max}(A + BKC)|^t \|s_0\| + \frac{2\gamma}{1 - \lambda_{\max}(A + BKC)} \end{aligned}$$

where we used $r_0 = s_0$ and the squared matrix property of $\|A^t v\| \leq |\lambda_{\max}(A)|^t \|v\|$ for any $v \in \mathbb{R}^n$ and $A \in \mathbb{R}^{n \times n}$. From the proof of Theorem 2, we obtain

$$\begin{aligned} \text{KL} \left(\mathbf{Q}(Y_0^a : Y_{T(\alpha+R_S, s_0)}^a) \middle\| \mathbf{P}(Y_0 : Y_{T(\alpha+R_S, s_0)}) \right) \\ \leq \sum_{i=0}^{T(\alpha+R_S, s_0)} \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \|r_i\|^2 \\ = \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \sum_{i=0}^{T(\alpha+R_S, s_0)} \|r_i\|^2 \\ = \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \left(\frac{2\gamma T(\alpha+R_S, s_0) + \|s_0\|}{1 - \lambda_{\max}(A + BKC)} \right). \end{aligned}$$

Since the matrix A is unstable, choosing $s_0 = cq_i$ where $c > 0$ is a scalar and q_i is the unstable eigenvector with associated eigenvalue $|\lambda_i| > 1$, results in $s_t = c\lambda_i^t q_i$; thus, s_t becomes arbitrarily large for large enough t .

Now, as $T(\alpha+R_S, s_0)$ is the time such that $\|s_t\| \geq \alpha + R_S$, for all $t \geq T(\alpha+R_S, s_0)$, it holds that

$$\begin{aligned} \|s_t\| - \|x_t^a\| &\leq \|x_t^a - s_t\| = \|e_t\| \leq R_S \\ \Rightarrow \|x_t^a\| &\geq R_S + \alpha - R_S = \alpha. \end{aligned} \quad (31)$$

■

Finally, note that for LTI plants, the attacker will be effective if and only if s_0 is not orthogonal to all unstable eigenvectors of the matrix A . Therefore, if γ is small enough, by choosing such s_0 that is also arbitrarily close to zero, the attacker can be ϵ -stealthy with ϵ being arbitrarily small. On the other hand, γ is not controlled by the attacker (rather, it is property of the controller design) and if γ is large, there may be no stealthiness guarantee for the attacker. While this implies that having large γ improves the resiliency of the systems, it should be noted that large γ is not desirable from the control's perspective, as it would degrade the control performance when the system is free of attack.

The instability condition for LTI plants in Theorem 5 is inline with the results from the work in [11] and [14] focused on LTI systems with linear controllers and without perception-based control, as well as with the notion of perfect attackability introduced there for the specific ADs (i.e., χ^2) considered in those works. On the other hand, in this work, we consider a general notion of stealthiness (being stealthy from any AD) and provide analysis for perception-based control systems.

Remark 8: It is important to note that the results obtained in this work capture sufficient conditions for the existence of (ϵ, α) attacks on a system. However, this does not imply that

systems whose model parameters do not satisfy these conditions are secure against attacks. While it can be shown that the open-loop stability of a system can improve the system's resiliency against stealthy attacks, a formal analysis of this aspect is beyond the scope of this work.

V. SIMULATION RESULTS

We illustrate and evaluate our methodology for vulnerability analysis of perception-based control systems on two case studies, 1) inverted pendulum and 2) AVs.

A. Inverted Pendulum

We consider a fixed-base inverted pendulum equipped with an end-to-end controller and a perception module that estimates the pendulum angle from camera images. By using $x_1 = \theta$ and $x_2 = \dot{\theta}$, the inverted pendulum dynamics can be modeled in the state-space form as

$$\begin{aligned} \dot{x}_1 &= x_2 \\ \dot{x}_2 &= \frac{g}{r} \sin x_1 - \frac{b}{mr^2} x_2 + \frac{L}{mr^2} \end{aligned} \quad (32)$$

where θ is the angle of the pendulum rod from the vertical axis measured clockwise, b is the Viscous friction coefficient, r is the radius of inertia of the pendulum about the fixed point, m is the mass of the pendulum, g is the acceleration due to gravity, and L is the external torque that is applied at the fixed base [51]. Finally, we assumed $m = 0.2\text{kg}$, $b = 0.1 \frac{\text{m} \times \text{kg}}{\text{s}}$, and $r = 0.3\text{ m}$ and discretized the model with $T_s = 0.01\text{s}$.

Using Lyapunov's indirect method, one can show that the origin of the previous system is unstable because the linearized model has an unstable eigenvalue. However, the direct Lyapunov method can help us to find the whole unstable region $-\pi < \theta < \pi$ (see [40]). We used a data set \mathcal{S} with 500 sample pictures of the fixed-base inverted pendulum with different angles in $(-\pi, \pi)$ to train a DNN P (perception module) to estimate the angle of the rod. The angular velocity is also measured directly by the sensor with a noise variance of $\Sigma_v = 0.0001$. We also trained a deep-reinforcement-learning-based controller directly mapping the image pixels and angular velocity values to the input control. For anomaly detection, we designed a standard χ^2 extended Kalman filter AD that receives the perception module's P output and angular velocity and outputs the residue/anomaly alarm.

After discretizing the model with sampling rate of $dt = 0.01\text{s}$, one can verify that $L'_f = 0.33$, $\|B\| = 0.556$, and we also found that $L'_{\Pi} = 0.12$ for the controller (note here that the inputs of Π , as defined in Section III-B2, are $x = [\theta \ \dot{\theta}]^T$ and v_t^s). We also considered a Lyapunov function with $c_1 = c_2 = 0.5$, $c_3 = 0.057$, and $c_4 = 1$ that satisfies the inequalities from (11). On the other hand, in more than 100 experiments without the attack, each simulated for a period of 5 s (i.e., 500 time steps), we found that for $b_x = 0.2$ and $b_v = 0.05$, the inequalities in Definition 5 were always satisfied—i.e., holding with the probability of almost one. Now, one can verify that the condition in Theorem 2 is satisfied if $b_{\zeta} < 0.051$.

We first chose $s_0 = [0.001 \ 0.001]^T$ and used Attack Strategy I to design attacks. To derive the current adversarial image at each time step, the attacker exploits the actual image and compromises it by deviating the pendulum rod by s_t degrees. This compromised image is then delivered to the controller, to be

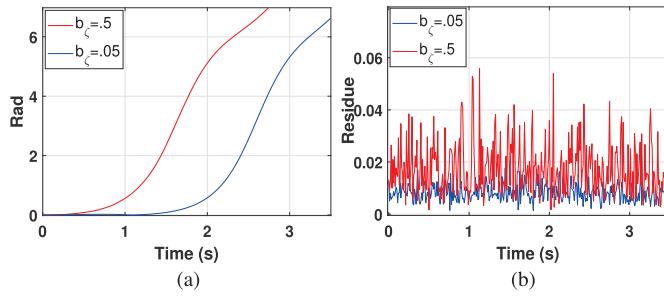


Fig. 3. (a) Evolution of the angle's (θ) absolute value over time for different levels of b_ζ . (b) Norm of the residue over time when the attack starts at time $t = 0$.

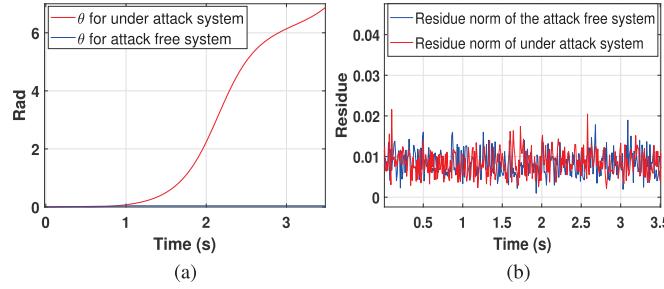


Fig. 4. (a) Angle's (θ) absolute value over time for Attack Strategy II (red) and normal condition (blue). (b) Residue norm over time for both underattack and attack-free systems.

used by the perception module to evaluate the system state. The other sensor's measurements are also compromised accordingly. Again, the attacker does not need to have access to the perception map P ; the knowledge about the dynamics f and estimate of the current plant state \hat{x}_t^a is sufficient to craft the perturbed images.

Fig. 3(a) shows the actual pendulum rod angle for different estimation uncertainty levels b_ζ (by the attacker) when the attack starts at $t = 0$. In both cases, the attacker can drive the pendulum rod into an unsafe region. Fig. 3(b) shows the residue signal over time; the attack stealthiness level decreases as b_ζ increases, consistent with our results in Section IV.

Fig. 4(b) presents the residue of the system in normal operating conditions (i.e., without attack) as well as under Attack Strategy II. We can see that the residue level of both Attack Strategy I with $b_\zeta = 0.05$ and Attack Strategy II are the same as for the system without attack. The red and blue lines in Fig. 4(a) also show the pendulum rod angle trajectory for Attack Strategy II and normal conditions, respectively.

B. Autonomous Vehicle

We consider the nonlinear dynamical model of an AV from the work in [52], with four states $[x \ y \ \psi \ v]^T$ in the form

$$\begin{aligned} \dot{x} &= v \cos(\psi + \beta), \quad \dot{y} = v \sin(\psi + \beta) \\ \dot{\psi} &= \frac{v}{l_r} \sin(\beta), \quad \dot{v} = a, \quad \beta = \tan^{-1} \left(\frac{l_r}{l_f + l_r} \tan(\delta_f) \right) \end{aligned} \quad (33)$$

where x and y represent the position of the center of mass in x - and y -axis, respectively, ψ is the inertial heading, v is the velocity of the vehicle, and β is the angle of the current velocity

of the center of mass with respect to the longitudinal axis of the car. Also, $l_f = 1.1$ and $l_r = 1.73$ m are the distance from the center of the mass of the vehicle to the front and rear axles, respectively, and a is the acceleration. The control inputs are the front steering angle δ_f and acceleration a . We assume only ψ is measured directly using noisy sensors, with zero-mean noise with variance $\Sigma_v = 0.0001$. Moreover, there is a camera affixed to a dashboard that receives the image of the scene in front of the car. The system noise is assumed to be zero-mean, with covariance $\Sigma_w = 0.0001$.

We consider the scenario where the car has a constant speed of 25 m/s. Since the control objective is to keep the vehicle within the lane (i.e., between the lines of the road), the relative position of the camera with respect to the two lines is the essential information used by the controller. Therefore, even if the image has some other backgrounds, we assume that the lines of the road would be detected, to allow for extracting the relevant information about the states of the system (including the position of the car with respect to the lane center). Using deep reinforcement learning, we trained a controller that takes the images (containing the lines) and the inertial heading measurements ψ as the observations and maps them to the control input δ_f that keeps the car between the lanes. As the reward function, we assign a higher reward when the car is between the lines and a lower reward when the car is moving further from the center of the lane. To find the position of the car with respect to the lane center, we trained another DNN (perception map) that takes the images and renders the position y in state space as follows $P(z_t) = y_t + v^P(y_t)$.

This information is used in AD, where we used $P(z_t)$ and measured ψ to find the residue for both χ^2 and CUSUM detectors. We assume the system is equipped with χ^2 [12] and CUSUM ADs [53], where an extended Kalman filter is used to find the residue signal. The thresholds are set to have $p^{\text{FA}} = 0.05$ as a false alarm rate. Although the system dynamics (33) is not in the form of (7), the car's kinematics in lateral movement can be approximated well by an LTI model with matrix A around

$$\text{operating point } (x_e, 0, 0, 0) \text{ as } A = \begin{bmatrix} 1 & 0 & 0 & dt \\ 0 & 1 & 25dt & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}; \text{ the model}$$

was obtained by linearizing the state dynamics with sampling time of $dt = 0.01$ s, where x_e is any arbitrary value. Then, we considered the LTI-based attack design using as $s_{t+1} = As_t$ for some small nonzero initial condition $s_0 = 0.001 [0 \ 1 \ 1 \ 0]^T$. Since the open-loop model is linear, we verify the conditions of Corollary 2; the only condition that needs to be satisfied in this case is $L_3 \|B\| < c_3$.

We obtained $L'_\Pi = 0.23$ (here, Π 's arguments are y , ψ , and v_t^s) and considered a Lyapunov function with $c_1 = c_2 = 0.5$, $c_3 = 0.032$ and $c_4 = 1$ that satisfies the conditions from (11). Also, when running 100 experiments without the attack for, each for a simulation period of 30 seconds (i.e., 3000 time steps) in all experiments, we obtained $\sup_{0 \leq t \leq 3000} \|x_t\| < 0.2$, $\sup_{0 \leq t \leq 3000} \|v_t^s\| < 0.05$, which means that for $b_x = 0.2$ and $b_v = 0.05$, with a probability of almost one the system stays in the zero center ball, as in Definition 5. Now, we have $L_3 = L'_\Pi(b_x + b_v + \phi) = 0.0532$ where the upper bound on the norm of the initial condition of attack states s is considered to be 0.0014. Thus, having $\|B\| = 0.556$ (obtained by linearizing and discretization)

$$L_3 \|B\| = 0.0106 < 0.032.$$

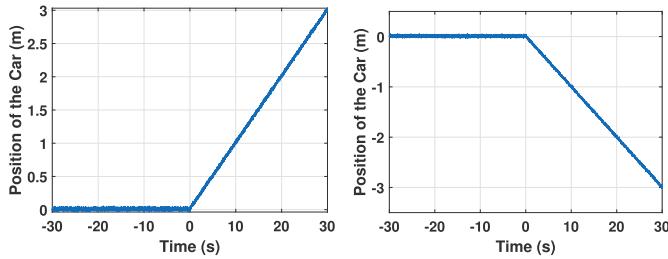


Fig. 5. (Left) Position of the car with respect to the road center over time for the initial condition $s_0 = 0.001 [0 \ 1 \ 1 \ 0]^T$ when the attack starts at time $t = 0$. (Right) Position of the car with respect to the road center over time for the initial condition $s_0 = -0.001 [0 \ 1 \ 1 \ 0]^T$ when the attack starts at time $t = 0$.

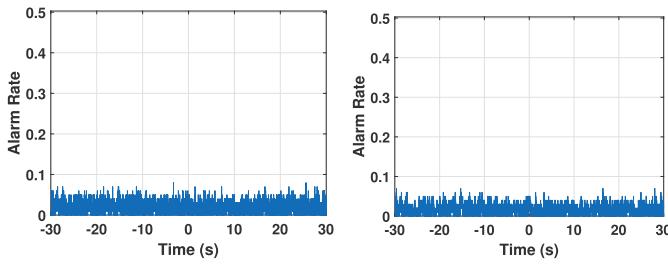


Fig. 6. (Left) Average number of alarms at each time step for χ^2 -based intrusion detector when the attack starts at time $t = 0$. (Right) Average number of alarms at each time step for CUSUM-based AD when the attack starts at time $t = 0$.

Now, $y_t^a - s_{t,2}$ provides the desired position of the camera (center of the car) in the Y -axis (where $s_{t,2}$ is the second element of s_t). To find the attacked image, we distorted the current image by shifting some pixels to right or left (depending on the sign of $s_{t,2}$) in order to have the camera placed in $y_t^a - s_{t,2}$. Another approach to find the attacked image is to use some prerecorded images of the road where the car is placed at different distances from the lane center. Then, the goal would be to find an image from those prerecorded images whose distance from the lane center is closest to $y_t^a - s_{t,2}$.

Fig. 5 shows the position of the center of the car with respect to the road center when the attack starts at time zero. If the attacker chooses the initial condition $s_0 = 0.001 [0 \ 1 \ 1 \ 0]^T$ the car will deviate to the left-hand side of the road (left figure) while choosing $s_0 = -0.001 [0 \ 1 \ 1 \ 0]^T$ will push the car to the right-hand side (right figure).

Fig. 6 illustrates the average number of alarms at each time step for both χ^2 (left figure) and CUSUM (right figure) ADs in 1000 experiments when the attack starts at time $t = 0$. As shown, the values of true alarm averages (for $t > 0$) are the same as the false alarm averages (for $t \leq 0$) for both ADs, which indicates the stealthiness of the attack according to Definition 3.

VI. CONCLUSION

In this work, we have considered the problem of resiliency under sensing and perception attacks for perception-based control systems, focusing on a general class of nonlinear dynamical plants. We have assumed that the noiseless closed-loop system

equipped with an end-to-end controller and AD is exponentially stable on a set around the equilibrium point. We have introduced a general notion of ϵ -stealthiness as a measure of difficulty in attack detection from the set of perception measurements and sensor values. Furthermore, we have derived sufficient conditions for effective yet ϵ -stealthy attack sequences to exist. Here, the control performance degradation has been considered as moving the system state outside of the safe region defined by a bounded ball with radius α , resulting in an (ϵ, α) -successful attack. Finally, we illustrated our results on two case studies: 1) fixed-base inverted pendulum and 2) vehicle lane-keeping.

A Proof of Theorem 2

Theorem 2: We need to show that the sequence of falsified perception (e.g., images) and physical sensor values $\{z_0^{c,a}, y_0^{c,s,a}\}, \{z_1^{c,a}, y_1^{c,s,a}\}, \dots$, obtained by Attack Strategy I is (ϵ, α) -successful attack for arbitrarily large α . By defining $e_t = x_t^a - s_t$ and $r_t = e_t - x_t$, we get $z_t^{c,a} = G(e_t)$, $y_t^{c,s,a} = C_s e_t + v_t^s$ and

$$\begin{aligned} e_{t+1} &= f(x_t^a) - f(x_t^a + \zeta_t) + f(e_t + \zeta_t) + B\Pi(e_t, v_t^s) \\ &\quad + w_t^a \\ r_{t+1} &= h(r_t, 0) + f(x_t^a) - f(x_t^a + \zeta_t) + f(e_t + \zeta_t) - f(x_t) \\ &\quad - f(r_t) + B\Pi(e_t, v_t^s) - B\Pi(x_t, v_t^s) - B\Pi(r_t, 0) \\ &= h(r_t, 0) + \sigma'_1 + \sigma'_2 \end{aligned} \quad (34)$$

with $\sigma'_1 = f(x_t^a) - f(x_t^a + \zeta_t) + f(e_t + \zeta_t) - f(x_t) - f(r_t)$ and $\sigma'_2 = B\Pi(G(e_t), y_t^{s,a}) - B\Pi(G(x_t), y_t^s) - B\Pi(G(r_t), C_s r_t)$.

Using the mean-value theorem and equality $x_t^a = r_t + s_t + x_t$, we have that

$$\begin{aligned} f(x_t^a + \zeta_t) &= f(x_t^a) + \frac{\partial f}{\partial x}|_{\Theta(x^a, x^a + \zeta)} \zeta_t \\ f(e_t + \zeta_t) &= f(x_t) + \frac{\partial f}{\partial x}|_{\Theta(x, x + r + \zeta)} (r_t + \zeta_t) \\ f(r_t) &= f(0) + \frac{\partial f}{\partial x}|_{\Theta(0, r)} r_t = \frac{\partial f}{\partial x}|_{\Theta(0, r)} r_t \end{aligned} \quad (35)$$

where for instance, $\frac{\partial f}{\partial x}|_{\Theta(x, y)} = \begin{bmatrix} \nabla f_1(c_1 x + (1 - c_1)y) \\ \vdots \\ \nabla f_n(c_n x + (1 - c_n)y) \end{bmatrix}$ for some $0 \leq c_1, \dots, c_n \leq 1$ where f_i is the i th element of function f . Therefore, we get that

$$\begin{aligned} \sigma'_1 &= \left(\frac{\partial f}{\partial x}|_{\Theta(x^a, x^a + \zeta)} - \frac{\partial f}{\partial x}|_{\Theta(x, x + r + \zeta)} \right) \zeta_t \\ &\quad + \left(\frac{\partial f}{\partial x}|_{\Theta(x, x + r + \zeta)} - \frac{\partial f}{\partial x}|_{\Theta(0, r)} \right) r_t. \end{aligned} \quad (36)$$

Similarly, using the mean-value theorem and $e_t = r_t + x_t$, we obtain

$$\begin{aligned} \Pi(e_t, v_t^s) &= \Pi(x_t, v_t^s) + \frac{\partial \Pi}{\partial e_t}|_{\Theta(x, x + r)} r_t \\ \Pi(r_t, 0) &= \Pi(0, 0) + \frac{\partial \Pi}{\partial r_t}|_{\Theta(0, r)} r_t \end{aligned} \quad (37)$$

where $\Pi(0, 0) = 0$. Therefore, we get

$$\sigma'_2 = B \left(\frac{\partial \Pi}{\partial e_t} \Big|_{\Theta((x, v), (x+r, v))} - \frac{\partial \Pi}{\partial r_t} \Big|_{\Theta((0, 0), (r, 0))} \right) r_t. \quad (38)$$

Using the fact that the functions f and $\frac{\partial f}{\partial x}(x)$ are Lipschitz, for all $0 \leq t \leq T(\alpha + b + b_x, s_0)$, it holds that

$$\begin{aligned} & \left\| \frac{\partial f}{\partial x} \Big|_{\Theta(x^a, x^a + \zeta)} - \frac{\partial f}{\partial x} \Big|_{\Theta(x, x+r+\zeta)} \right\| \\ & \leq L'_f (\Theta(x^a, x^a + \zeta) - \Theta(x, x+r+\zeta)) \leq L'_f (\|x_t^a\| + \|x_t\| + \|r_t\| \\ & \quad + \|\zeta_t\|) \leq L'_f (\alpha + b_x + b_\zeta) + L'_f \|r_t\| \end{aligned} \quad (39)$$

where we used our assumption that $\|x_t\| \leq b_x$ with probability $\delta(T(\alpha + b + b_x, s_0), b_x, b_v)$. Moreover, Lipschitz's assumption on f and the boundedness of ζ also results in

$$\left\| \frac{\partial f}{\partial x} \Big|_{\Theta(x^a, x^a + \zeta)} - \frac{\partial f}{\partial x} \Big|_{\Theta(x, x+r+\zeta)} \right\| \leq 2L_f. \quad (40)$$

Therefore, we get $\left\| \frac{\partial f}{\partial x} \Big|_{\Theta(x^a, x^a + \zeta)} - \frac{\partial f}{\partial x} \Big|_{\Theta(x, x+r+\zeta)} \right\| \leq \min\{2L_f, L'_f(\alpha + b_x + b_\zeta) + L'_f \|r_t\|\} \leq \min\{2L_f, L'_f(\alpha + b_x + b_\zeta) + L'_f \|r_t\|\}$. Similarly, we have

$$\begin{aligned} & \left\| \frac{\partial f}{\partial x} \Big|_{\Theta(x, x+r+\zeta)} - \frac{\partial f}{\partial x} \Big|_{\Theta(0, r)} \right\| \leq L'_f (\Theta(x, x+r+\zeta) - \Theta(0, r)) \\ & \leq L'_f (\|x_t\| + \|r_t\| + \|\zeta_t\|) \leq L'_f (b_x + b_\zeta) + L'_f \|r_t\|. \end{aligned} \quad (41)$$

Therefore, $\|\sigma'_1\| \leq \min\{2L_f, L'_f(\alpha + b_x + b_\zeta)\} b_\zeta + L'_f b_\zeta \|r_t\| + L'_f (b_x + b_\zeta) \|r_t\| + L'_f \|r_t\|^2$. Now, consider the bounded set of B_ϕ satisfying $B_\phi \subset \mathcal{D}$, where we know $\|r_t\|^2 \leq \phi \|r_t\|$ for all $r_t \in B_\phi$, and we get

$$\begin{aligned} & \|\sigma'_1\| \leq \min\{2L_f, L'_f(\alpha + b_x + b_\zeta)\} b_\zeta \\ & \quad + L'_f (b_x + 2b_\zeta + \phi) \|r_t\| = L_1 \|r_t\| + L_2 b_\zeta. \end{aligned} \quad (42)$$

In addition, since the function Π' is Lipschitz, for all $x \in \mathcal{D}$ and $r_t \in B_\phi$ with probability $\delta(T(\alpha + b + b_x, s_0), b_x, b_v)$ it holds that

$$\begin{aligned} & \left\| \frac{\partial \Pi}{\partial e_t} \Big|_{\Theta((x, v), (x+r, v))} - \frac{\partial \Pi}{\partial r_t} \Big|_{\Theta((0, 0), (0, r))} \right\| \leq L'_\Pi (\|x_t + r_t + v_t^s\| \\ & \leq L'_\Pi (\phi + b_x + b_v) \end{aligned}$$

and we get $\|\sigma'_2\| \leq L_3 \|B\| \|r_t\|$ with $L_3 = L'_\Pi (b_x + \phi + b_v)$; this results in

$$\|\sigma'_1 + \sigma'_2\| \leq L_2 b_\zeta + (L_1 + L_3 \|B\|) \|r_t\|. \quad (43)$$

Since for $r_t \in B_\phi$, we have that $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ and $L_2 b_\zeta < \frac{c_3 - (L_1 + L_3 \|B\|) c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$, using Lemma 8 for all $\|r_0\| = \|s_0\| < \min\{\sqrt{\frac{c_1}{c_2}} d, \phi\}$, there exists $t_1 > 0$, such that for all $t < t_1$, it holds that $\|r_t\| \leq \sqrt{\frac{c_2}{c_1}} e^{-\beta t} \|s_0\|$ with $\beta = \frac{(1-\theta)(c_3 - (L_1 + L_3 \|B\|) c_4)}{2c_2}$, and $\|r_t\| \leq b$ with $b = \frac{c_4}{c_3 - (L_1 + L_3 \|B\|) c_4} \sqrt{\frac{c_2}{c_1} \frac{L_2 b_\zeta}{\theta}}$ for $t \geq t_1$. It should be noted that our assumption $\phi > b$ (or equivalently $B_b \subset B_\phi$) guarantees that the trajectory of r_t will remain in the set B_ϕ .

Now, we need to show that for $t \geq T(\alpha + b_x + b, s_0)$, we get that $\|x_t^a\| \geq \alpha$. Since the function f is differentiable, using the

mean-value theorem, we obtain

$$\begin{aligned} s_{t+1} &= f(s_t + \zeta_t + e_t) - f(\zeta_t + e_t) \\ &= f(s_t) + \frac{\partial f}{\partial x} \Big|_{(s_t, s_t + e_t + \zeta_t)} (e_t + \zeta_t) - f(e_t + \zeta_t). \end{aligned} \quad (44)$$

Since $e_t = r_t + x_t$, for $0 \leq t \leq T(\alpha + b_x + b, s_0)$ with probability $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$, we have that

$$\begin{aligned} & \left\| \frac{\partial f}{\partial x} \Big|_{(s_t, s_t + e_t + \zeta_t)} (e_t + \zeta_t) - f(e_t + \zeta_t) \right\| \\ & \leq 2L_f \| (r_t + x_t + \zeta_t) \| \\ & \leq 2L_f (b_x + b + b_\zeta). \end{aligned}$$

Since we assume that $f \in \mathcal{U}_\rho$ with $\rho = 2L_f (b_x + b + b_\zeta)$, there exists s_0 such that s_t becomes arbitrarily large. Using Definition 2, $T(\alpha + b_x + b, s_0)$ is defined as the first time step that satisfies $\|s_t\| \geq \alpha + b_x + b$. On the other hand, using $e_t = x_t^a - s_t$ and $e_t = x_t + r_t$, it holds that

$$\|x_t^a\| \geq \|s_t\| - \|e_t\| \geq \|s_t\| - \|x_t\| - \|r_t\| \geq \alpha.$$

We now need to show that the designed Y_t^a for $t \geq 0$ satisfies the stealthiness condition; i.e., we need to show that $\text{KL}(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_t^a) \parallel \mathbf{P}(Y_{-\infty} : Y_t)) \leq \log(\frac{1}{1-\epsilon^2})$ for some $\epsilon > 0$. Since the sequences $Y_0^a : Y_t^a$, $Y_0 : Y_t$, and $Y_{-\infty}^{-1}$ are generated by $e_0 : e_t$, $x_0 : x_t$, and $x_{-\infty} : x_{-1}$, respectively, using the data-processing inequality of the KL divergence, it holds that

$$\begin{aligned} & \text{KL}(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_t^a) \parallel \mathbf{P}(Y_{-\infty} : Y_t)) \\ & \leq \text{KL}(\mathbf{Q}(x_{-\infty} : x_{-1}, e_0 : e_t) \parallel \mathbf{P}(x_{-\infty} : x_t)). \end{aligned} \quad (45)$$

On the other hand, by defining $Z_t = \begin{bmatrix} x_t \\ y_t^{c,s} \end{bmatrix}$ and $Z_t^e = \begin{bmatrix} e_t \\ y_t^{c,s,a} \end{bmatrix}$, and using monotonicity from Lemma 2 it holds that

$$\begin{aligned} & \text{KL}(\mathbf{Q}(x_{-\infty} : x_{-1}, e_0 : e_t) \parallel \mathbf{P}(x_{-\infty} : x_t)) \\ & \leq \text{KL}(\mathbf{Q}(Z_{-\infty} : Z_{-1}, Z_0^e : Z_t^e) \parallel \mathbf{P}(Z_{-\infty} : Z_t)) \end{aligned} \quad (46)$$

Then, we apply the chain-rule property of KL-divergence on the right-hand side of (46) to obtain the following:

$$\begin{aligned} & \text{KL}(\mathbf{Q}(Z_{-\infty} : Z_{-1}, Z_0^e : Z_t^e) \parallel \mathbf{P}(Z_{-\infty} : Z_t)) \\ & = \text{KL}(\mathbf{Q}(Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(Z_{-\infty} : Z_{-1})) \\ & \quad + \text{KL}(\mathbf{Q}(Z_0^e : Z_t^e | Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(Z_0 : Z_t | Z_{-\infty} : Z_{-1})) \\ & = \text{KL}(\mathbf{Q}(Z_0^e : Z_t^e | Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(Z_0 : Z_t | Z_{-\infty} : Z_{-1})) \end{aligned} \quad (47)$$

where we used the fact that the KL-divergence of two identical distributions (i.e., $\mathbf{Q}(Z_{-\infty} : Z_{-1})$ and $\mathbf{P}(Z_{-\infty} : Z_{-1})$ since the system is not under attack for $t < 0$) is zero. Using the chain rule property of the KL divergence, we have that

$$\begin{aligned} & \text{KL}(\mathbf{Q}(Z_0^e : Z_t^e | Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(Z_0 : Z_t | Z_{-\infty} : Z_{-1})) \\ & = \text{KL}(\mathbf{Q}(e_0 | Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(x_0 | Z_{-\infty} : Z_{-1})) \\ & \quad + \text{KL}(\mathbf{Q}(y_0^{c,s,a} | e_0, Z_{-\infty} : Z_{-1}) \parallel \mathbf{P}(y_0^{c,s} | x_0, Z_{-\infty} : Z_{-1})) \\ & \quad + \dots + \text{KL}(\mathbf{Q}(e_t | Z_{-\infty} : Z_{t-1}) \parallel \mathbf{P}(x_t | Z_{-\infty} : Z_{t-1})) \\ & \quad + \text{KL}(\mathbf{Q}(y_t^{c,s,a} | e_t, Z_{-\infty} : Z_{t-1}) \parallel \mathbf{P}(y_t^{c,s} | x_t, Z_{-\infty} : Z_{t-1})). \end{aligned}$$

Given $Z_{-\infty} : Z_{t-1}$, the distribution of x_t is a Gaussian with some mean $\mu(x_{t-1})$ and covariance Σ_w written as $x_t = \mu(x_{t-1}) +$

w_{t-1} . Similarly using (34) given $Z_{-\infty} : Z_{-1}, Z_0^e : Z_{t-1}^e$, the distribution of e_t is a Gaussian with some mean $\mu(e_{t-1})$ and covariance Σ_w written as $e_t = \mu(e_{t-1}) + w_{t-1}$. Therefore, we get $r_t = e_t - x_t = \mu(e_{t-1}) - \mu(x_{t-1})$. Using Lemmas 3 and 4 for all $t \geq 0$, it holds that

$$\begin{aligned} \text{KL}(\mathbf{Q}(e_t | Z_{-\infty} : Z_{t-1}^e) || \mathbf{P}(x_t | Z_{-\infty} : Z_{t-1})) \\ = \mathbf{E}_{\mathbf{Q}(Z_{-\infty} : Z_{t-1}^e)} \{r_t^T \Sigma_w^{-1} r_t\} \\ \leq \mathbf{E}_{\mathbf{Q}(e_t | Z_{-\infty} : Z_{t-1}^e)} \{\lambda_{\max}(\Sigma_w^{-1}) \|r_t\|^2\}. \end{aligned} \quad (48)$$

On the other hand, given x_t , the distribution of $y_t^{c,s}$ is a Gaussian with mean $C_s x_t$ and covariance Σ_{v^s} . Similarly, given e_t , the distribution of $y_t^{c,s,a}$ is a Gaussian with mean $C_s e_t$ and covariance Σ_{v^s} . Using Lemmas 4 and 3 for all $t \geq 0$

$$\begin{aligned} \text{KL}(\mathbf{Q}(y_t^{c,s,a} | e_t, Z_{-\infty} : Z_{t-1}^e) || \mathbf{P}(y_t^{c,s} | x_t, Z_{-\infty} : Z_{t-1})) \\ = \text{KL}(\mathbf{Q}(y_t^{c,s,a} | e_t) || \mathbf{P}(y_t^{c,s} | x_t)) \\ \leq \mathbf{E}_{\mathbf{Q}(y_t^{c,s,a} | e_t)} \{r_t^T C_s^T \Sigma_{v^s}^{-1} C_s r_t\}. \end{aligned} \quad (49)$$

The arguments inside the expectations in the right-hand side of the inequalities (48) and (49) are upper bounded by $\lambda_{\max}(\Sigma_w^{-1}) \|r_t\|^2$ and $\lambda_{\max}(C_s^T \Sigma_{v^s}^{-1} C_s) \|r_t\|^2$, respectively, using the norm property $x^T Q x \leq \lambda_{\max}(Q) \|x\|^2$. On the other hand, since $\|r_t\|$ is bounded from above, using Lemma 5, we can find an upper bound for (48) and (49).

Specifically, if $T(\alpha + b_x + b, s_0) < t_1$, then

$$\begin{aligned} \sum_{i=0}^{T(\alpha+b_x+b,s_0)} \|r_i\|^2 &\leq \min \left\{ T(\alpha + b_x + b, s_0) + 1, \sqrt{\frac{c_2}{c_1}} \frac{e^{-\beta}}{1 - e^{-\beta}} \right\} \\ &\quad \times \|s_0\|^2 \end{aligned}$$

with probability $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$. However, if $T(\alpha + b_x + b, s_0) \geq t_1$, then

$$\begin{aligned} \sum_{i=0}^{T(\alpha+b_x+b,s_0)} \|r_i\|^2 &\leq \min \left\{ t_1, \sqrt{\frac{c_2}{c_1}} \frac{e^{-\beta}}{1 - e^{-\beta}} \right\} \|s_0\|^2 \\ &\quad + (T(\alpha + b_x + b, s_0) + 1 - t_1)b \end{aligned}$$

with probability $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$. Using the inequalities (45)–(49), we get

$$\begin{aligned} \text{KL}(\mathbf{Q}(Y_{-\infty}^{-1}, Y_0^a : Y_{T(\alpha+b_x+b,s_0)}^a) || \mathbf{P}(Y_{-\infty} : Y_{T(\alpha+b_x+b,s_0)})) \\ \leq \lambda_{\max}(C_s^T \Sigma_{v^s}^{-1} C_s + \Sigma_w^{-1}) \\ \times \max \{ \min \{T(\alpha + b_x + b, s_0) + 1, \\ \sqrt{\frac{c_2}{c_1}} \frac{e^{-\beta}}{1 - e^{-\beta}}\} \|s_0\|^2 \\ \min \left\{ t_1, \sqrt{\frac{c_2}{c_1}} \frac{e^{-\beta}}{1 - e^{-\beta}} \right\} \|s_0\|^2 \\ + (T(\alpha + b_x + b, s_0) + 1 - t_1)b \} \\ = b_\epsilon \end{aligned}$$

which means that the system is (ϵ, α) -attackable with the probability of $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$ and $\epsilon = \sqrt{1 - e^{-b_\epsilon}}$.

B. Proof of Theorem 4

Theorem 4: We need to show that the sequence of compromised perception (e.g., images) and sensor values $\{z_0^a, y_0^{s,a}\}, \{z_1^a, y_1^{s,a}\}, \dots$, obtained by Attack Strategy II are (ϵ, α) -successful attack. By defining $e_t = x_t^a - s_t$ and $r_t = e_t - x_t$, we get $z_t^a = G(e_t)$, $y_t^{s,a} = C_s e_t + v_t^{s,a}$, and

$$\begin{aligned} e_{t+1} &= f(x_t^a) - f(s_t) + B\Pi(e_t, v_t^s) + w_t^a \\ r_{t+1} &= f(r_t) + B\Pi(r_t, 0) + f(x_t^a) - f(s_t) - f(x_t) \\ &\quad - f(r_t) + B\Pi(e_t, v_t^s) - B\Pi(x_t, v_t^s) - B\Pi(r_t, 0) \\ &= h(r_t, 0) + \sigma_1 + \sigma_2 \end{aligned}$$

with $\sigma_1 = f(x_t^a) - f(s_t) - f(x_t) - f(r_t)$ and $\sigma_2 = B\Pi(e_t, v_t^s) - B\Pi(x_t, v_t^s) - B\Pi(r_t, 0)$. Using the mean-value theorem and equality $x_t^a = r_t + s_t + x_t$, we obtain

$$\begin{aligned} f(x_t^a) &= f(x_t + s_t + r_t) = f(s_t) + \frac{\partial f}{\partial x}|_{\Theta_{(s,s+x+r)}}(x_t + r_t) \\ f(r_t) &= f(0) + \frac{\partial f}{\partial x}|_{\Theta_{(0,r)}} r_t = \frac{\partial f}{\partial x}|_{\Theta_{(0,r)}} r_t \\ f(x_t) &= f(0) + \frac{\partial f}{\partial x}|_{\Theta_{(0,x)}} x_t = \frac{\partial f}{\partial x}|_{\Theta_{(0,x)}} x_t. \end{aligned}$$

Therefore, we get

$$\begin{aligned} \sigma_1 &= \left(\frac{\partial f}{\partial x}|_{\Theta_{(s,s+x+r)}} - \frac{\partial f}{\partial x}|_{\Theta_{(0,x)}} \right) x_t \\ &\quad + \left(\frac{\partial f}{\partial x}|_{\Theta_{(s,s+x+r)}} - \frac{\partial f}{\partial x}|_{\Theta_{(0,r)}} \right) r_t. \end{aligned} \quad (50)$$

Similarly, using the mean-value theorem and $e_t = r_t + x_t$, we obtain that

$$\begin{aligned} \Pi(e_t, v_t^s) &= \Pi(x_t, v_t^s) + \frac{\partial \Pi}{\partial e_t}|_{\Theta_{(x,x+r)}} r_t \\ \Pi(r_t, 0) &= \Pi(0, 0) + \frac{\partial \Pi}{\partial r_t}|_{\Theta_{(0,r)}} r_t. \end{aligned} \quad (51)$$

Therefore, we get

$$\sigma_2 = B \left(\frac{\partial \Pi}{\partial e_t}|_{\Theta_{((x,v),(x+r,v))}} - \frac{\partial \Pi}{\partial r_t}|_{\Theta_{((0,0),(r,0))}} \right) r_t. \quad (52)$$

Using the fact that the function $\frac{\partial f}{\partial x}(x)$ is Lipschitz, for all $0 \leq t \leq T(\alpha + b_x + b, s_0)$, we have

$$\begin{aligned} \left\| \frac{\partial f}{\partial x}|_{\Theta_{(s,s+x+r)}} - \frac{\partial f}{\partial x}|_{\Theta_{(0,x)}} \right\| &\leq L'_f(\|x_t^a\| + \|x_t\|) \\ &\leq L'_f(\alpha + b_x) \end{aligned}$$

and

$$\begin{aligned} \left\| \frac{\partial f}{\partial x}|_{\Theta_{(s,s+x+r)}} - \frac{\partial f}{\partial x}|_{\Theta_{(0,r)}} \right\| &\leq L'_f(\|x_t^a\| + \|r_t\|) \\ &\leq L'_f(\alpha + \phi). \end{aligned}$$

Therefore, $\|\sigma_1\| \leq L_2 b_x + L_1 \|r_t\|$.

Similarly, $\left\| \frac{\partial \Pi}{\partial e_t}|_{\Theta_{((x,v),(x+r,v))}} - \frac{\partial \Pi}{\partial r_t}|_{\Theta_{((0,0),(r,0))}} \right\| \leq L'_\Pi(\|x_t + r_t\|) \leq L'_\Pi(\phi + b_x)$, and we get $L_3 = L'_\Pi(b_x + \phi + b_v)$ and $\|\sigma_2\| \leq L_3 \|r_t\|$; this results in $\|\sigma_1 + \sigma_2\| \leq L_2 b_x + (L_1 + L_3 \|B\|) \|r_t\|$. Since we have $L_1 + L_3 \|B\| < \frac{c_3}{c_4}$ and $L_2 b_x < \frac{c_3 - (L_1 + L_3 \|B\|) c_4}{c_4} \sqrt{\frac{c_1}{c_2}} \theta d$, using Lemma 8

for all $\|r_0\| = \|s_0\| < \min\{\phi, \sqrt{\frac{c_1}{c_2}}d\}$, there exists $t_1 > 0$, such that for all $t \geq t_1$, we have $\|r_t\| \leq b$ with $b = \frac{c_4}{c_3 - (L_1 + L_3\|B\|)c_4} \sqrt{\frac{c_2}{c_1} \frac{L_2 b_x}{\theta}}$; also, for all $0 \leq t < t_1$, it holds that $\|r_t\| \leq \sqrt{\frac{c_2}{c_1} e^{-\beta t}} \|s_0\|$ with $\beta = \frac{(1-\theta)(c_3 - (L_1 + L_3\|B\|)c_4)}{2c_2}$.

On the other hand, the dynamics $s_{t+1} = f(s_t)$ with nonzero s_0 will reach $\|s_t\| \geq \alpha + b_x + b$ for some $t \geq T(\alpha + b_x + b, s_0)$ as $f \in \mathcal{U}_0$. Using the reverse triangle inequality, we obtain

$$\|s_t\| - \|x_t^a\| \leq \|x_t^a - s_t\| = \|e_t\| = \|x_t + r_t\|$$

$$\leq b_x + b \Rightarrow -b - b_x + b + b_x + \alpha = \alpha \leq \|x_t^a\|$$

with probability $\delta(T(\alpha + b_x + b, s_0), b_x, b_v)$.

Now, we need to show that the designed y_t^a satisfies the stealthiness condition. In other words, we need to show that $\text{KL}(\mathbf{Q}(Y_0^a : Y_{T(\alpha+b_x+b, s_0)}^a) || \mathbf{P}(Y_0 : Y_{T(\alpha+b_x+b, s_0)})) \leq \log(\frac{1}{1-\epsilon^2})$ for some $\epsilon > 0$. Since the sequences Y_0^a, \dots, Y_t^a and Y_0, \dots, Y_t are generated by e_0, \dots, e_t and x_0, \dots, x_t , respectively, using the data-processing inequality of KL divergence and following the same procedure as for Theorem 2, we obtain

$$\begin{aligned} & \text{KL} \left(\mathbf{Q}(Y_0^a : Y_{T(\alpha+b_x+b, s_0)}^a) || \mathbf{P}(Y_0 : Y_{T(\alpha+b_x+b, s_0)}) \right) \\ & \leq \sum_{i=0}^{T(\alpha+b_x+b, s_0)} \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \|r_i\|^2. \end{aligned} \quad (53)$$

Similar argument as in the proof of Theorem 2 results in

$$\begin{aligned} & \text{KL} \left(\mathbf{Q}(Y_0^a : Y_{T(\alpha+b_x+b, s_0)}^a) || \mathbf{P}(Y_0 : Y_{T(\alpha+b_x+b, s_0)}) \right) \\ & \leq \lambda_{\max}(\Sigma_w^{-1}) \|s_0\|^2 + \lambda_{\max}(C_s^T \Sigma_v^{-1} C_s + \Sigma_w^{-1}) \\ & \quad \times \max \left\{ \min \left\{ T(\alpha + b_x + b, s_0) + 1, \right. \right. \\ & \quad \left. \left. \sqrt{\frac{c_2}{c_1} \frac{e^{-\beta}}{1 - e^{-\beta}}} \right\} \|s_0\|^2, \right. \\ & \quad \left. \min \left\{ t_1, \sqrt{\frac{c_2}{c_1} \frac{e^{-\beta}}{1 - e^{-\beta}}} \right\} \|s_0\|^2 \right. \\ & \quad \left. + (T(\alpha + b_x + b, s_0) - t_1)b \right\} \\ & = b_\epsilon. \end{aligned}$$

This means that the system is (ϵ, α) -attackable with the probability of $\delta(T(\alpha + b_x + b, s_0), b_x)$ with $\epsilon = \sqrt{1 - e^{-b_\epsilon}}$.

REFERENCES

- [1] A. Khazraei, H. Pfister, and M. Pajic, "Resiliency of perception-based controllers against attacks," in *Proc. Learn. Dyn. Control Conf.*, 2022, pp. 713–725.
- [2] A. Tampuu, T. Matiisen, M. Semikin, D. Fishman, and N. Muhammad, "A survey of end-to-end driving: Architectures and training methods," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 33, no. 4, pp. 1364–1384, Apr. 2022.
- [3] V. Rausch, A. Hansen, E. Solowjow, C. Liu, E. Kreuzer, and J. K. Hedrick, "Learning a deep neural net policy for end-to-end control of autonomous vehicles," in *Proc. Amer. Control Conf.*, 2017, pp. 4914–4919.
- [4] M. Jaritz, R. De Charette, M. Toromanoff, E. Perot, and F. Nashashibi, "End-to-end race driving with deep reinforcement learning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 2070–2075.
- [5] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 4693–4700.
- [6] S. Dean and B. Recht, "Certainty equivalent perception-based control," in *Proc. Learn. Dyn. Control Conf.*, 2021, pp. 399–411.
- [7] S. Dean, N. Matni, B. Recht, and V. Ye, "Robust guarantees for perception-based control," in *Proc. Learn. Dyn. Control Conf.*, 2020, pp. 350–360.
- [8] A. A. Al Makdah, V. Katewa, and F. Pasqualetti, "Accuracy prevents robustness in perception-based control," in *Proc. Amer. Control Conf.*, 2020, pp. 3940–3946.
- [9] A. Boloor, K. Garimella, X. He, C. Gill, Y. Vorobeychik, and X. Zhang, "Attacking vision-based perception in end-to-end autonomous driving models," *J. Syst. Architecture*, vol. 110, 2020, Art. no. 101766.
- [10] Y. Jia, Y. Lu, J. Shen, Q. A. Chen, Z. Zhong, and T. Wei, "Fooling detection alone is not enough: First adversarial attack against multiple object tracking," in *Proc. Int. Conf. Learn. Representations*, 2020.
- [11] Y. Mo and B. Sinopoli, "False data injection attacks in control systems," in *Proc. 1st Workshop Secure Control Syst.*, 2010, pp. 1–6.
- [12] I. Jovanov and M. Pajic, "Relaxing integrity requirements for attack-resilient cyber-physical systems," *IEEE Trans. Autom. Control*, vol. 64, no. 12, pp. 4843–4858, Dec. 2019.
- [13] A. Khazraei and M. Pajic, "Perfect attackability of linear dynamical systems with bounded noise," in *Proc. Amer. Control Conf.*, 2020, pp. 749–754.
- [14] C. Kwon, W. Liu, and I. Hwang, "Analysis and design of stealthy cyber attacks on unmanned aerial systems," *J. Aerosp. Inf. Syst.*, vol. 11, no. 8, pp. 525–539, 2014.
- [15] I. Jovanov and M. Pajic, "Sporadic data integrity for secure state estimation," in *Proc. IEEE 56th Annu. Conf. Decis. Control*, 2017, pp. 163–169.
- [16] A. Khazraei and M. Pajic, "Attack-resilient state estimation with intermittent data authentication," *Automatica*, vol. 138, 2022, Art. no. 110035.
- [17] C.-Z. Bai, F. Pasqualetti, and V. Gupta, "Data-injection attacks in stochastic control systems: Detectability and performance tradeoffs," *Automatica*, vol. 82, pp. 251–260, 2017.
- [18] C. Szegedy et al., "Intriguing properties of neural networks," in *Proc. 2nd Int. Conf. Learn. Representations*, 2014.
- [19] K. Eykholt et al., "Robust physical-world attacks on deep learning visual classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 1625–1634.
- [20] N. Papernot, P. McDaniel, I. Goodfellow, S. Jha, Z. B. Celik, and A. Swami, "Practical black-box attacks against machine learning," in *Proc. ACM Asia Conf. Comput. Commun. Secur.*, 2017, pp. 506–519.
- [21] J. Sun, Y. Cao, Q. A. Chen, and Z. M. Mao, "Towards robust LiDAR-based perception in autonomous driving: General black-box adversarial sensor attack and countermeasures," in *Proc. 29th USENIX Secur. Symp.*, 2020, pp. 877–894.
- [22] A. Wong, S. Cicek, and S. Soatto, "Targeted adversarial perturbations for monocular depth prediction," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 8486–8497.
- [23] H.-J. Yoon, H. J. Sani, and P. Voulgaris, "Learning image attacks toward vision guided autonomous vehicles," 2021, *arXiv:2105.03834*.
- [24] R. S. Hallyburton, Y. Liu, Y. Cao, Z. M. Mao, and M. Pajic, "Security analysis of camera-LiDAR fusion against black-box attacks on autonomous vehicles," in *Proc. 31st USENIX Secur. Symp.*, 2022, pp. 1903–1920.
- [25] F. Cai and X. Koutsoukos, "Real-time out-of-distribution detection in learning-enabled cyber-physical systems," in *Proc. ACM/IEEE 11th Int. Conf. Cyber-Phys. Syst.*, 2020, pp. 174–183.
- [26] Y. Mo and B. Sinopoli, "Secure control against replay attacks," in *Proc. 47th Allerton Conf. Commun., Control, Comput.*, 2009, pp. 911–918.
- [27] A. Teixeira, I. Shames, H. Sandberg, and K. H. Johansson, "Revealing stealthy attacks in control systems," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput.*, 2012, pp. 1806–1813.
- [28] R. S. Smith, "Covert misappropriation of networked control systems: Presenting a feedback structure," *IEEE Control Syst. Mag.*, vol. 35, no. 1, pp. 82–92, 2015.
- [29] T. Sui, Y. Mo, D. Marelli, X. Sun, and M. Fu, "The vulnerability of cyber-physical system under stealthy attacks," *IEEE Trans. Autom. Control*, vol. 66, no. 2, pp. 637–650, Feb. 2020.
- [30] A. Khazraei, R. S. Hallyburton, Q. Gao, Y. Wang, and M. Pajic, "Learning-based vulnerability analysis of cyber-physical systems," in *Proc. 13th IEEE/ACM Int. Conf. Cyber-Phys. Syst.*, 2022, pp. 259–269.

- [31] M. Pajic, I. Lee, and G. J. Pappas, "Attack-resilient state estimation for noisy dynamical systems," *IEEE Trans. Control Netw. Syst.*, vol. 4, no. 1, pp. 82–92, Mar. 2017.
- [32] Y.-C. Liu, G. Bianchin, and F. Pasqualetti, "Secure trajectory planning against undetectable spoofing attacks," *Automatica*, vol. 112, 2020, Art. no. 108655.
- [33] G. Bianchin, Y.-C. Liu, and F. Pasqualetti, "Secure navigation of robots in adversarial environments," *IEEE Control Syst. Lett.*, vol. 4, no. 1, pp. 1–6, Jan. 2020.
- [34] D. Gadginmath, V. Krishnan, and F. Pasqualetti, "Direct vs indirect methods for behavior-based attack detection," in *Proc. IEEE 61st Conf. Decis. Control*, 2022, pp. 7090–7096.
- [35] Q. Hu, D. Fooladivanda, Y. H. Chang, and C. J. Tomlin, "Secure state estimation and control for cyber security of the nonlinear power systems," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 3, pp. 1310–1321, Sep. 2018.
- [36] Y. Polyanskiy and Y. Wu, *Information Theory: From Coding to Learning*. Cambridge, MA, USA: Cambridge Univ. Press, 2024.
- [37] T. v. Dijk and G. d. Croon, "How do neural networks see depth in single images?", in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, 2019, pp. 2183–2191.
- [38] A. Lambert, A. Shaban, A. Raj, Z. Liu, and B. Boots, "Deep forward and inverse perceptual models for tracking and prediction," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2018, pp. 675–682.
- [39] H. J. Kushner, "A partial history of the early development of continuous-time nonlinear stochastic systems theory," *Automatica*, vol. 50, no. 2, pp. 303–334, 2014.
- [40] H. K. Khalil, *Nonlinear Systems*. Hoboken, NJ, USA: Prentice-Hall, 2002.
- [41] V. Lesi, I. Jovanov, and M. Pajic, "Network scheduling for secure cyber-physical systems," in *Proc. IEEE Real-Time Syst. Symp.*, 2017, pp. 45–55.
- [42] V. Lesi, I. Jovanov, and M. Pajic, "Integrating security in resource-constrained cyber-physical systems," *ACM Trans. Cyber-Phys. Syst.*, vol. 4, no. 3, May 2020, Art. no. 28.
- [43] R. S. Hallyburton, Q. Zhang, Z. M. Mao, and M. Pajic, "Partial-information, longitudinal cyber attacks on LiDAR in autonomous vehicles," 2023, *arXiv:2303.03470*.
- [44] MITRE, 2023. [Online]. Available: https://cve.mitre.org/cgi-bin/cve_name.cgi?name=CVE-2023-26246
- [45] MITRE, 2019. [Online]. Available: https://cve.mitre.org/cgi-bin/cve_name.cgi?name=CVE-2019-9977
- [46] MITRE, 2017. [Online]. Available: https://cve.mitre.org/cgi-bin/cve_name.cgi?name=CVE-2016-9337
- [47] I. A. Sumra, J.-L. Ab Manan, and H. Hasbullah, "Timing attack in vehicular network," in *Proc. 15th WSEAS Int. Conf. Comput., World Sci. Eng. Acad. Soc.*, 2011, pp. 151–155.
- [48] A. Krishnamurthy, "Lecture 21: Minimax theory," 2017.
- [49] Y. Polyanskiy and Y. Wu, "Lecture notes on information theory," *Lecture Notes ECE563*, vol. 6, no. 2012–2016, p. 7, 2014.
- [50] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge, U.K.: Cambridge Univ. Press, 2020.
- [51] A. Formal'skii, "An inverted pendulum on a fixed and a moving base," *J. Appl. Math. Mechanics*, vol. 70, no. 1, pp. 56–64, 2006.
- [52] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli, "Kinematic and dynamic vehicle models for autonomous driving control design," in *Proc. IEEE Intell. Veh. Symp.*, 2015, pp. 1094–1099.
- [53] D. Umsonst, H. Sandberg, and A. A. Cárdenas, "Security analysis of control system anomaly detectors," in *Proc. IEEE Amer. Control Conf.*, 2017, pp. 5500–5506.



Amir Khazraei received the B.S. degree from the Amirkabir University of Technology, Tehran, Iran, in 2014, and the M.S. degree from the University of Tehran, Tehran, in 2017, both in electrical engineering. He is currently working toward the Ph.D. degree in cyber-physical systems with the Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA.

His research interest includes the vulnerability analysis of cyber-physical systems, control theory, and machine learning.



Henry D. Pfister received the Ph.D. degree in electrical engineering from the University of California, San Diego, CA, USA, in 2003.

He is currently a Professor with the Electrical and Computer Engineering Department, Duke University, Durham, NC, USA, with a secondary appointment in mathematics. He was an Associate Professor with Texas A&M University from 2006 to 2014, a Postdoctoral Fellow with the École Polytechnique Fédérale de Lausanne during 2005–2006, and a Senior Engineer with Qualcomm Corporate R&D, San Diego, during 2003–2004. His current research interests include information theory, error-correcting codes, quantum computing, and machine learning.



Miroslav Pajic (Senior Member, IEEE) received the Dipl. Ing. and M.S. degrees from the University of Belgrade, Serbia, in 2003 and 2007, respectively, and the M.S. and Ph.D. degrees from the University of Pennsylvania, Philadelphia, PA, USA, in 2010 and 2012, respectively, all in electrical engineering.

He is currently the Dickinson Family Associate Professor with the Department of Electrical and Computer Engineering, Duke University, Durham, NC, USA. He also holds a secondary appointment with the Computer Science and Mechanical Engineering and Material Science Departments. His research interests include the design and analysis of high-assurance cyber-physical systems with varying levels of autonomy and human interaction, at the intersection of (more traditional) areas of embedded systems, AI, learning and controls, formal methods, and robotics.