REQUIEM FOR A DRONE: A MACHINE-LEARNING BASED FRAMEWORK FOR STEALTHY ATTACKS AGAINST UNMANNED AUTONOMOUS VEHICLES

Kyo Hyun Kim
University of Illinois at Urbana-Champaign
kkim103@illinois.edu

Denizhan Kara
University of Illinois at Urbana-Champaign
kara4@illinois.edu

Vineetha Paruchuri
George Washington University
vineetha.paruchuri@gwu.edu

Sibin Mohan
George Washington University
sibin.mohan@gwu.edu

Greg Kimberly
Boeing Research & Technology
Greg.Kimberly@boeing.com

Jae Kim
Boeing Research & Technology
jae.h.kim@boeing.com

Josh Eckhardt
Boeing Research & Technology
josh.d.eckhardt@boeing.com

Abstract—There is a space of uncertainty in the modeling of vehicular dynamics of autonomous systems due to noise in sensor readings, environmental factors or modeling errors. We present REQUIEM, a *software-only, blackbox approach* that exploits this space in a *stealthy* manner causing target systems, *e.g.,* unmanned aerial vehicles (UAVs), to significantly deviate from their mission parameters. Our system achieves this by modifying sensor values, all while avoiding detection by onboard anomaly detectors (hence, "stealthy"). The REQUIEM framework uses a combination of multiple deep learning models (that we refer to as "surrogates" and "spoofers") coupled with extensive, realistic simulations on a software-in-the-loop quadrotor UAV system. REQUIEM makes no assumptions about either the (types of) sensors or the onboard state estimation algorithm(s) — it works so long as the latter is "learnable"[1], [2], [3].

We demonstrate the effectiveness of our system using various attacks across multiple missions as well as multiple sets of statistical analyses. We show that REQUIEM successfully exploits the *modeling errors* (*i.e.,* causes significant deviations from planned mission parameters) while *remaining stealthy* (no detection even after tens of meters of deviations) and are *generalizable* (REQUIEM has potential to work across different attacks and sensor types).

I. INTRODUCTION

"Life is a set of choices. Lead or be led astray" – *Beyonce*.

Autonomous Vehicles (AVs) try to navigate through the world using sensors and state estimation algorithms; they may find use in a variety of applications, from agriculture and logistics to safety critical domains such as military[4], [5] and disaster response[6]. An increasing proportion of safety-critical vehicular operations (such as collision avoidance, obstacle detection/avoidance during lane-change, identifying stop signs, staying under the speed-limit *etc.*) that, until recently, exclusively depended on manual intervention/control are now increasingly being integrated into the intelligent navigation and control systems of contemporary vehicles with the eventual goal of achieving safe, reliable and law-abiding autonomous vehicular control across the gamut.

State estimation algorithms (*e.g.,* Extended Kalman Filter [7]) in AVs are used to model the state[1] of the system. They may have to account for deviations from the "ideal" state, that could be the result of system malfunctions (such as engine failure, GPS failure, *etc.*) and/or environmental factors (obstacle, micro-burst wind shear, *etc.*) or, in what could be a critical issue, *errors/noise in the measurements*. In practice, even the most meticulously designed control systems *always* operate under a certain amount of error because of the unavoidable observational/measurement errors involved in both, sensor measurement and complexity of modeling vehicular dynamics. Note that all AVs require sensors to perceive their environments. It is what enables an AV to correct itself when deviating from expected operational parameters. However, inaccurate sensor observations can prevent such corrections.

All of this *creates a space that can be exploited by an adversary without risk of detection*. REQUIEM, our general framework, *exploits* this space to generate spoofed input values in order to attack AVs, causing them to deviate away from mis-

---

[1]State means different things in different contexts, but we can think of it as a description of the *physical system* using a set of parameters (called "state variables"). State is static, *i.e.,* it does not offer any information about how/whether a system evolves over time.

sion objectives. Figure 1 shows two examples of an unmanned aerial vehicle (UAV)[2] deviating away from its expected mission and the state estimation algorithms not realizing that a disruption has occurred! Ideally, the algorithms (or anomaly detectors) on the UAV should have detected these deviations and corrected for it. However, the reported sensor value during the attack remains *very close* to the normal/expected values, as shown in Figure 2 — the difference between the nominal[3] and attack sensor values is within the margin of error.
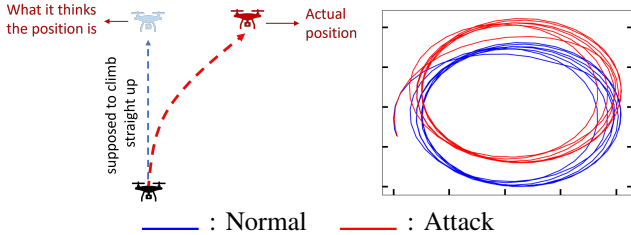

Fig. 1: Example of stealthy attack: the vehicle thinks it is following the mission path (blue) while in actuality, it is deviating (red). A realistic example is shown on the right where the attack trajectory is offset to north.
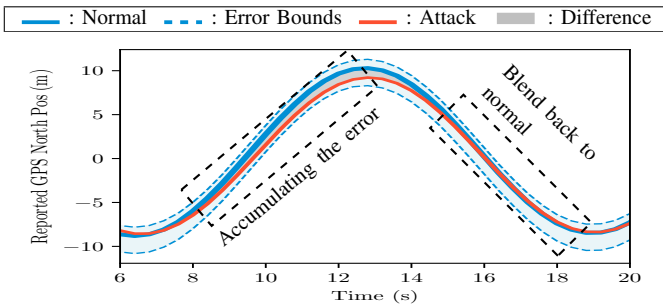

Fig. 2: Error space exploitation. Nominal mission trace (blue) vs same mission under stealthy attack (red). Grey region is the difference being exploited (*i.e.,* accumulated error). By the end of the accumulation, the vehicle's true position differs from the reported GPS position. To remain stealthy, the reported GPS position becomes similar to the nominal.

Such attack spaces have been shown to exist for EKF, the de-facto standard for vehicles with non-linear dynamics [8]. Pajic *et al.* [9] showed that stealthy attacks under linear time-invariant systems is not feasible. However, Hallyburton *et al.* [8] showed that finding optimal stealthy false data injection (FDI) values for non-linear dynamical systems is computationally feasible. Arnstrom and Teixeira [10] demonstrated stealthy FDI for an EKF through numerical simulation. Khazraei *et al.* [11] demonstrated such an attack against the CARLA[12] autonomous driving simulator. However, *vehicular dynamics for a UAV is more complex than a car since the former has to persistently balance itself in the air to be stable*. They also demonstrated an attack that works only on their own implementation of a theoretical UAV model [13], [14]. Other work shows attacks against the perception module on the PX4[4] system [15], [16]. However, this (whitebox) approach

requires that an attacker has full knowledge of the state estimator*, which may not be very feasible. ARES [17] proposed a reinforcement learning based approach to find adversarial values for the control task to move the vehicle away from its planned path. It is also a whitebox approach and the objective isn't stealthy deviations like the examples shown in Figure 1. Overall, existing literature in this space show that EKF-based algorithms are susceptible to stealthy attacks — this includes EKF-based anomaly detectors, *e.g.,* SAVIOR [18]. However, they are either specific to a particular state estimation algorithm, or require extensive knowledge about implementation details.

We present REQUIEM, a *general framework* that can be used to target state estimation models (*e.g.,* EKF) in autonomous systems and generate spoofed input/sensor values that can lead the systems astray. REQUIEM presents a *blackbox* attack — *i.e.,* there is no knowledge required about the internal details of the state estimator — the only requirement is that the state estimation function be "learnable" (explained in §VII-A) from observation of the inputs and outputs[5]. Hence, it is *not* specific to any particular estimation algorithm or even any specific sensor! The final result of a REQUIEM-based attack is a deviation of the physical system's trajectory while the system itself doesn't notice anything untoward (and believes that it is following the original mission parameters).

REQUIEM, at a high level, includes two components: *(a)* "surrogate" models and *(b)* a "spoofer". "Surrogate" models (often more than one as detailed in §V-A) try to emulate the target functions (*i.e.,* state estimators and anomaly detectors)[6] based on the observed I/O behavior. A surrogate model is a deep neural network (DNN) [19] that can capture both, the operational capabilities as well as the vulnerabilities of the target model; hence, the *vulnerabilities that are in the target can be transferred to the surrogate*. The use of DNNs allows us to use generative adversarial networks (GANs) [20] based training methods to develop a "spoofer". The spoofer *carefully crafts values for data injections in a deliberate manner* (*i.e.,* not random spraying) to *influence the movement of the UAV*, while *evading detection*. **Note:** we are not trying to take direct control of the vehicle, rather *influence* it to stray away from the mission parameters in a *specific* manner. Hence, **our contributions** are:

1) A novel, stealthy attack approach aimed at UAVs.
2) REQUIEM, a framework to implement said approach.
3) Demonstration of the attacks on a widely used, real-world flight controller, *viz.* PX4 using a software-in-the-loop (SITL) simulation framework.
4) Insights gathered from developing REQUIEM.

The REQUIEM code has been opensourced and is available online[7]. We will first discuss our system model, assumptions as well as some required background material next.

---

[2]We use UAVs and "drones" interchangeably in this paper.

[3]We use 'nominal' and 'normal' interchangeably in this paper.

[4]PX4 is an opensource flight controller suite for various types of vehicles
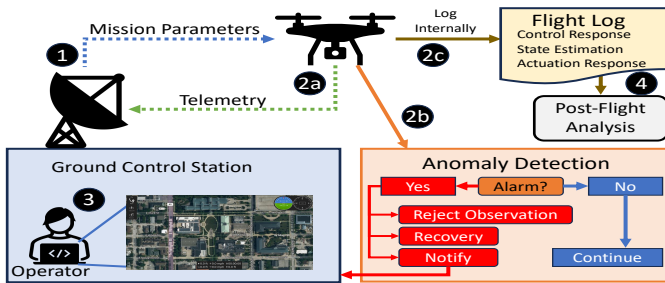
Fig. 3: Process of UAV deployment.

II. BACKGROUND AND SYSTEM MODEL

This section provides an overview of the process of deploying Unmanned Aerial Vehicles (UAVs) and describes the system model. It is followed by the description of Extended Kalman Filter (EKF) [7] (the de facto state estimation algorithm in autonomous vehicles) along with a sequence-of-events, to provide context on where such a state estimator fits in the operation of UAVs.

A. UAV Deployment Scenario

The operation of a UAV involves an operator, a Ground Control Station (GCS) and a UAV as shown in Figure 3. The vehicle is not always in the line-of-sight of the operator after the take-off.

At ❶ the GCS sends the mission parameters to the UAV, after which ❷a ❷b and ❷c occur simultaneously. At ❷a the UAV sends telemetry information back to the GCS, after which, at ❸, the operator uses the telemetry information to 'observe' the reported position and ensure that the vehicle follows its mission. At ❷b an Anomaly Detector (AD) runs on the UAV to make sure that the sensor readings are within acceptable tolerance limits. If the sensor readings trigger an alarm, then the AD rejects the sensor data and notifies the GCS. At ❷c the UAV internally logs detailed information about itself such as control response, state estimation status, actuation response etc., so that at ❹ the post-flight analysis is done to evaluate its performance at the end of the mission.

B. Our System Model and Assumptions

The Software architecture of the system is depicted in Figure 4, which illustrates how the information from the sensor disseminates across various tasks via middleware. Autonomous systems commonly utilize middleware as a software system architecture due to the benefit of being able to add a new components to the system without affecting the existing functionality (*i.e.,* essentially scalability). This is popular in robotics (*e.g.,* Robot Operating System (ROS) [21]) and in flight controllers (*e.g.,* PX4 [22]) where a collection of flight control tasks coordinate to execute the mission. The middleware often has a publish-subscribe architecture that allows each component of the system to publish its output and for other components

to subscribe to that output. The software tasks on the system are activated when the input message is available for use. We assume that the two stages of the state estimation, *predict* and *update*, are two separate tasks (elaborated in §II-C).

We assume a *more restrictive system*[8] than traditional flight controllers — the latter assumes that all tasks share the same user space environment. In this paper, we assume that the tasks required for operating the UAV are isolated from each other in user space, communicating via a middleware. Additionally, the integrity of certain calculations that are crucial for the operation of the vehicle (*e.g.,* anomaly detection and control) are *protected* using techniques such as triple redundancy modules[23] or trusted execution environment such as Arm TrustZone[24], Intel SGX[25] or SecureCore[26]. This results in an attacker being unable to tamper with the *protected* functions. For instance, the state estimation calculations are integral to the operation of UAVs. Therefore, in this paper, we assume that the state estimation update function is protected.

Sensors. The system has an Inertial Measurement Unit (IMU), the Global Positioning System (GPS), magnetometer, and barometer that are provided by autopilot hardware such as Pixhawk [27] or Navio2 [28] . Although camera can be used to track movement (*i.e.,* optical flow[29]), we consider scenarios where there is low visibility (*e.g.,* fog, smoke) or at night, making the camera useless. IMU is used during the prediction stage to predict the altitude, velocity and position. GPS provides information regarding the position and velocity. Magnetometer provides information regarding the orientation of the vehicle *w.r.t.* the environment. Barometer measures the relative altitude of the vehicle. Under our system model, the sensors are used when the update task receives them.

C. Extended Kalman Filter

A state estimator *never* knows the true state of such systems — it can only estimate the state based on observed evidence (*i.e.,* from sensors). The Extended Kalman Filter (EKF) [7] can handle the estimation of nonlinear dynamics, such as drone movement, via linearization (*i.e.,* linear approximation) and is the de facto standard in UAVs.

EKF has two stages: *predict* and *update*. The *predict* stage provides a meaningful state estimation based on the actuation command in between sensor updates. The *update* stage adjusts the state estimation based on sensors. Due to the space constraints, we explain the general process of the update and the remaining details are elaborated in the Appendix D.

Predict. EKF uses the previous actuation command to predict its current state. For example, if the previous command was to move the vehicle north by a meter, the current position should be a meter north from the previous one.

Update. When new sensor data becomes available, the estimator calculates the *expected* sensor value based on the predicted state using an observation function and compares the observed value with the expected value. The disparity between the two is known as the *residual*, and determines how much the predicted state needs to adjust. Additionally, the residual is used to gauge the degree of abnormality in sensor measurements.

---

[5]As we see later, this is a very lightweight requirement and doesn't need anything beyond a simple user-level access on the system or communication with a middleware.

[6]**Note:** we use "target function" and "state estimator" interchangeably in this paper.

[7]https://projrequiem.github.io

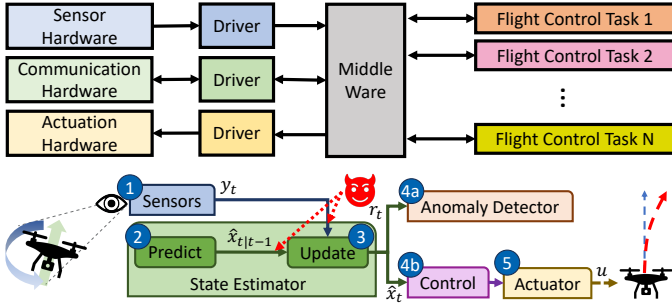[8]It is more restrictive for the *attacker*, therefore, harder for *our* objectives.

Fig. 4: (Top) Standard software architecture of common flight controllers. (Bottom) Sequence of events during the operation of a UAV.

$$r_t = y_t - h\left(\hat{x}_{t|t-1}\right) \quad (1) \qquad \hat{x}_t = \hat{x}_{t|t-1} + K_t r_t \quad (2)$$

(residual) (sensor) (observation function) (predicted state) (updated state)

Formally, when the sensor, y_t, is available, the *residual* is calculated based on Equation 1 using the observation function h, which happens to be non-linear. Therefore, linearization is performed to construct Kalman gain K that is used to update the state, $\hat{x}$, as shown in Equation 2. K adjusts how much influence the residual r should have on the update. EKF implementations often use χ^2 anomaly detector as default (explained in §II-D).

D. Anomaly Detector

χ^2 [30] is a statistical anomaly detection (AD) method often deployed in CPS [31], [8]. EKF implementations in autopilots also use χ^2 as shown in Equation 3:

$$r_t^T \; C^{-1} \; r_t = z_t \geq \eta \to \frac{z_t}{\eta} \geq 1 \quad (3)$$

(inverse residual covariance) (normalized residual) (χ^2 threshold)

z_t is the normalized residual assumed to be in χ^2 distribution and η is the user specified threshold in standard deviations. This paper refers to $\frac{z_t}{\eta}$ as χ^2 score. In this paper, we show that REQUIEM can easily bypass the implemented χ^2, with the (autopilot provided) default value of $\eta = 5$. Therefore, we use a more stringent AD, τ (described in §II-D), and show that REQUIEM can bypass that as well.

Therefore, in addition to χ^2, we also use threshold based AD as shown on Equation 4 where if the residual is outside the acceptable range τ_{min} and τ_{max}, the observed sensor is considered to be an anomaly, *i.e.,*;

$$r \leq \tau_{min} \; || \; r \geq \tau_{max} \quad (4)$$

(residual) (anomaly threshold)

The thresholds are determined by the distribution of residual observed during the normal mission. We will refer to it as τ-AD.

E. Sequence-of-Events

A state estimator goes through the following sequence of events as shown in Figure 4. ① The kinematics of the

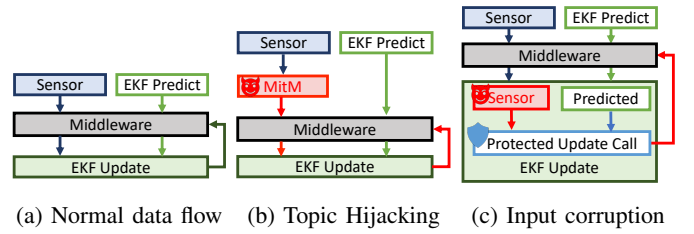(a) Normal data flow (b) Topic Hijacking (c) Input corruption

Fig. 5: Examples of adversary entry vectors. (a) shows the nominal flow. (b) is an example of malicious module in the system hijacking the sensor topic. (c) is an example of the attacker exploiting vulnerability in the task to corrupt the input but the integrity of the target function is protected.

vehicle are captured by sensors. ② Based on the previous actuation, the state estimation predicts the current state. ③ The sensor values are observed by the state estimator that is used to compare against the predicted observation. The *residuals* calculated during the update are sent to the anomaly detector and the updated state is sent to the control. ④a If the residual is too high, then the sensor value is rejected and an alarm is triggered. ④b With the newly updated state, the control module calculates the actuation command needed to keep the vehicle within the mission parameters. ⑤ The actuation command is then translated to a pulse-width modulation (PWM) signal and sent to the motors, resulting in the mission trajectory shown in blue. The motors affect the kinematics and the cycle continues.

However under the false data injection (FDI) spoofing attack, an adversary can observe the predicted state and manipulate a sensor values before ③. With REQUIEM, the attack results in trajectory shown in red while the vehicle thinks that it is following the original mission trajectory shown in blue dotted line. We discuss the threat model in the next section elaborating further about an attacker's capability to achieve such results.

III. THREAT MODEL

The main objective of an attacker (*i.e.,* our system) is to cause changes in vehicle movements while the mission operation *seems* normal without the anomaly detector raising any alarms. Hence, it is in *the attacker's best interest **not** to crash the vehicle or be detected too soon*.

Our threat model is similar to that of ARES [17] [9]. We assume realistic attack vectors (*e.g.,* firmware bug exploitation of a particular version, malicious component in the middleware, *etc.*) with the means to collect the target's input/output values and to inject values during deployment. However we broaden the threat model even further: the target function/module is a blackbox (*i.e.,* attacker can only view the input and the output of the target, not the internal logic). Due to the typical multi-vendor model-based development of avionics systems (even for unmanned systems) [32], [33], [34], there is a good chance that *(a)* not all vendors who developed the system have followed good software security practices [35], [36], [37], *(b)* there

---

[9]They use reinforcement learning with knowledge of the implementation details whereas we use treat the system as a blackbox. Also we are more focused on stealth.

could be many latent, undiscovered bugs in the system and *(c)* the entire system could have been developed by a combination of commodity/open-source components and proprietary software *(e.g.,* the flight controller). We also assume that the target function/module does *not* need root access; REQUIEM (realistically) assumes that the vehicle relies on the target function/module for its operation.

As a result, our threat model is more general compared to prior work [38], [39], [40] since *(i)* attackers have additional avenues of exploiting the middleware security vulnerabilities, as shown on Figure 5, *(e.g.,* message man-in-the-middle (MitM) attack[41], [42], [43]) on top of firmware exploitation *(e.g.,* memory corruption[44], [45], buffer overflow[46]); *(ii)* the blackboxing enables the targeting of protected (§II-B) functions (Figure 5c) and the lack of necessity for reverse engineering efforts of the target. *(iii)* if the attacker is not aware when the target executes, there are existing methods [47], [48] to calculate, and then predict, the timing behavior of the target functions. *(iv)* the certification process [49], [50] for such systems involves integration testing, often using simulations; hence it is possible to estimate the I/O behavior as well as expectations and configuration information for the system.

Note: Physical attacks such as GPS signal spoofing *(e.g.,* [51], [52], [53], [54], *etc.)* or using sound to spoof IMUs [55], [56], [57] have significantly different requirements such as additional hardware, shaping of the physical properties of the attacker/target, precise timing requirements as well as physical proximity to the target. These are *very different* threat models and, hence, physical attacks are out-of-scope and not considered in this paper.

IV. REQUIEM

State estimators must account for the errors in sensor values due to *(a)* noise from the environment, *(b)* hardware variations in sensor design, *(c)* placement of sensors, *(d)* calibration errors, *(e)* incomplete modeling of non-linear vehicular dynamics, etc. The main idea of stealthy attacks is to exploit such errors to cause deviations from mission parameters that are not easily detectable. Our stealthy attack framework REQUIEM leverages the space (or the margin) due to such errors. Challenges exist in using these errors to our advantage *viz.*, different vehicles may deploy with different sensors, estimation parameters or estimation algorithms. Hence the challenges are: *(a)* understanding the estimation algorithm behavior, *(b)* finding spoofing values to induce deviations without being detected and *(c)* evaluating the effectiveness of such attacks.

REQUIEM handles challenges through the blackbox approach by constructing deep models that emulate particular aspects of the state estimation algorithm (*i.e.,* by creation of *"surrogate models"*[10] §V-A). Using such *"surrogates"* as a proxy, a *"spoofer"* (§V-B) is trained to stealthily accumulate the error, over time, in the estimator. Note that REQUIEM is agnostic to the state estimation algorithm used as REQUIEM *only cares about how the sensor values impact the state estimation*, thereby removing the burden of manually reverse engineering the algorithm. As a result, the attack on the surrogate transfers to the estimation algorithm to induce deviated trajectories like

[10]A deep neural network that mimics the state estimation algorithm

Figure 1 where the vehicle thinks it is following the mission path while deviating, without triggering the on-board AD.

Furthermore, depending on the mission, certain trajectory deviations can be more desirable for the attacker. REQUIEM allows customization of deviations and, for this paper, presents two types of stealthy spoofing attacks, as shown in Table I: *(i)* "No correction" (NC) causes the vehicle to behave as if sensors do not exist; it navigates only by its heading and the distance traveled from the initial position *i.e.,* dead-reckoning while *(ii)* "Direction bias" (DB) results in deviations *towards a direction of the attack's choosing*.

TABLE I: REQUIEM stealthy attacks

Name	Description
No Correction (NC)	Removes the effect of sensors on state estimation, resulting in a dead-reckoning like behavior.
Direction Bias (DB)	The vehicle deviates towards a particular direction

TABLE II: Naïve attacks types where the attack strategy does not change based on the state of the vehicle.

Name	Description
Position Constant Offset (PCO)	Inject a constant value to the GPS north position
Position Random Offset (PRO)	Random value is injected to the GPS north position
Position Boiling Frog (PBF)	Gradually increase the value injected into the GPS north position
Velocity Constant Offset (VCO)	Add a constant value to the GPS north velocity
Velocity Random Offset (VRO)	Add a random value to the GPS north velocity
Velocity Boiling Frog (VBF)	Gradually increase the value added into the GPS north velocity

The REQUIEM framework is composed of three stages as shown in Figure 6: data collection (§IV-C), surrogate training (§V-A) and spoofer training (§V-B). First, we describe the overall design of REQUIEM.

A. Setup

REQUIEM uses flight simulations to train the surrogate and the spoofer models. Specifically, the simulation is tuned to best represent the real-world counterpart so that the trained spoofer can transfer to the vehicle in deployment. The first step is to extract the information required to set up the simulation. Then the simulation is played out so that the adversary can collect the input/output of the target function (*i.e.,* data collection stage). The collected data is then used to train the surrogate model where the latter will learn to emulate the target. This model will then be used to adversarially train the spoofer. The spoofer learns to exploit the target function through repeated interactions with the surrogate model; resulting in a successful stealthy deviation during deployment. In the following subsections, we provide detailed descriptions for each of the stages, starting with the information required to use REQUIEM.
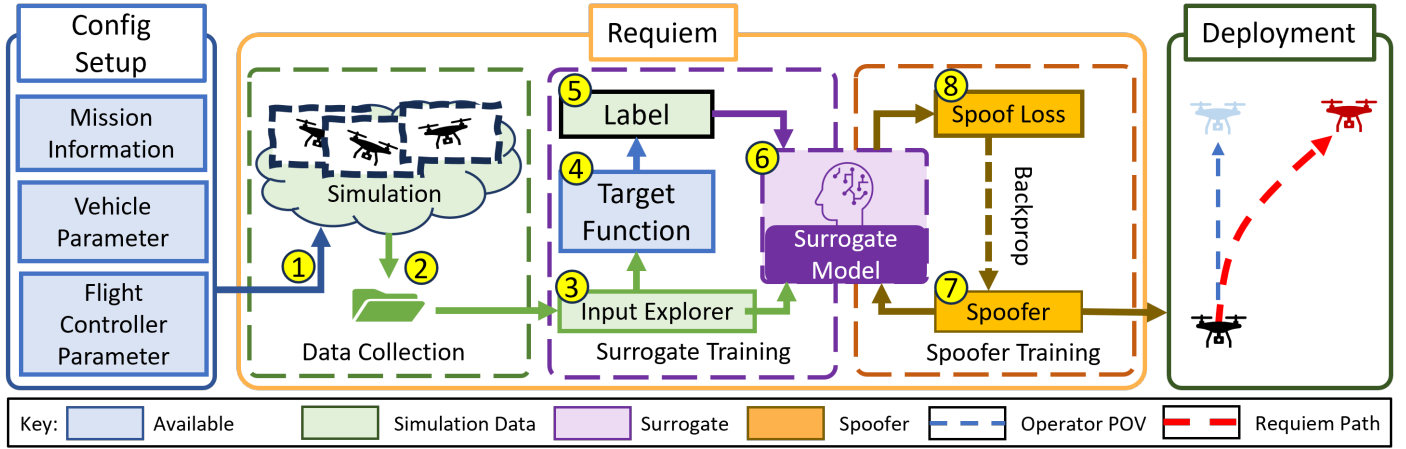
Fig. 6: REQUIEM Pipeline. ① Simulation is configured to best reflect the deployment scenario. ② data is collected over multiple simulation runs. ③ generate more data similar to collected. ④ generated data is sent to the target function to ⑤ get the label. The input-output pairs are used to ⑥ train the *surrogate* model. ⑦ spoofer model is then trained to optimize against the surrogate where the loss function ⑧ provides feedback to the spoofer, shaping its attack behavior.

B. Prerequisite: Configuration Setup

Different types of vehicles, missions and flight controllers incorporate sensors into state estimation algorithms in different ways; and the same is true for estimation of error bounds. The first step in our process ("config setup" in Figure 6) is to gather information from a system that is equivalent to the victim UAV. We leverage the fact that flight controllers go through simulation testing for certification [49], [50], hence, include an interface to run simulations. Even commercial, proprietary flight controllers provide hardware-in-the-loop simulation functionality [58], [59], [60]. Although simulation is *strictly not needed* to collect data since the information can be obtained during the physical operation of the UAV, it makes the data collection process more convenient. As a result, we use the available information about the UAV (§III) to configure the flight simulator to generate and collect data to train the surrogate and the spoof generator.

C. Data Collection

The mission information and vehicle parameters determine the distribution of a vehicle's state and sensor values in the training set; this set will be used to train the surrogate and spoofer. Simulation setup is tuned in Figure 6: ① using the configuration setup stage. The data collection components focus on collecting the data pertaining to the target sensor chosen by the adversary.

Algorithm 1 Data Collection Procedure

1: $s \leftarrow$ Snapshot() /* *Collect input* */
2: $EKF\_update\_state()$ /* *Update State* */
3: $s_h \leftarrow$ Snapshot() /* *Collect output* */

---

② While simulating a mission, the inputs and outputs of the state estimation function are collected as shown in Algorithm 1. The collection starts when the mission begins (at takeoff) only when the target sensor is available. Lines 1 and 3 capture snapshots of the variables before and after the

state update. The Snapshot after the state update collects the response when the state update is queried.

The adversary analyzes the snapshot data to determine which parts pertain to the target sensor by comparing the snapshots before and after state estimation. Narrowing down the relevant data assists the *surrogate* model to converge accurately during training since it narrows down the "search space" for the model to optimize.

Implementation. We target real-world flight controllers. Therefore, we implemented the data collection in the PX4 autopilot, an open-source flight controller [22].

In PX4, we targeted the `ekf2` module that use the EKF library class to perform the predict and update steps. The collection procedure is implemented by taking the snapshots of the EKF object, *i.e.,* collecting variables used as input for the update stage, during a simulation of a mission run. Snapshots can be taken by parsing the variables in memory and exfiltrating it as a JSON file. A full snapshot consists of 2020 variables of the EKF object that is sufficient to run the PX4's implementation of EKF offline. However, the snapshot size for training REQUIEM is narrowed down to those variables presented in Table V - obtained by analyzing how a change in inputs influences the outputs during the course of a mission.

To capture all types of kinematics for a potential mission, we utilize the random walk mission where the mission positions are chosen randomly (according the Brownian motion) where, at each position, the vehicle faces a random direction. Due to the space constraints, the details of the mission implementation of this data collection phase are elaborated in appendix B.

Overall, we ran 20 separate missions for data collection. Twelve of that are used for training and 8 for testing. resulting in 2275 examples. Training a deep model, however, require more examples. Therefore we use *input exploration* to generate more meaningful examples in the next stage to train a surrogate model.

V. Surrogate vs Spoofer: Adversarial Deep Models

The construction of a surrogate and its role (Figure 6) in training the spoofer is crucial. The surrogate model is optimized to emulate the target function where as the spoofer is optimized to induce a particular behavior from the surrogate (for instance, deviation). We implement both as deep neural network (DNN) models that recursively adjust their parameters (*i.e.,* via back-propagation) based on how each of the models' parameters contribute to its output (*i.e.,* gradient). Both models' usage of gradients in back-propagation is crucial as generative adversarial networks (GANs)[11] allow the spoofer to be cast as an antagonist to the surrogates. First we describe the construction of the surrogate model.

A. Surrogate Model

Our surrogate is a DNN that is trained to emulate the behavior of a target function (EKF in this case). However, the target function may either be too complex or the attacker may care only about a subset of the functionality. Therefore, *multiple smaller DNNs* are used where each captures a particular aspect of the target function (*i.e.,* a **slice** of the target function). This section elaborates on the creation of the surrogate models and our implementation.

Overview. DNNs requires large training sets. To ensure there are sufficient training examples, as shown in Figure 6, we use ③ *input exploration* where new examples are generated by adding noise to the input and ④ query the target function to retrieve corresponding labels. ⑤ The expanded I/O pairs are used to ⑥ train a surrogate model to prevent overfitting (*i.e.,* regularizer; see below for details).

Objective. The objective is to generate a function that demonstrates behavior close to, if not the same, the target function. Formally, let $Q : \mathcal{X} \rightarrow \mathcal{Y}$ be the target function that maps the input set $\mathcal{X}$ (*i.e.,* input to the EKF state update) to the output set $\mathcal{Y}$ (*i.e.,* output of the EKF state update). Capturing the exact mapping requires having access to all input output pairs which is not feasible. Instead, we look into a subset of the input $S \subset \mathcal{X}$ collected in the data collection stage that captures the input distribution during a mission. Therefore the objective function for constructing the surrogate model, D, is:

$$
\forall \underset{\text{snapshot}}{s} \in S, \underset{\text{noise}}{\omega} \sim \underset{\text{uniform dist.}}{U} (\omega_l, \omega_u)
$$
$$
\min_{D} \left[E \left[\| \underset{\text{surrogate output}}{D(s+\omega)} - \underset{\text{target output}}{Q(s+\omega)} \|_2 \right] \right] \quad (5)
$$

The noise ω serves as a regularizer to prevent the model from overfitting and is sampled uniformly between (ω_l, ω_u) for each input s. Hence, *the mean difference between surrogate and target is minimized*. The surrogate training procedure is in Appendix E.

Slices. Remaining stealthy while causing trajectory deviation is the end goal of REQUIEM. Therefore we create surrogate models for each of residual, velocity and position estimation updates as "slices" of the target function:
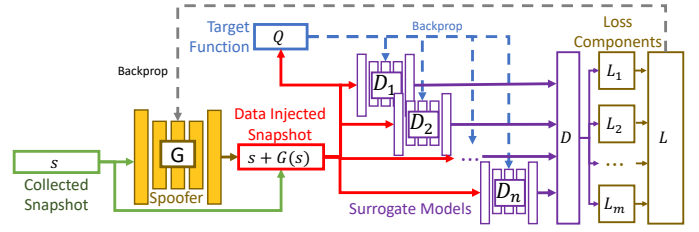

Fig. 7: Spoofer training process. Spoofer adds its output to the collected snapshot. The n surrogate models are fine tuned to the injected input by querying the target function. The outputs of the surrogate models are consolidated for loss calculation. The m loss components each quantize a facet of the surrogates' behavior. The loss components are consolidated to provide feedback to the spoofer via backpropagation.

D_1 : Residual surrogate (Equation 1)
D_2 : Velocity update surrogate (Equation 2)
D_3 : Position update surrogate (Equation 2)

Equation 2 is used for both velocity and position update in EKF.

Implementation. Target function can be queried by topic hijacking the input and subscribing to the target function's output (§III). To avoid running other components of the flight controller, we achieved the same effect by isolating the code for the target function and created a query server that takes snapshot as input and outputs the resulting snapshot. The isolated code was validated using the snapshot collected in the previous process.

The two surrogate models were implemented using PyTorch. All surrogates (*i.e.,* D_1, D_2 and D_3) are trained using the Algorithm 3. The training parameters are specified in Table VII in Appendix E. Overall, the final output of the surrogate models is defined as Equation 6 that shows it to be a concatenation of the other surrogates:

$$
D(s) : \underset{\text{ResSurModel output}}{D_1(s)} \| \underset{\text{VelSurModel output}}{D_2(D_1(s), s)} \| \underset{\text{PosSurModel output}}{D_3(D_1(s), s)} \quad (6)
$$

B. Spoofer

The Spoofer model optimizes against the surrogate model to cause deviations during deployment. Using DNNs as surrogate models allows us to calculate how each parameter of the spoofer affected the "spoofed" result of the state estimation process. The adversary may want to bias the state estimation a particular way (*e.g.,* moving north). Therefore, we control the effects of the spoofer via *loss components* that score the results on the surrogate (and on the target function by proxy).

Overview. The spoofer is trained by simulating its effect on the resulting state estimation for each collected snapshot. As a result, the final loss, L, is a function of both the spoofer and the surrogates (as shown in Figure 7). From Figure 6, we see that the spoof generator is ⑦ trained *against the surrogate model* using a GAN framework. ⑧ the attack performance score, *i.e., loss*, from the output of the surrogate model provides feedback to the spoofer. This allows the generator to find *adversarial*

---

[11]Details of GAN are in Appendix C

examples to induce the desired behavior from the surrogate model (*e.g.,* prevent the anomaly response...) that can be transferred to the target function.

Objective. The spoofer must induce a *specific* behavior (*i.e.,* deviation) from the target function via surrogate models without triggering the anomaly alarms. Therefore, the *behavior of the spoofer is shaped by its loss function*. We create the following attacks (Table I):

- No Correction (NC) attack where the residual being minimized results in a dead-reckoning trajectory (*i.e.,* as if the sensors do not exist)
- Direction Bias (DB) attack where deviation is biased towards a particular direction

We formalize the loss function for each of the attack.

Designing the No Correction (NC) Attack. To achieve a dead-reckoning trajectory, the *sensor values should not induce any correction to the predicted state*. Since the *residual* serves as the anomaly loss in Equation 7, minimizing the *anomaly level* (Equation 8) generates an NC attack and also evades detection.

$$\mathcal{L}_a = \underset{\text{anomaly function}}{A} \big(\underset{\text{surrogate output}}{D(G(s))} \big) \quad (7) \qquad \min_G \big[(\mathcal{L}_a)^2 \big] \quad (8)$$

In this case, function A filters out the output of the residual surrogate model, D_1, from the final output of D (Equation 6). Therefore, surrogates D_2 and D_3 are not used.

Designing the Direction Bias (DB) Attack. To cause a *specific* bias in deviations without triggering the anomaly detector, the spoofer must consider*(a)* the direction of deviation as well as *(b)* alarm triggers. The two facets are formulated as *deviation loss* (Equation 9) and *budget loss* (Equation 10) respectively. We construct the loss functions such that the deviation must taken into account only when the detector doesn't raise any alarms (Equation 11).

$$\mathcal{L}_d = ReLU(\underset{\text{predicted state position}}{s[pos]} - \underset{\text{updated state resulting from spoof}}{D(s + G(s))}) \quad (9)$$

$$\mathcal{L}_b = ReLU(\underset{\text{anomaly threshold}}{T} - |\mathcal{L}_a|) \quad (10)$$

Deviation bias loss ($\mathcal{L}_d$; Equation 9) measures the difference between the estimated state and the true state. $ReLU$[61] is used to ensure that the loss increases only if the spoofed values cause the state to update in favor of a desired direction. This loss function encourages the spoofer to trade off some residual for more deviation.

Budget loss ($\mathcal{L}_b$; Equation 10) is the slack between the current anomaly level and the anomaly threshold, T. T should be set below the target anomaly detector's threshold to account for the noise. Therefore, the maximization of $\mathcal{L}_d \cdot \mathcal{L}_b$ results in maximizing deviation *per anomaly level*. We still need to incorporate $\mathcal{L}_a$ to remain stealthy; hence, we end up with the following objective function:

$$\min_G \big[(\mathcal{L}_a)^2 - \mathcal{L}_b \cdot \mathcal{L}_d \big] \quad (11)$$

Notice that $\mathcal{L}_d$ only has an effect if $\mathcal{L}_b$ is greater than zero (*i.e.,* if the slack exists). Therefore, $\mathcal{L}_b$ acts as a "budget". It ensures that the spoofer does not tradeoff between gaining for more deviation at the cost of triggering the alarm since ($\mathcal{L}_a$ is the only loss left when the alarm is triggered). Therefore, optimizing for Equation 11 allows the spoofer to trade off between deviating more and the cost of increasing the anomaly level when the alarm isn't triggered.

Implementation. The spoofer was implemented using PyTorch. Due to the space constraints, the details are elaborated in the Appendix E

VI. EVALUATION AND RESULTS

A successful stealthy attack must show that the vehicle is operating in a seemingly normal state while causing the actual trajectory to deviate. We outline our evaluation methods and results in this section to show that REQUIEM successfully carries out the stealthy attack. Evaluating the effectiveness of REQUIEM raises the following set of research questions (RQ):

RQ1 (§VI-C) Did REQUIEM succeed in the stealthy attack?
RQ2 (§VI-D) What determines an attack to be stealthy?
RQ3 (§VI-E) What determines an attack to be "meaningful"?
RQ4 (§VI-F) How do we measure how well the surrogate model *"learns"* a target function?
RQ5 (§VI-G) How do environmental factors affect the attack?

We define the evaluation metrics, experiment parameters and processes to answer each of these research questions.

A. Metrics

To evaluate an attack, there are four trajectories to consider: *(i) planned*, *(ii) normal*, *(iii) attack* and *(iv) system point-of-view (POV)*. *Planned* trajectory is the ideal path the vehicle should take on a mission. *Normal* trajectory is the resulting path of a mission without any attacks. *Attack* trajectory is the *true* trajectory of the vehicle *during an attack* while the *system POV* trajectory is the path the vehicle (and potentially, a remote the operator) observes during the attack. Larger the difference between the *attack* trajectory and the *system POV* trajectory, larger the deviation.

We must also consider stealthiness when evaluating these attacks where the attack's success depends on achieving *(a) seemingly* normal operations of the vehicle (*i.e.,* **RQ2**) and *(b)* significant deviation (*i.e.,* **RQ3**) *simultaneously*. Therefore, we provide metrics for the attack *w.r.t.* the anomaly detector response and deviations.

As mentioned in §II-D, two types of anomaly detectors are considered: *(i) onboard AD* (*i.e.,* χ^2 §II-C) and the more stringent *(ii) threshold AD* (*i.e.,* τ, §II-D). The threshold of τ-AD was set to two standard deviations from the mean [12] based on the nominal residual distribution for each mission as shown in the appendix Figure VI. A stealthy attack mainly cares about whether the anomaly detector raises an alarm during a mission and the deviations achieved without triggering.

---

[12]Conservative bound since the bounds need to be more lax to achieve zero false positives under a normal mission – *i.e.,* we used a tighter (more restrictive) bound for *our* system since the default values were too easy to spoof

Table IVa describes the corresponding metrics, *viz.*: M1 and M3 check the alarm response; *i.e.,* when the detector raises alarms, the attack is considered to be *overt*. An attack is considered *partially successful* if the attack becomes overt at a much later stage in the attack. However if an attack becomes *overt* within a second of the attack, it is considered a failure *w.r.t.* the chosen anomaly detector. *Stealthy attacks should not trigger any of the anomaly detectors.*

M2 and M4 measure the maximum stealthy deviation achieved in meters (*i.e.,* maximum distance between the attack and system POV trajectory without becoming *overt*) *w.r.t.* the two anomaly detectors. In a partially successful attack, the maximum stealthy deviation is the largest deviation until the attack becomes *overt*. An effective attack should have as much stealthy deviation as possible within the given mission.

The adversary may care more about changing the shape (or path) of the trajectory while retaining the appearance that the UAV is following the planned path in missions (*e.g.,* where UAVs are used to survey an area). Therefore, we must quantify the trajectory shape difference between the planned path and the system POV trajectory, respectively both compared to the 'normal' paths. To do that, we use total Bregman divergence (TBD) [62] that compares the gaussian mixture model (GMM) [63] of one shape to another. Due to inherent randomness of GMM, we generate the TBD value 10 times and graph the distribution as box plots. *A good stealthy attack should have a high TBD value for the attack trajectory and a low TBD value for the system POV.*

The metrics will show that REQUIEM induces significant deviations while remaining stealthy across various mission scenarios. The justification of the metrics *w.r.t.* capturing the stealthiness and the significance of the attack are answered by **RQ2** (§VI-D) and **RQ3** (§VI-E) respectively. Learnability of EKF is answered by **RQ4** (§VI-F). The next subsection specifies the experiment parameters and the process to show that REQUIEM is necessary for stealthy attacks.

B. Experiment Parameters

We design the experiments to show the effectiveness of REQUIEM as an stealthy attack under various mission scenarios. This subsection establishes the baseline attacks that we will compare REQUIEM against as well as the missions. To demonstrate the feasibility of REQUIEM, we test our approach in simulation with an open-source flight controller PX4 [22] with gazebo [64] physics simulation.

Baseline Attacks. To show that REQUIEM is necessary, we establish baseline attacks as naïve attacks where the attack strategy is agnostic to the input of target function. We consider six naïve attacks specified in Table II where either the position or the velocity values are manipulated with one of three naïve strategies: constant offset, random offset, and boiling frog.

Constant offset attack adds a constant value of 1 to the GPS position (*i.e.,* PCO) or the velocity (*i.e.,* VCO) values respectively. Random offset adds random values (uniform between -1 and 1) to the position (*i.e.,* PRO) or the velocity (*i.e.,* VRO) values every GPS update. Boiling-frog attack gradually adds values to position (*i.e.,* PBF) or velocity (*i.e.,* VBF) such that value 1 in injected over a minute. Boiling-frog attack is

designed to show that even minuscule injection can cause the attack to be *overt*.

Missions. Attack performance is affected by the kinematics of a mission since different mission implies different input distribution of the state estimator. To test the attack's robustness against various vehicular kinematics, we specify the following missions with different trajectories as described in Table III: circle, linear and hold.

TABLE III: Mission scenarios with an example mission trajectory and average estimation error over 10 trials, $E[\epsilon]$. Due to space constraints, exact parameters of the mission are defined in the Appendix B

Name	Description	Trajectory	$E[\epsilon]$
Circle	Vehicle moves forward with fixed yaw-rate for a specified time. The resulting trajectory forms multiple circles in north and east plane.		0.06147
Linear	Set-points are set along the lateral plane. The resulting trajectory forms a triangle		0.0418
Hold	Maintain a specified altitude above the starting point for a specified time.		0.0159

Each mission have different challenges for the attacker. Circle mission represent movement where the vehicle rotates while moving. The attack becomes complex as the effect of vehicle attitude and angular velocity must be considered. Linear mission represent the standard mission of going from point A to point B without rotation. Attacker must ensure that the vehicle thinks that it converged to the mission setpoints. Hold mission represent scenarios where there is no little to no movement making it harder for the attacker to exploit estimation error due to the movement.

AD Parameter. We compare the attack performance on both default onboard AD, χ^2 and the threshold based AD, τ, with mission tuned threshold to show that bypassing τ has stricter requirement. It highlights the significance of REQUIEM by demonstrating that only REQUIEM can bypass τ whereas naïve methods cannot.

χ^2 AD has default threshold of $\eta = 5$. However for τ AD, the threshold range, τ_{min} and τ_{max}, were chosen based on the nominal distribution of the residual values for each mission. Specifically, the standard deviation of the distribution was calculated and removed outliers outside two standard deviations. Then the 2.5% and 97.5% percentile were selected as τ_{min} and τ_{max} respectively (due to space constraints, the details are in the appendix).

Experiment Setup. Experiments ran in PX4 autopilot's software-in-the-loop (SITL) simulation with Gazebo[64] plugin with the default environment (*i.e.,* no environmental hazards such as wind). The simulation ran on quad-core intel i5, Ubuntu 22.04, 16GB RAM, AMD Radeon Vega 20 GPU. During the course of simulation, the autopilot's state estimation module exfiltrates the snapshot with the implementation specified in Table V whenever the sensor value is ready to update the state. The module receives the spoof value after the exfiltration and injects the value into the GPS sensor.

Process. We examine each of the attack approaches (6 naïve and 2 REQUIEM) against the three missions (*i.e.,* hold, linear, circle) each representing different mission kinematics. Each experiment is repeated 10 times and evaluated using the previously mentioned metric. Results of the experiments are organized in Table IV which shows the average performance of the attacks for each mission measured using the previously mentioned metrics.

C. RQ1: Did REQUIEM succeed in Stealthy Attacks?

The objective of a stealthy attack is to avoid tripping any anomaly detectors and causing the system POV trajectory to seem nominal *w.r.t.* the planned mission while the *true trajectory* due to the attack significantly deviates from the mission parameters (*i.e.,* a meaningful deviation). Justification and definitions of 'stealthiness' and 'meaningfulness' are discussed in the subsequent section §VI-D and §VI-E respectively. On both fronts, REQUIEM *successfully* achieved stealth since it *caused significant deviation without triggering anomaly detectors* especially for the circle and linear missions. For the hold mission, REQUIEM was successful until an average of 17.72m of deviation was achieved after which it was no longer stealthy. The reason for the break in stealth is that the hold mission has a very limited movement errors for REQUIEM to exploit. However all of the naïve attacks (Table II) also failed to be a stealthy for all missions, displaying the non-triviality of the problem. In the rest of this section, we discuss the attack *w.r.t.* the circle mission while the details and results from the other missions are in appendix H due to the space constraints.

REQUIEM was indeed stealthy since it did not trigger any alarms –for both anomaly detectors –for the duration of the mission and the system POV was nominal[13]. The performance of REQUIEM's direction bias (DB) attack in Figure 8c shows that the *residual stayed within the threshold range for the duration of the mission*. Notice that in Figure 8a, the system POV trajectory is near exact to the normal trajectory. Consequently in Figure 10, the total Bregman divergence (TBD) value of the system POV trajectory is within the normal trajectory range.

> **Insight.** The stealthy attack performance is bounded by the training set distribution as well as the allowable error for a mission. When the spoofer encounters sensor values or estimated state values that are outside of the training set distribution, the performance worsens. Therefore, it is important for the data collection stage (§IV-C) to collect snapshots representative of the deployment mission.

We can see that although the attack residual remained stealthy for DB attack illustrated in Figure 8c, the residual is slowly reaching the lower threshold as the deviation increases. The effect is due to the larger discrepancy between the estimated state and the true sensor value when compared to the ones captured in the during the data collection stage (*i.e.,* training set). The effect is more prominent in REQUIEM's NC attack as shown in Figure 9 where the attack exploits the downward momentum, causing faster deviation at the cost of becoming *overt* (around $70m$ of deviation). For DB, had the vehicle kept deviating away from the planned trajectory for a longer time, the attack would eventually become *overt*. However, the fact

---

[13]we use 'nominal' and 'expected' interchangeably

that spoofer still remained stealty even after about $25m$ of deviation, far exceeding the scenario captured in the training set, indicates that the model generalized well.

REQUIEM achieved a 23m northward deviation for a helictical trajectory, significantly different from the intended circle as shown in Figure 8a. As a result, the TBD value of the attack trajectory (in Figure 10) is significantly higher than the system POV. In contrast, naïve position boiling-frog attack (PBF), the only naïve attack method that had partial success (shown in Figure 11) achieved maximum of about $0.111m$ *stealthy* deviation, twice the average nominal error as shown in Table III. This is not comparable to REQUIEM which is two orders of magnitude higher. Even at the cost of becoming *overt* and reaching about $1m$ of deviation, the naïve PBF's *overt* attack trajectory is similar to the nominal trajectory, thus not causing any real deviations.

The effect of the loss function on the spoofer is noticeable as the Direction Bias successfully caused the trajectory to shift northward. Notice at the start of the circle mission, there is a counter clockwise momentum causing the vehicle to head south (visible in Figure 11a). Therefore, the REQUIEM's NC attack causes the trajectory to go downward since it prevents the sensor value from correcting the momentum. The dominating southward momentum increases even more for each southward motion in the nominal mission and decreases for each northward movement but not enough to counter the initial build up. In contrast, Direction Bias counters the southward momentum early on and steadily builds northward momentum.

D. RQ2: What is a "Stealthy" Attack?

Recall from Figure 3 where the deployment process includes an operator, anomaly detector and post-flight analysis. By the time of post-flight analysis, detecting the attack is too late. However, if the mission ran seemingly normal, post-flight analysis may not be warranted. Therefore, stealthiness means that the attack should seem nominal *w.r.t.* the anomaly detector and the operator (*i.e.,* system POV trajectory):

C1 Attack does not trigger the anomaly detector
C2 From the operator's perspective, the trajectory reported by the UAV must not be distinguishable from the normal trajectory.

Demonstrating *C1* requires recording the anomaly detection response using the previously mentioned Table IVa which describes the metrics. The table measures the AD responses of χ^2 and τ (*i.e.,* M1 and M3 respectively). Because τ is more stringent anomaly detector, we consider the attack to be stealthy if τ-AD does not alarm.

To show that the attack satisfies *C2*, the system POV trajecotry resulting from the attack is visually compared against the planned trajectory. The trajectory shape comparison is then quantified using total Bregman divergence (elaborated in Appendix F). The nominal value for the TBD is calculated by comparing trajectory between two sets of nominal run of the mission. Therefore, TBD value for the system POV trajectory staying within the nominal value indicates that the reported trajectory is seemingly normal.

MID	Name	Description
M1	Stealth χ^2	Determines if the attack triggered the alarm under the default AD (*i.e.*, χ^2) (Equation 3).
M2	χ^2 Max Stealthy ϵ	Maximum stealthy deviation achieved under χ^2 AD in meters.
M3	Stealth τ	Checks if τ-AD (Equation 4) raised an alarm.
M4	τ Max Stealthy ϵ	Maximum stealthy deviation achieved under τ-AD in meters

(a) Evaluation metrics

Mis.	MID	Naïve						REQUIEM	
		PCO	PRO	PBF	VCO	VRO	VBF	NC	DB
○	M1	✓	✓	✓	✓	✓	✓	✓	✓
	M2	1.097	0.365	1.122	2.730	0.934	2.756	110.62	23.78
	M3	✗	✗	(✓/✗)	✗	✗	✗	(✓/✗)	✓
	M4	✗	✗	0.111	✗	✗	✗	77.09	23.78
△	M1	✓	✓	✓	✓	✓	✓	✓	✓
	M2	1.030	0.300	0.356	2.745	0.891	0.893	3.56	11.10
	M3	✗	✗	(✓/✗)	✗	✗	✗	✓	✓
	M4	✗	✗	0.0936	✗	✗	✗	3.56	11.10
(hold)	M1	✓	✓	✓	✓	✓	✓	✓	✓
	M2	1.018	0.357	1.135	2.756	0.925	2.497	50.62	289.26
	M3	✗	✗	✗	✗	✗	✗	(✓/✗)	(✓/✗)
	M4	✗	✗	✗	✗	✗	✗	17.72	9.47

(b) Attack evaluation for each mission (Mis.)

TABLE IV: Attack evaluation on the three mission scenarios: circle (○), linear (△) and hold (▾). The attack that resulted with the most deviation given a metric is shaded with purple. ✓indicates that the attack remained stealthy throughout the mission in all trials, (✓/✗) indicates that attacked start out stealthy for all trials but became overt for at least one of the trials, and ✗indicates that the attack was not stealthy to begin with for at-least one of the trials.

(a) *DB Mission Trajectory*   (b) *DB Observed GPS History*   (c) *DB Residual History*   (d) *DB χ^2 Score*
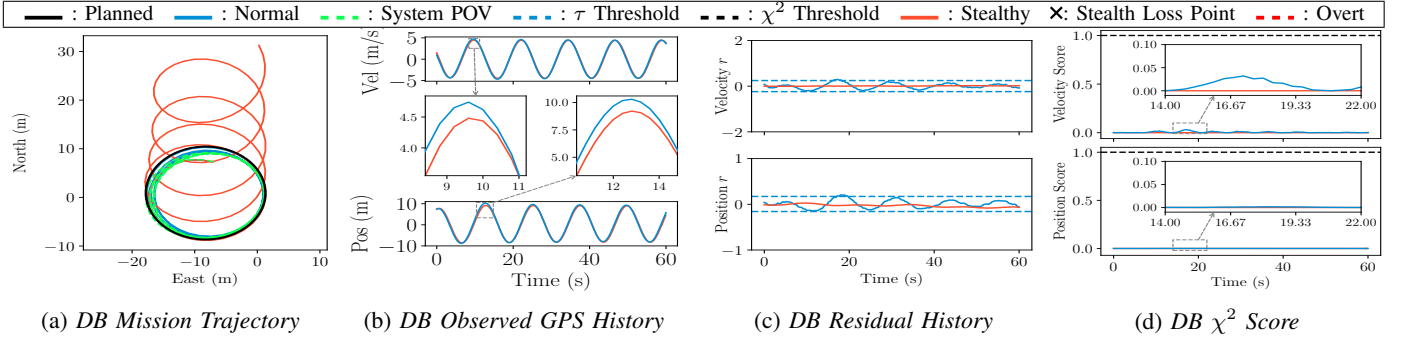
Fig. 8: REQUIEM's directional bias (DB) attack on circle mission. *(a)* shows that system POV trajectory (green) is not distinguishable from the normal where as the attack trajectory, colored in red, deviated far away. During the attack, GPS velocity and position readings are seemingly normal as shown in *(b)* where the difference is minuscule. *(c)* Resulting residual stayed within the τ threshold, avoiding the activation of the alarm. Onboard detector's inefficacy is demonstrated in *(d)* where χ^2 score remained near-zero throughout.

(a) *NC Mission Trajectory*   (b) *NC Residual*
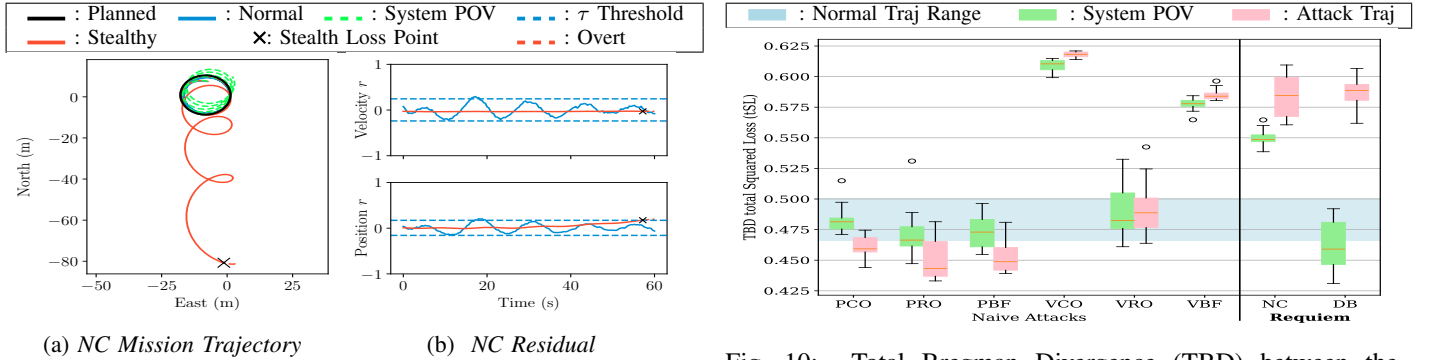
Fig. 9: REQUIEM's no correction (NC) attack on circle mission. (a) shows the UAV deviating a significant amount but became overt near the end. (b) shows the point of the stealth loss. The loss is due to receiving large position values (*i.e.,* >77m) as input which is far outside of the training distribution.
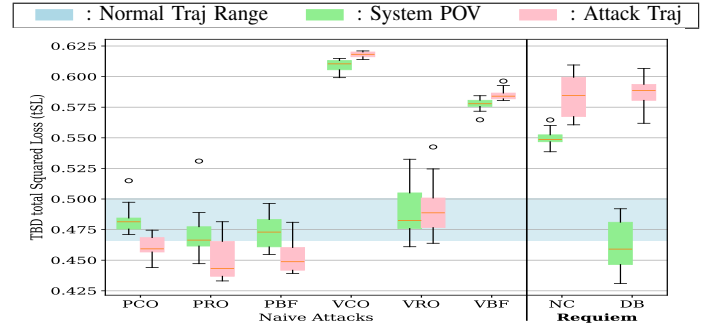
Fig. 10: Total Bregman Divergence (TBD) between the planned to the attack and the system POV trajectories. Larger difference between the green and the pink box indicates larger deviation. The plot shows that REQUIEM attacks achieve significant deviation whereas naïve attacks fails to do so.

(a) PBF Mission Trajectory (b) PBF Residual
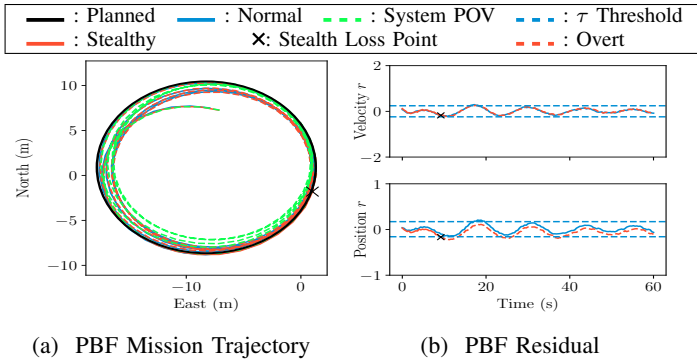
Fig. 11: Naïve position boiling frog (PBF) attack on circle mission. (a) shows negligible effect of PBF on the trajectory but becomes overt after traveling half circle. (b) shows the position residual gradually straying away from the nominal residual resulting in an alarm.

E. **RQ3**: *What is a "Meaningful" Attack?*

We have seen in the evaluation Table IV that the naïve PBF achieves around $1m$ maximum overt deviation. However, the resulting attack trajectory is similar compared to the system POV as well as the planned in Figure 11a. In contrast, an attack can have the same maximum deviation but result in a trajectory similar to Figure 1 where attack trajectory is an offset of the system POV. Such offset can be critical for missions such as surveillance or reconnaissance. Therefore, we consider an attack to be "meaningful" when the attack trajectory has a "different" shape compared to the planned trajectory.

> **Insight.** It implies that only comparing the maximum ϵ between the system POV and the attack trajectories does not tell the full story since the attack trajectory can be similar to the planned while having a "significant" deviation.

To demonstrate that the shapes are different, we use total Bregman divergence (elaborated in Appendix F) for shape comparison; comparing the attack to the planned trajectory to indicate that the two shapes are significantly different.

F. **RQ4** : *How do we measure how well the surrogate model "learn" the target function?*

A standard process in machine learning for evaluating a model's performance is by testing the model against a set of data that were not in the training set. For this paper, the test set is the 8 separate mission runs addressed in §IV-C and the performance measurement is the mean squared error loss.

Due to the space constraints, the results are explained in Appendix G. It provides the following insight:

> **Insight.** The surrogate models managed to learn the target function since important functions in safety critical systems, such as state estimator, are designed to be predictable, especially *w.r.t.* the correlation between their ins and outs.

We continue the discussion of *"learnability"* of the state estimation algorithm in §VII-A, reinforcing why surrogate models can *"learn"* for a large number of safety critical systems including autonomous vehicles.

(a) *Normal trajectory [NORTH WIND]* (b) *Zoomed trajectory from Figure 9a [NO WIND]*
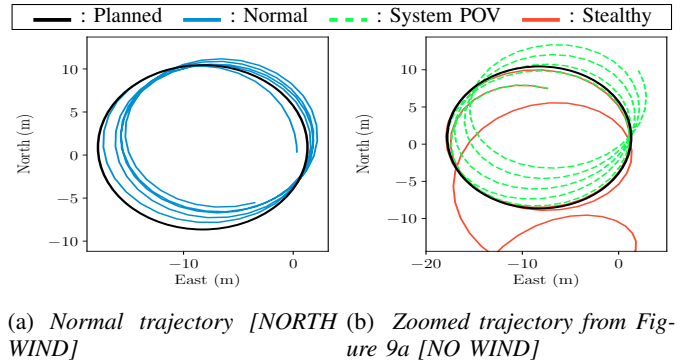
Fig. 12: Notice that the normal trajectory (blue) from a system with wind (a) does not seem out of place when compared to the System POV (green) without wind in (b). The latter is what the operator sees and it is hard to distinguish it from an attack, in the presence of wind or otherwise.

(a) *DB [NORTH WIND]* (b) *DB [SOUTH WIND]*
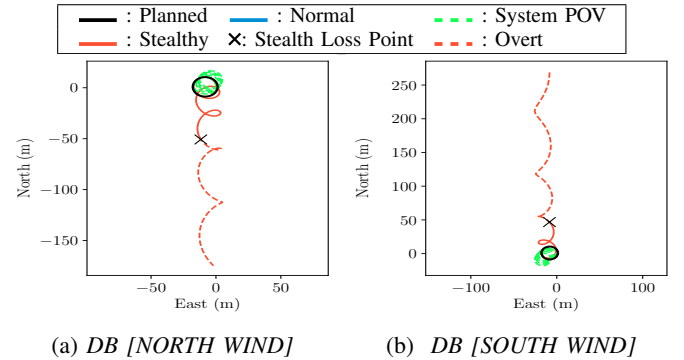
Fig. 13: Comparison of wind direction affecting the DB attack shows that direction of deviation is opposite to the wind in both cases (a) and (b) but achieves greater deviation under southern wind. [NORTH WIND] On average over 10 trials, DB achieves 48.63m of stealthy deviation and 173.13m of overt deviation. [SOUTH WIND] On average, DB achieves 38.73m of stealthy deviation and 286.84m of overt deviation.

G. **RQ5** *How do environmental factors affect the attack?*

Environment factors (such as wind and other weather conditions) can cause the UAV to stray slightly away from its planned path as shown in Figure 12. However, it can also be unpredictable and affect the mission performance of the UAV. To test the robustness of REQUIEM, we tested the spoofer, that was trained under the default environment (*i.e.,* no wind), under a windy weather going steadily at 4m/s north. We expected REQUIEM to exploit the wind due to the natural momentum it provides to the UAV and deviate *along* the direction of the wind. However, the opposite was true:

> **Insight.** REQUIEM exploits the control system's *response* to the wind, resulting in a stealthy deviation in a direction opposite to the wind direction (*i.e.,* along the direction of the control's response).
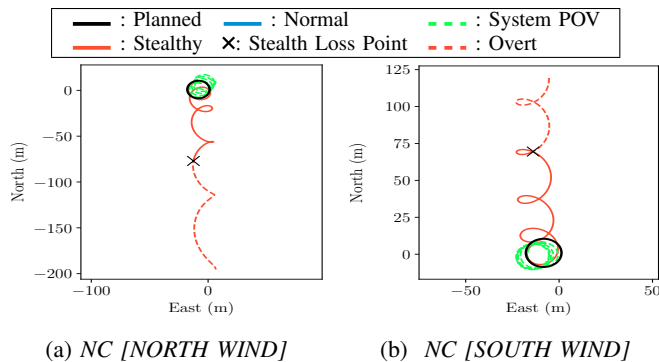
Fig. 14: Wind also makes NC attack to cause deviation opposite to the wind as shown in (a) and (b). [NORTH WIND] On average over 10 trials, NC achieves 82.44m of stealthy deviation and 183.67m of overt deviation. [SOUTH WIND] On average, NC under southern wind achieves 73.85m of stealthy deviation and 140.04m of overt deviation.

The result makes sense because the control already accounts for the wind during the takeoff (*i.e.,* before the mission start and the attack). The NC attack's prevention of the sensor from correcting the state estimation results in the control system overcompensation, thus causing deviations as shown in Figure 14.

Comparing among the DB and NC attacks, the latter fared better since it remains stealthy for a longer time than DB. However, DB achieves a larger deviation, at the cost of stealthiness: we can see the effect of the DB's bias effort in Figure 13a where the UAV deviates less towards south and Figure 13b where it deviates more towards north. Both methods of REQUIEM *still resulted in a significant stealthy deviation prior to becoming overt*.

VII. DISCUSSION

We discuss the conditions required for REQUIEM to work (*i.e.,* learnability of the target function), potential defenses against REQUIEM and limitation of our methods.

A. Learnability

For a target function to be *learnable*, it must be differentiable or have bounded continuity: current ML theory assumes *k-lipshitz continuity* that places upper bounds on how fast a function output can change *w.r.t.* the input [1], [2], [3]. Differentiabliliy and the bounded continuity also apply for state estimator analysis to provide guarantees about its behavior [65], [66], [67], [68]. Safety critical systems (*e.g.,* autonomous vehicles or avionics) require predictability guarantees. Hence, *our approach exploits the analyzability of the estimation algorithms*.

Furthermore, our approach is slightly more relaxed: the continuity constraint is *only a subset of the input domain reflective of the vehicle's mission* (*i.e.,* locally Lipschitz). This is because the state estimator will only observe regions of the input domain corresponding to the deployment mission throughout the attack.

B. Defending against REQUIEM

Because we exploit errors that naturally exist in vehicular systems, defending against REQUIEM is very tricky. Even analyzing the history of sensor values will be insufficient as the spoofed values *look like* they are drawing from the same distribution as the original sensor values. One could try to prevent the payload from reaching the target system, *i.e.,*; making it difficult to 'learn' the target function with surrogate models or increasing sensor redundancy where multiple sensors provide similar overlapping information. Another potential solution is via randomization [69], [70] since it makes the mapping between the input and the output seem more disjoint. Though in autonomous vehicle domains, applying randomization to state estimation may be difficult since the vehicular motion is bounded by the standard model of physics.

C. Limitations

Our approach is stealthy only with-in the sensor perspective: it will be stealthy against any anomaly detector that detects problems based only on sensors. If there are redundant sensors (*e.g.,* at least two of a same kind), then the spoofer must spoof all, which makes the attack harder; though we can handle this by using multiple surrogates and spoofers. We do not address stealthiness outside of that scope. In the threat model, we assume that the adversary is aware of the state of the module. We acknowledge that it is difficult to map every variable. However, it can be alleviated in some cases: identifying regions of memory corresponding to module parameters can be achieved by setting unique parameter values with MAVLink [71] commands and searching for them. Regions of memory for sensors can be discovered by manipulating sensor messages that reach the estimation module. We do not consider stealthiness under post-flight analysis where a full detailed log of the vehicle is used to diagnose any potential issues with the vehicle.

VIII. RELATED WORK

We now explore some relevant contemporary strategies for creating Adversarial Examples (AEs) and their deployment in autonomous systems, focusing especially on stealthy spoofing attacks, and note how REQUIEM differs from and/or is better than comparable work.

Spoofing Attacks. Recent studies[72], [73], [74], [75], [76], [77] involving adaptive adversarial attack approaches *etc.* highlight the inherent brittleness of Deep Neural Networks (DNNs), and consequently the susceptibility of anomaly detectors in various autonomous systems including UAV Autopilot systems. The Adversarial Examples (AEs) produced by REQUIEM are adaptable and broadly applicable across a variety of systems since they do not depend on specific system states or sensors.

Detection and Defense Against Spoofing Attacks. More recent anomaly detectors, such as SAVIOR [18], use EKF-based state estimation to detect anomalies using the residual. The essence of their EKF is similar to that used in PX4. Hence, we can apply our methods to generate another spoofer for this EKF and tune the loss functions to evade detection by such methods. Most of the current methods to detect[78], [79], [80] and defend against spoofing attacks such as sensor

watermarking[81], [82] *etc.* work best with fixed spoofing strategies (*i.e.,* naïve attacks and replay attacks). REQUIEM makes such defenses moot since our spoof generation is a function of the actuation command and we demonstrated that a single model is capable of minimizing the residual under various kinematics.

GANs in Crafting Adversarial Examples. The advent of Generative Adversarial Networks (GANs)[83], [84] has paved the way for their application in Adversarial Example (AE) generation including AEs for vision and perception-based Deep Neural Networks (DNNs)[85], [86], [87], [88], [89], [90]. AdvGAN[89] is perhaps the first to demonstrate AE crafting, and AC-GAN[91] builds on AdvGAN to generate AEs that bias the victim model toward a specific classification. However, the use of GANs in autonomous vehicles presents unique challenges, especially with non-vision sensors such as IMU or GPS that are typically governed by non-DNN, regression-based models. This is a departure from the traditional focus of GANs on vision models that are more susceptible due to their deep learning nature. REQUIEM differs from AdvGAN and AC-GAN in several key aspects: *(i)* we target the modeling errors in the target function, whereas AdvGAN focuses on the generalizability issues of the deep target model. *(ii)* The target function in AdvGAN is categorical, while in REQUIEM, it is a regression function. *(iii)* REQUIEM introduces surrogate models to learn specific 'slices' of the target function, aiming at particular subsets of this function.

REFERENCES

[1] Ulrike von Luxburg and Olivier Bousquet. Distance-based classification with lipschitz functions. *J. Mach. Learn. Res.*, 5(Jun):669–695, 2004.

[2] Damek Davis, Dmitriy Drusvyatskiy, Sham Kakade, and Jason D Lee. Stochastic subgradient method converges on tame functions. *Foundations of computational mathematics*, 20(1):119–154, 2020.

[3] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. *Advances in neural information processing systems*, 30, 2017.

[4] Luana Steffen. Military drone autonomously tracks targets using infrared. *Intelligent Living*, 2021.

[5] David Roza. The army is testing drones that can deliver life-saving blood to the battlefield. *Task and Purpose*, 2021.

[6] Isabella Lee, Vignesh Babu, Matthew Caesar, and David Nicol. Deep reinforcement learning for uav-assisted emergency response. In *MobiQuitous 2020-17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, pages 327–336, 2020.

[7] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960.

[8] R Spencer Hallyburton, Amir Khazraei, and Miroslav Pajic. Optimal myopic attacks on nonlinear estimation. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 5480–5485. IEEE, 2022.

[9] Miroslav Pajic, Paulo Tabuada, Insup Lee, and George J Pappas. Attack-resilient state estimation in the presence of noise. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 5827–5832. IEEE, 2015.

[10] Daniel Arnström and André MH Teixeira. Stealthy deactivation of safety filters. *arXiv preprint arXiv:2403.17861*, 2024.

[11] Amir Khazraei, Spencer Hallyburton, Qitong Gao, Yu Wang, and Miroslav Pajic. Learning-based vulnerability analysis of cyber-physical systems. In *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPS)*, pages 259–269. IEEE, 2022.

[12] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Conference on robot learning*, pages 1–16. PMLR, 2017.

[13] Samir Bouabdallah and Roland Siegwart. Full control of a quadrotor. In *2007 IEEE/RSJ international conference on intelligent robots and systems*, pages 153–158. Ieee, 2007.

[14] Amir Khazraei and Miroslav Pajic. Vulnerability analysis of nonlinear control systems to stealthy false data injection attacks. *arXiv preprint arXiv:2310.04516*, 2023.

[15] Amir Khazraei, Henry Pfister, and Miroslav Pajic. Resiliency of perception-based controllers against attacks. In *Learning for Dynamics and Control Conference*, pages 713–725. PMLR, 2022.

[16] Amir Khazraei, Haocheng Meng, and Miroslav Pajic. Stealthy perception-based attacks on unmanned aerial vehicles. *arXiv preprint arXiv:2303.02112*, 2023.

[17] Aolin Ding, Matthew Chan, Amin Hass, Nils Ole Tippenhauer, Shiqing Ma, and Saman Zonouz. Get your cyber-physical tests done! data-driven vulnerability assessment of robotic aerial vehicles. In *2023 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 67–80. IEEE, 2023.

[18] Raul Quinonez, Jairo Giraldo, Luis Salazar, Erick Bauman, Alvaro Cardenas, and Zhiqiang Lin. {SAVIOR}: Securing autonomous vehicles with robust physical invariants. In *29th {USENIX} Security Symposium ({USENIX} Security 20)*, pages 895–912, 2020.

[19] Christopher M Bishop et al. *Neural networks for pattern recognition*. Oxford university press, 1995.

[20] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.

[21] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, Andrew Y Ng, et al. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.

[22] Lorenz Meier, Dominik Honegger, and Marc Pollefeys. Px4: A node-based multithreaded open source robotics framework for deeply embedded platforms. In *2015 IEEE international conference on robotics and automation (ICRA)*, pages 6235–6240. IEEE, 2015.

[23] Robert E Lyons and Wouter Vanderkulk. The use of triple-modular redundancy to improve computer reliability. *IBM journal of research and development*, 6(2):200–209, 1962.

[24] Sandro Pinto and Nuno Santos. Demystifying arm trustzone: A comprehensive survey. *ACM computing surveys (CSUR)*, 51(6):1–36, 2019.

[25] Frank McKeen, Ilya Alexandrovich, Ittai Anati, Dror Caspi, Simon Johnson, Rebekah Leslie-Hurd, and Carlos Rozas. Intel® software guard extensions (intel® sgx) support for dynamic memory management inside an enclave. In *Proceedings of the Hardware and Architectural Support for Security and Privacy 2016*, pages 1–9. 2016.

[26] Man-Ki Yoon, Sibin Mohan, Jaesik Choi, Jung-Eun Kim, and Lui Sha. Securecore: A multicore-based intrusion detection architecture for real-time embedded systems. In *2013 IEEE 19th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 21–32. IEEE, 2013.

[27] K Priandana, M Hazim, and B Kusumoputro Wulandari. Development of autonomous uav quadcopters using pixhawk controller and its flight data acquisition, 2020 int. In *Conf. Comput. Sci. Its Appl. Agric. ICOSICA*, volume 2020, 2020.

[28] Navio2 – autopilot hat for raspberry pi. https://navio2.hipi.io/.

[29] Steven S. Beauchemin and John L. Barron. The computation of optical flow. *ACM computing surveys (CSUR)*, 27(3):433–466, 1995.

[30] Yilin Mo and Bruno Sinopoli. False data injection attacks in control systems. In *Preprints of the 1st workshop on Secure Control Systems*, volume 1, 2010.

[31] Ilija Jovanov and Miroslav Pajic. Relaxing integrity requirements for attack-resilient cyber-physical systems. *IEEE Transactions on Automatic Control*, 64(12):4843–4858, 2019.

[32] Tzu-Ching Horng. *A comparative analysis of supply chain management practices by Boeing and Airbus: long-term strategic implications*. PhD thesis, Massachusetts Institute of Technology, 2006.

[33] Muhammad Rashid and Irfan Manarvi. Vendor assessment and procurement decision making through data mining in aviation industry. In

2009 International Conference on Computers & Industrial Engineering, pages 1470–1474. IEEE, 2009.

[34] Yaxin Li. Supply chain management of aircraft industry. In *2018 8th International Conference on Management, Education and Information (MEICI 2018)*, pages 218–223. Atlantis Press, 2018.

[35] Emma Loxton, Henry Marcil, Mike Parkins, and Andrew Tingley. Addressing continued turbulence: The commercial-aerospace supply chain, April 2024.

[36] Sibin Mohan, Man Ki Yoon, Rodolfo Pellizzoni, and Rakesh Bobba. Real-time systems security through scheduler constraints. In *2014 26th Euromicro Conference on Real-Time Systems*, pages 129–140. IEEE, 2014.

[37] Rodolfo Pellizzoni, Neda Paryab, Man-Ki Yoon, Stanley Bak, Sibin Mohan, and Rakesh B Bobba. A generalized model for preventing information leakage in hard real-time systems. In *21st IEEE Real-Time and Embedded Technology and Applications Symposium*, pages 271–282. IEEE, 2015.

[38] Taegyu Kim, Chung Hwan Kim, Junghwan Rhee, Fan Fei, Zhan Tu, Gregory Walkup, Xiangyu Zhang, Xinyan Deng, and Dongyan Xu. {RVFuzzer}: Finding input validation bugs in robotic vehicles through {Control-Guided} testing. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 425–442, 2019.

[39] Hongjun Choi, Wen-Chuan Lee, Yousra Aafer, Fan Fei, Zhan Tu, Xiangyu Zhang, Dongyan Xu, and Xinyan Deng. Detecting attacks against robotic vehicles: A control invariant approach. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 801–816, 2018.

[40] Mahsa Foruhandeh, Yanmao Man, Ryan Gerdes, Ming Li, and Thidapat Chantem. Simple: Single-frame based physical layer identification for intrusion detection and prevention on in-vehicle networks. In *Proceedings of the 35th annual computer security applications conference*, pages 229–244, 2019.

[41] Se-Yeon Jeong, I-Ju Choi, Yeong-Jin Kim, Yong-Min Shin, Jeong-Hun Han, Goo-Hong Jung, and Kyoung-Gon Kim. A study on ros vulnerabilities and countermeasure. In *Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, pages 147–148, 2017.

[42] Bernhard Dieber, Ruffin White, Sebastian Taurer, Benjamin Breiling, Gianluca Caiazza, Henrik Christensen, and Agostino Cortesi. Penetration testing ros. *Robot Operating System (ROS) The Complete Reference (Volume 4)*, pages 183–225, 2020.

[43] Rafael R Teixeira, Igor P Maurell, and Paulo LJ Drews. Security on ros: analyzing and exploiting vulnerabilities of ros-based systems. In *2020 Latin American robotics symposium (LARS), 2020 Brazilian symposium on robotics (SBR) and 2020 workshop on robotics in education (WRE)*, pages 1–6. IEEE, 2020.

[44] Federico Maggi, Davide Quarta, Marcello Pogliani, Mario Polino, Andrea M Zanchettin, and Stefano Zanero. Rogue robots: Testing the limits of an industrial robot's security. *Trend Micro, Politecnico di Milano, Tech. Rep*, pages 1–21, 2017.

[45] Karel Domin, Iraklis Symeonidis, and Eduard Marin. Security analysis of the drone communication protocol: Fuzzing the mavlink protocol. 2016.

[46] Javid Habibi, Aditi Gupta, Stephen Carlsony, Ajay Panicker, and Elisa Bertino. Mavr: Code reuse stealthy attacks and mitigation on unmanned aerial vehicles. In *2015 IEEE 35th International Conference on Distributed Computing Systems*, pages 642–652. IEEE, 2015.

[47] Chien-Ying Chen, Sibin Mohan, Rodolfo Pellizzoni, Rakesh B Bobba, and Negar Kiyavash. A novel side-channel in real-time schedulers. In *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 90–102. IEEE, 2019.

[48] Songran Liu, Nan Guan, Dong Ji, Weichen Liu, Xue Liu, and Wang Yi. Leaking your engine speed by spectrum analysis of real-time scheduling sequences. *Journal of Systems Architecture*, 97:455–466, 2019.

[49] Stephen Jacklin. Certification of safety-critical software under do-178c and do-278a. In *Infotech@ Aerospace 2012*, page 2473. 2012.

[50] Markus Hochstrasser, Stephan Myschik, and Florian Holzapfel. Application of a process-oriented build tool for flight controller development along a do-178c/do-331 process. In *Model-Driven Engineering and Software Development: 6th International Conference, MODELSWARD

2018, Funchal, Madeira, Portugal, January 22-24, 2018, Revised Selected Papers 6*, pages 380–405. Springer, 2019.

[51] Andrew J Kerns, Daniel P Shepard, Jahshan A Bhatti, and Todd E Humphreys. Unmanned aircraft capture and control via gps spoofing. *Journal of Field Robotics*, 31(4):617–636, 2014.

[52] Juhwan Noh, Yujin Kwon, Yunmok Son, Hocheol Shin, Dohyun Kim, Jaeyeong Choi, and Yongdae Kim. Tractor beam: Safe-hijacking of consumer drones with adaptive gps spoofing. *ACM Transactions on Privacy and Security (TOPS)*, 22(2):1–26, 2019.

[53] Kexiong Curtis Zeng, Shinan Liu, Yuanchao Shu, Dong Wang, Haoyu Li, Yanzhi Dou, Gang Wang, and Yaling Yang. All your {GPS} are belong to us: Towards stealthy manipulation of road navigation systems. In *27th USENIX security symposium (USENIX security 18)*, pages 1527–1544, 2018.

[54] Junjie Shen, Jun Yeon Won, Zeyuan Chen, and Qi Alfred Chen. Drift with devil: Security of {Multi-Sensor} fusion based localization in {High-Level} autonomous driving under {GPS} spoofing. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 931–948, 2020.

[55] Yunmok Son, Hocheol Shin, Dongkwan Kim, Youngseok Park, Juhwan Noh, Kibum Choi, Jungwoo Choi, and Yongdae Kim. Rocking drones with intentional sound noise on gyroscopic sensors. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 881–896, 2015.

[56] Timothy Trippel, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu. Walnut: Waging doubt on the integrity of mems accelerometers with acoustic injection attacks. In *2017 IEEE European symposium on security and privacy (EuroS&P)*, pages 3–18. IEEE, 2017.

[57] Xiaoyu Ji, Yushi Cheng, Yuepeng Zhang, Kai Wang, Chen Yan, Wenyuan Xu, and Kevin Fu. Poltergeist: Acoustic adversarial machine learning against cameras and computer vision. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 160–175. IEEE, 2021.

[58] DJI. Dji onboard sdk (osdk) 4.1.0, 2021.

[59] Rafael Perez-Segui, Pedro Arias-Perez, Javier Melero-Deza, David Perez-Saura, Miguel Fernandez-Cortizas, and Pascual Campoy. Simulation-guided testing for autonomous aerial robotics applications.

[60] Parrot. Parrot sphinx guide book, 2023.

[61] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Icml*, 2010.

[62] Meizhu Liu, Baba C Vemuri, Shun-Ichi Amari, and Frank Nielsen. Total bregman divergence and its applications to shape retrieval. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3463–3468. IEEE, 2010.

[63] Geoffrey J McLachlan, Sharon X Lee, and Suren I Rathnayake. Finite mixture models. *Annual review of statistics and its application*, 6:355–378, 2019.

[64] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 3, pages 2149–2154. IEEE, 2004.

[65] Robert Hermann and Arthur Krener. Nonlinear controllability and observability. *IEEE Transactions on automatic control*, 22(5):728–740, 1977.

[66] Guoquan P Huang, Anastasios I Mourikis, and Stergios I Roumeliotis. Genearalized analysis and improvement of the consistency for ekf-based slam. *University of Minnesota*, pages 2008–0001, 2008.

[67] José A Castellanos, José Neira, and Juan D Tardós. Limits to the consistency of ekf-based slam. *IFAC Proceedings Volumes*, 37(8):716–721, 2004.

[68] Douglas A Allan, James Rawlings, and Andrew R Teel. Nonlinear detectability and incremental input/output-to-state stability. *SIAM Journal on Control and Optimization*, 59(4):3017–3039, 2021.

[69] Man-Ki Yoon, Sibin Mohan, Chien-Ying Chen, and Lui Sha. Taskshuffler: A schedule randomization protocol for obfuscation against timing inference attacks in real-time systems. In *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 1–12. IEEE, 2016.

[70] Chien-Ying Chen, Debopam Sanyal, and Sibin Mohan. Indistinguishability prevents scheduler side channels in real-time systems. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 666–684, 2021.

[71] Lorenz Meier. Introduction-mavlink developer guide.

[72] Changhong Fu, Sihang Li, Xinnan Yuan, Junjie Ye, Ziang Cao, and Fangqiang Ding. Ad2attack: Adaptive adversarial attack on real-time uav tracking. *2022 International Conference on Robotics and Automation (ICRA)*, pages 5893–5899, 2022.

[73] David I Urbina, Jairo A Giraldo, Alvaro A Cardenas, Nils Ole Tippenhauer, Junia Valente, Mustafa Faisal, Justin Ruths, Richard Candell, and Henrik Sandberg. Limiting the impact of stealthy attacks on industrial control systems. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 1092–1105, 2016.

[74] Ibrahim Sobh, Ahmed Hamed, V. Kumar, and S. Yogamani. Adversarial attacks on multi-task visual perception for autonomous driving. *ArXiv*, abs/2107.07449, 2021.

[75] Shang-Tse Chen, Cory Cornelius, Jason Martin, and Duen Horng Chau. Shapeshifter: Robust physical adversarial attack on faster r-cnn object detector. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2018, Dublin, Ireland, September 10–14, 2018, Proceedings, Part I 18*, pages 52–68. Springer, 2019.

[76] Giulio Lovisotto, Henry Turner, Ivo Sluganovic, Martin Strohmeier, and Ivan Martinovic. {SLAP}: Improving physical adversarial examples with {Short-Lived} adversarial perturbations. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1865–1882, 2021.

[77] Dawn Song, Kevin Eykholt, Ivan Evtimov, Earlence Fernandes, Bo Li, Amir Rahmati, Florian Tramer, Atul Prakash, and Tadayoshi Kohno. Physical adversarial examples for object detectors. In *12th USENIX workshop on offensive technologies (WOOT 18)*, 2018.

[78] Thomas Hickling, Nabil Aouf, and Phillippa Spencer. Robust adversarial attacks detection based on explainable deep reinforcement learning for uav guidance and planning. *IEEE Transactions on Intelligent Vehicles*, 2023.

[79] Kyo Kim, Siddhartha Nalluri, Ashish Kashinath, Yu Wang, Sibin Mohan, Miroslav Pajic, and Bo Li. Security analysis against spoofing attacks for distributed uavs. *2020 NDSS Decentralized IoT Systems and Security Workshop*, 2020.

[80] Hyungsub Kim, Rwitam Bandyopadhyay, Muslum Ozgur Ozmen, Z Berkay Celik, Antonio Bianchi, Yongdae Kim, and Dongyan Xu. A systematic study of physical sensor attack hardness. In *2024 IEEE Symposium on Security and Privacy (SP)*, pages 143–143. IEEE Computer Society, 2024.

[81] Yilin Mo and Bruno Sinopoli. Secure control against replay attacks. In *2009 47th annual Allerton conference on communication, control, and computing (Allerton)*, pages 911–918. IEEE, 2009.

[82] Arunava Naha, André MH Teixeira, Anders Ahlén, and Subhrakanti Dey. Quickest detection of deception attacks on cyber–physical systems with a parsimonious watermarking policy. *Automatica*, 155:111147, 2023.

[83] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.

[84] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.

[85] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

[86] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z Berkay Celik, and Ananthram Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, pages 372–387. IEEE, 2016.

[87] Fnu Suya, Jianfeng Chi, David Evans, and Yuan Tian. Hybrid batch attacks: Finding black-box adversarial examples with limited queries. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1327–1344, 2020.

[88] Yandong Li, Lijun Li, Liqiang Wang, Tong Zhang, and Boqing Gong. Nattack: Learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. In *International Conference on Machine Learning*, pages 3866–3876. PMLR, 2019.

TABLE V: Types of data collected in a snapshot.

Type	Name	Description
State	Quaternion	Complex numbers to represent attitude of the vehicle
	Position and Velocity	Represented in (m) and (m/s) respectively in North, East and Down coordinate frame.
	Δ Angle Bias	Bias in IMU angle measurement in radians
	Δ Velocity Bias	Bias in IMU velocity measurement in (m/s)
	Earth magnetic field	Represented in gauss
	Body magnetic field	Magnetic field of the vehicle in gauss
	Wind velocity	Used to calculate drag force typically used in fixed wing vehicles.
GPS	Position and Velocity	Represented in (m) and (m/s) respectively in North, East and Down coordinate frame.
	Horizontal Accuracy	Horizontal position noise variance
	Vertical Accuracy	Vertical position noise variance
	Speed Accuracy	Velocity observation noise variance
IMU	Angular Rate	Rotation speed of the vehicle in (Rad/s)
Prediction Output	Body-to-earth frame	Rotation matrix that transforms from body frame to earth frame.

[89] Chaowei Xiao, Bo Li, Jun-Yan Zhu, Warren He, Mingyan Liu, and Dawn Song. Generating adversarial examples with adversarial networks. *arXiv preprint arXiv:1801.02610*, 2018.

[90] Yulong Cao, Chaowei Xiao, Benjamin Cyr, Yimeng Zhou, Won Park, Sara Rampazzi, Qi Alfred Chen, Kevin Fu, and Z Morley Mao. Adversarial sensor attack on lidar-based perception in autonomous driving. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 2267–2281, 2019.

[91] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International conference on machine learning*, pages 2642–2651. PMLR, 2017.

[92] https://github.com/mavlink/MAVSDK-Python.

[93] Ardupilot - versatile, trusted, open.

APPENDIX

This appendix includes the referenced details in the main content of this paper.

A. Implementation Specifications

This subsection describes the implementation detail to run the experiments from §VI. We use PX4 Software-In-The-Loop (SITL) with Gazebo for physics simulation. The simulation uses default world with the default quadrotor model in PX4. We use MAVSDK-python[92] to send commands to the UAV. We use a combination of TCP and mmap for exfiltration of the snapshots from the PX4 to the attack server. where TCP packets were used to signal that data is ready for in the mmap. The specification is elaborated in more detail on the website??.

B. Referenced Missions

This section describes the exact procedure for each of the referenced missions from §IV-C and §VI-B: brownian motion, circle, linear, and hold. The exact implementation of each are specified on the website.

Brownian Motion. Brownian Motion is the motion resulting from a random walk. It is used to capture as much diverse movement as possible to ensure that the collected data can be used to train the model meaningfully. We get the setpoints for Brownian Motion by using a random walk algorithm on a 2D plane where the next setpoint is the current setpoint moved by a value sampled from 2D Gaussian with zero mean and unit variance as shown in Equation 12:

$$p_0 = (0,0); p_{t+1} = p_t + \epsilon_{t+1}; \epsilon_{t+1} \sim \mathcal{N}((0,0), I) \quad (12)$$

p_0 is the origin position with $0m$ east and $0m$ north respectively. Identity matrix I is the covariance matrix (*i.e.,* no covariance between east and north). To add rotational kinematics into the dataset, we randomize the heading for each setpoint. The UAV is considered converged to the setpoint if the vehicle is within $0.3m$ and $1°$ of the setpoint. Example trajectories for a Brownian Motion Mission are shown in Figure 15. Notice that the upper and lower trajectories cover different areas with different trajectory shape.

Fig. 15: Two example trajectories of Brownian motion mission used in the data collection stage (§IV-C).

Circle. Circle Mission simulates a surveillance mission where a UAV circles a designated area with a part of the UAV constantly facing toward the center. Specifically, the UAV moves forward at $5m/s$ while rotating $30°$ counter-clockwise resulting in a circle with $9.549m$ radius centered at $(-8.25m, 0.9m)$ east-north coordinate.

Linear. Linear Mission represents a mission involving logistics such as delivery (*e.g.,* Amazon Air) where the UAV reaches each of the check points in a straight line. Three setpoints are used to cover the movement solely along north axis, east axis and both axes. At altitude of 10m, the UAV moves from the origin to (10m east, 0m north), then to (10m east, 10m north), and returning back to the origin (0m east, 0m north).

Hold. In all types of missions, the UAV may need to briefly hold its position. The Hold Mission is accomplished by only sending a single setpoint: 0m north, 0m east of the starting point at 10m altitude.

C. Generative Adversarial Network

This subsection describes Generative Adversarial Network (GAN) [20] to provide a context into the interaction between the surrogate and the spoofer models mentioned in §V. GAN is composed of two models: generator and discriminator. There

are variations and improvements of GAN in various fields and/or applications. We only present the essence of GAN and we leave it up to the reader to investigate the variations if interested.

$$\min_G \max_D \left[E_z \big[1 - ln(D(G(z))) \big] + E_x \big[ln(D(x)) \big] \right] \quad (13)$$

Discriminator. The discriminator is trained against the generator to distinguish examples generated from the generator from the examples from the dataset. The model is optimization function is to adjust its parameters to maximize Equation 13 where x is a real example from the dataset and z is a noise.

Generator. This model is responsible for generating an example such that it causes the discriminator to misclassify the generated example as "real". Therefore, the generator optimizes Equation 13 opposite to the discriminator (*i.e.,* adjusting its parameter to maximize $E_z \big[1 - ln(D(G(z))) \big]$).

Adversarial GAN. Following the same paradigm, the discriminator in an Adversarial GAN is the victim model and the generator is the spoofer that adds negligible artifacts to an image such that the a deep model classifies an image differently from a human.

Take Away. GAN training is used to build a deep model (generator) that can cause a target deep-model (discriminator) to misclassify. We use the same paradigm as GAN but targeted towards non-deep-model functions where the discriminator is the proxy for the target function (*i.e.,* surrogate model), and the generator is the spoofer.

D. Extended Kalman Filter Detail

This subsection provides supplemental details about Extended Kalman Filter (EKF) from §II-C.

Predict. This stage utilizes the previous actuation command to provide the latest state estimation. For example, if the actuation command cause a full rotation of the wheels of a car, then the distance that the vehicle travels should be equal to the circumference of the wheel. The function that handles such calculations is the *transition function*.

Formally, the function f is the State Transition Function that predicts the State $\hat{x}_{t|t-1}$ at Time t, based on the Previous Estimation Update $\hat{x}_{t-1}$ and actuation command, as shown in Equation 14.

$$\underset{\text{predicted state}}{\hat{x}_{t|t-1}} = \underset{\text{transition function}}{f} (\; \underset{\text{previous state update}}{\hat{x}_{t-1}}, \underset{\text{actuation command}}{u_t} \;) \quad (14)$$

The formulation of f differs from vehicle to vehicle based on the hardware(*e.g.,* motor) and the physical characteristics(*e.g.,* wheel size).

There can be an error in the in the modeling of the vehicle dynamics due to expensiveness of accurate modeling in vehicles with complex dynamics (*e.g.,* quadrotor drones) or due to noise in the system (*e.g.,* due to friction or environmental factors). To capture how such an error can affect the prediction, error covariance matrix P is constructed via linearizing the

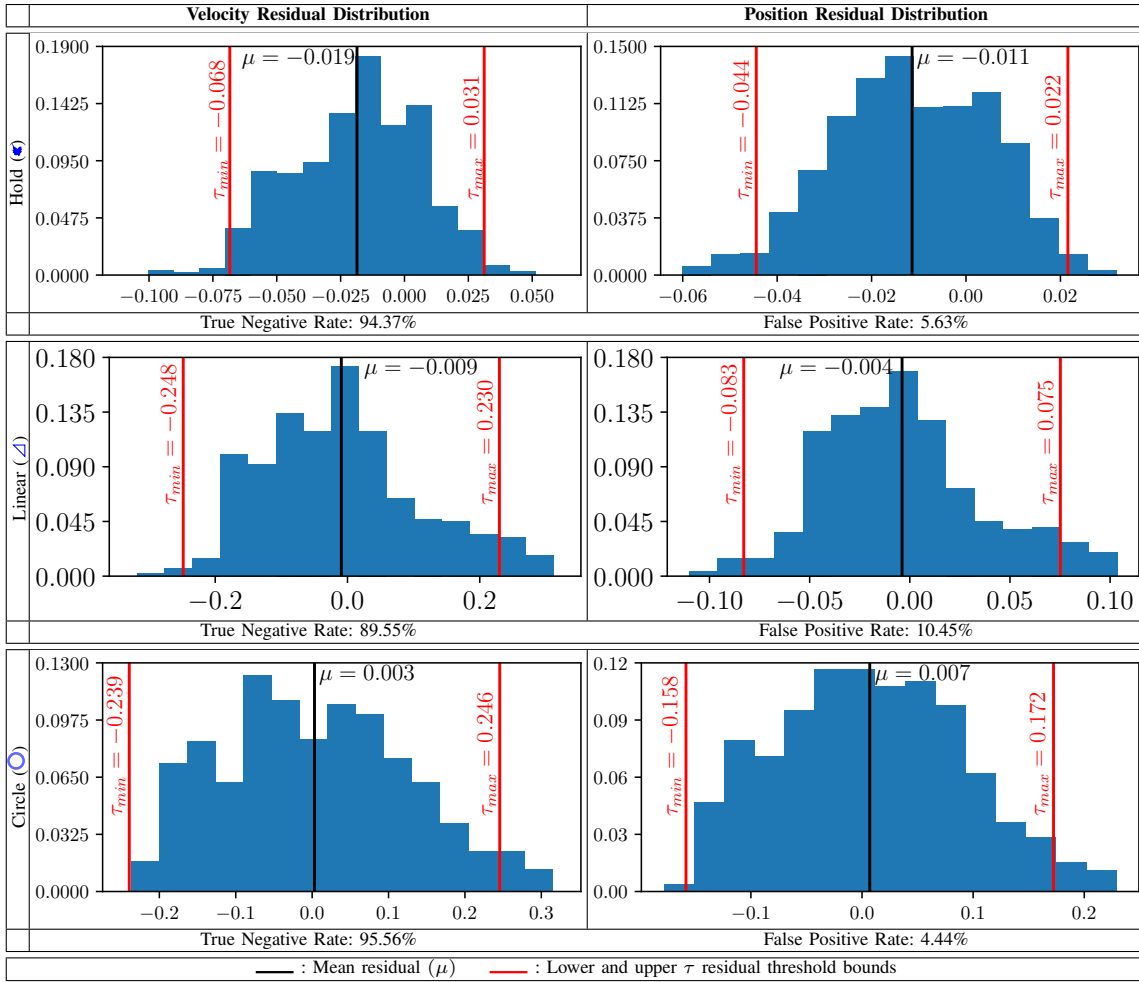| Velocity Residual Distribution	Position Residual Distribution

TABLE VI: Empirical probability density of the residual distribution represented as a histogram (bins=15) of the residual distributions over 10 trials of the normal run of the mission. X-axis is the residual and the y-axis is the probability density. The black line indicates the mean(μ). The red lines indicate the threshold range for τ anomaly detector which are two standard deviations away from the mean. The true positive rate (TPR) and false positive rate (FPR) resulting from using the τ thresholds on each of the normal missions are shown below the plot. This shows that the threshold bounds were very conservative: the thresholds would need to be more lax to achieve zero FPR under normal mission.

transition function by differentiating it *w.r.t.* the State, forming the Jacobian Matrix F, and *w.r.t.* the process noise, forming the process noise covariance matrix Q. as shown in Equation 15.

$$P_t = F\,P_{t-1}F^T + Q_t \tag{15}$$

where P_t is the error covariance, F is the transition Jacobian, and Q_t is the process noise covariance.

Due to the linearization process, true error covariance is approximated as an error covariance matrix. The error covariance matrix is used to create Kalman Gain in the next stage, responsible for how the state should update based on the difference between the predicted and the observed. *The differential between the approximation and the true covariance matrix as well as the error in the modeling of transition contributes to creating a space that the attacker can exploit.*

Update. Observation function h is differentiated *w.r.t.* the State as well as noise to create *Residual Covariance* Matrix C as shown in Equation 16.

$$C = H\,P_tH^T + R_t \tag{16}$$

where C is the residual covariance, H is the observation Jacobian, and R_t is the observation noise covariance.

The covariance matrices are used to construct Kalman Gain K_t that weighs how much influence the sensors have on state update, as shown in Equation 17.

$$K_t = P_t\,HC^{-1} \tag{17}$$

where P_t is the error covariance and K_t is the Kalman gain.

The next Error Covariance, P is updated according to Equation 15.

Algorithm 2 Input Exporation

Input: Target function, dataset, lower bound, upper bound, index
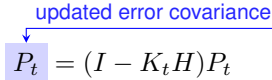**Output**: Expanded dataset
1: **func** EXPLOREINPUT(Q, S, ω_l, ω_u, I)
2:    $S' \leftarrow \{\}$
3:    **for** $i \in I$ **do**
4:       **for** $s \in S$ **do**
5:          $\omega \sim U(\omega_l, \omega_u)$
6:          $s' \leftarrow s$
7:          $s'[i] \leftarrow \omega + s[i]$
8:          $S' \leftarrow S' \cup \{s'\}$
9:       **end for**
10:    **end for**
11:    **return** S'
12: **end func**

$$P_t = (I - K_t H)P_t \qquad (18)$$

updated error covariance

Theory vs. Implementation. It is worth noting that the classic EKF assumes that all sensors are synchronous (*i.e.*, all sensors arrive and used at same time), while in practice, the arrival of sensors is not synchronous due to heterogeneity of the hardware. Therefore, there are adjustments to the algorithm that synchronizes the sensors *w.r.t.* the fastest sensor (*i.e.,* IMU) while keeping the core algorithm the same. This means that after a sensor is read, its values are stored in a buffer until they are ready to be used. Under our system model, the sensor value that is popped off the buffer is the value that is input into the EKF update task.

Flight controllers such as PX4[22] and Ardupilot[93], state transition function (f) do not directly map between the actuation command and previous state to predicted state. Instead, the flight controllers use Inertial Measurement Units (IMUs) as a proxy to the actuation command since an IMU can measure the actuation. This approach makes a flight controller compatible across various types of vehicle configurations where the actuators are located at different areas of the vehicle, and thus obviates the complex modeling of the State Transition Function f that is specific to each vehicle.

Process Noise Covariance Q is determined by the IMU driver and the parameters collected during IMU calibration. Observation Noise Covariance R is determined by the corresponding sensor driver. Notice that the Kalman Gain K, is only the function of the Covariance Matrices (P, H, and C) that are functions of predefined parameters and sensor hardware. This means that, the manipulation of the position and velocity sensor values do not affect the covariances directly. Therefore, an advantage of the software attack is that the adversary's only concern is to manipulate the Residual r_t to control how the vehicle estimates its State. In contrast, physical spoofing of GPS will risk affecting the Kalman Gain K, which makes manipulation more complex.

E. Surrogate and Spoofer Training Detail

Surrogate. To achieve the objective, we use Algorithm 3 to train the surrogate model. Because we have access to the input of the target function, we can always generate new training data by adding noise to the input and executing the target function. To ensure that the expanded inputs are varied,

Algorithm 3 Surrogate Training

1: $S \leftarrow$ EXPLOREINPUT() /* *Static Expansion* */
2: $range \leftarrow \max S - \min S$
3: $\omega_l \leftarrow \min S - 0.1range$
4: $\omega_u \leftarrow \max S + 0.1range$
5: $I \leftarrow \{\{\}\}$
6: **for** $e \in Epoch$ **do**
7:    $S' \leftarrow$ EXPLOREINPUT($Q, S, \omega_l, \omega_u,$)
8:    **for** $s \in S$ **do**
9:       $\omega \sim U(\omega_l, \omega_u)$ /* *Dynamic input exploration* */
10:       $\tilde{s} \leftarrow s_f + \omega$ /* *Add noise for better generalization* */
11:       $\tilde{s}_q \leftarrow Q(\tilde{s}_f)$ /* *Get corresponding label* */
12:       $\hat{s}_d \leftarrow D(\tilde{s}_f)$ /* *Surrogate's output* */
13:       $L_s \leftarrow ||\tilde{s}_q - \hat{s}_d||_2$
14:       $D \leftarrow Backprop(L_s)$
15:    **end for**
16: **end for**
17: **return** D

TABLE VII: Surrogate and spoof generator training parameters.

Trained Models	D_1	D_2, D_3	G
Model Description	Surrogate model for residual calculation	Surrogate model for velocity and position calculation respectively	Spoofer
Input Size	52		52
Input Description	Table. V	Table. V with output from D_1	Table. V
Hidden Size	50	50	50
# of Hidden Layers	2	10	5
Activation Func	ReLU	ReLU	ReLU
Output Size	8	1	2
Output Description	Velocity and position residual and residual variance during the state update	State estimation of north velocity or north position resulting from the state update	Injection values for GPS velocity and position
# of parameters	8158	28652	15502
Epochs	50	50	140
Training Time (Hrs)	16	24	74
% time spent on Query	99.3	99.1	92.5

we develop two methods to perform input expansion prior to training (*i.e., static expansion*) and during training (*i.e., dynamic expansion*) where *dynamic expansion* expands on top of already expanded input.

Line 9 in Algorithm 3 samples a uniformly random value which is the Regularizer ω. The resulting snapshot in line 10 is the input to both the surrogate and the target function. The loss(line 13) is the mean squared error of the two outputs(line 12 and 11). As a result, the surrogate model learns the mapping for the distribution collected.

Spoofer. We specify the implementation details to trains the two types of spoofers specified in §V-B: No Correction (NC) and Direction Bias (DB). For NC, the final loss ($\mathcal{L}$) in line 12 in the spoofer training algorithm is obtained from Equation 8 and for DB from Equation 11. In lines 6-10, the surrogate model is fined-tuned to ensure that its output (when used to

Algorithm 4 Spoofer Training

```
1: for e ∈ Epoch do
2:     /* Enumerate over the collected data */
3:     for s ∈ S do
4:         s̃ ← s + ω /* Noise added for regularization */
5:         /* Ensure that the surrogate model is within δ */
6:         while L_s > δ do
7:             s_q ← Q(s̃)
8:             L_s ← s_q − D(s̃)
9:             Backprop(L_s) /* Update the surrogate model */
10:        end while
11:        /* Train the spoofer */
12:        l ← 𝓛(D(s̃ + G(s̃))) /* Calculate the spoofer loss */
13:        Backprop(l) /* Updates the spoofer */
14:    end for
15: end for
16: return G
```


Fig. 16: Average test loss of the surrogate models during training.

trained the spoofer) is close to the target function. Parameters $\delta = 0.025$ for fine-tuning and $T = 0.1$ for L_b were used for this process.

F. Total Bregman Divergence

As mentioned in §VI-A, Total Bregman Divergence (TBD) is used to to show that an attack resulted in deviated trajectory is significant by comparing the trajectory of the normal mission to the true trajectory. To quantify the difference, we use the Total Squared Loss (TSL) of the TBD that is used to compare shapes [62]. In this paper we compare the shape of the trajectory.

TSL of the TBD is calculated by comparing the distribution between the two trajectories. It starts with finding a Gaussian representation of the trajectory by training Gaussian mixture model (*i.e.,* a set of n Gaussian distributions), p, for a trajectory as shown on Equation 19.

$$p(x) = \sum_i^n a_i \mathcal{N}(x; \mu_i, \Sigma_i) \qquad (19)$$

Each of the Gaussians is defined by mean μ, co-variance Σ and weight a. Therefore, the TSL of the TBD between two distributions is calculated as shown in equation Equation 20

$$tSL(p_r, p_q) = \frac{d_{r,r} + d_{q,q} - 2d_{r,q}}{\sqrt{q + 4d_{q,q}}} \qquad (20)$$

where

$$d_{k,h} = \sum_i^n a_i^k a_i^h \mathcal{N}(0, \mu_i^k - \mu_i^h, \Sigma_i^k + \Sigma_i^h) \qquad (21)$$

Furthermore, we calculate the TBD for the reported position to show that the observed trajectory is similar to the mission trajectory

G. Surrogate Models Test Performance

The result of the surrogate model evaluation process (§VI-F) is explained here. The test loss converging close to zero indicates that the model learned target function. It is corollary to the case where overfitting to the training set cause the test loss to gradually increase. Figure 16 shows the test loss plots for the three models. The right plot is the zoomed plot
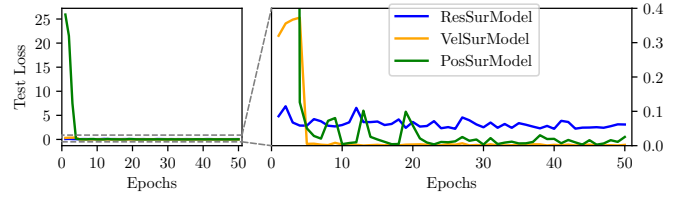
of the left. We see that VelSurModel (D_2) and PosSurModel (D_3) converged close to zero significantly after 5 epochs. For ResSurModel (D_1), the test loss at the start already is low. That indicates the residual calculation can be learned easily. The reduced fluctuation over each epoch indicates that the model converged.

H. Results for other missions

This section continues the discussion of REQUIEM on linear (L) and the hold (H) missions from §VI-C.

1) Attack on Linear Mission Performance: REQUIEM also succeeds for the linear mission since it remained stealthy while causing deviation up to approximately $11m$. When examining the Total Bregman Divergence (TBD) for each of the attacks in Figure 17, REQUIEM's direction bias attack has the largest difference between the attack trajectory and the system POV trajectory. The second largest deviation is naïve Position Constant Offset (PCO) attack, but PCO is not stealthy since it triggered both alarms as shown in Table IV.
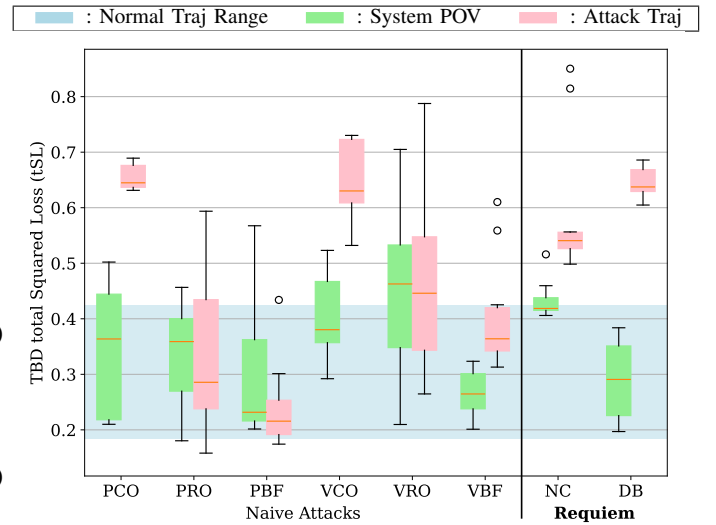

Fig. 17: Total Bregman Divergence (TBD) for a Linear Mission. In this context, only REQUIEM attacks are stealthy as shown in Table IV. The plot shows that REQUIEM attacks achieve significant deviation whereas naïve attacks fails to do so.

The trajectory in Figure 18 shows that REQUIEM's Direction Bias (DB) did not cause much deviation during the eastward movement towards the first setpoint but exploited the northward movement toward the second setpoint. Similar
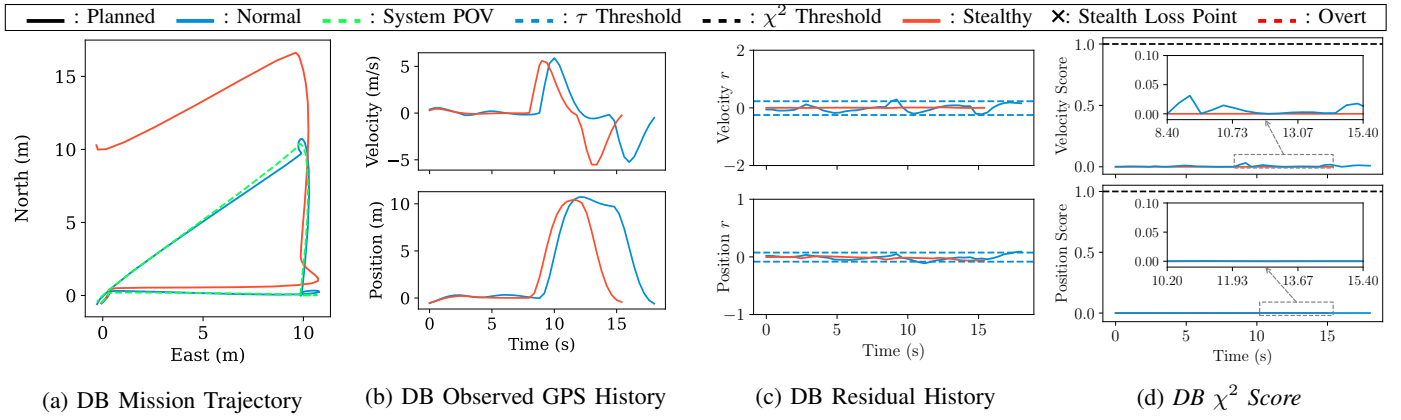
(a) DB Mission Trajectory   (b) DB Observed GPS History   (c) DB Residual History   (d) *DB χ^2 Score*

Fig. 18: Analysis of REQUIEM's directional bias (DB) attack on linear mission. (a) shows the observed trajectory converging to each setpoint. Notice the side effect of slightly faster convergence towards the second setpoint in (b) caused by the system not noticing the overshoot due to the influence of DB. DB manages to stay within the threshold as shown on (c) but the position is slowly converging towards the threshold due to the spoofer observing position values far outside the training distribution.

case for REQUIEM's No Correction (NC) attack in Figure 19 where most of the deviation was induced during the northward movement but not as much compared to DB. DB and NC demonstrates REQUIEM's opportunistic nature. Therefore, if the vehicle were to only move eastward, perpendicular to the target axis, deviation would be much slower. In contrast, Position Boiling Frog (PBF) in Figure 20 tries to cause stealthy deviation by slowly injecting small values but fails when the vehicle is converging at the first setpoint. PBF's failure illustrates that the physical context of the vehicle at the time of the attack is important.

> **Insight.** Attacks have a minor side-effect on the duration of the setpoint convergence.

REQUIEM caused the vehicle to believe it is operating nominally resulting in faster convergence. Where as the naïve PBF attack caused the vehicle to take longer to converge, as shown in Figure 20, since it slowly shifts the GPS position even when the vehicle tries to converge. Under REQUIEM, the behavior of the side-effect varies based on the control. For example, if the control was not tuned correctly and caused the state prediction to think that the vehicle overshot the setpoint, it would take longer for the vehicle to think that it converged. However, the environmental factors in real-world (*e.g.,* wind) can cause these side effects to seem nominal (hence a minor side-effect).

2) Attack on Hold Mission Performance: Hold Mission demonstrates REQUIEM's capability of stealthily deviating a stationary vehicle. The mission is more difficult to attack compared to the other two due to the lack of movement, hence less exploitable space and very low residual and τ threshold. Despite the challenges as shown in Figure 21c and Figure 23b, REQUIEM partially succeeds in hold mission with max stealthy deviation up to $17.72m$ average and up to $289m$ overt deviation. Even when the attack becomes overt, system POV trajectory still remains well within $1m$ from the mission parameter. REQUIEM's DB attack has less stealthy deviation but higher overt deviation compared to REQUIEM's

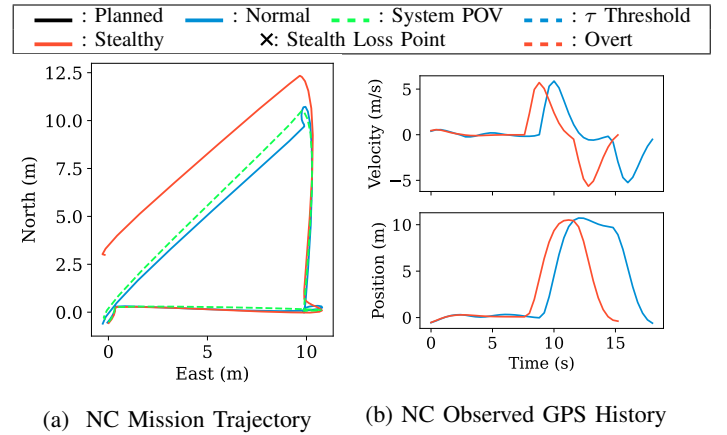(a) NC Mission Trajectory   (b) NC Observed GPS History

Fig. 19: Analysis of REQUIEM's no correction (NC) attack on linear mission. (a) shows that NC also cause deviation only after the first setpoint. NC also caused the system to not notice the overshoot in the second setpoint, resulting in the side effect of converging slightly faster towards the second setpoint as shown on (b).

no correction attack due to trading off some stealthiness for northern bias deviation.

Mission parameters affect the attack difficulty in stealthy attack since the nominal noise change based on the movement. A corollary is that the nominal residual distribution is reflective of the motion.

> **Insight.** However, building a momentary detection mechanism around by correlating motion to the residual is not feasible since it will result in high false positives.

In such a case, defender is tasked to decide between whether the latest residual is indeed 'too good' to be nominal or whether it is actually nominal. Correctly making that decision will take repeated observation since the defender needs to look at the overall trend (*i.e.,* post-flight analysis) otherwise

(a) PBF Mission Trajectory
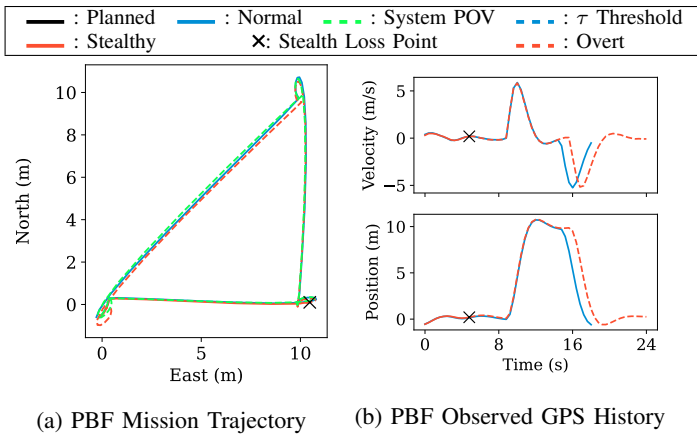
(b) PBF Observed GPS History

Fig. 20: Analysis of naïve position boiling frog (PBF) on linear mission. (a) caused very small change in trajectory between the second and third setpoint where the trajectory is slightly more south than the mission. (b) also shows that the side effect is visible even when a very small value is added.

will risk severe false positive alarm. By the time a thorough decision is made, the vehicle has already deviated from the mission parameters, and hence the attack has already been accomplished. Even then, the general approach of REQUIEM gives the potential to counter such mechanism by incorporating the correlation profiling into the spoofer's loss function (§V-B) as REQUIEM has demonstrated to control the residual to the value of the attacker's choosing (*i.e.,* zero).

(a) DB Mission Trajectory

(b) DB Observed GPS History
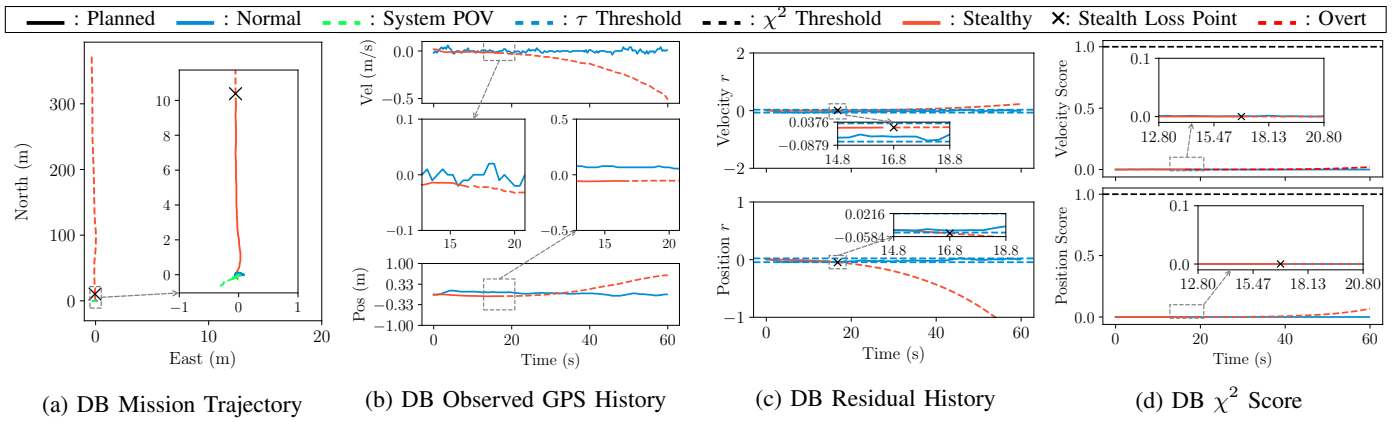
(c) DB Residual History

(d) DB χ^2 Score

Fig. 21: Analysis of REQUIEM's directional bias (DB) attack on hold mission. The observed trajectory in (a) shows that it stayed within a 1 m of the planned mission while the overall deviation of greater than 300m was achieved for that run. (b) stayed similar to the normal mission until the attack became overt. (c) the position residual rapidly increased after the attack became overt faster than the velocity residual. Another indication of large position value outside of the training distribution of the spoofer causing the loss of stealth. Despite τ AD triggering, χ^2 score in (d) stayed near zero throughout.
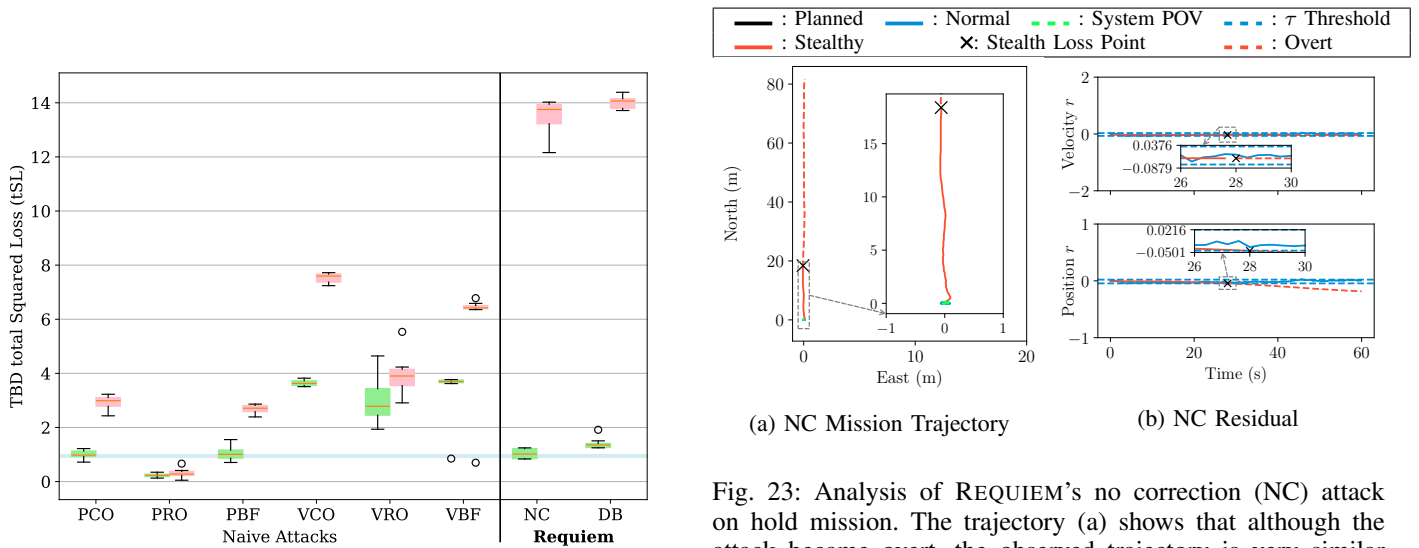

Fig. 22: The plot shows that REQUIEM attacks achieve significant deviation whereas naïve attacks fails to do so.
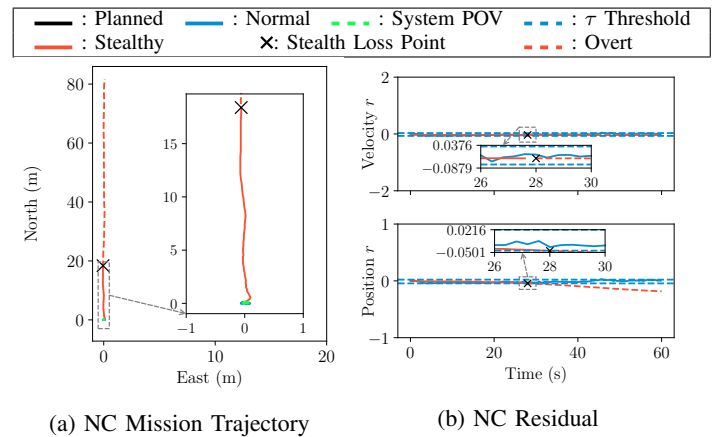

(a) NC Mission Trajectory

(b) NC Residual

Fig. 23: Analysis of REQUIEM's no correction (NC) attack on hold mission. The trajectory (a) shows that although the attack became overt, the observed trajectory is very similar to the planned trajectory. (b) managed to keep both residuals as close to zero as possible. Although the observed position value at the moment of stealth loss is around 20, the stringent threshold for hold mission resulted in the alarm.